

Monadic Second-Order Model Theory

Achim Blumensath

Monadic Second-Order Model Theory

Achim Blumensath

ab

BRNO 2025

Achim Blumensath
blumens@fi.muni.cz

This document was last updated 2025-02-23.
The latest version can be found at

www.fi.muni.cz/~blumens

COPYRIGHT 2025 Achim Blumensath



This work is licensed under the *Creative Commons Attribution 4.0 International License*. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Contents

A. Fundamentals	1
I Logics and Their Expressive Powers	3
1. Structures and logics	4
2. Simple translations between logics	11
3. Theories and back-and-forth arguments	16
4. Operations for monadic second-order logic	23
5. Operations for first-order logic	51
II Finite Words	61
1. Words and languages	61
2. Semigroups and Green's relations	63
3. Simon's Lemma	73
4. Regular languages of finite words	86
5. First-order logic	93
III Infinite Words	107
1. Ramsey theory	107
2. The theory of ω	117
3. ω -semigroups	120
4. ω -automata	130
IV Parity Games	145
1. Positional games	145
2. Reachability games	147

3. Gale-Stewart games	157
4. Regular games and parity games	159
5. The modal μ -calculus	179
V Trees	193
1. Composition theorems	193
2. Tree automata	206
3. The Muchnik Iteration	228
4. Löwenheim-Skolem theorems	237
5. The Cantor Topology	243
6. Counting quantifiers	259
B. Structure Theory	281
VI Linear Orders	283
1. Dense and scattered orders	283
2. Partition theorems	299
3. Interpretations	308
4. Regular linear orders	313
5. Modest linear orders	327
6. Choice functions	393
7. Uniformisation	414
8. First-order logic	441
VII Sparse Structures	461
1. Spanning forests	461
2. Sparse hypergraphs	481
3. Translating GSO into MSO	490
4. Sparse distributions	497
VIII Tree-Width and Graph Minors	531
1. Tree-decompositions	531
2. Minors	543

3. Brambles	553
4. The Excluded Grid Theorem	560
5. Branch-decompositions and tangles	578
6. Well-quasi-orderings	592
IX Crossing-Width	599
1. Partitions and ranks	599
2. Decompositions	610
3. Terms	615
4. Tree-width and crossing-width	622
5. Interpretations	631
6. Non-standard crossing-width	639
7. The monadic independence property	644
X Guarded Second-Order Transductions	653
1. Transductions	653
2. Tree-decompositions	665
3. Trees of bounded height	667
4. The Transduction Hierarchy	685
5. Defining tree-decompositions	692
C. Applications	713
XI The Gaucal Hierarchy	715
1. Finitely Presentable Structures	715
2. Prefix-recognisable structures	725
3. VR-equational and HR-equational structures	732
4. Regular path-interpretations	739
5. Pushdown graphs	751
6. Context-free groups	777
7. Graphs and paths	796
8. Axiomatisations	816

XII Automatic Structures	833
1. Automatic presentations	833
2. Interpretations	841
3. Closure properties	849
4. Undecidability	853
5. Injective presentations	857
6. Partition theorems	863
7. Counting quantifiers	883
8. Proving non-automaticity	892
9. Automatic groups	914
10. Automatic semirings	943
11. Automatic orders	957
12. Subclasses of automatic structures	970
Bibliography	1005
Symbol Index	1021
Index	1027

Part A

Sundamentals

I Logics and Their Expressive Powers

THE TWO CENTRAL TOPICS of this book are (I) the model checking problem for specific structures and (II) the study of the expressive power of various logics. To this end we will develop techniques to compute and compare the theories of given structures. This obviously solves the model checking problem since, if we know the theory of a structure \mathcal{A} , we can decide whether a formula is satisfied by it. But this also helps us to prove that certain things are not expressible in a given logic L . If we can find two structures \mathcal{A} and \mathcal{B} with the same L -theory such that \mathcal{A} has a given property P , but \mathcal{B} does not, then the property P cannot be expressed in L .

Notation. The following basic notation will be used throughout the book. For $n < \omega$, we set $[n] := \{0, \dots, n-1\}$. We tacitly identify a tuple $\vec{a} = \langle a_0, \dots, a_{n-1} \rangle$ with the set $\{a_0, \dots, a_{n-1}\}$ of its components. This allows us to write $\vec{a} \subseteq C$ or $c \in \vec{a}$. The empty tuple is $\langle \rangle$.

$\wp(A)$ denotes the power set of A , and $A + B$ is the disjoint union of A and B . For a function $f : A \rightarrow B$, we denote the domain by $\text{dom } f := A$ and its range by $\text{rng } f \subseteq B$. We write $f \upharpoonright X : X \rightarrow B$ for the restriction of f to the set X .

For a partial order $\langle A, \leq \rangle$ and a subset $X \subseteq A$, we set

$$\uparrow X := \{ a \in A \mid a \geq x \text{ for some } x \in X \},$$

$$\downarrow X := \{ a \in A \mid a \leq x \text{ for some } x \in X \}.$$

We denote the infimum and the supremum of two elements x and y by, respectively, $x \sqcap y$ and $x \sqcup y$.

1 Structures and Logics

Logics are formal languages designed to talk about mathematical objects. As we will deal with several different logics in the course of this book it is useful to adopt an abstract point of view. In general a logic consists of (I) a class of objects to talk about; (II) a set of statements we can make about them; and (III) a relation telling us which statements hold for a given object.

Definition 1.1. A logic is a triple $\langle L, \mathcal{M}, \models \rangle$ consisting of a set L of formulae, a class \mathcal{M} of models, and a satisfaction relation $\models \subseteq \mathcal{M} \times L$. To keep notation light, we usually identify a logic with its set of formulae L . \dashv

For instance, we can define first-order logic as a triple $\langle \text{FO}[\Sigma], \text{STR}[\Sigma], \models \rangle$ where $\text{FO}[\Sigma]$ is the set of all first-order formulae (without free variables) over the signature Σ and $\text{STR}[\Sigma]$ is the class of all Σ -structures. For formulae with free variables we can use the logic $\langle \text{FO}[\Sigma, X], \text{STR}[\Sigma, X], \models \rangle$ where $\text{FO}[\Sigma, X]$ is the set of all first-order formulae with free variables in the set X and $\text{STR}[\Sigma, X]$ is the class of all pairs $\langle \mathcal{A}, \beta \rangle$ consisting of a Σ -structure \mathcal{A} and a variable assignment $\beta : X \rightarrow A$.

The logics we consider in this book are mostly variants of first-order logic and monadic second-order logic. Let us quickly recall their definitions. A signature Σ is a set of relation symbols and function symbols, each of which has an associated arity. A Σ -structure $\mathcal{A} = \langle A, (\xi^{\mathcal{A}})_{\xi \in \Sigma} \rangle$ consists of a set A together with

- one n -ary relation $R^{\mathcal{A}} \subseteq A^n$, for every relation symbol $R \in \Sigma$ of arity n ,
- one n -ary function $f^{\mathcal{A}} : A^n \rightarrow A$, for every function symbol $f \in \Sigma$ of arity n .

Note that we allow functions of arity 0, which correspond to *constants*. Most of the time in this book we assume that all signatures are finite and purely relational.

Most of the time we will only work with 1-sorted structures, but sometimes many-sorted ones are more convenient. An S -sorted structure

$$\mathcal{A} = \langle (A_s)_{s \in S}, (\xi^{\mathcal{A}})_{\xi \in \Sigma} \rangle$$

has one domain A_s , for each sort $s \in S$, and each relation symbol and function symbol has an associated *type*. For an n -ary relation symbol R , this type is an n -tuple $\bar{s} \in S^n$, for an n -ary function symbol f , it is an $(n+1)$ -tuple $\bar{s}t \in S^{n+1}$, which we will usually write at $\bar{s} \rightarrow t$. If R has type \bar{s} , the corresponding relation is of the form $R^{\mathcal{A}} \subseteq \prod_i A_{s_i}$. Similarly, if f has type $\bar{s} \rightarrow t$, we are given a function $f^{\mathcal{A}} : \prod_i A_{s_i} \rightarrow A_t$.

Example. (a) The field of real numbers $\langle \mathbb{R}, +, \cdot, 0, 1, \leq \rangle$ is a structure with signature $\{+, \cdot, 0, 1, \leq\}$, where $+$ and \cdot are binary function symbols, 0 and 1 are 0-ary function symbols, and \leq is a binary relation symbol.

(b) A graph is a structure $\langle V, E \rangle$ with a single binary relation $E \subseteq V \times V$.

(c) We can represent a vector space V over a field K either as a 1-sorted structure of the form $\mathfrak{V} = \langle V, +, 0, (f_a)_{a \in K} \rangle$ where scalar multiplication is split into separate functions $f_a : V \rightarrow V$, for each $a \in K$, or we can use a two-sorted structure $\mathfrak{V} = \langle V, K, +, 0, \cdot \rangle$ with $+: V \times V \rightarrow V$ and $\cdot : K \times V \rightarrow V$. \dashv

The main logics we are concerned with in this book are first-order logic and various variants of monadic second-order logic. Recall that *first-order logic* $\text{FO}[\Sigma]$ consist of formulae that are built up from *atomic formulae* of the form $s = t$ and $Rt_0 \dots t_{n-1}$, where $R \in \Sigma$ is an n -ary relation symbol and $s, t, t_0, \dots, t_{n-1}$ are terms built up from variables and the function symbols in Σ . Such atomic formulae can be combined using boolean operations \wedge (conjunction), \vee (disjunction), \neg (negation), and first-order quantifiers $\exists x$ and $\forall x$.

Definition 1.2. Let Σ be a signature. The formulae of *monadic second-order logic* $\text{MSO}[\Sigma]$ are built up from atomic formulae of the form $s = t$, Zt , and $Rt_0 \dots t_{n-1}$, where $R \in \Sigma$ is an n -ary relation symbol, Z is a *set variable*, and $s, t, t_0, \dots, t_{n-1}$ are terms built up from first-order variables and the function symbols in Σ . Such atomic formulae can be combined using boolean operations \wedge (conjunction), \vee (disjunction), \neg (negation), and quantifiers $\exists x, \forall x, \exists Z$, and $\forall Z$, where x is a first-order variable and Z is a set variable.

The semantics of such a formula is defined as follows. Given a formula $\varphi(\bar{x}, \bar{Z}) \in \text{MSO}[\Sigma]$ with free first-order variables \bar{x} and free set variables \bar{Z}

and given a Σ -structure \mathcal{A} , a tuple of elements \bar{a} of \mathcal{A} , and a tuple of subsets \bar{P} of \mathcal{A} , we define the *satisfaction relation*

$$\mathcal{A} \models \varphi(\bar{a}, \bar{P})$$

by induction on φ . The definition is analogous to that for first-order logic. An atomic formula Zt holds in \mathcal{A} if the element denoted by the term t belongs to the set denoted by Z . A formula of the form $\exists Z\psi$ holds if there exists a set satisfying ψ , and $\forall Z\psi$ holds if every set satisfies ψ .

Throughout we use lower case letters for first-order variables and upper case ones for set variable. For readability we will sometimes use common short-hands such as, $s \neq t$ instead of $\neg(s = t)$, or $t \in Z$ instead of Zt .

As above we write $\text{MSO}[\Sigma, X]$ for the set of MSO-formulae with free variables in a given set X . A model of such a formula consists of a Σ -structure \mathcal{A} and a variable assignment β . We usually write $\varphi(\bar{x}, \bar{Z})$ for a formula φ to indicate that the free variables of φ are among the variables $\bar{x}\bar{Z}$. This allows us to use the more common notation

$$\mathcal{A} \models \varphi(\bar{a}, \bar{P}) \quad \text{iff} \quad \langle \mathcal{A}, \beta \rangle \models \varphi,$$

where β is the variable assignment mapping x_i to a_i and Z_i to P_i . Since Σ and X can usually be inferred from the context, we will frequently simplify notation by writing MSO instead of $\text{MSO}[\Sigma, X]$, and similarly for other logics.

Example. (a) For a linear order $\mathcal{A} = \langle A, \leq \rangle$, we can say that y is the immediate successor of x by the FO-formula

$$\varphi(x, y) := x \leq y \wedge x \neq y \wedge \forall z[x \leq z \wedge z \leq y \rightarrow (z = x \vee z = y)].$$

(b) For a tree $\mathcal{T} = \langle T, \leq \rangle$ where \leq is the predecessor order, we can express that a set variable X contains an infinite branch by the MSO-formula

$$\begin{aligned} \exists Z[Z \subseteq X \wedge Z \neq \emptyset \wedge \forall x \forall y[Zx \wedge Zy \rightarrow (x \leq y \vee y \leq x)) \\ \wedge \forall x \exists y[Zx \rightarrow x < y \wedge Zy]]]. \end{aligned}$$

(c) Given a graph $\mathcal{G} = \langle V, E \rangle$, the MSO-formula

$$\varphi(x, y) := \forall Z[Zx \wedge \forall u \forall v(Zu \wedge Euv \rightarrow Zv) \rightarrow Zy]$$

expresses that there exists a path from x to y .

(d) We can say that a graph $\mathcal{G} = \langle V, E \rangle$ is connected by the formula

$$\forall x \forall y \varphi(x, y),$$

where φ is the formula from (c).

We will also study the following variants of monadic second-order logic.

Definition 1.3. Let Σ be a signature.

(a) *Weak monadic-second order logic* $\text{WMSO}[\Sigma]$ has the same syntax as $\text{MSO}[\Sigma]$, but all set variables range over finite sets only.

(b) *Monadic-second order logic with first-order counting* $\text{CMSO}[\Sigma]$, or *counting monadic-second order logic* for short, is the extension of $\text{MSO}[\Sigma]$ by statements of the form

$$|X| < \aleph_0 \quad \text{and} \quad |X| \equiv k \pmod{m},$$

for a set variable X and finite numbers $k, m < \omega$. A statement of the form $|X| < \aleph_0$ holds if X is a finite set, and $|X| \equiv k \pmod{m}$ is true if, X is finite and its size is congruent k modulo m . We write $\text{MSO}[\text{inf}]$ if we only allow predicates of the first form.

(c) Let \mathcal{A} be a Σ -structure. A tuple $\bar{a} \in A^n$ is *guarded* if there exists a relation R of \mathcal{A} containing a tuple $\bar{c} \in R$ with $\bar{a} \subseteq \bar{c}$. Here, we allow R to be the equality relation $=$, even though it is not present in the signature. A relation $S \subseteq A^n$ is *guarded* if every tuple $\bar{a} \in S$ is guarded.

(d) *Guarded second-order logic* $\text{GSO}[\Sigma]$ extends first-order logic by atomic formulae of the form $Zt_0 \dots t_{n-1}$, where t_0, \dots, t_{n-1} are terms and Z is a relation variable of arity n , and by quantifiers $\exists Z$ and $\forall Z$ over relation variables. A formula of the form $\exists Z\psi$ holds if there exists a guarded relation satisfying ψ , and $\forall Z\psi$ holds if every guarded relation satisfies ψ .

Example. We consider undirected graphs $\mathfrak{G} = \langle V, E \rangle$ as structures over the signature $\{E\}$ consisting of one binary edge relation (irreflexive and symmetric).

(a) To express that a graph has a Hamiltonian cycle we can write down a GSO-formula stating that there is a guarded binary relation Z (i.e., a set of edges) such that

- for every vertex x there are unique vertices y and z with $(y, x) \in Z$ and $(x, z) \in Z$,
- every two vertices are connected by a sequence of Z -edges.

(b) A *minor* of a graph \mathfrak{G} is a graph \mathfrak{H} obtained from the first graph by deleting vertices and edges and by contracting edges. To say that a fixed finite graph \mathfrak{H} is a minor of the given graph, we can use an MSO-formula stating that, for each vertex v of \mathfrak{H} , there exists a set X_v such that

- the subgraph induced by X_v is connected and
- for every edge $\langle u, v \rangle$ of \mathfrak{H} there is an edge connecting some vertex of X_u with some vertex of X_v .

As defined above the logic MSO is not always convenient to use in proofs. Therefore, we introduce a simplified version that still has the same expressive power.

Definition 1.4. Let Σ be a relational signature. The logic $\text{MSO}^\circ[\Sigma]$ has atomic formulae of the form

$$\begin{aligned} X \subseteq Y, \quad \text{sing}(X), \quad RX_0 \dots X_{n-1}, \\ X \cap Y = \emptyset, \quad \text{cover}(X_0, \dots, X_{n-1}), \end{aligned}$$

where $R \in \Sigma$ is an n -ary relation symbol and $X, Y, X_0, \dots, X_{n-1}$ are set variables. The logic is closed under boolean operations and set quantifiers. The formulae $X \subseteq Y$ and $X \cap Y = \emptyset$ have the obvious meaning. $\text{sing}(X)$ states that $|X| = 1$. An atomic formula of the form $\text{cover}(X_0, \dots, X_{n-1})$ holds if the union $X_0 \cup \dots \cup X_{n-1}$ contains the whole universe, while a formula of the form $RX_0 \dots X_{n-1}$ holds if each set X_i is a singleton $\{a_i\}$ and the tuple $\langle a_0, \dots, a_{n-1} \rangle$ of elements belongs to R .

Remark. (a) We frequently use abbreviations like

$$\begin{aligned} (X = Y) &:= (X \subseteq Y) \wedge (Y \subseteq X), \\ (X \subset Y) &:= (X \subseteq Y) \wedge \neg(Y \subseteq X), \\ (X = \emptyset) &:= (X \cap X = \emptyset). \end{aligned}$$

(b) Note that every MSO° -formula is equivalent to one that does not contain atomic formulae of the form $X \cap X' = \emptyset$, $\text{sing}(X)$, or $\text{cover}(\bar{X})$ since we can define these in terms of \subseteq .

$$\begin{aligned} X \cap X' = \emptyset &\equiv \forall Y [Y \subseteq X \wedge Y \subseteq X' \rightarrow \forall Z (Y \subseteq Z)], \\ \text{sing}(X) &\equiv X \neq \emptyset \wedge \forall Y [Y \subset X \rightarrow Y = \emptyset], \\ \text{cover}(\bar{X}) &\equiv \forall Z [\text{sing}(Z) \rightarrow \bigvee_i Z \subseteq X_i]. \end{aligned}$$

But note that this translation does increase the quantifier rank.

Lemma 1.5. Let Σ be a relational signature. For every formula $\varphi(\bar{x}, \bar{Z}) \in \text{MSO}[\Sigma]$, there is a formula $\varphi^\circ(\bar{X}, \bar{Z}) \in \text{MSO}^\circ[\Sigma]$ such that

$$\mathfrak{A} \models \varphi^\circ(\{a_0\}, \dots, \{a_{m-1}\}, \bar{P}) \quad \text{iff} \quad \mathfrak{A} \models \varphi(a_0, \dots, a_{m-1}, \bar{P}),$$

for every Σ -structure \mathfrak{A} and all parameters \bar{a} and \bar{P} .

Proof. We define φ° by induction as follows.

$$\begin{aligned} (x = y)^\circ &:= \text{sing}(X) \wedge \text{sing}(Y) \wedge X \subseteq Y \wedge Y \subseteq X, \\ (RX_0 \dots X_{n-1})^\circ &:= RX_0 \dots X_{n-1}, \\ (\varphi \wedge \psi)^\circ &:= \varphi^\circ \wedge \psi^\circ, & (\exists x \psi)^\circ &:= \exists X [\text{sing}(X) \wedge \psi^\circ], \\ (\varphi \vee \psi)^\circ &:= \varphi^\circ \vee \psi^\circ, & (\forall x \psi)^\circ &:= \forall X [\text{sing}(X) \rightarrow \psi^\circ], \\ (\neg \varphi)^\circ &:= \neg \varphi^\circ, & (\exists Z \psi)^\circ &:= \exists Z \psi^\circ, \\ & & (\forall Z \psi)^\circ &:= \forall Z \psi^\circ. \end{aligned} \quad \square$$

Analogous statements hold for the other variants of MSO.

Exercise 1.1. We consider coloured linear orders of the form $\langle A, \leq, P \rangle$ where $P \subseteq A$ is a unary predicate. Find MSO-formulae expressing the following statements:

- (a) The set P is dense, i.e., it is non-empty and between any two elements of A there is an element of P .
- (b) The set P contains infinitely many elements.
- (c) The set P is finite and it has an even number of elements.

Exercise 1.2. An $(m \times n)$ -grid is a graph $\mathfrak{G} = \langle V, E \rangle$ where

$$V := [m] \times [n],$$

$$E := \{ \langle \langle i, k \rangle, \langle j, l \rangle \rangle \mid |i - j| + |k - l| = 1 \}.$$

- (a) Construct an MSO-formula expressing that a graph is a grid.
- (b) For each of the following functions $f : \omega \rightarrow \omega$, find an MSO-formula stating that the given graph is an $(n \times f(n))$ -grid, for some n .

$$(i) \quad f(n) = n, \quad (ii) \quad f(n) = n^2, \quad (iii) \quad f(n) = 2^n.$$

Exercise 1.3. We can encode a finite word $w = a_0 \dots a_{n-1} \in \Sigma^*$ over the alphabet Σ by a *word structure*

$$\hat{w} := \langle [n], \leq, (P_a)_{a \in \Sigma} \rangle,$$

where the universe $[n] = \{0, \dots, n-1\}$ is the set of positions in the word w and the predicates

$$P_a := \{ i < n \mid a_i = a \}$$

contain all positions carrying the corresponding letter. Prove that, for every regular expression α , there exists an MSO-formula φ such that

$$\hat{w} \models \varphi \quad \text{iff} \quad w \in L(\alpha).$$

Hint. First construct, for each regular expression α , an MSO-formula $\varphi(x, y)$ such that

$$\hat{w} \models \varphi(x, y) \quad \text{iff} \quad w[x, y] \in L(\alpha),$$

where $w[x, y]$ denotes the factor of w between positions x and y .

2 Simple Translations Between Logics

In this section we relate the various logics introduced above to each other, and we provide translations between them. We start with MSO and FO.

Definition 2.1. Let Σ be a relational signature.

(a) The *power-set structure* of a Σ -structure \mathfrak{A} is the structure $\wp(\mathfrak{A})$ with signature $\Sigma + \{\subseteq\}$ whose universe is the power set $\wp(A)$ of the universe of \mathfrak{A} . The relation symbol \subseteq denotes the usual subset relation on $\wp(A)$. For each n -ary relation symbol $R \in \Sigma$, $\wp(\mathfrak{A})$ has the relation

$$R^{\wp(\mathfrak{A})} := \{ \bar{P} \in \wp(A)^n \mid \text{each } P_i = \{a_i\} \text{ is a singleton and } \bar{a} \in R^{\mathfrak{A}} \}.$$

(b) The *finite power-set structure* of a Σ -structure \mathfrak{A} is the substructure $\wp_{\text{fin}}(\mathfrak{A})$ of $\wp(\mathfrak{A})$ consisting of all *finite* subsets of A .

It is straightforward to check that MSO over Σ -structures corresponds to FO over their power-set structures.

Proposition 2.2. Let Σ be a relational signature.

(a) For every MSO $[\Sigma]$ -formula $\varphi(\bar{X})$, there exists an FO $[\Sigma + \{\subseteq\}]$ -formula $\varphi'(\bar{x})$ such that

$$\mathfrak{A} \models \varphi(\bar{P}) \quad \text{iff} \quad \wp(\mathfrak{A}) \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all sets \bar{P} in \mathfrak{A} .

(b) For every FO $[\Sigma + \{\subseteq\}]$ -formula $\varphi(\bar{x})$, there exists an MSO $[\Sigma]$ -formula $\varphi'(\bar{X})$ such that

$$\wp(\mathfrak{A}) \models \varphi(\bar{P}) \quad \text{iff} \quad \mathfrak{A} \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all sets \bar{P} in \mathfrak{A} .

Proof. (a) By Lemma 1.5 and the remark after Definition 1.4, we may assume that φ is an MSO 0 -formula without subformulae of the form $\text{sing}(X)$, $X \cap Y = \emptyset$, or $\text{cover}(\bar{X})$. Then we obtain the desired formula φ' from φ by replacing every set variable X by a corresponding first-order variable x .

(b) It is sufficient to construct an MSO 0 -formula. We obtain it from φ by replacing every first-order variable x by a corresponding set variable X . \square

We obtain the analogous result for finite power-sets and weak MSO. The proof is identical to the one above.

Proposition 2.3. *Let Σ be a relational signature.*

(a) *For every WMSO[Σ]-formula $\varphi(\bar{X})$, there exists an FO[$\Sigma + \{\subseteq\}$]-formula $\varphi'(\bar{x})$ such that*

$$\mathfrak{A} \models \varphi(\bar{P}) \quad \text{iff} \quad \wp_{\text{fin}}(\mathfrak{A}) \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all finite sets \bar{P} in \mathfrak{A} .

(b) *For every FO[$\Sigma + \{\subseteq\}$]-formula $\varphi(\bar{x})$, there exists an WMSO[Σ]-formula $\varphi'(\bar{X})$ such that*

$$\wp_{\text{fin}}(\mathfrak{A}) \models \varphi(\bar{P}) \quad \text{iff} \quad \mathfrak{A} \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all finite sets \bar{P} in \mathfrak{A} .

There is also a variant for CMSO, again with basically the same proof.

Proposition 2.4. *Let Σ be a relational signature and $m < \omega$ a constant.*

(a) *Let $\varphi(\bar{X})$ be a CMSO[Σ]-formula such that, in every subformula of the form $|X| \equiv k \pmod{n}$, the constant n divides m . Then there exists an FO[$\Sigma + \{\subseteq, \bar{M}\}$]-formula $\varphi'(\bar{x})$ such that*

$$\mathfrak{A} \models \varphi(\bar{P}) \quad \text{iff} \quad \langle \wp(\mathfrak{A}), \bar{M} \rangle \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all sets \bar{P} in \mathfrak{A} , where

$$M_k := \{ S \subseteq A \mid |S| \equiv k \pmod{m} \}.$$

(b) *For every FO[$\Sigma + \{\subseteq, \bar{M}\}$]-formula $\varphi(\bar{x})$, there exists a CMSO[Σ]-formula $\varphi'(\bar{X})$ such that*

$$\langle \wp(\mathfrak{A}), \bar{M} \rangle \models \varphi(\bar{P}) \quad \text{iff} \quad \mathfrak{A} \models \varphi'(\bar{P}),$$

for all Σ -structures \mathfrak{A} and all sets \bar{P} in \mathfrak{A} , where the predicates $\bar{M} = (M_k)_{k < m}$ are defined as above.

Finally, we can also relate MSO to GSO via a suitable operation.

Definition 2.5. Let Σ be a relational signature. The *incidence structure* of a Σ -structure \mathfrak{A} is the 2-sorted Σ_{in} -structure

$$\mathfrak{A}_{\text{in}} := \langle A, E, (P_c)_c, \text{in}_o, \text{in}_i, \dots \rangle$$

with domains A and

$$E := \{ \bar{c} \in A^{<\omega} \mid \bar{c} \text{ guarded, all components distinct} \}.$$

For every relation $R \in \Sigma$ of arity n and every surjective monotone function $\sigma : [n] \rightarrow [k]$ with $k \leq n$, we have a unary predicate

$$P_{R,\sigma} := \{ \bar{c} \in E \mid \bar{d} \in R^{\mathfrak{A}}, d_i = c_{\sigma(i)} \},$$

containing all tuples that are guarded by some tuple in $R^{\mathfrak{A}}$. In addition, for every $k < \omega$, there is an *incidence relation*

$$\text{in}_k := \{ \langle a, \bar{c} \rangle \in A \times E \mid c_k = a \}.$$

Example. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph. Then

$$\mathfrak{G}_{\text{in}} = \langle V, E', (P_{E,\sigma})_\sigma, \text{in}_o, \text{in}_i \rangle$$

where

$$E' := V \cup \{ \langle u, v \rangle \mid u \neq v, \langle u, v \rangle \in E \text{ or } \langle v, u \rangle \in E \},$$

$$P_{E,\sigma} := \{ \langle u, v \rangle \mid u \neq v, \langle u, v \rangle \in E \}, \quad \sigma : [2] \rightarrow [2],$$

$$P_{E,\tau} := \{ \langle v \rangle \mid \langle v, v \rangle \in E \}, \quad \tau : [2] \rightarrow [1].$$

Let us check that GSO over Σ -structures corresponds to MSO over their incidence structures.

Proposition 2.6. *Let Σ be a finite relational signature.*

(a) For every $\text{GSO}[\Sigma]$ -formula φ , there is an $\text{MSO}[\Sigma_{\text{in}}]$ -formula φ' such that

$$\mathcal{A} \models \varphi \quad \text{iff} \quad \mathcal{A}_{\text{in}} \models \varphi', \quad \text{for all } \Sigma\text{-structures } \mathcal{A}.$$

(b) For every $\text{MSO}[\Sigma_{\text{in}}]$ -formula φ , there is a $\text{GSO}[\Sigma]$ -formula φ' such that

$$\mathcal{A}_{\text{in}} \models \varphi \quad \text{iff} \quad \mathcal{A} \models \varphi', \quad \text{for all } \Sigma\text{-structures } \mathcal{A}.$$

Proof. (a) For every n -ary relation variable Z of φ , the formula φ' will use a tuple $\bar{Z} = (Z_\sigma)_\sigma$ of variables that is indexed by all surjective functions $\sigma : [n] \rightarrow [k]$ with $k < \omega$. We define φ' by induction on φ as follows.

$$(Rx_0 \dots x_{n-1})' := \exists z \bigvee_\sigma \left[P_{R,\sigma} z \wedge \bigwedge_{k < n} \text{in}_{\sigma(k)}(x_k, z) \right],$$

$$(Zx_0 \dots x_{n-1})' := \exists z \bigvee_\sigma \left[Z_\sigma z \wedge \bigwedge_{k < n} \text{in}_{\sigma(k)}(x_k, z) \right],$$

$$(x = y)' := x = y, \quad (\exists x \psi)' := \exists x \psi',$$

$$(\varphi \wedge \psi)' := \varphi' \wedge \psi', \quad (\forall x \psi)' := \forall x \psi',$$

$$(\varphi \vee \psi)' := \varphi' \vee \psi', \quad (\exists Z \psi)' := \exists \bar{Z} \psi',$$

$$(\neg \varphi)' := \neg \varphi', \quad (\forall Z \psi)' := \forall \bar{Z} \psi'.$$

(b) We may assume that φ is an MSO° -formula without subformulae of the form $\text{sing}(X)$, $X \cap Y = \emptyset$, or $\text{cover}(\bar{X})$. We denote variables of φ representing a set of elements by X^e and variables representing a set of guarded tuples by X^g . For every variable X^g of φ the formula φ' will use a tuple $\bar{X} = \langle X_1, \dots, X_m \rangle$ where X_n is an n -ary relation variable and m is the maximal arity of a relation symbol in Σ . Variables X^e remain unchanged. We define φ' by induction on φ as follows.

$$(P_{R,\sigma} X^g)' := \exists \bar{x} \exists \bar{y} \left[X_n \bar{x} \wedge R \bar{y} \wedge \bigwedge_i y_i = x_{\sigma(i)} \right],$$

$$\text{where } \sigma : [k] \rightarrow [n],$$

$$(\text{in}_k(X^e, Y^g))' := \bigvee_{n=k}^m \exists \bar{y} \left[Y_n = \{\bar{y}\} \wedge X = \{y_k\} \wedge \bigwedge_{i \neq n} Y_i = \emptyset \right],$$

$$(X^e \subseteq Y^e)' := X \subseteq Y, \quad (X^g \subseteq Y^g)' := \bigwedge_{n=1}^m X_n \subseteq Y_n,$$

$$(\varphi \wedge \psi)' := \varphi' \wedge \psi', \quad (\exists X^e \psi)' := \exists X \psi',$$

$$(\varphi \vee \psi)' := \varphi' \vee \psi', \quad (\forall X^e \psi)' := \forall X \psi',$$

$$(\neg \varphi)' := \neg \varphi', \quad (\exists X^g \psi)' := \exists \bar{X} \psi',$$

$$(\forall X^g \psi)' := \forall \bar{X} \psi'. \quad \square$$

Using these two lemmas we could in theory reduce any question we have about MSO or GSO to FO. In practice this is not always the most convenient thing to do since it does make the underlying structures more complicated. In particular working with power-set structures can be quite unwieldy. Nevertheless the operations \wp and $-\text{in}$ will have many uses throughout this book.

There are many operations that behave like the above ones. This first chapter and, to a somewhat lesser extend, the whole book is devoted to their study. Let us give a name to capture the general situation.

Definition 2.7. Let $\langle L, \mathcal{M}, \models \rangle$ and $\langle L', \mathcal{M}', \models \rangle$ be two logics. A unary operation $f : \mathcal{M} \rightarrow \mathcal{M}'$ is (L, L') -compatible if, for every formula $\varphi' \in L'$, we can effectively compute a formula $\varphi \in L$ such that

$$f(\mathcal{A}) \models \varphi' \quad \text{iff} \quad \mathcal{A} \models \varphi, \quad \text{for every } \mathcal{A} \in \mathcal{M}.$$

We call f (L, L') -bicompatible if, furthermore, for every formula $\varphi \in L$, we can effectively compute a formula $\varphi' \in L'$ such that

$$\mathcal{A} \models \varphi \quad \text{iff} \quad f(\mathcal{A}) \models \varphi', \quad \text{for every } \mathcal{A} \in \mathcal{M}.$$

For the case that $L = L'$ we simply speak of L -compatible and L -bicompatible operations. J

Example. In this terminology, Proposition 2.2 states that the operation \wp is (MSO, FO) -bicompatible, and Proposition 2.6 states that $-\text{in}$ is (GSO, MSO) -bicompatible. J

Note that it follows immediately from the definition that compatible operations compose.

Lemma 2.8. *If f is (L, L') -compatible operation and g (L', L'') -compatible, then $g \circ f$ is (L, L'') -compatible. If f and g are bicompatible, so is $g \circ f$.*

3 Theories and Back-and-Forth Arguments

As already remarked above we are interested in computing the theory of certain structures. The problem is that such theories are infinite objects since our logics have infinitely many different formulae. The usual way around this issue is to write the logic in question as a union of finite sublogics indexed by some complexity parameter. In our case, the standard such parameter is the quantifier-rank of a formula.

Definition 3.1. (a) The *quantifier rank* $qr(\varphi)$ of a formula φ is the number of nested (first-order and second-order) quantifiers in φ .

(b) We denote by $FO_m[\Sigma, X]$, $MSO_m[\Sigma, X]$, ... the corresponding sublogic consisting of all formulae of quantifier-rank at most m . For CMSO, we use a slightly different definition: $CMSO_m$ denotes the sublogic consisting of all formulae of quantifier-rank at most m that only use counting predicates $|X| \cong k \pmod p$ with $p \leq m$. \lrcorner

Example. The formula $\forall x[\exists yRxy \wedge \exists zRzx]$ has quantifier-rank 2 since the two innermost quantifiers are not nested. \lrcorner

Let us formally define our notion of a theory. We also introduce an equivalence relation on structures for ‘having the same theory’, which will be central to much of this book.

Definition 3.2. Let $\langle L, \mathcal{M}, \models \rangle$ be a logic.

(a) The *L-theory* of a model $M \in \mathcal{M}$ is

$$\text{Th}_L(M) := \{ \varphi \in L \mid M \models \varphi \}.$$

If L is a logic for which the notion of a quantifier-rank is defined, we also set

$$\text{Th}_L^m(M) := \{ \varphi \in L \mid qr(\varphi) \leq m, M \models \varphi \}.$$

(b) A set $\Phi \subseteq L$ of formulae is a *theory* if it is of the form $\Phi = \text{Th}_L(M)$, for some $M \in \mathcal{M}$.

(c) For two models M and N , we define

$$M \sqsubseteq_L N \quad : \text{iff} \quad \text{Th}_L(M) \subseteq \text{Th}_L(N),$$

$$M \equiv_L N \quad : \text{iff} \quad \text{Th}_L(M) = \text{Th}_L(N).$$

If $M \equiv_L N$, we call M and N *L-equivalent*. Again, for logics that have a notion of a quantifier-rank, we use the notation

$$M \sqsubseteq_L^m N \quad : \text{iff} \quad \text{Th}_L^m(M) \subseteq \text{Th}_L^m(N),$$

$$M \equiv_L^m N \quad : \text{iff} \quad \text{Th}_L^m(M) = \text{Th}_L^m(N).$$

If the logic L is understood, we will speak of *m-equivalence* in this case. If we want to indicate the logic in question, we will use the terms *first-order m-equivalence*, *monadic m-equivalence*, or *guarded m-equivalence* instead.

(d) The class of *models* of a formula $\varphi \in L$ is the set

$$\text{Mod}(\varphi) := \{ M \in \mathcal{M} \mid M \models \varphi \}.$$

A class $\mathcal{C} \subseteq \mathcal{M}$ is *L-definable* if $\mathcal{C} = \text{Mod}(\varphi)$, for some $\varphi \in L$.

(e) A logic L is *lattice closed* if the collection of all *L-definable* classes is closed under finite intersections and unions. \lrcorner

Remark. In the common case where the logics in question are closed under negation, the relations \sqsubseteq_L and \equiv_L coincide. The more general definitions above are only needed to support logics like, e.g., existential first-order logic that are not closed under negation. \lrcorner

We can use *L-equivalence* to give a simple but useful conditions for when a class of models is definable.

Lemma 3.3. *Let $\langle L, \mathcal{M}, \models \rangle$ be a lattice-closed logic. A class $\mathcal{C} \subseteq \mathcal{M}$ is L -definable if, and only if, there exists a finite subset $\Delta \subseteq L$ such that*

$$M \in \mathcal{C} \quad \text{and} \quad M \sqsubseteq_{\Delta} N \quad \text{implies} \quad N \in \mathcal{C}.$$

Proof. (\Rightarrow) Let $\varphi \in L$ be a formula defining \mathcal{C} and set $\Delta := \{\varphi\}$. Suppose that $M \in \mathcal{C}$ and $M \sqsubseteq_{\Delta} N$. Then $M \models \varphi$, which implies that $N \models \varphi$. Hence, $N \in \mathcal{C}$.

(\Leftarrow) Set

$$\varphi := \bigvee \left\{ \bigwedge \text{Th}_{\Delta}(M) \mid M \in \mathcal{C} \right\}.$$

Note that this disjunction is finite since there are only finitely many subsets of Δ . For $N \in \mathcal{M}$, it follows that

$$\begin{aligned} N \models \varphi & \quad \text{iff} \quad N \models \bigwedge \text{Th}_{\Delta}(M), \quad \text{for some } M \in \mathcal{C} \\ & \quad \text{iff} \quad M \sqsubseteq_{\Delta} N, \quad \text{for some } M \in \mathcal{C} \\ & \quad \text{iff} \quad N \in \mathcal{C}. \end{aligned} \quad \square$$

As explained above, the reason why we consider bounded-quantifier theories is that they are finite objects that can be manipulated algorithmically. Let us prove this fact.

Proposition 3.4. *Let L be one of the logics defined above, Σ a finite relational signature, $k, m, r < \omega$, and let $L_{m,r}^k[\Sigma]$ be the set of all $L[\Sigma]$ -formulae of quantifier-rank at most m with at most k free variables and such that all constants i, p appearing in counting predicates $|X| \cong i \pmod{p}$ are bounded by r .*

- (a) *Up to logical equivalence, there are only finitely many $L_{m,r}^k[\Sigma]$ -formulae. Furthermore, given $m, k, r < \omega$, we can compute a finite set $\Phi_{m,r}^k$ of $L_{m,r}^k[\Sigma]$ -formulae such that every $L_{m,r}^k[\Sigma]$ -formula is equivalent to one in $\Phi_{m,r}^k$.*
- (b) *There are only finitely many $L_{m,r}^k[\Sigma]$ -theories.*
- (c) *For every $L_{m,r}^k[\Sigma]$ -theory T , there exists a single $L_{m,r}^k[\Sigma]$ -formula θ such that*

$$\mathfrak{A} \models \theta \quad \text{iff} \quad \text{Th}_{L_{m,r}^k}(\mathfrak{A}) \subseteq T.$$

Remark. Note that, in (a), it is in general undecidable which formula in $\Phi_{m,r}^k$ is equivalent to a given $L_{m,r}^k[\Sigma]$ -formula. We only know there is at least one. J

Proof. (a) We prove the claim by induction on m . First, consider the case where $m = 0$. Every quantifier-free L -formula can be written in disjunctive normal form. Since the signature, the number of variables, and the number of possible counting predicates $|X| \cong i \pmod{p}$ are all finite, there are only finitely many atomic formulae and only finitely many negated atomic formulae. (For $L = \text{MSO}^0$, we may assume w.l.o.g. that all variables in an atom cover(\bar{X}) are distinct.) Since, up to logical equivalence, a conjunction of such formulae is uniquely determined by the set of formulae appearing in it, it follows that there are only finitely many such conjunctions. In the same way, we see that, up to logical equivalence, there are only finitely many disjunctions of such conjunctions. Hence, there are only finitely many quantifier-free formulae in disjunctive normal form.

For the inductive step, suppose that $m > 0$. As above, every L -formula of quantifier-rank at most m can be written as a boolean combination of (i) atomic formulae and (ii) formulae of the form $\exists x\psi$ or $\exists X\psi$ with $\text{qr}(\psi) < m$. By inductive hypothesis, there are only finitely many formulae of these two forms. Writing the boolean combination of them in disjunctive normal form, we can use the same argument as above to show that there are only finitely many such combinations.

For the desired set $\Phi_{m,r}^k$ of representatives, we can take the set of all formulae in disjunctive normal form (without repetitions) built up from subformulae of the form $\exists x\psi$ and $\exists X\psi$ with $\psi \in \Phi_{m-1,r}^{k+1}$.

(b) By (a), we can fix a finite set $\Phi_{m,r}^k$ of $L[\Sigma]$ -formulae such that every $L_{m,r}^k[\Sigma]$ -formula is equivalent to one in $\Phi_{m,r}^k$. Then every $L_{m,r}^k[\Sigma]$ -theory T is uniquely determined by the intersection $T \cap \Phi_{m,r}^k$. Since there are only finitely many sets of the form $T \cap \Phi_{m,r}^k$, the number of theories is finite.

(c) By (a), we can compute a finite set $\Phi_{m,r}^k$ of $L_{m,r}^k[\Sigma]$ -formulae such that every $L_{m,r}^k[\Sigma]$ -formula is equivalent to one in $\Phi_{m,r}^k$. For every $L_{m,r}^k[\Sigma]$ -

theory T , the formula

$$\theta := \bigwedge (T \cap \Phi_{m,r}^k)$$

has the desired property. \square

Example. Inspecting the proofs of Propositions 2.2 and 2.3, we see that the quantifier-rank does not increase during the translation (if we work with MSO° -formulae). But the translation in Proposition 2.6 does introduce new quantifiers, although only a bounded number of them that depends only on the maximal arity of a relation in the signature. Consequently, we can restate these three lemmas in our new terminology as follows.

$$\begin{aligned} \mathfrak{A} \equiv_{\text{MSO}^\circ}^m \mathfrak{B} & \quad \text{iff} \quad \wp(\mathfrak{A}) \equiv_{\text{FO}}^m \wp(\mathfrak{B}), \\ \mathfrak{A} \equiv_{\text{WMSO}^\circ}^m \mathfrak{B} & \quad \text{iff} \quad \wp_{\text{fin}}(\mathfrak{A}) \equiv_{\text{FO}}^m \wp_{\text{fin}}(\mathfrak{B}), \\ \mathfrak{A} \equiv_{\text{GSO}}^m \mathfrak{B} & \quad \Rightarrow \quad \mathfrak{A}_{\text{in}} \equiv_{\text{MSO}^\circ}^{m+k} \mathfrak{B}_{\text{in}}, \\ \mathfrak{A}_{\text{in}} \equiv_{\text{MSO}^\circ}^m \mathfrak{B}_{\text{in}} & \quad \Rightarrow \quad \mathfrak{A} \equiv_{\text{GSO}}^{m+l} \mathfrak{B}, \end{aligned}$$

for some constants k, l depending on the signature, but independent of m . \downarrow

The most basic way to compute the theory of a structure is by induction on the quantifier rank. When doing so, we have to deal with formulae with free variables. We call the values provided for these variables *parameters*.

Definition 3.5. Let \mathfrak{A} be a structure. A *parameter* in \mathfrak{A} is a value α that can be assigned to a variable. There are three kinds of parameters:

- ♦ *first-order parameters* are elements $\alpha \in A$;
- ♦ *monadic parameters* are sets $\alpha \subseteq A$; and
- ♦ *guarded parameters* are guarded relations $\alpha \subseteq A^n$.

We say that α is a

- ♦ *second-order parameter*, if it is a monadic parameter or a guarded parameter;

- ♦ *FO-parameter*, if it is a first-order parameter;
- ♦ *MSO $^\circ$ -parameter*, if it is a monadic parameter;
- ♦ *MSO-parameter*, if it is a first-order parameter or a monadic parameter;
- ♦ *GSO-parameter*, if it is a first-order parameter, a monadic parameter, or a guarded parameter. \downarrow

Our main tool to inductively compute a theory consists in the following kind of argument.

Proposition 3.6. Let L be one of the logics defined above, Σ a finite signature, \mathfrak{A} and \mathfrak{B} Σ -structures with L -parameters $\tilde{\alpha}$ and $\tilde{\beta}$, and $m < \omega$. Then

$$\mathfrak{A}, \tilde{\alpha} \equiv_L^{m+1} \mathfrak{B}, \tilde{\beta}$$

if, and only if, the following two properties are satisfied.

(Forth Property) For every L -parameter α' in \mathfrak{A} , there exists an L -parameter β' in \mathfrak{B} such that

$$\mathfrak{A}, \tilde{\alpha}\alpha' \equiv_L^m \mathfrak{B}, \tilde{\beta}\beta'.$$

(Back Property) For every L -parameter β' in \mathfrak{B} , there exists an L -parameter α' in \mathfrak{A} such that

$$\mathfrak{A}, \tilde{\alpha}\alpha' \equiv_L^m \mathfrak{B}, \tilde{\beta}\beta'.$$

Proof. (\Leftarrow) Suppose that both properties are satisfied. We have to show that

$$\mathfrak{A} \models \varphi(\tilde{\alpha}) \quad \text{iff} \quad \mathfrak{B} \models \varphi(\tilde{\beta}),$$

for all L -formulae $\varphi(\tilde{x})$ of quantifier rank at most $m+1$. Every such formula is a boolean combination of formulae of the form $\exists x' \psi(\tilde{x}, x')$ where $\text{qr}(\psi) \leq m$ and the variable x' is either first-order, monadic, or guarded. Therefore, it

is sufficient to prove the claim for such formulae. By symmetry, it is further enough to prove that

$$\mathfrak{A} \models \exists x' \psi(\bar{\alpha}, x') \quad \text{implies} \quad \mathfrak{B} \models \exists x' \psi(\bar{\beta}, x'),$$

for all $\psi(\bar{x}, x')$ with $\text{qr}(\psi) \leq m$. Hence, suppose that $\mathfrak{A} \models \exists x' \psi(\bar{\alpha}, x')$. Then there exists an L -parameter α' in \mathfrak{A} such that $\mathfrak{A} \models \psi(\bar{\alpha}, \alpha')$. By assumption, we can find an L -parameter β' in \mathfrak{B} with

$$\mathfrak{A}, \bar{\alpha} \alpha' \equiv_L^m \mathfrak{B}, \bar{\beta} \beta'.$$

Consequently, $\mathfrak{B} \models \psi(\bar{\beta}, \beta')$, which implies that $\mathfrak{B} \models \exists x' \psi(\bar{\beta}, x')$.

(\Rightarrow) By Proposition 3.4 (a), there exists a finite set Φ of L -formulae of quantifier rank at most m such that every formula of quantifier rank at most m is equivalent to some formula in Φ .

Suppose that there exists an L -parameter α' in \mathfrak{A} such that

$$\mathfrak{A}, \bar{\alpha} \alpha' \not\equiv_L^m \mathfrak{B}, \bar{\beta} \beta', \quad \text{for all } \beta' \text{ in } B.$$

We have to prove that $\mathfrak{A}, \bar{\alpha} \not\equiv_L^{m+1} \mathfrak{B}, \bar{\beta}$. Set

$$\Theta := \{ \psi(\bar{x}, x') \in \Phi \mid \mathfrak{A} \models \psi(\bar{\alpha}, \alpha') \},$$

and let $\vartheta := \bigwedge \Theta$ be the conjunction of all formulae in Θ . It is sufficient to show that

$$\mathfrak{A} \models \exists x' \vartheta(\bar{\alpha}, x') \quad \text{and} \quad \mathfrak{B} \not\models \exists x' \vartheta(\bar{\beta}, x').$$

Since $\mathfrak{A} \models \vartheta(\bar{\alpha}, \alpha')$, we have $\mathfrak{A} \models \exists x' \vartheta(\bar{\alpha}, x')$. Furthermore, for every β' in B , we have

$$\mathfrak{A}, \bar{\alpha} \alpha' \not\equiv_L^m \mathfrak{B}, \bar{\beta} \beta'.$$

Hence, there exists a formula $\eta(\bar{x}, x')$ of quantifier rank at most m such that

$$\mathfrak{A} \models \eta(\bar{\alpha}, \alpha') \quad \text{and} \quad \mathfrak{B} \not\models \eta(\bar{\beta}, \beta').$$

By choice of the set Φ , we may choose $\eta \in \Phi$. Hence, $\eta \in \Theta$ and

$$\mathfrak{B} \not\models \eta(\bar{\beta}, \beta') \quad \text{implies} \quad \mathfrak{B} \not\models \vartheta(\bar{\beta}, \beta').$$

We have shown that

$$\mathfrak{B} \not\models \vartheta(\bar{\beta}, \beta'), \quad \text{for every } L\text{-parameter } \beta' \text{ in } \mathfrak{B}.$$

Hence, $\mathfrak{B} \not\models \exists x' \vartheta(\bar{\beta}, x')$. \square

Exercise 3.1. Let $m < \omega$ and suppose that $\bar{a}, \bar{b} \in \mathbb{Q}^n$ are tuples such that

$$a_i \leq a_j \iff b_i \leq b_j, \quad \text{for all } i, j < n.$$

Prove that $\langle \mathbb{Q}, \leq \rangle, \bar{a} \equiv_{\text{FO}}^m \langle \mathbb{Q}, \leq \rangle, \bar{b}$.

Exercise 3.2. Find MSO-formulae distinguishing the following structures:

$$\mathfrak{N} := \langle \mathbb{N}, \leq \rangle, \quad \mathfrak{Z} := \langle \mathbb{Z}, \leq \rangle, \quad \mathfrak{Q} := \langle \mathbb{Q}, \leq \rangle, \quad \mathfrak{R} := \langle \mathbb{R}, \leq \rangle.$$

Which of these structures are FO-equivalent?

Exercise 3.3. Let $\mathfrak{A} = \langle A, \leq \rangle$ and $\mathfrak{B} = \langle B, \leq \rangle$ be finite linear orders and $m < \omega$. Prove that

$$\mathfrak{A} \equiv_{\text{FO}}^m \mathfrak{B} \quad \text{iff} \quad |A| = |B| \quad \text{or} \quad |A|, |B| \geq 2^m - 1.$$

4 Operations for Monadic Second-Order Logic

In non-trivial cases the complexities involved in carrying out a back-and-forth argument quickly become unmanageable. Instead of computing a theory from first principles, it is often easier to preform a reduction to another theory which is already known. This approach is known as the *composition method*. In this section and the next one, we present several operations on structures that can be used for such reductions and we establish so-called *composition theorems* for them: statements to the effect that they are compatible with the logic in question. We start with operations that are compatible with MSO. Those compatible with FO we defer to the next section.

Disjoint Unions

One of the most basic operations, but a surprisingly versatile one, is that of forming a disjoint union.

Definition 4.1. Let Σ be a relational signature. The *disjoint union* of two Σ -structures \mathfrak{A} and \mathfrak{B} is the $(\Sigma + \{\text{Left}, \text{Right}\})$ -structure $\mathfrak{A} \oplus \mathfrak{B}$ with universe $A + B$ and relations

$$\begin{aligned} R^{\mathfrak{A} \oplus \mathfrak{B}} &:= R^{\mathfrak{A}} + R^{\mathfrak{B}}, \quad \text{for } R \in \Sigma, \\ \text{Left}^{\mathfrak{A} \oplus \mathfrak{B}} &:= A, \\ \text{Right}^{\mathfrak{A} \oplus \mathfrak{B}} &:= B. \end{aligned}$$

The composition theorem for disjoint unions reads as follows.

Proposition 4.2. Let L be one of the logics defined above, let Σ be a finite relational signature, $\mathfrak{A}, \mathfrak{A}', \mathfrak{B}$ and \mathfrak{B}' Σ -structures with first-order parameters $\bar{a}, \bar{a}', \bar{b}, \bar{b}'$ and second-order parameters $\bar{P}, \bar{P}', \bar{Q}, \bar{Q}'$, respectively, and let $m < \omega$. Then

$$\mathfrak{A}, \bar{P}, \bar{a} \equiv_L^m \mathfrak{A}', \bar{P}', \bar{a}' \quad \text{and} \quad \mathfrak{B}, \bar{Q}, \bar{b} \equiv_L^m \mathfrak{B}', \bar{Q}', \bar{b}'$$

implies

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q}, \bar{a} \bar{b} \equiv_L^m \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}', \bar{a}' \bar{b}'.$$

(We write $\bar{P} \cup \bar{Q}$ for the tuple whose i -th component is $P_i \cup Q_i$. We assume that the parameters are appropriate for the logic L , i.e., if $L = \text{MSO}^\circ$ there are no first-order parameters and if $L = \text{FO}$, there are only first-order parameters.)

Proof. We prove the claim by induction on m . First, consider the case where $m = 0$. Since quantifier-free formulae are boolean combinations of atomic formulae, it is sufficient to consider such formulae. By symmetry, we therefore only need to show that

$$\mathfrak{A} \oplus \mathfrak{B} \models \varphi(\bar{P} \cup \bar{Q}, \bar{a}, \bar{b}) \quad \text{implies} \quad \mathfrak{A}' \oplus \mathfrak{B}' \models \varphi(\bar{P}' \cup \bar{Q}', \bar{a}', \bar{b}'),$$

for every atomic formula $\varphi(\bar{X}, \bar{x}, \bar{y})$. We distinguish several cases. (I encourage the reader to skip most of them.)

If φ is an equality $z = z'$ and $\mathfrak{A} \oplus \mathfrak{B} \models \varphi(\bar{P} \cup \bar{Q}, \bar{a}, \bar{b})$, then the variables z and z' are either both among the \bar{x} or among the \bar{y} . By symmetry, we may assume the former, i.e., $\varphi = (x_i = x_j)$. Then $a_i = a_j$ and

$$\mathfrak{A}, \bar{P}, \bar{a} \equiv_L^0 \mathfrak{A}', \bar{P}', \bar{a}'$$

implies that $a'_i = a'_j$. Hence, $\mathfrak{A}' \oplus \mathfrak{B}' \models \varphi(\bar{P}' \cup \bar{Q}', \bar{a}', \bar{b}')$.

If $\varphi = R\bar{z}$ for $R \in \Sigma$, then again \bar{z} must be a subtuple of \bar{x} or of \bar{y} . Say it is the former. Then $\mathfrak{A}, \bar{P}, \bar{a} \models R\bar{z}$ and

$$\mathfrak{A}, \bar{P}, \bar{a} \equiv_L^0 \mathfrak{A}', \bar{P}', \bar{a}'$$

implies that $\mathfrak{A}', \bar{P}', \bar{a}' \models R\bar{z}$. Hence, $\mathfrak{A}' \oplus \mathfrak{B}' \models \varphi(\bar{P}' \cup \bar{Q}', \bar{a}', \bar{b}')$.

For $\varphi = \text{Left}(x)$ or $\varphi = \text{Right}(x)$, the proof is similar.

If $\varphi = X_i \subseteq X_j$, then

$$\begin{aligned} \mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q} \models X_i \subseteq Y_j & \\ \Rightarrow P_i \cup Q_i \subseteq P_j \cup Q_j & \\ \Rightarrow P_i \subseteq P_j \text{ and } Q_i \subseteq Q_j & \\ \Rightarrow P'_i \subseteq P'_j \text{ and } Q'_i \subseteq Q'_j & \\ \Rightarrow P'_i \cup Q'_i \subseteq P'_j \cup Q'_j & \\ \Rightarrow \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}' \models X_i \subseteq Y_j. & \end{aligned}$$

The proofs for $\varphi = X \cap Y = \emptyset$ and $\text{cover}(\bar{X})$ are analogous.

Suppose that $\varphi = RX_0 \dots X_{n-1}$, i.e., that

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q} \models R\bar{X}.$$

Then there are elements $a_i \in P_i \cup Q_i$ such that $\bar{a} \in R^{\mathfrak{A} \oplus \mathfrak{B}}$. Since $R^{\mathfrak{A} \oplus \mathfrak{B}} = R^{\mathfrak{A}} \cup R^{\mathfrak{B}}$, it follows that $\bar{a} \in R^{\mathfrak{A}}$ or $\bar{a} \in R^{\mathfrak{B}}$. By symmetry, we may assume the former. Then $\mathfrak{A}, \bar{P} \models R\bar{X}$, which implies that $\mathfrak{A}', \bar{P}' \models R\bar{X}$. It follows that

$$\mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}' \models R\bar{X}.$$

For $\varphi = \text{Left}(X)$ or $\varphi = \text{Right}(X)$, the proof is similar.

If $\varphi = |X_i| < \infty$, then

$$\begin{aligned} \mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q} &\models |X_i| < \infty \\ \Rightarrow P_i \cup Q_i &\text{ is finite} \\ \Rightarrow P_i &\text{ is finite and } Q_i \text{ is finite} \\ \Rightarrow P'_i &\text{ is finite and } Q'_i \text{ is finite} \\ \Rightarrow P'_i \cup Q'_i &\text{ is finite} \\ \Rightarrow \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}' &\models |X_i| < \infty. \end{aligned}$$

Finally, suppose that $\varphi = |X_i| \equiv k \pmod{p}$. then

$$\begin{aligned} \mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q} &\models |X_i| \equiv k \pmod{p} \\ \Rightarrow |P_i \cup Q_i| &\equiv k \pmod{p} \\ \Rightarrow |P_i| \equiv k_1 \pmod{p} \quad \text{and} \quad |Q_i| \equiv k_2 \pmod{p} \\ &\quad \text{with } k_1 + k_2 = k \\ \Rightarrow |P'_i| \equiv k_1 \pmod{p} \quad \text{and} \quad |Q'_i| \equiv k_2 \pmod{p} \\ &\quad \text{with } k_1 + k_2 = k \\ \Rightarrow |P'_i \cup Q'_i| &\equiv k \pmod{p} \\ \Rightarrow \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}' &\models |X_i| \equiv k \pmod{p}. \end{aligned}$$

For the inductive step, suppose that we have already established the claim for m , and consider structures

$$\mathfrak{A}, \bar{P}, \bar{a} \equiv_L^{m+1} \mathfrak{A}', \bar{P}', \bar{a}' \quad \text{and} \quad \mathfrak{B}, \bar{Q}, \bar{b} \equiv_L^{m+1} \mathfrak{B}', \bar{Q}', \bar{b}'.$$

We have to show that

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q}, \bar{a}\bar{b} \equiv_L^{m+1} \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}', \bar{a}'\bar{b}'.$$

By symmetry and Proposition 3.6, it is sufficient to prove that, for every L -parameter α in $\mathfrak{A} \oplus \mathfrak{B}$, there exists an L -parameter α' in $\mathfrak{A}' \oplus \mathfrak{B}'$ such

that

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q}, \bar{a}\bar{b}, \alpha \equiv_L^m \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}', \bar{a}'\bar{b}', \alpha'.$$

Hence, consider a parameter α in $\mathfrak{A} \oplus \mathfrak{B}$. We distinguish two cases.

If α is a first-order parameter, then $\alpha \in A$ or $\alpha \in B$. By symmetry, we may assume the former. According to Proposition 3.6, we can find an L -parameter α' in \mathfrak{A}' such that

$$\mathfrak{A}, \bar{P}, \bar{a}\alpha \equiv_L^m \mathfrak{A}', \bar{P}', \bar{a}'\alpha'.$$

By inductive hypothesis, this implies that

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q}, \bar{a}\bar{b}\alpha \equiv_L^m \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}', \bar{a}'\bar{b}'\alpha',$$

as desired.

Hence, suppose that α is a second-order parameter. According to Proposition 3.6, we can find parameters α'_0 in \mathfrak{A}' and α'_1 in \mathfrak{B}' such that

$$\mathfrak{A}, \bar{P}, \bar{a}, \alpha \upharpoonright A \equiv_L^m \mathfrak{A}', \bar{P}', \bar{a}', \alpha'_0$$

$$\text{and } \mathfrak{B}, \bar{Q}, \bar{b}, \alpha \upharpoonright B \equiv_L^m \mathfrak{B}', \bar{Q}', \bar{b}', \alpha'_1.$$

By inductive hypothesis, this implies that

$$\mathfrak{A} \oplus \mathfrak{B}, \bar{P} \cup \bar{Q}, \bar{a}\bar{b}, \alpha \equiv_L^m \mathfrak{A}' \oplus \mathfrak{B}', \bar{P}' \cup \bar{Q}', \bar{a}'\bar{b}', \alpha'_0 \cup \alpha'_1.$$

Hence, we can set $\alpha' := \alpha'_0 \cup \alpha'_1$. □

Let us present some applications. We start with structures over the empty signature.

Proposition 4.3. *Let \mathfrak{A} and \mathfrak{B} be structures over the empty signature. Then*

- (a) $\mathfrak{A} \equiv_{\text{FO}}^m \mathfrak{B}$ iff $|A| = |B|$ or $|A|, |B| \geq m$.
- (b) $\mathfrak{A} \equiv_{\text{MSO}}^{m+1} \mathfrak{B}$ iff $|A| = |B|$ or $|A|, |B| \geq 2^m$.

Proof. (a) (\Rightarrow) Suppose that $k := |A| < |B|$ with $k < m$. Then the formula

$$\exists x_0 \cdots \exists x_k \bigwedge_{0 \leq i < j \leq k} x_i \neq x_j$$

holds in \mathfrak{B} , but not in \mathfrak{A} . It quantifier-rank is $k + 1 \leq m$.

(\Leftarrow) Clearly, if $|A| = |B|$ then $\mathfrak{A} \cong \mathfrak{B}$ and both structures have the same theory. Hence, it remains to consider the case where $|A|, |B| \geq m$. We have to show that $\mathfrak{A} \equiv_{\text{FO}}^m \mathfrak{B}$. For $m = 0$ the claim is trivial, since there are no quantifier-free $\text{FO}[\emptyset]$ -formulae without free variables. Hence, we may assume that $m > 0$. In this case, according to Proposition 3.6, it is enough to prove the Back-and-Forth Property. As usually it is further sufficient to consider only the Forth Property. The Back Property then follows by symmetry. Hence, let $a \in A$. We pick an arbitrary element $b \in B$. Let \mathfrak{A}_0 and \mathfrak{B}_0 be the substructures of \mathfrak{A} and \mathfrak{B} induced by the sets $A_0 := A \setminus \{a\}$ and $B_0 := B \setminus \{b\}$, and let \mathfrak{A}_1 and \mathfrak{B}_1 be the substructures induced by $\{a\}$ and $\{b\}$. By inductive hypothesis,

$$|A_0|, |B_0| \geq m - 1 \quad \text{implies} \quad \mathfrak{A}_0 \equiv_{\text{FO}}^{m-1} \mathfrak{B}_0.$$

Since $\mathfrak{A}_1 \cong \mathfrak{B}_1$, we also have $\mathfrak{A}_1, a \equiv_{\text{FO}}^{m-1} \mathfrak{B}_1, b$. Consequently, it follows by Proposition 4.2 that

$$\mathfrak{A}, a \cong \mathfrak{A}_0 \oplus \mathfrak{A}_1, a \equiv_{\text{FO}}^{m-1} \mathfrak{B}_0 \oplus \mathfrak{B}_1, b \cong \mathfrak{B}, b.$$

(Strictly speaking, instead of the disjoint unions $\mathfrak{A}_0 \oplus \mathfrak{A}_1$ and $\mathfrak{B}_0 \oplus \mathfrak{B}_1$, we have to take their reducts that omit the new relations Left and Right.)

(b) (\Rightarrow) Suppose that $k := |A| < |B|$ with $k < 2^m$ and let n be the largest number such that $2^n \leq k$. For $i \leq 2^{n+1}$, set

$$\vartheta_i(X_0, \dots, X_n) := \exists y \bigwedge_{0 \leq j \leq n} \chi_{i,j}(\vec{X}, y)$$

where

$$\chi_{i,j}(\vec{X}, y) := \begin{cases} X_j y & \text{if the } j\text{-th bit of } i \text{ is } 1, \\ \neg X_j y & \text{if the } j\text{-th bit of } i \text{ is } 0. \end{cases}$$

The formula

$$\exists X_0 \cdots \exists X_n \bigwedge_{0 \leq i \leq k} \vartheta_i$$

holds in \mathfrak{B} , but not in \mathfrak{A} . It has quantifier rank $n + 2 \leq m + 1$.

(\Leftarrow) Again it is sufficient to consider the case where $|A|, |B| \geq 2^m$. We prove the claim by induction on m . First, suppose that $m = 0$. Every formula φ of quantifier rank 1 contains only one bound variable. As the signature is empty, it follows that the only atomic formulae appearing in φ are of the form $x = x$. Consequently, φ either states that the structure is non-empty, or that it is empty. Since \mathfrak{A} and \mathfrak{B} are both non-empty, such formulae therefore hold either in both structures, or in none of them.

For the inductive step, suppose that $m > 0$. Again it is sufficient to check the Forth Property. We distinguish two cases, depending on whether we deal with a first-order parameter or with a monadic one. First, consider a monadic parameter $P \subseteq A$. If $P = \emptyset$, we choose $Q := \emptyset$. If $P = A$, we choose $Q := B$. In both cases it follows by inductive hypothesis that

$$\mathfrak{A}, P \equiv_{\text{MSO}}^m \mathfrak{B}, Q.$$

Hence, we may assume that P is neither empty, nor all of A . If $|P| \leq 2^{m-1}$, choose a subset $Q \subseteq B$ of size $|Q| = |P|$. Otherwise, choose a subset $Q \subseteq B$ with $|B \setminus Q| = |A \setminus P|$. Let \mathfrak{A}_0 and \mathfrak{B}_0 be the substructures of \mathfrak{A} and \mathfrak{B} induced by P and Q , and let \mathfrak{A}_1 and \mathfrak{B}_1 be the substructures induced by $A \setminus P$ and $B \setminus Q$. It follows that

- ♦ $|P| = |Q|$ or $|P|, |Q| \geq 2^{m-1}$;
- ♦ $|A \setminus P| = |B \setminus Q|$ or $|A \setminus P|, |B \setminus Q| \geq 2^{m-1}$.

By inductive hypothesis, this implies that

$$\mathfrak{A}_0 \equiv_{\text{MSO}}^m \mathfrak{B}_0 \quad \text{and} \quad \mathfrak{A}_1 \equiv_{\text{MSO}}^m \mathfrak{B}_1.$$

By Proposition 4.2, it follows that

$$\mathfrak{A}, P \cong \mathfrak{A}_0, P \oplus \mathfrak{A}_1, \emptyset \equiv_{\text{MSO}}^m \mathfrak{B}_0, Q \oplus \mathfrak{B}_1, \emptyset \cong \mathfrak{B}, Q.$$

(Again, we have to omit the relations Left and Right.)

For a first-order parameter $a \in A$, we choose an arbitrary element $b \in B$. We denote by \mathfrak{A}_o and \mathfrak{B}_o the substructures of \mathfrak{A} and \mathfrak{B} induced by $\{a\}$ and $\{b\}$, and we write \mathfrak{A}_1 and \mathfrak{B}_1 for the substructures induced by $A \setminus \{a\}$ and $B \setminus \{b\}$. Then

$$\mathfrak{A}_o, a \cong \mathfrak{B}_o, b \quad \text{implies} \quad \mathfrak{A}_o, a \equiv_{\text{MSO}}^m \mathfrak{B}_o, b.$$

Furthermore, it follows by inductive hypothesis that

$$\mathfrak{A}_1 \equiv_{\text{MSO}}^m \mathfrak{B}_1.$$

By Proposition 4.2, this implies that

$$\mathfrak{A}, a \cong \mathfrak{A}_o, a \oplus \mathfrak{A}_1 \equiv_{\text{MSO}}^m \mathfrak{B}_o, b \oplus \mathfrak{B}_1 \cong \mathfrak{B}, b. \quad \square$$

Example. There is no $\text{MSO}[\Sigma]$ -formula φ such that, for every finite Σ -structure \mathfrak{A} ,

$$\mathfrak{A} \models \varphi \quad \text{iff} \quad |A| \text{ is even.}$$

For the proof, let $m := \text{qr}(\varphi)$ and let \mathfrak{A} and \mathfrak{B} be Σ -structures of size 2^m and $2^m + 1$, respectively, where every relation is empty. By Proposition 4.3, we have $\mathfrak{A} \equiv_{\text{MSO}}^{m+1} \mathfrak{B}$. Consequently,

$$\mathfrak{A} \models \varphi \quad \text{iff} \quad \mathfrak{B} \models \varphi.$$

A contradiction. J

Exercise 4.1. We consider structures of the form $\mathfrak{A} = \langle A, E \rangle$ where E is an equivalence relation. For an equivalence relation E , we denote by $N_k^-(E)$ the number of E -classes $[a]_E$ of size $|[a]_E| = k$ and $N_k^+(E)$ denotes the number of classes of size $|[a]_E| > k$. We write $m =_k n$ iff $m = n$ or $m, n \geq k$.

Let E and F be equivalence relations on the sets A and B , respectively. Prove that $\langle A, E \rangle \equiv_{\text{FO}}^m \langle B, F \rangle$ if, and only if, for all $k \leq m$,

$$N_k^-(E) =_{m-k} N_k^-(F) \quad \text{and} \quad N_k^+(E) =_{m-k} N_k^+(F). \quad \text{J}$$

Interpretations

Disjoint unions alone are not that interesting as they cannot be used to modify the relations of a structure. The next operation, called an *interpretation*, fills that hole. We will present the definition for of many-sorted structures since this more general case is what is needed in Chapter VII below.

Definition 4.4. Let L be one of the logics FO, MSO, WMSO, or CMSO, let Σ and Γ be relational signatures, and assume that Γ is S -sorted. An L -interpretation from Σ to Γ is an operation τ transforming Σ -structures into Γ -structures that is defined by a list

$$\langle (\delta_s(x))_{s \in S}, (\varphi_R(\vec{x}))_{R \in \Gamma} \rangle$$

of L -formulae over the signature Σ as follows. We assume that the formulae δ_s have one free variable, while the number of free variables of φ_R matches the arity of R . Then τ maps a Σ -structure \mathfrak{A} to the Γ -structure

$$\tau(\mathfrak{A}) := \langle (\delta_s^{\mathfrak{A}})_s, (\varphi_R^{\mathfrak{A}})_{R \in \Gamma} \rangle$$

whose domain of sort s is the set

$$\delta_s^{\mathfrak{A}} := \{ a \in A \mid \mathfrak{A} \models \delta_s(a) \}$$

defined by δ_s and whose relations are

$$\varphi_R^{\mathfrak{A}} := \{ \vec{a} \mid \mathfrak{A} \models \varphi_R(\vec{a}) \}, \quad \text{for } R \in \Sigma.$$

We call the list $\langle (\delta_s)_s, (\varphi_R)_{R \in \Gamma} \rangle$ the *definition scheme* of τ . The *quantifier rank* of τ is the maximal quantifier rank of a formula in its definition scheme. J

Let us show that L -interpretations are L -compatible.

Proposition 4.5. Let L be one of the logics FO, MSO, WMSO, or CMSO, and let $\tau = \langle (\delta_s(x))_{s \in S}, (\varphi_R(\vec{x}))_{R \in \Gamma} \rangle$ be an L -interpretation from Σ to Γ with

quantifier rank m . For every $L[\Gamma]$ -formula $\psi(\bar{X})$, there exists an $L[\Sigma]$ -formula $\psi^\tau(\bar{X})$ with quantifier rank at most $\text{qr}(\psi) + m$ such that

$$\tau(\mathfrak{A}) \models \psi(\bar{\alpha}) \quad \text{iff} \quad \mathfrak{A} \models \psi^\tau(\bar{\alpha}),$$

for all Σ -structures \mathfrak{A} and all parameters $\bar{\alpha}$ in $\tau(\mathfrak{A})$.

Proof. We define ψ^τ by induction on ψ as follows.

$$\begin{aligned} (x = y)^\tau &:= x = y, & (\varphi \wedge \psi)^\tau &:= \varphi^\tau \wedge \psi^\tau, \\ (Xy)^\tau &:= Xy, & (\varphi \vee \psi)^\tau &:= \varphi^\tau \vee \psi^\tau, \\ (R\bar{x})^\tau &:= \varphi_R(\bar{x}), & (\neg\varphi)^\tau &:= \neg\varphi^\tau, \\ (\exists y\psi)^\tau &:= \exists y[\delta_s(y) \wedge \psi^\tau], & (\exists Y\psi)^\tau &:= \exists Y\psi^\tau, \\ (\forall y\psi)^\tau &:= \forall y[\delta_s(y) \rightarrow \psi^\tau], & (\forall Y\psi)^\tau &:= \forall Y\psi^\tau, \end{aligned}$$

where s is the sort of the variable y . \square

Remark. Note that this statement fails for $L = \text{GSO}$ since guarded tuples in $\tau(\mathfrak{A})$ are not necessarily guarded in \mathfrak{A} . \lrcorner

Corollary 4.6. Let τ be an L -interpretation from Σ to Γ with quantifier rank m .

$$\mathfrak{A} \equiv_L^{k+m} \mathfrak{A}' \quad \text{implies} \quad \tau(\mathfrak{A}) \equiv_L^k \tau(\mathfrak{A}'). \quad \text{for all } \Sigma\text{-structures } \mathfrak{A}, \mathfrak{A}'.$$

Proof. By symmetry, it is sufficient to prove that

$$\tau(\mathfrak{A}) \models \varphi \quad \text{implies} \quad \tau(\mathfrak{A}') \models \varphi, \quad \text{for all } \varphi \text{ with } \text{qr}(\varphi) \leq k.$$

Hence, suppose that $\tau(\mathfrak{A}) \models \varphi$ and let φ^τ be the formula from Proposition 4.5. Then

$$\mathfrak{A} \models \varphi^\tau \quad \text{and} \quad \text{qr}(\varphi^\tau) \leq k + m.$$

Thus, $\mathfrak{A} \equiv_L^{k+m} \mathfrak{A}'$ implies that $\mathfrak{A}' \models \varphi^\tau$. It follows that $\tau(\mathfrak{A}') \models \varphi$. \square

Lemma 4.7. Let σ and τ be L -interpretations. Then so is $\tau \circ \sigma$.

Proof. Suppose that $\tau = \langle (\delta_s(x))_{s \in S}, (\varphi_R(\bar{x}))_{R \in \Gamma} \rangle$. We claim that $\tau \circ \sigma$ has the definition scheme

$$\langle (\delta_s^\sigma(x))_{s \in S}, (\varphi_R^\sigma(\bar{x}))_{R \in \Gamma} \rangle.$$

Note that, given a structure \mathfrak{A} , the elements $\tau(\sigma(\mathfrak{A}))$ of sort s are exactly those $a \in A$ satisfying

$$\sigma(\mathfrak{A}) \models \delta_s(a).$$

By Proposition 4.5, this condition is equivalent to

$$\mathfrak{A} \models \delta_s^\sigma(a).$$

Similarly, a tuple \bar{a} belongs to a relation R if, and only if,

$$\sigma(\mathfrak{A}) \models \varphi_R(\bar{a}) \quad \text{iff} \quad \mathfrak{A} \models \varphi_R^\sigma(\bar{a}). \quad \square$$

Frequently, disjoint unions and interpretations are all one needs to compute a theory. As an example, let us show how to generalise Proposition 4.3 to structures with unary predicates.

Proposition 4.8. Let $\Sigma = \{U_0, \dots, U_{m-1}\}$ be a signature consisting of unary predicates only. Over the class of all Σ -structures, every GSO-formula $\varphi(\bar{X}, \bar{x})$ is equivalent to an FO-formula.

Proof. Since the only guarded tuples over a unary signature are singletons, every GSO-formula can trivially be translated to an MSO-formula. Hence, by Lemma 3.3, it is sufficient to prove that, for every quantifier-rank $r < \omega$, there exists some $p < \omega$ such that

$$\mathfrak{A}, \bar{P}\bar{a} \equiv_{\text{FO}}^p \mathfrak{B}, \bar{Q}\bar{b} \quad \text{implies} \quad \mathfrak{A}, \bar{P}\bar{a} \equiv_{\text{MSO}}^r \mathfrak{B}, \bar{Q}\bar{b},$$

for all $\mathfrak{A}, \mathfrak{B}$ with parameters $\bar{P}\bar{a}$ and $\bar{Q}\bar{b}$. To simplify notation, we will not work with parameters but with structures $\mathfrak{A} = \langle A, \bar{P}, \bar{c} \rangle$ where the parameters are part of the structure itself. Hence, let $\mathcal{U}_{m,n}$ be the class of all

such structures with m unary predicates $P_0, \dots, P_{m-1} \subseteq A$ and n constant symbols c_0, \dots, c_{n-1} . Given $\mathfrak{A} \in \mathcal{U}_{m,n}$ and a set $\theta \subseteq [m]$, we set

$$P_\theta := \bigcap_{i \in \theta} P_i \setminus \bigcup_{i \in [m] \setminus \theta} P_i,$$

we denote by \mathfrak{A}_θ the substructure induced by $P_\theta \setminus \bar{c}$, and $\langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}}$ is the substructure generated by \bar{c} . Then we can write \mathfrak{A} as a disjoint union

$$\mathfrak{A} \cong \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \oplus \bigoplus_{\theta \subseteq [m]} \mathfrak{A}_\theta.$$

Let us make the following observations.

(I) For every structure \mathfrak{C} of size at most k , there exists a first-order formula of quantifier-rank $k+1$ that characterises \mathfrak{C} up to isomorphisms. Consequently, for $\bar{a} \in A^k$ and $\bar{b} \in B^k$,

$$\langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \equiv_{\text{FO}}^{k+1} \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}} \text{ implies } \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \cong \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}}.$$

(II) For every MSO-formula φ , we can use Proposition 4.3 to find a finite set $H \subseteq \omega$ and a number $N < \omega$ such that,

$$\mathfrak{C} \models \varphi \quad \text{iff} \quad |C| \in H,$$

$$\text{or } \mathfrak{C} \models \varphi \quad \text{iff} \quad |C| \in H \quad \text{or} \quad |C| \geq N,$$

for all $\mathfrak{C} \in \mathcal{U}_{\omega,0}$. Since, for every $k < \omega$, we can construct an FO-formula ψ_k stating that the structure has at least k elements, it follows that there exists some number $f(r)$ such that

$$\mathfrak{A} \equiv_{\text{FO}}^{f(r)} \mathfrak{B} \text{ implies } \mathfrak{A} \equiv_{\text{MSO}}^r \mathfrak{B}, \quad \text{for all } \mathfrak{A}, \mathfrak{B} \in \mathcal{U}_{\omega,0}.$$

(III) For every $\theta \subseteq [m]$, there exists a quantifier-free interpretation σ_θ mapping $\mathfrak{C} \in \mathcal{U}_{\omega,0}$ to a structure $\sigma_\theta(\mathfrak{C}) \in \mathcal{U}_{m,0}$ with predicates

$$P_i := \begin{cases} C & \text{if } i \in \theta, \\ \emptyset & \text{if } i \notin \theta. \end{cases}$$

In particular,

$$\mathfrak{A}_\theta = \sigma_\theta((\mathfrak{A}_\theta)|_\emptyset), \quad \text{for } \mathfrak{A} \in \mathcal{U}_{m,n} \text{ and } \theta \subseteq [m],$$

where $\mathfrak{C}|_\emptyset$ denotes the reduct to the empty signature. For $\mathfrak{A}, \mathfrak{B} \in \mathcal{U}_{m,n}$ and $\theta \subseteq [m]$, it follows that

$$(\mathfrak{A}_\theta)|_\emptyset \equiv_{\text{MSO}}^r (\mathfrak{B}_\theta)|_\emptyset \text{ implies } \mathfrak{A}_\theta \equiv_{\text{MSO}}^r \mathfrak{B}_\theta.$$

(IV) There exists quantifier-free interpretations τ° and τ_θ such that

$$\tau^\circ(\mathfrak{A}) = \langle\langle \bar{a} \rangle\rangle_{\mathfrak{A}} \quad \text{and} \quad \tau_\theta(\mathfrak{A}) = (\mathfrak{A}_\theta)|_\emptyset.$$

Hence,

$$\mathfrak{A} \equiv_{\text{FO}}^k \mathfrak{B} \text{ implies } \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \equiv_{\text{FO}}^k \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}} \text{ and } (\mathfrak{A}_\theta)|_\emptyset \equiv_{\text{FO}}^k (\mathfrak{B}_\theta)|_\emptyset.$$

We can conclude the proof as follows. Set $p := \max\{f(r), n+1\}$ and let $\mathfrak{A}, \mathfrak{B} \in \mathcal{U}_{m,n}$. Combining the above observations it follows by Proposition 4.2 that

$$\begin{aligned} & \mathfrak{A} \equiv_{\text{FO}}^p \mathfrak{B} \\ \Rightarrow & \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \equiv_{\text{FO}}^p \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}} \quad \text{and} \quad (\mathfrak{A}_\theta)|_\emptyset \equiv_{\text{FO}}^p (\mathfrak{B}_\theta)|_\emptyset, \quad \text{for all } \theta, \\ \Rightarrow & \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \cong \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}} \quad \text{and} \quad (\mathfrak{A}_\theta)|_\emptyset \equiv_{\text{MSO}}^r (\mathfrak{B}_\theta)|_\emptyset, \quad \text{for all } \theta, \\ \Rightarrow & \langle\langle \bar{c} \rangle\rangle_{\mathfrak{A}} \equiv_{\text{MSO}}^r \langle\langle \bar{d} \rangle\rangle_{\mathfrak{B}} \quad \text{and} \quad \mathfrak{A}_\theta \equiv_{\text{MSO}}^r \mathfrak{B}_\theta, \quad \text{for all } \theta, \\ \Rightarrow & \mathfrak{A} \equiv_{\text{MSO}}^r \mathfrak{B}. \end{aligned} \quad \square$$

As a second example, let us give the example of an *ordered sum*, which corresponds to concatenations of words.

Definition 4.9. Let C be a set of colours.

(a) A C -coloured order is a structure of the form $\mathfrak{A} = \langle A, \leq, (P_c)_{c \in C} \rangle$ where \leq is a linear ordering on A and the P_c are unary predicates.

(b) Let $\mathfrak{I} = \langle I, \leq \rangle$ be a linear order and let $\mathfrak{A}_i := \langle A_i, \leq_i, \bar{P}_i \rangle$, $i \in I$, be a family of C -coloured linear orders indexed by I . The *ordered sum*

$$\sum_{i \in I} \mathfrak{A}_i$$

is the linear order with universe

$$L := \{ \langle i, a \rangle \mid i \in I, a \in A_i \}$$

and order

$$\langle i, a \rangle \leq \langle j, b \rangle \quad \text{iff} \quad i < j \quad \text{or} \quad (i = j \text{ and } a \leq_i b).$$

The colour predicates are

$$P_c := \bigcup_{i \in I} (P_i)_c.$$

If $I = [2]$, we simply write $\mathfrak{A}_0 + \mathfrak{A}_1$ for the ordered sum.

Proposition 4.10. *Let $\mathfrak{A}_0, \mathfrak{A}_1, \mathfrak{B}_0, \mathfrak{B}_1$ be C -coloured linear orders and let L be FO, MSO, WMSO, or CMSO. Then*

$$\mathfrak{A}_0 \equiv_L^m \mathfrak{B}_0 \quad \text{and} \quad \mathfrak{A}_1 \equiv_L^m \mathfrak{B}_1 \quad \text{implies} \quad \mathfrak{A}_0 + \mathfrak{A}_1 \equiv_L^m \mathfrak{B}_0 + \mathfrak{B}_1.$$

Proof. We have

$$\mathfrak{A}_0 + \mathfrak{A}_1 \cong \tau(\mathfrak{A}_0 \oplus \mathfrak{A}_1),$$

where τ is a quantifier-free L -interpretation that corrects the order relation. It has the definition scheme

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= x \leq y \vee (\text{Left}(x) \wedge \text{Right}(y)), \\ \varphi_{P_c}(x) &:= P_c x. \end{aligned}$$

□

As a final application let us show that first-order logic cannot compute the length of a linear order.

Proposition 4.11. *Let \mathfrak{A} be a C -coloured linear order and $m < \omega$ a constant. Then*

$$k \times \mathfrak{A} \equiv_{\text{FO}}^m l \times \mathfrak{A}, \quad \text{for all } k, l \geq 2^m - 1,$$

where $k \times \mathfrak{A} := \sum_{i < k} \mathfrak{A}$ denotes the k -fold ordered sum of \mathfrak{A} with itself.

Proof. We proceed by induction on m . For $m = 0$, we have $k \times \mathfrak{A} \equiv_{\text{FO}}^0 l \times \mathfrak{A}$, for all k, l , since there are no quantifier-free first-order formulae without free variables (over the signature of C -coloured linear orders). For the inductive step, suppose that $m > 0$. We check the forth property. (As usual the back property follows by symmetry.) Hence, let a be an element of $k \times \mathfrak{A}$ and suppose that a belongs to the i -th copy of \mathfrak{A} . In $l \times \mathfrak{A}$, we choose the same element in the j -th copy of \mathfrak{A} where

$$j := \begin{cases} i & \text{if } i \leq 2^{m-1}, \\ l - (k - i) & \text{otherwise.} \end{cases}$$

Let us denote this element by b . By inductive hypothesis it follows that

$$\begin{aligned} k \times \mathfrak{A}, a &\cong (i - 1) \times \mathfrak{A} + \mathfrak{A}, a + (k - i) \times \mathfrak{A} \\ &\equiv_{\text{FO}}^{m-1} (j - 1) \times \mathfrak{A} + \mathfrak{A}, b + (l - j) \times \mathfrak{A} \cong l \times \mathfrak{A}, b. \end{aligned} \quad \square$$

For MSO and CMSO, we obtain the following result.

Proposition 4.12. *For every $m < \omega$, there exist numbers $k, k', l', l' < \omega$ such that*

$$\begin{aligned} \langle A, \leq \rangle \equiv_{\text{MSO}}^m \langle B, \leq \rangle &\quad \text{iff} \quad |A| = |B| < k, \quad \text{or} \\ &\quad |A|, |B| \geq k \quad \text{and} \quad |A| \equiv |B| \pmod{l}, \\ \langle A, \leq \rangle \equiv_{\text{CMSO}}^m \langle B, \leq \rangle &\quad \text{iff} \quad |A| = |B| < k', \quad \text{or} \\ &\quad |A|, |B| \geq k' \quad \text{and} \quad |A| \equiv |B| \pmod{l'}, \end{aligned}$$

for all finite linear orders $\langle A, \leq \rangle$ and $\langle B, \leq \rangle$.

Proof. Let L be one of MSO or CMSO. Let Θ be the set of all L_m -theories of finite linear orders. It follows by Proposition 4.10 that we can define a binary operation $+$ on Θ such that

$$\text{Th}_L^m(\mathfrak{A}) + \text{Th}_L^m(\mathfrak{B}) = \text{Th}_L^m(\mathfrak{A} + \mathfrak{B}),$$

for all finite linear orders \mathfrak{A} and \mathfrak{B} . This turns Θ into a finite semigroup. Let σ be the theory of the 1-element linear order. Since Θ is finite, there is some

number $n > 1$ such that $\sigma^n = \sigma^k$, for some $k < n$. We choose n minimal. Let $l := n - k$. Then

$$\sigma^{k+il+j} = \sigma^{k+j}, \quad \text{for all } i, j,$$

and it follows that

$$\sigma^i = \sigma^j \quad \text{iff} \quad i = j \quad \text{or} \quad i, j \geq k \text{ and } i - k \equiv j - k \pmod{l}.$$

Since $\text{Th}_L^m(\mathfrak{A}) = \sigma^{|\mathfrak{A}|}$ the claim follows. \square

Corollary 4.13. *For each CMSO-formula φ there exists an CMSO-formula φ^* such that*

$$\langle A, \leq \rangle \models \varphi \quad \text{iff} \quad \langle A \rangle \models \varphi^*.$$

Proof. Let m be the quantifier-rank of φ and k', l' the constants from the preceding lemma. By the lemma, there exist sets $K \subseteq [k']$ and $L \subseteq [l']$ such that

$$\langle A, \leq \rangle \models \varphi \quad \text{iff} \quad |A| \in K, \quad \text{or} \quad |A| \geq k' \quad \text{and} \quad |A| \bmod l' \in L.$$

This is a condition that can be expressed in CMSO. \square

Remark. We can rephrase this statement by saying that the reduct operation $\langle A, \leq \rangle \mapsto \langle A \rangle$ is CMSO-compatible. \lrcorner

Example. (a) There does not exist an FO-formula φ that holds in an undirected graph if, and only if, the graph is connected.

For a contradiction, suppose that such a formula φ exists. We will construct a new formula ψ that holds in a finite linear order if, and only if, this order has an even number of elements. Let m be the quantifier rank of φ and let \mathfrak{A} and \mathfrak{B} be linear orders of size 2^m and $2^m + 1$, respectively. Then


$$\mathfrak{A} \models \psi \quad \text{and} \quad \mathfrak{B} \not\models \psi,$$

in contradiction to the statement in the above exercise.

To construct the desired formula ψ , we define an FO-interpretation $\tau = \langle \delta, \varphi_E \rangle$ mapping linear orders to undirected graphs as follows. The formula δ is true while $\varphi_E(x, y)$ states that

- ♦ in the order \leq there is exactly one element between x and y , or
- ♦ x is the first element and y is the last one, or
- ♦ y is the first element and x is the last one.

Then τ maps finite linear orders of even size to paths and finite linear orders of odd size (at least 3) to the disjoint union of a path and a cycle. Orders of size 1 are mapped to a loop.



Hence,

$$\tau(\mathfrak{A}) \models \varphi \quad \text{iff} \quad \mathfrak{A} \text{ has either exactly one, or an even number of elements.}$$

Consequently, the formula $\psi := \varphi^\tau \wedge \exists x y (x \neq y)$ has the desired properties.

(b) There does not exist an FO-formula $\varphi(x, y)$ such that

$$\mathfrak{G} \models \varphi(u, v) \quad \text{iff} \quad \text{the graph } \mathfrak{G} \text{ contains a path from } u \text{ to } v.$$

Otherwise, the formula

$$\forall x \forall y \varphi(x, y)$$

would express that the graph is connected. \lrcorner

Example. We consider undirected graphs as structures over the signature $\{E\}$.

(a) There does not exist an MSO-formula φ that holds in a finite complete bipartite graph $K_{m,n}$ if, and only if, $m = n$. The proof is similar to that of Proposition 4.10. Suppose that such a formula φ exists and let k be its quantifier rank. Let \mathfrak{A} and \mathfrak{B} be graphs without any edges that have, respectively, $m := 2^k$ and $n := 2^k + 1$ vertices. Then

$$K_{m,m} := \tau(\mathfrak{A} \oplus \mathfrak{A}) \quad \text{and} \quad K_{m,n} := \tau(\mathfrak{A} \oplus \mathfrak{B}),$$

where τ is a quantifier-free interpretation that adds all edges between a vertex in Left and a vertex in Right. Since $\mathcal{A} \equiv_{\text{MSO}}^k \mathcal{B}$ it follows that

$$K_{m,m} = \tau(\mathcal{A} \oplus \mathcal{A}) \equiv_{\text{MSO}}^k \tau(\mathcal{A} \oplus \mathcal{B}) = K_{m,n}.$$

A contradiction, since φ distinguishes between these two graphs.

(b) There does not exist an MSO-formula φ that holds in a finite graph if, and only if, all vertices have the same number of neighbours. For a contradiction, suppose that such a formula φ exists. For a complete bipartite graph $K_{m,n}$ it follows that

$$K_{m,n} \models \varphi \quad \text{iff} \quad m = n.$$

This contradicts (a).

(c) There does not exist an MSO-formula φ that holds in a finite undirected graph if, and only if, the graph has a Hamiltonian cycle. For a contradiction, suppose that such a formula φ exists. Since a complete bipartite graph $\mathfrak{K}_{m,n}$ contains an Hamiltonian cycle if, and only if, $m = n$, it follows that

$$\mathfrak{K}_{m,n} \models \varphi \quad \text{iff} \quad m = n.$$

This contradicts (a). J

Quotients

In some contexts it is usual to combine interpretations with a quotient operation. To simplify the presentation we present these operations separately.

Definition 4.14. Let \mathcal{A} be a $(\Sigma + \approx)$ -structure where $\approx^{\mathcal{A}}$ is an equivalence relation on A . The *quotient* of \mathcal{A} by \approx is the Σ -structure \mathcal{A}/\approx with universe

$$A/\approx := \{ [a]_{\approx} \mid a \in A \}$$

and relations

$$R^{\mathcal{A}/\approx} := \{ \langle [a_0]_{\approx}, \dots, [a_{n-1}]_{\approx} \rangle \mid \bar{a} \in R^{\mathcal{A}} \}.$$
J

Proposition 4.15. Let L be one of FO, MSO, WMSO, or GSO, let Σ be a signature with relations of arity at most r , and let \mathcal{A}, \mathcal{B} be a $(\Sigma + \approx)$ -structures such that $\approx^{\mathcal{A}}$ and $\approx^{\mathcal{B}}$ are equivalence relations. Then

$$\mathcal{A} \equiv_L^{m+r} \mathcal{B} \quad \text{implies} \quad \mathcal{A}/\approx \equiv_L^m \mathcal{B}/\approx.$$

Proof. Given $\varphi \in L_m$, we construct a formula $\varphi' \in L_{m+r}$ such that

$$\mathcal{A}/\approx \models \varphi \quad \text{iff} \quad \mathcal{A} \models \varphi'.$$

We obtain φ' by

- replacing each atomic subformula of the form $x = y$ by $x \approx y$,
- replacing each atomic subformula of the form $R\bar{x}$ (where R is either a relation symbol or a guarded second-order variable) by

$$\exists \bar{y} \left[\bigwedge_i y_i \approx x_i \wedge R\bar{y} \right].$$

To reach the desired quantifier-rank for $L = \text{GSO}$, we have to make sure in this translation that the arity of guarded variables is bounded by r . But note that, since every guarded tuple has at most r distinct components, we can replace each guarded variable Z of arity $n > r$ by one of arity r (or rather a tuple $(Z_\sigma)_\sigma$ of such variables indexed by all surjective functions $\sigma : [n] \rightarrow [r]$). □

In the above lemma the quantifier rank increases when going from a structure to its quotient. Sometimes this can be avoided by using the following simple version of a quotient.

Definition 4.16. Let \mathcal{A} be a $(\Sigma + \{P\})$ -structure where P is a unary predicate. The *fusion* $\text{fuse}_P(\mathcal{A})$ of \mathcal{A} is the Σ -reduct of the quotient \mathcal{A}/\approx where

$$a \approx b \quad : \text{iff} \quad a = b \text{ or } a, b \in P.$$
J

Since a fusion is a quotient by an equivalence relation with just one non-trivial class, we can avoid increasing the quantifier rank by annotating the structure by information about the elements in this class.

Proposition 4.17. Let L be one of FO, MSO, WMSO, CMSO, or GSO, and let \mathcal{A} and \mathcal{B} be $(\Sigma + \{P\})$ -structures. Then

$$\langle \mathcal{A}|_{A \setminus P}, \bar{U} \rangle \equiv_L^m \langle \mathcal{B}|_{B \setminus P}, \bar{V} \rangle \text{ implies } \text{fuse}_P(\mathcal{A}) \equiv_L^m \text{fuse}_P(\mathcal{B}),$$

where the parameters $\bar{U} = (U_{R,w})_{R,w}$ and $\bar{V} = (V_{R,w})_{R,w}$ contain, for every $R \in \Sigma$ of arity n and every set $w \subseteq [n]$, the predicate

$$U_{R,w} := \{ \bar{a}|_w \mid \bar{a} \in R^{\mathcal{A}}, a_i \in P \Leftrightarrow i \notin w \}$$

(and similarly for \bar{V}).

Proof. Let \mathcal{C} be the Σ -structure with one element and empty relations. There exists a quantifier-free interpretation τ such that

$$\text{fuse}_P(\mathcal{A}) = \tau(\langle \mathcal{A}|_{A \setminus P}, \bar{U} \rangle \oplus \mathcal{C}).$$

Hence, the result follows from Propositions 4.2 and 4.5. \square

It turns out that when constructing a structure \mathcal{A} from smaller parts, we can often construct the annotated substructure $\langle \mathcal{A}|_{A \setminus P}, \bar{U} \rangle$ instead, which then allows us to compute $\text{fuse}_P(\mathcal{A})$ by the above proposition.

Exercise 4.2. Prove that, for every L -interpretation τ , there is some L -interpretation σ such that

$$\tau(\mathcal{A}/\approx) = \sigma(\mathcal{A})/\approx.$$

Exercise 4.3. Prove that, for every L -interpretation τ , there is some L -interpretation σ such that

$$\wp(\tau(\mathcal{A})) = \sigma(\wp(\mathcal{A})).$$

The Copying Operation

Next, let us introduce a variant of the disjoint union that will be used extensively in Chapter X.

Definition 4.18. The k -copy operation is of the form

$$\text{copy}_k(\mathcal{A}) := \langle \mathcal{A} \oplus \dots \oplus \mathcal{A}, H_0, \dots, H_{k-1}, I \rangle.$$

That is $\text{copy}_k(\mathcal{A})$ consists of k disjoint copies of \mathcal{A} with unary predicates

$$H_i := \{ \langle i, a \rangle \mid a \in A \}$$

containing the i -th copy, and a binary relation

$$I := \{ \langle \langle i, a \rangle, \langle j, a \rangle \rangle \mid a \in A, i, j < k \}$$

that relates all copies of the same element. J

Proposition 4.19. Let L be one of FO, MSO, WMSO, CMSO, or GSO, and let $m, k < \omega$. For every L_m -formula $\varphi(\bar{x}, \bar{Y})$ with $|\bar{x}| = r$ and $|\bar{Y}| = s$ and every tuple $\bar{u} \in [k]^s$, there exists an L_{mk} -formula $\varphi'_u(\bar{x}, \bar{Y}')$ with $|\bar{Y}'| = sk$ such that

$$\begin{aligned} \text{copy}_k(\mathcal{A}) \models \varphi(a_0, \dots, a_{r-1}, P_0, \dots, P_{s-1}) \\ \text{iff} \quad \mathcal{A} \models \varphi'_u(a'_0, \dots, a'_{r-1}, \bar{P}'_0, \dots, \bar{P}'_{s-1}), \end{aligned}$$

for all structures \mathcal{A} , elements \bar{a}, \bar{a}' , and sets \bar{P}, \bar{P}' that are related via

$$a_i = \langle u_i, a'_i \rangle \quad \text{and} \quad (P'_i)_v = \{ b \in A \mid \langle v, b \rangle \in P_i \}.$$

Proof. We can construct φ'_u by induction on φ . We replace each variable Y_i by a k -tuple $\bar{Y}'_i = \langle Y'_{i,0}, \dots, Y'_{i,k-1} \rangle$.

$$\begin{aligned} (x_i = x_j)'_{\bar{u}} &:= \begin{cases} x_i = x_j & \text{if } u_i = u_j, \\ \text{false} & \text{otherwise,} \end{cases} \\ (Rx_{i_0} \dots x_{i_{n-1}})'_{\bar{u}} &:= \begin{cases} Rx_{i_0} \dots x_{i_{n-1}} & \text{if } u_{i_0} = \dots = u_{i_{n-1}}, \\ \text{false} & \text{otherwise,} \end{cases} \end{aligned}$$

$$\begin{aligned}
(Y_i x_j)'_{\bar{u}} &:= Y'_{i,u_j} x_j, & (\exists z \varphi)'_{\bar{u}} &:= \bigvee_{v < k} \exists z \neg \varphi'_{\bar{u}v}, \\
(\varphi \vee \psi)'_{\bar{u}} &:= \varphi'_{\bar{u}} \vee \psi'_{\bar{u}}, & (\forall z \varphi)'_{\bar{u}} &:= \bigwedge_{v < k} \forall z \neg \varphi'_{\bar{u}v}, \\
(\varphi \wedge \psi)'_{\bar{u}} &:= \varphi'_{\bar{u}} \wedge \psi'_{\bar{u}}, & (\exists Z \varphi)'_{\bar{u}} &:= \exists \bar{Z}' \neg \varphi'_{\bar{u}}, \\
(\neg \varphi)'_{\bar{u}} &:= \neg \varphi'_{\bar{u}}, & (\forall Z \varphi)'_{\bar{u}} &:= \forall \bar{Z}' \neg \varphi'_{\bar{u}}, \\
(|Y_i| < \infty)'_{\bar{u}} &:= \bigwedge_{v < k} |Y'_{i,v}| < \infty, \\
(|Y_i| \equiv n \pmod{p})'_{\bar{u}} &:= \bigvee \left\{ \bigwedge_{v < k} |Y'_{i,v}| \equiv g(v) \pmod{p} \mid \right. \\
&\quad \left. g : [k] \rightarrow [p] \text{ with } \sum_{v < k} g(v) \equiv n \pmod{p} \right\}. \quad \square
\end{aligned}$$

Let us also note that copying operations commute with interpretations.

Lemma 4.20. *Let L be one of FO, MSO, WMSO, or CMSO, and let $k < \omega$. For every L -interpretation τ , there exists an L -interpretation τ' such that*

$$\text{copy}_k \circ \tau = \tau' \circ \text{copy}_k.$$

Proof. The transduction τ' applies τ separately to each copy H_i . That is, if τ is defined by $\langle \delta(x), (\varphi_R(\bar{x})) \rangle$, we define τ' by

$$\begin{aligned}
\delta'(x) &:= \bigvee_{i < k} [H_i x \wedge \delta^{(H_i)}(x)], \\
\text{and } \varphi'_R(\bar{x}) &:= \bigvee_{i < k} \left[\bigwedge_j H_i x_j \wedge \varphi_R^{(H_i)}(\bar{x}) \right],
\end{aligned}$$

where $\psi^{(P)}$ denotes the relativisation of ψ to the set P . \square

Exercise 4.4. Let $k, l < \omega$. Find a quantifier-free interpretation τ such that

$$\text{copy}_k(\text{copy}_l(\mathfrak{A})) = \tau(\text{copy}_{kl}(\mathfrak{A})).$$

Generalised Sums

So far, we have only looked finite unions. It is straightforward to generalise the arguments from Section 4 to infinite ones. But it turns out that, for many applications, an even more general approach is needed where we allow the number of structures in the two unions to differ. We use a union of Σ -structures where the index set is itself a Γ -structure, for some signature Γ . The result then has all relations in Σ and Γ . This leads to the following definition.

Definition 4.21. Let Γ and Σ be relational signatures, \mathfrak{I} a Γ -structure and, for every element $i \in I$, let \mathfrak{A}_i be a Σ -structure. The *generalised sum* of the \mathfrak{A}_i over \mathfrak{I} is the $(\Gamma + \Sigma + \{\sim\})$ -structure $\sum_{i \in \mathfrak{I}} \mathfrak{A}_i$ with universe

$$U := \{ \langle i, a \rangle \mid i \in I, a \in A_i \}$$

and the following relations. For every n -ary relation symbol $R \in \Sigma$, it has the relation

$$R := \{ \langle \langle i, a_0 \rangle, \dots, \langle i, a_{n-1} \rangle \rangle \mid i \in I, \langle a_0, \dots, a_{n-1} \rangle \in R^{\mathfrak{A}_i} \},$$

for every n -ary relation symbol $R \in \Gamma$, it has the relation

$$R := \{ \langle \langle i_0, a_0 \rangle, \dots, \langle i_{n-1}, a_{n-1} \rangle \rangle \mid \langle i_0, \dots, i_{n-1} \rangle \in R^{\mathfrak{I}}, a_j \in A_{i_j} \},$$

and additionally it has the equivalence relation

$$\sim := \{ \langle \langle i, a \rangle, \langle i, b \rangle \rangle \mid i \in I, a, b \in A_i \}.$$

Example. (a) The disjoint union $\mathfrak{A}_0 \oplus \mathfrak{A}_1$ can be written (up to a quantifier-free interpretation) as a generalised sum $\sum_{i \in \mathfrak{I}} \mathfrak{A}_i$ where

$$\mathfrak{I} = \langle [2], \text{Left}, \text{Right} \rangle \quad \text{with} \quad \text{Left} = \{0\} \quad \text{and} \quad \text{Right} = \{1\}.$$

(b) Given linear orders $\mathfrak{I} = \langle I, \sqsubseteq \rangle$ and $\mathfrak{A}_i := \langle A_i, \leq_i \rangle$, for $i \in I$, the generalised sum is the structure $\langle U, \sqsubseteq, \leq, \sim \rangle$ with relations

$$\begin{aligned}
\langle i, a \rangle \sqsubseteq \langle j, b \rangle &\quad \text{iff} \quad i \sqsubseteq j, \\
\langle i, a \rangle \leq \langle j, b \rangle &\quad \text{iff} \quad i = j \text{ and } a \leq_i b, \\
\langle i, a \rangle \sim \langle j, b \rangle &\quad \text{iff} \quad i = j.
\end{aligned}$$

We can interpret the ordered sum of the \mathfrak{A}_i in this structure via the interpretation $\tau = \langle \delta, \varphi_{\leq} \rangle$ with

$$\delta(x) := \text{true} \quad \text{and} \quad \varphi_{\leq}(x, y) := x \subseteq y \wedge [x \sim y \rightarrow x \leq y].$$

For an application in Section III.2 below, we will need to prove the composition theorem for generalised sums using a finer notion than just the quantifier rank: we will not only need to count the number of quantifiers, but also their alternation. We call this combined measure the *quantifier structure* of a formula.

Definition 4.22. Let $\bar{n} \in \omega^*$.

(a) We define the set $\text{MSO}_{\bar{n}}^{\circ}[\Sigma]$ of all MSO° -formulae with *quantifier structure* \bar{n} as follows. $\text{MSO}_{\emptyset}^{\circ}[\Sigma]$ contains all quantifier-free $\text{MSO}^{\circ}[\Sigma]$ -formulae, and $\text{MSO}_{m\bar{n}}^{\circ}[\Sigma]$ contains all formulae that can be written as boolean combinations of formulae of the form

$$\exists X_0 \cdots \exists X_{m-1} \psi \quad \text{with} \quad \psi \in \text{MSO}_{\bar{n}}^{\circ}[\Sigma].$$

(b) We denote by $\text{Th}_{\text{MSO}^{\circ}}^{\bar{n}}(\mathfrak{A}, \bar{P})$ the $\text{MSO}_{\bar{n}}^{\circ}$ -theory of \mathfrak{A}, \bar{P} and we set

$$\mathfrak{A}, \bar{P} \equiv_{\text{MSO}^{\circ}}^{\bar{n}} \mathfrak{B}, \bar{Q} \quad \text{iff} \quad \text{Th}_{\text{MSO}^{\circ}}^{\bar{n}}(\mathfrak{A}, \bar{P}) = \text{Th}_{\text{MSO}^{\circ}}^{\bar{n}}(\mathfrak{B}, \bar{Q}).$$

The composition theorem for generalised sums not only states that the theory of the resulting structure only depends on the theories of the arguments, but also that we can compute this theory by evaluating a formula on the index structure.

Definition 4.23. Let $\sum_{i \in \mathfrak{I}} \mathfrak{A}_i$ be a generalised sum and let \bar{P} be a tuple of monadic parameters. For an $\text{MSO}^{\circ}[\Sigma]$ -formula $\chi(\bar{X})$, we define

$$\llbracket \chi(\bar{P}) \rrbracket := \{ i \in I \mid \mathfrak{A}_i \models \chi(\bar{P} \upharpoonright A_i) \}.$$

Theorem 4.24. Let Γ and Σ be relational signatures. Given a formula $\varphi(\bar{X}) \in \text{MSO}_{\bar{n}}^{\circ}[\Gamma + \Sigma + \{\sim\}]$, we can compute a tuple $\bar{r} \in \omega^*$ of length $|\bar{r}| = |\bar{n}|$ and formulae

$$\chi_0(\bar{X}), \dots, \chi_{m-1}(\bar{X}) \in \text{MSO}_{\bar{n}}^{\circ}[\Sigma]$$

and $\psi(Z_0, \dots, Z_{m-1}) \in \text{MSO}_{\bar{r}}^{\circ}[\Gamma]$

such that

$$\sum_{i \in \mathfrak{I}} \mathfrak{A}_i \models \varphi(\bar{P}) \quad \text{iff} \quad \mathfrak{I} \models \varphi'(\llbracket \chi_0(\bar{P}) \rrbracket, \dots, \llbracket \chi_{m-1}(\bar{P}) \rrbracket),$$

for all Γ -structures \mathfrak{I} , Σ -structures \mathfrak{A}_i , and monadic parameters \bar{P} .

Proof. We construct $\bar{r}, \psi(\bar{Z})$, and $\chi_0(\bar{X}), \dots, \chi_{m-1}(\bar{X})$ by induction on φ . First, suppose that φ is atomic. We distinguish several cases. For $\varphi = (X \subseteq Y)$ we set

$$\psi := \text{cover}(Z_0) \quad \text{and} \quad \chi_0(X, Y) := X \subseteq Y.$$

For $\varphi = (X \cap Y = \emptyset)$ we set

$$\psi := \text{cover}(Z_0) \quad \text{and} \quad \chi_0(X, Y) := X \cap Y = \emptyset.$$

For $\varphi = \text{cover}(\bar{X})$ we set

$$\psi := \text{cover}(Z_0) \quad \text{and} \quad \chi_0(\bar{X}) := \text{cover}(\bar{X}).$$

For $\varphi = R\bar{X}$ with $R \in \Sigma$, we set

$$\psi := \text{sing}(Z_0) \quad \text{and} \quad \chi_0(\bar{X}) := R\bar{X}.$$

For $\varphi = R\bar{X}$ with $R \in \Gamma$ and R of arity m , we set

$$\begin{aligned} \psi &:= R\bar{Z} \wedge \text{cover}(Z_m), \\ \chi_i(\bar{X}) &:= \text{sing}(X_i), \quad \text{for } i < m, \\ \chi_m(\bar{X}) &:= \bigwedge_{i < m} [\text{sing}(X_i) \vee X_i = \emptyset]. \end{aligned}$$

For $\varphi = (X \sim Y)$, we set

$$\begin{aligned} \psi &:= Z_0 \cap Z_1 \neq \emptyset \wedge \text{sing}(Z_0) \wedge \text{sing}(Z_1) \wedge \text{cover}(Z_2), \\ \chi_0(X, Y) &:= \text{sing}(X), \\ \chi_1(X, Y) &:= \text{sing}(Y), \\ \chi_2(X, Y) &:= [\text{sing}(X) \vee X = \emptyset] \wedge [\text{sing}(Y) \vee Y = \emptyset]. \end{aligned}$$

For the inductive step, suppose that we have already computed formulae $\psi', \chi'_0, \dots, \chi'_{m'-1}$ and $\psi'', \chi''_0, \dots, \chi''_{m''-1}$ for, respectively, φ' and φ'' . For the conjunction $\varphi = \varphi' \wedge \varphi''$, we set

$$\psi(\bar{Z}, \bar{Z}') := \psi'(\bar{Z}) \wedge \psi''(\bar{Z}') \quad \text{and} \quad \chi_i := \begin{cases} \chi'_i & \text{if } i < m', \\ \chi''_{i-m'} & \text{if } i \geq m'. \end{cases}$$

The construction for $\varphi = \varphi' \vee \varphi''$ is analogous. For a negation $\varphi = \neg\varphi'$, we set

$$\psi(\bar{Z}) := \neg\psi'(\bar{Z}) \quad \text{and} \quad \chi_i := \chi'_i.$$

It remains to consider the case where $\varphi(\bar{X}) = \exists \bar{Y} \varphi'(\bar{X}, \bar{Y})$. Again we may assume that we have already constructed formulae $\psi', \chi'_0, \dots, \chi'_{m-1}$ for φ' . A first attempt might be to use the formulae $\psi := \psi'$ and $\chi_j := \exists \bar{Y} \chi'_j$ for φ . But this does not work since, for example, the sets \bar{Y} we use to make χ'_0 true might be different from those we take for χ'_1 . Instead, we have to know which of the χ'_j we can satisfy at the same time. Consequently, we set

$$\chi_w(\bar{X}) := \exists \bar{Y} \left[\bigwedge_{j \in w} \chi'_j(\bar{X}, \bar{Y}) \wedge \bigwedge_{j \in [m] \setminus w} \neg \chi'_j(\bar{X}, \bar{Y}) \right], \quad \text{for } w \subseteq [m],$$

$$\text{and} \quad \psi := \exists \bar{Z}' \left[\psi'(\bar{Z}') \wedge \forall x \bigvee_{w \subseteq [m]} \left[Z_w x \wedge \bigwedge_{j \in w} Z'_j x \wedge \bigwedge_{j \notin w} \neg Z'_j x \right] \right].$$

It remains to check that these formulae have the desired properties. We have

$$\begin{aligned} \sum_{i \in \mathfrak{I}} \mathfrak{A}_i &\models \exists \bar{Y} \varphi'(\bar{P}, \bar{Y}) \\ \text{iff} \quad &\sum_{i \in \mathfrak{I}} \mathfrak{A}_i \models \varphi'(\bar{P}, \bar{Q}), \quad \text{for some } \bar{Q}, \\ \text{iff} \quad &\mathfrak{I} \models \psi'(\llbracket \chi'_0(\bar{P}, \bar{Q}) \rrbracket, \dots, \llbracket \chi'_{m-1}(\bar{P}, \bar{Q}) \rrbracket), \quad \text{for some } \bar{Q}. \end{aligned}$$

We claim that the latter is equivalent to

$$\mathfrak{I} \models \psi(\llbracket \chi_w(\bar{P}) \rrbracket_{w \subseteq [m]}).$$

(\Rightarrow) Given sets \bar{Q} as above, we set

$$U_j := \llbracket \chi'_j(\bar{P}, \bar{Q}) \rrbracket \quad \text{and} \quad w_i := \{ j \mid i \in U_j \}, \quad \text{for } i \in I.$$

Then

$$\mathfrak{I} \models \psi'(\bar{U}), \quad i \in \llbracket \chi_w(\bar{P}) \rrbracket_{w \subseteq [m]} \quad \text{and} \quad i \in U_j \Leftrightarrow j \in w_i, \quad \text{for } i \in I,$$

which implies that

$$\mathfrak{I} \models \psi(\llbracket \chi_w(\bar{P}) \rrbracket_{w \subseteq [m]}).$$

(\Leftarrow) Suppose that $\mathfrak{I} \models \psi(\llbracket \chi_w(\bar{P}) \rrbracket_{w \subseteq [m]})$. Then there exists sets \bar{U} and $w_i \subseteq [m]$, for $i \in I$, such that

$$\mathfrak{I} \models \psi'(\bar{U}), \quad i \in \llbracket \chi_w(\bar{P}) \rrbracket, \quad \text{and} \quad w_i = \{ j \mid i \in U_j \}.$$

For each $i \in I$, we can therefore choose sets \bar{Q}^i in \mathfrak{A}_i such that

$$\mathfrak{A}_i \models \bigwedge_{j \in w_i} \chi'_j(\bar{P} \upharpoonright A_i, \bar{Q}^i) \wedge \bigwedge_{j \notin w_i} \neg \chi'_j(\bar{P} \upharpoonright A_i, \bar{Q}^i).$$

Setting $\bar{Q} := \bigcup_i \bar{Q}^i$, it follows that

$$\mathfrak{I} \models \psi'(\bar{U}) \quad \text{and} \quad U_j = \llbracket \chi'_j(\bar{P}, \bar{Q}) \rrbracket. \quad \square$$

We have seen above that an ordered sum can be expressed as a generalised sum followed by an interpretation. Therefore, we obtain a composition theorem for ordered sum as an easy application of the one for generalised sums.

Proposition 4.25. *Let \mathfrak{I} be a linear order and let $(\mathfrak{A}_i)_{i \in I}$ and $(\mathfrak{B}_i)_{i \in I}$ be two families of C-coloured linear orders indexed by I . If*

$$\mathfrak{A}_i \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{B}_i, \quad \text{for all } i \in I,$$

then $\sum_{i \in I} \mathfrak{A}_i \equiv_{\text{MSO}}^{\bar{n}} \sum_{i \in I} \mathfrak{B}_i$.

Proof. Let $\varphi \in \text{MSO}^{\circ}_n$, and let φ' and $\chi_0, \dots, \chi_{n-1}$ be the formulae obtained via Theorem 4.24. By assumption, the sets $\llbracket \chi_i \rrbracket$ have the same value when evaluated for the sequence $(\mathfrak{A}_i)_{i \in I}$ and for $(\mathfrak{B}_i)_{i \in I}$. Consequently, we have

$$\begin{aligned} \sum_{i \in \mathfrak{I}} \mathfrak{A}_i \models \varphi & \quad \text{iff} \quad \mathfrak{I} \models \varphi'(\llbracket \chi_0 \rrbracket, \dots, \llbracket \chi_{n-1} \rrbracket) \\ & \quad \text{iff} \quad \sum_{i \in \mathfrak{I}} \mathfrak{B}_i \models \varphi. \end{aligned} \quad \square$$

So far, we have only considered disjoint unions. Let us give an example showing how to extend our approach to certain unions that are not disjoint.

Example. Let \mathfrak{A} be a Σ -structure and $k < \omega$ a constant. We consider a non-disjoint decomposition $A = C \cup \bigcup \mathcal{H}$ of the following form where the set C is called the *center* of the decomposition and the sets in \mathcal{H} its *petals*.

- Every guarded tuple of \mathfrak{A} is entirely contained in C or in one of the $H \in \mathcal{H}$.
- $|H \cap C| \leq k$, for all $H \in \mathcal{H}$.
- $H \cap K \subseteq C$, for all $H \neq K$ in \mathcal{H} .
- For every $H \in \mathcal{H}$, there is some element $c_H \in C$ that belongs to H but not to any other petal $K \in \mathcal{H}$.

For $H \in \mathcal{H}$, let $\bar{a}_H \in A^k$ be an enumeration of $H \cap C$ that starts with the element c_H . (If $H \cap C$ has fewer than k elements, we repeat some of them to obtain a k -tuple.) Let $\mathfrak{C}_o := \mathfrak{A}|_C$ be the restriction of \mathfrak{A} to the set C and define

$$\mathfrak{C} := \langle \mathfrak{C}_o, (U_\theta)_\theta \rangle,$$

where, for every MSO° -theory θ of quantifier rank m , we have added the relation

$$U_\theta := \{ \bar{a}_H \mid \text{Th}^m_{\text{MSO}^{\circ}}(\mathfrak{A}|_H, \bar{a}_H) = \theta \}.$$

Then, for every MSO -formula φ , we can construct an MSO -formula φ' such that

$$\mathfrak{A} \models \varphi \quad \text{iff} \quad \mathfrak{C} \models \varphi'.$$

To see this, note that we can express \mathfrak{A} as a generalised sum followed by an interpretation

$$\mathfrak{A} = \tau \left(\sum_{c \in \mathfrak{C}} \mathfrak{B}_c \right),$$

where

$$\mathfrak{B}_c := \begin{cases} \langle \mathfrak{A}|_H, \bar{a}_H \rangle & \text{if } c = c_H, \\ \mathfrak{A}|_{\{c\}} & \text{if } c \notin \{c_H \mid H \in \mathcal{H}\}. \end{cases}$$

Exercise 4.5. Given a graph $\mathfrak{G} = \langle V, E \rangle$, we call a pair $\langle A, B \rangle$ of subsets $A, B \subseteq V$ a *separation* of \mathfrak{G} if $A \cup B = V$ and there is no edge between a vertex in $A \setminus B$ and one in $B \setminus A$.

Let $\mathfrak{G} = \langle V, E \rangle$ and $\mathfrak{G}' = \langle V', E' \rangle$ be two graphs with separations $\langle A, B \rangle$ and $\langle A', B' \rangle$, respectively. Let $\mathfrak{A}, \mathfrak{B}, \mathfrak{A}'$, and \mathfrak{B}' be the subgraphs of \mathfrak{G} and \mathfrak{G}' induced by the sets A, B, A' , and B' , respectively, let \bar{c} be an enumeration of $A \cap B$, and let \bar{c}' be one of $A' \cap B'$. Suppose that $A \cap B$ and $A' \cap B'$ are finite. Prove that

$$\mathfrak{A}, \bar{c} \equiv^m_{\text{MSO}} \mathfrak{A}', \bar{c}' \quad \text{and} \quad \mathfrak{B}, \bar{c} \equiv^m_{\text{MSO}} \mathfrak{B}', \bar{c}' \quad \text{implies} \quad \mathfrak{G} \equiv^m_{\text{MSO}} \mathfrak{G}'.$$

Hint. Express \mathfrak{G} and \mathfrak{G}' as a generalised sums followed by a quantifier-free interpretation.

5 Operations for First-Order Logic

The operations in the previous section are compatible for a wide variety of logics. In this section we take a look at operations that are FO-compatible, but not MSO-compatible.

Products

Most of the operations introduced in the previous section were based on sums. Here, we present analogous operations that are based on products instead. We start with the simplest one: the binary direct product.

Definition 5.1. Let \mathcal{A} and \mathcal{B} be two Σ -structures. The *direct product* of \mathcal{A} and \mathcal{B} is the Σ -structure $\mathcal{A} \times \mathcal{B}$ with universe $A \times B$ and relations

$$R^{\mathcal{A} \times \mathcal{B}} = \{ \langle \langle a_0, b_0 \rangle, \dots, \langle a_{n-1}, b_{n-1} \rangle \rangle \mid \bar{a} \in R^{\mathcal{A}}, \bar{b} \in R^{\mathcal{B}} \},$$

for $R \in \Sigma$.

Proposition 5.2. $\mathcal{A} \equiv_{\text{FO}}^m \mathcal{A}'$ and $\mathcal{B} \equiv_{\text{FO}}^m \mathcal{B}'$ implies $\mathcal{A} \times \mathcal{B} \equiv_{\text{FO}}^m \mathcal{A}' \times \mathcal{B}'$.

Proof. For every FO_m -formula $\varphi(\bar{x})$, we construct two finite sequences of FO_m -formulae $\psi_0(\bar{x}), \dots, \psi_{l-1}(\bar{x})$ and $\vartheta_0(\bar{x}), \dots, \vartheta_{l-1}(\bar{x})$ such that

$$\begin{aligned} \mathcal{A} \times \mathcal{B} \models \varphi(\langle a_0, b_0 \rangle, \dots, \langle a_{n-1}, b_{n-1} \rangle) \\ \text{iff } \mathcal{A} \models \psi_i(\bar{a}) \text{ and } \mathcal{B} \models \vartheta_i(\bar{b}), \text{ for some } i < l. \end{aligned}$$

We proceed by induction on φ . If φ is of the form $x = y$ or \bar{R} , we can take

$$\psi_0 := \varphi \text{ and } \vartheta_0 := \varphi.$$

For the inductive step, suppose that we have already constructed the sequences of formulae $\psi_0, \dots, \psi_{l-1}$, $\vartheta_0, \dots, \vartheta_{l-1}$ corresponding to φ , and $\psi'_0, \dots, \psi'_{l'-1}$, $\vartheta'_0, \dots, \vartheta'_{l'-1}$ corresponding to φ' . Then the sequences corresponding to $\varphi \vee \varphi'$ are

$$\psi_0, \dots, \psi_{l-1}, \psi'_0, \dots, \psi'_{l'-1} \text{ and } \vartheta_0, \dots, \vartheta_{l-1}, \vartheta'_0, \dots, \vartheta'_{l'-1}.$$

For $\varphi \wedge \varphi'$, we can take

$$\psi_i \wedge \psi'_j, \text{ and } \vartheta_i \wedge \vartheta'_j, \text{ for all } i < l \text{ and } j < l'.$$

For $\neg\varphi$, we use all formulae of the form

$$\psi_w^- \text{ and } \vartheta_w^-, \text{ for } w \subseteq [l],$$

where

$$\psi_w^- := \bigwedge_{i \in [l] \setminus w} \neg\psi_i \text{ and } \vartheta_w^- := \bigwedge_{i \in w} \neg\vartheta_i.$$

Finally, for $\exists x\varphi$, we can take

$$\exists x\psi_0, \dots, \exists x\psi_{l-1} \text{ and } \exists x\vartheta_0, \dots, \exists x\vartheta_{l-1}. \quad \square$$

Example. The *ordered product* of two linear orders $\mathcal{A} = \langle A, \leq^{\mathcal{A}} \rangle$ and $\mathcal{B} = \langle B, \leq^{\mathcal{B}} \rangle$ is the linear order $\mathcal{A} \cdot \mathcal{B}$ with universe $A \times B$ and ordering

$$\langle a, b \rangle \leq \langle c, d \rangle \quad \text{iff} \quad b < d \quad \text{or} \quad b = d \text{ and } a \leq c.$$

We can express this product using the direct product $\mathcal{A} \times \mathcal{B}$ followed by an FO-interpretation.

Exercise 5.1. Find FO-interpretations ρ, σ, τ such that

$$\wp(\mathcal{A} \oplus \mathcal{B}) = \tau(\rho(\wp(\mathcal{A})) \times \sigma(\wp(\mathcal{B}))). \quad \square$$

Finite Powers

Our analogue for the copying operation is the finite power of a structure.

Definition 5.3. Let \mathcal{A} be a Σ -structure and $k < \omega$. The *k-th power* of \mathcal{A} is the structure \mathcal{A}^k obtained from the k -fold direct product $\mathcal{A} \times \dots \times \mathcal{A}$ by adding the relations

$$I_{ij} := \{ \langle \bar{a}, \bar{b} \rangle \in A^k \times A^k \mid a_i = b_j \}, \text{ for } i, j < k.$$

Proposition 5.4. $\mathcal{A} \equiv_{\text{FO}}^m \mathcal{B}$ implies $\mathcal{A}^k \equiv_{\text{FO}}^{mk} \mathcal{B}^k$.

Proof. Given an FO_m -formula $\varphi(x_0, \dots, x_{n-1})$, we construct an FO_{mk} -formula $\varphi'(\bar{x}_0, \dots, \bar{x}_{n-1})$ such that

$$\mathcal{A}^k \models \varphi(\bar{a}_0, \dots, \bar{a}_{n-1}) \quad \text{iff} \quad \mathcal{A} \models \varphi'(\bar{a}_0, \dots, \bar{a}_{n-1}).$$

We construct φ' by induction on φ .

$$\begin{aligned} (x = y)' &:= \bigwedge_{i < k} x_i = y_i, & (I_{ij}xy)' &:= x_i = y_j, \\ (Rx^0 \dots x^{n-1})' &:= \bigwedge_{i < k} Rx_i^0 \dots x_i^{n-1}, & (\neg\varphi)' &:= \neg\varphi', \\ (\varphi \wedge \psi)' &:= \varphi' \wedge \psi', & (\exists x\varphi)' &:= \exists \bar{x}\varphi', \\ (\varphi \vee \psi)' &:= \varphi' \vee \psi', & (\forall x\varphi)' &:= \forall \bar{x}\varphi'. \end{aligned} \quad \square$$

Exercise 5.2. Find an FO-interpretation τ such that

$$\wp(\text{copy}_k(\mathfrak{A})) = \tau(\wp(\mathfrak{A})^k).$$

Finite powers are frequently combined with first-order interpretations.

Definition 5.5. Let $k < \omega$. A k -dimensional FO-interpretation is an operation of the form

$$\tau = \rho \circ \tau_o \circ (-)^k,$$

where ρ is a quotient operation and τ_o a normal FO-interpretation.

Remark. We can compactly specify a k -dimensional FO-interpretation by a list of formula

$$\langle \delta(\bar{x}), \varepsilon(\bar{x}, \bar{y}), (\varphi_R(\bar{x}_o, \bar{x}_1, \dots))_R \rangle$$

where each tuple $\bar{x}, \bar{y}, \bar{x}_o, \dots$ consists of k -variables. The formula δ defines the universe of the new structure, ε defines the new equality relation, and the formulae φ_R define the relations.

Examples. (a) Let $\mathfrak{R} = \langle \mathbb{R}, +, \cdot, o, 1 \rangle$ and $\mathfrak{C} = \langle \mathbb{C}, +, \cdot, o, 1 \rangle$ be the fields of real and complex numbers. There exists a 2-dimensional FO-interpretation τ mapping \mathfrak{R} to \mathfrak{C} which is given by

$$\begin{aligned} \delta(xx') &:= \text{true}, \\ \varepsilon(xx', yy') &:= \text{true}, \\ \varphi_+(xx', yy', zz') &:= z = x + y \wedge z' = x' + y', \\ \varphi_*(xx', yy', zz') &:= z = xy - x'y' \wedge z' = x'y + xy', \\ \varphi_o(xx') &:= x = o \wedge x' = o, \\ \varphi_1(xx') &:= x = 1 \wedge x' = o. \end{aligned}$$

(b) There exists a 2-dimensional FO-interpretation mapping $\mathfrak{N} := \langle \mathbb{N}, + \rangle$ to $\mathfrak{Z} := \langle \mathbb{Z}, + \rangle$. It is given by the formulae

$$\begin{aligned} \delta(xx') &:= \text{true}, \\ \varepsilon(xx', yy') &:= x + y' = y' + x, \\ \varphi_+(xx', yy', zz') &:= z = x + y \wedge z' = x' + y'. \end{aligned}$$

Generalised Products

Similarly to generalised sums for MSO, we can define a general version of a product for FO where the index set is equipped with additional relations.

Definition 5.6. Let I be a set, \mathfrak{I} a Γ -structure with universe $\wp(I)$, and $(\mathfrak{A}_i)_{i \in I}$ a family of Σ -structure. The *generalised product*

$$\prod_{i \in \mathfrak{I}} \mathfrak{A}_i$$

is obtained from the $(\Sigma + \Gamma + \{\uparrow\})$ -structure with universe

$$\sum_{K \subseteq I} \prod_{i \in K} A_i \cong \prod_{i \in I} (A_i + \{\perp\})$$

relations

$$\begin{aligned} R &:= \{ \langle a_o, \dots, a_{n-1} \rangle \mid \text{dom}(a_o) = \dots = \text{dom}(a_{n-1}) \text{ and} \\ &\quad \langle a_o(i), \dots, a_{n-1}(i) \rangle \in R^{\mathfrak{A}_i}, \text{ for all } i \in \text{dom}(a_o) \}, \end{aligned}$$

for every $R \in \Sigma$,

$$S := \{ \langle a_o, \dots, a_{n-1} \rangle \mid \langle \text{dom}(a_o), \dots, \text{dom}(a_{n-1}) \rangle \in S^{\mathfrak{I}} \},$$

for every $S \in \Gamma$, and one binary function

$$a \uparrow b := a|_{\text{dom}(a) \cap \text{dom}(b)}.$$

Definition 5.7. Let \mathfrak{I} be a Γ -structure and $(\mathfrak{A}_i)_{i \in I}$ a family of Σ -structures. Given an FO-formula $\chi(\bar{x})$ and a tuple \bar{a} in $\prod_{i \in I} (A_i + \perp)$, we define

$$\llbracket \chi(\bar{a}) \rrbracket := \{ i \in I \mid \mathfrak{A}_i \oplus 1 \models \chi(\bar{a}_i) \},$$

where \bar{a}_i denotes the projection of \bar{a} to the i -th component, and 1 denotes the 1-element Σ -structure where all relations are non-empty, i.e., the terminal object in the category of all Σ -structures.

The following composition theorem and its proof closely follow that for generalised sums. In fact, the case of generalised sums can be derived from the one for products using the power-set operation.

Theorem 5.8 (Feferman, Vaught). *For every formula $\varphi(\bar{x}) \in \text{FO}_m$, there exist formulae $\psi(\bar{x}, \bar{y}) \in \text{FO}$ and $\chi_o(\bar{x}), \dots, \chi_n(\bar{x}) \in \text{FO}_m$ such that*

$$\prod_{i \in \mathfrak{S}} \mathfrak{A}_i \models \varphi(\bar{a}) \quad \text{iff} \quad \langle \mathfrak{S}, \subseteq \rangle \models \psi(\llbracket \chi_o(\bar{a}) \rrbracket, \dots, \llbracket \chi_{n-1}(\bar{a}) \rrbracket).$$

Proof. We prove the claim by induction on φ . In the formulae χ_i (which are evaluated in a disjoint union $\mathfrak{A}_i \oplus 1$), we will for readability use the notation $x = \perp$ in instead of $\text{Right}(x)$. If $\varphi = (x_i = x_j)$, we can set

$$\psi(p) := (p = I) \quad \text{with} \quad \chi(\bar{x}) := (x_i = x_j).$$

If $\varphi = (x_k = x_i \upharpoonright x_j)$, we can set

$$\begin{aligned} \psi(p, p', p'', p_*) &:= p = p' \cap p'' \wedge p = p_*, \\ \chi(\bar{x}) &:= (x_k \neq \perp), \quad \chi'(\bar{x}) := (x_i \neq \perp), \\ \chi''(\bar{x}) &:= (x_j \neq \perp), \quad \chi_*(\bar{x}) := (x_k = x_i). \end{aligned}$$

If $\varphi = (x_i \sim x_j)$, we can set

$$\psi(p, p') := p = p', \quad \chi(\bar{x}) := (x_i \neq \perp), \quad \chi'(\bar{x}) := (x_j \neq \perp).$$

If $\varphi = Rx_{i_o} \dots x_{i_{k-1}}$ with $R \in \Sigma$, we can set

$$\psi(p) := (p = I) \quad \text{and} \quad \chi_o(\bar{x}) := Rx_{i_o} \dots x_{i_{k-1}}.$$

If $\varphi = Sx_{i_o} \dots x_{i_{k-1}}$ with $S \in \Gamma$, we can set

$$\psi(p_o, \dots, p_{k-1}) := Sp_o \dots p_{k-1},$$

and $\chi_j(\bar{x}) := (x_{i_j} \neq \perp)$, for $j < k$.

For boolean operations, we can simply take the corresponding boolean combination of the formulae ψ (after renaming the predicates p_i to make

them disjoint), and we take the union of all formulae χ_i . For instance, for $\varphi = \varphi_o \wedge \varphi_1$, we can set

$$\psi(\bar{p}_o, \bar{p}_1) := \psi_o(\bar{p}_o) \wedge \psi_1(\bar{p}_1).$$

Finally, suppose that $\varphi = \exists x' \varphi'(\bar{x}, x')$. Let

$$\psi'(\bar{y}) \quad \text{and} \quad \chi'_o(\bar{x}x'), \dots, \chi'_{n-1}(\bar{x}x')$$

be the formulae obtained by applying the inductive hypothesis to φ' . We set

$$\begin{aligned} \psi(\bar{y}) &:= \exists \bar{z} \left[\psi'(\bar{z}) \wedge \bigcup_{w \subseteq [m]} \left[y_w \cap \bigcap_{j \in w} z_j \setminus \bigcup_{j \notin w} z_j \right] = I \right], \\ \chi_w(\bar{x}) &:= \exists x' \left[\bigwedge_{j \in w} \chi'_j(\bar{x}x') \wedge \bigwedge_{j \in [m] \setminus w} \neg \chi'_j(\bar{x}x') \right], \quad \text{for } w \subseteq [n]. \end{aligned}$$

It remains to check that these formulae have the desired properties. We have

$$\begin{aligned} \prod_{i \in \mathfrak{S}} \mathfrak{A}_i \models \exists x' \varphi'(\bar{a}, x') \\ \text{iff} \quad \prod_{i \in \mathfrak{S}} \mathfrak{A}_i \models \varphi'(\bar{a}, b), \quad \text{for some } b, \\ \text{iff} \quad \langle \mathfrak{S}, \subseteq \rangle \models \psi'(\llbracket \chi'_o(\bar{a}, b) \rrbracket, \dots, \llbracket \chi'_{n-1}(\bar{a}, b) \rrbracket), \quad \text{for some } b. \end{aligned}$$

We claim that the latter is equivalent to

$$\langle \mathfrak{S}, \subseteq \rangle \models \psi(\llbracket \chi_w(\bar{a}) \rrbracket_{w \subseteq [m]}).$$

(\Rightarrow) Given an element b as above, we set

$$p_j := \llbracket \chi'_j(\bar{a}, b) \rrbracket \quad \text{and} \quad w_i := \{ j \mid i \in p_j \}, \quad \text{for } j < n \text{ and } i \in I.$$

Then

$$\langle \mathfrak{S}, \subseteq \rangle \models \psi'(\bar{p}), \quad i \in \llbracket \chi_{w_i}(\bar{a}) \rrbracket, \quad \text{and} \quad i \in p_j \Leftrightarrow j \in w_i,$$

for $j < m$ and $i \in I$, which implies that

$$i \in \llbracket \chi_{w_i}(\bar{a}) \rrbracket \cap \bigcap_{j \in w_i} p_j \setminus \bigcup_{j \notin w_i} p_j, \quad \text{for all } i \in I.$$

Hence,

$$\langle \mathfrak{I}, \subseteq \rangle \models \psi(\llbracket \chi_w(\bar{a}) \rrbracket_{w \subseteq [m]}).$$

(\Leftarrow) Suppose that $\mathfrak{I} \models \psi(\llbracket \chi_w(\bar{a}) \rrbracket_{w \subseteq [m]})$. Then there exists sets \bar{p} such that

$$\langle \mathfrak{I}, \subseteq \rangle \models \psi'(\bar{p})$$

and, for every $i \in I$, there is some $w_i \subseteq [m]$ such that

$$i \in \llbracket \chi_{w_i}(\bar{a}) \rrbracket \cap \bigcap_{j \in w_i} p_j \setminus \bigcup_{j \notin w_i} p_j.$$

This implies that $w_i = \{j \mid i \in p_j\}$. The fact that $i \in \llbracket \chi_{w_i}(\bar{a}) \rrbracket \cap$ therefore implies that there is some element $b_i \in A_i + 1$ such that

$$\mathfrak{A}_i \models \bigwedge_{j \in w_i} \chi'_j(\bar{a}_i, b_i) \wedge \bigwedge_{j \notin w_i} \neg \chi'_j(\bar{a}_i, b_i).$$

Setting $b := (b_i)_i$, it follows that

$$\langle \mathfrak{I}, \subseteq \rangle \models \psi'(\bar{p}) \quad \text{and} \quad p_j = \llbracket \chi'_j(\bar{a}, b) \rrbracket. \quad \square$$

Example. We can use generalised products to construct ultraproducts and, more generally, quotients of products by arbitrary filters. To this end, let $F \subseteq \wp(I)$ be a filter on I . Recall that

$$\prod_{i \in I} \mathfrak{A}_i / F$$

denotes the quotient of the direct product $\prod_{i \in I} \mathfrak{A}_i$ by the relation

$$a \approx b \quad : \text{iff} \quad \{i \in I \mid a_i = b_i\} \in F.$$

Using the index structure $\mathfrak{I} := \langle \wp(I), F \rangle$ we can define such a product in the generalised product

$$\prod_{i \in \mathfrak{I}} \mathfrak{A}_i$$

using an FO-interpretation defining the relation \approx , followed by a quotient by \approx . \square

Notes

Composition arguments were popularised by Shelah [193], although their use is much older. For instance, the composition theorem for generalised products is from [80]. Good surveys include [136, 22, 60].

The logic MSO° was invented by Thomas to simplify composition arguments [203], while guarded second-order logic was introduced in [89]. The latter generalises a logic for graphs defined by Courcelle, which is usually called MSO_2 or MS_2 , cf. [60].

The power-set construction was first systematically investigated in [52]. The example with the composition theorem for overlapping unions is taken from [67]. Exercise 1.2 on axiomatisations of grids was inspired by [140].

II Finite Words

1 Words and Languages

OUR FIRST DEEPER STUDY of the expressive power of monadic second-order logic concerns the definability of languages of finite or infinite words. We start by setting up our terminology regarding such languages.

Definition 1.1. (a) A binary relation $\leq \subseteq A \times A$ is a *partial ordering* if it has the following properties:

Reflexivity. $a \leq a$, for all $a \in A$.

Anti-Symmetry. $a \leq b$ and $b \leq a$ implies $a = b$, for all $a, b \in A$.

Transitivity. $a \leq b \leq c$ implies $a \leq c$, for all $a, b, c \in A$.

A *partial order* is a structure $\langle A, \leq \rangle$ where \leq is a partial ordering on A .

(b) A *linear order* is a partial order $\langle A, \leq \rangle$ where

$$a \leq b \text{ or } b \leq a, \quad \text{for all } a, b \in A.$$

(c) A linear order $\langle A, \leq \rangle$ is a *well-order* if every non-empty subset $X \subseteq A$ has a minimal element, that is, if there exists no infinite, strictly descending sequence $a_0 > a_1 > a_2 > \dots$. An *ordinal* is the isomorphism type of a well-order.

(d) We denote by ω the first infinite ordinal. It is isomorphic to the linear order of the natural numbers. J

Formal language theory deals with linear orders whose positions are labelled with elements of a given set Σ .

Definition 1.2. (a) An *alphabet* is a finite set Σ whose elements are called *letters*.

(b) A (finite) *word* over an alphabet Σ is a finite sequence

$$w = \langle a_0, \dots, a_{n-1} \rangle$$

of elements $a_i \in \Sigma$. The *empty word* is the empty sequence $\langle \rangle$. We denote the *length* of a word w by $|w|$. We write Σ^* for the set of all finite words over Σ , and Σ^+ for the set of all non-empty finite words.

(c) An ω -*word* over an alphabet Σ is an infinite sequence $w = (a_i)_{i < \omega}$ of elements $a_i \in \Sigma$. The set of all ω -words is denoted by Σ^ω . We also set $\Sigma^\infty := \Sigma^* \cup \Sigma^\omega$.

(d) A *language* is a set of words or a set of ω -words. \lrcorner

In order to define languages in some logic, we have to encode words as structures.

Definition 1.3. Let Σ be an alphabet.

(a) We can associate with every word $w \in \Sigma^*$ a relational structure

$$\langle W, \leq, (P_a)_{a \in \Sigma} \rangle$$

over the signature $\{\leq\} \cup \{P_a \mid a \in \Sigma\}$ where W is the set of positions of w , \leq is the ordering of the positions, and P_a is a set containing all positions labelled by the letter a . Structures of this form are called *word structures*.

(b) A language $K \subseteq \Sigma^\infty$ is *definable* in a logic L if there exists a formula $\varphi \in L$ such that

$$w \in K \quad \text{iff} \quad w \models \varphi.$$

(In the right-hand side, we have identified w with the associated word structure.) \lrcorner

Example. The language of all words over the alphabet $\Sigma := \{a, b, c\}$ with an even number of letters a can be defined by the MSO-formula

$$\begin{aligned} \varphi := & \exists X \left[\forall x \forall y [x < y \wedge P_a x \wedge P_a y \wedge \neg \exists z (x < z \wedge z < y \wedge P_a z) \right. \\ & \left. \rightarrow (Xx \leftrightarrow \neg Xy) \right] \\ & \wedge \forall x [P_a x \wedge \neg \exists y (y < x \wedge P_a y) \rightarrow Xx] \\ & \wedge \forall x [P_a x \wedge \neg \exists y (x < y \wedge P_a y) \rightarrow \neg Xx] \end{aligned} \quad \lrcorner$$

Exercise 1.1. Prove that the language in the above example is not FO-definable. \lrcorner

Definition 1.4. Let Σ be an alphabet.

(a) If $w \in \Sigma^\infty$ and $i \leq k < |w|$, we write $w(i)$ for the element of w at position i and

$$w[i, k] := \langle w(i), w(i+1), \dots, w(k-1) \rangle$$

for the *factor* of w from position i to $k-1$.

(b) The *concatenation* of two words $u \in \Sigma^*$ and $v \in \Sigma^\infty$ is the word $u \hat{\ } v$ that consists of the elements of u , followed by the elements of v . Formally,

$$(u \hat{\ } v)(i) = \begin{cases} u(i) & \text{if } i < |u|, \\ v(i - |u|) & \text{if } i \geq |u|. \end{cases}$$

Frequently, we omit the symbol $\hat{\ }$ and simply write uv instead.

(c) A word u is a *prefix* of a word $w \in \Sigma^\infty$ if $w = u \hat{\ } v$, for some $v \in \Sigma^\infty$. Similarly, u is a *suffix* of w if $w = v \hat{\ } u$, for some v . Finally, u is a *factor* of w if $w = x \hat{\ } u \hat{\ } y$, for some x, y . \lrcorner

2 Semigroups and Green's Relations

When studying languages of words, an algebraic approach based on semigroup theory sometimes proves to be quite convenient. The starting point is the observation that the set Σ^* of all finite words together with the concatenation operation $\hat{\ }$ forms a monoid.

Definition 2.1. (a) A *semigroup* is a structure $\mathfrak{S} = \langle S, \cdot \rangle$ where the multiplication $\cdot : S \times S \rightarrow S$ is associative:

$$a \cdot (b \cdot c) = (a \cdot b) \cdot c.$$

We usually omit the dot and simply write ab for the product.

(b) A *monoid* $\mathfrak{M} = \langle M, \cdot, 1 \rangle$ is a semigroup with a neutral element 1 , i.e., an element satisfying

$$1 \cdot a = a = a \cdot 1, \quad \text{for all } a \in M.$$

To each semigroup \mathfrak{S} we can associate a monoid \mathfrak{S}^1 by adding a new neutral element 1 .

(c) An element e of a semigroup is *idempotent* if $ee = e$.

Examples. (a) For every alphabet Σ , we have the *free semigroup* $\langle \Sigma^+, \cdot \rangle$ over Σ and the *free monoid* $\langle \Sigma^*, \cdot, \langle \rangle \rangle$ over Σ . The only idempotent element is the empty word $\langle \rangle$.

(b) The natural numbers form the monoid $\langle \mathbb{N}, +, 0 \rangle$ and the monoid $\langle \mathbb{N}, \cdot, 1 \rangle$. The former is called *Presburger arithmetic*, the latter *Skolem arithmetic*.

(c) The set of functions $f : X \rightarrow X$ forms a monoid $\langle X^X, \circ, \text{id} \rangle$ where \circ denotes function composition. A function $f : X \rightarrow X$ is idempotent if, and only if, it is a projection, i.e., $f \upharpoonright \text{rng } f = \text{id}$.

(d) Similarly, the set of all binary relations $R \subseteq X \times X$ forms a monoid $\langle \mathcal{P}(X \times X), \circ, \text{id} \rangle$ where \circ denotes the composition of relations and id is the diagonal. If $R \subseteq X \times X$ is reflexive and transitive, it is idempotent. Conversely, every idempotent element $R \subseteq X \times X$ is transitive, but not necessarily reflexive.

(e) Every semilattice $\langle L, \sqcup \rangle$ is a semigroup where all elements are idempotent.

Exercise 2.1. Prove that every semigroup has at most one neutral element.

We will be mostly dealing with finite semigroups or, more generally, ones that are *cyclically finite*.

Definition 2.2. A semigroup \mathfrak{S} is *cyclically finite* if

$$\langle\langle a \rangle\rangle_{\mathfrak{S}} := \{ a^n \mid 0 < n < \omega \} \text{ is finite, for every } a \in S.$$

While a semigroup does not need to have a neutral element, a cyclically finite one will always contain at least some idempotent (which can be regarded as a weak form of a neutral element).

Lemma 2.3. If \mathfrak{S} is a cyclically finite semigroup, every element $a \in S^1$ has an idempotent power.

Proof. The unit 1 is itself idempotent. Let $a \in S$. As \mathfrak{S} is cyclically finite, there are exponents $0 < i < j < \omega$ such that $a^i = a^j$. Set $k := j - i$. Iterating the equation $a^i = a^{i+k}$, we obtain

$$a^i = a^{i+k} = a^{i+2k} = \dots = a^{i+ik}.$$

Hence,

$$a^{ik} = a^i a^{ik-i} = a^{i+ik} a^{ik-i} = a^{ik+ik}$$

and a^{ik} is our desired idempotent element. \square

Example. If X is a non-empty, finite set, this lemma tells us that there is some number n such that, for every function $f : X \rightarrow X$, the n -th power $f^n : X \rightarrow X$ is a projection. In particular, f^n has a fixed point.

Exercise 2.2. Let \mathfrak{S} be a finite semigroup. Prove that there exists a number $0 < n < \omega$ such that a^n is idempotent, for every $a \in S$.

The main technical result the following material is based on, is the following property of cyclically finite semigroups.

Lemma 2.4. Let \mathfrak{S} be a cyclically finite semigroup, and $a \in S$ and $s, t, u, v \in S^1$ elements. Then

$$a = stauv \quad \text{implies} \quad xta = a = auy, \quad \text{for some } x, y \in S^1.$$

Proof. Iterating the equation $a = (st)a(uv)$, we obtain

$$a = (st)^i a(uv)^i, \quad \text{for all } i < \omega.$$

By Lemma 2.3, there is some exponent $0 < n < \omega$ such that $(st)^n$ and $(uv)^n$ are both idempotents. It follows that

$$(st)^n a = (st)^n (st)^n a(uv)^n = (st)^n a(uv)^n = a.$$

Similarly, we obtain $a(uv)^n = a$. Hence, we can set $x := (st)^{n-1}s$ and $y := v(uv)^{n-1}$. \square

Green's Relations

There exists a rich structure theory for cyclically finite semigroups (and more generally for compact semigroups). The starting point are the following divisibility relations. There are several of them since semigroups do not need to be commutative and we therefore have to distinguish between dividing from the left, from the right, or from somewhere in the middle.

Definition 2.5. *Green's relations* consist of the divisibility relations

$$\begin{aligned} a \leq_L b &: \text{iff } a = xb \quad \text{for some } x \in S^1, \\ a \leq_R b &: \text{iff } a = bx \quad \text{for some } x \in S^1, \\ a \leq_J b &: \text{iff } a = xby \quad \text{for some } x, y \in S^1. \end{aligned}$$

together with the associated equivalence relations

$$\begin{aligned} a \equiv_L b &: \text{iff } a \leq_L b \quad \text{and} \quad b \leq_L a, \\ a \equiv_R b &: \text{iff } a \leq_R b \quad \text{and} \quad b \leq_R a, \\ a \equiv_J b &: \text{iff } a \leq_J b \quad \text{and} \quad b \leq_J a. \end{aligned}$$

Furthermore, we set


$$a \equiv_H b : \text{iff } a \equiv_L b \quad \text{and} \quad a \equiv_R b.$$

We call the equivalence classes of the relations $\equiv_L, \equiv_R, \equiv_J, \equiv_H$ *L-classes*, *R-classes*, *J-classes*, and *H-classes*, respectively, and we use the notation $[a]_L, [a]_R, [a]_J$, and $[a]_H$ for the class of a . \square

Examples. (a) Let $M = \{1, a, b, ab, ba, aba, o\}$ be the monoid with unit 1, zero o and relations

$$aa = 1, \quad bab = o, \quad bb = o.$$

The Green's relations are depicted in the following schema:





Here each field represents a single H-class, each group of fields a J-class, each column inside a group an L-class, and each row an R-class. So we have three J-classes

$$\{o\} \leq_J \{b, ba, ab, aba\} \leq_J \{1, a\},$$

where the middle J-class is divided into two L-classes: $\{ba, aba\}$ and $\{b, ab\}$; and into two R-classes: $\{b, ba\}$ and $\{ab, aba\}$. The only non-trivial H-class is $\{1, a\}$.

(b) In the monoid of all relations $R \subseteq [2] \times [2]$ we have the following classes.





Again there is only one non-trivial H-class.

The diagram in the previous example is typical for Green's relations. The next series of lemmas shows that it is always the case that every J-class consists of a rectangular grid of H-classes where the columns are the L-classes and the rows the R-classes.

Lemma 2.6. *Let \mathcal{S} be a semigroup and $a, b, c \in S$.*

- (a) $a \leq_L b$ implies $ac \leq_L bc$.
- (b) $a \leq_R b$ implies $ca \leq_R cb$.
- (c) $\leq_L \circ \leq_R = \leq_R \circ \leq_L$.

Proof. (a) $a = xb$ implies $ac = xbc$.

(b) $a = bx$ implies $ca = cbx$.

(c) Suppose that $a \leq_L b \leq_R c$. Then there are elements $x, y \in S^1$ such that $a = xb$ and $b = cy$. Hence, $a = xcy \leq_R xc \leq_L c$. Therefore, $\leq_L \circ \leq_R \subseteq \leq_R \circ \leq_L$. The other inclusion follows in the same way. \square

Proposition 2.7. *Let \mathcal{S} be a finite semigroup. Then*

$$\equiv_J = \equiv_L \circ \equiv_R.$$

Proof. Suppose that $a \equiv_L c \equiv_R b$. Then $a \leq_L c \leq_R b$ and $a \geq_L c \geq_R b$ which, by Lemma 2.6 (c), implies that $a \leq_J b$ and $a \geq_J b$.

Conversely, suppose that $a \equiv_J b$. Then there are elements $s, t, u, v \in S^1$ such that $a = sbu$ and $b = tav$. Hence, $a = stauv$ and it follows by Lemma 2.4 that $xta = a = auv$, for some $x, y \in S^1$. In particular, $a \leq_L ta$ and $a \leq_R au$.

Since, trivially, $ta \leq_L a$ and $au \leq_R a$, it follows that $ta \equiv_L a \equiv_R au$. Using Lemma 2.6 (a) we obtain $b = tau \equiv_L au \equiv_R a$. \square

Since L-equivalence and R-equivalence both imply J-equivalence, we can partition every J-class into L-classes and into R-classes. By the above lemma, these two partitions intersect such that every L-class meets every R-class. Hence, we always obtain a picture as in the above example. The following lemma states that all the L-classes have the same size and the same holds for the R-classes.

Lemma 2.8 (Green). *Let \mathcal{S} be a cyclically finite semigroup and $a, b \in S$. Any two elements $s, t \in S^1$ with $b = as$ and $a = bt$ induce bijections*

$$\varphi : [a]_L \rightarrow [b]_L : x \mapsto xs \quad \text{and} \quad \psi : [b]_L \rightarrow [a]_L : x \mapsto xt$$

which are inverses of each other.

Proof. Let $c \equiv_L a$. First, note that $cs \equiv_L as = b$. Hence, φ maps $[a]_L$ to $[b]_L$. Furthermore, if $x \in S^1$ is an element with $c = xa$, then

$$\psi(\varphi(c)) = cst = xast = xbt = xa = c.$$

In the same way, we can show that ψ maps $[b]_L$ to $[a]_L$ and that $\varphi \circ \psi$ is the identity on $[b]_L$. \square

Remarks. (a) Applying this result to the dual semigroup \mathcal{S}^{op} (with product $a \cdot^{\text{op}} b := ba$) gives the analogous statement for R-classes.

(b) The maps φ and ψ above preserve H-equivalence (see the proof of Lemma 2.14 below). \square

The Structure of J-Classes

Next, let us take a closer look at some properties of the J-relation that turn out to be useful when studying FO-definability.

Lemma 2.9. *Let \mathcal{S} be a cyclically finite semigroup and $a, b \in S$.*

- (a) $b \leq_J ab$ implies $b \equiv_L ab$.

(b) $a \leq_J ab$ implies $a \equiv_R ab$.

(c) If $a \equiv_J ab \equiv_J b$, there is some idempotent $e \in S$ with $a \equiv_L e \equiv_R b$.

Proof. (a) Suppose that $b \leq_J ab$. Then there are elements $s, u \in S^1$ such that $b = sabu$. By Lemma 2.4, it follows that $xab = b$, for some $x \in S^1$. Hence, $b \leq_L ab$. Since, trivially, $ab \leq_L b$, the claim follows.

(b) follows in exactly the same way.

(c) Suppose that $a \equiv_J ab \equiv_J b$. By (a) and (b), it follows that $a \equiv_R ab \equiv_L b$. Fix an element $s \in S^1$ with $abs = a$. We have seen in Lemma 2.8 that the function $\varphi : x \mapsto xb$ is a bijection between the L-classes of a and ab with inverse $\psi : x \mapsto xs$. Setting $e := \psi(b)$, it therefore follows that $e \equiv_L a$ and

$$ee = e\psi(b) = ebs = \varphi(e)s = \varphi(\psi(b))s = bs = \psi(b) = e.$$

Thus, e is idempotent. Finally, we have $e \equiv_R b$ since, by definition, $\varphi(x) \leq_R x$ and $\psi(x) \leq_R x$, for all x . \square

Let us note a few consequences of this lemma that will turn out to be particularly useful.

Corollary 2.10. *Let J be a J-class of a cyclically finite semigroup \mathcal{S} and let $a, b \in J$ be elements such that $ab \in J$. Then*

$$a \equiv_R ab \equiv_L b$$

and there exists an idempotent $e \in J$ such that

$$a \equiv_L e \equiv_R b.$$

Corollary 2.11. *Let \mathcal{S} be a cyclically finite semigroup and $a, b, c \in S$.*

(a) $ab \equiv_J b \equiv_J bc$ implies $abc \equiv_J b$.

(b) $ab = a \equiv_J b$ implies $ac = a \Leftrightarrow bc = b$.

Proof. (a) By Lemma 2.9 (b), we have $b \equiv_R bc$. This implies that $abc \equiv_R ab$. In particular, $abc \equiv_J ab \equiv_J b$.

(b) (\Leftarrow) $bc = b$ implies $ac = abc = ab = a$.

(\Rightarrow) By Lemma 2.9 (a), $b \equiv_J ab$ implies $b \equiv_L ab = a$. Hence, $b = xa$, for some $x \in S^1$, and we have $bc = xac = xa = b$. \square

We obtain the following structure result for J-classes.

Definition 2.12. Let \mathcal{S} be a semigroup. An element $a \in S$ is *regular* if there exists some $c \in S$ with $aca = a$. A J-class is *regular* if it contains a regular element. \downarrow

Proposition 2.13. *Let J be a J-class of a cyclically finite semigroup \mathcal{S} . The following statements are equivalent.*

- (1) J is regular.
- (2) J contains an idempotent.
- (3) Every L-class in J contains an idempotent.
- (4) There are elements $a, b \in J$ such that $ab \in J$.

Proof. (3) \Rightarrow (2) is trivial and (4) \Rightarrow (2) follows by Lemma 2.9.

(2) \Rightarrow (1) If $e \in J$ is idempotent, then $eee = e$. Hence, e is regular.

(1) \Rightarrow (4) Suppose that there are elements $a \in J$ and $c \in S$ with $aca = a$. Then $a \leq_J ac$ and $acac = ac$. Since $ac \leq_J a$ holds trivially, it follows that $ac \in J$ and that it is an idempotent.

(2) \Rightarrow (3) Let $e \in J$ be idempotent. Given an element $a \in J$, we have to find an idempotent that is L-equivalent to a . By Proposition 2.7, there exists some b with $e \equiv_R b \equiv_L a$. Consequently, there are $x, y \in S^1$ with $ex = b$ and $e = by$. It follows that $eb = eex = ex = b$ and $yb = yeb = ybyb$. Hence, yb is idempotent. Furthermore, $byb = eb = b$ implies $yb \leq_L b$. Since we trivially have $b \leq_L yb$, it follows that $yb \equiv_L b \equiv_L a$. \square

The Structure of H-Classes

We conclude this section with a similar look at H-classes. Let us start by noting that all H-classes in a given J-class have the same size.

Lemma 2.14. *Let \mathcal{S} be a cyclically finite semigroup and let H, H' be two H-classes that belong to the same J-class. Then there are elements $s, t \in S^1$ such that the function*

$$\varphi : H \rightarrow H' : x \mapsto sxt$$

is bijective.

Proof. Fix $a \in H$ and $a' \in H'$. Since $a \equiv_J a'$, it follows by Proposition 2.7 that $a \equiv_L b \equiv_R a'$ for some b . Fix elements $s, t \in S^1$ with $b = sa$ and $a' = bt$. By Lemma 2.8 and the corresponding statement for R-classes, the functions

$$\psi : [a]_L \rightarrow [b]_L : x \mapsto sx \quad \text{and} \quad \psi' : [b]_R \rightarrow [a']_R : x \mapsto xt$$

are bijective. To show that $\varphi = \psi \circ \psi'$ is a bijection $H \rightarrow H'$ it is therefore sufficient to prove that ψ and ψ' preserve H-equivalence.

Hence, suppose that $c \equiv_H d$. Then $c \equiv_R d$ implies $sc \equiv_R sd$. Since all elements in the range of ψ are L-equivalent, we also have $sc \equiv_L sd$, as desired. In the same way it follows that $tc \equiv_H td$. \square

Lemma 2.15. $a \equiv_H b \equiv_J ab$ implies $a \equiv_H ab$.

Proof. Since $a \equiv_J ab$ and $b \equiv_J ab$, it follows by Lemma 2.9 that $ab \equiv_R a$ and $ab \equiv_L b$. As $a \equiv_H b$, the latter implies $a \equiv_L b \equiv_L ab$. Thus, $ab \equiv_H a$. \square

Proposition 2.16. Let H be an H-class of a cyclically finite semigroup \mathfrak{S} . The following statements are equivalent.

- (1) H contains an idempotent element.
- (2) There are $a, b \in H$ with $ab \equiv_J a$.
- (3) There are $a, b \in H$ with $ab \in H$.
- (4) H is closed under multiplication.
- (5) H is closed under multiplication and the subsemigroup $\langle H, \cdot \rangle$ is a group.

Proof. The implications (5) \Rightarrow (4) and (1) \Rightarrow (2) are trivial, while (2) \Rightarrow (3) follows immediately by Lemma 2.15.

(4) \Rightarrow (1) Given $a \in H$, we can use Lemma 2.3 to find a number $0 < n < \omega$ such that a^n is idempotent. Furthermore, we have $a^n \in H$ since H is closed under multiplication.

(3) \Rightarrow (4) Fix $a, b \in H$ with $ab \in H$ and consider two arbitrary elements $c, d \in H$. Since $a \leq_R c$ and $b \leq_L d$, there are $s, t \in S^1$ with $a = sc$ and $b = dt$. It follows that

$$cd \geq_J scdt = ab \in H,$$

which implies that $cd \equiv_J c$. Hence, we have $cd \in H$ by Lemma 2.15.

(4) \Rightarrow (5) Suppose that H is closed under multiplication. Fix $c \in H$. By Lemma 2.3, there exists some number $0 < n < \omega$ such that $e := c^n$ is idempotent. By assumption, it follows that $e \in H$. We start by showing that e acts as a neutral element. Hence, consider $a \in H$. Then $a \equiv_L e$ implies that $a = se$, for some $s \in S^1$. Therefore,

$$ae = see = se = a.$$

In the same way $a \equiv_R e$ implies that $ea = a$.

To conclude the proof, it remains to show that each element $a \in H$ has an inverse. Given a we use Lemma 2.3 to find some number $k > 0$ such that $b := a^k$ is idempotent. By Corollary 2.11 (b),

$$be = b \equiv_J e \quad \text{and} \quad bb = b \quad \text{implies} \quad eb = e.$$

Consequently, $a^k = b = eb = e$ and a^{k-1} is the desired inverse of a . \square

3 Simon's Lemma


Since the semigroup operation \cdot is associative, we can evaluate a product $a_0 \cdots a_{n-1}$ in many different ways, depending on where we put the parenthesis. For instance, we can do the evaluation left-to-right or right-to-left:

$$(\cdots((a_0 a_1) a_2) \cdots a_{n-1}), \quad (a_0(a_1(\cdots(a_{n-2} a_{n-1}) \cdots))).$$

If we want to do as much of the computation as possible in parallel, we can instead use the following scheme:

$$[((a_0 a_1)(a_2 a_3))((a_4 a_5)(a_6 a_7))]\cdots$$

Each of these possible ways of putting the parentheses can be visualised as a tree.




We call such trees *factorisation trees*.

Definition 3.1. Let \mathfrak{S} be a semigroup and $w \in S^+$ a sequence of elements. A *factorisation tree* of w is an S -labelled, successor-ordered tree where

- the leaves are labelled by the elements in w (in left-to-right order) and
- each internal vertex is labelled by the product of the labels of its successors.

Note that we do not require a factorisation tree to be binary. A typical application of factorisation trees is the following problem. Suppose we are given a sequence $w = \langle a_0, \dots, a_{n-1} \rangle$ of semigroup elements and want to compute the product $a_i \cdots a_j$ of some factor. If we know a factorisation tree of w , we can do so by traversing the subtree corresponding to the subsequence $\langle a_i, \dots, a_j \rangle$. For instance, given the tree



we can evaluate the product $a_3 \cdots a_{11}$ by multiplying the labels of the marked vertices. The amount of work we have to do for this depends both on the height of the tree and on its branching factor. Thus, we would ideally like to minimise both. If we prioritise the branching and take a binary tree, the

height is at least logarithmic. If, on the other hand, we minimise the height and take a tree of height one, the branching is the same as the length of our sequence and the tree becomes useless. It turns out that there is a middle ground where the height is still bounded by a constant, while all vertices with more than two successors have a labelling that makes it easy to compute the product: we can ensure that every vertex of this kind corresponds to a product of the form $e \cdots e$ where e is idempotent. This leads to the following definition.

Definition 3.2. A *Simon tree* of a sequence w is a factorisation tree where

- no vertex has exactly one successor, and
- for each vertex v with more than two successors u_0, \dots, u_{n-1} , there exists some idempotent $e \in S$ such that u_0, \dots, u_{n-1} (and thus also v) are all labelled by e .

We call an internal vertex v of a Simon tree *binary* if it has two successors and *idempotent* if it has more than two.

For technical reasons, we will not work with factorisation trees directly but a different encoding which behaves better with respect to compositions. Since we will use these results also in Chapter VI below, we will present the definition and the proofs for arbitrary linear orders instead of just finite ones.

Definition 3.3. Let \mathfrak{A} be a linear order and \mathfrak{S} a finite semigroup (written additively).

(a) An *additive labelling* of \mathfrak{A} is a function λ mapping each pair $i < j$ of elements of A to some element $\lambda(i, j) \in S$ such that

$$\lambda(i, k) = \lambda(i, j) + \lambda(j, k), \quad \text{for all } i < j < k.$$

(b) A *split* of \mathfrak{A} is a function $\sigma : A \rightarrow [n]$ mapping each element $a \in A$ to some number $\sigma(x) < n$. We call n the *height* of σ .


(c) A split $\sigma : A \rightarrow [n]$ is *right-guarded* if $\sigma^{-1}(n-1)$ either contains the maximal element of A , or it is unbounded from above. Similarly, σ is

left-guarded if $\sigma^{-1}(n-1)$ either contains the minimal element of A , or it is unbounded from below.

(d) Given a split $\sigma : A \rightarrow [n]$ of \mathfrak{A} , we define a binary relation \sqsubset_σ on A by

$$x \sqsubset_\sigma y \quad \text{iff} \quad x < y, \quad \sigma(x) = \sigma(y), \quad \text{and} \\ \sigma(z) \leq \sigma(x), \quad \text{for all } x \leq z \leq y.$$


As usual, \sqsubseteq_σ denotes the reflexive version of \sqsubset_σ .



(e) A split $\sigma : A \rightarrow [n]$ of \mathfrak{A} is *Ramseyan* for an additive labelling λ if

$$x \sqsubset_\sigma y \quad \text{and} \quad x \sqsubseteq_\sigma u \sqsubset_\sigma v \quad \text{implies} \quad \lambda(x, y) = \lambda(u, v).$$

Clearly, splits are just another way to encode factorisation trees. There is also a correspondence between Simon trees and Ramseyan splits, although it is less direct.



Lemma 3.4. Let \mathfrak{S} be a semigroup, $w = \langle a_0, \dots, a_{n-1} \rangle \in S^+$ a sequence of elements, and let λ be the additive labelling of $[n+1]$ defined by

$$\lambda(i, k) := a_i \cdots a_k, \quad \text{for } 0 \leq i < k \leq n.$$

(a) If there exists a Ramseyan split of λ of height k , the word w has a Simon tree of height at most $3k-2$.

(b) If w has a Simon tree of height k , there exists a Ramseyan split of λ of height at most k .

Proof. (a) We prove the claim by induction on k . If $k=1$, then we have

$$\lambda(x, y) = \lambda(u, v), \quad \text{for all cuts } x < y \text{ and } u < v.$$


This means that either $n \leq 2$ and w has a Simon tree of height 0 or 1, or $n > 2$ and there is some element e such that

$$a_0 = \cdots = a_{n-1} = e = a_0 a_1.$$

Hence, e is idempotent and w has a Simon tree of height 1 where the root is an idempotent vertex.

For the inductive step, suppose that $k > 1$. Let $z_1 < \cdots < z_m$ be an enumeration of $\sigma^{-1}(k-1)$ and set $u_i := w[z_i, z_{i+1})$, for $0 \leq i \leq m$ (where we have set $z_0 := 0$ and $z_{m+1} := n$). Then σ induces a Ramseyan split of height $k-1$ for each factor u_i . Hence, we can use the inductive hypothesis to get a Simon tree t_i of height $3(k-1)-2$ for each u_i . We define the desired tree for w as follows. If $m=1$, we use a binary vertex to combine t_0 and t_1 . Similarly, if $m=2$, we use two binary vertices to combine first t_0 and t_1 and then the resulting tree with t_2 . Finally, for $m > 2$, we use an idempotent vertex to combine t_1, \dots, t_{m-1} , and then two binary vertices to combine the resulting tree, first with t_0 and then with t_m . Note that the products $c_1 := \lambda(u_1), \dots, c_{m-1} := \lambda(u_{m-1})$ are indeed equal and idempotent since

$$z_i \sqsubset_\sigma z_{i+1} \sqsubset_\sigma z_{i+2} \quad \text{implies} \quad c_i = c_{i+1} \quad \text{and} \quad c_i = c_i c_{i+1}.$$



(b) Let t be a Simon tree of w of height k and let p_0, \dots, p_{n-1} be its leaves in left-to-right order. We set

$$\sigma(x) := \begin{cases} 0 & \text{if } x = 0, \\ k-1 - |p_{x-1} \sqcap p_x| & \text{if } 0 < x < n, \\ 0 & \text{if } x = n. \end{cases}$$

where as usual we have identified the vertices of t with words over some set D of directions, and $u \sqcap v$ denotes the longest common prefix of u and v . We claim that σ is Ramseyan. Hence, suppose that $x \sqsubset_\sigma y$ and $x \sqsubseteq_\sigma u \sqsubset_\sigma v$.

We first consider the case where all cuts x, y, u, v are proper. By definition of σ , there exists a vertex z with $|z| = k-1 - \sigma(x)$ such that

$$z = p_{x-1} \sqcap p_x = p_{y-1} \sqcap p_y = p_{u-1} \sqcap p_u = p_{v-1} \sqcap p_v.$$

One of the subtrees attached to z corresponds to the factor $w[x, y)$ and one to $w[u, v)$. Let us denote these by, respectively, r and s . The roots of r and s are labelled by, respectively, $\lambda(x, y)$ and $\lambda(u, v)$. As z has at least one successor to the left of r (the one above p_{x-1}), and one successor to the right of r and s (the one containing p_y or p_v), it follows that z has more than 2 successors. Consequently, z is an idempotent vertex of t . Since $\lambda(x, y)$ and $\lambda(u, v)$ are the labels of successors of z , it follows in particular that their values coincide.

It remains to consider the cases where $x = 0$, $y = n$, or $v = n$. Suppose that $x = 0$. (The other two cases follow analogously.) Then $\sigma(x) = 0$ and it

follows by definition of σ that there is some vertex z with $|z| = k-1$ such that p_x, p_y, p_u , and p_v are all successors of z . If z is binary, we have $x = u$ and $y = v$ and there is nothing to do. Otherwise, z is idempotent with label e and $\lambda(x, y) = e = \lambda(u, v)$. \square

After these preparations let us finally prove that Ramseyan splits of bounded height exist. We start with two special cases.

Lemma 3.5. *Let λ be an additive labelling of a linear order \mathfrak{A} by a finite semi-group \mathfrak{S} and let H be an H -class of \mathfrak{S} containing an idempotent. If*

$$\lambda(a, b) \in H, \quad \text{for all } a < b,$$

then λ has a right-guarded Ramseyan split σ of height at most $|H|$.

Proof. By Proposition 2.16, H forms a group. Hence, we can define

$$\lambda_*(a, b) := \begin{cases} \lambda(a, b) & \text{if } a < b, \\ 0 & \text{if } a = b, \\ -\lambda(b, a) & \text{if } a > b. \end{cases}$$

It follows that

$$\lambda_*(a, b) + \lambda_*(b, c) = \lambda_*(a, c), \quad \text{for all } a, b, c \in A.$$

We fix an element $a_0 \in A$ and we choose some $c \in H$ such that the set

$$\{ b \in A \mid \lambda_*(a_0, b) = c \}$$

either contains the maximal element of A , or it is unbounded. Let $\mu : H \rightarrow [n]$ be some bijection with $\mu(c) = n-1$. We claim that the split

$$\sigma(a) := \mu(\lambda_*(a_0, a)), \quad \text{for } a \in A,$$

is Ramseyan and right-bounded. Right-boundedness follows by choice of μ . To see that it is Ramseyan, consider elements $x \sqsubset_\sigma y$ and $x \sqsubseteq_\sigma u \sqsubset_\sigma v$. Then

$\sigma(x) = \sigma(y)$ implies that

$$\begin{aligned}\lambda(x, y) &= \lambda_*(x, y) = \lambda_*(x, a_o) + \lambda_*(a_o, y) \\ &= \lambda_*(a_o, y) - \lambda_*(a_o, x) = o.\end{aligned}$$

Similarly, it follows that $\lambda(u, v) = o$. In particular, $\lambda(x, y) = \lambda(u, v)$. \square

Lemma 3.6. *Let λ be an additive labelling of a linear order \mathfrak{A} by a finite semi-group \mathfrak{S} and let J be a regular J -class of \mathfrak{S} . If*

$$\lambda(a, b) \in J, \quad \text{for all } a < b,$$

then λ has a right-guarded Ramseyan split σ of height at most $|J|$.

Proof. If $|A| \leq 1$, the claim is trivial. Hence, we may assume that A has at least two elements. For every non-maximal $a \in A$, we fix some element $b > a$ and we set

$$R(a) := [\lambda(a, b)]_{\mathbb{R}}.$$

Note that $R(a)$ does not depend on b since, given $a < b < b'$, it follows by Corollary 2.10 that

$$\lambda(a, b') = \lambda(a, b) + \lambda(b, b') \equiv_J \lambda(a, b) \equiv_J \lambda(b, b')$$

implies $\lambda(a, b') \equiv_{\mathbb{R}} \lambda(a, b)$. Similarly, for every non-minimal $a \in A$, we set

$$L(a) := [\lambda(b, a)]_{\mathbb{L}}, \quad \text{for } b < a.$$

If \mathfrak{A} has a maximal element a , we set

$$R(a) := [e]_{\mathbb{R}}, \quad \text{for some idempotent } e \in L(a).$$

(Such an idempotent exists by Proposition 2.13.) Similarly, if \mathfrak{A} has a minimal element a , we set

$$L(a) := [e]_{\mathbb{L}}, \quad \text{for some idempotent } e \in R(a).$$

Let

$$H(a) := R(a) \cap L(a).$$

We start by proving that $H(a)$ forms a group, for every $a \in A$. If a is maximal or minimal, this follows from Proposition 2.16 and the fact that, by definition, $H(a)$ contains an idempotent. Otherwise, fix elements $b < a < b'$. Then

$$\lambda(b, a) \in L(a) \quad \text{and} \quad \lambda(a, b') \in R(a).$$

Then $\lambda(b, a) + \lambda(a, b') = \lambda(b, b') \in J$ and we can use Corollary 2.10 to find an idempotent e such that

$$\lambda(b, a) \equiv_{\mathbb{L}} e \equiv_{\mathbb{R}} \lambda(a, b').$$

Consequently, $e \in L(a) \cap R(a) = H(a)$ and it follows by Proposition 2.16 that $H(a)$ forms a group.

Let H_0, \dots, H_{k-1} be an enumeration of all H -classes in J that form groups and set

$$B_i := \{a \in A \mid H(a) = H_i\}, \quad \text{for } i < k.$$

We choose the enumeration H_0, \dots, H_{k-1} such that B_i either contains the maximal element of A , or such that it is unbounded in A . By Corollary 2.10, it follows that

$$\lambda(a, b) \in H_i, \quad \text{for all } a < b \text{ in } B_i.$$

Hence, we can use Lemma 3.5 to construct a right-guarded Ramseyan split τ_i of B_i of height at most $|H_i|$. Note that, according to Lemma 2.14, all classes H_i have the same size n . We set

$$\sigma(a) := ni + \tau_i(a), \quad \text{for } a \in B_i.$$

Then σ is a split of height $nk \leq |J|$. Furthermore, it is right-guarded since B_i is unbounded in A . To see that σ is also Ramseyan, consider elements $x \sqsubset_{\sigma} y$ and $x \sqsubset_{\sigma} u \sqsubset_{\sigma} v$. Then $x, y, u, v \in B_i$, for some i . Since τ_i is Ramseyan for λ , it follows that $\lambda(x, y) = \lambda(u, v)$. \square

Theorem 3.7. Let λ be an additive labelling of a linear order \mathfrak{A} by a finite semigroup \mathfrak{S} . Then λ has a Ramseyan split of height at most $2 \cdot |S|$.

Proof. Let J be a maximal J -class of \mathfrak{S} . We call a set $C \subseteq A$ a J -set if

$$\lambda(i, j) \in J, \quad \text{for all } i < j \text{ in } C.$$

We use two different orders on J -sets: inclusion \subseteq and the order

$$C < D \quad : \text{iff} \quad \text{there is some } c \in C \text{ with } c < d, \text{ for all } d \in D.$$

We call a J -set *maximal* if it is maximal with respect to \subseteq . Let us start by establishing the following claims.

- (I) The convex hull of a J -set is a J -set.
 - (II) The restriction of $<$ to the maximal J -sets forms a strict linear order.
 - (III) Every J -set is contained in a maximal one.
 - (IV) Two different maximal J -sets have at most one element in common.
 - (V) Every element $a \in A$ is contained in at least one, but at most two maximal J -sets.
- (i) Let D be the convex hull of a J -set C . To show that it is a J -set, consider elements $a < b$ of D . Then there are elements $a', b' \in C$ with $a' \leq a < b \leq b'$. By additivity, it follows that

$$\lambda(a', b') = \lambda(a', a) + \lambda(a, b) + \lambda(b, b').$$

(If $a' = a$ or $b = b'$, we omit the corresponding terms.) Hence,

$$\lambda(a', b') \leq_J \lambda(a, b).$$

By maximality of J , it follows that $\lambda(a', b') \in J$ implies $\lambda(a, b) \in J$.

(ii) Irreflexivity and transitivity are immediate. For linearity, fix two maximal J -sets C and D . Then there are elements $a \in C \setminus D$ and $b \in D \setminus C$. If $a < b$ it follows by convexity that $C < D$. Otherwise, $D < C$.

(iii) By Zorn's Lemma, it is sufficient to prove that the class of J -sets is closed under unions of increasing sequences. Hence, let $(C_i)_{i < \alpha}$ be such a

sequence and let $D := \bigcup_{i < \alpha} C_i$ be its union. Fix $a < b$ in D . Then there is some index $i < \alpha$ with $a, b \in C_i$. As C_i is a J -set, it follows that $\lambda(a, b) \in J$.

(iv) Let C and D be maximal J -sets with $|C \cap D| \geq 2$. By (ii) and symmetry, we may assume that $C < D$. We will prove that $C \cup D$ is also a J -set. By maximality, this implies that $C = C \cup D = D$. Hence, fix elements $a < b$ in $C \cap D$ and consider two elements $a' < b'$ in $C \cup D$. If $a', b' \in C$ or $a', b' \in D$, we have $\lambda(a', b') \in J$ by assumption. Otherwise, $C < D$ implies that $a' \in C \setminus D$ and $b' \in D \setminus C$, and that $a' < a < b < b'$. Since

$$\lambda(a', a) + \lambda(a, b) \equiv_J \lambda(a, b) \equiv_J \lambda(a, b) + \lambda(b, b'),$$

it follows by Corollary 2.11 (a) that

$$\lambda(a', b') = \lambda(a', a) + \lambda(a, b) + \lambda(b, b') \in J.$$

(v) Let $a \in A$. The singleton $\{a\}$ is a J -set since it does not contain two elements $b < b'$. By (iii), it is therefore contained in some maximal J -set C . To conclude the proof, suppose for a contradiction that a is contained in at least three different maximal J -sets C_0, C_1, C_2 . By maximality, none of these sets can be a singleton. It follows that either at least two of them contain an element less than a , or at least two contain an element greater than a . By symmetry, suppose that there are $b_0, b_1 < a$ with $b_0 \in C_0$ and $b_1 \in C_1$. Furthermore, we may assume by symmetry that $b_0 \leq b_1$. As C_0 is convex, this implies that $b_1 \in C_0$. Hence, $a, b_1 \in C_0 \cap C_1$, which implies by (iv) that $C_0 = C_1$. A contradiction.

Having established the above claims we now construct the desired split σ by induction on $|S|$. Fix a maximal set $B \subseteq A$ that contains at most 1 element of each maximal J -set. (We can use Zorn's Lemma to prove that such a set exists.) For $a < b$ in B it follows that $\lambda(a, b) \in S \setminus J$. As J is maximal, the complement $S \setminus J$ forms a subsemigroup of \mathfrak{S} . Therefore, we can use the inductive hypothesis to construct a Ramseyan split σ' of B of height at most $2(|S| - |J|)$.

To turn σ' into a split of all of A , we consider a maximal convex subset $C \subseteq A$ with $C \cap B = \emptyset$. It is sufficient to construct a Ramseyan split τ_C of C

of height at most $2|J|$. Then we obtain the desired split σ of A by

$$\sigma(a) := \begin{cases} \sigma'(a) + 2|J| & \text{if } a \in B, \\ \tau_C(a) & \text{if } a \in C \text{ and } C \text{ as above.} \end{cases}$$



To see that σ is Ramseyan, consider elements $x \sqsubset_{\sigma} y$ and $x \sqsubset_{\sigma} u \sqsubset_{\sigma} v$. If $\sigma(x) < 2|J|$, it follows that $x, y, u, v \in C$ for some set C as above and that $x \sqsubset_{\tau_C} y$ and $x \sqsubset_{\tau_C} u \sqsubset_{\tau_C} v$. Since τ_C is Ramseyan, this implies that $\lambda(x, y) = \lambda(u, v)$. Hence, suppose that $\sigma(x) \geq 2|J|$. Then $x, y, u, v \in B$ and $x \sqsubset_{\sigma'} y$ and $x \sqsubset_{\sigma'} u \sqsubset_{\sigma'} v$. Since σ' is Ramseyan, this implies that $\lambda(x, y) = \lambda(u, v)$.

Hence, it remains to construct the splits τ_C . Let B_- be the set of all elements of B that are smaller than those of C and let B_+ be those that are larger. By (i) and (v), there is at most one maximal J -set that intersects both B_- and Z . Similarly, there is at most one maximal J -set intersecting Z and B_+ . We denote these two sets by D_- and D_+ (setting $D_- := \emptyset$ or $D_+ := \emptyset$ if such sets do not exist). We claim that $C \subseteq D_- \cup D_+$. For a contradiction, suppose that there is some element $c \in C \setminus (D_- \cup D_+)$. By (v), the element c is contained in some maximal J -set E . If E intersects B_- , (v) would imply have $E = C_-$. Similarly, if E intersects B_+ , we would have $E = C_+$. Hence, $E \subseteq C$. It follows that

$$|(B \cup \{c\}) \cap X| \leq 1, \quad \text{for every } J\text{-set } X.$$

A contradiction to the maximality of B .

If $|C| \leq 2$, we can take any injective function $\tau_C : C \rightarrow [2]$ as the desired Ramseyan split. Its height is $2 \leq 2|J|$. Hence, suppose that C has at least 3 elements. By the above claim, it follows that one of $C \cap D_-$ or $C \cap D_+$ has at least 2 elements. By symmetry, we may assume that it is the former. Since $D_- \cap B \neq \emptyset$, the set D_- also has at least 3 elements. Fix elements $a < b < c$ in D_- . Then

$$\lambda(a, b) \in J, \quad \lambda(b, c) \in J, \quad \text{and} \quad \lambda(a, b) + \lambda(b, c) = \lambda(a, c) \in J.$$

Consequently, it follows by Proposition 2.13 that the J -class J is regular. Hence, we can use Lemma 3.6 to construct a right-guarded Ramseyan split τ_- of D_- of height at most $|J|$. In the same way we can use (the dual version of) Lemma 3.6 to construct a left-guarded Ramseyan split τ_+ of $D_+ \setminus D_-$ of height at most $|J|$. We set

$$\tau_C(a) := \begin{cases} \tau_-(a) & \text{if } a \in C \cap D_-, \\ \tau_+(a) + 1 & \text{if } a \in C \cap D_+ \setminus D_-. \end{cases}$$

To see that τ_C is Ramseyan, consider elements $x \sqsubset_{\tau_C} y$ and $x \sqsubset_{\tau_C} u \sqsubset_{\tau_C} v$. If $x, y, u, v \in C \cap D_-$ or $x, y, u, v \in C \cap D_+$, we have $\lambda(x, y) = \lambda(u, v)$ since τ_- and τ_+ are Ramseyan. Hence, suppose otherwise. Then $x \in D_-$ while $w := \max\{y, v\} \in D_+ \setminus D_-$. Since τ_+ is left-guarded, there is some element $z \in D_+ \setminus D_-$ with $x < w < z$ and $\tau_+(w) = |J| - 1$. Hence,

$$\tau_C(w) = \tau_C(x) = \tau_-(a) < |J| = \tau_+(z) + 1 = \tau_C(z),$$

which implies that $x \not\sqsubset_{\tau_C} w$. A contradiction. \square

The following rephrasing of this result is frequently more convenient in applications.

Corollary 3.8. *Let $\varphi : \mathfrak{S} \rightarrow \mathfrak{T}$ be a semigroup homomorphism where \mathfrak{S} is finitely generated and \mathfrak{T} is finite. If $f : \omega \rightarrow \omega$ and $\mu : S \rightarrow \omega$ are functions such that*

$$\mu(a_0 \cdots a_{n-1}) \leq f\left(\max_{i < n} \mu(a_i)\right),$$

holds for all elements $a_0, \dots, a_{n-1} \in S$ satisfying

- ♦ $n = 2$, or
- ♦ $\varphi(a_1) = \cdots = \varphi(a_n) = e$ for some idempotent $e \in T$,

then $\text{rng } \mu$ is finite.

Proof. Let G be a set of generators of S . Given $w = \langle a_0, \dots, a_{n-1} \rangle \in G^+$, let t be a Simon tree of the sequence $\langle \varphi(a_0), \dots, \varphi(a_{n-1}) \rangle$ of height at most

$3|S|$. By induction on the height k of t , it follows that

$$\mu(\pi(w)) \leq f^k\left(\max_{g \in G} \mu(g)\right),$$

which is independent of w . As k is bounded by $3|S|$ the claim follows. \square

Exercise 3.1. Let $\varphi : \Sigma^+ \rightarrow \mathfrak{S}$ be a semigroup homomorphism where Σ and \mathfrak{S} are finite. Prove that there exists a constant $k > 0$ such that every word $w \in \Sigma^+$ of length n^k with $n > 2$ has a factorisation

$$w = w_0 \dots w_{n+1} \quad \text{with} \quad \varphi(w_1) = \dots = \varphi(w_n).$$

4 Regular Languages of Finite Words

Before considering infinite words, we start with finite ones. For these it is quite simple to characterise which languages are MSO-definable. We present several equivalent ways to describe such languages. The first one is in terms of automata.

Definition 4.1. (a) A *nondeterministic automaton* $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$ consists of a finite set Q of *states*, a finite *input alphabet* Σ , an *initial state* $q_0 \in Q$, a set $F \subseteq Q$ of *final states*, and a *transition relation* $\Delta \subseteq Q \times \Sigma \times Q$. Instead of $\langle p, a, q \rangle \in \Delta$, we also write $p \xrightarrow{a} q$.

(b) A *run* of an automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$ on an input word $w = \langle a_0, \dots, a_{n-1} \rangle \in \Sigma^*$ is a sequence p_0, \dots, p_n of states starting with $p_0 = q_0$ such that

$$\langle p_i, a_i, p_{i+1} \rangle \in \Delta, \quad \text{for all } i < n.$$

A run p_0, \dots, p_n is *accepting* if $p_n \in F$.


(c) An automaton \mathcal{A} *accepts* a word w if there exists an accepting run of \mathcal{A} on w . The *language recognised* by \mathcal{A} is the set $L(\mathcal{A})$ of all words it accepts. \square

Example. The language of all words over the alphabet $\Sigma := \{a, b, c\}$ with an even number of letters a is recognised by the automaton

$$\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$$

with states $Q = \{q_0, q_1\}$, final state $F = \{q_0\}$, and transitions

$$q_0 \xrightarrow{a} q_1, \quad q_1 \xrightarrow{a} q_0, \quad q_0 \xrightarrow{b,c} q_0, \quad q_1 \xrightarrow{b,c} q_1.$$



The second way to describe languages is in terms of semigroups.

Definition 4.2. Let $L \subseteq \Sigma^*$.

(a) A homomorphism $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ into a semigroup \mathfrak{S} *recognises* L if

$$L \setminus \{\epsilon\} = \eta^{-1}[P], \quad \text{for some } P \subseteq \mathfrak{S}.$$

(b) The *syntactic congruence* of L is the relation on Σ^* defined by

$$x \sim_L y \quad \text{iff} \quad uxw \in L \Leftrightarrow uyw \in L \quad \text{for all } u, w \in \Sigma^*.$$

Note that in the definition of recognition by a homomorphism we have ignored the empty word. This can be avoided by using monoids instead of semigroups. But later on when we study infinite words, monoids would cause technical problems of their own.

Example. Let L be the language of all words over the alphabet $\Sigma := \{a, b, c\}$ with an even number of letters a . L is recognised by the homomorphism $\eta : \Sigma^+ \rightarrow \mathbb{Z}/2\mathbb{Z}$ that maps a to 1 and b, c to 0. Its syntactic congruence is

$$\begin{aligned} x \sim_L y & \quad \text{iff} \quad \text{modulo 2, } x \text{ and } y \text{ have the same number of} \\ & \quad \text{letters } a \\ & \quad \text{iff} \quad \eta(x) = \eta(y). \end{aligned}$$

Example. Let us show that every language recognised by an automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_o, F \rangle$ can also be recognised by some finite semigroup. This semigroup is $\mathfrak{S} := \langle \mathcal{P}(Q \times Q), \circ \rangle$ consisting of all relations $R \subseteq Q \times Q$ with the usual composition

$$R \circ T := \{ \langle a, c \rangle \mid \langle a, b \rangle \in R, \langle b, c \rangle \in T \}$$

as multiplication. To define the homomorphism $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ we only need to say what it does on single letters. For $a \in \Sigma$, we set

$$\eta(\langle a \rangle) := \{ \langle p, q \rangle \mid \langle p, a, q \rangle \in \Delta \}.$$

For words $w = \langle a_o, \dots, a_{n-1} \rangle$ it then follows that

$$\eta(w) = \eta(\langle a_o \rangle) \circ \dots \circ \eta(\langle a_{n-1} \rangle).$$

To see that η recognises the desired language, note that $L(\mathcal{A}) = \eta^{-1}[P]$ where

$$P := \{ R \subseteq Q \times Q \mid \langle q_o, p \rangle \in R \text{ for some } p \in F \}.$$

The subsemigroup of \mathfrak{S} induced by the elements in the range of η is also called the *transition semigroup* of \mathcal{A} . J

Example. The construction in the previous example simplifies for deterministic automata. Such automata take the form $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, F \rangle$ where, instead of a transition relation $\Delta \subseteq Q \times \Sigma \times Q$, we now have a transition function $\delta : Q \times \Sigma \rightarrow Q$. In this case, we can take the semigroup $\mathfrak{S} := \langle Q^Q, \circ^{\text{op}} \rangle$ of all functions $Q \rightarrow Q$ with ‘mirrored’ function composition as multiplication:

$$f \circ^{\text{op}} g := g \circ f.$$

The homomorphism $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ maps a letter $a \in \Sigma$ to the function

$$\delta_a(q) := \delta(q, a).$$

Then $L(\mathcal{A}) = \eta^{-1}[P]$ for $P := \{ f : Q \rightarrow Q \mid f(q_o) \in F \}$. J

Lemma 4.3. *The syntactic congruence \sim_L is a congruence relation (of the free semigroups (Σ^+, \cdot)).*

Proof. \sim_L is obviously reflexive and symmetric. For transitivity, suppose that $x \sim_L y \sim_L z$. To show that $x \sim_L z$, let $u, w \in \Sigma^*$. Then

$$uxw \in L \Leftrightarrow uyw \in L \Leftrightarrow uz w \in L.$$

Finally, suppose that $x \sim_L x'$ and $y \sim_L y'$. To show that $xy \sim_L x'y'$, consider words $u, w \in \Sigma^*$. Then

$$uxyw \in L \Leftrightarrow ux'yw \in L \Leftrightarrow ux'y'w \in L. \quad \square$$

We obtain the following characterisations of the class of MSO-definable languages.

Theorem 4.4. *Let $L \subseteq \Sigma^*$. The following statements are equivalent:*

- (1) *L is definable in MSO.*
- (2) *L is recognised by a homomorphism to a finite semigroup.*
- (3) *L is recognised by an automaton.*
- (4) *The syntactic congruence of L has finite index.*

Proof. (1) \Rightarrow (4) Suppose that there is an MSO-formula φ defining L , and set $m := \text{qr}(\varphi)$. Since

$$v \equiv_{\text{MSO}}^m v' \quad \text{and} \quad w \equiv_{\text{MSO}}^m w' \quad \text{implies} \quad v \hat{\sim} w \equiv_{\text{MSO}}^m v' \hat{\sim} w',$$

for $v, v', w, w' \in \Sigma^+$, the relation \equiv_{MSO}^m is a congruence relation on Σ^+ . Furthermore, if $x \equiv_{\text{MSO}}^m y$ then $uxw \equiv_{\text{MSO}}^m uyw$ implies that

$$uxw \in L \quad \text{iff} \quad uxw \models \varphi \quad \text{iff} \quad uyw \models \varphi \quad \text{iff} \quad uyw \in L.$$

Hence, $\equiv_{\text{MSO}}^m \subseteq \sim_L$. We have seen in Proposition I.3.4 that there are only finitely many \equiv_{MSO}^m -classes. As every \equiv_{MSO}^m -class is contained in a \sim_L -class, it follows that \sim_L also has only finitely many classes.

(4) \Rightarrow (2) If \sim_L has only finitely many classes, the quotient Σ^+/\sim_L is a finite semigroup. We claim that $L \setminus \{\langle \rangle\} = \pi^{-1}(P)$ where $\pi : \Sigma^+ \rightarrow \Sigma^+/\sim_L$ is the canonical projection mapping every word w to its \sim_L -class $[w]$ and

$$P := \{ [w] \mid w \in L \}.$$

Clearly, $w \in L$ implies $\pi(w) \in P$. Conversely, if $\pi(w) \in P$, there is some $v \sim_L w$ with $v \in L$. By definition of \sim_L , it follows that $w \in L$.

(2) \Rightarrow (3) Suppose that $L \setminus \{\langle \rangle\} = \eta^{-1}(P)$, where $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ is a homomorphism to a finite semigroup \mathfrak{S} and $P \subseteq S$. We obtain an automaton \mathcal{A} recognising L by setting $\mathcal{A} := \langle S^1, \Sigma, \Delta, \mathbf{1}, F \rangle$ where

$$F := \begin{cases} P & \text{if } \langle \rangle \notin L, \\ P \cup \{\mathbf{1}\} & \text{if } \langle \rangle \in L, \end{cases}$$

and $\Delta := \{ \langle s, a, s \cdot \eta(a) \rangle \mid s \in S, a \in \Sigma \}$.

(3) \Rightarrow (1) Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$ be an automaton recognising L . We obtain a formula φ defining L as follows. φ guesses a run of \mathcal{A} on the given word. It encodes this run by a tuple $(Z_q)_{q \in Q}$ of set variables, where Z_q contains all positions such that the automaton is in state q after having read that position. In the case where $\langle \rangle \notin L$, we set

$$\varphi := \exists (Z_q)_{q \in Q} [\text{ADM} \wedge \text{INIT} \wedge \text{TRANS} \wedge \text{ACC}]$$

where ADM states that every position is labelled by at most one state:

$$\text{ADM} := \forall x \bigwedge_{p \neq q} \neg (Z_p x \wedge Z_q x),$$

INIT states that the first state is correct:

$$\text{INIT} := \exists x \left[\forall y (x \leq y) \wedge \bigvee_{\langle q_0, a, q \rangle \in \Delta} (Z_q x \wedge P_a x) \right],$$

TRANS states that at every position a valid transition is used:

$$\text{TRANS} := \forall x \forall y \left[\text{suc}(x, y) \rightarrow \bigvee_{\langle p, a, q \rangle \in \Delta} (Z_p x \wedge Z_q y \wedge P_a y) \right],$$

and ACC states that the last state is final:

$$\text{ACC} := \bigvee_{q \in F} \exists x [Z_q x \wedge \forall y (y \leq x)].$$

If $\langle \rangle \in L$, we use the formula $\forall z \varphi$ instead, where φ is defined as above. \square

Example. Let us use the preceding theorem to show that the language

$$L := \{ a^n b^n \mid n < \omega \}$$

is not MSO-definable. We give three different proofs, one for each of the above characterisations.


The easiest way is to argue in terms of the syntactic congruence. For $m < n$, we have $a^m b^m \in L$ and $a^n b^m \notin L$. Hence, $a^m \sim_L a^n$ for all $m < n$ and the index of \sim_L is infinite.

Using semigroups we can proceed as follows. For a contradiction, suppose that L were MSO-definable. Then it would be recognised by a homomorphism $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ to some finite semigroup \mathfrak{S} . It follows that there must be two numbers $m < n$ with $\eta(a^m) = \eta(a^n)$. Consequently,

$$\eta(a^m b^m) = \eta(a^m) \cdot \eta(b^m) = \eta(a^n) \cdot \eta(b^m) = \eta(a^n b^m).$$

Hence, $a^m b^m \in L$ implies $a^n b^m \in L$. A contradiction.

Finally, let us argue in terms of automata. Again, suppose that L is regular. Then there exists an automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$ recognising L . Let $n := |Q|$ be its number of states. Then $a^n b^n \in L$ implies that there exists an accepting run $(p_i)_{i \leq 2n}$ of \mathcal{A} on $a^n b^n$. As there are only n states, we can find two indices $0 \leq j < k \leq n$ with $p_j = p_k$. We construct a new input word and a corresponding accepting run on it by taking the given word and its run and repeating the part between the indices j and k . (In automata theory, this process is called *pumping*.)



This produces the word $a^{n+(k-j)}b^n$ (note that $k \leq n$, so both indices correspond to positions in the first half of the word) and the run $(p'_i)_{i \leq 2n+(k-j)}$ where

$$p'_i := \begin{cases} p_i & \text{if } i \leq k, \\ p_{i-(k-j)} & \text{if } i > k. \end{cases}$$

It follows that \mathcal{A} also accepts the input $a^{n+(k-j)}b^n \notin L$. A contradiction. \dashv

Exercise 4.1. Give direct proofs for the following implications from Theorem 4.4:

- (1) \Rightarrow (2), (2) \Rightarrow (1),
 (3) \Rightarrow (4), (4) \Rightarrow (3), (4) \Rightarrow (1). \dashv

Exercise 4.2. For each of the following languages over the alphabet $\Sigma := \{a, b, c\}$, find (i) an MSO-formula defining them, (ii) an automaton recognising them, and (iii) a homomorphism recognising them.

- (a) The language of all words of the form $a^m b^n$ with $m, n < \omega$.
- (b) The language of all words such that, after every letter a , there is some later position with a b .
- (c) The language of all words such that between any two a there is at least one b .
- (d) The language of all words with exactly 2 occurrences of the letter a .
- (e) The language of all words of the form xay with $x, y \in \Sigma^*$ and $|y| = n$. \dashv

Exercise 4.3. Prove that the following languages over $\Sigma := \{a, b\}$ are not MSO-definable.

- (a) The language of all words of the form $a^m b^n$ with $m > n$.
- (b) The language of all words of the form ww for $w \in \Sigma^*$.
- (c) The language of all words of length n^2 for $n < \omega$.
- (d) The language of all words with the same number of letters a and b .
- (e) The language of all well-bracketed words where we consider a as an opening bracket and b as a closing one. \dashv

Exercise 4.4. Let L be an MSO-definable language over the alphabet Σ . Prove that there exists a constant $0 < n < \omega$ such that every word $w \in L$ of length $|w| \geq n$ has a factorisation $w = xyz$ satisfying

$$|xy| \leq n, \quad y \neq \langle \rangle, \quad \text{and} \quad xy^kz \in L \quad \text{for all } k < \omega. \quad \dashv$$

5 First-Order Logic

We can derive a similar characterisation of the class of first-order definable word languages. The goal of this section is to prove the following theorem, which contains logical and algebraic descriptions of this class. We omit the automata-theoretic characterisation, as it is more technical.

Theorem 5.1 (Schützenberger, McNaughton, Papert, Kamp). *Let $L \subseteq \Sigma^*$. The following statements are equivalent.*

- (1) L is definable in FO.
- (2) L is definable in LTL.
- (3) $L = L(\alpha)$, for some star-free regular expression α .
- (4) L is recognised by a homomorphism into a finite aperiodic semigroup.

We have not yet defined all the notions figuring in this statement. This will be done below.

Semigroups

We start with the algebraic characterisation.

Definition 5.2. A semigroup \mathfrak{S} is *aperiodic* if every H-class of \mathfrak{S} has a single element. \square

Let us mention several alternative definitions of aperiodicity that sometimes are more convenient.

Lemma 5.3. Let \mathfrak{S} be a finite semigroup. The following statements are equivalent.

- (1) \mathfrak{S} is aperiodic.
- (2) For every $a \in S$, there is some number $n > 0$ such that $a^{n+1} = a^n$.
- (3) No subsemigroup $\mathfrak{H} \subseteq \mathfrak{S}$ with more than one element forms a group.

Proof. (2) \Rightarrow (1) Suppose that $a \equiv_H b$. Then there are elements $u, v, s, t \in S^1$ such that

$$a = sb = bu \quad \text{and} \quad b = ta = av.$$

By assumption, there is some n with $s^{n+1} = s^n$. Consequently,

$$b = av = sbv = s^n bv^n = s^{n+1} bv^n = sb = a.$$

(1) \Rightarrow (3) Suppose that $\mathfrak{H} \subseteq \mathfrak{S}$ is a subsemigroup with at least two elements that forms a group. For all $a, b \in H$, it follows that $a \leq_L b$ and $a \leq_R b$. Consequently, all elements of H are contained in the same H-class and at least one H-class of \mathfrak{S} has more than one element.

(3) \Rightarrow (2) Given $a \in S$, we consider the sequence a, a^2, a^3, a^4, \dots . Since \mathfrak{S} is finite there are numbers $n, k > 0$ such that $a^n = a^{n+k}$. We choose them minimal. If $k = 1$, we are done. Hence, suppose that $k > 1$. Let m be the number such that $n \leq m < n + k$ and $m \equiv 1$ modulo k . It follows that

$$(a^m)^i = a^{mi} = a^{m+i-1}.$$

Hence, the element a^m generates the subsemigroup $\{a^{m-1}, a^m, \dots, a^{m+k-1}\}$ which is isomorphic to $\mathbb{Z}/k\mathbb{Z}$, a group with $k > 1$ elements. \square

We can already prove the following part of Theorem 5.1.

Proposition 5.4. A language $L \subseteq \Sigma^*$ is FO-definable if, and only if, it is recognised by a homomorphism into a finite aperiodic semigroup.

Proof. (\Rightarrow) Suppose that L is defined by an FO-formula of quantifier-rank m . It follows by Proposition I.4.10 that we can define a binary operation on the set Θ_m of all FO_m -theories that turns Θ_m into a semigroup and the theory map $\text{Th}_{\text{FO}}^m : \Sigma^+ \rightarrow \Theta_m$ into a semigroup homomorphism. As this homomorphism recognises every FO_m -definable language, it is therefore sufficient to show that Θ_m is aperiodic. By Proposition I.4.11, we have

$$\text{Th}_{\text{FO}}^m(w^{n+1}) = \text{Th}_{\text{FO}}^m(w^n),$$

for every $w \in \Sigma^+$ and every $n \geq 2^m - 1$. Consequently, aperiodicity follows by Lemma 5.3.

(\Leftarrow) Let $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ be a homomorphism recognising L where \mathfrak{S} is finite and aperiodic. We will construct FO-formulae $\varphi_a(x, y)$, for $a \in S$, such that

$$w \models \varphi_a(i, k) \quad \text{iff} \quad \eta(w[i, k]) = a.$$

We proceed by induction on the J-class J of a . By inductive hypothesis, we have already constructed formulae φ_c for all $c >_J a$. First, we construct a formula ϑ_J such that

$$w \models \vartheta_J \quad \text{iff} \quad \eta(w) \in J.$$

Let us call a factor $w[i, k]$ of a word w an *J-factor* if

- ♦ $\eta(w[i, k]) >_J a$,
- ♦ either $k = |w|$ or $\eta(w[i, k+1]) \leq_J a$,
- ♦ either $i = 0$ or $\eta(w[i-1, k]) \leq_J a$.

We can define a formula $\psi_c^J(x, y)$ stating that $x[x, y-1]$ is a J-factor and $\eta(w[x, y]) = c$ by expressing that

- ♦ $\varphi_b(x, y-1)$ holds for some $b >_J a$,
- ♦ $P_d y$ holds for some d with $b \cdot \eta(d) = c \leq_J a$, and

- ♦ x is either the least element, or we have $P_{d'}(x - 1)$ for some d' with $\eta(d') \cdot b \leq_J a$.

Let ϑ_J be the formula saying that

- ♦ every J -factor u of w satisfies $\eta(u) \geq_J a$,
- ♦ if u and v are consecutive J -factors of w , then $\eta(uv) \geq_J a$.

Then it follows by Corollary 2.11 (b) that $w \models \vartheta_J$ implies $\eta(w) \geq_J a$, as desired.

To conclude the proof note that, \mathfrak{S} being aperiodic, we have

$$\eta(w) = a \quad \text{iff} \quad \eta(w) \equiv_H a \quad \text{iff} \quad \eta(w) \equiv_L a \text{ and } \eta(w) \equiv_R a.$$

Furthermore,

$$\eta(w) \equiv_L \eta(v) \quad \text{and} \quad \eta(w) \equiv_R \eta(u),$$

where u is the first J -factor of w and v the last one.

Consequently, we obtain the desired formula $\varphi_a(x, y)$ by stating the following three conditions:

- ♦ $\eta(w) \in J$.
- ♦ $\eta(u) \equiv_R a$, where u is the first J -factor of w ,
- ♦ $\eta(v) \equiv_L a$, where v is the first J -factor of w .

By the above remarks, each of them can be expressed in first-order logic. \square

Remark. This result can be used to decide whether a given regular language L is first-order definable. Given an automaton for L , we start by computing a semigroup recognising it using the construction from the proof above. Unfortunately, simply checking this semigroup for aperiodicity is not enough since we need to know whether *some* semigroup recognising L is aperiodic. One can show that among all semigroups recognising a given language L there is always a minimal one, the so-called *syntactic semigroup* of L . This semigroup can be computed from any other semigroup recognising L by taking a suitable quotient. As aperiodicity is preserved under quotients it follows that, if any semigroup recognising L is aperiodic, so is its syntactic semigroup. Hence, from the semigroup we computed above we can construct the syntactic semigroup and check it for aperiodicity. \square

Star-Free Expressions

We can also characterise the first-order definable languages via a certain kind of regular expressions.

Definition 5.5. (a) A *star-free regular expression* α over an alphabet Σ is a term built up from binary operations \cdot , \cap , \cup , a unary operation \sim , and constant symbols \emptyset and a , for each letter $a \in \Sigma$.

(b) The language $L(\alpha) \subseteq \Sigma^*$ of such an expression α is defined inductively as follows.

$$L(\emptyset) := \emptyset,$$

$$L(a) := \{a\}, \quad \text{for } a \in \Sigma,$$

$$L(\alpha \cap \beta) := L(\alpha) \cap L(\beta),$$

$$L(\alpha \cup \beta) := L(\alpha) \cup L(\beta),$$

$$L(\sim \alpha) := \Sigma^* \setminus L(\alpha),$$

$$L(\alpha \cdot \beta) := L(\alpha) \cdot L(\beta).$$

Examples. (a) $\sim \emptyset \cdot a \cdot \sim \emptyset \cdot a \cdot \sim \emptyset$ describes the language of all words containing at least two occurrences of the letter a .

(b) $\sim(\sim \emptyset \cdot (aa \cup bb) \cdot \sim \emptyset) \cap (a \cdot \sim \emptyset \cdot b)$ defines $(ab)^+$.

To show the equivalence of star-free expressions and first-order logic, we use the following variant of the back-and-forth property for FO.

Lemma 5.6. For words $u, v \in \Sigma^*$ and a number $m < \omega$, we have

$$u \equiv_{\text{FO}}^{m+1} v \quad \text{iff} \quad (u \in EaF \Leftrightarrow v \in EaF) \\ \text{for all } a \in \Sigma \text{ and all } \equiv_{\text{FO}}^m\text{-classes } E, F \in \Sigma^* / \equiv_{\text{FO}}^m.$$

Proof. By Proposition I.3.6, we have

- $$u \equiv_{\text{FO}}^{m+1} v$$
- iff for every $i < |u|$ there is some $k < |v|$ with $u, i \equiv_{\text{FO}}^m v, k$, and for every $k < |v|$ there is some $i < |u|$ with $u, i \equiv_{\text{FO}}^m v, k$
 - iff for every factorisation $u = u_0 a u_1$ (with $a \in \Sigma$) there is some factorisation $v = v_0 a v_1$ with $u_0 \equiv_{\text{FO}}^m v_0$ and $u_1 \equiv_{\text{FO}}^m v_1$, and for every factorisation $v = v_0 a v_1$ (with $a \in \Sigma$) there is some factorisation $u = u_0 a u_1$ with $u_0 \equiv_{\text{FO}}^m v_0$ and $u_1 \equiv_{\text{FO}}^m v_1$
 - iff $u \in EaF \Rightarrow v \in EaF$, for all \equiv_{FO}^m -classes E, F and $a \in \Sigma$
 $v \in EaF \Rightarrow u \in EaF$, for all \equiv_{FO}^m -classes E, F and $a \in \Sigma$
 - iff $u \in EaF \Leftrightarrow v \in EaF$, for all \equiv_{FO}^m -classes E, F . \square

Exercise 5.1. Show that, over the class of all finite words, every first-order formula is equivalent to a formula that uses only three variables (which can be quantified several times). \square

Proposition 5.7. A language $L \subseteq \Sigma^*$ is FO-definable if, and only if, it can be expressed by a star-free regular expression.

Proof. (\Leftarrow) Given a star-free expression α we construct an FO-formula $\varphi_\alpha(x, y)$ such that

$$w \models \varphi_\alpha(i, j) \quad \text{iff} \quad w[i, j] \in L(\alpha).$$

As usual the definition proceeds by induction on α .

$$\begin{aligned} \varphi_\emptyset(x, y) &:= \text{false}, \\ \varphi_a(x, y) &:= x = y \wedge P_a x, \\ \varphi_{\alpha \cap \beta}(x, y) &:= \varphi_\alpha(x, y) \wedge \varphi_\beta(x, y), \\ \varphi_{\alpha \cup \beta}(x, y) &:= \varphi_\alpha(x, y) \vee \varphi_\beta(x, y), \\ \varphi_{\sim \alpha}(x, y) &:= x \leq y \wedge \neg \varphi_\alpha(x, y), \end{aligned}$$

$$\begin{aligned} \varphi_{\alpha \cdot \beta}(x, y) &:= \exists u \exists v [x \leq u \wedge u + 1 = v \wedge v \leq y \\ &\quad \wedge \varphi_\alpha(x, u) \wedge \varphi_\beta(v, y)] \\ &\quad \vee \psi_{\alpha, \beta}(x, y) \vee \psi_{\beta, \alpha}(x, y), \end{aligned}$$

where in the last definition we have used the formula

$$\psi_{\alpha, \beta}(x, y) := \begin{cases} \varphi_\beta(x, y) & \text{if } \langle \rangle \in L(\alpha), \\ \text{false} & \text{otherwise.} \end{cases}$$

(\Rightarrow) It is sufficient to construct, for every \equiv_{FO}^m -class $K \in \Sigma^+ / \equiv_{\text{FO}}^m$, a star-free expression defining K . We do so by induction on m .

If $m = 0$, all words are \equiv_{FO}^m -equivalent. Hence, $K = \Sigma^+$ and we can use the star-free expression $\sim \emptyset$. For the inductive step, let K be an \equiv_{FO}^{m+1} -class. By Lemma 5.6, it follows that K can be written as a finite boolean combination of languages of the form EaF where $a \in \Sigma$ and E, F are \equiv_{FO}^m -classes. We can use the inductive hypothesis to obtain expressions α and β for, respectively, E and F . Hence, $\alpha \cdot a \cdot \beta$ defines EaF . As star-free expressions are closed under boolean operations, we can combine these expressions to get one for K . \square

Linear Temporal Logic

Finally, we can also use a certain form of modal logic.

Definition 5.8. Let Σ be an alphabet. The formulae of *linear temporal logic* LTL are built up from atomic formulae of the form P_a with $a \in \Sigma$ using (i) boolean operations and (ii) a binary modal operator \mathbf{U} . We read $\varphi \mathbf{U} \psi$ as ‘ φ until ψ ’. The semantics is defined as follows. Given a word $w \in \Sigma^+$ of length $n > 0$, we set

$$\begin{aligned} w \models P_a & \quad \text{iff} \quad w(0) = a, \\ w \models \varphi \mathbf{U} \psi & \quad \text{iff} \quad \text{there is some } 0 < k < n \text{ such that } w[k, n] \models \psi \\ & \quad \text{and } w[i, n] \models \varphi \text{ for all } 0 < i < k. \end{aligned}$$

The boolean operations are interpreted in the usual way. In addition we use the following abbreviations:

$$\begin{aligned} X\varphi &:= \text{false} \cup \varphi && (\text{'next } \varphi'), \\ F\varphi &:= \text{true} \cup \varphi && (\text{'finally } \varphi'), \\ G\varphi &:= \neg F\neg\varphi && (\text{'generally } \varphi'). \end{aligned}$$

We also introduce reflexive versions of \cup , F , G :

$$\begin{aligned} \varphi \cup_* \psi &:= \psi \vee (\varphi \cup \psi), \\ F_*\varphi &:= \varphi \vee F\varphi, \\ G_*\varphi &:= \varphi \wedge G\varphi. \end{aligned}$$

Examples. (a) $F_*(P_a \wedge F_*P_a)$ defines the language of all words containing at least two occurrences of the letter a .

(b) $G_*(P_a \rightarrow F_*P_b)$ says that every letter a is followed (not necessarily immediately) by a b .

(c) $\neg X\text{true}$ states that the word consists of a single letter.

(d) $P_a \wedge P_a \cup GP_b$ defines the language a^*b^+ .

Clearly, the logic LTL can be embedded into first-order logic.

Lemma 5.9. *For every LTL-formula φ , there exists an FO-formula ψ such that*

$$w \models \varphi \quad \text{iff} \quad w \models \psi, \quad \text{for all words } w.$$

Proof. Given φ we construct an FO-formula $\varphi^*(x)$ such that

$$w \models \varphi \quad \text{iff} \quad uw \models \varphi^*(|u|), \quad \text{for all } w, u \in \Sigma^*.$$

The definition proceeds by induction on φ .

$$\begin{aligned} P_a^*(x) &:= P_ax, \\ (\varphi \wedge \psi)^*(x) &:= \varphi^*(x) \wedge \psi^*(x), \\ (\neg\varphi)^*(x) &:= \neg\varphi^*(x), \\ (\varphi \cup \psi)^*(x) &:= \exists y[x < y \wedge \psi^*(y) \wedge \forall z[x < z < y \rightarrow \varphi^*(z)]]. \quad \square \end{aligned}$$

To conclude the proof of Theorem 5.1 it is now sufficient to show that first-order definable languages can also be defined in LTL. This is the hardest part of the theorem and requires a bit of preparation.

Lemma 5.10. *For every LTL-formula φ and every set $\Delta \subseteq \Sigma$, there exists an LTL-formula $\varphi^{(\Delta)}$ such that*

$$w \models \varphi^{(\Delta)} \quad \text{iff} \quad u \models \varphi \quad \text{for the maximal prefix } u \text{ of } w \text{ with } u \in \Delta^+.$$

Proof. We start by transforming the given formula φ into negation normal form where negations are only allowed in front of the atomic predicates P_a . This can be done using the laws of de Morgan and the equivalences

$$\begin{aligned} \neg F\psi &\equiv G\neg\psi, \\ \neg G\psi &\equiv F\neg\psi, \\ \neg(\psi \cup \vartheta) &\equiv G(\psi \wedge \neg\vartheta) \vee (\psi \wedge \neg\vartheta) \cup (\neg\psi \wedge \neg\vartheta). \end{aligned}$$

After this simplification, we can construct $\varphi^{(\Delta)}$ by induction on φ as follows. Setting $P_\Delta := \bigvee_{c \in \Delta} P_c$, we define

$$\begin{aligned} P_a^{(\Delta)} &:= \begin{cases} P_a & \text{if } a \in \Delta, \\ \text{false} & \text{otherwise,} \end{cases} \\ \neg P_a^{(\Delta)} &:= \begin{cases} \neg P_a \wedge P_\Delta & \text{if } a \in \Delta, \\ \text{true} & \text{otherwise,} \end{cases} \end{aligned}$$

$$\begin{aligned} (\varphi \wedge \psi)^{(\Delta)} &:= \varphi^{(\Delta)} \wedge \psi^{(\Delta)}, \\ (\varphi \vee \psi)^{(\Delta)} &:= \varphi^{(\Delta)} \vee \psi^{(\Delta)}, \\ (\varphi \cup \psi)^{(\Delta)} &:= P_\Delta \wedge [\varphi^{(\Delta)} \wedge P_\Delta] \cup \psi^{(\Delta)}. \quad \square \end{aligned}$$

The second construction we need is the following analogue of an interpretation for LTL.

Definition 5.11. Let Σ and Γ be alphabets, $\square \notin \Gamma$ a new letter, and let $(\psi_c)_{c \in \Gamma \cup \{\square\}}$ be a family of LTL-formulae such that, for every $w \in \Sigma^+$, there exists exactly one $c \in \Gamma \cup \{\square\}$ with $w \models \psi_c$.

The LTL-transduction $\tau : \Sigma^+ \rightarrow \Gamma^+$ defined by $(\psi_c)_c$ is the following function. Given a word $w \in \Sigma^+$ of length $n := |w|$, let $c_i \in \Gamma \cup \{\square\}$, for $i < n$, be the letters such that

$$w, i \models \psi_{c_i}.$$

Then $\tau(w)$ is the word obtained from $\langle c_0, \dots, c_{n-1} \rangle$ by deleting all letters that are equal to \square . \lrcorner

Lemma 5.12. *Let $\tau : \Sigma^+ \rightarrow \Gamma^+$ be an LTL-transduction. For every LTL-formula φ , there exists an LTL-formula φ^τ such that*

$$\tau(w) \models \varphi \quad \text{iff} \quad w \models \varphi^\tau, \quad \text{for all } w \in \Sigma^+.$$

Proof. Given φ , we will define a formula φ^* such that

$$\tau(w) \models \varphi \quad \text{iff} \quad w \models \varphi^*, \quad \text{for all } w \in \Sigma^+ \text{ with } w \models \neg\psi_\square.$$

Then we can set

$$\varphi^\tau := \psi_\square \cup (\neg\psi_\square \wedge \varphi^*).$$

To define φ^* we proceed by induction on φ .

$$\begin{aligned} P_a^* &:= \psi_a, \\ (\varphi \wedge \vartheta)^* &:= \varphi^* \wedge \vartheta^*, \\ (\neg\varphi)^* &:= \neg(\varphi^*), \\ (\varphi \cup \vartheta)^* &:= (\neg\psi_\square \rightarrow \varphi^*) \cup (\neg\psi_\square \wedge \vartheta^*). \end{aligned} \quad \square$$

As an application, let us show how to compute products in aperiodic semigroups using LTL.

Proposition 5.13. *Let \mathfrak{S} be a finite, aperiodic semigroup. For every element $d \in S$, there exists an LTL-formula φ_d such that*

$$w \models \varphi_d \quad \text{iff} \quad \pi(w) = d, \quad \text{for all } w \in S^+.$$

Proof. We will prove the following more general claim. Given a finite, aperiodic semigroup \mathfrak{S} , a non-empty subset $C \subseteq S$, and an element $d \in S$, there exists an LTL-formula φ_d such that

$$w \models \varphi_d \quad \text{iff} \quad \pi(u) = d \quad \text{where } u \text{ is the maximal prefix of } w \text{ with } u \in C^+.$$

The proof proceeds by induction on $|S|$ and $|C|$.

If $S = \{c\}$, we can set $\varphi_c := \text{true}$. If $C = \{c\}$, we have to check whether $w = c^n v$ where $c^n = d$ and v does not start with c . As \mathfrak{S} is aperiodic, there exists some number k such that $c^n = c^k$, for all $n \geq k$. Setting

$$\psi_1 := P_c \quad \text{and} \quad \psi_{n+1} := P_c \wedge \neg\psi_n,$$

we obtain formulae such that

$$w \models \psi_n \quad \text{iff} \quad w = c^n v, \quad \text{for some } v \in S^+.$$

Since there exists at most one number $n \leq k$ with $c^n = d$, we can now set

$$\varphi_d := \begin{cases} \psi_n \wedge \neg\psi_{n+1} & \text{if } d = c^n \text{ with } n < k, \\ \psi_k & \text{if } d = c^k, \\ \text{false} & \text{if } c^n \neq d \text{ for all } n \leq k. \end{cases}$$

For the inductive step, suppose that we have already proved the claim for all semigroups \mathfrak{S}' and all subsets $C' \subseteq S'$ such that either $|S'| < |S|$, or $|S'| = |S|$ and $|C'| < |C|$. We first consider the case where, for every element $c \in C$, left-multiplication $\sigma_c := a \mapsto ca$ by c is bijective. Since \mathfrak{S} is aperiodic, there exists some number k such that $c^{k+1} = c^k$. Consequently, $\sigma_c^{k+1} = \sigma_c^k$. As σ_c is bijective, we can divide this equation by σ_c^k and obtain $\sigma_c = \text{id}$. Hence, we have $ca = a$, for all $c \in C$ and $a \in S$, and it follows that

$$\pi(w) = d \quad \text{iff} \quad \text{the last element of } w \text{ is equal to } d, \quad \text{for } w \in C^+.$$

Thus, we can set

$$\varphi_d := \left[\bigvee_{c \in C} P_c \right] \cup \left[P_d \wedge \neg\bigvee_{c \in C} P_c \right].$$

It remains to consider the case where there is some $c \in C$ such that the function $a \mapsto ca$ is not bijective. Set $T := cS$ and $D := C \setminus \{c\}$. By assumption, $T \subset S$. Furthermore, T induces a subsemigroup of \mathfrak{S} since $ca \cdot cb = c(acb) \in T$.

Let us define a *block* of $w \in C^+$ to be a maximal factor of the form $c^n u$ with $n < \omega$ and $u \in D^*$. To compute $\pi(w)$ we will proceed in two steps: first we multiply every block of w and then we multiply the results.

To accomplish the former we define formulae ψ_d , for $d \in S$, such that

$$w \models \psi_d \quad \text{iff} \quad w = ac^n uv \quad \text{and} \quad \pi(c^n u) = d \\ \text{where } c^n u \text{ is a block of } w, \quad a \in D, \quad v \in C^*.$$

Note that this formula is supposed to be evaluated at the position *preceding* the block in question. This is because we need to verify that we are at the beginning of a block and we cannot look backwards in LTL. By inductive hypothesis, we can construct formulae $(\varphi_a^c)_a$ and $(\varphi_a^D)_a$ evaluating products of sequences in, respectively, $\{c\}^+$ and D^+ . We set

$$\psi_d := \neg P_c \wedge X(P_c \wedge \psi'_d)$$

where

$$\psi'_d := [\varphi_d^c \wedge GP_c] \vee \bigvee_{\substack{a, b \in S \\ ab=d}} [\varphi_a^c \wedge [P_c \cup \varphi_b^D]].$$

(The first part deals with the special case where we are in the last block and this block is of the form c^n without elements from D . The second part is for the more common case where the current block does contain elements from D .) Together with the formula

$$\psi_\square := P_c \vee \neg X P_c$$

we obtain a family $(\psi_d)_{d \in S \cup \{\square\}}$ that defines an LTL-transduction τ that maps w to the sequence of products of the blocks (excluding the first block which we treat separately). This sequence belongs to T^+ . Since T is a proper

subsemigroup of S we can use the inductive hypothesis to obtain formulae φ_d^T for evaluating the resulting product. This leads to the following definition.

$$\varphi_d := [\psi'_d \wedge G\psi_\square] \vee \bigvee_{\substack{a, b \in S \\ ab=d}} [\psi'_a \wedge [\psi_\square \cup (\varphi_b^T)^\tau]].$$

(The first clause is for the case where there is only one block, the second one if there are more.) \square

As we have already established the equivalence between FO-definability and recognisability in an aperiodic semigroup, we now immediately obtain the last missing piece for the proof of Theorem 5.1.

Corollary 5.14. *Every FO-definable language is LTL-definable.*

Proof. Let $L \subseteq \Sigma^+$ be FO-definable. By Proposition 5.4, we can find a homomorphism $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ to a finite aperiodic semigroup \mathfrak{S} such that $L = \eta^{-1}[P]$ for some $P \subseteq S$. We construct an LTL-formula ψ defining L as follows. Let φ_d , $d \in S$, be the LTL-formulae from Proposition 5.13. Given a word $w = a_0 \dots a_{n-1} \in \Sigma^+$, let $w^\eta = \eta(a_0) \dots \eta(a_{n-1}) \in S^+$ be the word obtained from w by replacing each letter by its image under η . It follows that

$$w \in L \quad \text{iff} \quad \eta(w) \in P \quad \text{iff} \quad w^\eta \models \bigvee_{a \in P} \varphi_a \quad \text{iff} \quad w \models \psi,$$

where ψ is the formula obtained from $\bigvee_{a \in P} \varphi_a$ by replacing every predicate P_c with $c \in S$ by the formula

$$\vartheta_c := \bigvee_{b \in \eta^{-1}(c) \cap \Sigma} P_b.$$

\square

Notes

Ramseyan splits were introduced by Colcombet, extending earlier results by Simon [195] on factorisation trees. Their existence for arbitrary linear orders

is due to [50]. Our presentation follows expositions by Bojańczyk [26] and Colcombet [51].

The equivalence between monadic second-order logic and automata was independently discovered by Büchi [34], Elgot [78], and Trakhtenbrot [206]. The equivalence between star-free regular expressions and aperiodic monoids is due to Schützenberger [190], the one between star-free regular expressions and first-order logic due to McNaughton and Papert [142], and the equivalence to LTL is due to Kamp [113].

III Infinite Words

1 Ramsey Theory

OUR NEXT AIM IS TO DO what we just did in Section II.4 for languages of infinite words. Unfortunately this entails a bit of technical overhead. In particular, we need a few results from a branch of combinatorics called *Ramsey Theory*. We have already seen one result of this kind in Section II.3: the Lemma of Simon. In this section we will derive several more. The simplest example of such a result is the statement that every infinite undirected graph contains an infinite clique or an infinite independent set.

Definition 1.1. Let A be a linear order.

- (a) We denote by $[A]^2$ the set of all pairs $\langle i, k \rangle \in A^2$ with $i < k$.
- (b) A *finite colouring* of A is a function $\lambda : [A]^2 \rightarrow C$ where C is a finite set of colours.
- (c) Let \mathfrak{S} be a finite semigroup. A finite colouring $\lambda : [A]^2 \rightarrow S$ is *additive* if

$$\lambda(x, y) \cdot \lambda(y, z) = \lambda(x, z), \quad \text{for all } x < y < z. \quad \text{J}$$

Theorem 1.2 (Ramsey). *Let $\lambda : [\omega]^2 \rightarrow C$ be a finite colouring of ω . There exists an infinite subset $I \subseteq \omega$ such that*

$$\lambda(i, k) = \lambda(j, l), \quad \text{for all } i < k \text{ and } j < l \text{ in } I.$$

Proof. We construct an increasing sequence $n_0 < n_1 < \dots$ of indices, a sequence $c_0, c_1, \dots \in C$ of colours, and a decreasing sequence $J_0 \supseteq J_1 \supseteq \dots$ of infinite sets such that, for every $i < \omega$,

$$n_i \in J_i \quad \text{and} \quad \lambda(n_i, k) = c_i, \quad \text{for all } k \in J_{i+1}.$$

We start with $n_0 := 0$ and $J_0 := \omega$. By induction, suppose that we have already defined n_i and J_i . For $c \in C$, set

$$L_c := \{k \in J_i \mid k > n_i \text{ and } \lambda(n_i, k) = c\}.$$

Then $J_i \setminus [n_i + 1] = \bigcup_{c \in C} L_c$. As J_i is infinite and C is finite, there is some element $c_i \in C$ such that L_{c_i} is infinite. We set

$$J_{i+1} := L_{c_i} \quad \text{and} \quad n_{i+1} := \min J_{i+1}.$$

Having defined $(n_i)_{i < \omega}$, $(c_i)_{i < \omega}$, and $(J_i)_{i < \omega}$, we consider the sets

$$M_c := \{i < \omega \mid c_i = c\}, \quad \text{for } c \in C.$$

Note that $n_j \in J_j \subseteq J_{i+1}$, for $j > i$, implies that

$$\lambda(n_i, n_j) = c, \quad \text{for all } i < j \text{ in } M_c.$$

Since $\bigcup_{c \in C} M_c = \omega$, there is some $c \in C$ such that M_c is infinite. We set $I := \{n_i \mid i \in M_c\}$. \square

Remark. This theorem holds more generally for colourings of k -tuples instead of pairs. The proof is a straightforward induction on k using the argument from the above proof for the inductive step. \lrcorner

Exercise 1.1. (a) Let $\mathfrak{G} = \langle V, E \rangle$ be an infinite undirected graph. Prove that there exists an infinite set $X \subseteq V$ such that either all vertices in X are adjacent, or none of them are.

(b) Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph with at least 6 vertices. Prove that there exists three vertices $x, y, z \in V$ that are either all connected by an edge, or none of them are. \lrcorner

Exercise 1.2. Let \mathfrak{S} be a finite semigroup and a_0, a_1, \dots an infinite sequence of elements of S . Prove that there exists an increasing sequence $k_0 < k_1 < \dots$ of indices and two elements $b, e \in S$ such that

$$be = b, \quad ee = e, \quad b = a_0 \cdots a_{k_0-1}, \quad \text{and} \quad e = a_{k_i} \cdots a_{k_{i+1}-1},$$

for all $i < \omega$. \lrcorner

Exercise 1.3. Let \mathfrak{S} be a finite semigroup, $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ a homomorphism, and let $L \subseteq \Sigma^\omega$ be a language of the form

$$L = \bigcup_{i < n} \eta^{-1}(b_i) (\eta^{-1}(e_i))^\omega, \quad \text{for } n < \omega \text{ and } b_i, e_i \in S,$$

where $X^\omega := \{x_0 x_1 x_2 \dots \mid x_i \in X\}$.

Prove that the complement $\Sigma^\omega \setminus L$ is also of the form

$$\Sigma^\omega \setminus L = \bigcup_{i < m} \eta^{-1}(c_i) (\eta^{-1}(f_i))^\omega,$$

for suitable $m < \omega$ and $c_i, f_i \in S$. \lrcorner

Exercise 1.4. A *well-quasi-order* is a partial order $\langle A, \leq \rangle$ that does not contain any infinite descending sequence and any infinite antichain (i.e., a set of pairwise incomparable elements).

(a) Prove that every infinite partial order contains an infinite set that is either an ascending chain, a descending chain, or an antichain.

(b) Prove that $\langle A, \leq \rangle$ is a well-quasi-order if, and only if, for every infinite sequence a_0, a_1, a_2, \dots in A there are indices $i < k$ with $a_i \leq a_k$.

(c) Let Σ be a finite alphabet. We define an ordering on Σ^* by setting $x \leq y$ if the word x can be obtained from y by deleting some letters. Prove that $\langle \Sigma^*, \leq \rangle$ is a well-quasi-order.

Hint. Assume that $\langle \Sigma^*, \leq \rangle$ is not a well-quasi-order and find words w_0, w_1, \dots such that, for every $n < \omega$, the sequence w_0, \dots, w_n can be continued to an infinite sequence violating the condition in (b). \lrcorner

For additive colourings, we can improve the Theorem of Ramsey. One such result is the Lemma of Simon that we proved in Section II.3. We can interpret Simon's Lemma as a recursive version of the Theorem of Ramsey where we partition the input word not only once, but each of the resulting factors recursively until only single letters are left. Of course we could just repeatedly use the Theorem of Ramsey to get such a decomposition. The point of Simon's Lemma is that a bounded number of iterations is sufficient for this.

While Simon's Lemma is a powerful result, it does have one drawback: the split we obtain depends on the whole input word. Below we will need a way to compute a split while reading the word from left-to-right in a deterministic fashion without having to know how the part we have not seen yet looks like. In the following we will prove a result of Colcombet which shows how this can be done. The resulting split will unfortunately not be fully Ramseyan, it will only satisfy a slightly weaker property, which is nonetheless still sufficient for many applications.

The starting point is the following problem: given an additive colouring $\lambda : [I]^2 \rightarrow S$ of a linear order \mathfrak{I} , we would like to find a colouring $\chi : I \rightarrow C$ of the elements of \mathfrak{I} such that we can recover λ from χ . That is, we want to reduce a labelling of *pairs* to a labelling of *singletons*. The proof uses techniques from semigroup theory and Green's relations.

Definition 1.3. Let \mathfrak{S} be a finite semigroup.

(a) A *right action* of \mathfrak{S} on a set Q is a function $\triangleright : Q \times S \rightarrow Q$ satisfying the equation

$$q \triangleright (ab) = (q \triangleright a) \triangleright b, \quad \text{for all } q \in Q \text{ and } a, b \in S.$$

(b) A *J-chain* of \mathfrak{S} is a tuple $\bar{a} = \langle a_0, \dots, a_m \rangle \in S^+$ such that

$$a_0 <_J \dots <_J a_m.$$

We denote the set of all J-chains of \mathfrak{S} by $\text{Chain}_J(\mathfrak{S})$.

(c) We define a right action \triangleright of \mathfrak{S} on $\text{Chain}_J(\mathfrak{S})$ by

$$\langle a_0, \dots, a_m \rangle \triangleright b := \langle a_0, \dots, a_{k-1}, (a_k \cdots a_m b) \rangle,$$

where $0 \leq k \leq m+1$ is the maximal index such that the above tuple is a J-chain.

(d) Let $\alpha \leq \omega$ and let $\lambda : [\alpha]^2 \rightarrow S$ be an additive colouring. A function $\chi : \alpha \rightarrow \text{Chain}_J(\mathfrak{S})$ is a *J-chain labelling* for λ if, for all $0 < i < \alpha$,

$$\chi(i) = \chi(i-1) \triangleright \lambda(i-1, i).$$

Example. Let $M = \{1, a, b, ab, ba, aba, o\}$ be the monoid from the example on page 67 with three J-classes:

$$\{o\} \leq_J \{b, ba, ab, aba\} \leq_J \{1, a\}.$$

For the colouring $\lambda : [10]^2 \rightarrow M$ with

$$0 \xrightarrow{b} 1 \xrightarrow{a} 2 \xrightarrow{a} 3 \xrightarrow{b} 4 \xrightarrow{b} 5 \xrightarrow{a} 6 \xrightarrow{a} 7 \xrightarrow{a} 8 \xrightarrow{b} 9,$$

we obtain a J-chain labelling

$$\begin{array}{c} |a| \xrightarrow{b} |ab| \xrightarrow{a} \begin{array}{|c|} \hline a \\ \hline ab \\ \hline \end{array} \xrightarrow{a} \begin{array}{|c|} \hline 1 \\ \hline ab \\ \hline \end{array} \xrightarrow{b} |o| \xrightarrow{b} \begin{array}{|c|} \hline b \\ \hline o \\ \hline \end{array} \xrightarrow{a} \\ \xrightarrow{a} \begin{array}{|c|} \hline a \\ \hline b \\ \hline o \\ \hline \end{array} \xrightarrow{a} \begin{array}{|c|} \hline 1 \\ \hline b \\ \hline o \\ \hline \end{array} \xrightarrow{a} \begin{array}{|c|} \hline a \\ \hline b \\ \hline o \\ \hline \end{array} \xrightarrow{b} |o| \end{array}$$

where we have written a J-chain $\langle a_0, \dots, a_{m-1} \rangle$ as a column

$$\begin{array}{|c|} \hline a_{m-1} \\ \hline \vdots \\ \hline a_0 \\ \hline \end{array}.$$

We will prove that we can recover λ from a J-chain labelling.

Definition 1.4. Let $\chi : \alpha \rightarrow \text{Chain}_J(\mathfrak{S})$ be a J-chain labelling for $\lambda : [\alpha]^2 \rightarrow S$.

(a) Let $\mu : S^+ \rightarrow S$ be the function mapping a tuple $\langle a_0, \dots, a_{m-1} \rangle$ to its last element a_{m-1} , and let $\pi : S^* \rightarrow S^1$ be the function mapping a tuple $\langle a_0, \dots, a_{m-1} \rangle$ to the product $a_0 \cdots a_{m-1}$ of its components.

(b) For positions $i, k < \alpha$, we define

$$i <_\chi^\circ k \quad : \text{iff} \quad i < k \quad \text{and} \quad \mu\chi(j) \not\leq_J \mu\chi(i), \quad \text{for all } i \leq j \leq k,$$

$$i \leq_\chi k \quad : \text{iff} \quad i <_\chi^\circ k \quad \text{and} \quad \mu\chi(i) = \mu\chi(k).$$

Lemma 1.5. Let $\alpha \leq \omega$, let $\chi : \alpha \rightarrow \text{Chain}_J(\mathfrak{S})$ be a J -chain labelling for $\lambda : [\alpha^2] \rightarrow S$, and let $i \leq k < \alpha$ be positions with

$$i \leq_\chi^\circ k, \quad \chi(i) = sa, \quad \text{and} \quad \chi(k) = t, \quad \text{for } s, t \in S^*, a \in S.$$

(a) There exist $b \in S$ and $x \in S^*$ such that

$$t = s \frown b \frown x, \quad b \equiv_J a, \quad \text{and} \quad \pi(b \frown x) = a \cdot \lambda(i, k).$$

(b) If $t = s \frown b \frown x$ with $x \neq \langle \rangle$, there is some position $i \leq j < k$ such that $\lambda(j, k) = \pi(x)$.

(c) If $t = s \frown b$ and $i < k$, then $a \equiv_J \lambda(i, k)$.

Proof. (a) We prove the statement by induction on the number of positions between i and k . If $i = k$, then

$$t = \chi(k) = \chi(i) = s \frown a,$$

and we can set $b := a$ and $x := \langle \rangle$.

For the inductive step, let k' be the immediate predecessor of k and assume that $i \leq k'$. By inductive hypothesis, it follows that

$$\chi(k') = s \frown b \frown x,$$

for some $b \in S$ and $x \in S^*$ such that

$$b \equiv_J a \quad \text{and} \quad \pi(b \frown x) = a \cdot \lambda(i, k').$$

By definition of χ , there is a factorisation $u \frown v$ of $s \frown b \frown x$ such that

$$t = u \frown \pi(v \frown c), \quad \text{where} \quad c := \lambda(k', k).$$

We claim that s is a prefix of u . For a contradiction, suppose otherwise. Then $s = u \frown d \frown y$ and $v = d \frown y \frown b \frown x$ where $y \in S^*$ and $d \in S$ is the first element of v . Setting $d' := \pi(v \frown c)$ we obtain

$$d' = \pi(d \frown y \frown b \frown x \frown c) \leq_J d.$$

Since $s \frown b \frown x = u \frown d \frown y \frown b \frown x$ is a valid configuration, it follows that $d <_J b$. Consequently,

$$\mu\chi(k') = d' \leq_J d <_J b \equiv_J a.$$

A contradiction.

We have shown that t is of the form

$$t = s \frown z \frown d' \quad \text{where} \quad d' := \pi(v \frown c) \quad \text{and} \quad z \frown v = b \frown x.$$

By definition of χ and by inductive hypothesis, it further follows that

$$\begin{aligned} a \cdot \lambda(i, k) &= a \cdot \lambda(i, k') \cdot \lambda(k', k) \\ &= \pi(b \frown x) \cdot c = \pi(z \frown v \frown c) = \pi(z) \cdot \pi(v \frown c) = \pi(z \frown d'). \end{aligned}$$

It therefore remains to prove that the first element of $z \frown d'$ is J -equivalent to a . If $z \neq \langle \rangle$, then $z = b \frown z'$, for some $z' \in S^*$. Hence, $t = s \frown b \frown z' \frown d'$ where $b \equiv_J a$. If $z = \langle \rangle$, then $t = s \frown d'$ where

$$d' = \pi(v \frown c) = \pi(z \frown v \frown c) = \pi(b \frown x \frown c) \leq_J b \equiv_J a.$$

Since $d' = \mu\chi(k') \not\leq_J a$, it follows that $d' \equiv_J a$.

(b), (c) We prove both statements by induction on the number of positions between i and k . If $i = k$, then $b = a$, $x = \langle \rangle$, and (b) and (c) hold trivially.

For the inductive step, let k' be the immediate predecessor of k and suppose that $i \leq k'$. By (a), it follows that $\chi(k') = s \frown b' \frown x'$ where

$$b' \equiv_J a \equiv_J b \quad \text{and} \quad \pi(b \frown x) = a \cdot \lambda(i, k).$$

Set $c := \lambda(k', k)$. The definition of χ implies that either

$$x = \langle \rangle \quad \text{and} \quad b = \pi(b' \frown x' \frown c),$$

or $b = b'$, $x' = y \frown z$, and $x = y \frown \pi(z \frown c)$, for some $y, z \in S^*$.

To prove (b), suppose that $x \neq \langle \rangle$. Then $\pi(x) = \pi(y \frown z \frown c) = \pi(x' \frown c)$. Thus, it is sufficient to find a position $i \leq j < k$ such that $\lambda(j, k) = \pi(x' \frown c)$.

If $x' = \langle \rangle$, we can take $j := k'$. If $x' \neq \langle \rangle$, we can use the inductive hypothesis to obtain a position $i \leq j < k'$ with $\lambda(j, k') = \pi(x')$. Then $\lambda(j, k) = \lambda(j, k') \cdot c = \pi(x' \hat{\ } c)$.

To prove (c), suppose that $x = \langle \rangle$. Set $d' := \pi(x' \hat{\ } c)$. Then $b = b' \cdot d'$ and, by definition of χ , the sequence $s \hat{\ } b' \hat{\ } d'$ is not a J-chain, while $s \hat{\ } b$ is one. Furthermore, we have $b' \equiv_J a \equiv_J b = b' d'$, which implies that $b' \leq_J d'$. Hence, the only possible reason for $s \hat{\ } b' \hat{\ } d'$ not being a J-chain is that $b' \not\leq_J d'$. Consequently, $d' \equiv_J b' \equiv_J a$. We choose a position $i \leq j < k$ as follows. If $x' = \langle \rangle$, we set $j := k'$. If $x' \neq \langle \rangle$, we use (b) to choose j such that $\lambda(j, k') = \pi(x')$. In both cases it follows that

$$\lambda(i, k) \leq_J \lambda(j, k) = \pi(x' \hat{\ } c) = d' \equiv_J a.$$

Moreover, we have seen above that $b = \pi(b \hat{\ } x) = a \cdot \lambda(i, k)$. Therefore, $\lambda(i, k) \geq_J b \equiv_J a$ and it follows that $a \equiv_J \lambda(i, k)$. \square

Corollary 1.6. *Let $\alpha \leq \omega$ and let $\chi : \alpha \rightarrow \text{Chain}_J(\mathfrak{S})$ be a J-chain labelling for $\lambda : [\alpha^2] \rightarrow S$. If $i \leq_\chi k$, then*

$$\mu\chi(i) \cdot \lambda(i, k) = \mu\chi(i) \quad \text{and} \quad \mu\chi(i) \equiv_J \lambda(i, k).$$

Proof. By Lemma 1.5 (a), we have

$$\chi(i) = s \hat{\ } a \quad \text{and} \quad \chi(k) = s \hat{\ } b \hat{\ } x,$$

for some $s, x \in S^*$ and $a, b \in S$ such that $b \equiv_J a$ and $\pi(b \hat{\ } x) = a \cdot \lambda(i, k)$. Note that, if $x \neq \langle \rangle$, then $x = y \hat{\ } a$, for some $y \in S^*$, and $b \not\leq_J a$ implies that $t \hat{\ } a = s \hat{\ } b \hat{\ } y \hat{\ } a$ is not a J-chain. Hence, $x = \langle \rangle$. This implies that

$$t = s, \quad b = a, \quad \text{and} \quad a \cdot \lambda(i, k) = \pi(b \hat{\ } x) = b = a.$$

Furthermore, it follows by Lemma 1.5 (c) that $a \equiv_J \lambda(i, k)$. \square

We are finally able to state our deterministic version of Simon's Lemma. Recall the notion of a split σ and the corresponding preorder \sqsubseteq_σ from Section II.3.

Definition 1.7. Let A and B be linear orders and $\lambda : [A]^2 \rightarrow S$ an additive colouring. A function $\sigma : A \rightarrow B$ is a *weak Ramseyan split* of λ if

$$\lambda(x, y) = \lambda(x, y) \cdot \lambda(x', y'), \quad \text{for all } x \sqsubseteq_\sigma y \text{ and } x' \sqsubseteq_\sigma y' \\ \text{such that } x \sqsubseteq_\sigma x' \text{ or } x' \sqsubseteq_\sigma x.$$

The above results allow us to construct a weak Ramseyan split as follows.

Lemma 1.8. *Let \mathfrak{S} be a finite semigroup of size $N := |S|$ and $\chi : \alpha \rightarrow \text{Chain}_J(\mathfrak{S})$ a J-chain labelling for an additive colouring $\lambda : [\alpha]^2 \rightarrow S$. If $v : S \rightarrow [N]$ is a bijection such that*

$$a <_J b \quad \text{implies} \quad v(a) > v(b),$$

then $\sigma := v \circ \mu \circ \chi$ is a weak Ramseyan split for λ .

Proof. Consider positions $x < y$ and $x' < y'$ with $x \approx_\sigma y \approx_\sigma x' \approx_\sigma y'$. Then

$$a := \mu\chi(x) = \mu\chi(y) = \mu\chi(x') = \mu\chi(y')$$

and $\mu\chi(z) \not\leq_J a$, for all z between any two of x, y, x', y' .

Consequently, $x \leq_\chi y$ and $x' \leq_\chi y'$. By Corollary 1.6 it follows that

$$a \cdot \lambda(x, y) = a \quad \text{and} \quad a \equiv_J \lambda(x, y), \\ a \cdot \lambda(x', y') = a \quad \text{and} \quad a \equiv_J \lambda(x', y').$$

Applying Corollary II.2.11 (a) to the values $b := \lambda(x, y)$ and $c := \lambda(x', y')$, we obtain

$$\lambda(x, y) \cdot \lambda(x', y') = b \cdot c = b = \lambda(x, y). \quad \square$$

We can compute weak Ramseyan splits by an automaton.

Definition 1.9. (a) A *deterministic finite-state transducer*

$$\mathcal{T} = \langle Q, \Sigma, \Gamma, q_0, \delta, \eta \rangle$$

consists of a finite set Q of states, an input alphabet Σ , an output alphabet Γ , an initial state $q_o \in Q$, an output function $\eta : Q \rightarrow \Gamma$, and a transition function $\delta : Q \times \Sigma \rightarrow Q$.

Let $\mathcal{T} = \langle Q, \Sigma, \Gamma, q_o, \delta, \eta \rangle$ be a transducer. The run of \mathcal{T} on a word $w = (c_i)_{i < \alpha} \in \Sigma^\omega$ is the sequence $(q_i)_{i < \beta}$ of states where q_o is the initial state and

$$q_{i+1} := \delta(q_i, c_i), \quad \text{for all } i, \quad \text{and} \quad \beta := \begin{cases} \alpha + 1 & \text{if } \alpha < \omega, \\ \omega & \text{if } \alpha = \omega. \end{cases}$$

Every transducer \mathcal{T} defines a function $\mathcal{T} : \Sigma^\omega \rightarrow \Gamma^\omega$ that maps an input word $w \in \Sigma^\omega$ to the word

$$\mathcal{T}(w) = (\eta(q_i))_{i < \beta}, \quad \text{where } (q_i)_{i < \beta} \text{ is the run of } \mathcal{T} \text{ on } w.$$

(b) Let \mathfrak{S} a finite semigroup and $N < \omega$ a natural number. We say that a transducer $\mathcal{T} = \langle Q, S, [N], q_o, \delta, \eta \rangle$ computes weak Ramseyan splits for \mathfrak{S} if, for every additive colouring $\lambda : [\alpha]^2 \rightarrow S$ with $\alpha \leq \omega$, the function $\sigma : \alpha \rightarrow [N]$ defined by

$$\sigma = \mathcal{T}(w) \quad \text{where} \quad w := (\lambda(i, i+1))_{i+1 < \alpha},$$

is a weak Ramseyan split of λ . J

Theorem 1.10 (Colcombet). *Given a finite semigroup \mathfrak{S} of size $N := |S|$, we can effectively construct a deterministic finite-state transducer*

$$\mathcal{T} = \langle Q, S, [N], q_o, \delta, \eta \rangle$$

that computes weak Ramseyan splits for \mathfrak{S} .

Proof. We use the set $Q := \text{Chain}_J(\mathfrak{S})$ of all J-chains as states of the transducer. Note that this set is finite, since there are at most $|S|$ J-classes. The initial state q_o is an arbitrary J-chain. We define the transition function δ using the right action $\triangleright : \text{Chain}_J(\mathfrak{S}) \times S \rightarrow \text{Chain}_J(\mathfrak{S})$ by

$$\delta(q, a) := q \triangleright a.$$

Then the run of \mathcal{T} on a given input λ is a J-chain labelling χ for λ . Fixing a bijection $v : S \rightarrow [N]$ as in Lemma 1.8, we can define the output function $\eta : Q \rightarrow [N]$ by $\eta(q) := v(\mu(q))$. □

Exercise 1.5. (a) Let Σ be a finite alphabet, $w \in \Sigma^\omega$, and $k < \omega$. Prove that there are sets $Q_0, \dots, Q_{n-1} \subseteq \omega$ such that, for every MSO-theory θ of quantifier rank k , there exists an FO-formula $\varphi_\theta(x, y)$ such that

$$\langle \omega, \bar{Q} \rangle \models \varphi_\theta(x, y) \quad \text{iff} \quad \text{Th}_{\text{MSO}}^k(w[x, y]) = \theta.$$

(b) Let τ be an MSO-interpretation τ and $P_0, \dots, P_{m-1} \subseteq \omega$ monadic parameters. Prove that there exist an FO-interpretation σ and sets $Q_0, \dots, Q_{n-1} \subseteq \omega$ such that

$$\tau(\langle \omega, \bar{P} \rangle) = \sigma(\langle \omega, \bar{Q} \rangle).$$

2 The Theory of ω

Büchi was the first one to show that the monadic second-order theory of $\langle \omega, \leq \rangle$ is decidable. His original proof uses automata for ω -words. In this section we present an alternative model-theoretic proof due to Shelah. We start with computing the theories of finite linear orders.

Definition 2.1. For $\bar{n} \in \omega^*$ and $m < \omega$, we denote by $\Theta_{\bar{n}}(m)$ the set of all sets of $\text{MSO}_0^{\bar{n}}[\leq]$ -formulae with free variables X_0, \dots, X_{m-1} . And we set

$$\Phi_{\bar{n}}(m) := \left\{ \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathfrak{A}, P_0 \dots P_{m-1}) \mid \mathfrak{A} \text{ a finite linear order and } P_0, \dots, P_{m-1} \subseteq A \right\}.$$

Lemma 2.2. *We can equip $\Theta_{\bar{n}}(m)$ with two operations \cdot and $^\omega$ such that*

$$\text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathfrak{A}, \bar{P} + \mathfrak{B}, \bar{Q}) = \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathfrak{A}, \bar{P}) \cdot \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathfrak{B}, \bar{Q})$$

$$\text{and} \quad \text{Th}_{\text{MSO}_0}^{\bar{n}}\left(\sum_{i < \omega} \mathfrak{A}_i, \bar{P}\right) = \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathfrak{A}, \bar{P})^\omega,$$

for all linear orders \mathfrak{A} and \mathfrak{B} and all parameters \bar{P}, \bar{Q} .

Proof. This follows immediately from Proposition I.4.25. \square

Proposition 2.3. *Given $\bar{n} \in \omega^*$ and $m < \omega$, we can compute $\Phi_{\bar{n}}(m)$.*

Proof. The set

$$\Psi := \{ \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{A}, \bar{P}) \mid \mathcal{A} = \langle A, \leq \rangle \text{ a linear order with } |A| \leq 1 \}$$

is a finite subset of $\Theta_{\bar{n}}(m)$ which we can compute from \bar{n} and m . As every finite linear order can be written as a finite ordered sum of one-element linear orders, it follows that $\Phi_{\bar{n}}(m)$ is the subsemigroup of $\langle \Theta_{\bar{n}}(m), \cdot \rangle$ generated by the set Ψ . Hence, we can compute $\Phi_{\bar{n}}(m)$ from $\Theta_{\bar{n}}(m)$ and Ψ . \square

The next lemma allows us to go from finite orders to infinite ones.

Lemma 2.4. *Let $\bar{n} \in \omega^*$ and $m < \omega$. There exists a tuple $\bar{r} \in \omega^*$ of length $|\bar{r}| = |\bar{n}|$ such that, for every $\sigma \in \Phi_{\bar{n}}(m)$, we can compute $\sigma^\omega \in \Theta_{\bar{n}}(m)$ from σ and $\text{Th}_{\text{MSO}_0}^{\bar{r}}(\omega, \leq)$.*

Proof. Note that

$$\sigma^\omega = \text{Th}_{\text{MSO}_0}^{\bar{n}}(\sum_{i < \omega} \mathcal{A}, \bar{P}),$$

where \mathcal{A} is any finite linear order with $\text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{A}, \bar{P}) = \sigma$. Let $\tau \in \Theta_{\bar{n}}(m)$ and let ϑ_τ be an $\text{MSO}_0^{\bar{n}}$ -formula equivalent to τ . It follows that

$$\sigma^\omega = \tau \quad \text{iff} \quad \sum_{i < \omega} \mathcal{A}, \bar{P} \models \vartheta_\tau.$$

According to Theorem I.4.24, we can compute formulae $\vartheta'_\tau(\bar{Z}) \in \text{MSO}_0^{\bar{r}}$ and $\chi_0, \dots, \chi_{l-1} \in \text{MSO}_0^{\bar{n}}$ such that

$$\sum_{i < \omega} \mathcal{A}, \bar{P} \models \vartheta_\tau \quad \text{iff} \quad \langle \omega, \leq \rangle \models \vartheta'_\tau(\llbracket \chi_0 \rrbracket, \dots, \llbracket \chi_{l-1} \rrbracket).$$

As all terms in the sum above are equal, we have

$$\llbracket \chi_i \rrbracket = \begin{cases} \omega & \text{if } \chi_i \in \sigma, \\ \emptyset & \text{if } \chi_i \notin \sigma. \end{cases}$$

Let ϑ''_τ be the formula obtained from $\vartheta'_\tau(\bar{Z})$ by replacing every variable Z_i by

$$\begin{cases} \text{true} & \text{if } \chi_i \in \sigma, \\ \text{false} & \text{if } \chi_i \notin \sigma. \end{cases}$$

Then it follows that

$$\sigma^\omega = \tau \quad \text{iff} \quad \langle \omega, \leq \rangle \models \vartheta''_\tau \quad \text{iff} \quad \vartheta''_\tau \in \text{Th}_{\text{MSO}_0}^{\bar{r}}(\omega, \leq). \quad \square$$

The key argument in our decidability proof below is the following lemma, which states that every labelling of ω is equivalent to an ultimately periodic one. It is a direct application of the Theorem of Ramsey.

Lemma 2.5. *Let $\bar{n} \in \omega^*$ and $m < \omega$. Then*

$$\{ \text{Th}_{\text{MSO}_0}^{\bar{n}}(\omega, \leq, \bar{P}) \mid P_0, \dots, P_{m-1} \subseteq \omega \} = \{ \sigma \tau^\omega \mid \sigma, \tau \in \Phi_{\bar{n}}(m) \}.$$

Proof. (\supseteq) Given $\sigma, \tau \in \Phi_{\bar{n}}(m)$, fix finite linear orders \mathcal{A} and \mathcal{B} and parameters \bar{P} and \bar{Q} with theories σ and τ , respectively. Then

$$\mathcal{A}, \bar{P} + \sum_{i < \omega} \mathcal{B}, \bar{Q} \cong \langle \omega, \leq, \bar{S} \rangle, \quad \text{for some } S_0, \dots, S_{m-1} \subseteq \omega.$$

Consequently,

$$\sigma \tau^\omega = \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{A}, \bar{P}) \cdot \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{B}, \bar{Q})^\omega = \text{Th}_{\text{MSO}_0}^{\bar{n}}(\omega, \leq, \bar{S}).$$

(\subseteq) Let $P_0, \dots, P_{m-1} \subseteq \omega$. For $i < k < \omega$, we define

$$\mathcal{A}_{i,k} := \langle \{i, \dots, k-1\}, \leq, \bar{P} \upharpoonright \{i, \dots, k-1\} \rangle.$$

By the Theorem of Ramsey, there exist a theory $\tau \in \Phi_{\bar{n}}(m)$ and an infinite sequence $k_0 < k_1 < \dots$ of positions such that

$$\text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{A}_{k_i, k_j}) = \tau, \quad \text{for all } i < j < \omega.$$

Setting $\sigma := \text{Th}_{\text{MSO}_0}^{\bar{n}}(\mathcal{A}_{0, k_0})$, it follows that

$$\text{Th}_{\text{MSO}_0}^{\bar{n}}(\omega, \leq, \bar{P}) = \text{Th}_{\text{MSO}_0}^{\bar{n}}\left(\mathcal{A}_{0, k_0} + \sum_{i < \omega} \mathcal{A}_{k_i, k_{i+1}}\right) = \sigma \tau^\omega. \quad \square$$

Theorem 2.6 (Büchi). $\text{Th}_{\text{MSO}}(\omega, \leq)$ is decidable.

Proof. We prove by induction on $|\bar{n}|$ that, given $\bar{n} \in \omega^*$, we can compute $\text{Th}_{\text{MSO}_0}^{\bar{n}}(\omega, \leq)$. For $\bar{n} = \langle \rangle$, we have

$$\text{Th}_{\text{MSO}_0}^{\bar{n}}(\omega, \leq) = \emptyset.$$

Hence, suppose that $\bar{n} = m\bar{n}'$ and that we already know how to compute $\text{Th}_{\text{MSO}_0}^{\bar{n}'}(\omega, \leq)$, for all $\bar{r} \in \omega^*$ with $|\bar{r}| = |\bar{n}'|$.

To compute $\text{Th}_{\text{MSO}_0}^{m\bar{n}'}(\omega, \leq)$ it is sufficient to decide whether or not

$$\langle \omega, \leq \rangle \models \exists X_0 \dots \exists X_{m-1} \psi,$$

for all $\text{MSO}_0^{\bar{n}}$ -formulae $\psi(\bar{X})$. Hence, given an $\text{MSO}_0^{\bar{n}}$ -formula $\psi(\bar{X})$, we have to decide whether there are sets $P_0, \dots, P_{m-1} \subseteq \omega$ such that

$$\langle \omega, \leq \rangle \models \psi(\bar{P}).$$

By Lemma 2.5, this is equivalent to the question of whether there are theories $\sigma, \tau \in \Phi_{\bar{n}}(m)$ such that

$$\psi(\bar{X}) \in \sigma\tau^\omega.$$

Therefore, it is sufficient to compute $\sigma\tau^\omega$, for all of the finitely many possible choices of σ and τ . This we can do with the help of Proposition 2.3 and Lemma 2.4 since, by inductive hypothesis, we can compute $\text{Th}_{\text{MSO}_0}^{\bar{n}'}(\omega, \leq)$, for all $\bar{r} \in \omega^*$ with $|\bar{r}| = |\bar{n}'|$. \square

3 ω -Semigroups

To study languages of infinite words, we extend the notion of a semigroup by adding an infinite product.

Definition 3.1. (a) An ω -semigroup is a two-sorted structure $\mathfrak{S} = \langle S, S_\omega \rangle$ with three products

$$\cdot : S \times S \rightarrow S, \quad \cdot : S \times S_\omega \rightarrow S_\omega, \quad \text{and} \quad \pi : S^\omega \rightarrow S_\omega$$

that satisfy the following associative laws:

$$\begin{aligned} (ab)c &= a(bc), \\ (ab)u &= a(bu), \\ b \cdot \pi(a_0, a_1, \dots) &= \pi(b, a_0, a_1, \dots), \\ \pi(a_0, a_1, \dots) &= \pi((a_0 \cdots a_{k_0-1}), (a_{k_0} \cdots a_{k_1-1}), \dots) \end{aligned}$$

for all $a, b, c, a_0, a_1, \dots \in S$ and $u \in S_\omega$, and all increasing sequences $0 < k_0 < k_1 < \dots < \omega$. Informally, we refer to the element of S as the *finite elements* and to those of S_ω as the *infinite elements* of \mathfrak{S} .

(b) The ω -power of an element $a \in S$ is

$$a^\omega := \pi(a, a, a, \dots).$$

(c) A homomorphism $\eta : \mathfrak{S} \rightarrow \mathfrak{T}$ of ω -semigroups consists of two maps

$$\eta_0 : S \rightarrow T \quad \text{and} \quad \eta_\omega : S_\omega \rightarrow T_\omega$$

that commute with products, i.e., for $a, b, a_0, a_1, \dots \in S$ and $u \in S_\omega$,

$$\begin{aligned} \eta_0(a) \cdot \eta_0(b) &= \eta_0(ab), \\ \eta_0(a) \cdot \eta_\omega(u) &= \eta_\omega(au), \\ \pi(\eta_0(a_0), \eta_0(a_1), \dots) &= \eta_\omega(\pi(a_0, a_1, \dots)). \end{aligned}$$

Definition 3.2. Let Σ be a set.

(a) The *free ω -semigroup* over Σ is $\langle \Sigma^+, \Sigma^\omega \rangle$. By abuse of notation we also denote it simply by Σ^ω .

(b) A language $L \subseteq \Sigma^\omega$ is *recognised* by a homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ to an ω -semigroup \mathfrak{S} if there exists a set $P \subseteq S_\omega$ such that $L = \eta^{-1}[P]$. \dashv

Example. Let $\mathfrak{S} = \langle S, S_\omega \rangle$ be the ω -semigroup with $S := \{0, 1\}$ and $S_\omega := \{0, 1\}$ where

$$\begin{aligned} a \cdot b &:= \max\{a, b\}, & \text{for } a, b \in S, \\ a \cdot u &:= u, & \text{for } a \in S, u \in S_\omega, \\ \pi(a_0, a_1, \dots) &:= \limsup_{n \rightarrow \infty} a_n, & \text{for } a_0, a_1, \dots \in S. \end{aligned}$$

The language L of all ω -words w containing infinitely many letters a is recognised by the morphism sending the letter a to $\mathbf{1}$ and every other letter to $\mathbf{0}$. \lrcorner

Example. In the previous section we have already introduced the ω -semigroup $\langle \Theta_{\bar{n}}(m), \Theta_{\bar{n}}(m) \rangle$ of all $\text{MSO}_{\mathbf{0}}^{\bar{n}}$ -theories over the signature $\{\leq, P_0, \dots, P_{m-1}\}$. Note that the function $\eta : \Sigma^\omega \rightarrow \Theta_{\bar{n}}(m)$ mapping a word $u \in \Sigma^\omega$ to its theory, is a homomorphism.

If $L \subseteq \Sigma^\omega$ is a language defined by an $\text{MSO}_{\mathbf{0}}^{\bar{n}}$ -formula φ , then

$$L = \{ w \in \Sigma^\omega \mid \varphi \in \text{Th}_{\text{MSO}_{\mathbf{0}}}^{\bar{n}}(w) \} = \eta^{-1}[P]$$

where $P := \{ \theta \in \Theta_{\bar{n}}(m) \mid \varphi \in \theta \}$. Thus, every definable language is recognised by η . \lrcorner

Exercise 3.1. Find homomorphisms into finite ω -semigroups that recognise the following languages over the alphabet $\{a, b, c\}$.

- (a) The language of all ω -words containing infinitely many a , but only finitely many b .
- (b) The language of all ω -words where immediately after or immediately before every letter a there is another a .
- (c) The language of all ω -words containing an even (and finite) number of a .
- (d) The language of all ω -words where after every letter a there is a later position with a letter b .
- (e) The language of all ω -words where, for every prefix p , the numbers of the letters a, b, c differ by at most $\mathbf{1}$. \lrcorner

Exercise 3.2. Prove that a language $L \subseteq \Sigma^\omega$ is recognised by homomorphisms into a finite ω -semigroup if, and only if, it is of the form

$$L = \bigcup_{i < m} U_i V_i^\omega,$$

where $m < \omega$, $U_i, V_i \subseteq \Sigma^+$ are MSO -definable languages of finite words, and $V^\omega := \{ v_0 v_1 v_2 \dots \mid v_i \in V \}$. \lrcorner

Exercise 3.3. Let \mathfrak{S} be a finite semigroup, $\eta : \Sigma^+ \rightarrow \mathfrak{S}$ a homomorphism, and let $L \subseteq \Sigma^\omega$ be a language of the form

$$L = \bigcup_{i < n} \eta^{-1}(b_i) (\eta^{-1}(e_i))^\omega, \quad \text{for } n < \omega \text{ and } b_i, e_i \in S,$$

where $X^\omega := \{ x_0 x_1 x_2 \dots \mid x_i \in X \}$.

Prove that the complement $\Sigma^\omega \setminus L$ is also of the form

$$\Sigma^\omega \setminus L = \bigcup_{i < m} \eta^{-1}(c_i) (\eta^{-1}(f_i))^\omega,$$

for suitable $m < \omega$ and $c_i, f_i \in S$. \lrcorner

Exercise 3.4. Prove that the free ω -semigroup $\langle \Sigma^+, \Sigma^\omega \rangle$ really is free: show that, for every ω -semigroup \mathfrak{S} and every map $h : \Sigma \rightarrow S$, there exists a unique homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ such that $\eta(a) = h(a)$, for all $a \in \Sigma$. \lrcorner

Exercise 3.5. Let $\mathfrak{A} = \langle S, S_\omega \rangle$ be an ω -semigroup where each element is invertible, i.e., for every $a \in S$, there is some element $a^{-1} \in S$ such that $aa^{-1} = \mathbf{1} = a^{-1}a$ (for some fixed $\mathbf{1} \in S$). Prove that

$$\pi(a_0, a_1, \dots) = \pi(b_0, b_1, \dots), \quad \text{for all } a_i, b_i \in S. \quad \lrcorner$$

In order to prove that a language is recognisable precisely when it is MSO -definable, we employ the Theorem of Ramsey.

Definition 3.3. (a) Let \mathfrak{S} be a semigroup and $(a_n)_{n < \omega}$ a sequence of semigroup elements. For $i < k < \omega$, we write

$$a_{[i,k)} := a_i a_{i+1} \dots a_{k-1}.$$

A *Ramsey factorisation* of $(a_n)_{n < \omega}$ is a sequence of indices $0 < k_0 < k_1 < \dots < \omega$ such that

$$a_{[k_i, k_j)} = a_{[k_{i'}, k_{j'})}, \quad \text{for all } i < j \text{ and } i' < j'.$$

The type of such a factorisation is the pair

$$\langle a_{[0, k_1]}, a_{[k_1, k_2]} \rangle.$$

(b) Let \mathfrak{S} be an ω -semigroup and $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ a homomorphism. A *Ramsey factorisation* of a word $w \in \Sigma^\omega$ is a Ramsey factorisation of the sequence $(\eta(w(n)))_{n < \omega}$.

Lemma 3.4. *Let \mathfrak{S} be a finite semigroup and $(a_n)_{n < \omega}$ a sequence of semigroup elements.*

(a) $(a_n)_{n < \omega}$ has a Ramsey factorisation.

(b) If $\langle b, e \rangle$ is the type of a Ramsey factorisation of $(a_n)_{n < \omega}$ then

$$be = b \quad \text{and} \quad ee = e.$$

(c) If $\langle b, e \rangle$ and $\langle c, f \rangle$ are the types of two Ramsey factorisations of $(a_n)_{n < \omega}$ then there are elements $u, v \in S^1$ such that

$$c = bu, \quad e = uv, \quad \text{and} \quad f = vu.$$

Proof. (a) We define a colouring $\lambda : [\omega]^2 \rightarrow S$ by assigning to a pair $i < j$ of indices the colour $\lambda(i, j) := a_{[i, j]}$. By the Theorem of Ramsey, there exists an infinite set $I \subseteq \omega$ such that $\lambda(i, j) = \lambda(i', j')$, for all $i < j$ and $i' < j'$ in I . We can choose for $k_0 < k_1 < \dots$ an increasing enumeration of I .

(b) Let $k_0 < k_1 < \dots$ be a Ramsey factorisation with type $\langle b, e \rangle$. Set $c := a_{[0, k_0]}$ and $d_n := a_{[k_n, k_{n+1}]}$, for $n < \omega$. Then $b = cd_0$ and $e = d_1$. Furthermore,

$$d_i d_{i+1} \dots d_j = a_{[k_i, k_{j+1}]} = a_{[k_{i'}, k_{j'+1}]} = d_{i'} d_{i'+1} \dots d_{j'},$$

for all $i \leq j$ and $i' \leq j'$. In particular, $d_i = d_j$, for all i, j . Consequently, we have

$$be = cd_0 d_1 = cd_0 = b \quad \text{and} \quad ee = d_1 d_1 = d_1 d_2 = d_1 = e.$$

(c) Let $k_0 < k_1 < \dots$ and $l_0 < l_1 < \dots$ be Ramsey factorisations with types $\langle b, e \rangle$ and $\langle c, f \rangle$, respectively. Replacing $(k_n)_{n < \omega}$ and $(l_n)_{n < \omega}$ by suitable

subsequences we may assume without loss of generality that $k_0 \leq l_0 \leq k_1 \leq l_1 \leq \dots$. For $n < \omega$, set

$$u_n := a_{[k_n, l_n]} \quad \text{and} \quad v_n := a_{[l_n, k_{n+1}]}.$$

Then $c = bu_1$ and

$$e = u_0 v_0 = u_1 v_1 = \dots \quad \text{and} \quad f = v_0 u_1 = v_1 u_2 = \dots.$$

Since the number of possible pairs $\langle u_n, v_n \rangle$ is finite, there exist elements $u, v \in S$ and an infinite set $I \subseteq \omega$ such that

$$\langle u_n, v_n \rangle = \langle u, v \rangle, \quad \text{for all } n \in I.$$

Fix elements $m, n \in I$ with $n > m + 1$. Then

$$\begin{aligned} c &= cf^m = (bu_0)(v_0 u_1) \dots (v_{m-1} u_m) \\ &= b(u_0 v_0) \dots (u_{m-1} v_{m-1}) u_m = be^m u = bu, \end{aligned}$$

$$e = e^2 = u_m v_m e = uve,$$

$$\begin{aligned} f &= f^{n-m} = (v_m u_{m+1}) \dots (v_{n-1} u_n) \\ &= v_m (u_{m+1} v_{m+1}) \dots (u_{n-1} v_{n-1}) u_n = ve^{n-m-1} u = veu. \end{aligned}$$

Consequently, the elements u and ve have the desired properties. \square

Lemma 3.5. *Let $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ be a homomorphism into a finite ω -semigroup \mathfrak{S} . For $u \in S_\omega$, set*

$$F_u := \{ \langle a, b \rangle \in S^2 \mid ab^\omega = u \}.$$

Then $\eta(w) = u$ if, and only if, w has a Ramsey factorisation with type in F_u .

Proof. (\Leftarrow) If w has a Ramsey factorisation $k_0 < k_1 < \dots$ of type $\langle a, b \rangle \in F_u$ then

$$\begin{aligned} \eta(w) &= \eta(w[0, k_1] w[k_1, k_2] w[k_2, k_3] \dots) \\ &= \eta(w[0, k_1]) \cdot \eta(w[k_1, k_2]) \cdot \eta(w[k_2, k_3]) \dots \\ &= abb \dots = ab^\omega = u. \end{aligned}$$

(\Rightarrow) Suppose that $\eta(w) = u$. By Lemma 3.4 (a), w has a Ramsey factorisation $k_0 < k_1 < \dots$. Let $\langle b, e \rangle$ be its type. Then

$$\begin{aligned} u &= \eta(w) = \eta(w[0, k_1)) \cdot \eta(w[k_1, k_2)) \cdot \eta(w[k_2, k_3)) \cdots \\ &= bee \cdots = be^\omega. \end{aligned}$$

Hence, $\langle b, e \rangle \in F_u$. \square

Lemma 3.6. *Let \mathfrak{S} be a finite ω -semigroup. For every pair $\langle b, e \rangle \in S^2$, there exists an MSO-formula $\varphi_{b,e}$ defining the set of all ω -words with a Ramsey factorisation of type $\langle b, e \rangle$.*

Proof. Let $c \in S$. We start by defining a formula $\psi_c(x, y)$ stating that the factor from position x to $y - 1$ is mapped to c . We use set variables $(Z_d)_{d \in S}$ containing all positions $x \leq z < y$ such that the factor from x to z is mapped to d . The formula $\psi_c(x, y)$ states that there are sets $(Z_d)_{d \in S}$ such that

- ♦ $Z_d \cap Z_{d'} = \emptyset$ for $d \neq d'$,
- ♦ if a is the letter at position x , then $x \in Z_{\eta(a)}$,
- ♦ if $x < z < y$ and the letter at position z is a , then $z - 1 \in Z_d$ implies $z \in Z_{d\eta(a)}$, and
- ♦ $y - 1 \in Z_c$.

Clearly, each of these statements can be expressed in MSO.

Having defined the formulae ψ_c , the desired formula $\varphi_{b,e}$ states that there exists an infinite set Z such that

- ♦ $\psi_b(0, x)$ holds, where x is the second element of Z , and
- ♦ $\psi_e(x, y)$ holds for all elements $x < y$ of Z . \square

Theorem 3.7. *Let Σ be a finite alphabet. A language $L \subseteq \Sigma^\omega$ is MSO-definable if, and only if, there exists a homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ into a finite ω -semigroup \mathfrak{S} recognising L .*

Proof. (\Rightarrow) We have seen in the example after Definition 3.2 that every $\text{MSO}_0^{\bar{n}}$ -definable language is recognised by a homomorphism into the ω -semigroup $\Theta_{\bar{n}}(m)$ of all $\text{MSO}_0^{\bar{n}}$ -theories.

(\Leftarrow) Let $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ be a morphism such that $L = \eta^{-1}[P]$, for some $P \subseteq S_\omega$. We have seen in Lemma 3.5 that

$$\eta(w) = u \quad \text{iff} \quad w \text{ has a Ramsey factorisation with type in } F_u.$$

Consequently, the formula

$$\psi := \bigvee_{u \in P} \bigvee_{\langle b, e \rangle \in F_u} \varphi_{b,e}$$

defines L , where $\varphi_{b,e}$ are the formulae from Lemma 3.6. \square

If we want to compute with an ω -semigroup, we face the problem that we cannot write down the multiplication table of the infinite product since it is infinite. For algorithmic applications we need to represent this table in a finite way.

Definition 3.8. A *Wilke algebra* is a structure $\langle S, S_\omega, \cdot, \cdot^\omega \rangle$ with two products

$$\cdot : S \times S \rightarrow S \quad \text{and} \quad \cdot : S \times S_\omega \rightarrow S_\omega$$

and one unary ω -power operation

$$^\omega : S \rightarrow S_\omega.$$

These operations satisfy the following associative laws:

$$\begin{aligned} (ab)c &= a(bc), & (ab)^\omega &= a(ba)^\omega, \\ a(bu) &= (ab)u, & (a^n)^\omega &= a^\omega, \end{aligned}$$

for $a, b, c \in S, u \in S_\omega$, and $0 < n < \omega$. \square

Theorem 3.9. *Let \mathfrak{S} be a finite Wilke algebra. There exists a unique function $\pi : S^\omega \rightarrow S_\omega$ turning \mathfrak{S} into an ω -semigroup with*

$$\pi(a, a, a, \dots) = a^\omega.$$

Proof. Given a sequence $(a_n)_{n < \omega}$ we define

$$\pi(a_0, a_1, \dots) := be^\omega,$$

where $\langle b, e \rangle$ is the type of a Ramsey factorisation of $(a_n)_{n < \omega}$. To see that this is well-defined, first note that, according to Lemma 3.4 (a), every sequence does have a Ramsey factorisation. Furthermore, if $\langle b, e \rangle$ and $\langle c, f \rangle$ are the types of two Ramsey factorisations of $(a_n)_{n < \omega}$ then we can use Lemma 3.4 (c) to find elements $u, v \in S^1$ such that

$$c = bu, \quad e = uv, \quad \text{and} \quad f = vu.$$

Hence,

$$cf^\omega = bu(vu)^\omega = b(uv)^\omega = be^\omega.$$

To prove that this operation turns \mathfrak{S} into an ω -semigroup, we have to show associativity. For the first equation, let $(a_n)_{n < \omega}$ be a sequence of semigroup elements and let $c \in S$. If $k_0 < k_1 < \dots$ is a Ramsey factorisation of $(a_n)_{n < \omega}$, then $k_0 + 1 < k_1 + 1 < \dots$ is a Ramsey factorisation of the sequence c, a_0, a_1, \dots and we have

$$\begin{aligned} c \cdot \pi(a_0, a_1, \dots) &= c(a_0 \cdots a_{k_1-1})(a_{k_1} \cdots a_{k_2-1})^\omega \\ &= (ca_0 \cdots a_{k_1-1})(a_{k_1} \cdots a_{k_2-1})^\omega \\ &= \pi(c, a_0, a_1, \dots). \end{aligned}$$

For the second equation, let $(a_n)_{n < \omega}$ be a sequence of semigroup elements and let $l_0 < l_1 < \dots < \omega$ be a sequence of indices. Suppose that $k_0 < k_1 < \dots$ is a Ramsey factorisation of the sequence $(a_{l_n} \cdots a_{l_{n+1}-1})_{n < \omega}$. Then $l_{k_0} < l_{k_1} < \dots$ is a Ramsey factorisation of $(a_n)_{n < \omega}$ and we have

$$\begin{aligned} \pi(a_0 \cdots a_{l_0-1}, a_{l_0} \cdots a_{l_1-1}, \dots) &= (a_0 \cdots a_{l_{k_1}-1})(a_{l_{k_1}} \cdots a_{l_{k_2}-1})^\omega \\ &= \pi(a_0, a_1, \dots). \end{aligned}$$

It remains to show that the product π is unique. Suppose that $\pi' : S^\omega \rightarrow S_\omega$ is any associative operation such that $\pi'(a, a, a, \dots) = a^\omega$, for all $a \in S$.

To prove that $\pi' = \pi$, consider a sequence $(a_n)_{n < \omega}$ in S and let $k_0 < k_1 < \dots$ be a Ramsey factorisation of $(a_n)_{n < \omega}$ of type $\langle b, e \rangle$. Then it follows that

$$\begin{aligned} \pi'(a_0, a_1, \dots) &= \pi'(a_{[0, k_1)}, a_{[k_1, k_2)}, a_{[k_2, k_3)}, \dots) \\ &= \pi'(b, e, e, e, \dots) \\ &= b \cdot \pi'(e, e, e, \dots) = be^\omega = \pi(a_0, a_1, \dots). \quad \square \end{aligned}$$

Exercise 3.6. (a) Let L_0 and L_1 be languages that are recognised by homomorphisms $\eta_0 : \Sigma^\omega \rightarrow \mathfrak{S}_0$ and $\eta_1 : \Sigma^\omega \rightarrow \mathfrak{S}_1$ into finite ω -semigroups. Prove that the languages $L_0 \cap L_1$, $L_0 \cup L_1$, and $\Sigma^\omega \setminus L_0$ are also recognised by a homomorphism into some finite ω -semigroup.

(b) Let $L \subseteq \Sigma^\omega$ be recognised by a homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ into a finite ω -semigroup \mathfrak{S} and let $\pi : \Sigma \rightarrow \Gamma$ be a function. Prove that

$$\pi[L] := \{ \pi(w) \mid w \in L \}$$

is also recognised by a homomorphism into a finite ω -semigroup. (This exercise is a bit more involved.)

(c) Use (a) and (b) to give an alternative proof of the fact that every MSO-definable language is recognised by a homomorphism into a finite ω -semigroup. Show furthermore that, given a formula φ one can effectively construct a description of the corresponding homomorphism and Wilke algebra.

(d) Use (c) to give an alternative decidability proof for the monadic theory of $\langle \omega, \leq \rangle$. \square

Exercise 3.7. The *syntactic congruence* of a language $L \subseteq \Sigma^\omega$ is the relation

$$\begin{aligned} x \sim_L y &: \text{iff} \quad u(xv)^\omega \in L \Leftrightarrow u(yv)^\omega \in L \\ &\text{and} \quad ux \in L \Leftrightarrow uy \in L, \quad \text{for all } u, v \in \Sigma^*. \end{aligned}$$

(a) Prove that the syntactic congruence is a congruence of the free ω -semigroup.

(b) Prove that a language $L \subseteq \Sigma^\omega$ is MSO-definable if, and only if, \sim_L has only finitely many classes and L is a union of languages of the form $K_0 K_1 K_2 \dots$, where each K_i is an \sim_L -class.

(c) Show that $\sim_L = \sim_K$ where $L := (\mathbf{o}^* \mathbf{1})^\omega$ and

$$K := \{ \mathbf{o}^{n_0} \mathbf{1} \mathbf{o}^{n_1} \mathbf{1} \mathbf{o}^{n_2} \mathbf{1} \dots \mid n_0, n_1, n_2, \dots \text{ is unbounded} \}.$$

4 ω -Automata

As usual, instead of algebras we can also use automata to recognise languages. For ω -words, we have to modify our notion of acceptance since, when reading an infinite word, an automaton cannot simply use the final state to decide whether or not to accept its input. A workable alternative turns out to be to look at the states that appear infinitely often during the run of the automaton. One common choice is to decide acceptance based on the set of all these states. A simpler, better behaved, but equivalent alternative is to fix an ordering of the states and just use the least state appearing infinitely often. We will adopt this simpler method, although slightly reformulated for technical reasons. It is easy to check that the following definition is equivalent to using an ordering on the states.

Definition 4.1. (a) An (nondeterministic) ω -automaton is a tuple $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ where where Q is a finite set of states, Σ is a finite input alphabet, $q_0 \in Q$ is the initial state, $\Omega : Q \rightarrow \omega$ is a priority function, and $\Delta \subseteq Q \times \Sigma \times Q$ is the transition relation. Instead of $\langle p, a, q \rangle \in \Delta$, we also write $p \xrightarrow{a} q$.

(b) A run of an ω -automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ on an ω -word $w \in \Sigma^\omega$ is an ω -word $\rho \in Q^\omega$ such that

$$\langle \rho(n), w(n), \rho(n+1) \rangle \in \Delta, \quad \text{for all } n < \omega.$$

A run ρ is accepting if $\rho(0) = q_0$ and ρ satisfies the parity condition

$$\liminf_{n \rightarrow \infty} \Omega(\rho(n)) \text{ is even.}$$


(c) An ω -automaton \mathcal{A} accepts an ω -word w if there exists an accepting run of \mathcal{A} on w . The language recognised by \mathcal{A} is the set $L(\mathcal{A})$ of all ω -words it accepts.

(d) An ω -automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ is deterministic if, for all states $q \in Q$ and all letters $a \in \Sigma$, there is a unique state $q' \in Q$ with $\langle q, a, q' \rangle \in \Delta$. In this case, we can replace the transition relation Δ by a transition function $\delta : Q \times \Sigma \rightarrow Q$ such that


$$\Delta = \{ \langle q, a, \delta(q, a) \rangle \mid q \in Q, a \in \Sigma \}.$$

(e) A Büchi automaton is an ω -automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ where $\text{rng } \Omega \subseteq \{0, 1\}$.

Example. The language of all words $w \in \{a, b\}^\omega$ with infinitely many a is recognised by the ω -automaton



where the numbers labelling the states indicate their priority. We obtain an ω -automaton for the language of all words $w \in \{a, b\}^\omega$ with only finitely many a by changing the priorities:



Exercise 4.1. Find ω -automata recognising the following languages over the alphabet $\{a, b, c\}$.

- The language of all ω -words containing infinitely many a , but only finitely many b .
- The language of all ω -words where immediately after or immediately before every letter a there is another a .
- The language of all ω -words containing an even (and finite) number of a .

- (d) The language of all ω -words where after every letter a there is a later position with a letter b .
- (e) The language of all ω -words where, for every prefix p , the numbers of the letters a, b, c differ by at most 1. \downarrow

Exercise 4.2. (a) Prove that, for every ω -automaton \mathcal{A} , there exists a Büchi automaton \mathcal{B} recognising the same language.

(b) Prove that a language $L \subseteq \Sigma^\omega$ is recognised by a Büchi automaton if, and only if, it is of the form

$$L = \bigcup_{i < m} U_i V_i^\omega,$$

where $m < \omega$ and $U_i, V_i \subseteq \Sigma^+$ are MSO-definable languages of finite words.

(c) Find a Büchi automaton recognising the language of all ω -words with only finitely many letters a . Prove that this language is not recognised by a deterministic Büchi automaton, i.e., one where the transition relation Δ is the graph of a function $Q \times \Sigma \rightarrow Q$.

(d) Prove that the class of languages recognised by Büchi automata is closed under union, intersection, complement, and projection.

(e) Prove that a language is recognised by a Büchi automaton if, and only if, it is MSO-definable. \downarrow

Example. Let $F \subseteq \Sigma$. We will construct a (deterministic) automaton \mathcal{A} recognising all words $w \in \Sigma^\omega$ such that F is the set of letters appearing infinitely often in w . The idea is as follows. The automaton maintains a record of all letters it has seen and their ordering. More precisely it stores a permutation $c_0 \dots c_{n-1}$ of Σ such that letters on the right have been seen more frequently than those on the left. Everytime the automaton reads a new input letter a this letter is removed from its position in the permutation and moved to the end. In addition, the automaton remembers what the old position was, by adding a special marker $|$. Hence, the states of \mathcal{A} are permutations of $\Sigma \cup \{| \}$, which we will call *latest appearance records*. For instance, when reading the word

$a \ a \ b \ a \ b \ c \ c \ b \ a \ a \ b \dots$

the run of the automaton looks as follows. (The initial state does not matter.)

$$\begin{aligned} abc| &\rightarrow^a |bca \rightarrow^a bc|a \rightarrow^b |cab \rightarrow^a c|ba \rightarrow^b c|ab \rightarrow^c |abc \rightarrow^c ab|c \\ &\rightarrow^b a|cb \rightarrow^a |cba \rightarrow^a cb|a \rightarrow^b c|ab \dots \end{aligned}$$

We assign the priority 2 to all states of the form $c_0 \dots c_{k-1}|c_k \dots c_{n-1}$ where $\{c_k, \dots, c_{n-1}\} = F$, and priority 1 to all other states.

We claim that \mathcal{A} accepts w if, and only if, the set of letters appearing infinitely often in w is equal to F .

(\Leftarrow) We can factorise the input word as $w = u_0 u_1 u_2 \dots$ such that u_0 contains all letters that appear only finitely often in w and each u_i with $i > 0$ contain every letter in F at least once. After reading the prefix $u_0 u_1$, the automaton \mathcal{A} will only see states of the form $c_0 \dots c_{k-1}|c_k \dots c_{n-1}$ with $c_k, \dots, c_{n-1} \in F$. Furthermore, in each factor u_i , $i > 1$, there will be at least one state of the form $c_0 \dots c_{k-1}|c_k \dots c_{n-1}$ with $\{c_k, \dots, c_{n-1}\} = F$.

(\Rightarrow) Fix a state $q := c_0 \dots c_{k-1}|c_k \dots c_{n-1}$ that appears infinitely often in the run of \mathcal{A} on w , and let $w = u_0 u_1 \dots$ be the factorisation of w such that q occurs after each factor u_i . As \mathcal{A} accepts w , we have $\{c_k, \dots, c_{n-1}\} = F$. Furthermore, every letter c_j with $k \leq j < n$ must occur somewhere in u_i , for $i > 0$. Hence, every $c \in F$ occurs infinitely often in w . For a contradiction, suppose that there is some other letter $c \in \Sigma \setminus F$ that also occurs infinitely often. Then some state p occurs infinitely often in the run where c appears on the right of $|$. But such states have priority 1, which means the run cannot be accepting. \downarrow

Similarly to automata for finite words, we can associate a transition semigroup with an ω -automaton.

Definition 4.2. Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ be an ω -automaton and let $D := \text{rng } \Omega$ be the set of priorities used. The *transition ω -semigroup* $\mathfrak{S}_{\mathcal{A}} := \langle S, S_\omega \rangle$ of \mathcal{A} has domains

$$S := \mathcal{P}(Q \times D \times Q) \quad \text{and} \quad S_\omega := \mathcal{P}(Q).$$

Binary multiplication is defined by

$$\begin{aligned} A \cdot B &:= \{ \langle p, \min \{d, d'\}, r \rangle \mid \langle p, d, q \rangle \in A, \langle q, d', r \rangle \in B \}, \\ A \cdot U &:= \{ p \mid \langle p, d, q \rangle \in A, q \in U \}, \end{aligned}$$

for $A, B \in S$ and $U \in S_\omega$. The infinite product is given by

$$\pi(A_0, A_1, \dots) := \{ p_0 \mid \text{there are } \langle p_n, d_n, p_{n+1} \rangle \in A_n, \text{ for } n < \omega, \\ \text{such that } \liminf_{n \rightarrow \infty} d_n \text{ is even} \}.$$

Theorem 4.3. *Let Σ be a finite alphabet and $L \subseteq \Sigma^\omega$ a language of ω -words. The following statements are equivalent:*

- (1) *L is recognised by an ω -automaton.*
- (2) *L is recognised by a Büchi automaton.*
- (3) *L is recognised by a homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ into a finite ω -semigroup.*

Proof. (2) \Rightarrow (1) is trivial.

(1) \Rightarrow (3) Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ be an ω -automaton recognising L and let $\mathfrak{S}_{\mathcal{A}}$ be its transition ω -semigroup. We define a homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}_{\mathcal{A}}$ by mapping finite words $w \in \Sigma^*$ to the set of all triples $\langle p, d, q \rangle$ such that there exists a run of \mathcal{A} on w starting in state p , ending in state q , and having the minimal priority d . Similarly, infinite words $w \in \Sigma^\omega$ are mapped to the set of all states p such that there exists a run of \mathcal{A} on w starting in state p and satisfying the parity condition. Then

$$L = \eta^{-1}[P] \quad \text{where} \quad P := \{ U \subseteq Q \mid q_0 \in U \}.$$

Hence, η recognises L .

(3) \Rightarrow (2) Suppose that $L = \eta^{-1}[P]$ for some $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$. We have seen in Lemma 3.5 that

$$\eta(w) = u \quad \text{iff} \quad w \text{ has a Ramsey factorisation with type in } F_u.$$

Consequently, we can construct a Büchi automaton that, on input w , guesses a value $u \in P$ and a type $\langle b, e \rangle \in F_u$ and then checks that w has a Ramsey

factorisation with type $\langle b, e \rangle$. This can be done as follows. After reading a prefix v the automaton remembers the image $\eta(v)$. The automaton can do this since, if the current value is $\eta(v)$ and the next letter is c , the next value will be $\eta(vc) = \eta(v)\eta(c)$. Hence, when reading a letter c it only needs to multiply the current value by $\eta(c)$. If the current value is equal to b , the automaton can nondeterministically decide that it has read the first factor of the factorisation. In this case, it resets the stored value and reads letters until it reaches the value e . After having found a factor with value e , the automaton can again nondeterministically decide that it has found the next factor of the factorisation. It resets its stored value and reads the next factor. The automaton accepts if this reset was performed infinitely many times. Formally, we have states

$$Q := \{0, 1, 2\} \times S \cup \{q_0\}$$

and the following transitions, for $a \in S$ and $c \in \Sigma$,

$$\begin{aligned} q_0 &\xrightarrow{c} \langle 0, \eta(c) \rangle, \\ \langle 0, a \rangle &\xrightarrow{c} \langle 0, a\eta(c) \rangle, & \langle 0, b \rangle &\xrightarrow{c} \langle 1, \eta(c) \rangle, \\ \langle 1, a \rangle &\xrightarrow{c} \langle 2, a\eta(c) \rangle, & \langle 1, e \rangle &\xrightarrow{c} \langle 1, \eta(c) \rangle, \\ \langle 2, a \rangle &\xrightarrow{c} \langle 2, a\eta(c) \rangle, & \langle 2, e \rangle &\xrightarrow{c} \langle 1, \eta(c) \rangle. \end{aligned}$$

The initial state is q_0 and the priority function is

$$\Omega(q_0) := 1 \quad \text{and} \quad \Omega(\langle k, a \rangle) := \begin{cases} 0 & \text{if } k = 1, \\ 1 & \text{otherwise.} \end{cases} \quad \square$$

Our next aim is to show that every ω -automaton is equivalent to a deterministic one.

Lemma 4.4. *Let \mathfrak{S} be a finite ω -semigroup and let $e_0, e_1, \dots \in S$ be elements such that*

$$e_i e_k = e_i, \quad \text{for all } i, k < \omega.$$

Then $\pi(e_0, e_1, e_2, \dots) = e_0^\omega$.

Proof.

$$\begin{aligned}\pi(e_0, e_1, e_2, \dots) &= \pi(e_0, e_1 e_0, e_2 e_0, \dots) \\ &= \pi(e_0 e_1, e_0 e_2, e_0 e_3, \dots) \\ &= \pi(e_0, e_0, e_0, \dots) = e_0^\omega.\end{aligned}\quad \square$$

Theorem 4.5 (McNaughton). *For every homomorphism $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ into a finite ω -semigroup and every set $P \subseteq S_\omega$, we can construct a deterministic ω -automaton \mathcal{A} recognising $\eta^{-1}[P]$.*

Proof. Given an ω -word $w \in \Sigma^\omega$, we consider the colouring

$$\lambda(i, k) := \eta(w[i, k]), \quad \text{for } i < k < \omega.$$


This colouring is additive. Hence, we can use Theorem 1.10 to construct a deterministic finite-state transducer $\mathcal{T} = \langle Q, S, [N], q_0, \delta, \mu \rangle$ that, given λ , produces some weak Ramseyan split $\sigma : \omega \rightarrow [N]$ for \mathfrak{S} .

The idea of the construction is as follows. To compute $\eta(w)$ we find an infinite increasing sequence $z_0 \sqsubseteq_\sigma z_1 \sqsubseteq_\sigma \dots$. Then it follows by Lemma 4.4 that

$$\eta(w) = \lambda(o, z_0) \cdot \lambda(z_0, z_1)^\omega.$$

When trying to do this deterministically we face the problem that we do not know the value $k = \sigma(z_i)$. Therefore, the automaton has to do the above computation simultaneously for all possible values of k .

To make this idea precise, let us introduce some terminology. Given a position n in the input word, we call a position z *visible* (from n) if $o < z \leq n$ and there is no position $z < x \leq n$ with $\sigma(x) > \sigma(z)$. The *level* of a visible position z is the number $\sigma(z)$.



At each input position n , the sequence $z_0 \sqsubseteq_\sigma \dots \sqsubseteq_\sigma z_{l-1}$ of all visible positions at level k might be the start of the infinite factorisation we are looking for. Therefore, the automaton needs to remember the values $\lambda(o, z_0)$ and $\lambda(z_0, z_1)$. In order to update this information, it also needs to know the value $\lambda(z_1, n)$. Hence, for each $k < N$, the automaton will use three registers containing values from S^1 . Their values $\chi_k(n)$ at position n , can be defined as follows. Let $z_0 \sqsubseteq_\sigma \dots \sqsubseteq_\sigma z_{l-1}$ be an enumeration of all visible positions at level k . Since we have to support the cases where $l \leq 1$ or $z_1 = n$, we obtain

$$\chi_k(n) := \begin{cases} \langle 1, 1, \lambda(o, n) \rangle & \text{if } l = o, \\ \langle \lambda(o, z_0), 1, \lambda(z_0, n) \rangle & \text{if } l = 1, \\ \langle \lambda(o, z_0), \lambda(z_0, z_1), \lambda(z_1, n) \rangle & \text{if } l > 1 \end{cases}$$

(using the convention that $\lambda(x, x) = 1$).

The memory of the automaton will consist of the state of \mathcal{T} and the values $\chi_o(n), \dots, \chi_{N-1}(n)$. This is possible as we can compute $\chi_k(n+1)$ from $\chi_k(n)$, $\sigma(n+1)$ (which is provided by \mathcal{T}), and $\lambda(n, n+1) = \eta(w[n])$. To see this, note that

$$\chi_k(n) = \langle a_k, b_k, c_k \rangle \quad \text{and} \quad \lambda(n, n+1) = d,$$

implies that

$$\begin{aligned}\chi_k(n+1) &= \langle 1, 1, a_k b_k c_k d \rangle, & \text{for } k < \sigma(n+1), \\ \chi_k(n+1) &= \langle a_k, b_k, c_k d \rangle, & \text{for } k > \sigma(n+1), \\ \chi_k(n+1) &= \begin{cases} \langle c_k d, 1, 1 \rangle & \text{if } a_k = 1, b_k = 1, \\ \langle a_k, c_k d, 1 \rangle & \text{if } a_k \neq 1, b_k = 1, \\ \langle a_k, b_k, c_k d \rangle & \text{if } a_k \neq 1, b_k \neq 1, \end{cases} \\ & & \text{for } k = \sigma(n+1). \end{aligned}$$

After these preparations we are finally able to formally define the automaton \mathcal{A} we are looking for. Its states are

$$Q' := Q \times (S^1 \times S^1 \times S^1)^N.$$

After reading the first n letters of its input, \mathcal{A} is in the state

$$\langle q_n, \chi_0(n), \dots, \chi_{N-1}(n) \rangle,$$

where q_n is the state of \mathcal{T} . The initial state is

$$\langle q_0, \langle 1, 1, 1 \rangle, \dots, \langle 1, 1, 1 \rangle \rangle.$$

We assign to a state

$$p = \langle q, \langle a_0, b_0, c_0 \rangle, \dots, \langle a_{N-1}, b_{N-1}, c_{N-1} \rangle \rangle$$

with $\mu(q) = k$ the priority

$$\Omega(p) := \begin{cases} 2(N-k) & \text{if } a_k, b_k \neq 1 \text{ and } a_k b_k^\omega \in P, \\ 2(N-k) - 1 & \text{if } a_k = 1, b_k = 1, \text{ or } a_k b_k^\omega \notin P. \end{cases}$$

We claim that this automaton \mathcal{A} accepts an ω -word w if, and only if, $\eta(w) \in P$.

(\Leftarrow) Suppose that $\eta(w) \in P$. Let $k := \limsup_{n \rightarrow \infty} \sigma(n)$ and let $z_0 < z_1 < \dots$ be an increasing enumeration of all positions in $\sigma^{-1}(k)$. By choice of k , there is some index $l < \omega$ such that $\sigma(x) \leq k$, for all $x \geq z_l$. Set

$$a := \lambda(o, z_l) \quad \text{and} \quad e_i := \lambda(z_{l+i}, z_{l+i+1}), \quad \text{for } i < \omega.$$

At position z_{l+i} , the automaton is in a state of the form

$$\langle q, \dots, \underbrace{\langle a, e_0, c_k \rangle}_{k\text{-th component}}, \dots \rangle \quad \text{with} \quad \mu(q) = k.$$

Hence, the minimal priority seen infinitely often is either $2(N-k)$ or $2(N-k) - 1$ depending on whether or not $a e_0^\omega \in P$. As σ is a weak Ramseyan spilt, we have $e_i e_k = e_i$, for all i, k . Therefore, it follows by Lemma 4.4 that

$$a e_0^\omega = \pi(a, e_0, e_1, e_2, \dots) = \eta(w) \in P.$$

Hence, \mathcal{A} accepts w .

(\Rightarrow) Suppose that there exists an accepting run ρ of \mathcal{A} on w and let $2(N-k)$ be the minimal priority occurring infinitely often in it. Then ρ contains infinitely many states of the form

$$\langle q, \dots, \underbrace{\langle a_k, e_k, c_k \rangle}_{k\text{-th component}}, \dots \rangle \quad \text{where} \quad a_k e_k^\omega \in P \text{ and } \mu(q) = k.$$

Let $z_0 < z_1 < \dots$ be an enumeration of all positions with such a state. Since ρ does not contain infinitely many states with priority smaller than $2(N-k)$, it follows that there is some index $n < \omega$ such that $\sigma(x) \leq k$, for all $x \geq z_n$. Since $\sigma(z_i) = k$, for all i , we therefore have

$$z_i \approx_\sigma z_k, \quad \text{for } i, k \geq n.$$

Setting $a := \lambda(o, z_n)$ and $e_i := \lambda(z_{n+i}, z_{n+i+1})$ it follows that

$$\rho(z_{n+i}) = \langle q, \dots, \langle a, e_0, c_i \rangle, \dots \rangle, \quad \text{for some } c_i \in S,$$

and that $e_i e_k = e_i$, for all i, k . Hence, Lemma 4.4 implies that

$$\eta(w) = \pi(a, e_0, e_1, e_2, \dots) = a e_0^\omega \in P. \quad \square$$

The results of the previous sections are summarised in the following theorem. We also add one further logical characterisation.

Definition 4.6. *Weak monadic second-order logic* WMSO has the same syntax as MSO, but all set quantifiers range over finite sets only. \dashv

Theorem 4.7. *Let $L \subseteq \Sigma^\omega$ be a language of ω -words. The following statements are equivalent:*

- (1) L is definable in MSO.
- (2) L is definable in WMSO.
- (3) L is recognised by a homomorphism into a finite ω -semigroup.
- (4) L is recognised by a nondeterministic ω -automaton.

(5) L is recognised by a deterministic ω -automaton.

(6) L is recognised by a nondeterministic Büchi automaton.

Furthermore, all translations between these formalisms are effective.

Proof. The equivalences (1) \Leftrightarrow (3) \Leftrightarrow (4) \Leftrightarrow (6) were already proved in Theorems 3.7 and 4.3, respectively. (2) \Rightarrow (1) is trivial, and (3) \Rightarrow (5) follows from Theorem 4.5. Hence, it remains to prove (5) \Rightarrow (2).

Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ be a deterministic ω -automaton that recognises L . We start by constructing formulae $\text{STATE}_q(x)$, for $q \in Q$, stating that $\rho(x) = q$, where ρ is the unique run of \mathcal{A} on the input word. These formulae guess finite sets Z_p , $p \in Q$, containing all positions (up to x) with state p . We set

$$\text{STATE}_q(x) := \exists (Z_p)_{p \in Q} [\text{ADM} \wedge \text{INIT} \wedge \text{TRANS}(x) \wedge Z_q x],$$

where

$$\text{ADM} := \forall y \bigwedge_{p \neq p'} \neg (Z_p y \wedge Z_{p'} y)$$

states that every position is labelled by at most one state,

$$\text{INIT} := \exists y [\forall z (y \leq z) \wedge Z_{q_o} y]$$

states that the first state is q_o ,

$$\begin{aligned} \text{TRANS}(x) &:= \forall y \forall z [\text{suc}(y, z) \wedge z \leq x \\ &\quad \rightarrow \bigvee_{p \in Q} \bigvee_{a \in \Sigma} (Z_p y \wedge P_a y \wedge Z_{\delta(p, a)} z)] \end{aligned}$$

states that at every position a valid transition is used, and

$$\text{suc}(x, y) := x < y \wedge \neg \exists z [x < z \wedge z < y]$$

states that y is the immediate successor of x .

Using these formulae $\text{STATE}_q(x)$, we can construct a formula ψ defining L as follows. Let

$$H_k := \{ q \in Q \mid \Omega(q) \leq k \}$$

be the set of all states with priority at most k . We use the formula

$$\text{INF}_q := \forall x \exists y [x \leq y \wedge \text{STATE}_q(y)]$$

stating that the run contains infinitely many occurrences of the state q , and the formula

$$\text{MIN}_k := \bigvee_{q \in H_k} \text{INF}_q \wedge \bigwedge_{q \in H_{k-1}} \neg \text{INF}_q$$

stating that the minimal priority seen infinitely often is k . Then we can set

$$\psi := \bigvee_{k < n} \text{MIN}_{2k},$$

where n is chosen such that the maximal priority of \mathcal{A} is smaller than $2n$. \square

Together with the following result we obtain an alternative proof that the monadic theory of $\langle \omega, \leq \rangle$ is decidable.

Theorem 4.8. *Given an ω -automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_o, \Omega \rangle$, we can decide whether $L(\mathcal{A}) = \emptyset$.*

Proof. We claim that $L(\mathcal{A}) \neq \emptyset$ if, and only if, there exist two finite words $u, v \in \Sigma^*$ and a state $p \in Q$ such that

- ♦ there is a run of \mathcal{A} on u leading from the initial state q_o to p and
- ♦ there is a run of \mathcal{A} on v leading from p to p whose minimal priority is even.

(\Leftarrow) Clearly, if there are such words u and v , then \mathcal{A} has an accepting run on uv^ω . Hence, $L(\mathcal{A}) \neq \emptyset$.

(\Rightarrow) Let $w \in L(\mathcal{A})$ and let ρ be an accepting run of \mathcal{A} on w . Let d be the minimal priority occurring infinitely often in ρ and fix a state p that occurs infinitely often in ρ . Let $k < n < \omega$ be numbers such that

- ♦ $\rho(k) = p$,
- ♦ $\rho(n) = p$,
- ♦ no state $\rho(i)$ with $i \geq k$ has a priority smaller than d ,
- ♦ there is some $k \leq i < n$ such that $\rho(i)$ has priority d .

Let u be the prefix of w of length k and let v be the factor of w from position k to $n - 1$. These two words have the desired properties. \square

Exercise 4.3. Let Σ be a finite alphabet. The *Cantor topology* on Σ^ω is given by the following basis of open sets:

$$O_w := \{ x \in \Sigma^\omega \mid w \text{ is a prefix of } x \}, \quad \text{for } w \in \Sigma^*.$$

- (a) Show that every basic open set O_w is also closed.
- (b) Show that a set $U \subseteq \Sigma^\omega$ is open if, and only if, there exists a set $W \subseteq \Sigma^*$ such that

$$U = \{ x \in \Sigma^\omega \mid \text{some } w \in W \text{ is a prefix of } x \}.$$

- (c) Show that a set $C \subseteq \Sigma^\omega$ is closed if, and only if, there exists a set $W \subseteq \Sigma^*$ such that

$$C = \{ x \in \Sigma^\omega \mid \text{every finite prefix of } x \text{ belongs to } W \}.$$

- (d) Prove that Σ^ω is a compact Hausdorff space.
- (e) A set $U \subseteq \Sigma^\omega$ is a Π_2^0 -set if it can be written as a countable intersection of open sets. Prove that every language $L \subseteq \Sigma^\omega$ recognised by a deterministic ω -automaton is a finite boolean combination of Π_2^0 -sets. \square

Exercise 4.4. (a) Let \mathcal{A}_0 and \mathcal{A}_1 be ω -automata. Show that there are ω -automata recognising the languages $L(\mathcal{A}_0) \cap L(\mathcal{A}_1)$, $L(\mathcal{A}_0) \cup L(\mathcal{A}_1)$, and $\Sigma^\omega \setminus L(\mathcal{A}_0)$. (The case of the complement is a bit more involved.)

(b) Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ be an ω -automaton and let $\pi : \Sigma \rightarrow \Gamma$ be a function. Prove that there exists an ω -automaton recognising the language

$$\pi[L(\mathcal{A})] := \{ \pi(w) \mid w \in L(\mathcal{A}) \}.$$

(c) Use (a) and (b) to prove that, given an MSO-formula φ , we can effectively construct an ω -automaton recognising the language defined by φ .

(d) Show that we can decide whether a given ω -automaton \mathcal{A} recognises the empty language.

(e) Use (c) and (d) to give an alternative decidability proof for the monadic theory of $\langle \omega, \leq \rangle$. \square

Notes

For a good introduction to formal language theory for ω -words, ω -semigroups, and automata see [153].

The original proof of the decidability of the theory of ω is due to Büchi [35]. It combines automata-theoretic techniques with a Ramsey argument, see also [202] for a survey. The proof we presented is due to Shelah [193]. An exposition can be found in [203].

Ramseyan splits were introduced by Colcombet, extending results by Simon on so-called factorisation forests. An exposition that also includes a proof of Theorem 1.10 can be found in [51].

The Theorem of McNaughton (Theorem 4.5) is from [141]. A good exposition is [204]. Our proof is new and based on ideas by Colcombet.

IV Parity Games

1 Positional Games

BEFORE GENERALISING THE THEORY DEVELOPED in the two preceding chapters from words to trees, we need to develop a bit of combinatorial machinery. In particular, we need a substitute for the various versions of Ramsey's Theorem that works for trees. One such substitute is based on the notion of a combinatorial game, in particular, that of a parity game.

Many games like go, chess, and checkers can be modelled as a directed graph where the vertices represent the different *states*, or *positions*, of the game and the edges correspond to possible *moves*. Starting in a given initial state such a game consists of a sequence of moves which forms a path in the game graph. We call such a path a *partial play* of the game. The path is a (*complete*) *play* if it is either infinite or if it ends in some vertex without outgoing edges. For simplicity, we will only consider games with two players, *Player* \diamond and *Player* \square , and we assume that the *outcome* of each play is either a *win* for one of the two players, or a *draw*.

How exactly the actions of the players determine the next move depends on the specific kind of game we are considering. The three main options are: (i) *positional games* where the current position determines which player may choose where to move next; (ii) *alternating games* where the players alternate making this choice; and (iii) *simultaneous games* where both players act simultaneously and the resulting move is then determined by combining their choices in some way. In this book we only consider positional games.

Definition 1.1. A *positional game* is a tuple

$$\mathfrak{G} = \langle V_{\diamond}, V_{\square}, E, \Omega_{\diamond}, \Omega_{\square} \rangle,$$

where V_\diamond is the set of positions belonging to Player \diamond , V_\square are the positions for Player \square , $V := V_\diamond \cup V_\square$ is the set of all positions, $E \subseteq V \times V$ is the (directed) edge relation, and Ω_\diamond and Ω_\square are disjoint sets of infinite plays that determine the winning condition of the game as follows. A player that has to make a move, but cannot, loses. If this does not happen, the game results in an infinite play and player σ wins if this play belongs to Ω_σ . If the play does not belong to either set, the result is a draw.

The central problem associated with a positional game is which player wins when the game is started in a given initial position and to find a corresponding strategy.

Definition 1.2. Let \mathfrak{G} be a game.

(a) A *strategy* for player σ is a function s that, given a partial play p ending in a position $v \in V_\sigma$ chooses one of the outgoing edges. If the value of $s(p)$ does not depend on all of p , but only on the final position, we say that s is *positional* or *memory-free*.

(b) We say that a partial play p *conforms* to such a strategy s if, for every proper prefix p_\circ of p ending in a position for Player σ the extension of p_\circ by the edge $s(p_\circ)$ is again a prefix of p .

(c) Finally, a strategy s for Player σ is *winning* from a position $v \in V$ if he wins every (complete) play that starts in v and conforms to s . The *winning region* W_σ of Player σ is the set of all positions $v \in V$ from which he has a winning strategy.

Informally, we will say that a player *wins* a game if the given starting position belongs to his winning region. These notions of course only make sense if the game is *determined*.

Definition 1.3. A game \mathfrak{G} is *determined* if, from every initial position, either one of the players has a winning strategy, or both players have a strategy that guarantees at least a draw.

While most of the games one encounters ‘in the wild’ are in fact determined, we will see below that, although hard to find, there indeed exist games

which are not. The class of games we will consider below is even more well-behaved: they are not only determined but the corresponding strategies can always be chosen to be positional.

Definition 1.4. A game \mathfrak{G} is *positional determined* if, from every initial position, either one of the players has a *positional* winning strategy, or both players have a positional strategy that guarantees at least a draw.

In the terminology we have just established we can rephrase the central problem thus as:

Given a determined game \mathfrak{G} , find the winning regions and the corresponding winning strategies.


In the rest of this chapter we will consider this context for several classes of positional games.

2 Reachability Games

We start with the simplest form of a positional game.

Definition 2.1. A *reachability game* $\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega_\diamond, \Omega_\square \rangle$ is a positional game where the sets Ω_\diamond and Ω_\square are both empty, that is, a game where all infinite plays are considered draws.

As an example, in the game



(where the label \diamond denotes positions for Player \diamond and \square positions for Player \square) the shaded part constitutes the winning region for Player \diamond , while the remaining ones constitute the winning region for Player \square .

The winning regions of a reachability game are easy to compute recursively: a player wins from some position v if either v belongs to him and at least one outgoing edge leads to a winning position, or it belongs to his opponent and all the outgoing edges lead to winning positions. To define this formally we introduce the following notation. For $X \subseteq V$, we set

$$\begin{aligned}\Diamond X &:= \{ v \in V \mid \langle v, w \rangle \in E \text{ for some } w \in X \}, \\ \Box X &:= \{ v \in V \mid \langle v, w \rangle \in E \text{ implies } w \in X \}.\end{aligned}$$

We denote the opponent of Player σ by $\bar{\sigma}$, i.e., $\overline{\Diamond} := \Box$ and $\overline{\Box} := \Diamond$. Then we can define the winning region for Player σ as the least set W_σ such that

$$W_\sigma = (V_\sigma \cap \Diamond W_\sigma) \cup (V_{\bar{\sigma}} \cap \Box W_\sigma).$$

Hence, W_σ is the least fixed-point of the following function.

Definition 2.2. The *step function* associated with a game \mathfrak{G} is

$$\text{Step}_\sigma(X) := (V_\sigma \cap \Diamond X) \cup (V_{\bar{\sigma}} \cap \Box X).$$

It is easy to see that $\text{Step}_\sigma(X)$ contains all the positions from which Player σ can ensure that in the next step the game either reaches some position in X or the game ends with a win for him. Thus, by iterating the step function we obtain the set of all positions from which the player either wins or eventually reaches a position in X . This iteration of Step_σ is called the σ -attractor of the set X . The formal definition is as follows.

Definition 2.3. The σ -attractor of a set X is

$$\text{Attr}_\sigma(X) := \bigcup_{\alpha} \text{Step}_\sigma^\alpha(X),$$

where α ranges over all ordinals (actually, it is sufficient to take the union for all $\alpha < |V|^+$) and $\text{Step}_\sigma^\alpha$, the α -th iteration of the step function, is defined as follows

$$\begin{aligned}\text{Step}_\sigma^0(X) &:= X, \\ \text{Step}_\sigma^{\alpha+1}(X) &:= \text{Step}_\sigma(\text{Step}_\sigma^\alpha(X)), \\ \text{Step}_\sigma^\delta(X) &:= \bigcup_{\alpha < \delta} \text{Step}_\sigma^\alpha(X), \quad \text{for limit ordinals } \delta.\end{aligned}$$

Lemma 2.4. Let \mathfrak{G} be a positional game and $X \subseteq V$ a set of positions. Player σ has a positional strategy s such that, every play p that starts in some position $v \in \text{Attr}_\sigma(X)$ and that conforms to s is winning or contains some position from X .


Proof. Let $v \in \text{Attr}_\sigma(X)$. Then $v \in \text{Step}_\sigma^\alpha(X)$, for some ordinal α . We prove the claim by induction on α .

If $\alpha = 0$, then $v \in X$ and the claim is trivial. For the successor step, suppose that $\alpha = \beta + 1$. Then $v \in \text{Step}_\sigma(\text{Step}_\sigma^\beta(X))$ and Player σ has a strategy to either win in one step, or to reach some vertex of $\text{Step}_\sigma^\beta(X)$. In the first case we are done. In the second one, we can continue with the strategy from the inductive hypothesis. Finally, if α is a limit ordinal, then $v \in \text{Step}_\sigma^\beta(X)$, for some $\beta < \alpha$, and the claim follows immediately from the inductive hypothesis. \square

Using the notion of an attractor we can define a measure for how long it takes from a given position to win. The *rank* of a position v is the least ordinal α such that

$$v \in \text{Step}_\sigma^{\alpha+1}(\emptyset).$$

If there is no such ordinal, we set the rank to ∞ . In the above example the ranks for Player \Diamond are



Next let us take a look at complements of attractors.

Definition 2.5. We call a subset $U \subseteq V$ a σ -trap if $\text{Step}_\sigma(V \setminus U) \subseteq V \setminus U$, that is, if the opponent can ensure that Player σ never leaves the set U once the game has entered it. \square

An easy way to find traps is by computing attractors.

Lemma 2.6. *A set U is a σ -trap if, and only if, it is of the form*

$$U = V \setminus \text{Attr}_\sigma(X), \quad \text{for some } X.$$

Proof. (\Leftarrow) is obvious since $\text{Step}_\sigma(\text{Attr}_\sigma(X)) \subseteq \text{Attr}_\sigma(X)$. For (\Rightarrow), note that $\text{Step}_\sigma(V \setminus U) \subseteq V \setminus U$ implies $U = V \setminus \text{Attr}_\sigma(V \setminus U)$. \square

Lemma 2.7. *For every σ -trap U , Player $\bar{\sigma}$ has a positional strategy s ensuring that, starting from any vertex $v \in U$, the game never leaves U .*

Proof. Setting $A := V \setminus U$, we know that $A = \text{Attr}_\sigma(X)$, for some set X . Hence,

$$V_\sigma \setminus A = V_\sigma \setminus \text{Step}_\sigma(A) = V_\sigma \setminus \Diamond A$$

$$\text{and } V_{\bar{\sigma}} \setminus A = V_{\bar{\sigma}} \setminus \text{Step}_{\bar{\sigma}}(A) = V_{\bar{\sigma}} \setminus \Box A.$$

In particular, (I) no vertex $v \in V_\sigma \cap U$ has an outgoing edge leading to a vertex in A and (II) every vertex $v \in V_{\bar{\sigma}} \cap U$ has at least one outgoing edge $\langle v, u \rangle \in E$ with $u \in U$. Consequently, if the game starts in some vertex $v \in U$ Player σ can never move into A , while Player $\bar{\sigma}$ always has the option to stay in U . \square

Theorem 2.8. *Every reachability game is positionally determined with winning regions $W_\sigma := \text{Attr}_\sigma(\emptyset)$.*

Proof. Given σ , let s be the strategy for Player σ from Lemma 2.4 for $W_\sigma = \text{Attr}_\sigma(\emptyset)$, and let t be the strategy from Lemma 2.7 for $V \setminus W_\sigma$. As no play can ever reach a position in \emptyset , it follows that s is winning for Player σ from every position $v \in W_\sigma$. While, t ensures that Player $\bar{\sigma}$ does not lose when starting from a position in $V \setminus W_\sigma$. This implies that W_σ is the winning region for Player σ and that, for $V \setminus (W_\Diamond \cup W_\Box)$, each player has a positional strategy ensuring that he does not lose. \square

Having shown that winning regions exist, we next take a look at how to efficiently compute them.

```

function Win( $v, \sigma$ )    // Does Player  $\sigma$  win from position  $v$ ?
  if  $v \in V_\sigma$  then
    if there is an edge  $v \rightarrow u$  with  $\text{Win}(u, \sigma)$  then
      return true
    else
      return false
  if  $v \in V_{\bar{\sigma}}$  then
    if for every edge  $v \rightarrow u$  we have  $\text{Win}(u, \sigma)$  then
      return true
    else
      return false
end

```

Figure 1: Quadratic time algorithm

Theorem 2.9. *The winning regions of a finite reachability game can be computed in linear time.*

Before presenting the linear time algorithm we start with a simpler, non-linear version which is depicted in Figure 1 and which is just a direct translation of the definition of an attractor. There are several obvious problems with this algorithm. First of all, it might not terminate if the game graph contains a cycle. And secondly, it is very inefficient (exponential time) as it does not remember if it has already computed the winner of a position and recomputes this information every time. There is a rather straightforward fix for both of these issues: we introduce an array where we store whether we have already visited a position and who the winner is. (Thus each entry can have one of four values: (I) not visited yet, (II) already visited, but we do not know the winner yet, (III) Player \Diamond wins, and (IV) Player \Box wins.) With this modification, the algorithm will run in quadratic time.

To improve the runtime to linear, we need to be more clever. In the algorithm in Figure 2, we introduce two more arrays with auxiliary data that helps us to avoid unnecessary work. To see that this algorithm works in

Input: a game $\langle V_{\diamond}, V_{\square}, E \rangle$

Output: an array containing the winner for every position

// initialise auxiliary arrays

forall $v \in V$ **do**

$\text{win}[v] := \perp$ // the winner of the position

$P[v] := \emptyset$ // the set of predecessors of v

$n[v] := 0$ // the number of (not yet processed)

 // successors of v

end

forall $\langle u, v \rangle \in E$ **do**

$P[v] := P[v] \cup \{u\}$

$n[u] := n[u] + 1$

end

// compute the winning regions

forall $v \in V_{\diamond}$ **do**

if $n[v] = 0$ **then** $\text{Propagate}(v, \square)$

forall $v \in V_{\square}$ **do**

if $n[v] = 0$ **then** $\text{Propagate}(v, \diamond)$

return win

procedure $\text{Propagate}(v, \sigma) =$

if $\text{win}[v] \neq \perp$ **then return**

$\text{win}[v] := \sigma$

forall $u \in P[v]$ **do**

$n[u] := n[u] - 1$

if $u \in V_{\sigma}$ **or** $n[u] = 0$ **then** $\text{Propagate}(u, \sigma)$

end

end

Figure 2: Linear time algorithm

linear time (in the number of positions plus the number of edges), note that the body of the procedure *Propagate* (except for the first line) is executed exactly once for each vertex v . Furthermore, the loop in *Propagate* is executed once for each incoming edge which means that, in total, it is executed at most as many times as there are edges in the game. Since the precomputation steps are also linear in the number of vertices or edges, it follows that so is the total runtime of the algorithm.

It remains to show that the algorithm really computes the winning regions W_{\diamond} and W_{\square} . To see this it is sufficient to note that, every time the procedure $\text{Propagate}(v, \sigma)$ is called and we still have $\text{win}[v] = \perp$, then the vertex v really belongs to the winning region for Player σ . This is clear for the two calls of *Propagate* in the main part of the algorithm, where only vertices v without successors are considered. For the recursive call inside the body of *Propagate* we need to distinguish two cases. If $u \in V_{\sigma}$ and v already belongs to W_{σ} then Player σ can take the edge from u to v to win. Hence, $u \in W_{\sigma}$. Otherwise, we have $u \in V_{\bar{\sigma}}$ and $n[u] = 0$, which means that we already know for all successors w_0, \dots, w_n of u to which region they belong. If there is some $w_i \in W_{\bar{\sigma}}$ then we have already called $\text{Propagate}(u, \bar{\sigma})$ when processing w_i and $\text{win}[u]$ is already set. Otherwise, all successors belong to W_{σ} , which means that u belongs to it as well.

Horn Formulae

As an application of reachability games let us take a look at the satisfiability problem for propositional Horn formulae. Such a formula is an implication of the form

$$A_1 \wedge \dots \wedge A_n \rightarrow B,$$

where we allow both the left-hand side and the right-hand side to be empty, i.e., we allow implications of the form $\perp \rightarrow B$ and $A_1 \wedge \dots \wedge A_n \rightarrow \perp$. We are interested in deciding in whether a given set of such formulae is satisfiable. Note that such a set is always satisfiable if there are no implications where the right-hand side is \perp . We call such implications *purely negative*. It is not difficult to prove that every set of Horn formulae with no purely negative

implications has a *minimal model*, that is, there exists a unique variable assignment that satisfies all the formulae and that only assigns 1 to those variables that are true in *every* satisfying variable assignment. We will use games to show that this minimal model can be computed in linear time. Then we can check for satisfiability of a given set Φ of Horn formulae by

- (1) removing all the purely negative implications from Φ ,
- (2) computing the minimal model, and
- (3) checking that every of the removed implications is true in this model.

As an example, let us consider the following set of Horn formulae.

$$\begin{array}{lll} 1 \rightarrow A & A \wedge C \wedge F \rightarrow D & E \rightarrow G \\ A \wedge D \rightarrow B & B \wedge E \wedge G \rightarrow D & 1 \rightarrow E \\ F \rightarrow C & G \rightarrow D & \end{array}$$

The minimal model assigns the value 1 to A, B, D, E , and G , and the value 0 to C and F .

The game corresponding to a set Φ of Horn formulae looks as follows. The positions for Player \Diamond are of the form $\langle A \rangle$, where A is a variable appearing in Φ , the positions for Player \Box are of the form $[\varphi]$ with $\varphi \in \Phi$. For each formula $A_0 \wedge \dots \wedge A_{n-1} \rightarrow B \in \Phi$, we have edges


$$\langle B \rangle \longrightarrow [A_0 \wedge \dots \wedge A_{n-1} \rightarrow B]$$

$$\text{and } [A_0 \wedge \dots \wedge A_{n-1} \rightarrow B] \longrightarrow \langle A_i \rangle, \quad \text{for } i < n.$$

Intuitively, in the resulting game Player \Diamond tries to prove that a variable A must have value 1 by choosing an implication $B_0 \wedge \dots \wedge B_{n-1} \rightarrow A$ that forces it to be true, while Player \Box tries to prove that such an implication is not applicable by finding some condition B_i that is not met. With this intuition it is straightforward to show that our game has the desired properties.

Lemma 2.10. *Let Φ be a set of Horn formulae that does not contain any purely negative implications. A position of the form $\langle A \rangle$ belongs to the winning region of Player \Diamond if, and only if, the variable A is true in the minimal model of Φ .*

The game corresponding to the above set of formulae is the following one, whose winning regions we have already computed above.



Modal Logic

As a second application let us take a look at the model-checking problem for propositional modal logic. Recall that (*propositional*) *modal logic* is the extension of propositional logic by two modal operators $\Diamond\varphi$ and $\Box\varphi$. Formulae of this logic are evaluated over transition systems $\mathfrak{S} = \langle S, E, \bar{P} \rangle$, i.e., directed graphs expanded by additional unary predicates \bar{P} . The satisfaction relation for modal logic takes the form $\mathfrak{S}, s \models \varphi$, where \mathfrak{S} is a transition system, $s \in S$ a state, and φ a formula. It is defined by induction on φ as follows.

$$\begin{array}{ll} \mathfrak{S}, s \models P & : \text{iff } s \in P^{\mathfrak{S}}, \\ \mathfrak{S}, s \models \Diamond\varphi & : \text{iff there is an edge } s \rightarrow t \text{ with } \mathfrak{S}, t \models \varphi, \\ \mathfrak{S}, s \models \Box\varphi & : \text{iff } \mathfrak{S}, t \models \varphi, \text{ for every edge } s \rightarrow t, \end{array}$$

and the definition for boolean operations is the usual one.

Given a transition system \mathfrak{S} and a formula φ , we can construct a game that is won by Player \Diamond if, and only if, $\mathfrak{S}, s \models \varphi$.

Definition 2.11. Let $\mathfrak{S} = \langle S, E, \bar{P} \rangle$ be a transition system with starting state $s \in S$ and let φ be a modal formula in *negation normal form*. The *model-checking game* $\mathcal{G}(\mathfrak{S}, \varphi)$ is defined as follows. The positions are of the form $\langle t, \psi \rangle$ where t is a state of \mathfrak{S} and ψ is a subformula of φ . Intuitively, in such a position Player \Diamond tries to prove that $\mathfrak{S}, t \models \psi$, while Player \Box tries to show




Figure 3: Model-checking game for modal logic

that $\mathfrak{S}, t \models \psi$. The moves are as follows.

$$\begin{aligned}
 \langle t, \psi_0 \vee \psi_1 \rangle &\rightarrow \langle t, \psi_i \rangle, & \text{for } i = 0, 1, \\
 \langle t, \psi_0 \wedge \psi_1 \rangle &\rightarrow \langle t, \psi_i \rangle, & \text{for } i = 0, 1, \\
 \langle t, \diamond \vartheta \rangle &\rightarrow \langle u, \vartheta \rangle, & \text{if } t \rightarrow u \text{ is an edge of } \mathfrak{S}, \\
 \langle t, \square \vartheta \rangle &\rightarrow \langle u, \vartheta \rangle, & \text{if } t \rightarrow u \text{ is an edge of } \mathfrak{S}.
 \end{aligned}$$

Finally, positions of the form $\langle t, \psi_0 \vee \psi_1 \rangle$ and $\langle t, \diamond \vartheta \rangle$ belong to Player \diamond ; those of the form $\langle t, \psi_0 \wedge \psi_1 \rangle$ and $\langle t, \square \vartheta \rangle$ belong to Player \square ; and a position of the form $\langle t, P \rangle$ belongs to Player \diamond if $t \notin P$; otherwise, it belongs to Player \square . Similarly, a position of the form $\langle t, \neg P \rangle$ belongs to Player \diamond if $t \in P$.

Example. For the transition system \mathfrak{S} and the formula φ on the left of Figure 3, the resulting game (or at least its reachable part) is depicted on the right.

It is straightforward to check by induction on φ that the game $\mathcal{G}(\mathfrak{S}, \varphi)$ has the desired properties.

Theorem 2.12. *In the game $\mathcal{G}(\mathfrak{S}, \varphi)$, Player \diamond has a winning strategy from a position $\langle t, \psi \rangle$ if, and only if, $\mathfrak{S}, t \models \psi$.*

3 Gale-Stewart Games

Reachability games are very simple since we can ignore infinite plays. Let us now take a look at what happens for games where every infinite play is winning for one player or the other. We start by games played on a tree (which is not really a restriction since every game can be unravelled to one of this form).

Definition 3.1. A Gale-Stewart game is a game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega_\diamond, \Omega_\square \rangle$ where the set of positions is of the form $V = A^*$, for some set A . Every position $v \in V$ has outgoing edges to all vertices of the form va with $a \in A$. The players strictly alternate, that is, positions $v \in (A^2)^*$ of even length belong to Player \diamond and those $v \in A(A^2)^*$ of odd length to Player \square . Finally, we assume that there are no draws, that is, that $\Omega_\diamond \cup \Omega_\square = A^\omega$. We denote such a game by the pair $\langle A^*, \Omega_\diamond \rangle$.

Proposition 3.2. *There exist Gale-Stewart games that are not determined.*

Proof. We play on the complete binary tree with vertices $V := \{0, 1\}^*$. As the game graph is acyclic, every strategy is automatically positional. Thus, a strategy for Player σ is a function $V_\sigma \rightarrow \{0, 1\}$ where V_σ is either the set of all $v \in \{0, 1\}^*$ of even length, or the set of all v of odd length. There are $\kappa := 2^{\aleph_0}$ such functions. We fix enumerations $(s_\alpha)_{\alpha < \kappa}$ and $(t_\alpha)_{\alpha < \kappa}$ of all strategies for, respectively, Player \diamond and Player \square . To construct a non-determined game we have to find two disjoint sets $\Omega_\diamond, \Omega_\square \subseteq \{0, 1\}^\omega$ of infinite paths such that none of the s_α and t_α are winning strategies in the game with winning condition $\Omega_\diamond, \Omega_\square$.

We start with a bit of notation. Given a strategy s , we denote by $[s]$ the set of all infinite plays $p \in A^\omega$ that conform to s . With this notation we can say that a strategy s for Player σ is a winning if, and only if, $[s] \subseteq \Omega_\sigma$.

By induction on $i < \kappa$, we construct two sequences $(\xi_i)_{i < \kappa}$ and $(\zeta_i)_{i < \kappa}$ of elements of A^ω as follows. Suppose that we have already defined ξ_i and ζ_i , for all $i < \alpha$. Then we pick some element $\xi_\alpha \in [s_\alpha]$ that is different from ξ_i and ζ_i , for all $i < \alpha$, and we pick some element $\zeta_\alpha \in [t_\alpha]$ that is different from ξ_i and ζ_i , for all $i \leq \alpha$, and also from the ξ_α we have just chosen. Note

that we can do so since $[s_\alpha]$ has size κ , while the set

$$\{\xi_i \mid i < \alpha\} \cup \{\zeta_i \mid i < \alpha\}$$

has size $|\alpha| < \kappa$. The same holds for $[t_\alpha]$.

We claim that the game with winning conditions

$$\Omega_\diamond := \{\zeta_i \mid i < \kappa\} \quad \text{and} \quad \Omega_\square := A^\omega \setminus \Omega_\diamond$$

is not determined. For the proof, consider a strategy s for Player \diamond . Then $s = s_\alpha$, for some $\alpha < \kappa$. Since $\xi_\alpha \in [s_\alpha] \cap \Omega_\square$, we have $[s] \notin \Omega_\diamond$. Hence, s is not a winning strategy. In the same way it follows that no strategy t for Player \square is winning. \square

As Gale-Stewart games can be non-determined in general, we have to put restrictions on the allowed winning conditions to get positive results. One handy way to do so is by equipping the set of all infinite plays with a topology. We call a set O of infinite plays *open* if there exists a set P of finite partial plays such that contains all infinite plays starting with some $p \in P$. The complement of an open set is called *closed*. Note that the open sets are closed under arbitrary unions and finite intersections. Hence, they form a topology. A set is *Borel* if it is contained in the smallest class of sets that contains the open ones and that is closed under complement and countable unions.

Theorem 3.3 (Martin). *If $\Omega \subseteq A^\omega$ is Borel, then $\langle A^*, \Omega \rangle$ is determined.*

The proof is a bit involved. Instead of proving the result in its full generality, we will only consider the much simpler case of open and closed winning conditions, i.e., where Ω_\diamond is open and Ω_\square closed, or vice versa. By symmetry, we may assume that the winning condition for Player \diamond is open. Then Ω_\diamond is determined by some set $P \subseteq A^*$ of prefixes and every play containing a position from P is winning. Thus, open winning condition corresponds to a reachability game: Player \diamond has to reach a position in P . As we have seen in Section 2, the winning region for Player \diamond in such a game is given by $\text{Attr}_\diamond(P)$. The complement of $\text{Attr}_\diamond(P)$ is a \diamond -trap. Since every infinite

play inside this complement cannot contain positions from P , all such plays are winning for Player \square . Consequently, the game is determined with winning regions

$$\text{Attr}_\diamond(P) \quad \text{and} \quad A^* \setminus \text{Attr}_\diamond(P).$$

4 Regular Games and Parity Games

In general, Gale-Stewart games have no finite representation as the winning set Ω can be an arbitrary set of infinite sequences. In this section we take a look at a simple way to represent certain Gale-Stewart games in a finite way: if the winning set Ω is a *regular* set of infinite plays, we can use an automaton to represent it. Or we can use ω -semigroups instead of automata. This leads to the following definition.

Definition 4.1. A *regular game* over an ω -semigroup $\mathfrak{S} = \langle S, S_\omega \rangle$ is a game

$$\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, \Omega \rangle$$

where $\lambda : E \rightarrow S$ is an edge-labelling and the winning set is given by a subset $\Omega \subseteq S_\omega$. Player \diamond wins an infinite play p if the product of the corresponding edge labels evaluates to an element of Ω . \lrcorner

Example. Consider the ω -semigroup $\mathfrak{S} = \langle S, S_\omega \rangle$ where $S = \wp\{0, 1\}$ and $S_\omega = \wp\{0, 1\}^\omega$. We define the product by

$$\begin{aligned} a \cdot b &:= a \cup b, & \text{for } a, b \in S, \\ a \cdot u &:= u, & \text{for } a \in S \text{ and } u \in S_\omega, \\ \pi(a_0, a_1, \dots) &:= \bigcap_{i < \omega} \bigcup_{i \leq k < \omega} a_i, & \text{for } a_0, a_1, \dots \in S. \end{aligned}$$

In the game



with winning set $\Omega = \{\{0, 1\}\}$ Player \diamond has a winning strategy by alternating between the two edges. But note that he does not have a positional winning strategy as using only one of the edges will result in a loss.

Computing the winning regions of a regular game is more complicated than for open games. The difference is that, instead of reaching a certain set once, we have to be able to reach it over and over again. To simplify our task, let us start by considering a special case of regular games of the following form.


Definition 4.2. A *parity game* is a game of the form

$$\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega \rangle$$

where $\Omega : V \rightarrow D$ for some finite set $D \subseteq \omega$. We call Ω the *priority function* and $\Omega(v)$ the *priority* of the position v . Player \diamond wins an infinite play $p = (v_i)_{i < \omega}$ if it satisfies the *parity condition*:

$$\liminf_{i < \omega} \Omega(v_i) \text{ is even.}$$

Example. In the following parity game Player \diamond wins from every position except for the one in the lower right. (The numbers denote the priorities.)




As we will show below, parity games are simpler than general regular games. In particular, they are positionally determined. In fact, a regular game is positionally determined if, and only if, it is a parity game. Let us start our investigation of parity games by noting that they are regular. We can turn the set D of priorities into an ω -semigroup $\mathfrak{D} = \langle D, D_\omega \rangle$ with two infinite

elements $D_\omega := \{0, 1\}$ and the product

$$\begin{aligned} k \cdot l &:= \min\{k, l\}, & \text{for } k, l \in D, \\ k \cdot b &:= b, & \text{for } k \in D \text{ and } b \in D_\omega, \\ \pi(k_0, k_1, \dots) &:= (\liminf_{i < \omega} k_i) \bmod 2, & \text{for } k_0, k_1, \dots \in D. \end{aligned}$$

Then a parity game \mathfrak{G} can be turned into a regular game by labelling every edge $\langle u, v \rangle$ by $\Omega(v)$. In the above example we get:



Positional Determinacy

For the proof that parity games are positionally determined, we introduce some terminology.

Definition 4.3. (a) We say that a strategy s for Player σ is *winning on* some set $U \subseteq V$ if every play p that conforms to s and starts at a vertex in U is winning for σ and p never leaves the set U .

(b) We call a subset $U \subseteq V$ a σ -*domain* if it is a $\bar{\sigma}$ -trap and Player σ has a positional strategy that is winning on U .

Lemma 4.4. The union of a set of σ -domains is again a σ -domain.

Proof. Let $W = \bigcup_{i \in I} U_i$, where each U_i is a σ -domain. We start by proving that W is a $\bar{\sigma}$ -trap. Let $v \in W$. Then $v \in U_i$, for some $i \in I$. We distinguish two cases. If $v \in V_\sigma$ it follows that there is an edge $\langle v, u \rangle$ with $u \in U_i \subseteq W$. If $v \in V_{\bar{\sigma}}$ then every edge $\langle v, u \rangle$ leads to a vertex $u \in U_i \subseteq W$. This implies that W is a $\bar{\sigma}$ -trap.

It remains to construct a positional strategy s for Player σ on W . By assumption, there are positional strategies t_i that are winning on U_i . Fix a

$$s(v) := t_i(v), \quad \text{for the least } i \in I \text{ with } v \in U_i,$$
$$s(v) = t_i(v) \in U_i \subseteq W, \quad \text{for every vertex } v \in V_\sigma \cap W.$$

Theorem 4.5. *Parity games are positionally determined.*

Since $\text{Attr}_{\bar{\sigma}}(W_{\bar{\sigma}})$ is a $\bar{\sigma}$ -domain, it follows by definition of $W_{\bar{\sigma}}$ that $\text{Attr}_{\bar{\sigma}}(W_{\bar{\sigma}}) = W_{\bar{\sigma}}$. Consequently, Lemma 2.6 implies that the complement $W_{\sigma} = V \setminus \text{Attr}_{\bar{\sigma}}(W_{\bar{\sigma}})$ is a $\bar{\sigma}$ -trap. For a set $X \subseteq V$, we denote by $\mathfrak{G}[X]$ the subgame of \mathfrak{G} consisting of all positions in X . Define

$$K := W_\sigma \cap \Omega^{-1}(k) \quad \text{and} \quad U := W_\sigma \setminus \text{Attr}_\sigma(K/W_\sigma),$$

The diagram shows a directed graph with nodes represented by squares and diamonds. The nodes are grouped into three sets: U (shaded gray), K (dashed circle), and W (dashed circle). U contains nodes 3, 5, I, 2, 2, 2, 3. K contains nodes 0 and 2. W contains node 3. Edges show transitions between these sets. Labels $U = U_\diamond$ and W_\square are at the bottom.

It remains to find a positional strategy for Player \bar{o} on $W_{\bar{o}} \cup U_{\bar{o}}$. As $W_{\bar{o}}$ and $U_{\bar{o}}$ are \bar{o} -domains in, respectively, \mathfrak{G} and $\mathfrak{G}[U]$, there are positional strategies t_W and t_U for Player \bar{o} on these to sets. We define a strategy s by

$$s(v) := \begin{cases} t_U(v) & \text{if } v \in U, \\ t_W(v) & \text{otherwise.} \end{cases}$$

To show that s is winning, consider a play p conforming to s and starting in some position in $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$. If p enters $W_{\bar{\sigma}}$, it never leaves this set. Consequently, the rest of p conforms to t_W and is therefore winning. Otherwise, the play stays the whole time in $U_{\bar{\sigma}}$ and conforms to t_U . Thus, it is also winning.

We have shown that $U = U_\sigma$ which, by inductive hypothesis, is a σ -domain in $\mathfrak{G}[U]$. Let t be the corresponding strategy. To conclude the proof it is sufficient to show that $W_\sigma = V \setminus W_{\bar{\sigma}}$ is a σ -domain in \mathfrak{G} . We have already seen above that it is a $\bar{\sigma}$ -trap. Hence, it remains to construct a positional strategy for Player σ on W_σ . Note that $W_\sigma = U_\sigma \cup \text{Attr}_\sigma(K/W_\sigma)$. For positions $v \in U$, we use the strategy t for $\mathfrak{G}[U]$. For positions $v \in K$, we choose an arbitrary successor in W_σ . For the remaining positions, we use the attractor strategy that ensures that we visit K . Let s be the resulting strategy. To show that it is winning, consider a play p conforming to s and starting in some position in W_σ . If p enters U_σ it will stay in this set and the remainder of the play conforms to t . Hence, p is winning. Otherwise, p is entirely contained in $\text{Attr}_\sigma(K/W_\sigma)$. This implies that p either ends in a terminal vertex belonging to Player $\bar{\sigma}$, or it visits the set K infinitely often. In both cases Player σ wins. \square


Finite-Memory Strategies

One can show that all regular games are Borel. Hence, determinacy follows from the Theorem of Martin. But we can prove a stronger statement: regular games admit what is called *finite-memory* strategies.

Definition 4.6. A *finite-memory strategy* for Player σ is given by a finite set M (the *memory*) and two functions $s : M \times V_\sigma \rightarrow E$ and $\alpha : M \times E \rightarrow M$ such that, given a state $m \in M$ and a vertex $v \in V_\sigma$, $s(m, v)$ returns the outgoing edge e to be chosen by Player σ after which $\alpha(m, e)$ will be the new memory state. (If there is no outgoing edge, we let s remain undefined.) Formally, we say that a play $p = (e_i)_i$ (which, for regular games where the edge labelling matters, we consider as a sequence of edges) *conforms* to such a strategy if there exists a sequence $(m_i)_i$ of memory states such that, for every step i ,

- ♦ if $e_i = \langle v_i, v_{i+1} \rangle$ with $v_i \in V_\sigma$, then $e_i = s(m_i, v_i)$, and
- ♦ $m_{i+1} := \alpha(m_i, e_i)$.

Example. In the game



from above, Player \diamond has a finite-memory strategy with $M = \{0, 1\}$. If $m = 0$, he takes the left edge and sets the memory state to 1. Otherwise, he takes the right edge and sets the state to 0. This results in alternately taking the two edges, which is winning for him. \lrcorner

We can reformulate the definition of a finite-memory strategy as follows. For a regular game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, \Omega \rangle$ and a function $\alpha : M \times E \rightarrow M$, we define the *product game*

$$\mathfrak{G} \times_\alpha M := \langle V'_\diamond, V'_\square, E', \lambda', \Omega' \rangle$$

with positions


$$V'_\diamond := V_\diamond \times M \quad \text{and} \quad V'_\square := V_\square \times M,$$

edge relations

$$E' := \{ \langle \langle u, m \rangle, \langle v, n \rangle \rangle \mid \langle u, v \rangle \in E \text{ and } n = \alpha(u, \langle u, v \rangle) \},$$

edge labelling $\lambda'(\langle u, m \rangle, \langle v, n \rangle) := \lambda(\langle u, v \rangle)$, and the same winning condition $\Omega' := \Omega$. Then a strategy is finite-memory for \mathfrak{G} if, and only if, there exists a finite set M and a function $\alpha : M \times E \rightarrow M$ such that s is a positional strategy in the game $\mathfrak{G} \times_\alpha M$.

Example. In the above example, the product $\mathfrak{G} \times_\alpha M$ is the game



which clearly has a positional winning strategy. \lrcorner

Remark. Note that this operation of equipping a game with memory does not change the game much. There exist one-to-one correspondences between

- ♦ plays of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$;
- ♦ strategies of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$;

- ♦ winning strategies of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$.

The only difference between these two games is the amount of memory a strategy needs. In particular, some positional strategies of $\mathfrak{G} \times_\alpha M$ might correspond to strategies of \mathfrak{G} which are not positional. ,

This is exactly what we need to prove the following result.

Theorem 4.7 (Büchi, Landweber). *In every regular game both players have a finite-memory winning strategy on their respective winning regions.*

Proof. Let $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, \Omega \rangle$ be a regular game. We fix a deterministic parity automaton $\mathcal{A} = \langle Q, S, \delta, q_o, \Omega \rangle$ recognising the set Ω of winning plays and construct the product game $\mathfrak{G} \times_\delta \mathcal{Q}$. In this game a play p is winning for Player \diamond if, and only if, its projection to the second component produces an accepting run of \mathcal{A} . Consequently, we can turn $\mathfrak{G} \times_\delta \mathcal{Q}$ into a parity game by using as priority function the function Ω from \mathcal{A} applied to the second component. Since parity games are positionally determined, we obtain two positional winning strategies s_\diamond and s_\square for the two players in their respective winning regions. As we have seen in the remark before the theorem, these two strategies induce finite-memory strategies in the original game. \square

Positionally Determined Games

We have seen that parity games admit positional strategies while arbitrary regular games in general only admit finite-memory ones. One might wonder whether there exists a larger class of games with positional strategies. It turns out that this is not the case. We will prove below that (nearly) every regular, positionally determined game is equivalent to a parity game – with one notable caveat: we will only be able to establish this statement for games with a winning condition of the following form.

Definition 4.8. A winning condition $\Omega \subseteq S_\omega$ is called *prefix-invariant* if

$$w \in \Omega \Leftrightarrow aw \in \Omega, \quad \text{for all } a \in S \text{ and } w \in S_\omega. \quad \text{,}$$

Note that this is not much of a restriction since most of the common winning conditions used in game theory or automata theory are of this form. To analyse such conditions we start with an observation from semigroup theory. We have seen in Lemma III.3.4 that, in a *finite* ω -semigroup \mathfrak{S} , every infinite product $a_0 a_1 a_2 \dots$ has a factorisation of the form be^ω . If Ω is prefix-invariant, we have $be^\omega \in \Omega \Leftrightarrow e^\omega \in \Omega$. Thus, the set Ω is completely determined by the powers e^ω it contains.

Definition 4.9. Let $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, \Omega \rangle$ be a regular game over an ω -semigroup \mathfrak{S} .

- (a) We write $\Omega_\diamond := \Omega$ and $\Omega_\square := S_\omega \setminus \Omega$.
- (b) The *winning condition* of \mathfrak{G} is the pair $\langle \mathfrak{S}, \Omega \rangle$.
- (c) We call the set

$$P_\sigma := \{ e \in S \mid e^\omega \in \Omega_\sigma \}$$

the *period set* for Player σ .

- (d) Finally, let us say that $\langle \mathfrak{S}, \Omega \rangle$ is *equivalent* to a parity condition if there exists a function $\Omega : S \rightarrow \omega$ such that

$$\pi(a_0, a_1, \dots) \in \Omega \quad \text{iff} \quad \liminf_i \Omega(a_i) \text{ is even.} \quad \text{,}$$

If $\langle \mathfrak{S}, \Omega \rangle$ is equivalent to a parity condition, we can turn every game \mathfrak{G} with winning condition $\langle \mathfrak{S}, \Omega \rangle$ into a parity game as follows. We first replace all edge labels by their image under Ω . In this way we obtain a kind of parity game where the priorities are attached to the edges instead of the vertices. We can turn the resulting game into an ordinary parity game by adding intermediate vertices to the edges where we can put the priorities.

Theorem 4.10 (Colcombet, Niwiński). *Let \mathfrak{S} be an ω -semigroup (not necessarily finite) and $\Omega \subseteq S_\omega$ prefix-invariant. If all games with winning condition $\langle \mathfrak{S}, \Omega \rangle$ are positionally determined, then $\langle \mathfrak{S}, \Omega \rangle$ is equivalent to a parity condition.*

We split the proof of this theorem into two lemmas. The first one collects some basic properties of the period sets.

Lemma 4.11. Let \mathfrak{S} be an ω -semigroup and $\Omega \subseteq S_\omega$ prefix-invariant. If all games with winning condition (\mathfrak{S}, Ω) are positionally determined, then the following condition holds:

- (a) $a, b \in P_\sigma$ implies $ab \in P_\sigma$.
- (b) $ab \in P_\sigma$ implies $ba \in P_\sigma$.
- (c) Every element $w \in S_\omega$ that can be written as an infinite product of elements of P_σ belongs to Ω_σ .
- (d) For all $A, B \subseteq S$,

$$(\exists a \in A)(\forall b \in B)[ab \in P_\sigma] \Leftrightarrow (\forall b \in B)(\exists a \in A)[ab \in P_\sigma].$$


- (e) For every $a \in S$, there exists some $n > 0$ such that

$$ab_0, \dots, ab_k \in P_\sigma \Rightarrow a^n b_0 \dots b_k \in P_\sigma, \quad \text{for all } b_0, \dots, b_k \in S.$$

- (f) For $a \in P_\sigma$ and $B \subseteq S$,

$$aB \subseteq P_\sigma \text{ implies } aB^* \subseteq P_\sigma.$$

Proof. (c) Suppose that $w = \pi(a_0, a_1, \dots) \notin \Omega_\sigma$. We have to show that there is some index k with $a_k^\omega \notin \Omega_\sigma$. Consider the game \mathfrak{G} with a single position v belonging to Player $\bar{\sigma}$ and one a_i -labelled edge $v \rightarrow v$, for every $i < \omega$.




Since $w \notin \Omega_\sigma$, Player $\bar{\sigma}$ has a winning strategy in \mathfrak{G} by choosing in turn i the edge with label a_i . By assumption, he also has a positional winning strategy s . Let a_k be the label of the edge chosen by s . As the resulting play is winning, it follows that $a_k^\omega \notin \Omega_\sigma$. Hence, $a_k \notin P_\sigma$.

(a) If $a, b \in P_\sigma$, then $(ab)^\omega \in \Omega_\sigma$ by (c).

(b) Let $ab \in P_\sigma$. Then $a(ba)^\omega = (ab)^\omega \in \Omega_\sigma$ implies, by prefix-invariance, that $(ba)^\omega \in \Omega_\sigma$.

(d) (\Rightarrow) is trivial. For (\Leftarrow) , consider the game \mathfrak{G} with positions $V_\sigma = \{u\}$ and $V_{\bar{\sigma}} = \{v\}$. For every $a \in A$, we add an a -labelled edge $u \rightarrow v$, and for every $b \in B$, a b -labelled edge $v \rightarrow u$.



Then Player σ can win this game by playing as follows. Every time Player $\bar{\sigma}$ chooses a b -labelled edge, Player σ responds with an a -labelled edge for some $ab \in P$. By assumption, Player σ also has a positional winning strategy s . Let a be the label of the edge chosen by s . For every $b \in B$, there exists a play with labelling $(ab)^\omega$ conforming to s . Since these plays must be winning, it follows that $(ab)^\omega \in \Omega_\sigma$, i.e., $ab \in P_\sigma$, for all b .

(e) Fix $a \in S$ and set $B_0 := \{b \in S \mid ab \in P_\sigma\}$. Applying (d) to the sets

$$A := \{a^n \mid n > 0\} \quad \text{and} \quad B := \{b_0 \dots b_k \mid b_0, \dots, b_k \in B_0\},$$

we see that it is sufficient to prove that

$$a^{k+1}b_0 \dots b_k \in P_\sigma, \quad \text{for } b_0, \dots, b_k \in B_0.$$

We do so by induction on k . If $k = 0$, then $b_0 \in B_0$ implies that $a^1 b_0 \in P_\sigma$. Hence, suppose that $k > 0$. By inductive hypothesis and the fact that $b_k \in B_0$, we have $a^k b_0 \dots b_{k-1} \in P_\sigma$ and $ab_k \in P_\sigma$. Hence, $(ab_k)^\omega \in \Omega_\sigma$. Since $a(b_k a)^\omega = (ab_k)^\omega \in \Omega_\sigma$, prefix-invariance implies that $(b_k a)^\omega \in \Omega_\sigma$, i.e., $b_k a \in P_\sigma$. Consequently, it follows by (c) that $(a^k b_0 \dots b_{k-1} b_k a)^\omega \in \Omega_\sigma$ and we can again use prefix-invariance to show that

$$(a^{k+1}b_0 \dots b_k)^\omega = a(a^k b_0 \dots b_{k-1} b_k a)^\omega \in \Omega_\sigma.$$

Thus, $a^{k+1}b_0 \dots b_k \in P_\sigma$.

(f) Fix $a \in P_\sigma$ and $B \subseteq S$ with $aB \subseteq P_\sigma$. By (e) the set

$$N := \{n \geq 1 \mid a^n B^+ \subseteq \Omega_\diamond\}$$

is non-empty. Note that, if $n \in N$ then $a^n u \in P_\sigma$, for all $u \in B^+$. By (a) this implies that $aa^n u \in P_\sigma$. Consequently, $n+1 \in N$. Thus, $N = \{k, k+1, k+2, \dots\}$ for some $k < \omega$.

We claim that $k = 1$. Then $aB^+ \subseteq P_\sigma$. Since also $a \in P_\sigma$, it follows that $aB^* \subseteq P_\sigma$, as desired. Hence, it remains to prove the claim. For a contradiction, suppose that $k > 1$ and set $m := k - 1$. Then $m \notin N$, but $2m \in N$. Hence, there is some $u \in B^+$ with $a^m u \in P_{\bar{\sigma}}$. By (a) and (b), it follows that $ua^m \in P_{\bar{\sigma}}$, $a^m uua^m \in P_{\bar{\sigma}}$, and $a^{2m} uu \in P_{\bar{\sigma}}$. Hence, $2m \notin N$. A contradiction. \square

The second lemma now concludes the proof of Theorem 4.10.

Lemma 4.12. *Let \mathfrak{S} be an ω -semigroup and $\Omega \subseteq S_\omega$ a prefix-invariant set such that all games with winning condition $\langle \mathfrak{S}, \Omega \rangle$ are positionally determined. There exists a function $\Omega : S \rightarrow [2m + 1]$, for some $m < \omega$, such that*

- (1) Ω maps P_\Diamond to even numbers and P_\square to odd ones,
- (2) $\Omega(a) \leq \Omega(b)$ implies $\Omega(ab) \equiv \Omega(a) \pmod{2}$.
- (3) $\pi(a_0, a_1, \dots) \in \Omega_\Diamond$ iff $\liminf_i \Omega(a_i)$ is even.

Proof. Consider the relation $\sqsubseteq \subseteq P_\square \times P_\square$ defined by

$$a \sqsubseteq b \quad : \text{iff} \quad ac \in P_\square \Rightarrow bc \in P_\square, \quad \text{for all } c \in S.$$

We start by proving that it is a linear preorder of finite index.

Reflexivity and transitivity of \sqsubseteq follow immediately from the definition. For linearity, suppose that a and b are non-comparable. Then there are elements c and d such that

$$ac \in P_\square, \quad bc \notin P_\square, \quad ad \notin P_\square, \quad bd \in P_\square.$$

By Lemma 4.11 (a) and (b), it follows that

$$acbd \in P_\square, \quad da \in P_\Diamond, \quad cb \in P_\Diamond, \quad dacb \in P_\Diamond, \quad acbd \in P_\Diamond.$$

A contradiction.

It remains to prove that \sqsubseteq has finite index. For a contradiction, suppose otherwise. We distinguish two cases. If there exists an infinite strictly increasing chain $a_0 \sqsubset a_1 \sqsubset a_2 \sqsubset \dots$, we can fix elements $c_i \in S$ with $a_i c_i \notin P_\square$ and $a_{i+1} c_i \in P_\square$. Then Lemma 4.11 (a) and (c) implies that $c_i a_{i+1} \in P_\square$ and

$$c_0 a_1 c_1 a_2 c_2 a_3 \dots \in \Omega_\square.$$

But by the same argument as above, $a_i c_i \in P_\Diamond$ implies that

$$a_0 c_0 a_1 c_1 a_2 c_2 a_3 \dots \in \Omega_\Diamond.$$

A contradiction to prefix-invariance.

Similarly, if there exists an infinite strictly decreasing chain $a_0 \supset a_1 \supset a_2 \supset \dots$, we can fix elements $c_i \in S$ with $a_i c_i \in P_\square$ and $a_{i+1} c_i \notin P_\square$. In the same way as above it follows that

$$c_0 a_1 c_1 a_2 c_2 a_3 \dots \in \Omega_\Diamond \quad \text{and} \quad a_0 c_0 a_1 c_1 a_2 c_2 a_3 \dots \in \Omega_\square.$$

Again a contradiction.

To conclude the proof, let $B_0 \supset \dots \supset B_{m-1}$ be a decreasing enumeration of all \sqsubseteq -classes and set

$$A_i := \{c \in P_\Diamond \mid ac \in P_\square \text{ for some/all } a \in B_i\}, \quad \text{for } 0 \leq i < m.$$

In addition, we set $A_{-1} := P_\Diamond$ and $A_m := \emptyset$. Note that, by definition of \sqsubseteq , we have $P_\Diamond = A_{-1} \supseteq A_0 \supseteq \dots \supseteq A_{m-1} = A_m = \emptyset$. We claim that the function $\Omega : S \rightarrow [2m + 1]$ defined by

$$\Omega(a) := \begin{cases} 2k & \text{if } a \in A_{k-1} \setminus A_k, \\ 2k + 1 & \text{if } a \in B_k, \end{cases}$$

has the desired properties.

(1) Clearly, Ω maps each $A_k \subseteq P_\Diamond$ to an even number and each $B_k \subseteq P_\square$ to an odd one.

(2) Suppose that $\Omega(a) \leq \Omega(b)$. We distinguish four cases. If both $\Omega(a)$ and $\Omega(b)$ are even, then $a, b \in P_\Diamond$, which implies by Lemma 4.11 (a) that

$ab \in P_\Diamond$. Hence, $\Omega(ab)$ is also even. In the same way it follows that, if $\Omega(a)$ and $\Omega(b)$ are odd, then so is $\Omega(ab)$.

Suppose that $\Omega(a) = 2k + 1$ and $\Omega(b) = 2i$. Then $k < i$, $a \in B_k$ and $b \in A_{i-1} \setminus A_i \subseteq A_k$. By definition of A_k it follows that $ab \in P_\square$. Hence, $\Omega(ab)$ is odd.

Finally, suppose that $\Omega(a) = 2k$ and $\Omega(b) = 2i + 1$. Then $k \leq i$, $a \in A_{k-1} \setminus A_k$ and $b \in B_i \subseteq B_k$. By definition of A_k and the fact that $a \notin A_k$, it follows that $ba \in P_\Diamond$. Hence, Lemma 4.11 (a) implies that $ab \in P_\Diamond$ and $\Omega(ab)$ is even.

(3) Fix $a_0, a_1, \dots \in S$ and set $k := \liminf_i \Omega(a_i)$. Since Ω is prefix-invariant, we may assume w.l.o.g. that there is no i with $\Omega(a_i) < k$. Set

$$B := \Omega^{-1}(k) \quad \text{and} \quad C := \Omega^{-1}[\{k+1, \dots, 2m\}].$$

Then we can factorise the sequence $(a_i)_i$ into words $u_0, u_1, \dots \in BC^*$. Let c_i be the product of u_i . If k is even, it follows by (b) that $bc \in P_\Diamond$, for all $b \in B$ and $c \in C$. Consequently, we can use Lemma 4.11 (f) and (c) to show that $c_i \in P_\Diamond$ and $\pi(a_0, a_1, \dots) = \pi(c_0, c_1, \dots) \in \Omega_\Diamond$. If k is odd, it follows in the same way that $\pi(a_0, a_1, \dots) \in \Omega_\square$. \square

Solving Parity Games

To solve reachability games, we introduced the notion of a rank which, intuitively counts how far away from the goal we are. For parity games the situation is more complicated since we have to reach the goal not only once but repeatedly. It is possible to define ranks also for parity games. The difference is that, instead of a single ordinal, we have to use a tuple with one component for each priority k that counts how far away we are from a position of that priority. As this turns out to be a bit technical and not very enlightening, we will not do so.

Instead, we will present an algorithm for computing the winning regions of a parity game which is similar to the construction in the Theorem of Büchi and Landweber. We will prove that, for every parity game \mathfrak{G} , there exists an action $\alpha : M \times E \rightarrow M$ that turns the product game $\mathfrak{G} \times_\alpha M$ into a reachability game. Then we can use the linear time algorithm from

Section 2 to compute the winning regions. The memory M we will construct below has size $n^{\mathcal{O}(\log d)}$. As we can solve reachability games in linear time, we therefore obtain the following complexity bound.

Theorem 4.13. *The winning regions of a parity game \mathfrak{G} with n positions and d priorities can be computed in time $n^{\mathcal{O}(\log d)}$.*

The precise complexity of computing the winning regions of a parity game are still unknown. One can show that the problem belongs to the complexity class $U \cap \text{co-}U$, which means that it is probably not NP-complete. It might even belong to P, but no one has found a polynomial time algorithm so far.

Let us present the algorithm the above theorem is based on. Consider a parity game \mathfrak{G} and let p be a play of $\mathfrak{G} \times_\alpha M$. We say that p contains an *even cycle* if $p = xyz$ where y is a non-empty path starting and ending at the same position of \mathfrak{G} (the memory contents may differ) and such that the least priority seen along y is even. Note that, whether or not a given play p contains an even cycle is a reachability property: once we have found the end of the cycle, we do not need to look at the rest of p . We will design our memory M in such a way that detection of such cycles becomes easy. But first, let us show that the existence of even cycles is equivalent to winning.

Lemma 4.14. *Let \mathfrak{G} be a finite parity game and $\alpha : M \times E \rightarrow M$ an action. The following statements are equivalent.*

- (1) *Player \Diamond has a winning strategy s in \mathfrak{G} .*
- (2) *Player \Diamond has a strategy s' for $\mathfrak{G} \times_\alpha M$ such that all cycles in every play conforming to s' are even.*
- (3) *Player \Diamond has a strategy s' for $\mathfrak{G} \times_\alpha M$ such that every play conforming to s' has an even cycle.*

Proof. (2) \Rightarrow (3) is trivial.

(3) \Rightarrow (1) Suppose that Player \Diamond does not have a winning strategy for \mathfrak{G} . By determinacy, it then follows that Player \square has a positional winning strategy s in that game. This strategy induces a strategy s' for Player \square in the game $\mathfrak{G} \times_\alpha M$. Let p' be a play conforming to s' and let p be the corresponding play of \mathfrak{G} . Then p conforms to s . If p' contained an even cycle then, s being

positional, it would follow that p contained infinite repetitions of this cycle. In particular, the least priority seen infinitely often in p would be even and p would be winning for Player \Diamond . A contradiction to our choice of s .

(1) \Rightarrow (2) Let s be a winning strategy for Player \Diamond in \mathfrak{G} . W.l.o.g. we may assume that s is positional. Let s' be the strategy in $\mathfrak{G} \times_\alpha M$ induced by s . To show that it has the desired property, consider a play p' conforming to s' and let p be the corresponding play in \mathfrak{G} . Then p conforms to s and is, therefore, winning. Since s is positional, p consists of a path leading to a cycle which is repeated infinitely often. Let k be the minimal priority along this cycle. As p is winning, it follows that k is even. Hence, so is (every copy of) the cycle. \square

How can we detect an even cycle? The easiest way would be to store all the positions of \mathfrak{G} we have already seen. Once we see one of them for the second time, we have found a cycle. Unfortunately, storing that many positions requires too much memory. So instead, we resort to a counting trick.

Let \mathfrak{G} be a parity game with n positions and priorities $\{0, \dots, d-1\}$. We use the ω -semigroup \mathfrak{S} with domains $S := [d]$ and $S_\omega := [d]$ and product


$$\begin{aligned} k \cdot k' &:= \min\{k, k'\}, & \text{for } k, k' \in S, \\ k \cdot l &:= l, & \text{for } k \in S \text{ and } l \in S_\omega, \\ \pi(k_0, k_1, k_2, \dots) &:= \liminf_{i < \omega} k_i, & \text{for } k_i \in S, \end{aligned}$$

and we label an edge $u \rightarrow v$ of \mathfrak{G} by the semigroup element $\Omega(u) \in S$.

Given a finite word $w = k_0 \dots k_{n-1} \in [d]^*$, we call a sequence $z_0 < \dots < z_{m-1} < n$ an *even factorisation* of w if

- ♦ k_{z_i} is even, for all $i < m$,
- ♦ $k_j \geq \min\{k_{z_i}, k_{z_{i+1}}\}$, for all $z_i \leq j \leq z_{i+1}$, $i < m-1$,
- ♦ $k_j \geq k_{z_0}$, for $j \leq z_0$,
- ♦ $k_j \geq k_{z_{m-1}}$, for $j \geq z_{m-1}$.

We call m the *length* of the factorisation and the number $\min_i k_{z_i}$ its *value*.



Note that, if w has an even factorisation of length m and value k and w' has one of length m' and value k' , the ww' has an even factorisation of length $m + m'$ and value $\min\{k, k'\}$.

For each $l < \omega$, we will define a deterministic automaton \mathcal{A}_l over the alphabet $[d]$ that computes the length and value of an even factorisation of (some suffix of) its input. The precise definition is as follows. The set of states M_l consists of a special *accepting state* $*$ plus all triples

$$\langle s, \bar{k}, \bar{n} \rangle \in [l+1] \times [d]^* \times [l]^*,$$

where $s \leq l$ is the *size* of the state, \bar{k} is a non-decreasing sequence of priorities of length s , and \bar{n} a strictly decreasing sequence of *counters*, also of length s . To compute the cardinality of M_l , note that we can encode each state $\langle s, \bar{k}, \bar{n} \rangle$ as a word $c_0 \dots c_{l-1} \in ([d] + \square)^l$ where

$$c_i := \begin{cases} k_j & \text{if } n_j = i, \\ \square & \text{otherwise.} \end{cases} \quad \langle 3, \langle 0, 3, 4 \rangle, \langle 5, 2, 1 \rangle \rangle$$

	4	3			0
--	---	---	--	--	---

Below we will choose l such that $2^{l-1} \leq n < 2^l$. Then it follows

$$\begin{aligned} |M_l| &\leq (d+1)^l + 1 \leq (d+1)^{\log(n+1)} + 1 \\ &= (n+1)^{\log(d+1)} + 1 \in n^{\mathcal{O}(\log d)}, \end{aligned}$$

which is the right size for the theorem.

Before defining the transition relation of \mathcal{A}_l , let us state the intended behaviour of the automaton.

Lemma 4.15. Suppose that after having read a word $w \in [d]^*$ the automaton \mathcal{A}_l enters the state $\langle s, \bar{k}, \bar{n} \rangle$. Then

$$w = w' w_0 \dots w_{s-1},$$

where each w_i has an even factorisation of length at least 2^{n_i} with value k_i . In particular, the word w has a suffix with an even factorisation of length at least

$$\sum_{i < s} 2^{n_i}.$$

Before giving the proof, we have to finish the definition of \mathcal{A}_l . The initial state is the pair $\langle o, \langle \rangle, \langle \rangle \rangle$ consisting of two empty sequences. In a state $\langle s, \bar{k}, \bar{n} \rangle$ when reading the letter c , the automaton can enter the state $\langle s', \bar{k}', \bar{n}' \rangle$ if one of the following three conditions is met.

(I) $s = o$ or $k_{s-1} \leq c$,

$$s' = s, \quad \bar{k}' = \bar{k}, \quad \text{and} \quad \bar{n}' = \bar{n}.$$

(II) There is some $o < i < s$ such that $k_{i-1} < c < k_i$,

$$s' = i + 1, \quad \bar{k}' = \langle k_o, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_o, \dots, n_{i-1}, n_i \rangle.$$

(III) There is some $o < i < s$ such that $k_{i-1} \leq c$, the priorities k_i, \dots, k_{s-1} are even, $\langle n_i, \dots, n_{s-1} \rangle = \langle s - i - 1, s - i - 2, \dots, 1, o \rangle$,

$$s' = i + 1, \quad \bar{k}' = \langle k_o, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_o, \dots, n_{i-1}, n_i + 1 \rangle.$$

(IV) There is some $o < i < s$ such that $k_{i-1} \leq c$, the priority k_i is odd, while k_{i+1}, \dots, k_{s-1} are even,

$$\langle n_i, \dots, n_{s-1} \rangle = \langle s - i - 1, s - i - 2, \dots, 1, o \rangle,$$

$$s' = i + 1, \quad \bar{k}' = \langle k_o, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_o, \dots, n_{i-1}, n_i + 1 \rangle.$$

(V) $s = l$, each priority k_i is even, c is even, and the next state is $*$.

(VI) Once \mathcal{A}_l has reached the state $*$, it remains there.

If there are several possible transitions, we choose the one that leads to a state of minimal length.

Proof of Lemma 4.15. We prove the claim by induction on the length of w . For $w = \langle \rangle$, \mathcal{A}_l is in the initial state $\langle o, \langle \rangle, \langle \rangle \rangle$ and the claim is trivial.

For the inductive step, suppose that the input is $w c$ with $w \in [d]^*$ and $c \in [d]$, and let $\langle s, \bar{k}, \bar{c} \rangle$ be the state after reading w . By inductive hypothesis, w has a suffix of the form $w_0 \dots w_{s-1}$ where each w_i has an even factorisation of length at least 2^{n_i} with value k_i . We distinguish several cases, depending on which transition the automaton takes while reading the last letter c .

If the last transition is of the form (I), we can obtain the desired suffix $w'_0 \dots w'_{s-1}$ of $w c$ by setting $w'_i := w_i$, for $i < s - 1$, and $w'_{s-1} := w_{s-1} c$.

If the last transition is of the form (II), let i be the index such that $k_{i-1} < c < k_i$. We set $w'_j := w_j$, for $j < i$, and $w'_i := w_i \dots w_{s-1} c$. Then w'_i has an even factorisation of length at least $2^{n_i} + \dots + 2^{n_{s-1}} + 1 \geq 2^{n_i}$.

If the last transition is of the form (III) or (IV), let i be the index from the above definition. We set $w'_j := w_j$, for $j < i$, and $w'_i := w_i \dots w_{m-1} c$. Then w'_i has an even factorisation of length at least

$$2^{n_i} + \dots + 2^{n_{s-1}} + 1 = 2^{s-i-1} + 2^{s-i-2} + \dots + 2^1 + 2^0 + 1 = 2^{s-i}. \quad \square$$

With the help of Lemma 4.15, we are able to show that the information contained in the states of \mathcal{A}_l is sufficient to detect whether the input contains an even cycle.

Lemma 4.16. Let p be an infinite play in a parity game \mathfrak{G} with n positions, and let $w = (c_i)_{i < \omega}$ be the sequence of priorities along p . If $l > \log n$ and \mathcal{A}_l accepts w , then p contains an even cycle.

Proof. Let $\langle l, \bar{k}, \bar{n} \rangle \in M_l$ be the last state in the run of \mathcal{A}_l before it enters the state $*$, let w' be the prefix of w leading to this state, and let c be the next input letter. By Lemma 4.15, the word w' has a suffix with an even factorisation of length at least

$$2^{l-1} + \dots + 2^0 = 2^l - 1 \geq n,$$

while $w'c$ has a suffix with an even factorisation $z_0 < \dots < z_{m-1}$ of length $m \geq 2^l - 1 + 1 > n$. By the Pigeon Hole Principle it follows that, after reading the additional letter c , there must be two positions $z_i < z_j$ in the factorisation that correspond to the same vertex v of \mathfrak{G} . Let p_o be the part of p corresponding to the path between these two positions. To see that p_o is an even cycle, note that the minimal priority seen along the closed path is the minimal value of c_{z_i}, \dots, c_{z_j} . In particular, it is even. Hence, the play is winning. \square

Lemma 4.17. *The automaton \mathcal{A}_l accepts every word w satisfying the parity condition.*

Proof. By induction on l we will prove that, starting from an arbitrary state $\langle s, \bar{k}, \bar{n} \rangle$ the automaton \mathcal{A}_l accepts every infinite word w that satisfies the parity condition. Hence, fix l and a run ρ of \mathcal{A}_l on w . If ρ contains the accepting state $*$, we are done. Hence, suppose otherwise.

If every state of size l appears only finitely often in ρ , some suffix of ρ is a run of \mathcal{A}_{l-1} . By inductive hypothesis, this suffix is accepting. Hence, so is ρ .

Consequently, we may assume that some state $\langle l, \bar{k}', \bar{n}' \rangle$ of size l appears infinitely often in ρ . The way the transitions are defined it follows that, after the first appearance of $\langle l, \bar{k}', \bar{n}' \rangle$ every state $\langle s, \bar{k}'', \bar{n}'' \rangle$ in ρ satisfies $k''_o = k'_o$. Let ρ' be the sequence of states obtained from ρ by (i) removing the part before the first appearance of $\langle l, \bar{k}', \bar{n}' \rangle$ and (ii) removing the first components of all remaining states, i.e., replacing $\langle s, \bar{k}'', \bar{n}'' \rangle$ by

$$\langle s - 1, \langle k''_1, \dots, k''_{s-1} \rangle, \langle n''_1, \dots, n''_{s-1} \rangle \rangle.$$

Then ρ' is a run of \mathcal{A}_{l-1} on the corresponding suffix of w . Again it follows by inductive hypothesis that this run is accepting. Hence, so is ρ . \square

Proof of Theorem 4.13. Set $l := \lfloor \log n \rfloor + 1$ and let $\alpha : M_l \times E \rightarrow M_l$ be the action induced by \mathcal{A}_l . We claim that a position v of \mathfrak{G} belongs to the winning region of Player \Diamond in \mathfrak{G} if, and only if, Player \Diamond has a strategy in the game $\mathfrak{G} \times_\alpha M_l$ from the position $\langle v, \langle o, \langle \rangle, \langle \rangle \rangle$ to reach some position of the form $\langle u, * \rangle$. Since the latter game is a reachability game of size

$n \times \mathcal{O}(n^{\log n}) = \mathcal{O}(n^{\log n})$ and we can compute its winning regions in linear time, the theorem follows. Hence, it remains to prove the claim.

(\Leftarrow) follows immediately by Lemma 4.16 and the implication $(2) \Rightarrow (1)$ in Lemma 4.14, while (\Rightarrow) follows by Lemma 4.17 and the implication $(1) \Rightarrow (3)$ in Lemma 4.14. \square

5 The Modal μ -Calculus

We have seen above that the model-checking game for modal logic is a reachability game. There also exists a logic whose model-checking game is a parity game. As with every modal logic, this logic does not talk about arbitrary structures, but only about transition systems.

Definition 5.1. A *transition system* is a structure of the form

$$\mathfrak{S} = \langle S, (E_a)_{a \in A}, (P_c)_{c \in C} \rangle,$$

where the E_a are binary relations and the P_c are unary ones.

The logic we are interested in is obtained from basic modal logic by adding a fixed-point operator.

Definition 5.2. Let A and C be two sets of labels and V a set of propositional variables. The *modal μ -calculus* L_μ is the logic with formulae of the form

- ♦ *atomic propositions* P_c , for $c \in C$,
- ♦ *variables* $X \in V$,
- ♦ *boolean operations* $\varphi \vee \psi$, $\varphi \wedge \psi$, and $\neg \varphi$, for $\varphi, \psi \in L_\mu$,
- ♦ *modal operators* $\langle a \rangle \varphi$ and $[a] \varphi$, for $\varphi \in L_\mu$ and $a \in A$,
- ♦ *fixed point operators* $\mu X \varphi$ and $\nu X \varphi$, for a variable $X \in V$ and a formula $\varphi \in L_\mu$ where every occurrence of X is under an even number of negations.

For a transition system $\mathfrak{S} = \langle S, (E_a)_{a \in A}, (P_c)_{c \in C} \rangle$, a state $s \in S$, a formula $\varphi \in L_\mu$, and a variable assignment $\beta : V \rightarrow \mathcal{P}(S)$, we define the

satisfaction relation $\mathfrak{S}, s \models \varphi[\beta]$ by induction on φ as follows.

$$\begin{aligned}
\mathfrak{S}, s \models P_a[\beta] & : \text{iff } s \in P_a, \\
\mathfrak{S}, s \models X[\beta] & : \text{iff } s \in \beta(X), \\
\mathfrak{S}, s \models (\varphi \vee \psi)[\beta] & : \text{iff } \mathfrak{S}, s \models \varphi[\beta] \text{ or } \mathfrak{S}, s \models \psi[\beta], \\
\mathfrak{S}, s \models (\varphi \wedge \psi)[\beta] & : \text{iff } \mathfrak{S}, s \models \varphi[\beta] \text{ and } \mathfrak{S}, s \models \psi[\beta], \\
\mathfrak{S}, s \models (\neg\varphi)[\beta] & : \text{iff } \mathfrak{S}, s \not\models \varphi[\beta], \\
\mathfrak{S}, s \models (\langle a \rangle \varphi)[\beta] & : \text{iff there is some edge } \langle s, t \rangle \in E_a \text{ such} \\
& \text{that } \mathfrak{S}, t \models \varphi[\beta], \\
\mathfrak{S}, s \models ([a]\varphi)[\beta] & : \text{iff for every edge } \langle s, t \rangle \in E_a \text{ we have} \\
& \mathfrak{S}, t \models \varphi[\beta], \\
\mathfrak{S}, s \models (\mu X.\varphi)[\beta] & : \text{iff } s \text{ belongs to the least fixed-point of the} \\
& \text{operation } F_\varphi \text{ below,} \\
\mathfrak{S}, s \models (\nu X.\varphi)[\beta] & : \text{iff } s \text{ belongs to the greatest fixed-point of} \\
& \text{the operation } F_\varphi \text{ below,}
\end{aligned}$$

where the function $F_\varphi : \wp(S) \rightarrow \wp(S)$ in the last two lines is defined by

$$F_\varphi(U) := \{s \in S \mid \mathfrak{S}, s \models \varphi[\beta[X \mapsto U]]\}.$$

Here, the variable assignment $\beta[X \mapsto U]$ is given by

$$Y \mapsto \begin{cases} U & \text{if } Y = X, \\ \beta(Y) & \text{otherwise.} \end{cases}$$

If $|A| = 1$, we usually simplify notation by writing $\Diamond\varphi$ and $\Box\varphi$ without the edge label. Furthermore, if the formula φ has no free variables, we drop the variable assignment β from the notation and simply write $\mathfrak{S}, s \models \varphi$.

Remark. Note that the requirement on X occurring only positively in φ ensures that the function F_φ is monotone. Hence, the least and the greatest fixed-point do exist.

Examples. The formula

$$\mu X[P \vee \Diamond X]$$

states that there exists a path from the current state to some state in P .

$$\mu X\Box X$$

states that there is no infinite path starting at the current state.

$$\nu X[P \wedge \Diamond X]$$

states that there exists an infinite path from the current state where every visited state belongs to P .

$$\nu X\mu Y[\Diamond Y \vee (P \wedge \Diamond X)]$$

states that there is an infinite path containing infinitely many states in P .

Exercise 5.1. Show that, for every L_μ -formula $\varphi(\tilde{X})$, there exists an MSO-formula $\varphi^*(x, \tilde{X})$ such that

$$\mathfrak{S}, s \models \varphi(\tilde{P}) \quad \text{iff} \quad \mathfrak{S} \models \varphi^*(s, \tilde{P}),$$

for all transition systems \mathfrak{S} , states $s \in S$, and predicates \tilde{P} .

Next let us introduce the model-checking game for L_μ . As usual, Player \Diamond tries to prove that the formula holds in the given state, while Player \Box tries to prove that it does not.

Definition 5.3. Let $\mathfrak{S} = \langle S, (E_a)_{a \in A}, (P_c)_{c \in C} \rangle$ be a transition system and φ an L_μ -formula in negation normal form. The *model-checking game* $\mathcal{G}(\mathfrak{S}, \varphi)$ is the following parity game. As positions of \mathfrak{G} we use the pairs $\langle s, \psi \rangle \in S \times \Phi$ where Φ is the set of all subformulae of φ . The set V_\Diamond of positions for Player \Diamond consists of all pairs $\langle s, \psi \rangle$ where an existential choice has to be made to satisfy ψ , that is, where

- $\psi = P_c$ is atomic and $s \notin P_c$,

- ♦ $\psi = \neg P_c$ and $s \in P_c^\complement$,
- ♦ $\psi = \vartheta \vee \vartheta'$,
- ♦ $\psi = \langle a \rangle \vartheta$, or
- ♦ $\psi = \mu X \vartheta$.

The other positions belong to Player \square . The edge relation is defined as follows.

$$\begin{aligned}
 \langle s, \psi_o \vee \psi_i \rangle &\rightarrow \langle s, \psi_i \rangle, & \text{for } i = o, i, \\
 \langle s, \psi_o \wedge \psi_i \rangle &\rightarrow \langle s, \psi_i \rangle, & \text{for } i = o, i, \\
 \langle s, \langle a \rangle \psi \rangle &\rightarrow \langle t, \psi \rangle, & \text{for every transition } s \xrightarrow{a} t, \\
 \langle s, [a] \psi \rangle &\rightarrow \langle t, \psi \rangle, & \text{for every transition } s \xrightarrow{a} t, \\
 \langle s, \mu X \psi \rangle &\rightarrow \langle s, \psi \rangle, \\
 \langle s, \nu X \psi \rangle &\rightarrow \langle s, \psi \rangle, \\
 \langle s, X \rangle &\rightarrow \langle s, \psi \rangle, & \text{where } \psi \text{ is the definition of } X,
 \end{aligned}$$

where the *definition* of a fixed-point variable X is the formula ψ that appears as the body in the fixed-point formula $\mu X \psi$ or $\nu X \psi$ binding X .

Finally, the priorities are as follows. For a subformula of the form $\mu X \psi$ or $\nu X \psi$ that occurs inside of k other fixed-point operators, we set

$$\Omega(\langle s, \mu X \psi \rangle) := 2k + 1 \quad \text{and} \quad \Omega(\langle s, \nu X \psi \rangle) := 2k.$$


All other other priorities are larger than those.

Remark. Note that every cycle in $\mathcal{G}(\mathfrak{S}, \varphi)$ contains a position of the form $\langle s, \mu X \psi \rangle$ or $\langle s, \nu X \psi \rangle$. Since the priorities of such positions are smaller than all other priorities, these are the only positions that matter when determining the winner.

Example. Given the following transition system and formula

$$\mathfrak{S} = \begin{array}{c} \textcircled{s} \longrightarrow \textcircled{t} \\ \textcircled{s} \longrightarrow \textcircled{t} \end{array} P \quad \varphi = \mu X (P \vee \Diamond X)$$

we construct the game



where the priorities are the numbers in the circles. (All other priorities are larger than 1.)

Let us prove that this game has the desired properties.

Theorem 5.4. *In the game $\mathcal{G}(\mathfrak{S}, \varphi)$, Player \Diamond has a winning strategy from a position $\langle t, \varphi \rangle$ if, and only if, $\mathfrak{S}, t \models \varphi$.*

Proof. To prove the statement by induction on φ , we have to deal with formulae φ that have free variables. Hence, suppose that φ has free variables X_0, \dots, X_{n-1} and let \tilde{P} be a corresponding tuple of subsets $P_i \subseteq S$. We define a variant $\mathcal{G}(\mathfrak{S}, \varphi, \tilde{P})$ of the model-checking game in the same way as above where the variables X_i are treated as propositions with value P_i . That is, positions of the form $\langle u, X_i \rangle$ are considered to be winning for Player \Diamond if, and only if, $u \in P_i$.

For this more general version of the model-checking game we can now prove by induction on the formula $\varphi(\tilde{X})$ that Player \Diamond wins $\mathcal{G}(\mathfrak{S}, \varphi, \tilde{P})$ with starting position $\langle t, \varphi \rangle$ if, and only if, $\mathfrak{S}, t \models \varphi(\tilde{P})$.

If φ is an atomic formula or a negated atomic formula, the claim follows immediately from the definition of $\mathcal{G}(\mathfrak{S}, \varphi, \tilde{P})$. If φ starts with a boolean operation or a modal operator, the claim follows by inductive hypothesis. Hence, it remains to consider the case where $\varphi = \mu Y \psi(\tilde{X}, Y)$ or $\varphi = \nu Y \psi(\tilde{X}, Y)$.

First, suppose that φ consists of a least fixed-point. We have to prove two directions.

(\Leftarrow) Let $Q_\alpha := (F_\psi)^\alpha(\emptyset)$, for $\alpha < \kappa$, be the α -th stage of the corresponding fixed-point induction. Since $\mathfrak{S}, t \models \mu Y \psi(\bar{P}, Y)$, there is some ordinal α such that $t \in Q_{\alpha+1}$. We construct the desired winning strategy by induction on α . By the inductive hypothesis for ψ , we have a winning strategy s_* for the game $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q_\alpha)$ with starting position $\langle t, \psi \rangle$, and by the inductive hypothesis for α there exists, for every $u \in Q_\alpha$, a winning strategy s_u for the game $\mathcal{G}(\mathfrak{S}, \varphi, \bar{P})$ with starting position $\langle u, \mu Y \psi \rangle$. We combine these strategies into a single one as follows. Player \Diamond starts by following s_* until he reaches a position of the form $\langle u, Y \rangle$ with $u \in Q_\alpha$. Then he switches to the strategy s_u and follows it until the end of the game. By choice of s_* and s_u , this combined strategy is winning for the starting position $\langle t, \mu Y \psi \rangle$.

(\Rightarrow) Let $Q' \subseteq S$ be the set of all states t such that Player \Diamond has a winning strategy in the game $\mathcal{G}(\mathfrak{S}, \mu Y \psi, \bar{P})$ with starting position $\langle t, \mu Y \psi \rangle$. Furthermore, we inductively define sets $Q_\alpha \subseteq S$ as follows.

$$Q_0 := \emptyset \quad \text{and} \quad Q_\delta := \bigcup_{\alpha < \delta} Q_\alpha \quad \text{for limit ordinals } \delta.$$

For the successor step, let $Q_{\alpha+1}$ be the set of all states t such that Player \Diamond has a winning strategy in the game $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q_\alpha)$ with starting position $\langle t, \psi \rangle$. Since every position of the form $\langle t, \mu Y \psi \rangle$ has priority 1 and the game $\mathcal{G}(\mathfrak{S}, \mu Y \psi, \bar{P})$ has no positions with priority 0, it follows that

$$Q' = \bigcup_{\alpha} Q_\alpha.$$

Furthermore, it follows by inductive hypothesis that

$$\mathfrak{S}, t \models \psi(\bar{P}, Q_\alpha), \quad \text{for all } t \in Q_{\alpha+1}.$$

Consequently, the union $\bigcup_{\alpha} Q_\alpha = Q'$ is contained in the least fixed-point of the operator F_ψ . Hence,

$$\mathfrak{S}, t \models \mu Y \psi(\bar{P}, Y), \quad \text{for all } t \in Q'.$$

It remains to consider the case where $\varphi = \nu X \psi(\bar{X}, Y)$ is a greatest fixed-point. Again we have to prove two directions.

(\Leftarrow) Let $Q \subseteq S$ be the greatest fixed-point of F_ψ . By inductive hypothesis, there exists a strategy s for Player \Diamond in the game $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q)$ that is winning for every starting position $\langle t, \psi \rangle$ with $t \in Q$. We claim that the same strategy is also winning in the game $\mathcal{G}(\mathfrak{S}, \nu Y \psi, \bar{P})$ with starting position $\langle t, \nu Y \psi \rangle$, for $t \in Q$. Hence, let p be a play conforming to s starting in $\langle t, \nu Y \psi \rangle$ with $t \in Q$. If p contains infinitely many positions of the form $\langle u, \nu Y \psi \rangle$, it is winning since these positions have priority 0. Otherwise, the suffix of p after these positions is a play in $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q)$ and, therefore, also winning.

(\Rightarrow) Let $Q \subseteq S$ be the set of all states t such that Player \Diamond has a winning strategy for the starting position $\langle t, \nu Y \psi \rangle$. We claim that

$$\mathfrak{S}, t \models \psi(\bar{P}, Q), \quad \text{for all } t \in Q.$$

Then Q is contained in the greatest fixed-point of F_ψ , which implies that

$$\mathfrak{S}, t \models \nu Y \psi(\bar{P}, Y), \quad \text{for all } t \in Q,$$

as desired. For the proof, fix $t \in Q$. By assumption, there exists a winning strategy s for Player \Diamond starting in $\langle t, \nu Y \psi \rangle$. We claim that s is also a winning strategy in the game $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q)$ with starting position $\langle t, \psi \rangle$. Hence, let p be a play conforming to s in that game. If p is also a play in the original game $\mathcal{G}(\mathfrak{S}, \nu Y \psi, \bar{P})$, it is winning for Player \Diamond by choice of s . Suppose otherwise. Since the game $\mathcal{G}(\mathfrak{S}, \psi, \bar{P}Q)$ is obtained from $\mathcal{G}(\mathfrak{S}, \nu Y \psi, \bar{P})$ by deleting some edges, it follows that p is a partial play in the game $\mathcal{G}(\mathfrak{S}, \nu Y \psi, \bar{P})$ which ends in a position of the form $\langle u, Y \rangle$. By choice of s and Q , no partial play conforming to s can lead to such a position where $u \notin Q$. Consequently, we have $u \in Q$, which implies that the position $\langle u, Y \rangle$ is winning for Player \Diamond . \square

Intuitively, this theorem states that parity games are as expressive as the μ -calculus. Conversely, we can show that the winning condition for parity games can be expressed in L_μ . Before giving the proof, let us take a look at the simpler case of reachability games.

Example. Let \mathfrak{G} be a reachability game. The formula

$$\psi := (\bigvee_{\sigma} \Diamond X) \vee (\bigvee_{\bar{\sigma}} \Box X)$$

expresses the step function Step_σ . As the winning region for Player σ is the least fixed-point of this function, we can define it by the formula

$$\varphi := \mu X[(V_\sigma \wedge \Diamond X) \vee (V_{\bar{\sigma}} \wedge \Box X)].$$

Definition 5.5. Let $\mathfrak{G} = \langle V_\Diamond, V_\Box, E, \Omega \rangle$ and $\mathfrak{G}' = \langle V'_\Diamond, V'_\Box, E', \Omega' \rangle$ be two parity games. An *immersion* $\iota : \mathfrak{G} \rightarrow \mathfrak{G}'$ is a function $\iota : V \rightarrow V'$ with the following property.

For each Player σ , each position $v \in V_\sigma$, and each successor u of v , Player σ has a strategy s in the game \mathfrak{G}' with starting position $\iota(v)$ ensuring that every play conforming to s either

- is winning and does not contain any position in $\text{rng } \iota$ (except the one before the first move of course), or
- the first such position is equal to $\iota(u)$ and the least priority seen between $\iota(v)$ and $\iota(u)$ (inclusive) is equal to $\min\{\Omega(v), \Omega(u)\}$.

Example. A simple kind of immersion is a *homomorphism* of games, that is, a function $h : V \rightarrow V'$ such that

$$h[V_\Diamond] \subseteq V'_\Diamond, \quad h[V_\Box] \subseteq V'_\Box, \quad h[E] \subseteq E',$$

$$\Omega(h(v)) = \Omega(v), \quad \text{for all } v \in V,$$

and such that h is *locally surjective* in the sense that, for every $v \in V$, each successor of $h(v)$ belongs to $\text{rng } h$.

Lemma 5.6. Let $\iota : \mathfrak{G} \rightarrow \mathfrak{G}'$ be an immersion between two parity games $\mathfrak{G} = \langle V_\Diamond, V_\Box, E, \Omega \rangle$ and $\mathfrak{G}' = \langle V'_\Diamond, V'_\Box, E', \Omega' \rangle$ and let $W_\Diamond, W_\Box, W'_\Diamond, W'_\Box$ be the winning regions of the respective games. Then

$$\iota[W_\Diamond] \subseteq W'_\Diamond \quad \text{and} \quad \iota[W_\Box] \subseteq W'_\Box.$$

Proof. Let $v \in W_\sigma$ and let s be a corresponding winning strategy. We construct a winning strategy s' for Player σ in the game \mathfrak{G}' with starting position $\iota(v)$ as follows. For every vertex $w \in V_\sigma$, let t_w be a strategy in \mathfrak{G}' with starting position $\iota(w)$ as in the definition of an immersion. We construct s' by

combining these strategies. That is, when in a position of the form $\iota(w)$, Player σ follows t_w until the play is either won or it reaches a position of the form $\iota(u)$, for some successor u of w . Then the game continues using t_u and so on.

We claim that the resulting strategy s' is winning. Fix a play p' conforming to s' . If p' contains only finitely many positions in $\text{rng } \iota$, let w be the last of them. Then a suffix of p' conforms to the corresponding strategy t_w and is, therefore, winning for Player σ . Otherwise, p' induces a play $p = (v_i)_{i < \omega}$ in \mathfrak{G} and p' has a factorisation $p' = p'_0 p'_1 \dots$ such that the least priority seen along p'_i is equal to $\min\{\Omega(v_i), \Omega(v_{i+1})\}$. Since the least priority seen infinitely often in p is winning for Player σ , so is the least one in p' . \square

Theorem 5.7. Let \mathfrak{G} be a parity game that only uses priorities from the set $[2k]$. The winning region for Player \Diamond is defined by the formula

$$vX_\Diamond \mu X_1 \dots vX_{2k-2} \mu X_{2k-1} \bigwedge_{i < 2k} [P_i \rightarrow [(V_\Diamond \wedge \Diamond X_i) \vee (V_\Box \wedge \Box X_i)]],$$

where $P_i := \Omega^{-1}(i)$ is the set of positions of priority i .

Proof. For $j \leq 2k$, we set

$$\begin{aligned} \psi &:= \bigwedge_{i < 2k} [P_i \rightarrow [(V_\Diamond \wedge \Diamond X_i) \vee (V_\Box \wedge \Box X_i)]], \\ \varphi_j &:= \begin{cases} \psi & \text{if } j = 2k, \\ \mu X_j \varphi_{j+1} & \text{if } j < 2k \text{ is odd}, \\ vX_j \varphi_{j+1} & \text{if } j < 2k \text{ is even}. \end{cases} \end{aligned}$$

Let $\iota : \mathfrak{G} \rightarrow \mathcal{G}(\mathfrak{G}, \varphi)$ be the function mapping a position v of \mathfrak{G} to the position $\langle v, \varphi_{\Omega(v)} \rangle$. We claim that ι is an immersion. By Lemma 5.6 and Theorem 5.4, it then follows that

- Player \Diamond wins \mathfrak{G} with starting position v
- iff Player \Diamond wins \mathfrak{G} with starting position $\iota(v)$
- iff $\mathfrak{G}, v \models \varphi_{\Omega(v)}(\bar{P})$
- iff $\mathfrak{G}, v \models \varphi$,

where P_j is the set defined by φ_j .

Hence, it remains to prove the above claim. Let $v \in V$ be a position of \mathfrak{G} . Starting from position $\langle v, \varphi_i \rangle$ in the game $\mathcal{G}(\mathfrak{G}, \varphi)$ (ignoring positions with a unique outgoing edge), Player \square has to choose an index i in the conjunction and Player \diamond has to reply with a choice between $\neg P_i$, $V_\diamond \wedge \diamond X_i$, and $V_\square \wedge \square X_i$. Finally, one of the two players has to pick a successor of v for $\diamond X_i$ or $\square X_i$. If they do not want to lose immediately, Player \square has to pick $i := \Omega(v)$ and Player \diamond has to pick one of the two latter formulae depending on which player the vertex v belongs to. Thus, these two choices are forced, which means that the only relevant choices for Players \diamond and \square are picking the successor for the formulae $\diamond X_i$ or $\square X_i$.

To show that ι is an immersion, let u be a successor of v . By the remarks above, if $v \in V_\diamond$, Player \square has to pick $i := \Omega(v)$ and Player \diamond has to reply with the subformula $V_\diamond \wedge \diamond X_i$. Then Player \square has to choose the subformula $\diamond X_i$, after which Player \diamond can pick $\iota(u)$, as desired.

Similarly, if $v \in V_\square$, Player \square picks $i := \Omega(v)$, Player \diamond replies with $V_\square \wedge \square X_i$. Then \square chooses $\square X_i$ followed by the position $\iota(u)$.

Thus, in both cases the respective player has a strategy to reach $\iota(u)$ (or to win the game immediately). Furthermore, the priorities seen between $\iota(v)$ and $\iota(u)$ are $\Omega(v)$, $\Omega(v) + 1, \dots, 2k, \Omega(u)$. \square

Alternating Parity Games

The aim of this section is to derive a special version of the formula defining the winning regions of a parity game where the players take turns alternatingly. This will come in handy in the next chapter.

Definition 5.8. (a) A parity game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega \rangle$ is *alternating* if

$$u \in V_\diamond \Leftrightarrow v \in V_\square, \quad \text{for every edge } \langle u, v \rangle \in E.$$

(b) An alternating parity game is *normalised* if all positions of Player \square have maximal priority, while all positions of Player \diamond have some priority that is not maximal. \dashv

First, let us prove that we can normalise every alternating game.

Lemma 5.9. For every alternating parity game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega \rangle$, there exists a normalised alternative parity game $\mathfrak{G}' = \langle V'_\diamond, V'_\square, E', \Omega' \rangle$ and a surjective function $\rho : V' \rightarrow V$ such that Player \diamond has a winning strategy starting at a position v' in \mathfrak{G}' if, and only if, he has a winning strategy starting at $\rho(v')$ in \mathfrak{G} .

Proof. To be able to remove the priorities of V_\square -positions, we will store them in the V_\diamond -positions. Choose some number m larger than all priorities of \mathfrak{G} . To define \mathfrak{G}' , we set

$$\begin{aligned} V'_\diamond &:= V_\diamond \times [m], \\ V'_\square &:= V_\square, \\ E' &:= \left\{ \langle \langle u, k \rangle, v \rangle \mid \langle u, v \rangle \in E \cap (V_\diamond \times V_\square) \right\} \\ &\quad \cup \left\{ \langle u, \langle v, \Omega(u) \rangle \rangle \mid \langle u, v \rangle \in E \cap (V_\square \times V_\diamond) \right\}, \\ \Omega'(\langle v, k \rangle) &:= \min\{\Omega(v), k\}, \quad \text{for } \langle v, k \rangle \in V'_\diamond, \\ \Omega'(v) &:= m, \quad \text{for } v \in V'_\square. \end{aligned}$$

Let $\rho : V' \rightarrow V$ be the function that maps positions $v \in V'_\square = V_\square$ to itself and positions $\langle v, k \rangle \in V'_\diamond$ to the corresponding V_\diamond -position v . We claim that ρ is the desired surjection.

Let σ' be a positional winning strategy for Player \diamond in \mathfrak{G}' . We define a (non-positional) strategy σ in \mathfrak{G} by setting

$$\sigma(v_0, \dots, v_n) := \begin{cases} \sigma'(\langle v_n, \Omega(v_{n-1}) \rangle) & \text{if } n > 0, \\ \sigma'(\langle v_n, m-1 \rangle) & \text{if } n = 0. \end{cases}$$

If $p = (v_n)_{n < \omega}$ is a play in \mathfrak{G} conforming to σ , there is a unique play $p' = (v'_n)_{n < \omega}$ in \mathfrak{G}' such that $\rho(v'_n) = v_n$, for all n . This play conforms to σ' and is therefore winning for Player \diamond . By definition of ρ , it follows that the play p' has the form

$$\langle v_0, k \rangle, v_1, \langle v_2, \Omega(v_1) \rangle, v_3, \langle v_4, \Omega(v_3) \rangle, \dots, \quad \text{for some } k < m.$$

(For simplicity, we have assumed that the play starts in V'_\Diamond . Otherwise, we have to remove the first entry.) The corresponding sequence of priorities is

$$\min\{\Omega(v_o), k\}, m, \min\{\Omega(v_1), \Omega(v_2)\}, \\ m, \min\{\Omega(v_3), \Omega(v_4)\}, \dots$$

As p satisfies the parity condition, it follows that so does p' .

Conversely, consider a positional winning strategy σ for Player \Diamond in \mathfrak{G} . We define a positional strategy σ' in \mathfrak{G}' by setting

$$\sigma'(\langle v, k \rangle) := \sigma(v), \quad \text{for } \langle v, k \rangle \in V'_\Diamond, \\ \sigma'(v) := \langle \sigma(v), \Omega(v) \rangle, \quad \text{for } v \in V'_\Box.$$

Let $p' = (v'_n)_{n < \omega}$ be a play in \mathfrak{G}' following σ' . Then $p := (\rho(v'_n))_{n < \omega}$ is a play in \mathfrak{G} following σ . As above it follows that the fact that p satisfies the parity condition implies that so does p' . \square

For normalised alternating games we can compute the winning regions in a different way by taking two steps at a time.

Proposition 5.10. *Let \mathfrak{G} be a normalised alternating parity game with maximal priority m and let k be some constant with $m < 2k$. We can define the set $W_\Diamond \cap V_\Diamond$ of all positions for Player \Diamond where he has a winning strategy by the formula*

$$vX_o\mu X_1\cdots vX_{2k-2}\mu X_{2k-1}\bigwedge_{i < m} [P_i \rightarrow \Diamond\Box X_i],$$

where $P_i := \Omega^{-1}(i)$ is the set of positions of priority i .

Proof. We have shown in Theorem 5.7 that the formula

$$vX_o\mu X_1\cdots vX_{2k-2}\mu X_{2k-1}\bigwedge_{i < 2k} [P_i \rightarrow [(V_\Diamond \wedge \Diamond X_i) \vee (V_\Box \wedge \Box X_i)]]$$

defines W_\Diamond . Since \mathfrak{G} is normalised we have $P_m \subseteq V_\Box$ and $P_i \subseteq V_\Diamond$, for $i < m$. Hence, the above formula simplifies to

$$\varphi := vX_o\mu X_1\cdots vX_{2k-2}\mu X_{2k-1}\left[\bigwedge_{i < m} [P_i \rightarrow \Diamond\Box X_i] \wedge [P_m \rightarrow \Box X_m]\right].$$

We claim that, for positions $v \in V_\Diamond$, this formula is equivalent to

$$\psi := vX_o\mu X_1\cdots vX_{2k-2}\mu X_{2k-1}\bigwedge_{i < m} [P_i \rightarrow \Diamond\Box X_i].$$

By Theorem 5.4, it is sufficient to show that, for every $v \in V_\Diamond$, the position $\langle v, \varphi \rangle$ is winning for Player \Diamond in $\mathcal{G}(\mathfrak{G}, \varphi)$ if, and only if, the position $\langle v, \psi \rangle$ is winning in $\mathcal{G}(\mathfrak{G}, \psi)$.

(\Rightarrow) Let s be a winning strategy for Player \Diamond in $\mathcal{G}(\mathfrak{G}, \varphi)$ with starting position $\langle v, \varphi \rangle$. We define a strategy s' in $\mathcal{G}(\mathfrak{G}, \psi)$ as follows. In this game there are two non-trivial choices for Player \Diamond : (I) in a position of the form $\langle u, P_i \rightarrow \Diamond\Box X_i \rangle$, he has to choose between the formulae $\neg P_i$ and $\Diamond\Box X_i$; and (II) in a position of the form $\langle u, \Diamond\Box X_i \rangle$, he has to choose a successor of u .

Choice (I) is easy: if $u \in P_i$, Player \Diamond has to pick $\langle u, \Diamond\Box X_i \rangle$, otherwise he chooses $\langle u, \neg P_i \rangle$. For Choice (II), Player \Diamond follows the strategy s , i.e., he chooses the successor w of u such that

$$s(\langle u, \Diamond X_i \rangle) = \langle w, X_i \rangle.$$

We claim that the resulting strategy s' is winning. Hence, let p' be a play in $\mathcal{G}(\mathfrak{G}, \psi)$ conforming to s' with starting position $\langle v, \psi \rangle$. By construction of s' , this play corresponds to some play p in $\mathcal{G}(\mathfrak{G}, \varphi)$ conforming to s with starting position $\langle v, \varphi \rangle$ such that the sequences of first components of the positions in p and those in p' coincide. ('Coincide' here means that the two sequences contain the same positions of \mathfrak{G} in the same order, but the number of times each position appears may be different.) Note that the least priority seen infinitely often along p is determined by the positions of p containing a fixed-point formula. The same holds for p' . Comparing φ and ψ , we see that this sequence of fixed-point formulae for p contains twice as many entries, while that for p' omits every second one. But the omitted entries correspond to the fixed-point associated with the variable X_m , whose priorities are larger than all other ones. Hence, the least priority seen infinitely often is the same in both sequences. Since p conforms to s , it is winning. Hence, so is p' .

(\Leftarrow) Let s' be a winning strategy for Player \Diamond in $\mathcal{G}(\mathfrak{G}, \psi)$ with starting position $\langle v, \psi \rangle$. We define a strategy s in $\mathcal{G}(\mathfrak{G}, \varphi)$ as follows. In this game

there are two non-trivial choices for Player \Diamond : (I) in a position of the form $\langle u, P_i \rightarrow \Diamond X_i \rangle$, he has to choose between the formulae $\neg P_i$ and $\Diamond X_i$; and (II) in a position of the form $\langle u, \Diamond X_i \rangle$, he has to choose a successor of u .

Choice (I) is easy: if $u \in P_i$, Player \Diamond has to pick $\langle u, \Diamond X_i \rangle$, otherwise he chooses $\langle u, \neg P_i \rangle$. For Choice (II), Player \Diamond follows the strategy s' , i.e., he chooses the successor w of u such that

$$s'(\langle u, \Diamond \Box X_i \rangle) = \langle w, \Box X_i \rangle.$$

We claim that the resulting strategy s is winning. Hence, let p be a play in $\mathcal{G}(\mathfrak{G}, \varphi)$ conforming to s with starting position $\langle v, \varphi \rangle$. By construction of s , this play corresponds to some play p' in $\mathcal{G}(\mathfrak{G}, \psi)$ conforming to s' with starting position $\langle v, \psi \rangle$ such that the sequences of first components of the positions in p' and those in p coincide (with the same meaning of 'coincide' as above). As above, it follows that the least priority seen infinitely often in p is the same one as in p' . Since p' is winning, so is therefore p . \square

Notes

One of the first articles on games is by Zermelo [212], who proved the determinacy of chess. Gale and Stewart [84] proved the existence of undetermined games and the determinacy for open games. The full proof of Borel determinacy is by Martin [137].

The section on parity games follows [213] and [90]. The Theorem of Büchi and Landweber was originally proved in [36], and Theorem 4.10 is taken from [53]. The algorithm to solve parity games in Theorem 4.13 is from [37]. Our presentation owes much to a set of lecture notes by Bojańczyk and Czerwiński [27].

V Trees

1 Composition Theorems

TREES CAN BE NATURALLY DECOMPOSED. But before stating the corresponding composition theorems, let us fix our terminology regarding trees. We will use several different versions, depending on which one is most convenient at the time. We start with the graph-theoretic notion.

Definition 1.1. An *undirected tree* is an undirected graph \mathfrak{T} which is connected and acyclic. \lrcorner

Undirected trees will be mostly used in the more graph-theoretic chapters of this book. In the current chapter, we are mainly interested in directed ones. These come in several different variants. We start by defining them as plain sets.

Definition 1.2. Let D be a set of *directions*.

(a) The *prefix ordering* on D^* is defined by

$$x \leq_{\text{pf}} y \quad : \text{iff} \quad y = xz, \quad \text{for some } z \in D^*.$$

(b) A *tree domain* is a subset $T \subseteq D^*$ that is *prefix-closed*, i.e., such that $x \leq_{\text{pf}} y \in T$ implies $x \in T$. If T is a tree domain and $x, y \in T$, we call x an (immediate) *successor* of y if $y = xd$, for some $d \in D$. In this case, we also say that y is an (immediate) *predecessor* of x . We write $\text{Suc}(x)$ for the set of all successors of x and $\text{Suc}_*(x)$ for $\{x\} \cup \text{Suc}(x)$.

(c) A *branch* β of a tree domain T is a maximal linearly ordered set of vertices. For a branch β of T , we write $\beta(n)$ for the n -th vertex of β and we write $w \leq_{\text{pf}} \beta$ to indicate that w is some vertex of β . That is, we sometimes

identify a branch $\beta \subseteq T$ of length α with a function $\alpha \rightarrow T$ or with a word in D^α .

(d) The *subtree* of T attached at a vertex $v \in T$ is the tree domain

$$T|_v := \{ u \in D^* \mid vu \in T \}.$$

In language theory trees are usually labelled by some alphabet.

Definition 1.3. Let C be a set. A C -labelled tree is a function $t : T \rightarrow C$ where T is a tree domain. We usually denote the domain T by $\text{dom}(t)$, and we write $t(v)$ for the label at the vertex v .

There are several possible ways to encode a tree as a relational structure.

Definition 1.4. (a) An *order-tree* is a structure isomorphic to one of the form $\langle T, \leq_{\text{pf}} \rangle$ for some tree domain T . The elements of a tree are called *vertices*.

(b) A *successor-tree* is a structure of the form $\langle T, \text{suc} \rangle$ where T is a tree domain and

$$\text{suc} := \{ \langle u, v \rangle \mid u, v \in T \text{ and } v = ud \text{ for some } d \in D \}.$$

(c) An order-tree or a successor-tree \mathfrak{T} is *successor-ordered* if it is equipped with an additional partial order \leq_{so} such that

- ♦ \leq_{so} linearly orders the set of successors of every vertex $v \in T$, and
- ♦ successors of distinct vertices are incomparable.

Given a successor-ordered successor-tree $\mathfrak{T} = \langle T, E, \leq_{\text{so}} \rangle$ where the set $D = \{d_0, \dots, d_{n-1}\}$ of directions is finite, we will often use the format $\mathfrak{T} = \langle T, \text{suc}_0, \dots, \text{suc}_{n-1} \rangle$ where

$$\text{suc}_i := \{ \langle u, v \rangle \mid u, v \in T \text{ with } v = ud_i \}.$$

(d) A C -labelled tree is a structure obtained from an order-tree or a successor-tree by adding unary predicates $(P_c)_{c \in C}$ containing the vertices with label c . We will use the names *labelled tree* and *coloured tree* interchangeably.

We also consider trees equipped with additional relations between the successors of a vertex.

Definition 1.5. Let Σ be a relational signature. A Σ -enriched tree \mathfrak{T} is a $(\Sigma + \{\leq\})$ -structure $\langle T, \leq, \bar{R} \rangle$ such that $\langle T, \leq \rangle$ forms an order-tree and each relation R_i only contains tuples \bar{a} such that $\bar{a} \subseteq \{\langle \rangle\}$ or $\bar{a} \subseteq \text{Suc}(w)$, for some $w \in T$.

Remark. In particular, a C -labelled tree is \bar{P} -enriched, for a set of unary predicates \bar{P} .

Sometimes it is possible to give a single proof for a result that holds both for trees and linear orders. To do so, we have to introduce a generalised notion of a tree that also covers all linear orders. In a normal tree, the path from the root to a given vertex is always a finite chain. We relax this condition to allow arbitrary linear orders. This leads to the following definition.

Definition 1.6. (a) A *generalised tree* is a meet-semilattice $\mathfrak{T} = \langle T, \leq \rangle$ where every set of the form $\Downarrow v$ with $v \in T$ forms a chain. We denote the meet of $u, v \in T$ by $u \sqcap v$.

(b) A *branch* of a generalised tree \mathfrak{T} is a maximal chain $\beta \subseteq T$.

(c) A *subtree* of a generalised tree \mathfrak{T} is a subset $S \subseteq T$ that is upwards-closed with respect to the ordering and closed under \sqcap .

Example. Every linear order is a generalised tree.

We will prove several composition theorems for generalised trees. The first one concerns replacing subtrees.

Proposition 1.7. Let \mathfrak{T} be a coloured generalised tree and let \mathfrak{T}' be the generalised tree obtained from \mathfrak{T} by replacing an arbitrary number of subtrees \mathfrak{S}_i by generalised trees \mathfrak{S}'_i , for $i \in I$, such that

$$\mathfrak{S}_i \equiv_{\text{MSO}}^m \mathfrak{S}'_i, \quad \text{for all } i \in I.$$

Then

$$\mathfrak{T} \equiv_{\text{MSO}}^m \mathfrak{T}'.$$

Proof. We can write \mathfrak{T} as a generalised sum followed by a quantifier-free interpretation:

$$\mathfrak{T} \cong \tau \left(\sum_{k \in \mathfrak{R}} \mathfrak{C}_k \right),$$

where \mathfrak{R} is the generalised tree obtained from \mathfrak{T} by replacing every subtree \mathfrak{S}_i by a single vertex, and where each component \mathfrak{C}_k is either (I) a singleton or (II) isomorphic to \mathfrak{S}_i , for some i . The interpretation τ is given by

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= x \leq_K y \vee (x \sim y \wedge x \leq_C y), \\ \varphi_P(x) &:= Px \end{aligned}$$

(where \leq_K denotes the relation \leq of the index tree \mathfrak{R} and \leq_C the relation \leq of the components). Similarly, we can write

$$\mathfrak{T}' \cong \tau \left(\sum_{k \in \mathfrak{R}} \mathfrak{C}'_k \right),$$

where τ is the same interpretation as above and either

$$\mathfrak{C}'_k = \mathfrak{C}_k, \quad \text{or} \quad \mathfrak{C}'_k = \mathfrak{S}'_i \quad \text{and} \quad \mathfrak{C}_k = \mathfrak{S}_i, \quad \text{for some } i.$$

By assumption, we have $\mathfrak{C}_k \equiv_{\text{MSO}}^m \mathfrak{C}'_k$, for all k . Since τ is quantifier-free, the claim therefore follows by Theorem I.4.24 and Corollary I.4.6. \square

Instead of cutting a tree along an antichain, we can also decompose it along a branch.

Definition 1.8. Let I be a partial order. A subset $C \subseteq I$ is *convex* if

$$x, y \in C \quad \text{implies} \quad z \in C \quad \text{for all } x \leq z \leq y.$$

Similarly, we call an equivalence relation \sim on I *convex* if every \sim -class is convex. \square

Proposition 1.9. Let \mathfrak{T} and \mathfrak{T}' be coloured generalised trees, β and β' branches of \mathfrak{T} and \mathfrak{T}' , respectively, and \sim and \sim' convex equivalence relations on β and β' . For a convex set $I \subseteq T$, we set

$$\mathfrak{T}_{\beta}[I] := \langle \mathfrak{T}|_C, I \rangle, \quad \text{where} \quad C := \uparrow I \setminus \bigcup_{v \in \beta \setminus \downarrow I} \uparrow v.$$

If there exists an isomorphism $\sigma : \beta/\sim \cong \beta'/\sim'$ of linear orders such that

$$\mathfrak{T}_{\beta}[I] \equiv_{\text{MSO}}^m \mathfrak{T}'_{\beta'}[\sigma(I)], \quad \text{for every } \sim\text{-class } I,$$

then $\mathfrak{T} \equiv_{\text{MSO}}^m \mathfrak{T}'$.

Proof. We can write \mathfrak{T} as a generalised sum of the factors $\mathfrak{T}_{\beta}[I]$, followed by a quantifier free interpretation τ , and similarly for \mathfrak{T}' .

$$\mathfrak{T} \cong \tau \left(\sum_{I \in \beta/\sim} \mathfrak{T}_{\beta}[I] \right) \quad \text{and} \quad \mathfrak{T}' \cong \tau \left(\sum_{I \in \beta'/\sim'} \mathfrak{T}'_{\beta'}[I] \right),$$

where σ uses the formula

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= [x \sim y \wedge x \leq_T y] \vee [x <_{\beta} y \wedge Ix], \end{aligned}$$

where \leq_T denotes the ordering of the factors $\mathfrak{T}_{\beta}[I]$, \leq_{β} the ordering of the index set β/\sim , and I is the additional predicate added to $\mathfrak{T}_{\beta}[I] = \langle \mathfrak{T}|_C, I \rangle$. \square


Finally, there is a composition theorem that allows us to also replace interior parts of the tree.

Definition 1.10. Let \mathfrak{T} be a generalised tree and $S \subseteq T$ a finite set closed under meets (which then has to form a tree). Let $p : S \rightarrow S$ be the function mapping every vertex $s \in S$ to its parent in S (hence, $p(s)$ is undefined for the root s), and set

$$U(s) := \begin{cases} T & \text{if } s \text{ is the minimal element of } S, \\ \bigcup \{ \uparrow t \mid p(s) < t \leq s \} & \text{otherwise,} \end{cases}$$

for $s \in S$. The *decomposition* of \mathfrak{T} induced by S is the partition of T with the classes

$$T_s := U(s) \setminus \bigcup_{t > s} U(t), \quad \text{for } s \in S.$$



Remark. The above formula for $U(s)$ is made complicated by the fact that it has to work for generalised trees. If \mathfrak{T} is a tree, we can use the simpler definition

$$U(s) := \uparrow u,$$

where u is the successor of $p(x)$ with $p(x) < u \leq x$. (If $p(x)$ does not exist, we take for u the root of T .)

Proposition 1.11. Let \mathfrak{T} be a generalised tree, \bar{a} a finite tuple of vertices, S the closure of \bar{a} under meets, and $(T_s)_{s \in S}$ the decomposition of \mathfrak{T} induced by S . Then $\text{Th}_{\text{MSO}}^m(\mathfrak{T}, \bar{a})$ can be computed from (the isomorphism type of) the tree $\langle S, \leq, \bar{a} \rangle$ and the family $(\theta_s)_{s \in S}$ where

$$\theta_s := \text{Th}_{\text{MSO}}^m(\langle T_s, \leq, \downarrow s \rangle).$$

Proof. Set $\mathfrak{T}_s := \langle T_s, \leq, \downarrow s \rangle$ and $\mathfrak{S} := \langle S, \leq, \bar{a} \rangle$. Then

$$\langle \mathfrak{T}, \bar{a} \rangle \cong \sigma \left(\sum_{s \in \mathfrak{S}} \mathfrak{T}_s \right),$$

where σ is the quantifier-free interpretation with formulae

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= [x \sim y \wedge x \leq_T y] \vee [x \not\sim y \wedge x <_S y \wedge Px], \end{aligned}$$

where

- \sim is the equivalence relation whose classes are the components T_s of the generalised sum,
- \leq_S is the ordering of the index structure \mathfrak{S} ,
- \leq_T is the one of the components \mathfrak{T}_s , and
- P denotes the predicate $\downarrow s$ in \mathfrak{T}_s .

For every MSO_m-formula $\varphi(\bar{x})$ it follows by Theorem I.4.24 that there exists an MSO-formula ψ such that

$$\mathfrak{T} \models \varphi(\bar{a}) \quad \text{iff} \quad \langle \mathfrak{S}, \bar{U} \rangle \models \psi,$$

where $U_\theta := \{s \in S \mid \text{Th}_{\text{MSO}}^m(\mathfrak{T}_s) = \theta\}$. Fix an MSO_m-theory θ . For $\varphi(\bar{x}) := \bigwedge \theta$, we obtain some formula ψ_θ such that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{T}, \bar{a}) = \theta \quad \text{iff} \quad \langle \mathfrak{S}, \bar{U} \rangle \models \psi_\theta.$$

In particular, $\langle \mathfrak{S}, \bar{U} \rangle$ determines $\text{Th}_{\text{MSO}}^m(\mathfrak{T}, \bar{a})$. As the family $(\theta_s)_{s \in S}$ contains a complete description of $\langle \mathfrak{S}, \bar{U} \rangle$ the claim follows. \square

Exercise 1.1. Prove that, over the class of all generalised trees, every MSO-formula $\varphi(\bar{x})$ with $|\bar{x}| \geq 1$ is equivalent to a boolean combination of formulae of the form

$$\psi(x_i \sqcap x_j, x_k \sqcap x_l) \quad \text{and} \quad \vartheta(x_i),$$

where ψ is an MSO-formula with two free variables and ϑ an MSO-formula with one free variable. \square

Comparing FO and MSO

As an application let us use composition arguments to compare the expressive powers of FO and MSO over trees. Over most classes of structures, MSO is strictly stronger than FO. But there is a little trick that, for trees, allows us to translate MSO-formulae into FO, if we allow the FO-formulae to access additional information in the form of a certain colouring. We start with a lemma.

Lemma 1.12. *Let \mathfrak{S} be a finite semigroup and $(\varphi_c(x, y))_{c \in S}$ a family of MSO-formulae. There exist FO-formulae $(\psi_c(x, y; \bar{Z}))_{c \in S}$ and finitely many MSO-formulae $(\vartheta_i(x))_{i < n}$ such that, for every coloured order-tree \mathfrak{T} such that $(\varphi_c)_c$ defines an additive labelling λ on \mathfrak{T} , we have*

$$\mathfrak{T} \models \psi_c(u, v; \vartheta_0^{\mathfrak{T}}, \dots, \vartheta_{n-1}^{\mathfrak{T}}) \quad \text{iff} \quad \lambda(u, v) = c.$$

Proof. Fix a tree \mathfrak{T} and let λ be the labelling defined by $(\varphi_c)_c$. By Theorem III.1.10, there exists a weak Ramseyan split $\sigma : T \rightarrow [N]$ for λ . For $n < N$ and $c \in S$, we define

$$\begin{aligned} Q_n &:= \sigma^{-1}(n), \\ R_c &:= \{v \in T \mid \text{there is } u \sqsubset_\sigma v \text{ such that } \lambda(u, v) = c \text{ and} \\ &\quad \text{there is no } w \text{ with } u \sqsubset_\sigma w \sqsubset_\sigma v\}, \\ U_c &:= \{v \in T \mid \lambda(u, v) = c \text{ where } u \text{ is the predecessor of } v\}. \end{aligned}$$

Note that these predicates can be defined in MSO using the formulae $(\varphi_c)_c$.

We claim that λ can be defined in terms of these two predicates. For each $n < N$, we construct FO-formulae ψ_c^n such that

$$\mathfrak{T} \models \psi_c^n(u, v; \vartheta_0^{\mathfrak{T}}, \dots, \vartheta_{n-1}^{\mathfrak{T}}) \quad \text{iff} \quad \lambda(u, v) = c \text{ and} \\ \sigma(w) \leq n \text{ for all } u \leq w \leq v.$$

Then the formulae $\psi_c := \psi_c^{N-1}$ have the desired properties.

We define ψ_c^n by induction on n . If $n = 0$, we have $u \sqsubset_\sigma v$, which implies that

$$\lambda(u, v) = \lambda(u, u') \cdot \lambda(u', v) = \lambda(u, u'),$$

where u' is the immediate successor of u . Hence, can use

$$\begin{aligned} \psi_c^0(x, y) &:= x < y \wedge \forall z[x \leq z \leq y \rightarrow Q_0 z] \\ &\quad \wedge \exists x'[x < x' \leq y \wedge \neg \exists z[x < z < x'] \wedge R_c x']. \end{aligned}$$

For the inductive step, suppose that we have already defined ψ_c^n . Given $u < v$, let $w_0 < \dots < w_{k-1}$ be an enumeration of all vertices w with $u \leq w \leq v$ and $\sigma(w) = n + 1$. We distinguish several cases. If $k = 0$, we can compute $\lambda(u, v)$ by inductive hypothesis. If $k = 1$, we have

$$\begin{aligned} \lambda(u, v) &= \lambda(u, w_0) \cdot \lambda(w_0, v) \\ &= \lambda(u, w') \cdot \lambda(w', w_0) \cdot \lambda(w_0, w'') \cdot \lambda(w'', v), \end{aligned}$$

where the first and the last factor can be computed by inductive hypothesis, and the two middle factors can be obtained from the predicates \bar{U} . (If $u = w'$, or $u = w_0$, or $v = w''$, or $v = w_0$, we have to omit some of the factors above.)

Finally, suppose that $k > 1$. If $u < w_0$ and $w_{k-1} < v$, we have

$$\begin{aligned} \lambda(u, v) &= \lambda(u, w_0) \cdot \lambda(w_0, w_1) \cdot \lambda(w_1, w_{k-1}) \cdot \lambda(w_{k-1}, v) \\ &= \lambda(u, w_0) \cdot \lambda(w_0, w_1) \cdot \lambda(w_{k-1}, v). \end{aligned}$$

The first and the last factor can be computed as in the case $k = 1$ above, while the middle one can be obtained using the predicates \bar{R} .

If $u = w_0$ or $v = w_{k-1}$, we proceed similarly, just omitting the corresponding factors. \square

Theorem 1.13. *Given an MSO-formula $\varphi(\bar{x})$, we can compute an FO-formula $\varphi^*(\bar{x})$ and finitely many MSO-formulae $\psi_0(z), \dots, \psi_{n-1}(z)$ with a single free variable z such that*

$$\mathfrak{T} \models \varphi(\bar{v}) \quad \text{iff} \quad \langle \mathfrak{T}, \psi_0^{\mathfrak{T}}, \dots, \psi_{n-1}^{\mathfrak{T}} \rangle \models \varphi^*(\bar{v}),$$

for all coloured order-trees \mathfrak{T} .

Proof. Let $\mathfrak{T} = \langle T, \leq_{\text{pf}}, \bar{P} \rangle$ be a coloured order-tree (possibly successor-ordered) and let m be the quantifier-rank of φ . Consider a finite tuple $\bar{a} \subseteq T$,

let S be its closure under meets, and let $(U_s)_{s \in S}$ be the decomposition of \mathfrak{T} induced by S as in Definition 1.10. By Proposition 1.11, the theories

$$\theta_s := \text{Th}_{\text{MSO}}^m(\langle U_s, \leq_{\text{pf}}, \Downarrow s \rangle), \quad \text{for } s \in S,$$

determine $\text{Th}_{\text{MSO}}^m(\mathfrak{T}, \bar{a})$. Below we will construct MSO-definable unary predicates \tilde{Q} and FO-formula $(\vartheta_\theta)_\theta$ (that both do not depend on \bar{a}) such that

$$\langle \mathfrak{T}, \tilde{Q} \rangle \models \vartheta_\theta(s; \bar{a}) \quad \text{iff} \quad s \in S \text{ and } \theta_s = \theta.$$

Since the size of S is bounded by $|S| \leq 2 \cdot |\bar{a}|$ and since every element of S is definable from \bar{a} , we can therefore set

$$\varphi^*(\bar{x}) := \bigvee_{\bar{\theta}} (\forall s \in S) \vartheta_{\theta_s}(s; \bar{x}),$$

where the disjunction ranges over all families $(\theta_s)_s$ that imply the formula φ .

It remains to explain how to define the predicates \tilde{Q} . For vertices $u <_{\text{pf}} v$, we set

$$W_{uv} := \uparrow u' \setminus (\uparrow v \setminus \{v\}),$$

where u' is the successor of u with $u' <_{\text{pf}} v$. Let λ be the function mapping a pair of vertices $u <_{\text{pf}} v$ to

$$\lambda(u, v) := \text{Th}_{\text{MSO}}^m(\mathfrak{B}_{uv}) \quad \text{where} \quad \mathfrak{B}_{uv} := \langle W_{uv}, \leq, \Downarrow v, v \rangle.$$

For $u <_{\text{pf}} v <_{\text{pf}} w$, it follows that

$$\mathfrak{B}_{uw} = \sigma(\mathfrak{B}_{uv} \oplus \mathfrak{B}_{vw}),$$

for some fixed quantifier-free interpretation σ . Consequently, the labelling λ is additive. As it is also MSO-definable, it follows by Lemma 1.12 that it is FO-definable using suitable MSO-definable monadic parameters \tilde{Q} . Furthermore, for each MSO_m-theory θ , we define the sets

$$\begin{aligned} Q'_\theta &:= \{v \in T \mid \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v \setminus \{v\}}) = \theta\}, \\ Q''_\theta &:= \{v \in T \mid \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v}) = \theta\}, \\ Q'''_\theta &:= \{v \in T \mid \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus (\uparrow v \setminus \{v\})}) = \theta\}, \end{aligned}$$

We can express the theories θ_s as follows.

First, suppose that s is the minimal element of S . Then we can write U_s as a generalised sum

$$\mathfrak{T}|_{U_s} = \mathfrak{T}|_{T \setminus (\uparrow s \setminus \{s\})} + \sum_{v \in \text{Suc}(s) \setminus \Downarrow s} \mathfrak{T}|_{\uparrow v}$$

over the index structure $\text{Suc}_*(s) \setminus (\Downarrow s \cap \text{Suc}(s))$ (followed by a quantifier-free interpretation, which we omit for readability). Since this index structure is a tree of height 1 and since every MSO-formula over such a tree only depends on (I) the label of the root and (II) the number (up to some bound k) of leaves with a given label, it follows that we can compute θ_s from (I) the theory σ such that $s \in Q'''_\sigma$ and (II), for each theory τ , the number (up to k) of successors v of s with $v \in Q''_\tau$. This data can be determined by an FO-formula.

Similarly, if s is a maximal element of S with immediate predecessor t , we can write U_s as a generalised sum

$$\mathfrak{T}|_{U_s} = \mathfrak{T}|_{W_{ts}} + \mathfrak{T}|_{\uparrow s \setminus \{s\}}.$$

Consequently, we can compute θ_s from $\lambda(t, s)$ and \tilde{Q}' .

Finally, suppose that $s \in S$ is neither minimal nor maximal and let t be its immediate predecessor. Then

$$\mathfrak{T}|_{U_s} = \mathfrak{T}|_{W_{ts}} + \sum_{v \in \text{Suc}(s) \setminus \Downarrow s} \mathfrak{T}|_{\uparrow v}.$$

Hence, θ_s can be computed from $\lambda(s, t)$ and the statements $v \in Q''_\tau$, for $v \in \text{Suc}(s)$.

In each of these three cases, the computation of θ_s can be done by an FO-formula with parameters s and \bar{a} . \square

We can rephrase this result in terms of interpretations. Let us call an MSO-interpretation τ an MSO-colouring if τ only adds unary predicates, but otherwise leaves the input structure unchanged.

Corollary 1.14. *For every MSO-interpretation τ , there exists an FO-interpretation σ and an MSO-colouring ρ such that*

$$\tau(\mathcal{T}) = \sigma(\rho(\mathcal{T})), \quad \text{for every class } \mathcal{T} \text{ of coloured order-trees.}$$

Proof. By Theorem 1.13, we can replace each formula $\varphi(\bar{x})$ in τ by an FO-formula $\varphi^*(\bar{x})$. Let σ be the interpretation using these FO-formulae. Each formula in σ uses auxiliary MSO-formulae $\psi(z)$ with a single free variable. Let ρ be the interpretation adding all the relations defined by these formulae $\psi(z)$. Then $\tau = \sigma \circ \rho$ (over trees). \square

Lemma 1.15. *Let C be a finite set. There exists an FO-interpretation σ such that, for every class of C -labelled order-trees \mathcal{T} , we have*

$$\mathcal{T} = \sigma(\mathcal{T}_0), \quad \text{for some class } \mathcal{T}_0 \text{ of unlabelled order-trees.}$$

Proof. Suppose that $C = \{c_0, \dots, c_{n-1}\}$. Given a tree $\mathfrak{T} \in \mathcal{T}$, let \mathfrak{T}_0 be the uncoloured tree obtained from \mathfrak{T} by attaching $i+1$ new leaves as successors to every vertex v with label c_i . Then we can recover \mathfrak{T} from \mathfrak{T}_0 by the FO-interpretation $\sigma = \langle \delta, \varphi_{\leq}, (\psi_{P_{c_i}})_{i < n} \rangle$ defined by

$$\begin{aligned} \delta(x) &:= \text{'x is not a leaf'} \\ \varphi_{\leq}(x, y) &:= x \leq y \\ \psi_{P_{c_i}}(x) &:= \bigvee_k \text{'x has exactly } i+1 \text{ successors that are leaves.} \end{aligned}$$

It follows that $\mathcal{T} = \sigma(\mathcal{T}_0)$ where $\mathcal{T}_0 := \{\mathfrak{T}_0 \mid \mathfrak{T} \in \mathcal{T}\}$. \square

Corollary 1.16. *A class is MSO-interpretable in some class of order-trees if, and only if, it is FO-interpretable in such a class (but possibly a different one).*

Proof. Suppose that $\mathcal{C} = \tau(\mathcal{T})$ where τ is an MSO-interpretation and \mathcal{T} a class of trees. By Corollary 1.14, it follows that $\mathcal{C} = \sigma(\rho(\mathcal{T}))$, where σ is an FO-interpretation and ρ an MSO-colouring. Note that $\mathcal{S} := \rho(\mathcal{T})$ is a class of coloured trees. Let ν and \mathcal{S}_0 be the FO-interpretation and the class of trees from Lemma 1.15. Then $\mathcal{C} = (\sigma \circ \nu)(\mathcal{S}_0)$, as desired. \square

For trees of bounded height, we can do better by removing the need of a colouring.

Proposition 1.17. *Let $n < \omega$ and let Σ be a signature consisting of unary predicates only. Over the class of all Σ -enriched trees of height at most n , every MSO-formula $\varphi(\bar{X}, \bar{x})$ is equivalent to an FO-formula.*

Proof. The statement follows from the fact that, over the empty signature, all MSO can do is to count up to some constant depending on the quantifier-rank. To simplify notation, we include the parameters in the structure. That is, we work with structures of the form $\mathfrak{T} = \langle T, \leq, \bar{P}, \bar{c} \rangle$ where $\langle T, \leq \rangle$ is an order-tree of height at most n and \bar{P} and \bar{c} are parameters. We prove by induction on n that there exists some function $f_n : \omega \rightarrow \omega$ such that

$$\mathfrak{S} \equiv_{\text{FO}}^{f_n(m)} \mathfrak{T} \quad \text{implies} \quad \mathfrak{S} \equiv_{\text{MSO}}^m \mathfrak{T}, \quad \text{for every } m < \omega.$$

Then the claim follows by Lemma 1.3.3.

If $n = 0$, the trees consist of a single vertex. Structures of size 1 can be characterised up to isomorphism by an FO-formula of quantifier-rank 2. Hence,

$$\mathfrak{S} \equiv_{\text{FO}}^2 \mathfrak{T} \quad \text{implies} \quad \mathfrak{S} \equiv_{\text{MSO}}^m \mathfrak{T}, \quad \text{for all } m.$$

For the inductive step, suppose that $n > 0$. Given a tree \mathfrak{S} of height at most n , let $(\mathfrak{S}_i)_{i \in I}$ be an enumeration of the subtrees attached at the root and let \mathfrak{S}° be the substructure consisting only of the root. Then we can write \mathfrak{S} as

$$\mathfrak{S} \cong \sigma\left(\mathfrak{S}^\circ \oplus \bigoplus_{i \in I} \mathfrak{S}_i\right),$$

where σ is a quantifier-free interpretation that adds the order relations between the root and the other elements. We can replace the disjoint union by a generalised sum and we obtain

$$\mathfrak{S} \cong \sigma'\left(\mathfrak{S}^\circ \oplus \sum_{i \in I} \mathfrak{S}_i\right),$$

for a slightly different interpretation σ' . By Theorem I.4.24, it follows that there exists some function g such that

$$\mathfrak{C}^\circ \equiv_{\text{MSO}}^m \mathfrak{T}^\circ \quad \text{and} \quad \langle I, \bar{Q} \rangle \equiv_{\text{MSO}}^{g(m)} \langle I', \bar{Q}' \rangle \quad \text{implies} \quad \mathfrak{C} \equiv_{\text{MSO}}^m \mathfrak{T},$$

where

$$Q_\theta := \{ i \in I \mid \text{Th}_{\text{MSO}}^m(\mathfrak{C}_i) = \theta \}$$

and similarly for Q'_θ and the \mathfrak{T}_i .

By inductive hypothesis, there exists an FO-interpretation of quantifier-rank $f_{n-1}(m) + 2$ mapping \mathfrak{C} to the index structure $\langle I, \bar{Q} \rangle$. There also exists an FO-interpretation of quantifier-rank 1 mapping \mathfrak{C} to \mathfrak{C}° . Furthermore, by Proposition I.4.8, there exists some function h such that

$$\langle I, \bar{Q} \rangle \equiv_{\text{FO}}^{h(k)} \langle I', \bar{Q}' \rangle \quad \text{implies} \quad \langle I, \bar{Q} \rangle \equiv_{\text{MSO}}^k \langle I', \bar{Q}' \rangle.$$

Consequently,

$$\begin{aligned} & \mathfrak{C} \equiv_{\text{FO}}^{f_{n-1}(m)+2+h(g(m))} \mathfrak{T} \\ \Rightarrow & \mathfrak{C}^\circ \equiv_{\text{FO}}^2 \mathfrak{T}^\circ \quad \text{and} \quad \langle I, \bar{Q} \rangle \equiv_{\text{FO}}^{b(g(m))} \langle I', \bar{Q}' \rangle \\ \Rightarrow & \mathfrak{C}^\circ \cong \mathfrak{T}^\circ \quad \text{and} \quad \langle I, \bar{Q} \rangle \equiv_{\text{MSO}}^{g(m)} \langle I', \bar{Q}' \rangle \\ \Rightarrow & \mathfrak{C} \equiv_{\text{MSO}}^m \mathfrak{T}. \end{aligned} \quad \square$$

2 Tree Automata

The Theorem of Büchi about the decidability of $\langle \omega, \leq \rangle$ can be extended to the monadic second-order theory of the infinite complete binary tree $\mathfrak{T} := \langle \{0, 1\}^*, \text{succ}_0, \text{succ}_1 \rangle$. As there currently does not exist a purely model-theoretic proof of this fact, we present the standard automata-theoretic version here. In this section we prove that, over enriched trees, monadic second-order logic is equivalent to tree automata. In the next one, we will then prove several decidability results.

Nondeterministic Automata

Let us define our model of an automaton. We start with a nondeterministic version. Intuitively, an automaton labels each vertex of the given input tree by some state and this labelling has to satisfy two consistency conditions: (i) the labelling of the successors of a vertex v have to match the state at v and the letter at v , and (ii) each infinite branch has to satisfy the parity condition. For (i), the automaton is equipped with an MSO-formula that checks whether the states match. The details are as follows.

Definition 2.1. Let L be a logic.

(a) Let Σ be a relational signature and Q a set. We write

$$\Sigma_Q := \Sigma + \{\text{rt}\} + \{S_q \mid q \in Q\}$$

for the expanded signature where rt is a constant symbol and the S_q are unary predicates. The *transition logic* $\text{TL}_L[\Sigma, Q]$ consists of those $L[\Sigma_Q]$ -formulae where every subformula of the form $S_q x$, for $q \in Q$, only appears positively, i.e., under an even number of negation signs.

(b) Let \mathfrak{T} be a Σ -enriched tree and $\rho : T \rightarrow Q$ a function. The *successor structure* $\mathfrak{Cuc}_*(v; \rho)$ associated with a vertex $v \in T$ is the Σ_Q -structure

$$\mathfrak{Cuc}_*(v; \rho) := \langle \mathfrak{C}, (S_q)_{q \in Q}, v \rangle,$$

where \mathfrak{C} is the substructure of \mathfrak{T} induced by the set $\text{Suc}_*(v)$ and

$$S_q := \{ x \in \text{Suc}(v) \mid \rho(x) = q \}, \quad \text{for } q \in Q.$$

(c) A *nondeterministic L -automaton* is a tuple $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, \Omega \rangle$ where Q is a finite set of *states*, Σ is a finite *input signature* of enriched trees, $q_0 \in Q$ is the *initial state*, $\Omega : Q \rightarrow \omega$ is a *priority function*, and $\delta : Q \rightarrow \text{TL}_L[\Sigma, Q]$ is the *transition function*.

(d) A *run* of a tree automaton $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, \Omega \rangle$ on an enriched tree \mathfrak{T} over the signature Σ is a function $\rho : T \rightarrow Q$ such that

$$\mathfrak{Cuc}_*(v; \rho) \models \delta(\rho(v)), \quad \text{for all } v \in T.$$

A run ρ is *accepting* if $\rho(\langle \rangle) = q_o$ and, for every infinite branch β of \mathfrak{T} ,

$$\liminf_{n \rightarrow \infty} \Omega(\rho(\beta(n))) \text{ is even.}$$

(e) An L -automaton \mathcal{A} *accepts* a Σ -enriched tree \mathfrak{T} if there exists an accepting run of \mathcal{A} on \mathfrak{T} . The *language recognised* by \mathcal{A} is the set $L_{\text{nd}}(\mathcal{A})$ of all trees it accepts. \lrcorner

Example. Let K be the language of all trees over the alphabet $\{a, b\}$ the contain at least one letter a . We regard such trees as $\{P_a, P_b\}$ -enriched where P_a and P_b are two unary predicates encoding the labelling. We can recognise K by the following nondeterministic MSO-automaton.

$$\mathcal{A} := \langle Q, \Sigma, \delta, q, \Omega \rangle$$

where $Q := \{q\}$, $\Omega(q) := 1$, and

$$\delta(q) := P_a(\text{rt}) \vee \exists x[x \neq \text{rt} \wedge S_q x]. \quad \lrcorner$$

Exercise 2.1. Find tree automata recognising the following languages over the alphabet $\{a, b\}$.

- (a) The language of all trees containing infinitely many letters a .
- (b) The language of all trees such that below every vertex there is some vertex with the letter a .
- (c) The language of all trees such that there is some vertex below which there are only letters a .
- (d) The language of all trees such that some branch contains only letters a .
- (e) The language of all trees such that every branch contains at least one a .
- (f) The language of all trees such that every branch contains only finitely many a . \lrcorner

Exercise 2.2.

(a) Let $K \subseteq \Sigma^\omega$ be an MSO-definable language. Construct an MSO-automaton \mathcal{A} accepting the language of all trees where each infinite branch belongs to L .

(b) Let $K \subseteq \{0, 1\}^*$ be an MSO-definable language. Construct an MSO-automaton \mathcal{A} accepting the language of all trees over the alphabet $\{a, b\}$ such that a vertex v is labelled by a if, and only if, $v \in K$. \lrcorner

Exercise 2.3. Let \mathcal{A} be a nondeterministic MSO-automaton. Find an MSO-formula defining $L(\mathcal{A})$. \lrcorner

For the translation of MSO into tree automata, we have to establish several closure properties for languages recognised by automata. We start with the closure under union and projection.

Proposition 2.2. Let L be one of MSO, CMSO, GSO, or CGSO. Given two nondeterministic L -automata \mathcal{A} and \mathcal{A}' , we can compute an nondeterministic L -automaton that recognises the language $L_{\text{nd}}(\mathcal{A}) \cup L_{\text{nd}}(\mathcal{A}')$.

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ and $\mathcal{A}' = \langle Q', \Sigma, \delta', q'_o, \Omega' \rangle$. We set

$$\mathcal{B} := \langle Q + Q' + \{q_+\}, \Sigma, \delta_+, q_+, \Omega_+ \rangle$$

where

$$\delta_+(q) := \begin{cases} \delta(q) & \text{if } q \in Q, \\ \delta'(q) & \text{if } q \in Q', \\ \delta(q_o) \vee \delta'(q'_o) & \text{if } q = q_+, \end{cases}$$

$$\Omega(q) := \begin{cases} \Omega(q) & \text{if } q \in Q, \\ \Omega'(q) & \text{if } q \in Q', \\ 0 & \text{if } q = q_+. \end{cases}$$

We claim that $L_{\text{nd}}(\mathcal{B}) = L_{\text{nd}}(\mathcal{A}) \cup L_{\text{nd}}(\mathcal{A}')$.

(\supseteq) Given an accepting run ρ of \mathcal{A} on some tree \mathfrak{T} , we obtain an run of \mathcal{B} on \mathfrak{T} by replacing the initial state $\rho(\emptyset) = q_o$ by q_+ . This run is again accepting. In the same way, we can turn an accepting run of \mathcal{A}' on \mathfrak{T} into one of \mathcal{B} .

(\subseteq) Let ρ be an accepting run of \mathcal{B} on some tree \mathfrak{T} . By definition of $\delta_+(q_+)$, we have

$$\text{Suc}_*(\langle \rangle; \rho) \models \delta(q_o) \vee \delta'(q'_o).$$

By symmetry, we may suppose that

$$\mathcal{S}uc_*(\langle \rangle; \rho) \models \delta(q_o).$$

Then we can replace the state $\rho(\langle \rangle) = q_+$ by q_o and we obtain a run of \mathcal{A} on \mathfrak{T} which is again accepting. \square

Proposition 2.3. *Let L be one of MSO, CMSO, GSO, or CGSO. Let Σ be a signature and let $U \notin \Sigma$ be an additional unary predicate. For every non-deterministic L -automaton \mathcal{A} over the signature $\Sigma + \{U\}$, we can compute a nondeterministic L -automaton \mathcal{B} over the signature Σ such that*

$$L(\mathcal{B}) = \{ \mathfrak{T} \mid \langle \mathfrak{T}, U \rangle \in L(\mathcal{A}) \text{ for some } U \subseteq T \}.$$

Proof. Given $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$, we set

$$\mathcal{B} := \langle Q \times \{o, i\} + \{q_+\}, \Sigma, \delta', q_+, \Omega' \rangle$$

where the priority function is

$$\Omega'(\langle q, b \rangle) := \Omega(q) \quad \text{and} \quad \Omega'(q_+) := o,$$

and the transition function is

$$\begin{aligned} \delta'(q_+) &:= \exists U \delta(q_o)[S_p \mapsto \vartheta_p]_{p \in Q}, \\ \delta'(\langle q, o \rangle) &:= \exists U [\neg U(\text{rt}) \wedge \delta(q)][S_p \mapsto \vartheta_p]_{p \in Q}, \\ \delta'(\langle q, i \rangle) &:= \exists U [U(\text{rt}) \wedge \delta(q)][S_p \mapsto \vartheta_p]_{p \in Q}, \end{aligned}$$

where $\varphi[S_p \mapsto \vartheta_p]_p$ denotes the formula obtained from φ by replacing every atom of the form $S_p x$ by the formula

$$\vartheta_p(x) := [\neg Ux \wedge S_{(p,o)}x] \vee [Ux \wedge S_{(p,i)}x].$$

We claim that \mathcal{B} accepts a tree \mathfrak{T} if, and only if, \mathcal{A} accepts $\langle T, U \rangle$, for some $U \subseteq U$.

(\Leftarrow) Suppose that $\langle \mathfrak{T}, U \rangle \in L(\mathcal{A})$ and let ρ be a corresponding accepting run. We obtain an accepting run ρ' of \mathcal{B} on \mathfrak{T} by setting

$$\rho'(\langle \rangle) := q_+ \quad \text{and} \quad \rho'(v) := \begin{cases} \langle \rho(v), o \rangle & \text{if } v \notin U, \\ \langle \rho(v), i \rangle & \text{if } v \in U, \end{cases} \quad \text{for } v \neq \langle \rangle.$$

(\Rightarrow) Suppose that $\mathfrak{T} \in L(\mathcal{B})$ and let ρ be a corresponding accepting run. For every vertex $v \in T$ with $\rho(v) = \langle q, b \rangle$, there exists a set $U_v \subseteq \text{Suc}_*(v)$ such that

$$v \in U_v \Leftrightarrow b = i \quad \text{and} \quad \langle \mathcal{S}uc_*(v; \rho), U_v \rangle \models \delta(q)[S_p \mapsto \vartheta_p]_{p \in Q}.$$

Set $U := \bigcup_{v \in T} (U_v \setminus \{v\})$ and let $\rho' : T \rightarrow Q$ be the function $\rho' := f \circ \rho$ where

$$f(q_+) := q_o \quad \text{and} \quad f(\langle q, b \rangle) := q.$$

Then ρ' is an accepting run of \mathcal{A} on $\langle \mathfrak{T}, U \rangle$. \square

Alternating Automata

Closure under complement is difficult to prove using nondeterministic automata. To simplify the proof, we therefore introduce a second automaton model. An *alternating* automaton can make not only existential choices of states but also universal ones. Intuitively, one can think of the automaton splitting into several different copies, each of which reading the remainder of the tree independently.

Definition 2.4. Let L be a logic. An *alternating L -automaton*

$$\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$$

has the same form as a nondeterministic one, but the notions of a run and a successor structure are defined differently.

A *run* of an alternating automaton $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ on a Σ -enriched tree \mathfrak{T} is a function $\rho : T \rightarrow \wp(Q \times Q)$ such that

$$\mathcal{S}uc_*(v; \rho/q) \models \delta(q), \quad \text{for all } \langle p, q \rangle \in \rho(v) \text{ and all } v \in T,$$

where the *successor structure* $\mathfrak{Succ}_*(v; \rho/q)$ is obtained from the substructure of \mathfrak{T} induced by the set $\text{Suc}_*(v)$ by adding the predicates

$$S_p := \{x \in \text{Suc}(v) \mid \langle q, p \rangle \in \rho(x)\}, \quad \text{for } p \in Q.$$

A *trace* of a run ρ is a sequence $(p_n)_{n < \omega}$ of states such that, for some branch β of ρ ,

$$\langle p_n, p_{n+1} \rangle \in \rho(\beta(n)), \quad \text{for all } n < \omega.$$

A run ρ is *accepting* if $\langle q_o, q_o \rangle \in \rho(\langle \rangle)$ and

$$\liminf_{n \rightarrow \infty} \Omega(p_n) \text{ is even, for every trace } (p_n)_{n < \omega} \text{ of } \rho.$$

The *language recognised* by \mathcal{A} is denoted $L_{\text{alt}}(\mathcal{A})$. \downarrow

Note that the traces of a run correspond to the various copies of the automaton as it reads the tree. To reconstruct these traces from the run, we not only have to know the current state, but also the previous one. This is why we use pairs of states.

Example. The following alternating MSO-automaton recognises the language of all trees over the alphabet $\{a, b, c\}$ that contain at least one letter a and at least one letter b .

$$\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$$

where $Q := \{q_o, q_a, q_b\}$, $\Omega(q_o) = \Omega(q_a) = \Omega(q_b) = 1$, and

$$\begin{aligned} \delta(q_o) &:= [P_a(\text{rt}) \vee \exists x[x \neq \text{rt} \wedge S_{q_a}x]] \\ &\quad \wedge [P_b(\text{rt}) \vee \exists x[x \neq \text{rt} \wedge S_{q_b}x]], \\ \delta(q_a) &:= P_a(\text{rt}) \vee \exists x[x \neq \text{rt} \wedge S_{q_a}x], \\ \delta(q_b) &:= P_b(\text{rt}) \vee \exists x[x \neq \text{rt} \wedge S_{q_b}x]. \end{aligned} \quad \downarrow$$

When working with alternating automata it is often easier to base acceptance not on runs but on a certain parity game. The definition is very similar to the model-checking game for the modal μ -calculus.

Definition 2.5. Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ be an alternating automaton and let \mathfrak{T} be an enriched tree over Σ . The *automaton game* $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ for \mathcal{A} on \mathfrak{T} is the parity game where the positions of Player \Diamond , also called *Automaton*, are

$$V_{\Diamond} := T \times Q$$

and the positions of Player \square , also called *Tracer*, are

$$V_{\square} := \bigcup_{v \in T} \wp(\text{Suc}(v) \times Q).$$

The initial position is $\langle \langle \rangle, q_o \rangle$.

The edge relation is defined as follows. From a V_{\square} -position H there are edges to every V_{\Diamond} -position $\langle u, p \rangle \in H$. From a V_{\Diamond} -position $\langle v, q \rangle$ there are edges to every V_{\square} -position $H \subseteq \text{Suc}(v) \times Q$ such that

$$\langle \mathfrak{S}, \bar{S}, v \rangle \models \delta(q),$$

where, similarly to the definition of $\mathfrak{Succ}_*(v; \rho)$, \mathfrak{S} is the substructure of \mathfrak{T} induced by the set $\text{Suc}_*(v)$ and

$$S_p := \{u \in \text{Suc}(v) \mid \langle u, p \rangle \in H\}, \quad \text{for } p \in Q.$$

Finally, we assign to positions $\langle v, q \rangle \in V_{\Diamond}$ the priority $\Omega(q)$ and to positions $H \in V_{\square}$ an arbitrary priority greater than all priorities used by \mathcal{A} . \downarrow

Proposition 2.6. *Automaton has a winning strategy in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ if, and only if, $\mathfrak{T} \in L_{\text{alt}}(\mathcal{A})$.*

Proof. (\Leftarrow) Given an accepting run ρ of \mathcal{A} on \mathfrak{T} , we construct a winning strategy σ for Automaton in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ as follows. In a position $\langle v, q \rangle \in V_{\Diamond}$, Automaton chooses the new position

$$H := \{ \langle u, p \rangle \in \text{Suc}(v) \times Q \mid \langle q, p \rangle \in \rho(u) \}.$$

To see that this strategy is winning, consider a play

$$\langle v_o, q_o \rangle, H_o, \langle v_1, q_1 \rangle, H_1, \dots$$

conforming to σ . By definition of σ , it follows that the sequence q_0, q_1, \dots of states appearing in this play is a trace of ρ . As ρ is accepting, this trace satisfies the parity condition. Consequently, the above play also satisfies the parity condition and Automaton wins the game.

(\Rightarrow) Let σ be a winning strategy for Automaton in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$. We construct a run ρ of \mathcal{A} on \mathfrak{T} inductively as follows. We start with $\rho(\langle \rangle) := \{ \langle q_0, q_0 \rangle \}$. For the inductive step, suppose that $\rho(v)$ has already been defined. Let

$$U := \{ q \in Q \mid \langle p, q \rangle \in \rho(v) \text{ for some } p \}.$$

For $u \in \text{Suc}(v)$, set

$$\rho(u) := \{ \langle q, p \rangle \mid q \in U, \langle u, p \rangle \in \sigma(\langle v, q \rangle) \}.$$

Then ρ is a run of \mathcal{A} on \mathfrak{T} .

To see that it is accepting, consider a trace $(q_n)_{n < \omega}$ of ρ along a branch β . Let

$$\langle v_0, q_0 \rangle, H_0, \langle v_1, q_1 \rangle, H_1, \dots$$

be a play of $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ conforming to σ where Tracer chooses in step n some pair $\langle v_n, q_n \rangle \in H_{n-1}$ such that v_n is the successor of v_{n-1} that lies on the branch β . Since σ was assumed to be winning, this play satisfies the parity condition. Hence, so does the trace $(q_n)_{n < \omega}$. \square

After these preparations, we are finally able to prove closure under complement.

Definition 2.7. Let $\varphi \in \text{TL}_L[\Sigma, Q]$. The *dual* of φ is the formula $\bar{\varphi}$ obtained from $\neg\varphi$ by negating all atomic formulae of the form $S_q x$ with $q \in Q$.

Let us note the following property of this operation.

Lemma 2.8. Let \mathfrak{C} be a Σ -structure, $a \in C$, and $\varphi \in \text{TL}_L[\Sigma, Q]$. For a family \tilde{S} of subsets $S_q \subseteq C$, we have

$$\langle \mathfrak{C}, \tilde{S}, a \rangle \models \bar{\varphi} \quad \text{iff} \quad \text{for all } \tilde{S}' \text{ in } \mathfrak{C} \text{ with } \langle \mathfrak{C}, \tilde{S}', a \rangle \models \varphi \text{ there is} \\ \text{some } q \in Q \text{ with } S_q \cap S'_q \neq \emptyset.$$

Proof. Set $S_q^c := C \setminus S_q$. Then

$$\begin{aligned} \langle \mathfrak{C}, \tilde{S}, a \rangle \models \bar{\varphi} & \quad \text{iff} \quad \langle \mathfrak{C}, \tilde{S}^c, a \rangle \models \neg\varphi \\ & \quad \text{iff} \quad \text{there is no } \tilde{S}' \text{ in } \mathfrak{C} \text{ with } \langle \mathfrak{C}, \tilde{S}', a \rangle \models \varphi \\ & \quad \text{and } S'_q \subseteq S_q^c \text{ for all } q \in Q \\ & \quad \text{iff} \quad \text{for all } \tilde{S}' \text{ in } \mathfrak{C} \text{ with } \langle \mathfrak{C}, \tilde{S}', a \rangle \models \varphi, \text{ there} \\ & \quad \text{is some } q \in Q \text{ with } S_q \cap S'_q \neq \emptyset, \end{aligned}$$

where the second step follows from the fact that the formula φ is monotone in \tilde{S} . \square

Proposition 2.9. Let L be one of MSO, CMSO, GSO, or CGSO. Given an alternating L -automaton \mathcal{A} , we can compute an alternating L -automaton that recognises the complement of $L_{\text{alt}}(\mathcal{A})$.

Proof. Suppose that $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, \Omega \rangle$. We construct a new automaton $\mathcal{B} := \langle Q, \Sigma, \tilde{\delta}, q_0, \tilde{\Omega} \rangle$ where, for $q \in Q$,

$$\tilde{\delta}(q) := \overline{\delta(q)} \quad \text{and} \quad \tilde{\Omega}(q) := \Omega(q) + 1.$$

We claim that \mathcal{B} recognises the complement of $L_{\text{alt}}(\mathcal{A})$.

First, let $\mathfrak{T} \in L_{\text{alt}}(\mathcal{A})$ and let ρ be an accepting run of \mathcal{A} on \mathfrak{T} . Consider any run ρ' of \mathcal{B} on \mathfrak{T} . We have to show that ρ' is not accepting. We define a branch $(v_n)_{n < \omega}$ of \mathfrak{T} and a sequence $(p_n)_{n < \omega}$ of states such that

$$\langle p_n, p_{n+1} \rangle \in \rho(v_{n+1}) \cap \rho'(v_{n+1}), \quad \text{for all } n < \omega.$$

We start with $v_0 := \langle \rangle$ and $p_0 := q_0$. Suppose that v_n and p_n have already been defined. Then

$$\mathfrak{C}_{\text{uc}_*}(v_n; \rho/p_n) \models \delta(p_n) \quad \text{and} \quad \mathfrak{C}_{\text{uc}_*}(v_n; \rho'/p_n) \models \overline{\delta(p_n)}.$$

Hence, we can use Lemma 2.8 to find $v_{n+1} \in \text{Suc}(v_n)$ and $p_{n+1} \in Q$ with

$$\langle p_n, p_{n+1} \rangle \in \rho(v_{n+1}) \cap \rho'(v_{n+1}).$$

Note that the sequence $(p_n)_{n < \omega}$ we constructed is a trace of both runs. Since ρ is accepting, it follows that this trace satisfies the parity condition, that is,

$$\liminf_{n \rightarrow \infty} \Omega(p_n) \text{ is even.}$$

Consequently,

$$\liminf_{n \rightarrow \infty} \tilde{\Omega}(p_n) = \liminf_{n \rightarrow \infty} \Omega(p_n) + 1 \text{ is odd.}$$

We found a trace of ρ' that does not satisfy the parity condition. Consequently, ρ' is not accepting.

To conclude the proof we have to show that every tree $\mathfrak{T} \notin L_{\text{alt}}(\mathcal{A})$ has an accepting run for the automaton \mathcal{B} . By Proposition 2.6, Tracer has a winning strategy σ in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ and it is sufficient to construct a winning strategy σ' for Automaton in $\mathcal{G}(\mathcal{B}, \mathfrak{T})$. For $\langle v, q \rangle \in V_{\Diamond}$, we set

$$\sigma'(\langle v, q \rangle) := \{ \sigma(K) \mid \langle v, q \rangle \rightarrow K \text{ an edge in } \mathcal{G}(\mathcal{A}, \mathfrak{T}) \}.$$

First, let us prove that $H := \sigma'(\langle v, q \rangle)$ is actually a successor of $\langle v, q \rangle$ in $\mathcal{G}(\mathcal{B}, \mathfrak{T})$. Let \mathfrak{C} be the substructure of \mathfrak{T} induced by $\text{Suc}_*(v)$ and set

$$S_p := \{ u \in \text{Suc}(v) \mid \langle u, p \rangle \in H \}.$$

We have to show that $\langle \mathfrak{C}, \tilde{S}, v \rangle \models \overline{\delta(q)}$. By Lemma 2.8, it is sufficient to prove that, for all \tilde{S}' in \mathfrak{C} ,

$$\langle \mathfrak{C}, \tilde{S}', v \rangle \models \delta(q) \text{ implies } S_p \cap S'_p \neq \emptyset, \text{ for some } p \in Q.$$

Hence, suppose that $\langle \mathfrak{C}, \tilde{S}', v \rangle \models \delta(q)$. Set

$$K := \{ \langle u, p \rangle \in \text{Suc}(v) \times Q \mid u \in S'_p \}.$$

By choice of \tilde{S}' , the position K is a successor of $\langle v, q \rangle$ in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$. Let $\langle u, p \rangle := \sigma(K)$. By definition of H , it follows that $\langle u, p \rangle \in H$, i.e., $u \in S_p$.

As $\langle u, p \rangle$ is a successor of K in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$, we further have $\langle u, p \rangle \in K$, i.e., $u \in S'_p$. Consequently, $u \in S_p \cap S'_p \neq \emptyset$, as desired.

It remains to prove that σ' is a winning strategy for Automaton. Let

$$\langle v_0, q_0 \rangle, H_0, \langle v_1, q_1 \rangle, H_1, \dots$$

be a play in $\mathcal{G}(\mathcal{B}, \mathfrak{T})$ conforming to σ' . Then $\langle v_n, q_n \rangle \in H_{n-1}$, for $n < \omega$. By definition of σ' , it follows that in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ there are edges

$$\langle v_{n-1}, q_{n-1} \rangle \rightarrow K_{n-1} \quad \text{with} \quad \sigma(K_{n-1}) = \langle v_n, q_n \rangle, \quad \text{for } n < \omega.$$

Consequently,

$$\langle v_0, q_0 \rangle, K_0, \langle v_1, q_1 \rangle, K_1, \dots$$

is a play in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$ following the strategy σ . Since σ is winning for Tracer, it follows that the sequence $(q_n)_{n < \omega}$ does not satisfy the parity condition in $\mathcal{G}(\mathcal{A}, \mathfrak{T})$. As the priorities in $\mathcal{G}(\mathcal{B}, \mathfrak{T})$ are shifted by 1, this implies that $(q_n)_{n < \omega}$ does satisfy the parity condition in $\mathcal{G}(\mathcal{B}, \mathfrak{T})$. Hence, the play is winning for Automaton. \square

Equivalence Of The Automaton Models

It remains to prove that alternating automata are equivalent to nondeterministic ones. One direction is straightforward.

Proposition 2.10. *Let L be one of MSO, CMSO, GSO, or CGSO. For every nondeterministic L -automaton, we can compute an alternating L -automaton recognising the same language.*

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, \Omega \rangle$ be a nondeterministic automaton. We construct the automaton $\tilde{\mathcal{A}} := \langle Q, \Sigma, \tilde{\delta}, q_0, \Omega \rangle$ with transition function

$$\tilde{\delta}(q) := \exists \tilde{Z} [\text{refine}(\tilde{Z}, \tilde{S}) \wedge \delta(q)[\tilde{S} \mapsto \tilde{Z}]]$$

where the formula

$$\begin{aligned} \text{refine}(\tilde{Z}, \tilde{S}) := & \bigwedge_{p \in Q} Z_p \subseteq S_p \wedge \bigwedge_{p \neq q} Z_p \cap Z_q = \emptyset \\ & \wedge \forall x \left[x \neq \text{rt} \rightarrow \bigvee_{p \in Q} Z_p \right] \end{aligned}$$

states that \tilde{Z} is a partition of all non-root vertices that is contained in \tilde{S} . ($\varphi[\tilde{S} \mapsto \tilde{Z}]$ denotes the formula obtained from φ by replacing each atomic subformula of the form $S_i x$ by the corresponding formula $Z_i x$.) We claim that

$$L_{\text{nd}}(\mathcal{A}) = L_{\text{nd}}(\tilde{\mathcal{A}}) = L_{\text{alt}}(\tilde{\mathcal{A}}).$$

For the first equation, let $\rho : T \rightarrow Q$ be a run of a nondeterministic automaton. Then the predicates \tilde{S} of the structure $\mathfrak{Succ}_*(v; \rho)$ form a partition of the non-root vertices. Hence, $\tilde{Z} = \tilde{S}$ is the unique family of sets satisfying the formula $\text{refine}(\tilde{Z}, \tilde{S})$. Consequently,

$$\mathfrak{Succ}_*(v; \rho) \models \tilde{\delta}(q) \quad \text{iff} \quad \mathfrak{Succ}_*(v; \rho) \models \delta(q).$$

It follows that ρ is an accepting run of $\tilde{\mathcal{A}}$ if, and only if, it is an accepting run of \mathcal{A} .

For the second equation, let $\rho : T \rightarrow Q$ be an accepting run of $\tilde{\mathcal{A}}$, considered as a nondeterministic automaton. Then we obtain an accepting run ρ' of the corresponding alternating automaton by setting

$$\rho'(\langle \rangle) := \{ \langle q_o, q_o \rangle \} \quad \text{and} \quad \rho'(v) := \{ \langle \rho(u), \rho(v) \rangle \},$$

for every vertex v with immediate predecessor u .

Conversely, let $\rho : T \rightarrow \mathcal{P}(Q \times Q)$ be an accepting run of $\tilde{\mathcal{A}}$, considered as an alternating automaton. As we have already established the first equality, it is sufficient to construct an accepting run ρ' of \mathcal{A} . We do so by induction. We start with $\rho'(\langle \rangle) := q_o$. For the inductive step, suppose that $\rho'(v)$ has already been defined. Since

$$\mathfrak{Succ}_*(v; \rho/\rho'(v)) \models \tilde{\delta}(\rho'(v))$$

there exists a family \tilde{S}' of sets such that

$$\mathfrak{Succ}_*(v; \rho/\rho'(v)) \models \text{refine}(\tilde{S}', \tilde{S}) \wedge \delta(\rho'(v))[\tilde{S} \mapsto \tilde{S}'].$$

For each $u \in \text{Succ}(v)$, we define

$$\rho'(u) := q \quad \text{where} \quad u \in S'_q.$$

(As \tilde{S}' is a partition of $\text{Succ}(v)$, this is well-defined.) Then

$$\mathfrak{Succ}_*(v; \rho') \models \delta(\rho'(v)), \quad \text{for all } v \in T,$$

which means that ρ' is a run of \mathcal{A} . It is accepting since, in the above construction, the sets \tilde{S}' are chosen as subsets of \tilde{S} . \square

The converse is more involved since the obvious power-set construction produces an automaton whose acceptance condition is not a parity condition.

Definition 2.11. Let Q be a set and $\Omega : Q \rightarrow \omega$ a priority function. The *trace semigroup* is the semigroup $\mathfrak{S}_\Omega(Q) := \langle S, S_\omega \rangle$ where

$$S := \mathcal{P}(Q \times Q) \quad \text{and} \quad S_\omega := \mathcal{P}(Q).$$

For $A, B \in S$ and $P \in S_\omega$, the binary products are defined by

$$\begin{aligned} A \cdot B &:= \{ \langle p, r \rangle \mid \langle p, q \rangle \in A \text{ and } \langle q, r \rangle \in B \text{ for some } q \}, \\ A \cdot P &:= \{ p \in Q \mid q \in P \text{ for all } q \text{ with } \langle p, q \rangle \in A \}. \end{aligned}$$

Given an infinite sequence $A_o, A_1, \dots \in S$, we define the infinite product $\pi(A_o, A_1, \dots)$ as follows. We call a sequence $(p_n)_{n < \omega}$ of states a *trace* of $(A_n)_{n < \omega}$ if

$$\langle p_n, p_{n+1} \rangle \in A_n, \quad \text{for all } n < \omega.$$

The set $\pi(A_o, A_1, \dots)$ consists of all states $p \in Q$ such that every trace $(p_n)_{n < \omega}$ of $(A_n)_{n < \omega}$ starting with $p_o = p$ satisfies the parity condition Ω .

Proposition 2.12. *Let L be one of MSO, CMSO, GSO, or CGSO. For every alternating L -automaton, we can compute a nondeterministic L -automaton recognising the same language.*

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ be an alternating L -automaton. We start by constructing a nondeterministic L -automaton \mathcal{C} with set of states $Q' := \wp(Q \times Q)$ such that every run $\rho : T \rightarrow \wp(Q \times Q)$ of \mathcal{A} is also a run $T \rightarrow Q'$ of \mathcal{C} , and vice versa. This can be done by defining the transition function as

$$\delta'(A) := \bigwedge_{\langle p, q \rangle \in A} \delta(q)[S_r \mapsto \vartheta_r^q]$$

where

$$\vartheta_r^q(x) := \bigvee \{ S_{Bx} \mid B \in Q' \text{ with } \langle q, r \rangle \in B \},$$

and where we denote by $\varphi[S_r \mapsto \vartheta_r]_{r \in Q}$ the formula obtained from φ by replacing every subformula of the form $S_r x$, for $r \in Q$, by the corresponding formula $\vartheta_r(x)$.

Clearly, every run of \mathcal{A} is a run of \mathcal{C} and every run of \mathcal{C} is one of \mathcal{A} . Unfortunately, the same is not true for accepting runs. Therefore, we modify \mathcal{C} as follows.

Let $\mathfrak{S}_\Omega(Q) = \langle Q', \wp(Q) \rangle$ be the trace semigroup of Q . According to Theorem III.4.5 we can effectively construct a deterministic automaton $\mathcal{D} = \langle \tilde{Q}, Q', \tilde{\delta}, \tilde{q}_o, \tilde{\Omega} \rangle$ that recognises the language of all ω -words $w \in (Q')^\omega$ whose product $\pi(w) \in \wp(Q)$ contains the state q_o .

The automaton $\mathcal{B} = \langle Q'', \Sigma, \delta'', q_o'', \Omega'' \rangle$ is the product of \mathcal{C} and \mathcal{D} . The set of states is $Q'' := Q' \times \tilde{Q}$, the initial state is $q_o'' := \{ \langle q_o, q_o \rangle \}$, \tilde{q}_o , the priority function is

$$\Omega''(\langle A, p \rangle) := \tilde{\Omega}(p),$$

and the transition function is defined by

$$\delta''(\langle A, p \rangle) := \delta'(A)[S_B \mapsto \vartheta_B]_{B \in Q'} \wedge \forall x [x \neq \text{rt} \rightarrow \eta_{\tilde{\delta}(p, A)}(x)].$$

where

$$\vartheta_B(x) := \bigvee_{p \in \tilde{Q}} S_{\langle B, p \rangle} x \quad \text{and} \quad \eta_p(x) := \bigvee_{B \in Q'} S_{\langle B, p \rangle} x.$$

We claim that \mathcal{B} accepts the same trees as \mathcal{A} . Suppose that $\mathfrak{T} \in L(\mathcal{A})$ and let $\rho : T \rightarrow \wp(Q \times Q)$ be an accepting run of \mathcal{A} on \mathfrak{T} . Then ρ is also a run of \mathcal{C} . We define a function $\tau : T \rightarrow \tilde{Q}$ by setting

$$\tau(v) := \tilde{\delta}^*(\tilde{q}_o, v),$$

where $\tilde{\delta}^*$ is the extension of $\tilde{\delta} : \tilde{Q} \times Q' \rightarrow \tilde{Q}$ to a function $\tilde{Q} \times (Q')^* \rightarrow \tilde{Q}$. We obtain a run $\rho' : T \rightarrow Q''$ by setting

$$\rho'(v) := \langle \rho(v), \tau(v) \rangle.$$

To show that ρ' is accepting, consider a branch β of \mathfrak{T} . Since every trace of ρ satisfies the parity condition, the product

$$\pi(\rho(v))_{v \prec \beta}$$

evaluates to a set containing the state q_o . Consequently, \mathcal{D} accepts the word $(\rho(v))_{v \prec \beta}$ and the run $(\tau(v))_{v \prec \beta}$ is accepting, i.e., it satisfies the parity condition. By definition of Ω'' , it follows that $(\rho'(v))_{v \prec \beta}$ also satisfies the parity condition.

Conversely, suppose that $\mathfrak{T} \in L(\mathcal{B})$ and let $\rho : T \rightarrow Q''$ be a corresponding accepting run. Let $\rho' : T \rightarrow Q'$ and $\tau : T \rightarrow \tilde{Q}$ be the functions such that

$$\rho(v) = \langle \rho'(v), \tau(v) \rangle.$$

Then ρ' is a run of \mathcal{C} and, hence, one of \mathcal{A} . To show that it is accepting, let β be a branch of \mathfrak{T} . Since the sequence $(\rho(v))_{v \prec \beta}$ satisfies the parity condition, it follows by definition of Ω'' that so does the projection $(\tau(v))_{v \prec \beta}$. By definition of the trace semigroup, this implies that every trace of $(\rho'(v))_{v \prec \beta}$ satisfies the parity condition. \square

Equivalence to MSO

Using the closure properties established above, it is now straightforward to translate between MSO-formulae and automata.

Definition 2.13. A set of trees is *regular* if it is recognised by an MSO-automaton.

Theorem 2.14. Let L be one of MSO or CMSO. For every $L[\Sigma]$ -formula $\varphi(\bar{X})$ with monadic variables \bar{X} , we can effectively construct an alternating L -automaton \mathcal{A}_φ such that

$$\mathcal{A}_\varphi \text{ accepts a } \Sigma\text{-enriched tree } (\mathfrak{T}, \bar{P}) \quad \text{iff} \quad \mathfrak{T} \models \varphi(\bar{P}).$$

Proof. We construct \mathcal{A}_φ by induction on φ . Without loss of generality, we may assume that φ does not contain neither first-order variables, nor subformulae of the form $Y \cap Z = \emptyset$ or $\text{cover}(\bar{Z})$. Thus, there are the following cases.

If $\varphi = (X \subseteq Y)$, we use a single state q checking that $X \subseteq Y$ at each position. Hence, we set $\mathcal{A}_\varphi := \langle \{q\}, \Sigma, \delta, q, \Omega \rangle$ where $\Omega(q) := \circ$ and

$$\delta(q) := \forall x[Xx \rightarrow Yx] \wedge \forall x[x \neq \text{rt} \rightarrow S_q x].$$

If $\varphi = \text{sing}(X)$, we use two states: q looks for an element in X and $*$ makes sure that there are no other elements. Thus, we set $\mathcal{A}_\varphi := \langle \{q, *\}, \Sigma, \delta, q, \Omega \rangle$ where

$$\begin{aligned} \delta(q) &:= [X = \{\text{rt}\} \wedge \forall x(x \neq \text{rt} \rightarrow S_* x)] \\ &\quad \vee [\text{rt} \notin X \wedge \exists x[x \neq \text{rt} \wedge S_q x \wedge \forall y(y \neq \text{rt} \wedge y \neq x \rightarrow S_* y)]] \end{aligned}$$

$$\delta(*) := X = \emptyset \wedge \forall x[x \neq \text{rt} \rightarrow S_* x],$$

and $\Omega(q) := \mathbf{1}$, $\Omega(*) := \circ$.

If $\varphi = |X| < \infty$, we use an automaton \mathcal{A} with two states p, q with $\Omega(p) = \mathbf{1}$ and $\Omega(q) = \circ$, where q checks that the set X is empty in the corresponding subtree. The initial state is p and the transition formulae are

$$\begin{aligned} \delta(p) &:= \exists Z_p Z_q [Z_p \subseteq S_p \wedge Z_q \subseteq S_q \wedge \text{cover}(Z_p, Z_q, \{\text{rt}\}) \\ &\quad \wedge |Z_p| < \infty], \end{aligned}$$

$$\delta(q) := X = \emptyset \wedge \forall x[x \neq \text{rt} \rightarrow S_q x].$$

(The formula $\delta(p)$ is a bit complicated since it has to be monotone in S_p .)

If $\varphi = |X| \equiv k \pmod{m}$, we use an automaton \mathcal{A} with states $Q := \{*\} + [m]$ where $\Omega(*) = \circ$ and $\Omega(i) = \mathbf{1}$, for $i \in [m]$. The state $*$ checks that the set X is empty in the corresponding subtree while a state $i \in [m]$ checks that $|X| \equiv i \pmod{m}$ in this subtree. The initial state is k and the transition formulae are

$$\delta(*) := X = \emptyset \wedge \forall x[x \neq \text{rt} \rightarrow S_* x],$$

$$\begin{aligned} \delta(i) &:= \exists \bar{Z} \left[\bigwedge_{q \in Q} Z_q \subseteq S_q \wedge \text{cover}(\bar{Z}, \{\text{rt}\}) \right. \\ &\quad \left. \wedge [(\text{rt} \notin X \wedge \text{Mod}_i(\bar{Z})) \vee (\text{rt} \in X \wedge \text{Mod}_{i-1}(\bar{Z}))] \right], \end{aligned}$$

where

$$\text{Mod}_i(\bar{Z}) := \bigvee_g \bigwedge_{j < m} |Z_j| \equiv g(j) \pmod{m}$$

and the disjunction in this formula ranges over all functions $g : [m] \rightarrow [m]$ such that

$$g(\circ) + \dots + g(m-1) \equiv i \pmod{m}.$$

Suppose that $\varphi = X \leq_{\text{pf}} Y$. This formula is equivalent to

$$\text{sing}(X) \wedge \text{sing}(Y) \wedge \exists x y [Xx \wedge Yy \wedge x \leq_{\text{pf}} y].$$

Since we have already constructed an automaton recognising $\text{sing}(X)$ and since automata are closed under intersection, it is therefore sufficient to find an automaton for the formula $\exists x y [Xx \wedge Yy \wedge x \leq_{\text{pf}} y]$. We use two states: q looks for an element in X and p for one in Y . Hence, we set $\mathcal{A} := \langle \{q, p\}, \Sigma, \delta, q, \Omega \rangle$ where

$$\begin{aligned} \delta(q) &:= X = \{\text{rt}\} = Y \vee [X = \{\text{rt}\} \wedge \exists x(x \neq \text{rt} \wedge S_p x) \\ &\quad \vee \exists x[x \neq \text{rt} \wedge S_q x], \end{aligned}$$

$$\delta(p) := Y = \{\text{rt}\} \vee \exists x[x \neq \text{rt} \wedge S_p x],$$

and $\Omega(q) := \mathbf{1}$, $\Omega(p) := \mathbf{1}$.

If $\varphi = RX_0 \dots X_{n-1}$ with $R \in \Sigma$, it is again sufficient to find an automaton for the formula

$$\exists \bar{x} \left[\bigwedge_i X_i x_i \wedge R \bar{x} \right].$$

We set $\mathcal{A}_\varphi := \langle \{q\}, \Sigma, \delta, q, \Omega \rangle$ where $\Omega(q) := \mathbf{1}$ and

$$\delta(q) := \exists \bar{x} \left[R \bar{x} \wedge \bigwedge_{i < n} x_i \in X_i \right] \vee \exists x [x \neq \text{rt} \wedge S_q x].$$

Suppose that $\varphi = \psi \vee \vartheta$. By inductive hypothesis, we can construct automata \mathcal{A}_ψ and \mathcal{A}_ϑ for ψ and ϑ . Hence, the desired automaton for φ can be obtained via Proposition 2.2.

Suppose that $\varphi = \neg\psi$. By inductive hypothesis, we can construct an automaton for \mathcal{A}_ψ . Hence, the desired automaton for φ can be obtained via Proposition 2.9.

Finally, suppose that $\varphi = \exists X \psi$. By inductive hypothesis, we can construct an automaton for \mathcal{A}_ψ . Since we can translate between alternating automata and nondeterministic automata, we can therefore obtain the desired automaton for φ by Proposition 2.3. \square

Theorem 2.15. *Let L be one of MSO or CMSO. For every alternating L -automaton \mathcal{A} over the signature Σ , we can effectively construct an $L[\Sigma]$ -formula $\varphi_{\mathcal{A}}$ such that*

$$\mathfrak{T} \models \varphi \quad \text{iff} \quad \mathcal{A} \text{ accepts } \mathfrak{T}.$$

Proof. Let $\mathcal{A} := \langle Q, \Sigma, \delta, q_0, \Omega \rangle$. The formula $\varphi_{\mathcal{A}}$ guesses sets encoding a run of the automaton and then checks that the guessed run is accepting. We set

$$\varphi_{\mathcal{A}} := \exists (Z_{p,q})_{p,q \in Q} [\text{INIT} \wedge \text{TRANS} \wedge \text{ACC}]$$

where we use the following formulae.

$$\text{INIT} := \exists x [Z_{q_0, q_0} x \wedge \forall y (x \leq_{\text{pf}} y)]$$

states that the root contains the pair $\langle q_0, q_0 \rangle$.

$$\text{TRANS} := \forall x \bigwedge_{p,q \in Q} [Z_{p,q} x \rightarrow \hat{\delta}(q)(x)]$$

states that at every vertex a correct transition is used. Here, $\hat{\delta}(q)(x)$ denotes the restriction of the formula $\delta(q)$ to the set $\text{Suc}_*(x)$.

$$\text{ACC} := \forall Y [\text{BRANCH}(Y) \rightarrow \text{PARITY}(\bar{Z}, Y)]$$

checks the parity condition. The formula

$$\begin{aligned} \text{BRANCH}(Y) := \\ Y \neq \emptyset \wedge \forall x \forall y [Yx \wedge Yy \rightarrow (x \leq_{\text{pf}} y \vee y \leq_{\text{pf}} x)] \\ \wedge \forall x \exists y [Yx \rightarrow x <_{\text{pf}} y \wedge Yy] \end{aligned}$$

states that the elements in Y form an infinite branch and

$$\begin{aligned} \text{PARITY}(\bar{Z}, Y) := \\ \forall (X_{p,q})_{p,q \in Q} [\text{TRACE}(\bar{Z}, \bar{X}, Y) \rightarrow \bigvee_{k < n} \text{MIN}_{2k}(\bar{X}, Y)] \end{aligned}$$

states that every trace for the branch Y satisfies the parity condition. Here, n is any number such that the maximal priority of \mathcal{A} is smaller than $2n$, the formula

$$\begin{aligned} \text{TRACE}(\bar{Z}, \bar{X}, Y) := \\ \bigwedge_{p,q} [X_{p,q} \subseteq Z_{p,q} \cap Y] \wedge \forall x [Yx \rightarrow \bigvee_{p,q \in Q} X_{p,q} x] \end{aligned}$$

states that the sets \bar{X} encode a trace of the branch Y , and the formula

$$\text{MIN}_k(\bar{X}, Y) := \bigvee_{q \in H_k} \text{INF}_q(\bar{X}, Y) \wedge \bigwedge_{q \in H_{k-1}} \neg \text{INF}_q(\bar{X}, Y)$$

states that the minimal priority seen infinitely often in the trace encoded by X is equal to k . The set

$$H_k := \{ q \in Q \mid \Omega(q) \leq k \}$$

contains all states with priority at most k and the formula

$$\text{INF}_q(\tilde{X}, Y) := \forall x \exists y \left[x \leq y \wedge Yy \wedge \bigvee_{p \in Q} X_{p,q}y \right]$$

states that the trace contains infinitely many occurrences of the state q . \square

Regular Trees

As an application of the equivalence between automata and MSO-formulae, we prove that every non-empty regular tree language contains a regular tree.

Definition 2.16. A tree \mathfrak{T} is *regular* if, up to isomorphism, it has only finitely many subtrees. ,

Exercise 2.4. Prove that a finitely-branching tree is regular if, and only if, it is the unravelling of a finite directed graph. ,

Theorem 2.17. Let L be a regular language of enriched trees over the signature Σ . If L is non-empty, it contains a regular tree. Furthermore, if Σ consists of only unary predicates, this tree can be effectively constructed from a given automaton recognising L .

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ be an MSO-automaton recognising L and let m be the maximal quantifier-rank of the formulae $\delta(q)$, $q \in Q$. We use a variant of the automaton game where Player \Diamond not only chooses the next transition but also the corresponding input letter. Let $\tilde{S} = (S_q)_{q \in Q}$ be unary predicates encoding the states of \mathcal{A} . We denote by \mathcal{U}_n the class of all enriched trees \mathbb{U} over the signature $\Sigma + \tilde{S} + \{\text{rt}\}$ whose height is at most n and such that there is exactly one state $q \in Q$ with

$$\text{rt}^{\mathbb{U}} \in S_q^{\mathbb{U}}.$$

Let $\mathcal{U}_n^o \subseteq \mathcal{U}_n$ be a set of representatives containing exactly one structure of each MSO_m -theory. The game uses positions

$$V_{\Diamond} := \{*\} + Q \times \mathcal{U}_o^o \quad \text{and} \quad V_{\square} := \mathcal{U}_1^o.$$

The initial position is $* \in V_{\Diamond}$. In position $*$, Player \Diamond chooses some $\mathfrak{B} \in \mathcal{U}_1^o$ such that

$$\mathfrak{B} \models S_{q_o}(\text{rt}) \wedge \delta(q_o),$$

while, in position $\langle q, \mathbb{U} \rangle$, he chooses $\mathfrak{B} \in \mathcal{U}_1^o$ satisfying

$$\mathfrak{B} \models S_q(\text{rt}) \wedge \delta(q) \quad \text{and} \quad \mathfrak{B}|_{\{\text{rt}\}} \cong \mathbb{U}.$$

Player \square responds with some vertex $v \in V$ different from the root and some state $p \in Q$ with $v \in S_p$. The game continues in position $\langle p, \mathfrak{B}' \rangle$ where \mathfrak{B}' is the structure obtained from $\mathfrak{B}|_{\{v\}}$ by removing the root rt from all predicates S_q with $q \neq p$. Finally, a play is winning if the corresponding sequence of states satisfies the parity condition.

Note that a winning strategy for Player \Diamond is determined by a tree \mathfrak{T} and an accepting run of \mathcal{A} on \mathfrak{T} . Conversely, Player \square wins the game if no tree has an accepting run. Since $L \neq \emptyset$, it follows that Player \Diamond has a positional winning strategy σ . Let \mathfrak{G} be the graph obtained from the game by removing all edges that do not correspond to σ . The unravelling of \mathfrak{G} is a regular tree accepted by \mathcal{A} .

Finally, note that the construction of \mathfrak{G} is effective, provided that we can compute the sets of representatives \mathcal{U}_o and \mathcal{U}_1 . If all predicates in Σ are unary, we can do so since it is decidable whether an MSO_m -theory is satisfied by some extended tree of height 0 or 1 over Σ . (The easiest way to see this is using Theorem 3.5 below. Alternatively one can use a direct composition argument to build a tree using disjoint unions.) \square

Corollary 2.18. Let $m < \omega$. For every tree \mathfrak{T} , there exists a regular tree \mathfrak{S} with $\mathfrak{S} \equiv_{\text{CMSO}}^m \mathfrak{T}$.

Proof. Let $\theta := \text{Th}_{\text{MSO}}^m(\mathfrak{T})$ and set $\varphi := \bigwedge \theta$. The language of all trees satisfying φ is regular and, therefore, it contains some regular tree. \square

3 The Muchnik Iteration

Having defined tree automata, we now can use them to study decidability questions. We consider the theory of enriched trees of the following kind.

Definition 3.1. Let $\mathfrak{A} = \langle A, \bar{R} \rangle$ be a Σ -structure. The *Muchnik iteration* of \mathfrak{A} is the $(\Sigma + \{\text{suc}, \text{cl}\})$ -structure

$$\mathfrak{A}^* := \langle A^*, \bar{R}^*, \text{suc}, \text{cl} \rangle$$

where

$$\begin{aligned} R_i^* &:= \{ \langle wa_0, \dots, wa_{n-1} \rangle \mid w \in A^*, \bar{a} \in R_i \}, \\ \text{suc} &:= \{ \langle w, wa \rangle \mid w \in A^*, a \in A \}, \\ \text{cl} &:= \{ \langle waa \mid w \in A^*, a \in A \}. \end{aligned}$$

The predicate cl is called the *clone relation*. J

Remark. Note that \mathfrak{A}^* is a $(\Sigma + \{\text{cl}\})$ -enriched tree. Hence, we can use the automata-theoretic tools from the previous section to study the MSO-theory of such structures. J

Example. The unravelling of a directed graph $\mathfrak{G} = \langle V, E \rangle$ is MSO-interpretable in the Muchnik iteration \mathfrak{G}^* . To do so, we only need to find a formula $\varphi(x)$ stating that a sequence $x \in V^*$ of vertices corresponds to a path of \mathfrak{G} . Such a formula is given by

$$\varphi(x) := \forall y [y \leq_{\text{pf}} x \rightarrow \exists z (\text{cl}(z) \wedge Ezy)].$$
J

To show that the Muchnik iteration is MSO-compatible, we can use tree automata and the corresponding automaton games. Since trees of the form \mathfrak{A}^* are very regular, we can simplify these games considerably. In particular, we do not need to remember the precise vertex $v \in A^*$ we are in, but only its last letter.

Definition 3.2. Let $\mathcal{A} := \langle Q, \Sigma + \{\text{cl}\}, \delta, q_0, \Omega \rangle$ be an alternating L -automaton and let \mathfrak{A} be a Σ -structure. The *reduced game*

$$\mathcal{G}_o(\mathcal{A}, \mathfrak{A}) = \langle V_\diamond, V_\square, E, \Omega' \rangle$$

has positions

$$V_\diamond := A \times Q + \{ \langle \rangle, q_0 \} \quad \text{and} \quad V_\square := \wp(A \times Q).$$

The initial position is $\langle \rangle, q_0 \rangle$ and the edge relation E is defined as follows. From a V_\square -position H there are edges to all V_\diamond -positions $\langle b, q \rangle \in H$. From a V_\diamond -position $\langle a, q \rangle$ with $a \in A \cup \{ \langle \rangle \}$, there are edges to every V_\square -position H satisfying

$$\langle \mathfrak{A} \oplus a, \bar{S}, a \rangle \models \delta(q), \quad \text{where} \quad S_p := \{ b \in A \mid \langle b, p \rangle \in H \}$$

and $\mathfrak{A} \oplus a$ denotes the substructure $\mathfrak{Succ}_*(a)$ of \mathfrak{A}^* , that is, $\mathfrak{A} \oplus a$ is the disjoint union of \mathfrak{A} and a singleton structure $\mathfrak{A}|_{\{a\}}$ expanded by the successor relation suc and the clone predicate cl . Similarly, for $a = \langle \rangle$, we denote by $\mathfrak{A} \oplus \langle \rangle$ the substructure of \mathfrak{A}^* consisting of the root and its immediate successors.

The priority function is defined in the same way as for $\mathcal{G}(\mathcal{A}, \mathfrak{A}^*)$, i.e., we set

$$\Omega'(\langle v, q \rangle) := \Omega(q), \quad \text{for } \langle v, q \rangle \in V_\diamond,$$

while $\Omega'(H)$, for $H \in V_\square$, is an arbitrary number larger than all priorities used by \mathcal{A} . J

Lemma 3.3. Let \mathcal{A} be an alternating L -automaton and let \mathfrak{A} be a Σ -structure. Automaton has a winning strategy for $\mathcal{G}_o(\mathcal{A}, \mathfrak{A})$ if, and only if, he has one for $\mathcal{G}(\mathcal{A}, \mathfrak{A}^*)$.

Proof. Suppose that $\mathcal{A} = \langle Q, \Sigma \cup \{\text{cl}\}, \delta, q_0, \Omega \rangle$,

$$\mathcal{G}(\mathcal{A}, \mathfrak{A}^*) = \langle V_\diamond, V_\square, E, \Omega \rangle \quad \text{and} \quad \mathcal{G}_o(\mathcal{A}, \mathfrak{A}) = \langle V_\diamond^\circ, V_\square^\circ, E^\circ, \Omega^\circ \rangle.$$

Let $r : A^+ \rightarrow A$ be the function mapping a word to its last letter:

$$r(va) := a.$$

We define a projection function $h : V \rightarrow V^\circ$ by applying r to the first component of every position.

$$h(\langle v, q \rangle) := \begin{cases} \langle \langle \rangle, q_\circ \rangle & \text{if } v = \langle \rangle, \\ \langle r(v), q \rangle & \text{otherwise,} \end{cases} \quad \text{for } \langle v, q \rangle \in V_\diamond,$$

$$h(H) := \{ \langle r(v), q \rangle \mid \langle v, q \rangle \in H \}, \quad \text{for } H \in V_\square.$$

(Note that, in a play of $\mathcal{G}(\mathcal{A}, \mathfrak{A}^*)$, we will never see a position of the form $\langle \langle \rangle, q \rangle$ with $q \neq q_\circ$. Hence, in the above definition the replacement of q by q_\circ is harmless.)

The function h is a homomorphism from $\mathcal{G}(\mathcal{A}, \mathfrak{A}^*)$ to $\mathcal{G}_\circ(\mathcal{A}, \mathfrak{A})$. In particular, it is an immersion. Consequently, the claim follows by Lemma IV.5.6. \square

We can use the reduced game $\mathcal{G}_\circ(\mathcal{A}, \mathfrak{A})$ to compute the winning region inside the structure \mathfrak{A} .

Lemma 3.4. *Let \mathcal{A} be an alternating L -automaton where L is one of MSO or CMSO, and let \mathfrak{A} be a Σ -structure. Given an MSO-formula φ where each set quantifier ranges only over subsets of V_\diamond , we can compute an $L[\Sigma]$ -formula φ' such that*

$$\mathcal{G}_\circ(\mathcal{A}, \mathfrak{A}) \models \varphi \quad \text{iff} \quad \mathfrak{A} \models \varphi'.$$

Proof. Suppose that $\mathcal{A} = \langle \mathcal{Q}, \Sigma + \{\leq_{\text{pf}}, \text{cl}\}, \delta, q_\circ, \Omega \rangle$ and set $s := |\mathcal{Q}|$. Recall that $\mathcal{G}_\circ(\mathcal{A}, \mathfrak{A}) = \langle V_\diamond, V_\square, E, \Omega' \rangle$ where

$$V_\diamond = A \times \mathcal{Q} + \{ \langle \langle \rangle, q_\circ \rangle \},$$

$$V_\square = \wp(A \times \mathcal{Q}) = \wp(V_\diamond \setminus \{ \langle \langle \rangle, q_\circ \rangle \}),$$

$$E = \{ \langle \langle a, q \rangle, S \rangle \in V_\diamond \times V_\square \mid \mathfrak{A} \oplus a, (S_p)_p, a \models \delta(q) \}$$

$$\cup \{ \langle S, \langle a, q \rangle \rangle \in V_\square \times V_\diamond \mid \langle a, q \rangle \in S \},$$

where we have used the notation $S_p := \{ b \in A \mid \langle b, p \rangle \in S \}$.

To reduce this structure to \mathfrak{A} , we construct as an intermediate step the 2-sorted structure

$$\mathfrak{B} := \langle \wp(V_\diamond), V_\square, E', \bar{P} \rangle$$

where

$$E' := \{ \langle \{u\}, v \rangle \mid \langle u, v \rangle \in E \cap (V_\diamond \times V_\square) \}$$

$$\cup \{ \langle u, \{v\} \rangle \mid \langle u, v \rangle \in E \cap (V_\square \times V_\diamond) \},$$

$$P_i := \{ \{v\} \in \wp(V_\diamond) \mid \Omega(v) = i \} \cup \{ v \in V_\square \mid \Omega(v) = i \}.$$

That is, \mathfrak{B} is obtained from $\mathcal{G}(\mathcal{A}, \mathfrak{A})$ by applying the power-set operation \wp , but only to the positions of Player \diamond . Since the set quantifiers in φ only range over subsets of V_\diamond , we can use the construction in Proposition I.2.2 to translate φ to an FO-formula ψ satisfying

$$\mathcal{G}_\circ(\mathcal{A}, \mathfrak{A}) \models \varphi \quad \text{iff} \quad \mathfrak{B} \models \psi.$$

To conclude the proof, we shall construct an (L, FO) -compatible operation σ and an FO-interpretation τ such that

$$\mathfrak{B} \cong \tau(\sigma(\mathfrak{A})).$$

Then it follows that we can compute an L -formula φ' such that

$$\mathfrak{A} \models \varphi' \quad \text{iff} \quad \sigma(\mathfrak{A}) \models \psi^\tau \quad \text{iff} \quad \mathfrak{B} \models \psi \quad \text{iff} \quad \mathcal{G}_\circ(\mathcal{A}, \mathfrak{A}) \models \varphi,$$

as desired. Hence, it remains to construct σ and τ . The definition depends on the logic L .

First, suppose that $L = \text{MSO}$. We set

$$\sigma(\mathfrak{A}) := \wp(1 \oplus \text{copy}_{2s}(\mathfrak{A})).$$

This operation is (MSO, FO) -compatible since $1 \oplus -$ and copy_{2s} are MSO-compatible and \wp is (MSO, FO) -compatible. Hence, it remains to define τ .

Note that, given $\mathfrak{A} = \langle A, \tilde{R} \rangle$, we have

$$\begin{aligned} \text{copy}_{2s}(\mathfrak{A}) &= \langle [2s] \times A, \tilde{R}, \tilde{H}, I \rangle, \\ 1 \oplus \text{copy}_{2s}(\mathfrak{A}) &= \langle \{*\} + [2s] \times A, \tilde{R}, \tilde{H}, I, \text{Left}, \text{Right} \rangle, \\ \sigma(\mathfrak{A}) &= \langle \wp(\{*\} + [2s] \times A), \tilde{R}, \tilde{H}, I, \text{Left}, \text{Right}, \subseteq \rangle. \end{aligned}$$

where the predicates H_i denote the various copies of A , I is the binary relation indicating which elements are copies of the same element of A , and the predicates Left and Right denote the two parts of the disjoint union. We can identify the universe $[2s] \times A$ of $\text{copy}_{2s}(\mathfrak{A})$ with $(Q + Q) \times A$. We use the first copy of $Q \times A$ to encode the elements of V_\diamond and the second copy for V_\square . The unique element $*$ of 1 represents the initial position $\langle \langle \rangle, q_o \rangle$. Using this encoding we can express that an element x is a subset of $V_\diamond \setminus \{\langle \langle \rangle, q_o \rangle\}$ by the formula

$$\vartheta_\diamond(x) := \text{Right}(x) \wedge \forall z [\text{sing}(z) \wedge z \subseteq x \rightarrow \bigvee_{q \in Q} H_q z].$$

we can say that x is the initial position by

$$\vartheta_*(x) := \text{Left}(x) \wedge \text{sing}(x),$$

and that x is a position for Player \square by

$$\vartheta_\square(x) := \text{Right}(x) \wedge \forall z [\text{sing}(z) \wedge z \subseteq x \rightarrow \bigvee_{q \in Q} H_{q+s} z].$$

Using these formulae and the transition formulae $\delta(q)$ from \mathcal{A} , we can now define the edge relation E' of \mathfrak{B} by the FO-formula

$$\begin{aligned} \varphi_E(x, y) &:= [\vartheta_\square(x) \wedge \vartheta_\diamond(y) \wedge \text{sing}(y) \wedge \exists y' (Iy y' \wedge y' \subseteq x)] \\ &\quad \vee [\vartheta_\diamond(x) \wedge \text{sing}(x) \wedge \vartheta_\square(y) \wedge \bigvee_{q \in Q} [H_q x \wedge \hat{\delta}_q^\circ(x, y)]] \\ &\quad \vee [\vartheta_*(x) \wedge \vartheta_\square(y) \wedge \hat{\delta}_{q_o}^1(y)], \end{aligned}$$

where the formulae $\hat{\delta}_q^\circ(x, y)$ and $\hat{\delta}_q^1(y)$ are defined as follows. By Proposition I.2.2, there exists an FO-formula $\delta'_q(x, \bar{y})$ such that

$$\wp(\mathfrak{A}) \models \delta'_q(\{a\}, \bar{S}) \quad \text{iff} \quad \langle \mathfrak{A} \oplus a, \bar{S}, a \rangle \models \delta(q).$$

Since $\wp(\mathfrak{A})$ can be interpreted in $\wp(1 \oplus \text{copy}_{2s}(\mathfrak{A}))$, we obtain an FO-formula $\delta''_q(x, \bar{y})$ such that

$$\wp(1 \oplus \text{copy}_{2s}(\mathfrak{A})) \models \delta''_q(\{\langle p, a \rangle\}, \bar{S}) \quad \text{iff} \quad \wp(\mathfrak{A}) \models \delta'_q(\{a\}, \bar{S}'),$$

where $S'_i := \{b \in A \mid \langle q, b \rangle \in S_i \text{ for some } q \in Q\}$. Setting

$$\begin{aligned} \hat{\delta}_q^\circ(x, y) &:= \exists \bar{y}' [\forall z [\text{sing}(z) \rightarrow \bigwedge_{q' \in Q} [z \subseteq y'_{q'} \leftrightarrow z \subseteq y \wedge H_q z]]] \\ &\quad \wedge \delta''_q(x, \bar{y}'), \end{aligned}$$

it follows that

$$\wp(1 \oplus \text{copy}_{2s}(\mathfrak{A})) \models \hat{\delta}_q^\circ(\{\langle p, a \rangle\}, \bar{S}) \quad \text{iff} \quad \langle \mathfrak{A} \oplus a, \bar{S}, a \rangle \models \delta(q).$$

In the same way we obtain an FO-formula $\hat{\delta}_q^1(y)$ such that

$$\wp(1 \oplus \text{copy}_{2s}(\mathfrak{A})) \models \hat{\delta}_q^1(\bar{S}) \quad \text{iff} \quad \langle \mathfrak{A} \oplus \langle \rangle, \bar{S} \rangle \models \delta(q).$$

The desired FO-interpretation τ mapping \mathfrak{B} to \mathfrak{G} is given by

$$\tau := \langle \chi_\diamond(x), \chi_\square(x), \varphi_E(x, y), (\varphi_{P_k})_k \rangle,$$

where

$$\begin{aligned} \chi_\diamond(x) &:= \vartheta_*(x) \vee \vartheta_\diamond(x), \\ \chi_\square(x) &:= \vartheta_\square(x), \\ \varphi_{P_k}(x) &:= \begin{cases} \vartheta_\square(x) & \text{if } k \text{ is the maximal priority,} \\ \vartheta_\diamond(x) \wedge \bigvee_{q \in \Omega^{-1}(k)} H_q x & \text{otherwise.} \end{cases} \end{aligned}$$

It remains to consider the case where $L = \text{CMSO}$. Let m be the least common multiple of all numbers m' such that some transition formula $\delta(q)$ contains a subformula of the form $|X| \equiv k \pmod{m'}$. We set

$$\sigma(\mathfrak{A}) := \langle \wp(1 \oplus \text{copy}_{2s}(\mathfrak{A})), \bar{M} \rangle,$$

where

$$M_k := \{ S \mid |S| \equiv k \pmod{m} \}.$$

Then the proof proceeds as for the case of MSO except that we use Proposition I.2.4 instead of Proposition I.2.2. \square

Theorem 3.5 (Muchnik). *Let Σ be a finite relational signature and L one of MSO or CMSO. For every formula $\varphi \in L[\Sigma + \{\text{succ}, \text{cl}\}]$, we can effectively construct a formula $\varphi^* \in L[\Sigma]$ such that*

$$\mathfrak{A}^* \models \varphi \quad \text{iff} \quad \mathfrak{A} \models \varphi^*, \quad \text{for all } \Sigma\text{-structures } \mathfrak{A}.$$

Proof. Given the formula φ , we can use Theorem 2.14 to construct an L -automaton \mathcal{A} such that

$$\mathfrak{A}^* \models \varphi \quad \text{iff} \quad \mathcal{A} \text{ accepts } \mathfrak{A}^*.$$

By Proposition 2.6 and Lemma 3.3, the latter is equivalent to Automaton having a winning strategy for the game $\mathcal{G}_o(\mathcal{A}, \mathfrak{A})$. Note that this game is alternating and normalised. Hence, we can use Proposition IV.5.10 to construct an L_μ -formula χ such that

$$\text{Automaton wins } \mathcal{G}_o(\mathcal{A}, \mathfrak{A}) \quad \text{iff} \quad \mathcal{G}_o(\mathcal{A}, \mathfrak{A}), \langle \langle \rangle, q_o \rangle \models \chi.$$

Let $\chi'(x)$ be the translation of χ into MSO. Since we are only interested in which Automaton positions belong to the fixed-points computed by χ , we can choose χ' such that all set quantifiers range over subsets of V_o . By Lemma 3.4, there therefore exists an L -formula φ^* such that

$$\mathcal{G}_o(\mathcal{A}, \mathfrak{A}) \models \chi'(\langle \langle \rangle, q_o \rangle) \quad \text{iff} \quad \mathfrak{A} \models \varphi^*.$$

It follows that

$$\begin{aligned} \mathfrak{A}^* \models \varphi & \quad \text{iff} \quad \text{Automaton wins } \mathcal{G}_o(\mathcal{A}, \mathfrak{A}) \\ & \quad \text{iff} \quad \mathcal{G}_o(\mathcal{A}, \mathfrak{A}) \models \chi'(\langle \langle \rangle, q_o \rangle) \\ & \quad \text{iff} \quad \mathfrak{A} \models \varphi^*. \end{aligned} \quad \square$$

Corollary 3.6. *The Muchnik iteration is L -bicompatible, for $L = \text{MSO}$ and for $L = \text{CMSO}$.*

Proof. Compatibility has already been shown in Theorem 3.5. The converse direction follows from the fact that there exists an FO-interpretation mapping \mathfrak{A}^* to \mathfrak{A} . Note that \mathfrak{A}^* consists of countably many copies of \mathfrak{A} arranged in a tree and that we can define the copy attached to the root in \mathfrak{A}^* . Consequently, we can use the interpretation with formulae

$$\delta(x) := |x| = 1 \quad \text{and} \quad \varphi_R(\bar{x}) := R\bar{x}. \quad \square$$

In the tree \mathfrak{A}^* the root and the first level are special since there is no clone relation. For technical reasons, we sometimes need to consider a variant of \mathfrak{A}^* where the clone relation is also defined for these vertices. That is, we consider a subtree $\mathfrak{A}^*|_{\uparrow wa}$ where $w \in A^+$, $a \in A$, and $\uparrow wa = waA^*$. The following version of the Theorem of Muchnik is proved in the same way as Theorem 3.5, the only difference being that we start the game $\mathcal{G}_o(\mathcal{A}, \mathfrak{A})$ from the position $\langle a, q_o \rangle$ instead of $\langle \langle \rangle, q_o \rangle$.

Theorem 3.7. *Let Σ be a finite relational signature and L one of MSO or CMSO. For every formula $\varphi \in L[\Sigma + \{\text{succ}, \text{cl}\}]$, we can effectively construct a formula $\varphi^*(x) \in L[\Sigma]$ such that*

$$\begin{aligned} \mathfrak{A}^*|_{wa} \models \varphi & \quad \text{iff} \quad \mathfrak{A} \models \varphi^*(a), \quad \text{for all } \Sigma\text{-structures } \mathfrak{A} \\ & \quad \text{and all } w \in A^+. \end{aligned}$$

The Theorem of Muchnik is one of the strongest decidability results for MSO known. Let us collect a few immediate consequences, the most well-known being the result that the MSO-theory of the infinite binary tree is decidable.

Definition 3.8. The *infinite binary tree* is the successor-tree

$$\mathfrak{T}_{\text{bin}} := \langle \{0, 1\}^*, \text{suc}_0, \text{suc}_1 \rangle.$$

Corollary 3.9 (Rabin). *The MSO-theory of the infinite binary tree $\mathfrak{T}_{\text{bin}}$ is decidable.*

Proof. We consider the two-element structure $\mathfrak{A} := \langle \{0, 1\}, P_0, P_1 \rangle$ where $P_c := \{c\}$. The binary tree $\mathfrak{T}_{\text{fin}}$ can be obtained from

$$\mathfrak{A}^* = \langle \{0, 1\}^*, \text{suc}, \text{cl}, P_0, P_1 \rangle$$

by an FO-interpretation τ . It follows that

$$\mathfrak{T}_{\text{bin}} \models \varphi \quad \text{iff} \quad \mathfrak{A}^* \models \varphi^\tau \quad \text{iff} \quad \mathfrak{A} \models \varphi^*,$$

where φ^* is the formula obtained from φ^τ via Theorem 3.5. As \mathfrak{A} is finite, the latter property is decidable. \square

Corollary 3.10. $\text{Th}_{\text{MSO}}(\mathbb{Q}, \leq)$ is decidable.

Proof. It is sufficient to construct an interpretation

$$\tau = \langle \delta(x), \varphi_{\leq}(x, y) \rangle$$

mapping the binary tree $\mathfrak{T}_{\text{bin}}$ to the order of the rationals $\mathfrak{Q} := \langle \mathbb{Q}, \leq \rangle$. We take for the universe all words ending in the letter 1 and for the order \leq the lexicographic ordering. Formally, we set

$$\begin{aligned} \delta(x) &:= \exists y [\text{suc}_1(y, x)], \\ \varphi(x, y) &:= x \leq_{\text{pf}} y \vee \exists z \exists u \exists v [\text{suc}_0(z, u) \wedge \text{suc}_1(z, v) \\ &\quad \wedge u \leq_{\text{pf}} x \wedge v \leq_{\text{pf}} y], \end{aligned}$$

where \leq_{pf} can be computed by an MSO-formula as the transitive closure of $\text{suc}_0 \cup \text{suc}_1$. \square

Corollary 3.11. *The unravelling operation on directed graphs is MSO-compatible.*

Proof. As shown in the Example on page 228 we can interpret the unravelling of a graph \mathfrak{G} in its iteration \mathfrak{G}^* . \square

Finally, let us note that Muchnik iterations commute with interpretations.

Lemma 3.12. *For every MSO-interpretation τ , there exists an MSO-interpretation σ such that*

$$\tau(\mathfrak{A})^* = \sigma(\mathfrak{A}^*).$$

Proof. Given $\tau = \langle \delta(x), (\varphi_R(\bar{x}))_R \rangle$, we set

$$\sigma := \langle \delta'(x), \varphi'_{\text{suc}}(x, y), \varphi'_{\text{cl}}(x), (\varphi'_R(\bar{x}))_R \rangle,$$

where

$$\begin{aligned} \delta'(x) &:= \forall y \forall z [\text{suc}(y, z) \wedge z \leq_{\text{pf}} x \rightarrow \delta^y(z)], \\ \varphi'_{\text{suc}}(x, y) &:= \text{suc}(x, y), \\ \varphi'_{\text{cl}}(x) &:= \text{cl}(x), \\ \varphi'_R(\bar{x}) &:= \exists y \left[\bigwedge_i \text{suc}(y, x_i) \wedge \varphi_R^y(\bar{x}) \right], \end{aligned}$$

where ψ^y denotes the relativisation of ψ to the set $\{z \mid \text{suc}(y, z)\}$. \square

4 Löwenheim-Skolem Theorems

We can use tree automata to derive a number of Löwenheim-Skolem theorems for MSO over trees. Our key argument is contained in the following lemma.

Lemma 4.1. *Let \mathcal{A} be an MSO-automaton and \mathfrak{T} a coloured tree accepted by \mathcal{A} . There exist a number $N < \omega$ and a family $(W_v)_{v \in T}$ of subsets $W_v \subseteq T$ with the following properties.*

- $W_v \subseteq \text{Suc}(v)$ and $|W_v| < N$ for all $v \in T$.

- ♦ \mathcal{A} accepts $\mathfrak{T}|_P$, for every non-empty prefix-closed set $P \subseteq T$ that is closed under $(W_v)_v$ in the sense that $v \in P \Rightarrow W_v \subseteq P$.

Proof. Suppose that $\mathcal{A} = \langle Q, \Sigma, \delta, q_o, \Omega \rangle$ is non-deterministic and fix an accepting run ρ of \mathcal{A} on \mathfrak{T} . To define W_v , we fix a vertex $v \in T$. By definition of an accepting run, we know that

$$\mathfrak{Succ}_*(v; \rho) \models \delta(\rho(v)).$$

Let m be the maximal quantifier rank of the transition formulae $\delta(q)$, $q \in Q$, and set $K := 2^{m-1}$. We have shown in Proposition I.4.3 (b) that, for structures X and Y with the empty signature,

$$X \equiv_{\text{MSO}}^m Y \quad \text{iff} \quad |X| = |Y| \text{ or } |X|, |Y| \geq K.$$

Set $S_q := \rho^{-1}(q) \cap \text{Suc}(v)$. For every state $q \in Q$, we select a set $X_q \subseteq S_q$ as follows. If $|S_q| \leq K$, we set $X_q := S_q$. Otherwise, we choose an arbitrary subset $X_q \subseteq S_q$ of size $|X_q| = K$. Then we set

$$W_v := \bigcup_{q \in Q} X_q.$$

By construction, it follows that $|W_v| \leq K \cdot |Q| =: N$ and

$$\langle S_q, S_q \rangle \equiv_{\text{MSO}}^m \langle X_q, X_q \rangle, \quad \text{for every } q \in Q.$$

As $\mathfrak{Succ}_*(v; \rho)$ can be obtained from the disjoint union of these structures by adding a root, it follows by Proposition I.4.2 that

$$\mathfrak{Succ}_*(v; \rho) \equiv_{\text{MSO}}^m \mathfrak{Succ}_*(v; \rho)|_{W_v \cup \{v\}}.$$

Let $P \subseteq T$ be a prefix closed subset satisfying $v \in P \Rightarrow W_v \subseteq P$. We claim that the restriction $\rho|_P$ is an accepting run of \mathcal{A} on $\mathfrak{T}|_P$. Obviously, every infinite branch of $\rho|_P$ is an infinite branch of ρ and, hence, satisfies the parity condition. Furthermore, for every vertex $v \in P$,

$$\mathfrak{Succ}_*(v; \rho) \models \delta(\rho(v)) \quad \text{implies} \quad \mathfrak{Succ}_*(v; \rho)|_P \models \delta(\rho(v)). \quad \square$$

As an immediate consequence we obtain two downward Löwenheim-Skolem theorems, one for single formulae and one for MSO-theories.

Theorem 4.2. *For every MSO-formula φ , there exists a finite number $N < \omega$ such that every coloured tree \mathfrak{T} satisfying φ has a prefix-closed subset $P \subseteq T$ such that*

$$\mathfrak{T}|_P \models \varphi$$

and every vertex has at most N successors in P .

Proof. Let \mathcal{A} be an automaton equivalent to φ , let $(W_v)_{v \in T}$ be the family from Lemma 4.1, and let $P \subseteq T$ be the least (w.r.t. inclusion) subset of T that contains the root and satisfies $v \in P \Rightarrow W_v \subseteq P$, for all $v \in P$. Then $\mathfrak{T}|_P \models \varphi$ and the branching of P is bounded by N . \square

Theorem 4.3. *Let κ be an infinite cardinal and Σ a unary signature of size at most κ . For every Σ -enriched tree \mathfrak{T} and every cardinal $\kappa \leq \lambda \leq |T|$, there exists a prefix-closed set $P \subseteq T$ of size $|P| = \lambda$ such that $\mathfrak{T}|_P \equiv_{\text{MSO}} \mathfrak{T}$.*

Proof. For every MSO-formula $\varphi \in \text{Th}_{\text{MSO}}(\mathfrak{T})$, we fix an equivalent automaton \mathcal{A}_φ . Let $(W_v^\varphi)_{v \in T}$ be the family from Lemma 4.1 for \mathcal{A}_φ . Fix some subset $P_o \subseteq T$ of size $|P_o| = \lambda$ and let $P \subseteq T$ be the least (w.r.t. inclusion) subset of T that contains P_o and satisfies

$$v \in P \Rightarrow W_v^\varphi \subseteq P, \quad \text{for all } v \in P \text{ and all } \varphi.$$

Then $\mathfrak{T}|_P \models \varphi$, for all formulae φ . Furthermore, the set P has size $|P| = \lambda \cdot \kappa = \lambda$ it is the union of that many finite sets. \square

Corollary 4.4. *Let κ be an infinite cardinal and Σ a unary signature of size at most κ . Every MSO-axiomatisable non-empty class \mathcal{C} of Σ -enriched trees contains a tree of size at most κ .*

Finally, let us also prove corresponding upwards versions of these two theorems.

Theorem 4.5. *Let Σ be a unary signature and $\varphi \in \text{MSO}$ a formula. If there exists an uncountable Σ -enriched tree \mathfrak{T} satisfying φ , then φ is satisfied by Σ -enriched trees of size κ , for every infinite cardinal κ .*

Proof. Suppose that $\mathfrak{T} \models \varphi$ and fix an infinite cardinal κ . If $\kappa \leq |T|$, the claim follows by Theorem 4.3. Hence, suppose that $\kappa > |T|$. Let \mathcal{A} be an automaton equivalent to φ and let ρ be an accepting run of \mathcal{A} on \mathfrak{T} . Since \mathfrak{T} is uncountable, there exists some vertex $v \in T$ with uncountably many successors. Hence, we can find some state q of \mathcal{A} such that uncountably many successors of v are labelled by q . Fix one such successor u of v . Let \mathfrak{S} be the tree obtained from \mathfrak{T} by attaching κ additional copies of the subtree $\mathfrak{T}|_u$ to v and let σ be the run on \mathfrak{S} induced by ρ . Since in both runs the vertex v has infinitely many successors with label q , it follows as in the proof of Lemma 4.1 that

$$\mathfrak{S}uc_*(v; \rho) \equiv_{\text{MSO}} \mathfrak{S}uc_*(v; \sigma).$$

In particular,

$$\mathfrak{S}uc_*(v; \rho) \models \delta(\rho(v)) \quad \text{implies} \quad \mathfrak{S}uc_*(v; \sigma) \models \delta(\sigma(v)).$$

Hence, σ is an accepting run of \mathcal{A} on \mathfrak{S} , which implies that $\mathfrak{S} \models \varphi$. Since $|S| = \lambda$ the claim follows. \square

Theorem 4.6. *Let κ be an infinite cardinal and Σ a unary signature of size at most κ . For every Σ -enriched tree \mathfrak{T} of size $|T| > 2^{2^\kappa}$ and every $\lambda \geq \kappa$, there exists a Σ -enriched tree \mathfrak{S} such that*

$$\mathfrak{S} \equiv_{\text{MSO}} \mathfrak{T} \quad \text{and} \quad |S| = \lambda.$$

Proof. Fix a tree $\mathfrak{T} \in \mathcal{C}$ of size $|T| > 2^{2^\kappa}$ and let Θ be the MSO-theory of \mathfrak{T} . If $\lambda \leq |T|$, we obtain the desired tree \mathfrak{S} by Theorem 4.3. Hence, suppose that $\lambda > |T|$.

For every formula $\varphi \in \Theta$, let \mathcal{A}_φ be an MSO-automaton equivalent to φ and let ρ_φ be an accepting run of \mathcal{A}_φ on \mathfrak{T} . We define a combined labelling ρ of T by setting

$$\rho(v) := (\rho_\varphi(v))_{\varphi \in \Theta}, \quad \text{for } v \in T.$$

Since

$$|\text{rng } \rho| \leq \prod_{\varphi \in \Theta} |Q_\varphi| \leq 2^{|\Theta|} \leq 2^{2^\kappa} < |T|,$$

there is some vertex $v \in T$ with infinitely many successors that are labelled by ρ by the same value $(q_\varphi)_\varphi$. Let u be one such successor and let \mathfrak{S} be the tree obtained from \mathfrak{T} by attaching λ additional copies of the subtree $\mathfrak{T}|_u$ to v .

To prove that $\mathfrak{S} \equiv_{\text{MSO}} \mathfrak{T}$, it is sufficient to show that every automaton \mathcal{A}_φ accepts \mathfrak{S} . Hence, fix $\varphi \in \Theta$ and let σ be the labelling of S induced by ρ_φ . Since the vertex v has infinitely many successors labelled by the state q_φ it follows that

$$\mathfrak{S}uc_*(v; \rho) \equiv_{\text{MSO}} \mathfrak{S}uc_*(v; \sigma).$$

In particular,

$$\mathfrak{S}uc_*(v; \rho) \models \delta(\rho(v)) \quad \text{implies} \quad \mathfrak{S}uc_*(v; \sigma) \models \delta(\sigma(v)).$$

Hence, σ is an accepting run of \mathcal{A}_φ on \mathfrak{S} . \square

Examples. Let us show that these theorems do not hold for structures that are not trees.

(a) We consider 2-sorted structures of the form $\langle A, B, E, \leq \rangle$ with relations $E \subseteq A \times B$ and $\leq \subseteq A \times A$. There exists an MSO-formula φ expressing the following conditions.

- ♦ \leq is a linear order on A of order type ω .
- ♦ For every subset $X \subseteq A$, there exists a unique element $b \in B$ such that

$$\langle a, b \rangle \in E \quad \text{iff} \quad a \in X.$$

Then φ has a unique model where $|A| = \aleph_0$ and $B = \wp(A)$. This model has size 2^{\aleph_0} .

(b) We can construct formulae with larger models by iterating the construction from (a). Consider 2-sorted structures of the form $\langle A, I, E, p, h, \leq \rangle$ with two relations $E \subseteq A \times A$ and $\leq \subseteq I \times I$ and two functions $p : A \rightarrow I$, and $h : I \rightarrow A$. There exists an MSO-formula φ expressing the following conditions.

- \leq is a linear order on I of order type ω .
- $p : A \rightarrow I$ is surjective and $h : I \rightarrow A$ is injective with $\text{rng } h = p^{-1}(o)$.
- $\langle a, b \rangle \in E$ implies $p(b) = p(a) + 1$.
- For every $i \in I$ and every subset $X \subseteq p^{-1}(i)$, there exists a unique element $b \in p^{-1}(i + 1)$ such that

$$\langle a, b \rangle \in E \quad \text{iff} \quad a \in X.$$

Then φ has a unique model where $|p^{-1}(o)| = \aleph_o$ and $p^{-1}(i+1) = \wp(p^{-1}(i))$. This model has size \beth_ω .

Exercise 4.1. Can you use the constructions from the preceding two examples to find formulae with even larger unique models?

Finally, let us note that the theory of all trees is decidable.

Theorem 4.7. *The monadic second-order theory of the class \mathcal{T} of all order-trees is decidable.*

Proof. Let $\varphi \in \text{MSO}$. We have $\varphi \in \text{Th}_{\text{MSO}}(\mathcal{T})$ if, and only if,

$$\langle \kappa^{<\omega}, \leq_{\text{pf}} \rangle \models \forall X [X \text{ is prefix-closed} \rightarrow \varphi^X], \quad \text{for all } \kappa \geq \aleph_o.$$

Hence, it is sufficient to show that the MSO-theory of the class of all trees of the form $\langle \kappa^{<\omega}, \leq_{\text{pf}} \rangle$ is decidable.

We claim that

$$\langle \kappa^{<\omega}, \leq_{\text{pf}} \rangle \equiv_{\text{MSO}} \langle \aleph_o^{<\omega}, \leq_{\text{pf}} \rangle, \quad \text{for } \kappa \geq \aleph_o.$$

Let \mathfrak{C}_κ be the structure of size κ with empty signature. By Proposition I.4.3, it follows that

$$\mathfrak{C}_\kappa \equiv_{\text{MSO}} \mathfrak{C}_{\aleph_o}, \quad \text{for all } \kappa \geq \aleph_o.$$

Hence, the Theorem of Muchnik implies that

$$\mathfrak{C}_\kappa^* \equiv_{\text{MSO}} \mathfrak{C}_{\aleph_o}^*.$$

Since there exists an MSO-interpretation τ with

$$\langle \kappa^{<\omega}, \leq_{\text{pf}} \rangle = \tau(\mathfrak{C}_\kappa^*), \quad \text{for all } \kappa,$$

it follows that

$$\langle \kappa^{<\omega}, \leq_{\text{pf}} \rangle \equiv_{\text{MSO}} \langle \aleph_o^{<\omega}, \leq_{\text{pf}} \rangle,$$

as desired.

We have seen that $\varphi \in \text{Th}_{\text{MSO}}(\mathcal{T})$ if, and only if,

$$\langle \aleph_o^{<\omega}, \leq_{\text{pf}} \rangle \models \forall X (X \text{ is prefix-closed} \rightarrow \varphi^X).$$

Furthermore, by the Theorem of Muchnik, the decidability of $\text{Th}_{\text{MSO}}(\mathfrak{C}_{\aleph_o})$ implies the one of $\text{Th}_{\text{MSO}}(\aleph_o^{<\omega}, \leq_{\text{pf}})$. \square

By contrast, the theory of trees $\kappa^{<\alpha}$ with $\alpha > \omega$ is undecidable.

Theorem 4.8. *Let \mathcal{K} be any class that contains some tree of the form $\langle \kappa^{<\alpha}, \leq_{\text{pf}} \rangle$ where $\kappa > 1$ is a cardinal and $\alpha > \omega$ an ordinal. The monadic second-order theory of \mathcal{K} is undecidable.*

Proof. We will prove in Corollary VI.5.9 below that the monadic theory of the real line $\langle \mathbb{R}, \leq \rangle$ is undecidable. Since $\langle \mathbb{R}, \leq \rangle$ can be MSO-interpreted in $\langle 2^{<\omega+1}, \leq_{\text{pf}}, P_o, P_1 \rangle$ where $P_c := \{w \in 2^{<\omega} \mid w \in c\}$, it follows that the MSO-theory of $\langle 2^{<\omega+1}, \leq_{\text{pf}}, P_o, P_1 \rangle$ is also undecidable. Let $\psi(X, Y_o, Y_1) \in \text{MSO}$ be a formula stating that the substructure induced by the set X when expanded by the unary predicates Y_o and Y_1 is isomorphic to the tree $\langle 2^{<\omega+1}, \leq_{\text{pf}}, P_o, P_1 \rangle$. It follows that, for $\varphi \in \text{MSO}$,

$$\langle 2^{<\omega+1}, \leq_{\text{pf}} \rangle \models \varphi(P_o, P_1)$$

iff $\forall X \forall Y_o \forall Y_1 [\psi(X, Y_o, Y_1) \rightarrow \varphi^X(Y_o, Y_1)] \in \text{Th}_{\text{MSO}}(\mathcal{K})$. \square

5 The Cantor Topology

Sometimes it is advantageous to use topological tools when studying definability issues for trees. To do so, we can equip the set of branches of a tree with a topology.

Definition 5.1. Given a generalised tree \mathfrak{T} , we denote by $[\mathfrak{T}]$ the set of all branches of \mathfrak{T} . The *Cantor topology* on $[\mathfrak{T}]$ is the topology whose open sets are of the form

$$\langle W \rangle := \{ \beta \in [\mathfrak{T}] \mid \beta \cap W \neq \emptyset \}, \quad \text{for } W \subseteq T.$$

In particular, if \mathfrak{T} is a tree of the form D^* , the open sets are of the form WD^ω , for some $W \subseteq D^*$. For singletons $W = \{w\}$, we simply write $\langle w \rangle$.

Before using the topology on the set of branches, let us note that the space $[\mathfrak{T}]$ is well-behaved.

Proposition 5.2. *Let \mathfrak{T} be a generalised tree. The space $[\mathfrak{T}]$ is Hausdorff and zero-dimensional.*

Proof. To see that $[\mathfrak{T}]$ is Hausdorff, fix two different branches β and γ . Then there exist vertices $u \in \beta \setminus \gamma$ and $v \in \gamma \setminus \beta$. It follows that $\langle u \rangle$ and $\langle v \rangle$ are disjoint open sets with $\beta \in \langle u \rangle$ and $\gamma \in \langle v \rangle$.

To see that $[\mathfrak{T}]$ is zero-dimensional, note that its topology has a basis consisting of all sets of the form $\langle v \rangle$, for $v \in T$. We claim that these sets are clopen. Hence, fix $v \in T$ and let W be the set of all vertices of T that are incomparable to v . Then $\langle W \rangle \cap \langle v \rangle = \emptyset$ and $\langle W \rangle \cup \langle v \rangle = T$. Hence, $\langle v \rangle$ is closed. \square

Proposition 5.3. *Let \mathfrak{T} be a tree. The space $[\mathfrak{T}]$ is compact if, and only if, \mathfrak{T} is finitely branching.*

Proof. (\Rightarrow) Suppose that \mathfrak{T} has some vertex v with infinitely many successors, let U be the set of these successors, and let V be the set of all vertices that are incomparable to v . Then we obtain an open cover of \mathfrak{T} that consists of the sets $\langle V \rangle$ and $\langle u \rangle$, for $u \in U$. This cover has no finite subcover.

(\Leftarrow) Suppose that \mathfrak{T} is finitely branching and let $(U_i)_{i \in I}$ be an open cover of $[\mathfrak{T}]$. If there is some $n < \omega$ such that each $w \in T$ of length $|w| = n$ belongs to some $U_{i(w)}$, we obtain a finite subcover

$$(U_{i(w)})_{|w|=n}.$$

Suppose otherwise. Let $F \subseteq T$ be the set of all words $w \in T$ with $w \notin \bigcup_{i \in I} U_i$. By assumption, F contains words of arbitrarily high length. Since F is downwards closed, it follows that $\langle F, \leq \rangle$ is a finitely branching infinite tree. By the Lemma of König, this tree has an infinite branch β . Since $F \subseteq T$, β is also a branch of \mathfrak{T} . Furthermore, by definition of F , we have $\beta \notin \langle U_i \rangle$, for every $i \in I$. A contradiction to the fact that $(U_i)_i$ is a cover. \square

Thin Trees

One way to measure the complexity of a topological space is to distinguish between spaces with ‘many’ points and those with ‘few’ points. We can make this precise by using the following definition.

Definition 5.4. Let \mathfrak{X} be a topological space.

- (a) A point $x \in X$ is *isolated* if the set $\{x\}$ is clopen.
- (b) The *Cantor-Bendixson derivative* $\partial \mathfrak{X}$ of \mathfrak{X} is the subspace of \mathfrak{X} consisting of all non-isolated points.
- (c) The *Cantor-Bendixson rank* $\text{CB}(\mathfrak{X})$ of \mathfrak{X} is the smallest ordinal α such that $\partial^{\alpha+1} \mathfrak{X}$ is empty (where ∂^α denotes the α -th iteration of ∂). If no such ordinal exists, we set $\text{CB}(\mathfrak{X}) = \infty$ and we say that \mathfrak{X} has *no Cantor-Bendixson rank*.
- (c) The *Cantor-Bendixson rank* of a point $x \in X$, is the least ordinal α such that $x \notin \partial^{\alpha+1} \mathfrak{X}$. We denote it by $\text{CB}(x/\mathfrak{X})$.
- (d) A subset $C \subseteq X$ is *perfect* if it is closed and satisfies $\delta C = C$.

Example. Let us consider the following subsets of the real line \mathbb{R} .

- (a) The set $U := \{0\} \cup \{1/n \mid n \in \mathbb{N}, n > 0\}$ (or rather the subspace induced by it) has Cantor-Bendixson rank $\text{CB}(U) = 1$. The points $1/n$ are isolated and have rank 0, while the limit 0 has rank 1.
- (b) No interval $I \subseteq \mathbb{R}$ has a Cantor-Bendixson rank. The same holds for every non-trivial curve in \mathbb{R}^n .

Cantor-Bendixson ranks were originally introduced to prove a version of the Continuum Hypothesis for well-behaved sets in topological spaces. We present the proof for the simplest case, that of closed sets of real numbers.

Theorem 5.5 (Cantor, Bendixson). *Every closed subset of the real line \mathbb{R} can be written as a disjoint union of a countable set and a perfect set.*

Proof. Given a closed set $C \subseteq \mathbb{R}$, let α be the smallest ordinal such that $\partial^{\alpha+1}C = \partial^\alpha C$. Then $P := \partial^\alpha C$ is perfect and it remains to show that $S := C \setminus P$ is countable. Every point $x \in S$ has some rank $\alpha := \text{CB}(x/C) < \infty$. Thus, x is isolated in $\partial^\alpha C$, which means that there is some open neighbourhood I_x of x with $I_x \cap \partial^\alpha C = \{x\}$. As the topology of \mathbb{R} is generated by open intervals with rational end-points, we may assume that $I_x = (a, b)$, for some $a, b \in \mathbb{Q}$. Since there are only countably many such intervals, it is now sufficient to prove that $I_x \neq I_y$ for $x \neq y$.

Hence, consider distinct points $x \neq y$ in P with ranks $\alpha := \text{CB}(x/C)$ and $\beta := \text{CB}(y/C)$. By symmetry, we may assume that $\alpha \leq \beta$. Then $I_x \cap \partial^\alpha = \{x\}$ implies that $y \notin I_x$. Since $y \in I_y$, it follows that $I_x \neq I_y$. \square

Corollary 5.6. *Let $C \subseteq \mathbb{R}$ be closed. Then $|C| \leq \aleph_0$ or $|C| = 2^{\aleph_0}$.*

Proof. By the preceding theorem we can write $C = S + P$, where S is countable and P is perfect. If $P = \emptyset$, we have $|C| = |S| \leq \aleph_0$. Suppose otherwise. We claim that $|P| = 2^{\aleph_0}$. We construct a family $(U_w)_{w \in \{0,1\}^*}$ of non-empty closed intervals $U_w \subseteq P$ such that

- $u \leq_{\text{pf}} v$ implies $U_u \supseteq U_v$,
- $U_u \cap U_v = \emptyset$, for incomparable u, v .
- $P \cap \bigcap_{w <_{\text{pf}} \beta} U_w \neq \emptyset$, for every branch $\beta \in \{0,1\}^\omega$.

Then the function $\eta : [\mathfrak{T}_{\text{bin}}] \rightarrow P$ mapping each branch $\beta \in \{0,1\}^\omega$ to some element in $\bigcap_{w <_{\text{pf}} \beta} U_w$ is injective. In particular,

$$|P| \geq |\text{rng } \eta| = |[\mathfrak{T}_{\text{bin}}]| = 2^{\aleph_0},$$

as desired.

To construct the sets $(U_w)_w$, we will make use of the following observations.

- (I) If P is perfect and non-empty, then $|P| > 1$.
- (II) If U is open and P perfect, then $U \cap P$ is also perfect.

The proofs are straightforward.

(I) If $P = \{x\}$, the point x would be isolated.

(II) For a contradiction, suppose that $U \cap P$ is not perfect. Then it has some isolated point x . This means there is some open set V with $V \cap U \cap P = \{x\}$. Since $V \cap U$ is open, it follows that the point x is isolated in P . A contradiction.

We start the construction of $(U_w)_w$ with the set $U_{\langle \rangle} := \mathbb{R}$. For the inductive step, suppose that we have already defined U_w . By (I) and (II), the set $U_w \cap P$ has at least two elements. Fix different points $x, y \in U_w \cap P$. As \mathbb{R} is Hausdorff, we can choose disjoint open neighbourhoods V and W of x and y . We set $U_{w0} := V$ and $U_{w1} := W$.

It remains to show that $\bigcap_{w <_{\text{pf}} \beta} U_w \neq \emptyset$ for every branch $\beta \in \{0,1\}^\omega$. For a contradiction, suppose that there is some branch β with $P \cap \bigcap_{w <_{\text{pf}} \beta} U_w = \emptyset$. By induction on w , we chose a decreasing sequence $(K_w)_{w <_{\text{pf}} \beta}$ of closed intervals $K_w \subseteq U_w$ with non-empty interior. Since $K_{\langle \rangle}$ is compact and the family $(K_w)_w$ has the finite intersection property, it follows that $\bigcap_{w <_{\text{pf}} \beta} K_w \neq \emptyset$. \square

In the above prove, we have constructed an embedding of the binary tree into the family of open subsets of \mathbb{R} . This idea can be adapted to give a definition of when a tree has ‘few’ branches.

Definition 5.7. Let \mathfrak{S} and \mathfrak{T} be generalised trees.

(a) An *embedding* $\varphi : \mathfrak{S} \rightarrow \mathfrak{T}$ of generalised trees is a map such that

$$\varphi(u \sqcap v) = \varphi(u) \sqcap \varphi(v), \quad \text{for all } u, v \in S.$$

(b) A generalised tree \mathfrak{T} is *thin* if the complete binary tree $\mathfrak{T}_{\text{bin}}$ cannot be embedded into it. \perp

We can characterise thinness in terms of the Cantor-Bendixson rank of $[\mathfrak{T}]$. But it turns out that, for generalised trees, our definition of the rank has to be modified a bit.

Definition 5.8. Let \mathfrak{T} be a generalised tree.

- (a) A branch β of \mathfrak{T} is *isolated* if there exists a vertex $v \in \beta$ such that $\uparrow v \subseteq \beta$.
- (b) The *Lifsches-Shelah derivative* $\partial\mathfrak{T}$ of \mathfrak{T} is the prefix of \mathfrak{T} consisting of all vertices that lie on some non-isolated branch of \mathfrak{T} .
- (c) The *Lifsches-Shelah rank* $\text{LS}(\mathfrak{T})$ of \mathfrak{T} is the smallest ordinal α such that $\partial^{\alpha+1}\mathfrak{T}$ is the empty tree. Similarly, we define the *Lifsches-Shelah rank* $\text{LS}(v, \mathfrak{T})$ of a vertex $v \in T$ as the least ordinal α such that $v \notin \partial^{\alpha+1}\mathfrak{T}$.
- (d) A *subbranch* of \mathfrak{T} is an upwards closed subset of a branch.
- (e) A *skeleton* S of \mathfrak{T} is a partition of T such that every class forms a subbranch of \mathfrak{T} and every branch of \mathfrak{T} intersects only finitely many different classes.

Remark. Note that $\partial[\mathfrak{T}] \subseteq [\partial\mathfrak{T}]$, where the first ∂ denotes the Cantor-Bendixson derivative and the second one the Lifsches-Shelah derivative. This implies that $\text{CB}([\mathfrak{T}]) \leq \text{LS}(\mathfrak{T})$.

Exercise 5.1. Let \mathfrak{T} be a generalised tree and α an ordinal. Prove that

$$\text{LS}(v/\mathfrak{T}) \geq \alpha \quad \text{iff} \quad \begin{array}{l} \text{for every } \beta < \alpha, \text{ there are two incomparable} \\ \text{vertices } u, u' > v \text{ with } \text{LS}(u/\mathfrak{T}) \geq \beta \text{ and} \\ \text{LS}(u'/\mathfrak{T}) \geq \beta. \end{array}$$

Proposition 5.9. Let \mathfrak{T} be a generalised tree. The following statements are equivalent.

- (1) \mathfrak{T} is thin.
- (2) \mathfrak{T} has a skeleton.
- (3) $\text{LS}(\mathfrak{T}) < \infty$

Proof. (2) \Rightarrow (1) For a contradiction, suppose that there exists an embedding $\varphi : \mathfrak{T}_{\text{bin}} \rightarrow \mathfrak{T}$ and that \mathfrak{T} has a skeleton S . We define an infinite branch v_0, v_1, \dots of $\mathfrak{T}_{\text{bin}}$ as follows. We start with the root $v_0 := \langle \rangle$. For the inductive step, suppose that we have already defined v_n and let $\beta \in S$ be the subbranch containing $\varphi(v_n)$. There exists a least one direction $d \in \{0, 1\}$ with $\varphi(\text{suc}_d(v_n)) \notin \beta$. We set $v_{n+1} := \text{suc}_d(v_n)$. Let β be some branch of \mathfrak{T}

containing the image $\varphi(v_0), \varphi(v_1), \dots$ of the sequence constructed in this way. Then β intersects infinitely many different classes from S . A contradiction.

(1) \Rightarrow (3) Suppose that $\text{LS}(\mathfrak{T}) = \infty$ and fix an ordinal α such that $\partial^{\alpha+1}\mathfrak{T} = \partial^\alpha\mathfrak{T}$. By assumption, $\mathfrak{S} := \partial^\alpha\mathfrak{T}$ is not empty. We construct an embedding $\eta : \mathfrak{T}_{\text{bin}} \rightarrow \mathfrak{S} \subseteq \mathfrak{T}$ as follows. For the root, we pick an arbitrary element $\eta(\langle \rangle) \in S$. Inductively, suppose that we have already chosen $\eta(v) \in S$. Since $\partial\mathfrak{S} = \mathfrak{S}$, the subtree $\uparrow\eta(v) \cap S$ contains 2 incomparable vertices u_0 and u_1 . We set $\eta(\text{suc}_d(v)) := u_d$, for $d < 2$.

(3) \Rightarrow (2) We construct a skeleton S of \mathfrak{T} by induction on $\alpha := \text{LS}(\mathfrak{T})$. If $\text{LS}(\mathfrak{T}) = 0$, the tree \mathfrak{T} consists of a single branch β and we can set $S := \{\beta\}$. For the inductive step, suppose that $\alpha > 0$. Note that, for every vertex $v \in T$, the set of vertices $u \geq v$ with $\text{rank } \text{LS}(u/\mathfrak{T}) = \text{LS}(v/\mathfrak{T})$ forms a chain. Let β be some branch containing all vertices v with $\text{LS}(v/\mathfrak{T}) = \alpha$. (There might be none.) We call the maximal \sqcap -closed subsets $C \subseteq T \setminus \beta$ the *components* of $T \setminus \beta$. For every such component C of $T \setminus \beta$, we can find a skeleton S_C by inductive hypothesis. Let S be the set consisting of β and the union of all sets S_C . We claim that S is a skeleton of \mathfrak{T} . Clearly, S is a set of disjoint subbranches with union T . For the second condition, fix a branch γ of T . If $\gamma = \beta$, it intersects with only one class in S . Otherwise, there is some component C such that $\gamma \subseteq \beta \cup C$. Since S_C is a skeleton, γ intersects only finitely many classes in S_C . \square

Example. Note that this theorem does not hold if we replace the Lifsches-Shelah rank by the Cantor-Bendixson rank. Let $\mathfrak{T} = \langle T, \leq_{\text{pf}} \rangle$ be the generalised tree with domain $T := \{0, 1\}^{<\omega+1}$, that is, \mathfrak{T} is obtained from the binary tree $\mathfrak{T}_{\text{bin}}$ by adding a new maximal element to every branch. Then every branch is isolated and $\text{CB}(\mathfrak{T}) = 0$. But there clearly exists an embedding $\mathfrak{T}_{\text{bin}} \rightarrow \mathfrak{T}$.

For finitely branching trees \mathfrak{T} , matters are particularly simple. In particular, in this case we can use the following variant of the Lifsches-Shelah rank which is closer in value to the Cantor-Bendixson rank of $[\mathfrak{T}]$.

Definition 5.10. Let \mathfrak{T} be a finitely branching tree. The *Cantor-Bendixson rank* $\text{CB}(\mathfrak{T})$ of \mathfrak{T} is the least ordinal α such that the tree obtained from \mathfrak{T} by removing all vertices v with $\text{CB}(\mathfrak{T}|_v) < \alpha$ has only finitely many infinite branches. If no such ordinal exists, we set $\text{CB}(\mathfrak{T}) := \infty$. \lrcorner

Proposition 5.11. Let \mathfrak{T} be a finitely branching order-tree. The following statements are equivalent.

- (1) \mathfrak{T} is thin.
- (2) \mathfrak{T} has a skeleton.
- (3) $\text{LS}(\mathfrak{T}) < \omega_1$
- (4) $\text{CB}([\mathfrak{T}]) < \omega_1$
- (5) $\text{CB}(\mathfrak{T}) < \omega_1$
- (6) \mathfrak{T} has only countably many infinite branches.
- (7) \mathfrak{T} has less than 2^{\aleph_0} infinite branches.

Proof. (1) \Leftrightarrow (2) has already been proved in Proposition 5.9.

(6) \Rightarrow (7) is trivial.

(7) \Rightarrow (1) The infinite binary tree $\mathfrak{T}_{\text{bin}}$ has 2^{\aleph_0} infinite branches. If there exists an embedding $\mathfrak{T}_{\text{bin}} \rightarrow \mathfrak{T}$, the tree \mathfrak{T} has at least as many infinite branches as $\mathfrak{T}_{\text{bin}}$.

(1) \Rightarrow (3) If \mathfrak{T} is thin, we have $\alpha := \text{LS}(\mathfrak{T}) < \infty$, by Proposition 5.9. Consequently, we have $\partial^{\beta+1}\mathfrak{T} \setminus \partial^\beta\mathfrak{T} \neq \emptyset$, for all $\beta < \alpha$. This implies that $|T| \geq |\alpha|$. As finitely branching trees have only countably many vertices, it follows that $\alpha < \omega_1$.

(3) \Rightarrow (4) \Rightarrow (5) By definition, we have $\text{CB}(\mathfrak{T}) \leq \text{CB}([\mathfrak{T}]) \leq \text{LS}(\mathfrak{T})$.

(5) \Rightarrow (6) we prove the claim by induction on $\alpha := \text{CB}(\mathfrak{T})$. Let

$$U := \{v \in T \mid \text{CB}(\mathfrak{T}|_v) = \alpha\}.$$

Since T is countable, the complement $T \setminus U$ is a union of countably many trees of Cantor-Bendixson rank less than α . By inductive hypothesis, it therefore follows that $T \setminus U$ contains only countably many branches. Furthermore, by definition, there are only finitely many branches in U . Since

every branch of T is either a branch of U or it contains a suffix in $T \setminus U$, it follows that \mathfrak{T} has only countably many branches as well. \square

Exercise 5.2. Prove that a tree \mathfrak{T} is thin if, and only if, one can assign ordinal numbers to the vertices of \mathfrak{T} in such a way that the successors of a vertex with label α are assigned ordinals $\beta \leq \alpha$ and at most one of them gets the label α itself. \lrcorner

Next, let us take a look at definability questions related to the Cantor-Bendixson rank.

Proposition 5.12. The class of all finitely-branching order-trees that are thin is MSO-definable.

Proof. Note that, for ordinary trees, we can encode a skeleton S as a set of vertices by taking the first vertex of every subbranch in S . Using this encoding, we can express in MSO that the given tree has a skeleton by stating that there exists a set S of vertices such that

- ♦ every connected component of $T \setminus S$ forms a chain and
- ♦ no branch of T contains infinitely many elements from S . \square

Lemma 5.13. For every $n < \omega$, there exists an WMSO-formula φ such that

$$\mathfrak{T} \models \varphi \quad \text{iff} \quad \mathfrak{T} \text{ is a finitely branching tree with } \text{CB}(\mathfrak{T}) \leq n.$$

Proof. Let ϑ be the formula stating that \mathfrak{T} is a finitely branching tree. We can set $\varphi := \vartheta \wedge \psi_n$ where the formula ψ_n states that the set

$$U := \{v \in T \mid \mathfrak{T}|_v \models \psi_{n-1}\}$$

contains only finitely many infinite branches. (For simplicity, we have set $\psi_{-1} := \text{false}$.) To see that the latter can be expressed in WMSO let

$$W := \{v \in U \mid \uparrow v \cap U \text{ is infinite}\}.$$

Then U contains only finitely many branches if, and only if, there exists a finite downwards-closed set $P \subseteq W$ such that $W \setminus P$ is a disjoint union of infinite paths. This can be expressed in WMSO. \square

Let us show that skeletons of thin trees correspond to well-orders and vice versa. Going from a skeleton to a well-order is straightforward.

Proposition 5.14. *There exists an FO-formula $\varphi(x, y; Z)$ such that, for every finitely-branching successor-ordered thin tree \mathfrak{T} with skeleton S , $\varphi(x, y; S)$ defines a well-order on T .*

Proof. Given a skeleton S , let \sqsubseteq_{so} be the ordering obtained from the successor-ordering of T by rearranging the ordering of $\text{Suc}(v)$, for every vertex $v \in T$, such that the successor singled out by an edge in S is the maximal element, while the relative ordering of the other successors is left unchanged. Furthermore, let \sqsubseteq be the lexicographic ordering induced by this new successor-ordering \sqsubseteq_{so} . Note that, given S , we can define \sqsubseteq in FO. Hence, it remains to show that \sqsubseteq is a well-ordering on T .

For a contradiction, suppose that there exists a strictly descending chain $v_o \supset v_1 \supset \dots$. Note that, for every v_i , there are only finitely many $v_j \leq_{pf} v_i$. Replacing the above chain by a subchain, we may therefore assume that $v_i \not\leq_{pf} v_j$, for all i, j . Applying the Lemma of König to the tree with vertices $F := \bigcup_i \downarrow v_i$, we obtain an infinite branch $\beta \subseteq F$. Let w_i be the maximal vertex in β with $w_i \leq_{pf} v_i$. Since every vertex of β is less than some of the v_i , we can construct an infinite infinite subset $I \subseteq \omega$ such that the subsequence $(w_i)_{i \in I}$ is strictly increasing (with respect to \leq_{pf}). Let $g : \omega \rightarrow I$ be an increasing enumeration of I . Then

$$w_{g(i)} = v_{g(i)} \sqcap w_{g(i+1)}, \quad \text{for all } i.$$

Since $v_{g(i+1)} \sqsubset v_{g(i)}$, it follows that the edge going out from $w_{g(i)}$ in direction of $w_{g(i+1)}$ does not belong to S . Consequently, the branch containing the $w_{g(i)}$ contains infinitely many edges that do not belong to S . A contradiction to the definition of a skeleton. \square

Remark. If we impose additional restrictions to the possible orderings on each branch, we can extend this result to generalised trees. We will postpone the proof to Theorem VI.6.3 below. \dashv

The converse is slightly more involved.

Proposition 5.15. *Let $m < \omega$. There exists an MSO-formula $\varphi(x)$ such that, for every finitely-branching successor-ordered thin tree \mathfrak{T} with an MSO_m -definable well-ordering \sqsubseteq , $\varphi(x)$ defines a skeleton of \mathfrak{T} .*

Proof. We claim that the relation

$$x \sqsubseteq_{so} y \quad : \text{iff} \quad x = y \text{ or there is some } v \geq_{pf} y \text{ such that} \\ u \sqsubseteq v \text{ for all } u \geq_{pf} x$$

has the following properties.

- (I) The restriction of \sqsubseteq_{so} to $\text{Suc}(v)$ forms a partial order, for every $v \in T$.
- (II) If u and u' are successors of the same vertex, then $u \sqsubseteq_{so} u'$ or $u' \sqsubseteq_{so} u$.
- (III) \sqsubseteq_{so} induces a successor-ordering on \mathfrak{T} .
- (IV) The set

$$S := \{ \langle u, v \rangle \mid v \text{ is the } \sqsubseteq_{so}\text{-maximal vertex in } \text{Suc}(u) \}$$

forms a skeleton of \mathfrak{T} .

Since S is definable from \sqsubseteq , the proposition then follows.

(I) Reflexivity holds by definition. To prove transitivity, suppose that $x \sqsubseteq_{so} y \sqsubseteq_{so} z$. To show that $x \sqsubseteq_{so} z$, let $w \geq_{pf} z$ be the element such that

$$v \sqsubseteq w, \quad \text{for all } v \geq_{pf} y.$$

Similarly, let $v \geq_{pf} y$ be the element such that

$$u \sqsubseteq v, \quad \text{for all } u \geq_{pf} x.$$

For every $u \geq_{pf} x$, it follows that $u \sqsubseteq v \sqsubseteq w$, as desired.

For anti-symmetry, suppose that $x \sqsubseteq_{so} y$ and $y \sqsubseteq_{so} x$. Let $u \geq_{pf} x$ and $v \geq_{pf} y$ be the elements such that

$$u' \sqsubseteq v, \quad \text{for all } u' \geq_{pf} x, \\ v' \sqsubseteq u, \quad \text{for all } v' \geq_{pf} y.$$

It follows that $u \sqsubseteq v$ and $v \sqsubseteq u$. Hence, $u = v$. Since $u \geq_{\text{pf}} x$, $v \geq_{\text{pf}} y$, and x and y are successors of the same vertex, this implies that $x = y$.

(II) For a contradiction, suppose that $u \not\sqsubseteq_{\text{so}} u'$ and $u' \not\sqsubseteq_{\text{so}} u$. By definition of \sqsubseteq so, it follows that

- ♦ for every $v' \geq_{\text{pf}} u'$ there is some $v \geq_{\text{pf}} u$ with $v \sqsubseteq v'$,
- ♦ for every $v \geq_{\text{pf}} u$ there is some $v' \geq_{\text{pf}} u'$ with $v' \sqsubseteq v$.

Using these two conditions, we can inductively construct two sequences $(v_i)_i$ in $T|_u$ and $(v'_i)_i$ in $T|_{u'}$ such that

$$v_i \sqsubseteq v'_i \quad \text{and} \quad v'_i \sqsubseteq v_{i+1}, \quad \text{for all } i < \omega.$$

Fix $i < j < \omega$, with

$$\mathfrak{T}|_u, v_i \equiv_{\text{MSO}}^m \mathfrak{T}|_u, v_j.$$

Since \sqsubseteq is MSO_m -definable and $v'_i \notin T|_u$, it follows by Proposition 1.7 that

$$v_i \sqsubseteq v'_i \quad \text{implies} \quad v_j \sqsubseteq v'_i.$$

But the latter contradicts the fact that $v'_i \sqsubseteq v_j$.

(III) follows immediately from (I) and (II).

(IV) Clearly, every internal vertex u of \mathfrak{T} has exactly one out-going edge that belongs to S . So, S is a set of subbranches. For a contradiction, suppose that there exists an infinite branch β infinitely many edges of which do not belong to S . Let

$$u_0 \rightarrow v_0, \quad u_1 \rightarrow v_1, \dots$$

be the sequence of these edges and let w_i be the \sqsubseteq_{so} -maximal successor of u_i . By definition of S it follows that $v_i \sqsubseteq_{\text{so}} w_i$. For every i , we can therefore find some vertex $y_i \geq_{\text{pf}} w_i$ such that

$$x \sqsubseteq y_i, \quad \text{for every } x \geq_{\text{pf}} v_i.$$

Hence, $v_i \sqsubseteq_{\text{pf}} u_{i+1} \sqsubseteq_{\text{pf}} w_{i+1} \sqsubseteq_{\text{pf}} y_{i+1}$ implies that $y_{i+1} \sqsubseteq y_i$. Consequently, $y_0 \supset y_1 \supset \dots$ forms an infinite strictly descending sequence. A contradiction to the fact that \sqsubseteq is a well-order. \square

Finally, let us remark that we can define skeletons if the tree has finite Cantor-Bendixson rank.

Lemma 5.16. *For every $n < \omega$, there exists an MSO-formula $\varphi(Z)$ such that, for every finitely branching successor-ordered tree \mathfrak{T} with $\text{CB}(\mathfrak{T}) \leq n$,*

$$\mathfrak{T} \models \varphi(S) \quad \text{iff} \quad S \text{ encodes a skeleton of } \mathfrak{T}.$$

Proof. We define the following relations on T .

$$\begin{aligned} u \rightarrow_o v & : \text{iff} \quad v \text{ is a successor of } u \text{ with } \text{CB}(\mathfrak{T}|_u) = \text{CB}(\mathfrak{T}|_v), \\ u \rightarrow v & : \text{iff} \quad v \text{ is the } \leq_{\text{so}}\text{-minimal successor of } u \text{ with } u \rightarrow_o v. \end{aligned}$$

According to Lemma 5.13, these relations are definable. Hence, we can take for φ the formula stating that S consists of the connected components of \rightarrow . \square

Corollary 5.17. *For every finitely branching successor-ordered tree \mathfrak{T} with rank $\text{CB}(\mathfrak{T}) < \omega$, there exists an MSO-formula $\varphi(x, y)$ (without parameters) that defines a well-order on T .*

Borel Sets

Besides the Cantor-Bendixson rank, there is a second way of measuring the complexity of a subspace: we can look at how it is composed from simpler sets. The topological simplest sets are the open and the closed ones. As operation for composition we take countable unions. This leads to the following definition.

Definition 5.18. Let \mathfrak{X} be a topological space.

(a) A set $U \subseteq X$ is *Borel* if it belongs to the closure of the class of open sets under complement and countable unions.

(b) A function $f : \mathfrak{X} \rightarrow \mathfrak{Y}$ between topological spaces is *Borel* if

$$f^{-1}[U] \text{ is Borel, for every Borel set } U \subseteq X.$$

(c) Similarly, a relation $R \subseteq X^n$ is *Borel* if it is Borel as a subset of \mathfrak{X}^n . \square

Remark. (a) We can associate an ordinal rank to each Borel set U by counting how many unions we need to express it. This induces a hierarchy on the Borel sets of length ω_1 . We omit the definition, since we will not use this additional information.

(b) We have already used Borel sets in Section IV.3 where we showed that games with a Borel accepting condition are determined. \square

The following properties of Borel sets follow immediately from the definition.

Lemma 5.19. *Let \mathfrak{X} and \mathfrak{Y} be topological spaces.*

- (a) *The preimage of a Borel set under a continuous function $f : \mathfrak{X} \rightarrow \mathfrak{Y}$ is also Borel.*
- (b) *The class of Borel relations on \mathfrak{X} is closed under countable unions, countable intersections, complement, and finite direct products.*
- (c) *If $R \subseteq X \times Y$ is Borel, so is every fibre*

$$R_c := \{x \in X \mid \langle x, c \rangle \in R\}, \quad \text{for all } c \in Y.$$

Proof. (a) Let $f : \mathfrak{X} \rightarrow \mathfrak{Y}$ be continuous and $U \subseteq Y$ Borel. We prove the claim by induction on the number of operations needed to produce U from open sets. By assumption, the preimage of every open set is open. For the inductive step, note that

$$\begin{aligned} f^{-1}[Y^n \setminus U] &= X^n \setminus f^{-1}[U], \\ f^{-1}[\bigcup_{n < \omega} U_n] &= \bigcup_{n < \omega} f^{-1}[U_n]. \end{aligned}$$

(b) is immediate.

(c) Note that $R_c = f^{-1}[R]$, where $f : X \rightarrow X \times Y$ is the function mapping x to $\langle x, c \rangle$. By (b), it is therefore sufficient to prove that f is continuous. Every open set in $\mathfrak{X} \times \mathfrak{Y}$ is of the form $U \times V$, where $U \subseteq X$ and $V \subseteq Y$ are open. Furthermore, we have

$$f^{-1}[U \times V] = \begin{cases} U & \text{if } c \in V, \\ \emptyset & \text{if } c \notin V, \end{cases}$$

both of which are open. \square

We need some more substantial results whose proofs are too involved to be reproduced here. The first two can be found as, respectively, Theorem 14.12 and Theorem 13.6 in [115], while a proof of the third one is given in Example 1.6 (2) of [102].

Lemma 5.20. *A function $f : \mathfrak{X} \rightarrow \mathfrak{Y}$ is Borel if, and only if, its graph is Borel in $\mathfrak{X} \times \mathfrak{Y}$.*

Theorem 5.21. *Let \mathfrak{T} be a finitely-branching tree and $B \subseteq [\mathfrak{T}]$ Borel. Then B is uncountable if, and only if, there exists a subset $S \subseteq T$ such that $\langle S, \leq_{\text{pf}} \rangle$ is isomorphic to the infinite binary tree and every branch of $\langle S, \leq_{\text{pf}} \rangle$ belongs to B .*

Theorem 5.22. *There exists no Borel function $f : D^\omega \rightarrow D^\omega$ such that*

$$f(\alpha) = f(\beta) \quad \text{iff} \quad \alpha \approx_* \beta,$$

where

$$\alpha \approx_* \beta \quad \text{iff} \quad \text{there are only finitely many positions } x \text{ with } \alpha(x) \neq \beta(x).$$

The connection between logic and topology is given by the following two results. The first one concerns languages of ω -words, the second one branches in a tree.

Lemma 5.23. *Every regular language $L \subseteq \Sigma^\omega$ is Borel.*

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ be a deterministic parity automaton recognising L . Fix a number $m < \omega$ such that all priorities used by \mathcal{A} are smaller than $2m$. For $k < 2m$, let W_n^k be the set of all words $w \in \Sigma^*$ such that the (unique) run of \mathcal{A} on w contains exactly n states with priority k . Setting

$$U^k := \bigcap_{n < \omega} W_n^k \Sigma^\omega,$$

we then have

$$L = \bigcup_{k < m} \left[U^{2k} \setminus \bigcup_{i < 2k} U^i \right].$$

As each set $W_n^k \Sigma^\omega$ is open, it follows that L is Borel. \square

Proposition 5.24. *Let \mathfrak{T} be a finitely-branching tree, $\psi(X, \bar{Z}) \in \text{MSO}$, \bar{P} monadic parameters, and*

$$B := \{ \beta \mid \beta \text{ infinite branch with } \mathfrak{T} \models \psi(\beta, \bar{P}) \}.$$

- (a) B is Borel.
- (b) The following statements are equivalent.
 - (1) B is countable.
 - (2) $|B| < 2^{\aleph_0}$.
 - (3) There is no subset $S \subseteq T$ such that $\langle S, \leq_{\text{pf}} \rangle$ is isomorphic to the infinite binary tree and every branch of $\langle S, \leq_{\text{pf}} \rangle$ belongs to B .

Proof. (a) Given a branch β and unary predicates \bar{P} , we can write

$$\langle \mathfrak{T}, \beta, \bar{P} \rangle \cong \sigma \left[\sum_{v \in \langle \beta, \leq_{\text{pf}} \rangle} \mathfrak{S}_v \right]$$

as a generalised sum followed by an interpretation σ , where \mathfrak{S}_v is the substructure of $\langle \mathfrak{T}, \beta, \bar{P} \rangle$ with universe

$$S_v := \uparrow v \setminus \bigcup_{u \in \beta \setminus \downarrow v} \uparrow u.$$

Therefore, Theorem I.4.24 yields formulae $\vartheta, \chi_0, \dots, \chi_{n-1}$ such that

$$\mathfrak{T} \models \psi(\beta, \bar{P}) \quad \text{iff} \quad \langle \beta, \leq_{\text{pf}} \rangle \models \vartheta(\llbracket \chi_0(\bar{P}) \rrbracket, \dots, \llbracket \chi_{n-1}(\bar{P}) \rrbracket),$$

for all β and \bar{P} . By Lemma 5.23, the language K of ω -words defined by ϑ is Borel. Let f be the function mapping a branch β to the tuple

$$f(\beta) := (\llbracket \chi_i(\bar{P}) \rrbracket)_{i < n}.$$

This function is continuous since every prefix ρ of $f(\beta)$ is completely determined by the prefix $v \in \beta$ of the same length $|\rho| = |\nu|$. Consequently, $B = f^{-1}[K]$ is also Borel.

(b) (1) \Rightarrow (2) is trivial.

(2) \Rightarrow (3) If there exists a set S as in (3), the tree $\langle S, \leq_{\text{pf}} \rangle$ has 2^{\aleph_0} infinite branches. As each of them belongs to B , it follows that $|B| = 2^{\aleph_0}$.

(3) \Rightarrow (1) By (a), B is Borel. If B is uncountable, it therefore follows by Theorem 5.21 that we can find a set S as in (3). \square

Remark. Note that there are regular languages of infinite trees that are *not* Borel. \lrcorner

6 Counting Quantifiers

Let us take a look at the extension of MSO by various counting quantifiers. We have already defined CMSO in Chapter I. There is a second, more expressive version of this logic where we count not elements but sets. We will show below that in certain cases such quantifiers can be eliminated.

Definition 6.1. C_2MSO , monadic-second order logic with *second-order counting*, is the extension of MSO by *counting quantifiers* of the form

$$\exists^\kappa Y \varphi(\bar{X}, Y) \quad \text{‘There exists at least } \kappa \text{ sets satisfying } \varphi\text{.’}$$

$$\exists^{k,m} Y \varphi(\bar{X}, Y) \quad \text{‘The number of sets } Y \text{ satisfying } \varphi \text{ is finite and congruent } k \text{ modulo } m\text{.’}$$

If we want to explicitly specify which quantifiers are allowed, we write $\text{MSO}[\exists^{\aleph_0}]$, $\text{MSO}[\exists^{k,m}]$, $\text{MSO}[\exists^{\aleph_0}, \exists^{2^{\aleph_0}}]$, etc. \lrcorner

Remark. The expressive power of C_2MSO does not increase if we add cardinality quantifiers over *tuples* of sets since

$$(\exists^\kappa X \bar{Y}) \varphi \equiv \exists^\kappa X [\exists \bar{Y} \varphi] \vee \exists^\kappa \bar{Y} [\exists X \varphi], \quad \text{for infinite cardinals } \kappa. \lrcorner$$

We start with a few simple cases where (some of the) counting quantifiers can be eliminated. The first one is well-orders.

Lemma 6.2. *Over the class of all well-orders, CMSO is equivalent to MSO.*

Proof. Over a well-order, we can express the predicate $|X| \geq \aleph_0$ by the formula

$$\forall z \exists x [z < x \wedge Xx] \vee \exists y (\forall z < y) \exists x [z < x < y \wedge Xx],$$

where the first part handles the case where X is unbounded and the second one the case where X has a least upper bound y .

A predicate of the form $|X| \equiv k \pmod{m}$ holds if, and only if, $|X| < \aleph_0$ and there are sets Z_0, \dots, Z_{m-1} satisfying the following conditions.

- ♦ $\min X \in Z_0$ and $\max X \in Z_k$.
- ♦ Every element belongs to at most one set Z_i .
- ♦ If $x < y$ are elements of X such that no element between x and y belongs to X , then $x \in Z_i \Leftrightarrow y \in Z_{i+1}$ (indices modulo m).

By the first part of the proof, every of these conditions can be expressed in MSO. \square

Let us also note that the second-order version of the quantifier \exists^{\aleph_0} can always be replaced by the first-order one.

Proposition 6.3. $\text{MSO}[\exists^{\aleph_0}]$ is equivalent to $\text{MSO}[\text{inf}]$.

Proof. Clearly, we can translate every $\text{MSO}[\text{inf}]$ -formula into $\text{MSO}[\exists^{\aleph_0}]$. Conversely, note that, if there are only finitely many sets X satisfying a given formula φ , we can find a finite set Z such that any two sets X and X' satisfying φ differ in some element of Z . Hence,

$$\begin{aligned} \exists^{\aleph_0} X \varphi(X) &\equiv \\ &\neg \exists Z [|Z| < \aleph_0 \wedge (\forall X. \varphi(X)) (\forall X'. \varphi(X')) \\ &\quad [X \neq X' \rightarrow X \cap Z \neq X' \cap Z]]. \end{aligned} \quad \square$$

Finally, we turn to finitely-branching trees. Our goal is to establish the following two results.

Theorem 6.4.

- (a) *Over finitely-branching trees, $\text{MSO}[\exists^{\aleph_0}, \exists^{2^{\aleph_0}}]$ is equivalent to MSO.*
- (b) *Over finitely-branching trees, C_2MSO is equivalent to CMSO.*
- (c) *Over finitely-branching successor-ordered trees, C_2MSO is equivalent to MSO.*

Theorem 6.5. *Over the class of finitely-branching trees, we have*

$$\exists^{\aleph_1} Y \varphi(\bar{X}, Y) \equiv \exists^{2^{\aleph_0}} Y \varphi(\bar{X}, Y), \quad \text{for all } \text{C}_2\text{MSO-formulae } \varphi.$$

Proof of Theorems 6.4 and 6.5

For the proof, we need to be able to evaluate products in finite semigroups. The following lemma explains how this can be done.

Lemma 6.6. *Let \mathfrak{S} be a finite semigroup.*

- (a) *For every $c \in S$, there exists an MSO-formula φ_c such that, for every sequence $a_0, \dots, a_{n-1} \in S$,*

$$\langle [n], \leq, \bar{P} \rangle \models \varphi_c \quad \text{iff} \quad a_0 \cdots a_{n-1} = c,$$

where $\bar{P} = (P_b)_{b \in S}$ is the family with $P_b := \{ i < n \mid a_i = b \}$, for $b \in S$.

- (b) *If \mathfrak{S} is commutative, there exists CMSO-formulae φ_c such that, for every sequence $a_0, \dots, a_{n-1} \in S$,*

$$\langle [n], \bar{P} \rangle \models \varphi_c \quad \text{iff} \quad a_0 \cdots a_{n-1} = c.$$

Proof. (a) We guess a labelling of the positions with elements from \mathfrak{S} such that position i is labelled by the value of the product $a_0 \cdots a_i$. Then the label of the last position is the product of the whole sequence. Thus, φ_c states that there are sets Z_b , $b \in S$, satisfying the following conditions.

- ♦ The first position x satisfies $x \in Z_b \Leftrightarrow x \in P_b$.
- ♦ The last position belongs to Z_c .
- ♦ Every position belongs to exactly one set Z_b .

- ♦ If y is the successor of x then $x \in Z_a \wedge y \in P_b \Rightarrow y \in Z_{ab}$.

Clearly, all of this can be expressed in MSO.

(b) Since \mathfrak{S} is commutative, the value of a product $a_0 \cdot \dots \cdot a_{n-1}$ only depends on the cardinalities $|P_b|$, for $b \in S$. Furthermore, as \mathfrak{S} is finite, there exists a number $m < \omega$ such that $a^m = a^\pi$, for every $a \in S$. Then

$$a^{im+k} = a^{m+k}, \quad \text{for } i > 0,$$

implies that the value of the product only depends on the numbers

$$\min\{|P_b|, m\} \quad \text{and} \quad |P_b| \bmod m, \quad \text{for } b \in S.$$

These can be computed using modulo predicates. \square

We start with the simple parts of Theorem 6.4.

Proposition 6.7.

- (a) Over finitely-branching trees, $\text{MSO}[\exists^{\aleph_0}]$ is equivalent to MSO.
- (b) Over finitely-branching trees, $\text{MSO}[\exists^{\aleph_0}, \exists^{k,m}]$ is equivalent to CMSO.
- (c) Over finitely-branching successor-ordered trees, $\text{MSO}[\exists^{\aleph_0}, \exists^{k,m}]$ is equivalent to MSO.

Proof. (a) Proposition 6.3 provides a translation of $\text{MSO}[\exists^{\aleph_0}]$ to $\text{MSO}[\text{inf}]$. Hence, it remains to show how to eliminate the predicate inf . Over finitely-branching trees, we can use König's Lemma to do so. Since a set X is infinite if, and only if, its prefix closure contains an infinite path, we can use the formula

$$\text{inf}(X) := \exists Z [Z \neq \emptyset \wedge (\forall z \in Z)(\exists x \in X)(\exists y \in Z)[z < x \wedge z < y]].$$

(b) We can eliminate the counting quantifier \exists^{\aleph_0} as in (a). For a modulo-counting quantifier $\exists^{k,m} Y \varphi(\bar{X}, Y)$ we proceed as follows. Given values for \bar{X} , we pick a finite prefix P such that every set Y satisfying φ is determined

by its intersection with P . (As explained in the proof of Proposition 6.3, if there is no such prefix, then there are infinitely many sets Y satisfying φ .) Suppose that $\varphi \in \text{CMSO}_r$ and let Θ be the set of all CMSO_r -theories. For each vertex $v \in P$ and every $\tau \in \Theta$, let U_v^τ be the set of all $Y_0 \subseteq P \cap \uparrow v$ such that there exists some $Y \subseteq \uparrow v$ with $Y_0 = P \cap Y$ and $\text{Th}_{\text{CMSO}}^r(\mathfrak{T}|_v, Y) = \tau$. For each v and τ , we guess the number

$$g_v(\tau) := |U_v^\tau| \bmod m$$

and we label v with the function g_v . By choice of the set P , the formula $\exists^{k,m} Y \varphi(\bar{X}, Y)$ is equivalent to the statement that

$$\sum_{\substack{\tau \in \Theta \\ \varphi \in \tau}} |U_v^\tau| \equiv k \pmod{m},$$

which we can verify using the labelling $g_{(\cdot)}$.

To conclude the proof it is therefore sufficient to show that we can express the correctness of this guessed labelling in CMSO. We can do so going through the tree bottom-up. For $v \notin P$, we have $g_v(\tau) = 0$, for all τ . Hence it suffices to consider vertices $v \in P$. Let u_0, \dots, u_{n-1} be the successors of v . Note that we can express the subtree $\mathfrak{T}|_v$ as a generalised sum of the singleton $\{v\}$ and the subtrees $\mathfrak{T}|_{u_0}, \dots, \mathfrak{T}|_{u_{n-1}}$. Furthermore, the set Θ of theories forms a semigroup when equipped with the disjoint union operation \oplus . We can form a new semigroup with universe m^Θ and multiplication

$$(g \cdot h)(\tau) := \sum \{g(\rho) \cdot h(\sigma) \mid \rho \oplus \sigma = \tau\} \bmod m.$$

Since \oplus is commutative, so is this new semigroup. Consequently, we can use Lemma 6.6 (b) to construct CMSO-formulae ψ_g , for $g \in m^\Theta$, such that

$$([n], \bar{N}) \models \psi_g$$

if, and only if, the number of sets $Y \subseteq \uparrow u_0 \cup \dots \cup \uparrow u_{n-1}$ with

$$\text{Th}_{\text{CMSO}}^r(\mathfrak{T}|_{u_0}, Y \oplus \dots \oplus \mathfrak{T}|_{u_{n-1}}, Y) = \tau$$

is congruent $g(\tau)$ modulo m . Since the CMSO_r -theory of $\mathfrak{T}|_v, Y$ can be computed from the theories of $\{v\}, Y$ and $\mathfrak{T}|_{u_0}, Y \oplus \dots \oplus \mathfrak{T}|_{u_{n-1}}, Y$, we can use these formulae ψ_g to construct a formula $\vartheta_g(v)$ that uses the labelling of the vertices u_0, \dots, u_{n-1} and that satisfies

$$\begin{aligned} \mathfrak{T} \models \vartheta_g(v) \quad \text{iff} \quad & \text{the number of sets } Y \subseteq \uparrow v \text{ with} \\ & \text{Th}_{\text{CMSO}}^r(\mathfrak{T}|_v, Y) = \tau \\ & \text{is congruent } g(\tau) \text{ modulo } m. \end{aligned}$$

This formula can check that the labelling of v is correct.

(c) Note that we only use modulo predicates for sets of successors of a given vertex. If the tree is successor-ordered, the successors of every vertex v are well-ordered. Hence, we can use Lemma 6.2 to eliminate all modulo predicates. \square

It remains to consider quantifiers $\exists^* X \varphi(X, \tilde{Y})$ for uncountable cardinals κ . Without loss of generality, we may assume that we have already eliminated all cardinality quantifiers in φ , that is, $\varphi \in \text{CMSO}_m$ for some m . In the following, to simplify notation, we will fix values \tilde{Q} for the free variables \tilde{Y} and include them in the tree $\mathfrak{T} = \langle T, \leq_{\text{pf}}, \tilde{Q} \rangle$. That way, we can assume that $\varphi(X)$ has a single free variable X . Thus, below trees will always implicitly be assumed to have additional unary predicates \tilde{Q} . We will make use of the following composition operation for such trees.

Definition 6.8. Let \mathfrak{S} and \mathfrak{T}_s , for $s \in S$, be order trees. The *ordered tree sum* of the family $(\mathfrak{T}_s)_s$ is the tree

$$\sum_{s \in \mathfrak{S}} \mathfrak{T}_s$$

whose universe is the disjoint union $\sum_s T_s$ and whose order is given by

$$\begin{aligned} \langle s, u \rangle \leq \langle t, v \rangle \quad : \text{iff} \quad & s = t \text{ and } u \leq v, \text{ or} \\ & s \leq t \text{ and } u \text{ is the root of } \mathfrak{T}_s. \end{aligned}$$

Since an ordered tree sum can be written as a generalised sum followed by a quantifier-free interpretation, we obtain the following composition theorem.

Proposition 6.9. For each CMSO_m -formula $\varphi(\tilde{X})$, we can construct CMSO -formulae $\psi(\tilde{Z})$ and $\chi_0(\tilde{X}), \dots, \chi_{n-1}(\tilde{X})$ such that each χ_i has quantifier-rank at most m and

$$\sum_{s \in \mathfrak{S}} \mathfrak{T}_s \models \varphi(\tilde{P}) \quad \text{iff} \quad \mathfrak{S} \models \psi(\llbracket \chi_0(\tilde{P}) \rrbracket, \dots, \llbracket \chi_{n-1}(\tilde{P}) \rrbracket),$$

for all trees \mathfrak{S} and \mathfrak{T}_s , $s \in S$, where

$$\llbracket \chi(\tilde{P}) \rrbracket := \{ s \in S \mid \mathfrak{T}_s \models \chi(\tilde{P}|_{T_s}) \}.$$

To apply this result, we need a bit of notation and terminology for parts of a tree.

Definition 6.10. Let \mathfrak{T} be a finitely-branching tree.

(a) A *factor* of \mathfrak{T} is a connected subset $D \subseteq T$, i.e., a subset such that, for all $u, v \in D$, the path (ignoring edge directions) between u and v also belongs to D .

(b) For a vertex $v \in T$ and a set $D \subseteq T$, we set

$$T[v, D] := \{ u \in T \mid v \leq_{\text{pf}} u \text{ and } w \not\leq_{\text{pf}} u \text{ for all } w \in D \text{ with } v < w \}.$$

So $T[v, D]$ consists of the subtree attached at v where we have removed every subtree attached to some $w \in D$ with $v < w$. We denote the substructures of \mathfrak{T} induced by these sets by $\mathfrak{T}[v, D]$. \lrcorner

Remark. Given a prefix-closed subset $B \subseteq T$, we can express \mathfrak{T} as the ordered tree sum

$$\mathfrak{T} \cong \sum_{v \in B} \mathfrak{T}[v, B].$$

The following special version of Proposition 6.9 will be used below.

Corollary 6.11. *Given an CMSO_m-formula $\varphi(X)$, we can construct CMSO-formulae $\psi(\vec{Z})$ and $\chi_o(X), \dots, \chi_{n-1}(X)$ such that each χ_i has quantifier-rank at most m and*

$$\mathfrak{T} \models \varphi(P) \quad \text{iff} \quad \langle B, \leq_{\text{pf}} \rangle \models \psi(\llbracket \chi_o(P) \rrbracket, \dots, \llbracket \chi_{n-1}(P) \rrbracket),$$

for all trees \mathfrak{T} , sets $P \subseteq T$, and prefix-closed $B \subseteq T$, where

$$\llbracket \chi(P) \rrbracket := \{ v \in B \mid \mathfrak{T}[v, B] \models \chi(P \cap T[v, B]) \}.$$

Let us introduce the central notion the proof of Theorem 6.4 below is based on.

Definition 6.12. Let $\mathfrak{T} = \langle T, \leq_{\text{pf}}, \vec{Q} \rangle$ be a finitely-branching tree, $\varphi(X)$ an CMSO-formula of quantifier-rank m , and P a set satisfying $\varphi(X)$.

(a) A factor D is a *P-choice factor* if $P \cap D$ is not determined by its CMSO_m-theory, i.e., if there exists a set $P' \subseteq D$ with $P' \neq P \cap D$ such that

$$\mathfrak{T}|_D, P' \equiv_{\text{CMSO}}^m \mathfrak{T}|_D, P|_D.$$

(b) If $D = T|_v$, we call v a *P-choice vertex*. The set of all *P-choice* vertices is denoted by $\text{Ch}(P)$.

(c) An infinite branch $\beta \subseteq \text{Ch}(P)$ is called a *P-choice branch*. We denote by $\text{CB}(P)$ the set of all such branches. J

Remark. (a) The above definition of a *P-choice factor* depends on the formula φ . We have omitted it from the notation to keep the terminology light.

(b) Note that $\text{Ch}(P)$ is prefix closed. J

First, we check that these notions are definable.

Lemma 6.13. (a) *There exists an CMSO-formula $\psi(P, D)$ expressing that D is a P-choice factor.*

(b) *There exists an CMSO-formula $\psi(P, v)$ stating that $v \in \text{Ch}(P)$.*

(c) *There exists an CMSO-formula $\psi(P, D)$ stating that $B \in \text{CB}(P)$.*

Proof. (a) Let

$$\vartheta(X) := \forall Z \bigwedge_{\psi} [\psi(Z) \leftrightarrow \psi(X)] \rightarrow Z = X,$$

where the conjunction ranges over all formulae ψ of quantifier-rank m . Then we can set

$$\psi(P, D) := 'D \text{ is a factor}' \wedge \neg \vartheta^{(D)}(P \cap D),$$

where $\vartheta^{(D)}$ denotes the relativisation of ϑ to the set D .

(b) follows immediately by (a).

(c) follows by (b). □

In particular, we obtain the following consequence of Corollary 6.11.

Corollary 6.14. *There exists CMSO-formulae $\psi, \chi_o, \dots, \chi_{n-1}$ such that, for every tree \mathfrak{T} and every prefix-closed set $B \subseteq T$,*

$$P \in \varphi^{\mathfrak{T}} \quad \text{and} \quad B \subseteq \text{Ch}(P)$$

if, and only if,

$$\langle B, \leq_{\text{pf}} \rangle \models \psi(\llbracket \chi_o(P) \rrbracket, \dots, \llbracket \chi_{n-1}(P) \rrbracket).$$

Proof. By Lemma 6.13 (c), there exists an CMSO-formula stating that $P \in \varphi^{\mathfrak{T}}$ and $\beta \in \text{CB}(P)$. Hence, the claim follows by Corollary 6.11. □

Working towards a proof of Theorems 6.4 and 6.5, we derive a sequence of combinatorial conditions that are equivalent to the existence of uncountably many sets satisfying the a given formula φ .

Lemma 6.15. *Let $\varphi(X)$ be a formula and $P \in \varphi^{\mathfrak{T}}$.*

(a) *If $\text{Ch}(P)$ contains an infinite antichain, then $|\varphi^{\mathfrak{T}}| = 2^{\aleph_0}$.*

(b) *$\text{CB}(P) = \emptyset$ implies $|\varphi^{\mathfrak{T}}| < \aleph_0$.*

Proof. (a) Fix an infinite antichain $C \subseteq \text{Ch}(P)$. By the Lemma of König, the set $\downarrow C$ contains an infinite branch β . It is sufficient to prove that $\beta \in \text{CB}(P')$, for 2^{\aleph_0} different sets $P' \in \varphi^{\mathfrak{T}}$. Let

$$S := \{ v \in T \mid \text{no } u \leq_{\text{pf}} v \text{ belongs to } C \}.$$

Then we can decompose \mathfrak{T} as

$$\mathfrak{T} = \sum_{v \in S} \{v\} + \sum_{v \in C} \mathfrak{T}|_v,$$

(where \sum denotes the ordered tree sum from Definition 6.8 and the above expression is considered as a single sum, not two) where $\{v\}$ denotes the substructure of \mathfrak{T} consisting of the single vertex v . As every $v \in V$ is a P -choice vertex, there exist sets $P'_v \subseteq T|_v$, for $v \in V$, such that

$$P'_v \neq P \cap T|_v \quad \text{and} \quad \mathfrak{T}|_v, P'_v \equiv_{\text{MSO}}^m \mathfrak{T}|_v, P \cap T|_v.$$

For every function $\sigma : C \rightarrow [2]$, define $P_\sigma \subseteq T$ by

$$P_\sigma \cap S = P \cap S \quad \text{and} \quad P_\sigma \cap T|_v := \begin{cases} P \cap T|_v & \text{if } \sigma = 0, \\ P'_v & \text{if } \sigma = 1. \end{cases}$$

By Proposition 6.9 it follows that

$$\mathfrak{T}, P_\sigma \equiv_{\text{MSO}}^m \mathfrak{T}, P, \quad \text{for all } \sigma.$$

In particular, each set P_σ satisfies $\varphi(X)$. Furthermore, $\beta \subseteq C \subseteq \text{Ch}(P_\sigma)$.

(b) If $\text{Ch}(P)$ does not contain an infinite branch, it follows by König's Lemma that $\text{Ch}(P)$ is finite. Furthermore, the set P is completely determined by (i) the subset $P \cap \text{Ch}(P)$ and (ii) the MSO_m -theories of all subtrees attached at some successor of a vertex in $\text{Ch}(P)$. Since there are only finitely many such subtrees and only finitely many choices for their MSO_m -theory, it follows that there are only finitely many sets P . \square

Below we want to prove that there is no difference between the cardinals \aleph_0 and 2^{\aleph_0} . This is the reason we allow both choices of κ in the next characterisation.

Proposition 6.16. *Let \mathfrak{T} be a finitely-branching tree, $\varphi(X)$ an CMSO_m -formula, and let κ be either \aleph_1 or 2^{\aleph_0} . Then*

$$|\varphi^{\mathfrak{T}}| \geq \kappa$$

if, and only if, at least one of the following conditions is satisfied.

- (I) $\text{Ch}(P)$ contains an infinite antichain, for some $P \in \varphi^{\mathfrak{T}}$.
- (II) The set $\bigcup \{ \text{CB}(P) \mid P \in \varphi^{\mathfrak{T}} \}$ is uncountable.
- (III) There exists an branch β with $\beta \in \text{CB}(P)$, for at least κ many $P \in \varphi^{\mathfrak{T}}$.

Proof. (\Leftarrow) (I) implies $|\varphi^{\mathfrak{T}}| = 2^{\aleph_0} \geq \kappa$ by Lemma 6.15 (a), while (III) trivially implies that $|\varphi^{\mathfrak{T}}| \geq \kappa$.

It therefore remains to consider the case where (II) holds, but (I) and (III) do not. Then every infinite branch β belongs to less than κ sets $\text{CB}(P)$. If $|\varphi^{\mathfrak{T}}| < \kappa$, it would therefore follow that the size of $B := \bigcup \{ \text{CB}(P) \mid P \in \varphi^{\mathfrak{T}} \}$ would be less than κ . By Proposition 5.24 it then follows that B is countable. A contradiction.

(\Rightarrow) Suppose that $|\varphi^{\mathfrak{T}}| \geq \kappa$ and set $B := \bigcup \{ \text{CB}(P) \mid P \in \varphi^{\mathfrak{T}} \}$. If B contains uncountably many infinite branches, then (II) holds and we are done. Hence, suppose otherwise. By Lemma 6.15 (b), we have $\text{CB}(P) \neq \emptyset$, for all $P \in \varphi^{\mathfrak{T}}$. Consequently,

$$\kappa \leq |\varphi^{\mathfrak{T}}| \leq \sum_{P \in \varphi^{\mathfrak{T}}} |\text{CB}(P)| = \sum_{\beta \in B} |\{ P \in \varphi^{\mathfrak{T}} \mid \beta \in \text{CB}(P) \}|.$$

If there is some infinite branch β with $\beta \in \text{CB}(P)$, for at least κ many $P \in \varphi^{\mathfrak{T}}$, then (III) holds and we are done. Otherwise, the above bound on κ is given by a countable sum of cardinals less than κ . Since the cofinalities of \aleph_1 and 2^{\aleph_0} are both uncountable, this sum is therefore also less than κ . A contradiction. \square

It is not clear how to express condition (III) above in monadic second-order logic. The following proposition provides an alternative formulation that is easier to express in logic.

Definition 6.17. Let \mathfrak{T} be a finitely-branching tree, $B \subseteq T$ prefix-closed, and let $\psi, \chi_0, \dots, \chi_{n-1}$ be the MSO-formulae from the preceding corollary. By adding more formulae to the list $\chi_0, \dots, \chi_{n-1}$, if necessary, we may assume without loss of generality that, for every CMSO_m-theory θ there is some $i < n$ such that

$$\chi_i \equiv \bigwedge \theta.$$

(a) We set

$$\mathcal{P}_B := \{ P \in \varphi^{\mathfrak{T}} \mid B \subseteq \text{Ch}(P) \}$$

and we define a function $\tau : \mathcal{P}_B \rightarrow \wp(B)^n$ mapping each set $P \in \mathcal{P}_B$ to the tuple \bar{U} with

$$U_i := \llbracket \chi_i(P) \rrbracket, \quad \text{for } i < n.$$

(b) A family \bar{U} is an *admissible labelling* of B if $\bar{U} = \tau(P)$, for some $P \in \mathcal{P}_B$. In this case we say that \bar{U} is *associated* with P . $\quad \lrcorner$

Proposition 6.18. Let \mathfrak{T} be a finitely-branching tree, $\varphi(X)$ an CMSO_m-formula, and β an infinite branch such that \mathcal{P}_β is uncountable. At least one of the following conditions is satisfied.

- (I) There exists some $P \in \varphi^{\mathfrak{T}}$ and infinitely many $v < \beta$ such that the set $T[v, \beta)$ is a P -choice factor.
- (II) There are uncountably many admissible labellings of β .
- (III) There exists some $P \in \varphi^{\mathfrak{T}}$ and some $v < \beta$ such that there are uncountably many sets $P' \subseteq T[v, \beta)$ with

$$\mathfrak{T}[v, \beta) \models \chi_i(P') \iff \mathfrak{T}[v, \beta) \models \chi_i(P \cap T[v, \beta)),$$

for all $i < n$.

Proof. We distinguish two cases. First, suppose that $\tau^{-1}(\bar{U})$ is countable, for all admissible labellings \bar{U} of β . Since $\text{dom } \tau = \mathcal{P}_\beta$ is uncountable while $\tau^{-1}(\bar{U})$ is countable, it then follows that $\text{rng } \tau$ is countable. This proves (II).

Hence, it remains to consider the case where there exists a labelling \bar{U} such that $\tau^{-1}(\bar{U})$ is uncountable. Suppose that (I) does not hold and fix a set $P \in \tau^{-1}(\bar{U})$. We will prove (III). By assumption, there are only finitely many $v < \beta$ such that $T[v, \beta)$ is a P -choice factor. Let $C \subseteq \beta$ be the finite set of these vertices. Recall that we chose the formulae χ_i such that, for every CMSO_m-theory θ , there is some index i with $\chi_i \equiv \bigwedge \theta$. For $P' \in \tau^{-1}(\bar{U})$, it therefore follows that

$$\mathfrak{T}[v, \beta), P' \cap T[v, \beta) \equiv_{\text{CMSO}}^m \mathfrak{T}[v, \beta), P \cap T[v, \beta),$$

for all $v < \beta$. Consequently,

$$P' \cap T[v, \beta) = P \cap T[v, \beta), \quad \text{for all } v \in \beta \setminus C.$$

Since $\tau^{-1}(\bar{U})$ is uncountable, it follows that there is some vertex $v \in C$ such that

$$P' \cap T[v, \beta) \neq P \cap T[v, \beta), \quad \text{for uncountably many } P' \in \tau^{-1}(\bar{U}).$$

This proves (III). $\quad \square$

Lemma 6.19. There exists an CMSO-formula $\psi(\beta, \bar{U})$ stating that β is an infinite branch and \bar{U} an admissible labelling of β .

Proof. Given β , the formula ψ guesses some set P and checks that

- ♦ P satisfies φ ,
- ♦ $\beta \in \text{CB}(P)$,
- ♦ $U_i = \llbracket \chi_i(P) \rrbracket$.

Note that the second statement can be expressed in MSO by Lemma 6.13. $\quad \square$

The following is a special case of Theorem XII.6.7.

Definition 6.20. Let β be a branch of a tree \mathfrak{T} . For two labellings \bar{U} and \bar{V} of β , we write

$$\bar{U} \approx_* \bar{V} \quad : \text{iff} \quad \bar{U} \text{ and } \bar{V} \text{ differ in only finitely many positions.} \quad \lrcorner$$

Lemma 6.21. *There exists a constant k such that the following statements are equivalent for every finitely-branching tree \mathfrak{T} and every infinite branch β .*

- (1) *There are only countably many admissible labellings of β .*
- (2) *There are less than 2^{\aleph_0} admissible labellings of β .*
- (3) *There exist families $\bar{W}^0, \dots, \bar{W}^{k-1}$ of subsets of β such that every admissible labelling \bar{U} of β satisfies*

$$\bar{U} \approx_* \bar{W}^i, \quad \text{for some } i.$$

Proof. Let $\psi(\beta, \bar{U})$ be the formula from Lemma 6.19 defining the admissible labellings of β . Let m be its quantifier-rank, let d be the number of CMSO_m -theories, and set $k := d^2$.

- (1) \Rightarrow (2) is trivial.
- (3) \Rightarrow (1) For every tuple \bar{W} there are only countably many labellings \bar{U} with $\bar{U} \approx_* \bar{W}$.
- (2) \Rightarrow (3) Suppose that there are $k+1$ pairwise non- \approx_* -equivalent admissible labellings $\bar{W}^0, \dots, \bar{W}^k$ of β . By the Theorem of Ramsey, there exists an infinite subset $H \subseteq \beta$ such that

$$\begin{aligned} \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[\langle \rangle, u)} &\equiv_{\text{CMSO}_m} \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[\langle \rangle, u')}, \\ \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[u, v)} &\equiv_{\text{CMSO}_m} \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[u', v')}, \end{aligned}$$

for all $u < v$ and $u' < v'$ in H and all $i \leq k$, where

$$[u, v) := \{ w \in \beta \mid u \leq_{\text{pf}} w < v \}.$$

As there are only d CMSO_m -theories, we can find $i < j$ such that

$$\begin{aligned} \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[\langle \rangle, u)} &\equiv_{\text{CMSO}_m} \langle \beta, \leq_{\text{pf}}, \bar{W}^j \rangle|_{[\langle \rangle, u)}, \\ \langle \beta, \leq_{\text{pf}}, \bar{W}^i \rangle|_{[u, v)} &\equiv_{\text{CMSO}_m} \langle \beta, \leq_{\text{pf}}, \bar{W}^j \rangle|_{[u, v)}, \end{aligned}$$

for all $u < v$ in H . W.l.o.g. we may assume that $i = 0$ and $j = 1$. Since $\bar{W}^0 \not\approx_* \bar{W}^1$, there exists an infinite subset $I \subseteq H$ such that $\bar{W}^0|_{[u, v)} \neq \bar{W}^1|_{[u, v)}$,

for all $u < v$ in I . Let $u_0 < u_1 < \dots$ be an enumeration of I . For every $\sigma : \omega \rightarrow [2]$, we define a labelling \bar{U}_σ of β by

$$\begin{aligned} \bar{U}_\sigma|_{[\langle \rangle, u_0)} &:= \bar{W}^0|_{[\langle \rangle, u_0)} \\ \text{and } \bar{U}_\sigma|_{[u_i, u_{i+1})} &:= \bar{W}^{\sigma(i)}|_{[u_i, u_{i+1})}, \quad \text{for } i < \omega. \end{aligned}$$

Then $\bar{U}_\sigma \neq \bar{U}_\tau$, for $\sigma \neq \tau$, and

$$\langle \beta, \leq_{\text{pf}}, \bar{U}_\sigma \rangle \equiv_{\text{CMSO}_m} \langle \beta, \leq_{\text{pf}}, \bar{W}^0 \rangle, \quad \text{for all } \sigma.$$

In particular, \bar{U}_σ is an admissible labelling of β . Consequently, there are 2^{\aleph_0} such labellings. \square

Theorem 6.22. *Let \mathfrak{T} be a finitely-branching tree and $\varphi(X)$ an CMSO_m -formula. The following statements are equivalent.*

- (I) $|\varphi^\mathfrak{T}| \geq \aleph_1$
- (2) $|\varphi^\mathfrak{T}| \geq 2^{\aleph_0}$
- (3) *At least one of the following conditions is satisfied.*
 - (I) $\text{Ch}(P)$ contains an infinite antichain, for some $P \in \varphi^\mathfrak{T}$.
 - (II) The set $\bigcup \{ \text{CB}(P) \mid P \in \varphi^\mathfrak{T} \}$ contains uncountably many infinite branches.
 - (III) There exist a branch β and a set $P \in \varphi^\mathfrak{T}$ such that $T[v, \beta)$ is a choice factor, for infinitely many $v < \beta$.
 - (IV) There exists an infinite branch β with uncountably many admissible labellings.

Proof. (2) \Rightarrow (1) is trivial.

(3) \Rightarrow (2) If (I) or (II) hold, the claim follows by Proposition 6.16. If (III) holds, there are 2^{\aleph_0} different sets $P \in \varphi^\mathfrak{T}$ with $\beta \in \text{CB}(P)$. Finally, suppose that (IV) holds. Then Lemma 6.21 implies that β has 2^{\aleph_0} admissible labellings. Since each labelling corresponds to a different set $P \in \varphi^\mathfrak{T}$, it follows that $\varphi^\mathfrak{T}$ has size 2^{\aleph_0} .

(1) \Rightarrow (3) For a contradiction, suppose that $\varphi^\mathfrak{T}$ is uncountable but (I)–(IV) do not hold. Since (II) fails, the set $B := \bigcup \{ \text{CB}(P) \mid P \in \varphi^\mathfrak{T} \}$ contains only

countably many infinite paths. Furthermore, by the failure of (i), every set $\text{CB}(P)$ is finite. Since every countable set has only countably many finite subsets, it follows that there are only countably many sets of the form $\text{CB}(P)$, for $P \in \varphi^\omega$. Since there are uncountably many $P \in \varphi^\omega$, it follows that we can find a set S such that $S = \text{CB}(P)$, for uncountably many sets $P \in \varphi^\omega$. Let \mathcal{P} be the set of all $P \in \varphi^\omega$ with $\text{CB}(P) = S$. Let D be the set of vertices lying on some branch in S , let $D_1 \subseteq D$ be the set of vertices $v \in D$ with at most one successor in D and let $D_2 \subseteq D$ be the set of remaining vertices. Since S is finite, so is D_2 . Furthermore, for every set $P \in \varphi^\omega$ and every branch $\beta \subseteq D$, the induced labellings of D and of β coincide on all vertices in $\beta \cap D_1$. Since D has only finitely many infinite branches, since D_2 is finite, and since, by the failure of (iv), every branch has only countably many admissible labellings, it follows that there are only countably many admissible labellings of D . Hence, there exists some labelling \tilde{U} of D that is associated with uncountably many sets $P \in \varphi^\omega$. As D is countable, we can find some vertex $v \in D$ such that the set

$$H := \{ P \cap T[v, D) \mid P \in \mathcal{P}, \tilde{U} \text{ associated with } P \}$$

is uncountable. Note that $T[v, D)$ can be decomposed into the vertex v and a finite number of subtrees $T[u]$, for successors u of v with $u \notin S$. For $P \in \mathcal{P}$ and $u \notin D$, the fact that $\text{CB}(P) = S$ implies that $P \cap T[u]$ is uniquely determined by $\text{Th}_{\text{CMSO}}^m(\mathfrak{A}[u], P)$. Since there are only finitely many CMSO_m -theories, it follows that H is finite. A contradiction. \square

This theorem immediately implies Theorem 6.5. To prove Theorem 6.4, it remains to show that the above conditions can be expressed in CMSO.

Proof of Theorem 6.4. By Proposition 6.7 and Theorem 6.5, it is sufficient to eliminate the quantifier \exists^{\aleph_1} . We have shown in Lemma 6.13, that the set $\text{Ch}(P)$ is MSO-definable, and in Lemma 6.19 that admissibility is definable. We therefore obtain the following formulae expressing conditions (i)–(iv)

from Theorem 6.22.

$$\begin{aligned} \psi_I &:= \exists P[\varphi(P) \wedge \exists Z[‘Z \text{ antichain}’ \wedge |Z| \geq \aleph_0 \wedge Z \subseteq \text{Ch}(P)]], \\ \psi_{II} &:= \exists Z[‘\langle Z, \leq_{\text{pf}} \rangle \cong \langle 2^*, \leq_{\text{pf}} \rangle’ \\ &\quad \wedge \forall \beta[\beta \subseteq Z \wedge \beta \text{ branch}’ \rightarrow \exists P[\varphi(P) \wedge \beta \subseteq \text{Ch}(P)]]], \\ \psi_{III} &:= \exists \beta \exists P(\forall u < \beta)(\exists v < \beta)[u \leq_{\text{pf}} v \wedge v \in \text{Ch}(P)], \\ \psi_{IV} &:= \exists \beta \exists \tilde{U}^0 \dots \exists \tilde{U}^k[‘\beta \text{ branch}’ \wedge \bigwedge_{i \neq j} \tilde{U}^i \neq_* \tilde{U}^j \\ &\quad \wedge \bigwedge_{i \leq k} ‘\tilde{U}^i \text{ admissible labelling of } \beta’], \end{aligned}$$

where the number k in ψ_{IV} is the constant from Lemma 6.21 and correctness of ψ_{II} follows by Proposition 5.24. Consequently, we have

$$\exists^{\aleph_1} Y \varphi(\tilde{X}, Y) \equiv \psi_I \vee \psi_{II} \vee \psi_{III} \vee \psi_{IV}.$$

\square

Formulae Without Free Variables

For formulae without free variables, we can improve the above results. In this case, the problem is equivalent to checking whether a given regular language of trees is uncountable. We will derive a simple criterion for this to be the case. In the following, we will work with Σ -labelled finitely-branching trees, for some finite alphabet Σ . For various constructions of trees below, it will be convenient to consider not only trees but also *contexts*.

Definition 6.23. (a) We consider Σ -labelled trees as functions $t : \text{dom}(t) \rightarrow \Sigma$. For $v \in \text{dom}(t)$, we denote by $t|_v$ the subtree attached at v .

(b) A *context* is a Σ -labelled tree p where some of the leaves are labelled by special *variables* x_0, \dots, x_{n-1} (which are assumed not to belong to Σ). We say that p is n -ary if it contains at most the variables x_0, \dots, x_{n-1} . (Some of the variables may be missing.)

(c) For an n -ary context p and trees s_0, \dots, s_{n-1} , we write $p(s_0, \dots, s_{n-1})$ for the tree obtained from p by replacing every leaf labelled by a variable x_i by a copy of s_i .

For unary contexts, we omit the parentheses and simply write ps_o . Furthermore, we write p^ω for the tree obtained as the limit $p, pp, ppp, pppp, \dots$

(d) Let p and q be unary contexts. A vertex w *distinguishes* p and q if either

- $w \in \text{dom}(p)$, $w \in \text{dom}(q)$, and $p(w) \neq q(w)$, or
- $w \in \text{dom}(p)$, $w \notin \text{dom}(q)$, and there is no leaf v of q with $v < w$ and $q(v) = x_o$, or
- $w \in \text{dom}(q)$, $w \notin \text{dom}(p)$, and there is no leaf v of p with $v < w$ and $p(v) = x_o$.

If such a vertex exists, we call p and q *distinguishable*.

(e) We call a tree t *expandable* if it is isomorphic to a proper subtree of itself, that is, if $t = pt$, for some non-trivial context p . \downarrow

We start with a few lemmas.

Lemma 6.24. *Every non-regular tree t has an infinite branch β such that*

$$t|_u \not\equiv t|_v, \quad \text{for all } u < v < \beta.$$

Proof. Consider the set

$$P := \{ v \in T \mid t|_v \not\equiv t|_u \text{ for all } u <_{\text{lex}} v \}.$$

By the Lemma of Kőnig, the set $\Downarrow P$ contains an infinite branch β . We claim that β has the desired properties. Hence, fix two vertices $u < v < \beta$. Let w be the \leq_{lex} -least vertex in P with $v \leq_{\text{pf}} w$. Suppose that $w = vz$. Then $|uz| < |vz|$ implies that $uz <_{\text{lex}} vz = w \in P$. By definition of P , we therefore have $t|_{uz} \not\equiv t|_{vz}$. Hence, $t|_u \not\equiv t|_v$. \square

Lemma 6.25. *Let p, a_i, b_i , for $i < \omega$, be unary contexts that each have at least one occurrence of the variable x_o . Suppose that*

- *for every $i < \omega$, a_i and b_i are either distinguishable or equal,*
- *there is at least one index i with $a_i \neq b_i$.*

Then $pa_o a_1 a_2 \dots \neq pb_o b_1 b_2 \dots$

Proof. Fix the minimal index i with $a_i \neq b_i$. Let w be the vertex distinguishing a_i and b_i and let u be some leaf of $pa_o \dots a_{i-1} = pb_o \dots b_{i-1}$ labelled by x_o . Then uw distinguishes $pa_o a_1 a_2 \dots$ and $pb_o b_1 b_2 \dots$. \square

In the case of ω -words a regular language $K \subseteq \Sigma^\omega$ is countable if, and only if, it is a finite union of languages of the form Uv^ω with $U \subseteq \Sigma^*$ regular and $v \in \Sigma^+$. We next theorem contains a similar characterisation for languages of trees.

Theorem 6.26. *Let K be a regular language of finitely-branching Σ -labelled trees. The following statements are equivalent.*

- (1) *K is countable.*
- (2) *$|K| < 2^{\aleph_o}$*
- (3) *Every tree in K is regular.*
- (4) *There are only finitely many expandable trees t that are subtrees of some tree in K .*
- (5) *$K = \{ p(s_o, \dots, s_{n-1}) \mid p \in P \}$, for some set P of finite contexts and some finite tuple s_o, \dots, s_{n-1} of expandable trees.*

Proof. (5) \Rightarrow (1) As P is a set of finite contexts, it is countable. Hence, there are only countably many trees of the form $p(\bar{s})$ with $p \in P$.

(1) \Rightarrow (3) There are only countably many regular trees.

(3) \Rightarrow (2) Let φ be an MSO-formula defining K and let m be its quantifier-rank. Fix a non-regular tree $t \in K$ and let β be the branch from Lemma 6.24. By the Theorem of Ramsey, there exists an infinite subset $H \subseteq \beta$ such that

$$t[u, v] \equiv_{\text{MSO}}^m t[u', v'] \quad \text{for all } u < v \text{ and } u' < v' \text{ in } H.$$

Since H is infinite, we can find vertices $u < v$ and $u' < v'$ in H such that the intervals $t[u, v]$ and $t[u', v']$ are distinguishable. Suppose that $v = uz$ and $v' = u'z'$. For $\sigma \in [2]^\omega$ and $i < \omega$, let c_i^σ be the context with

$$c_i^\sigma := \begin{cases} t[u, v] & \text{if } \sigma(i) = 0, \\ t[u', v'] & \text{if } \sigma(i) = 1, \end{cases}$$

and let s_σ be the tree obtained by concatenating

$$t[o, u), c_o^\sigma, c_1^\sigma, c_2^\sigma, \dots$$

For $\sigma \neq \tau$, it follows by Lemma 6.25 that $s_\sigma \not\equiv s_\tau$. Furthermore, $c_i^\sigma \equiv_{\text{MSO}}^m c_i^\tau$ implies that $s_\sigma \equiv_{\text{MSO}}^m s_\tau$. Consequently, we have found 2^{\aleph_0} different trees $s_\sigma \in K$.

(2) \Rightarrow (4) Fix an MSO-formula φ defining K and let m be its quantifier rank. For a contradiction, suppose that there are infinitely many expandable trees as in (4). Then we can find two trees $s, t \in K$ of the form $s = pu$ and $t = qv$ where $u \neq v$, the subtrees $u = au$ and $v = bv$ are expandable, and $a \equiv_{\text{MSO}}^m b$. In particular, $s = pa^\omega$ and $a^\omega \neq b^\omega$. Consequently, there exists some vertex w distinguishing a^ω and b^ω . For a sufficiently large number n , the same vertex distinguishes a^n and b^n . For $\sigma \in [2]^\omega$, set

$$s_\sigma := pc_o c_1 c_2 \dots \quad \text{where} \quad c_i := \begin{cases} a^n & \text{if } \sigma(i) = o, \\ b^n & \text{if } \sigma(i) = 1. \end{cases}$$

Then $s_\sigma \equiv_{\text{MSO}}^m s$ implies $s_\sigma \in K$. Furthermore, it follows by Lemma 6.25 that $s_\sigma \neq s_\tau$ for all $\sigma \neq \tau$. Consequently, K is uncountable. A contradiction.

(4) \Rightarrow (5) Let S be the set of all expandable trees which are a subtree of some tree in K . By assumption, S is finite. Suppose that $S = \{s_o, \dots, s_{n-1}\}$ and let P be the set of all trees obtained from some tree in K by deleting every subtree isomorphic to some s_i . Then

$$K = \{p(s_o, \dots, s_{n-1}) \mid p \in P\}$$

and it remains to prove that every tree $p \in P$ is finite. For a contradiction, suppose that there is some infinite $p \in P$. Then p has some infinite branch β . Since p is regular, it has only finitely many different subtrees. Hence, there are two vertices $u < v$ on the branch β with $p|_u \cong p|_v$. Consequently, the subtree $p|_u$ is expandable. But no tree in P can have an expandable subtree. A contradiction. \square

As an application, we show how we can encode the trees in a countable language by finite ones.

Theorem 6.27. *Let K be a regular language of finitely-branching Σ -labelled trees. If K is countable, there exists an MSO-definable injective function σ mapping each tree in K to some finite tree.*

Proof. By Theorem 6.26 (5), we can write

$$K = \{p(s_o, \dots, s_{n-1}) \mid p \in P, s_o, \dots, s_{n-1} \in S\}.$$

Let σ be the function that maps each tree $t \in K$ to the prefix obtained from t by

- deleting ever subtree in S and
- replacing it with some leave that is labelled by the isomorphism type of the deleted subtree.

Since S is a finite set of regular trees, it is MSO-definable. Hence, so is σ . Furthermore, σ is clearly injective. \square

Open Questions.

- (a) Does Theorem 6.27 hold for formulae with parameters, i.e., for sets of the form $K = \varphi(\tilde{X}; \tilde{P})^\otimes$ with parameters $P_i \subseteq T$?
- (b) Is there an analogue of Theorem 6.22 for sets of the form φ^\otimes/E , where E is an MSO-definable equivalence relation (similar to the results of Section XII.6)?

Notes

Good introductions to the theory of automata on ω -words and infinite trees are [202, 204]. The latter also contains a proof that parity games are positionally determined. Corollary 3.9 is from Rabin [156].

The Theorem of Muchnik was announced in [192], but a proof was never published. The proof above was provided by Walukiewicz [208]. This paper is also the source of the automata constructions presented in Section 2. Corollary 3.11 was an earlier result by [64]. Theorem 4.7 can be found in [205, 193].

There is one general decidability result strictly subsuming the Theorem of Muchnik: a certain way of recursively decomposing trees and computing their theories inductively introduced by Puppis [155].

Theorem 1.13 is due to Colcombet [48]. The downward Löwenheim-Skolem theorems (Theorems 4.2 and 4.3) are taken from [25].

A standard reference for Borel complexity and Cantor-Bendixson ranks is [115]. Most of our results on thin trees are extracted from [133]. For further logical properties of thin trees we refer to [28, 196]. The material on counting quantifiers is based on [150, 10].

Part B

Structure Theory

VI Linear Orders

IN THIS SECOND PART OF THE BOOK we study ways to decompose structures, consider how these decompositions influence their logical properties, and show how to utilise them for composition arguments. We start with the simplest case, that of linear orders.

1 Dense and Scattered Orders

We can decompose linear orders using sums and products. Ordered sums were already defined in Section I.4. We repeat the definition here for convenience.

Definition I.I. (a) Let $(\mathcal{A}_i)_{i \in I}$ be a family of linear orders indexed by a linearly ordered set I . The *ordered sum* $\sum_{i \in I} \mathcal{A}_i$ is the linear order with universe

$$\sum_{i \in I} A_i$$

and ordering

$$\langle i, a \rangle \leq \langle j, b \rangle \quad : \text{iff} \quad i < j \text{ or } (i = j \text{ and } a \leq b).$$

(b) The *ordered product* $\mathcal{A} \times \mathcal{B}$ of two linear orders \mathcal{A} and \mathcal{B} is the linear order with universe

$$A \times B$$

and ordering

$$\langle a, b \rangle \leq \langle a', b' \rangle \quad : \text{iff} \quad b < b' \text{ or } (b = b' \text{ and } a \leq a').$$

(c) An *embedding* of linear orders is a function $h : \mathfrak{A} \rightarrow \mathfrak{B}$ such that

$$a \leq b \quad \text{iff} \quad h(a) \leq h(b), \quad \text{for all } a, b \in A.$$

(d) For $\mathfrak{A} = \langle A, \leq \rangle$, we denote by $\mathfrak{A}^{\text{op}} := \langle A, \geq \rangle$ the linear order with the opposite ordering.

(e) An *interval* is a set of the form

$$[a, b] := \{ c \in A \mid a \leq c \leq b \text{ or } b \leq c \leq a \},$$

$$[a, b) := \{ c \in A \mid a \leq c < b \text{ or } b \leq c < a \}.$$

Example. We have $\omega \times 2 \cong \omega + \omega$ but $2 \times \omega \cong \omega$.

Exercise 1.1. Let \mathfrak{A} , \mathfrak{B} , and \mathfrak{C} be linear orders. Prove the following equalities.

(a) $\mathfrak{A} \times \mathfrak{B} \cong \sum_{b \in B} \mathfrak{A}$

(b) $\mathfrak{A} \times (\mathfrak{B} + \mathfrak{C}) \cong (\mathfrak{A} \times \mathfrak{B}) + (\mathfrak{A} \times \mathfrak{C})$

Exercise 1.2. Find two countable ordinals α and β with $\alpha \equiv_{\text{MSO}} \beta$.

Exercise 1.3. Let Σ be a finite alphabet. An Σ -labelled countable linear order w of order type \mathbb{Z} is called *recurrent* if every finite word $u \in \Sigma^+$ occurs in every prefix and every suffix of w . Prove that $w \equiv_{\text{MSO}} w'$, for all recurrent orders w and w' .

Dense Orders

It is not always possible to decompose a linear order into parts that are ‘simpler’ than the given one. One such case is when the order is dense.

Definition 1.2. Let \mathfrak{A} be a linear order.

(a) Given two sets $P, Q \subseteq A$, we say that P is *dense in* Q if, for all $a < a'$ in Q , there is some $b \in P$ with $a < b < a'$. We call P *densely ordered* if it has at least 2 elements and it is dense in itself.

(b) \mathfrak{A} is *scattered* if no infinite subset $P \subseteq A$ is densely ordered.

(c) A family $(P_i)_{i \in I}$ of subsets $P_i \subseteq A$ is *mutually dense* if each P_i is infinite and, for all $i, j \in I$ and all $a < a'$ in P_i , there is some $b \in P_j$ with $a < b < a'$, that is, if every P_j is dense in $P_i \cup P_j$.

Similarly, we call a colouring $\chi : A \rightarrow C$ *dense* if the sets $(\chi^{-1}(c))_{c \in C}$ are mutually dense.

Let us start by classifying all countable dense linear orders.

Lemma 1.3. Every countable linear order can be embedded into the rationals.

Proof. Let $\mathfrak{A} = \langle A, \leq \rangle$ be countable and fix an enumeration $(a_n)_{n < \omega}$ of A . We construct an increasing sequence $h_0 \subseteq h_1 \subseteq \dots$ of partial embeddings $h_n : A \rightarrow \mathbb{Q}$ with

$$\text{dom } h_n = \{a_0, \dots, a_n\}.$$

The limit $h := \bigcup_n h_n$ is then the desired embedding $\mathfrak{A} \rightarrow \mathbb{Q}$.

We start with the function $h_0 : \{a_0\} \rightarrow \mathbb{Q}$ mapping a_0 to 0. For the inductive step, suppose that we have already defined h_n . Set

$$b := \min \{ a \in \text{dom } h_n \mid a < a_n \},$$

$$c := \max \{ a \in \text{dom } h_n \mid a > a_n \}$$

(if the minimum and the maximum exist). We consider three cases. Note that at least one of b and c exists. If both do, we set

$$h_{n+1}(a_n) := \frac{1}{2} [h_n(b) + h_n(c)].$$

If only b exists, we set

$$h_{n+1}(a_n) := h_n(b) + 1,$$

and if only c exists, we set

$$h_{n+1}(a_n) := h_n(c) - 1.$$

Clearly, the resulting function h_{n+1} is injective and monotone. \square

Proposition 1.4. Let C be a finite set of colours and $\mathfrak{A} = \langle A, \leq, \alpha \rangle$ and $\mathfrak{B} = \langle B, \leq, \beta \rangle$ two countable dense linear orders without end-points with dense colourings $\alpha : A \rightarrow C$ and $\beta : B \rightarrow C$. Then $\mathfrak{A} \cong \mathfrak{B}$.

Proof. We repeat the construction in the proof of the preceding lemma while making sure the resulting map is surjective and that it preserves the colours. Hence, fix enumerations $(a_n)_{n<\omega}$ and $(b_n)_{n<\omega}$ of A and B such that $\alpha(a_0) = \beta(b_0)$. We construct an increasing sequence $h_0 \subseteq h_1 \subseteq \dots$ of partial embeddings $h_n : \mathfrak{A} \rightarrow \mathfrak{B}$ such that $\text{dom } h_n$ is finite and

$$\{a_0, \dots, a_n\} \subseteq \text{dom } h_n \quad \text{and} \quad \{b_0, \dots, b_n\} \subseteq \text{rng } h_n.$$

We start with the function h_0 mapping a_0 to b_0 . For the inductive step, suppose that we have already defined h_n . Since β is dense and \mathfrak{B} has no end-points, there exists some element $b' \in B$ such that, $\beta(b') = \alpha(a_{n+1})$ and, for all $a \in \text{dom } h_n$,

$$\begin{aligned} h_n(a) \leq b' & \quad \text{iff} \quad a \leq a_{n+1}, \\ b' \leq h_n(a) & \quad \text{iff} \quad a_{n+1} \leq a. \end{aligned}$$

Let h' be the extension of h_n mapping a_{n+1} to b' . Since α is dense and \mathfrak{A} has no end-points, we can similarly find some element $a' \in A$ such that, $\alpha(a') = \beta(b_{n+1})$ and, for all $a \in \text{dom } h'$,

$$\begin{aligned} a \leq a' & \quad \text{iff} \quad h'(a) \leq b_{n+1}, \\ a' \leq a & \quad \text{iff} \quad b_{n+1} \leq h'(a). \end{aligned}$$

For h_{n+1} we take the extension of h' mapping a' to b_{n+1} . \square

Corollary 1.5 (Cantor). *Every countable dense linear order without end-points is isomorphic to the rationals.*

Corollary 1.6. *Up to isomorphism, there exist exactly 4 countable dense linear orders.*

It remains to prove the existence of dense colourings.

Lemma 1.7. *For every countable set C , there exists a dense colouring $\chi : \mathbb{Q} \rightarrow C$.*

Proof. Set $T := \{w1 \mid w \in \{0, 1\}^*\}$. Then the order $\langle T, \leq_{\text{lex}} \rangle$ is countable, dense, and it has no end-points. Consequently, it is isomorphic to $\langle \mathbb{Q}, \leq \rangle$

and it is sufficient to construct a dense colouring $\chi : T \rightarrow C$. Let $(c_n)_{n<\omega}$ be an enumeration of C (possibly with repetitions). We set

$$\chi(1^n) := c_n \quad \text{and} \quad \chi(w01^n) := c_n, \quad \text{for } n < \omega.$$

For every two words $u <_{\text{lex}} v$, there is some w with $u <_{\text{lex}} w0 <_{\text{lex}} v$. Consequently, $x := w01^n$ is an element of T with $u <_{\text{lex}} x <_{\text{lex}} v$ and $\chi(x) = c_n$. \square

We can generalise this result to uncountable orderings as follows.

Lemma 1.8. *Let \mathfrak{A} be a dense linear ordering and C a countable set. There exists a dense colouring $\chi : A \rightarrow C$.*

Proof. Let \mathcal{F} be the set of all pairs $\langle D, \lambda \rangle$ such that $D \subseteq A$ is densely ordered without end-points and $\lambda : D \rightarrow C$ is a dense colouring of D . We order \mathcal{F} by

$$\langle D, \lambda \rangle \leq \langle D', \lambda' \rangle \quad \text{iff} \quad D \subseteq D' \text{ and } \lambda = \lambda' \upharpoonright D.$$

We will prove the following claims.

- (I) \mathcal{F} is inductively ordered.
- (II) If $\langle D, \lambda \rangle \in \mathcal{F}$ is maximal, then the suborder $A \setminus D$ is scattered.
- (III) If $\langle D, \lambda \rangle \in \mathcal{F}$ is maximal and $\chi : A \rightarrow C$ is some function with $\chi \upharpoonright D = \lambda$, then χ is dense.

Then it follows by (I) and Zorn's Lemma, that \mathcal{F} has a maximal element $\langle D, \lambda \rangle$. Hence, we can use (III) to obtain the desired dense colouring of A .

- (I) Let $\langle D_i, \lambda_i \rangle_{i<\alpha}$ be an increasing sequence in \mathcal{F} . Set

$$E := \bigcup_{i<\alpha} D_i \quad \text{and} \quad \mu := \bigcup_{i<\alpha} \lambda_i.$$

Then E is densely ordered and $\mu : E \rightarrow C$ is a dense colouring of E . Hence, $\langle E, \mu \rangle \in \mathcal{F}$ is an upper bound of $\langle D_i, \lambda_i \rangle_{i<\alpha}$.

- (II) For a contradiction, suppose that $A \setminus D$ is non-scattered. Then there exists an embedding $e : \mathbb{Q} \rightarrow A \setminus D$. Hence, we can use Lemma 1.8 to

find a dense colouring $\mu : \text{rng } e \rightarrow C$ of $\text{rng } e$. Setting $D' := D \cup \text{rng } e$ and $\lambda' := \lambda \cup \mu$ we obtain a pair $\langle D', \lambda' \rangle \in \mathcal{F}$ with $\langle D, \lambda \rangle < \langle D', \lambda' \rangle$. A contradiction to the maximality of $\langle D, \lambda \rangle$.

(iii) Fix two elements $x < y$ of A and a colour $c \in C$. If there exist elements $x', y' \in D$ with $x < x' < y' < y$, then we can use the denseness of λ to find some $z \in D$ with $x' < z < y'$ and $\chi(z) = \lambda(z) = c$. For a contradiction, suppose that there is at most 1 element of D between x' and y' . Since A is densely ordered, so is the set $I := \{z \in A \mid x' < z < y'\}$. As $I \cap D$ contains at most 1 element, there exists a densely ordered subset $J \subseteq I \setminus D$. Hence, $A \setminus D$ is not scattered. A contradiction to (ii). \square

The Condensation Rank

Orders that are scattered can recursively be decomposed into simple orders. There are several ways to do this, each of which corresponding to a certain notion of a rank which, intuitively, measures the size and distribution of its gaps. Many of the proofs below will be by induction on one of these ranks. We start with one that measures how many 'limits' exist in the given order.

Definition 1.9. Let \mathcal{A} be a linear order.

(a) We denote by $\text{cn} : A \rightarrow A/\approx$ the quotient map for the equivalence relation

$$a \approx b \quad : \text{iff} \quad [a, b] \text{ is finite.}$$

We call \mathcal{A}/\approx the *condensation* of \mathcal{A} and cn is the corresponding *condensation map*.

(b) The *finite-condensation rank* $\text{FC}(\mathcal{A})$, or *FC-rank* for short, is the least ordinal α such that

$$\text{cn}^\alpha(\mathcal{A}) \text{ has only 1 element.}$$

If no such ordinal exists, we set $\text{FC}(\mathcal{A}) := \infty$.

(c) The *generalised finite-condensation rank* $\text{gFC}(\mathcal{A})$, or *generalised FC-rank* for short, is the least ordinal α such that

$$\text{cn}^{\alpha+1}(\mathcal{A}) = \text{cn}^\alpha(\mathcal{A}).$$

(d) We denote by $\text{FC}_*(\mathcal{A})$ and $\text{gFC}_*(\mathcal{A})$ the least ordinal α such that \mathcal{A} can be written as a finite ordered sum of orders of [generalised] FC-rank at most α . \square

Exercise 1.4. Let \mathcal{A} be a linear order and \approx an equivalence relation on A . Prove that the quotient \mathcal{A}/\approx is a linear order if, and only if, every \approx -class is convex. \square

Exercise 1.5. Show that a linear order \mathcal{A} is dense if, and only if, $\text{cn}(\mathcal{A}) \cong \mathcal{A}$. \square

Exercise 1.6. Prove that $\text{gFC}(\mathcal{A}) < |A|^+$, for all linear orders \mathcal{A} . \square

Proposition 1.10. A linear order \mathcal{A} is scattered if, and only if, $\text{FC}(\mathcal{A}) < \infty$.

Proof. (\Leftarrow) Suppose that \mathcal{A} contains a densely ordered subset C . It is sufficient to prove that, for every ordinal α , the restriction of the quotient map cn^α to the set C is injective. We proceed by induction on α . For $\alpha = 0$, the claim is trivial.

For the successor step, suppose that $\text{cn}^\alpha \upharpoonright C$ is injective. Then the image $D := \text{cn}^\alpha[C]$ is densely ordered (as a subset of $\text{cn}^\alpha(\mathcal{A})$). For all $a < b$ in D , it follows that the interval $[a, b]$ is infinite. Hence, the restriction of cn to D is injective. Consequently, so is the restriction of $\text{cn}^{\alpha+1}$ to C .

Finally, suppose that δ is a limit ordinal such that the restriction of cn^α to C is injective, for all $\alpha < \delta$. By definition of cn^δ , it then follows that so is the restriction of cn^δ to C .

(\Rightarrow) Let $\alpha := \text{FC}(\mathcal{A})$ and suppose that $\text{cn}^\alpha(\mathcal{A})$ is not a singleton. Let $\rho := \text{cn}^\alpha : \mathcal{A} \rightarrow \text{cn}^\alpha(\mathcal{A})$ be the quotient map and fix a set $C \subseteq A$ containing exactly one element from every set $\rho^{-1}(a)$. We claim that C is densely ordered. For a contradiction, suppose otherwise. Then there are $a < b$ in C such that C contains no element c with $a < c < b$. Hence, the interval $[\rho(a), \rho(b)]$ of $\text{cn}^\alpha(\mathcal{A})$ is finite, which implies that it is condensed into a single element by another application of cn . Thus, $\text{cn}^{\alpha+1}(\mathcal{A}) \neq \text{cn}^\alpha(\mathcal{A})$. A contradiction. \square

We obtain the following decomposition theorem for linear orders.

Theorem 1.11 (Hausdorff). *Every linear order \mathcal{A} can be written as an ordered sum*

$$\mathcal{A} \cong \sum_{i \in I} \mathfrak{B}_i,$$

where each factor \mathfrak{B}_i is scattered and the index order I is either densely ordered or of size at most 1.

Proof. Set $\alpha := \text{gFC}(\mathcal{A})$, let $\rho := \text{cn}^\alpha : \mathcal{A} \rightarrow \text{cn}^\alpha(\mathcal{A})$ be the corresponding quotient map and $\mathfrak{I} := \rho[\mathcal{A}]$ the quotient. Then

$$\mathcal{A} \cong \sum_{i \in I} \rho^{-1}(i).$$

By Proposition 1.10, each term $\rho^{-1}(i)$ is scattered. Furthermore, $\text{cn}(\mathfrak{I}) = \mathfrak{I}$ implies that $|I| \leq 1$ or \mathfrak{I} is densely ordered. \square

Lemma 1.12. $\text{gFC}(\mathcal{A}) = \text{FC}(\mathcal{A})$, for every scattered linear order \mathcal{A} .

Exercise 1.7. Prove the preceding lemma. J

The reason why we have introduced the ranks FC_* and gFC_* is that they are better behaved when considering decompositions of a given order. Let us present one such result.

Lemma 1.13. *Let \mathcal{A} be a scattered linear order and $P_0 \cup \dots \cup P_{n-1} = A$ a finite partition of A . Then $\text{gFC}_*(\mathcal{A}) = \text{gFC}_*(P_i)$, for some i .*

Proof. Clearly, $P_i \subseteq A$ implies that $\text{gFC}_*(P_i) \leq \text{gFC}_*(\mathcal{A})$. For the converse, it is sufficient to prove that

$$\text{gFC}_*(\mathcal{A}) > \alpha \quad \text{implies} \quad \text{gFC}_*(P_i) > \alpha \text{ for some } i.$$

We do so by induction on α . Hence, suppose that $\text{gFC}_*(\mathcal{A}) > \alpha$. Using Lemma 1.12, it follows that

$$\text{cn}^\alpha(\mathcal{A}) = I, \quad \text{for some infinite order } I.$$

Let $\rho := \text{cn}^\alpha : A \rightarrow I$ be the corresponding quotient map. For every $v \in I$, P_0, \dots, P_{n-1} induce a partition on the order $\rho^{-1}(v)$. Consequently, it follows by inductive hypothesis that there is some index $i_v < n$ such that

$$\text{gFC}_*(P_{i_v} \cap \rho^{-1}(v)) = \text{gFC}_*(\rho^{-1}(v)) = \alpha.$$

Since I is infinite, we can find an infinite subset $H \subseteq I$ such that $i_u = i_v$, for all $u, v \in H$. Let i be this common index. It follows that $\text{cn}^\alpha(P_i)$ is infinite, which implies that $\text{gFC}_*(P_i) > \alpha$. \square

The Hausdorff Rank

The second rank we introduce measures how far the given order is from being a well-order.

Definition 1.14. Let \mathcal{A} be a linear order.

(a) The *Hausdorff rank* $\text{HR}(\mathcal{A})$ of \mathcal{A} is inductively defined as follows. If \mathcal{A} is empty, it has rank -1 . If it is non-empty and finite, it has rank 0. If \mathcal{A} is infinite, the rank is the least ordinal α such that

$$\mathcal{A} \cong \sum_{i \in I} \mathfrak{B}_i,$$

where I is a well-ordering or the opposite of a well-ordering, and each \mathfrak{B}_i is a linear order with $\text{HR}(\mathfrak{B}_i) < \alpha$. If no such ordinal α exists, we set $\text{HR}(\mathcal{A}) := \infty$.

(b) We denote by $\text{HR}_*(\mathcal{A})$ the least ordinal α such that \mathcal{A} can be written as a finite ordered sum of orders of HR-rank at most α .

(c) The *length* $\text{len}(\mathcal{A})$ of \mathcal{A} is the least ordinal β such that there are linear orders $(\mathfrak{B}_i)_{i < \beta}$ with

$$\mathcal{A} \cong \sum_{i \in I} \mathfrak{B}_i \quad \text{and} \quad \text{HR}(\mathfrak{B}_i) < \text{HR}(\mathcal{A}),$$

where the index order I is either equal to β or to its inverse β^{op} . J

Exercise 1.8. Let $h : \mathcal{A} \rightarrow \mathfrak{B}$ be an embedding. Prove that $\text{HR}(\mathfrak{B}) \geq \text{HR}(\mathcal{A})$. J

The relation between the various ranks is as follows.

Lemma 1.15. *Let \mathfrak{A} be a linear order.*

- (a) $\text{HR}_*(\mathfrak{A}) \leq \text{HR}(\mathfrak{A}) \leq \text{HR}_*(\mathfrak{A}) + 1$
- (b) $\text{HR}_*(\mathfrak{A}) \leq \text{FC}(\mathfrak{A})$

Proof. (a) follows directly from the definitions.

(b) We prove the claim by induction on $\alpha := \text{FC}(\mathfrak{A})$. If $\alpha = 0$, \mathfrak{A} is a singleton and $\text{HR}_*(\mathfrak{A}) = 0$.

For the successor step, suppose that $\alpha = \beta + 1$. Then $\mathfrak{S} := \text{cn}^\alpha(\mathfrak{A})$ is an order with more than one element, such that $\text{cn}(\mathfrak{S})$ is a singleton. This implies that every interval $[a, b]$ in I is finite. Consequently, \mathfrak{S} is either finite or isomorphic to ω , ω^{op} , or \mathbb{Z} . Let $\rho := \text{cn}^\beta : \mathfrak{A} \rightarrow \mathfrak{S}$ be the quotient map. By inductive hypothesis, we have

$$\rho^{-1}(i) = \mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)}^i, \quad \text{for } n(i) < \omega \text{ and } \text{HR}_*(\mathfrak{C}_j^i) \leq \beta.$$

Depending on whether or not \mathfrak{S} is isomorphic to \mathbb{Z} , we obtain one of the following decompositions

$$\begin{aligned} \mathfrak{A} &= \sum_{i < 0} [\mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)}^i] + \sum_{i \geq 0} [\mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)}^i], & \text{if } \mathfrak{S} \cong \mathbb{Z}, \\ \mathfrak{A} &= \sum_{i \in I} [\mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)}^i], & \text{if } \mathfrak{S} \not\cong \mathbb{Z}. \end{aligned}$$

In both cases, it follows that $\text{HR}_*(\mathfrak{A}) \leq \beta + 1 = \alpha$.

Finally, suppose that α is a limit ordinal. Fixing an element $a \in A$, we can write

$$A = \bigcup_{i < \alpha} [a]_i, \quad \text{where } [a]_i := \{b \in A \mid \text{cn}^i(b) = \text{cn}^i(a)\}.$$

For $i < \alpha$, we set

$$\begin{aligned} B_i &:= \{b < a \mid b \in [a]_{i+1} \setminus [a]_i\}, \\ C_i &:= \{b > a \mid b \in [a]_{i+1} \setminus [a]_i\}. \end{aligned}$$

Let \mathfrak{B}_i and \mathfrak{C}_i be the corresponding suborders of \mathfrak{A} . Since

$$\text{FC}([a]_i) \leq i + 1 < \alpha,$$

it follows by inductive hypothesis that

$$\mathfrak{B}_i = \mathfrak{B}_0^i + \cdots + \mathfrak{B}_{m(i)-1}^i \quad \text{and} \quad \mathfrak{C}_i = \mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)-1}^i$$

where $\text{HR}(\mathfrak{B}_j^i), \text{HR}(\mathfrak{C}_j^i) < \alpha$ and $m(i), n(i) < \omega$. Then

$$\mathfrak{A} \cong \sum_{i \in \alpha^{\text{op}}} [\mathfrak{B}_0^i + \cdots + \mathfrak{B}_{m(i)-1}^i] + \{a\} + \sum_{i < \alpha} [\mathfrak{C}_0^i + \cdots + \mathfrak{C}_{n(i)-1}^i].$$

Hence, \mathfrak{A} is the sum of 3 orders of HR-rank at most α , which implies that $\text{HR}_*(\mathfrak{A}) \leq \alpha$. \square

Proposition 1.16. *A linear order \mathfrak{A} is scattered if, and only if, $\text{HR}(\mathfrak{A}) < \infty$.*

Proof. (\Rightarrow) If \mathfrak{A} is scattered, it follows by Proposition 1.10 and Lemma 1.15 that

$$\text{HR}(\mathfrak{A}) \leq \text{HR}_*(\mathfrak{A}) + 1 \leq \text{FC}(\mathfrak{A}) + 1 < \infty.$$

(\Leftarrow) Suppose that $\alpha := \text{HR}(\mathfrak{A}) < \infty$. Then

$$\mathfrak{A} \cong \sum_{i \in I} \mathfrak{B}_i,$$

where I is a well-ordering or an inverse well-ordering and each \mathfrak{B}_i is a linear order with $\text{HR}(\mathfrak{B}_i) < \alpha$.

We prove by induction on α that \mathfrak{A} is scattered. Suppose that $\eta : \mathbb{Q} \rightarrow \mathfrak{A}$ is an order-preserving injection. By inductive hypothesis, each \mathfrak{B}_i is scattered. Hence, the image of η contains at most one element of B_i . Consequently, we obtain an order-preserving injection $\eta' : \mathbb{Q} \rightarrow I$. But I is scattered. A contradiction. \square

As with the FC-rank, the variant HR_* has been introduced mostly because it is better behaved with respect to decompositions.

Lemma 1.17. Let $\mathfrak{A}, \mathfrak{B}, \mathfrak{A}_i$ be linear orders.

- (a) $\text{HR}_*(\mathfrak{A} + \mathfrak{B}) = \max \{ \text{HR}_*(\mathfrak{A}), \text{HR}_*(\mathfrak{B}) \}$.
- (b) $\text{HR}_*\left(\sum_{i \in I} \mathfrak{A}_i\right) \leq \gamma + \text{HR}_*(I)$ where $\gamma := \sup_i \text{HR}_*(\mathfrak{A}_i)$.

Proof. (a) Let α be the maximum of $\text{HR}_*(\mathfrak{A})$ and $\text{HR}_*(\mathfrak{B})$. By definition, we can write

$$\mathfrak{A} = \mathfrak{A}_0 + \cdots + \mathfrak{A}_{m-1} \quad \text{and} \quad \mathfrak{B} = \mathfrak{B}_0 + \cdots + \mathfrak{B}_{n-1},$$

for orders of Hausdorff rank at most α . Consequently,

$$\mathfrak{A} + \mathfrak{B} = \mathfrak{A}_0 + \cdots + \mathfrak{A}_{m-1} + \mathfrak{B}_0 + \cdots + \mathfrak{B}_{n-1},$$

which implies that $\text{HR}_*(\mathfrak{A} + \mathfrak{B}) \leq \alpha$. The other direction follows from the fact that \mathfrak{A} and \mathfrak{B} are embedded in $\mathfrak{A} + \mathfrak{B}$ (cf. Exercise 1.8).

(b) We prove the claim by induction on $\alpha := \text{HR}_*(I)$. If $\alpha = 0$, then $I = \{i_0, \dots, i_{n-1}\}$ is finite and the claim follows by (a). If $\alpha > 0$, then

$$I = \sum_{k \in K} J_k,$$

where $\text{HR}_*(J_k) < \alpha$ and K is a well-order or the opposite of one. By inductive hypothesis, we have

$$\text{HR}_*\left(\sum_{j \in J_k} \mathfrak{A}_j\right) \leq \gamma + \text{HR}_*(J_k) < \gamma + \alpha.$$

Hence,

$$\text{HR}_*\left(\sum_{i \in I} \mathfrak{A}_i\right) = \text{HR}_*\left(\sum_{k \in K} \sum_{j \in J_k} \mathfrak{A}_j\right) \leq \gamma + \alpha.$$

□

The Lexicographic Order

As an application, we consider linear orders of the form $\langle T, \leq_{\text{lex}} \rangle$ where T is a tree domain.

Definition 1.18. Let D be a linearly ordered set. The *lexicographic ordering* \leq_{lex} on D^* is defined by

$$u \leq_{\text{lex}} v \quad : \text{iff} \quad u \leq_{\text{pf}} v, \quad \text{or} \\ u = wax \text{ and } v = wby, \text{ for } a, b \in D, \quad w, x, y \in D^* \\ \text{with } a < b.$$

We start with a lemma.

Lemma 1.19. Let \mathfrak{A} be a linear order and $S \subseteq A$ a subset of the form $S = \sum_{i \in I} C_i$, where I is some infinite linear order and every C_i has a minimal element and is convex in A . Then

$$\text{gFC}_*(\mathfrak{A}) \geq \sup_{i \in I} (\text{gFC}_*(C_i) + 1).$$

Proof. We prove by induction on α that

$$\text{gFC}_*(C_i) \geq \alpha, \text{ for all } i, \quad \text{implies} \quad \text{gFC}_*(\mathfrak{A}) \geq \alpha + 1.$$

Suppose that $\text{gFC}_*(C_i) \geq \alpha$, for all i . Then $\text{cn}^\alpha(C_i)$ is not a finite sum of dense orders. We have to show that $\text{cn}^\alpha(\mathfrak{A})$ is infinite but not densely ordered.

For a contradiction, suppose otherwise. Note that $\text{cn}^\alpha(C_i)$ is a convex subset of $\text{cn}^\alpha(\mathfrak{A})$. Consequently, if $\text{cn}^\alpha(\mathfrak{A})$ were densely ordered, so would be every $\text{cn}^\alpha(C_i)$. A contradiction. Hence, we may assume that $\text{cn}^\alpha(\mathfrak{A})$ is finite. Then there exists an infinite convex subset $K \subseteq I$ such that cn^α contracts $\bigcup_{k \in K} C_k$ to a single element. This implies that $\text{cn}^\alpha(C_k) = 1$, for every $k \in K$. Since $\text{gFC}_*(C_k) \geq \alpha$, it follows that α is a successor ordinal and $\text{cn}^{\alpha-1}(C_k)$ is finite. Hence, $\text{gFC}_*(C_k) < \alpha$. A contradiction. □

It is possible to establish a relationship between the Lifschitz-Shelah rank (cf. Definition V.5.8) of a tree $\langle T, \leq_{\text{pf}} \rangle$ and the FC-rank of $\langle T, \leq_{\text{lex}} \rangle$. In the same way we have worked above with the generalised FC-rank instead of the ordinary one, we need a slight variant of the LS-rank that is also meaningful for trees that are not thin.

Definition 1.20. The generalised Lifsches-Shelah rank $\text{gLS}(\mathfrak{T})$ of a forest \mathfrak{T} is the least ordinal α such that $\partial^{\alpha+1}\mathfrak{T} = \partial^\alpha\mathfrak{T}$. \downarrow

Proposition 1.21. Let $\mathfrak{T} = \langle T, \leq_{\text{pf}} \rangle$ be a finitely-branching order-tree. Then

$$\text{gLS}(\mathfrak{T}) \leq \text{gFC}_*(\langle T, \leq_{\text{lex}} \rangle).$$

Proof. Set $L(\mathfrak{T}) := \langle T, \leq_{\text{lex}} \rangle$. We prove by induction on α that

$$\text{gLS}(\mathfrak{T}|_v) \geq \alpha \quad \text{implies} \quad \text{gFC}_*(L(\mathfrak{T}|_v)) \geq \alpha, \quad \text{for all } v \in T.$$

Fix a vertex $v \in T$ with $\text{gLS}(\mathfrak{T}|_v) \geq \alpha$. For $\alpha = 0$, the claim is trivial. If α is a limit ordinal, it follows by inductive hypothesis that

$$\text{gFC}_*(L(\mathfrak{T}|_v)) \geq \beta, \quad \text{for all } \beta < \alpha.$$

Consequently, $\text{gFC}_*(L(\mathfrak{T}|_v)) \geq \alpha$.

Hence, we only have to consider the successor step. Suppose that $\alpha = \beta + 1$, for some β . Then there exists an infinite branch ζ of $\mathfrak{T}|_v$ such that there are infinitely many vertices w on ζ with a successor u that is not on the branch and the attached subtree $\mathfrak{T}|_u$ has rank at least β . Then

$$L(\mathfrak{T}|_v) = \sum_{w <_{\text{pf}} \zeta} (1 + \mathfrak{A}_w) + \sum_{\substack{w <_{\text{pf}} \zeta \\ \text{in reverse order}}} \mathfrak{B}_w,$$

for where the orders \mathfrak{A}_w and \mathfrak{B}_w are finite sums of orders of the form $L(\mathfrak{T}|_z)$, for successors z of w . Furthermore, there are infinitely many vertices w such that the sum for \mathfrak{A}_w or that for \mathfrak{B}_w contains a term of the form $L(\mathfrak{T}|_u)$ with $\text{gLS}(\mathfrak{T}|_u) \geq \beta$. By symmetry, we may assume that this term belongs to \mathfrak{A}_w . Then $L(\mathfrak{T}|_v)$ contains a suborder of the form

$$\sum_{i < \omega} \mathfrak{C}_i \quad \text{where} \quad \text{gFC}_*(\mathfrak{C}_i) \geq \beta.$$

Consequently, the claim follows by Lemma 1.19. \square

Finally, let us show that the Hausdorff rank of a linear order tells us which trees we can embed into it.

Lemma 1.22. Let $\mathfrak{T}_n = \langle T_n, \leq_{\text{pf}} \rangle$ be the tree of height $n < \omega$ with domain

$$T_n := \{ w \in \mathbb{Z}^* \mid |w| \leq n, w(i) \geq 0 \text{ for } i \text{ even}, w(i) \leq 0 \text{ for } i \text{ odd} \},$$

For every linear order \mathfrak{A} with Hausdorff rank $\text{HR}(\mathfrak{A}) > n$, there exists an embedding $\eta : \langle T_n, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{A}$.

Proof. Let $\text{HR}(\mathfrak{A}) > n$. If \mathfrak{A} contains a densely ordered interval, we can use Lemma 1.3 to embed $\langle T_n, \leq_{\text{lex}} \rangle$. Hence, we may assume that \mathfrak{A} is scattered. Let \mathfrak{L}_n be the suborder of $\langle T_n, \leq_{\text{lex}} \rangle$ consisting of the leaves of \mathfrak{T}_n . That is,

$$\mathfrak{L}_0 := 1 \quad \text{and} \quad \mathfrak{L}_{n+1} := \sum_{i < \omega} \mathfrak{L}_n^{\text{op}}, \quad \text{for } n < \omega.$$

To simplify our task, we will split the embedding into two parts

$$\langle T_n, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{L}_n \quad \text{and} \quad \mathfrak{L}_n \rightarrow \mathfrak{A}.$$

Let $\mathfrak{S}_n = \langle S_n, \leq_{\text{pf}} \rangle$ be the tree obtained from \mathfrak{T}_n by reversing the successor ordering. That is,

$$S_n := \{ w \in \mathbb{Z}^* \mid |w| \leq n, w(i) \geq 0 \text{ for } i \text{ odd}, w(i) \leq 0 \text{ for } i \text{ even} \},$$

We start by constructing embeddings

$$\eta_n : \langle T_n, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{L}_n \quad \text{and} \quad \zeta_n : \langle S_n, \leq_{\text{lex}} \rangle \rightarrow 1 + \mathfrak{L}_n^{\text{op}}$$

by induction on n . For $n = 0$, we have

$$\langle T_0, \leq_{\text{lex}} \rangle = 1 = \mathfrak{L}_0 \quad \text{and} \quad \langle S_0, \leq_{\text{lex}} \rangle = 1 \rightarrow 1 + 1 = 1 + \mathfrak{L}_0^{\text{op}}.$$

For the inductive step, suppose that we have already constructed η_n and ζ_n .

We use these to obtain embeddings

$$\begin{aligned}
 \langle T_{n+1}, \leq_{\text{lex}} \rangle &\cong 1 + \sum_{n \in \omega} \langle S_n, \leq_{\text{lex}} \rangle \\
 &\rightarrow 1 + \sum_{n \in \omega} (1 + \mathfrak{L}_n^{\text{op}}) \\
 &\cong 2 + \sum_{n < \omega} \mathfrak{L}_n^{\text{op}} \\
 &\rightarrow \sum_{n < \omega} \mathfrak{L}_n^{\text{op}} \cong \mathfrak{L}_{n+1}, \\
 \langle S_{n+1}, \leq_{\text{lex}} \rangle &\cong 1 + \sum_{n \in \omega^{\text{op}}} \langle T_n, \leq_{\text{lex}} \rangle \rightarrow 1 + \sum_{n \in \omega^{\text{op}}} \mathfrak{L}_n \cong 1 + \mathfrak{L}_{n+1}^{\text{op}}.
 \end{aligned}$$

To conclude the proof it is now sufficient to show that $\mathfrak{L}_{n+1} \rightarrow \mathfrak{A}$ or $\mathfrak{L}_{n+1}^{\text{op}} \rightarrow \mathfrak{A}$. Then we obtain the desired embedding as

$$\langle T_n, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{L}_n \rightarrow \mathfrak{L}_{n+1} \rightarrow \mathfrak{A} \quad \text{or} \quad \langle T_n, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{L}_n \rightarrow \mathfrak{L}_{n+1}^{\text{op}} \rightarrow \mathfrak{A}.$$

We will establish the following slightly more precise statement by induction on n . Suppose that $\text{HR}(\mathfrak{A}) = n + 1$ and

$$\mathfrak{A} = \sum_{i \in I} \mathfrak{B}_i \quad \text{where} \quad \text{HR}(\mathfrak{B}_i) \leq n \text{ and } I \text{ is a well-order,}$$

then there exists an embedding $\mathfrak{L}_{n+1} \rightarrow \mathfrak{A}$. (Applying the operation $-^{\text{op}}$ it follows that, if I is the opposite of a well-order, we obtain an embedding $\mathfrak{L}_{n+1}^{\text{op}} \rightarrow \mathfrak{A}$.)

If $n = 0$, then \mathfrak{A} is an infinite well-ordering and there exists an embedding $\mathfrak{L}_1 \rightarrow \mathfrak{A}$.

Suppose that $n > 0$. By Lemma 1.17, there exists an infinite subset $I^0 \subseteq I$ such that

$$\text{HR}(\mathfrak{B}_i) = n, \quad \text{for all } i \in I^0.$$

Furthermore, for every $i \in I^0$, we can find a decomposition

$$\mathfrak{B}_i = \sum_{j \in J_i} \mathfrak{B}'_{ij} \quad \text{with} \quad \text{HR}(\mathfrak{B}'_{ij}) < n.$$

There exists an infinite subset $I^1 \subseteq I^0$ such that, for all $i \in I^1$, J_i is the opposite of a well-order since, otherwise, it would follow by Lemma 1.17 that

$$\mathfrak{A} \cong \sum_{i \in I} \sum_{j \in J_i} \mathfrak{B}'_{ij}$$

would have rank at most n . By inductive hypothesis, there exist embeddings $\mathfrak{L}_n^{\text{op}} \rightarrow \mathfrak{B}_i$, for $i \in I^1$. Combining them we obtain the desired embedding $\mathfrak{L}_{n+1} \rightarrow \sum_{i \in I^1} \mathfrak{B}_i \rightarrow \mathfrak{A}$. \square

Exercise 1.9. Let $C \subseteq \{0, 1\}^*$ be an antichain in the infinite binary tree such that $\langle C, \leq_{\text{lex}} \rangle$ is well-ordered. Prove that the tree with domain $\downarrow C$ has Cantor-Bendixson rank α if, and only if, $\langle C, \leq_{\text{lex}} \rangle$ has an order type in the interval $[\omega^\alpha, \omega^{\alpha+1})$. \square

2 Partition Theorems

The Theorem of Ramsey can be generalised from ω and finite linear orders to arbitrary ones. When doing so new phenomena arise as the homogeneous subset one obtains might be embedded in the order in different ways. This gives rise to several different variants of the theorem. We start with a version of the Pigeon Hole Principle for dense orders.

Proposition 2.1. Let $\lambda : A \rightarrow C$ be a labelling of a dense linear order A by a finite set C of colours. There exists an infinite convex subset $I \subseteq A$ such that, for every colour $c \in \lambda[I]$, the preimage $\lambda^{-1}(c) \cap I$ is infinite and dense in I .

Proof. Suppose that $C = \{c_0, \dots, c_{n-1}\}$. We inductively define an increasing sequence $C_0 \subseteq \dots \subseteq C_n \subseteq C$ of colours and a decreasing sequence $A \supseteq I_0 \supseteq \dots \supseteq I_n$ of convex subsets as follows. We start with $I_0 := A$ and $C_0 := \emptyset$. Once we have defined I_i and C_i , we distinguish two cases. If c_i occurs infinitely often in I_i and it is dense in I_i , we set

$$C_{i+1} := C_i \cup \{c_i\} \quad \text{and} \quad I_{i+1} := I_i.$$

Otherwise, we set $C_{i+1} := C_i$ and we choose a convex subset $I_{i+1} \subseteq I_i$ such that $c_i \notin \lambda[I_{i+1}]$ and $|I_{i+1}| > 1$.

It follows that I_n is a convex subset of A with at least 2 elements. Hence, I_n is infinite and every colour $c \in C_n = \lambda[I_n]$ is dense in I_n . \square

As usual we obtain stronger statements if we assume that the colours form a semigroup since in that case we can use Ramseyan splits. The following two Ramsey theorems provide an example.

Theorem 2.2. *Let $\mathcal{A} = \langle A, \leq \rangle$ be an infinite linear order, \mathcal{S} a finite semigroup, and λ an additive labelling of \mathcal{A} .*

- (a) *If \mathcal{A} is a well-order, there exists a cofinal subset $H \subseteq A$ such that the restriction of λ to H is constant.*
- (b) *If \mathcal{A} is densely ordered, there exists an interval $I \subseteq A$ containing a set $H \subseteq I$ such that H is dense in I and the restriction of λ to H is constant.*

Proof. By Theorem II.3.7, there exists a Ramseyan split $\sigma : A \rightarrow [N]$ of λ with $N < \omega$.

(a) Let $k < N$ be the largest number such that the set $\sigma^{-1}(k)$ is cofinal in A . Then there is some element $a_0 \in A$ such that

$$\sigma(b) \leq k, \quad \text{for all } b \geq a_0.$$

We set

$$H := \sigma^{-1}(k) \cap \uparrow a_0.$$

By choice of k , this set is cofinal in A . Furthermore, consider elements $a < b$ and $a' < b'$ in H . By symmetry, we may assume that $a \leq a'$. Then $a \sqsubset_\sigma b$ and $a \sqsubseteq_\sigma a' \sqsubset_\sigma b'$, and the fact that σ is Ramseyan implies that $\lambda(a, b) = \lambda(a', b')$.

(b) By Proposition 2.1, there exists an infinite convex set $I \subseteq A$ such that, for every $k \in \sigma[I]$, the preimage $\sigma^{-1}(k)$ is infinite and dense in I . Set

$$k := \max \sigma[I] \quad \text{and} \quad H := \sigma^{-1}(k) \cap I.$$

Then H is infinite and dense in I and, thus, also in H . Furthermore, consider elements $a < b$ and $a' < b'$ in H . By symmetry, we may assume that $a \leq a'$. Then $a \sqsubset_\sigma b$ and $a \sqsubseteq_\sigma a' \sqsubset_\sigma b'$, and the fact that σ is Ramseyan implies that $\lambda(a, b) = \lambda(a', b')$. \square

As an application, we present a variant of the Pigeon Hole Principle with infinitely many colours. To do so, we have to assume that the colours are definable.

Lemma 2.3. *Let \mathcal{A} be a linear order and \approx an MSO-definable equivalence relation on A (possibly definable using monadic parameters). Then \approx has only finitely many unbounded classes.*

Proof. Let m be the quantifier rank of the formula defining \approx and let \tilde{P} be the monadic parameters used. To simplify the notation, we set $\mathcal{A}^+ := \langle \mathcal{A}, \tilde{P} \rangle$. Consider the additive labelling

$$\lambda(i, j) := \text{Th}_{\text{MSO}}^{m+1}(\langle \mathcal{A}^+, i \rangle|_{[i, j]}).$$

We fix a strictly increasing cofinal sequence $(a_i)_{i < \delta}$ in A of minimal length. For a contradiction, suppose that \approx has infinitely many unbounded classes. Then δ must be a limit ordinal. By Theorem 2.2 (a), there exists a cofinal subset $I \subseteq \delta$ and a theory θ such that

$$\lambda(a_i, a_j) = \theta, \quad \text{for all } i < j \text{ in } I.$$

Let $(b_i)_{i < \delta}$ be an enumeration of $(a_i)_{i \in I}$. Given $i < \delta$ and $c < b_i$ such that $[c]_\approx$ is unbounded, let $i < j < \delta$ be some index such that

$$[b_i, b_j) \cap [c]_\approx \neq \emptyset.$$

This property can be expressed in MSO_{m+1} by the formula

$$\psi(b_i, b_j, c) := \exists x [b_i \leq x < b_j \wedge x \approx c].$$

Consequently,

$$\begin{aligned}
& \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+, b_i, b_j, c) \\
&= \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{\downarrow b_i}, c) + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{[b_i, b_j]}, b_i) \\
&\quad + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{[b_j, b_{j+1}]}, b_j) + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{\uparrow b_{j+1}}) \\
&= \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{\downarrow b_i}, c) + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{[b_i, b_{i+1}]}, b_i) \\
&\quad + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{[b_{i+1}, b_{j+1}]}, b_{i+1}) + \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+|_{\uparrow b_{j+1}}) \\
&= \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{A}^+, b_i, b_{i+1}, c),
\end{aligned}$$

implies that

$$[b_i, b_{i+1}) \cap [c]_{\approx} \neq \emptyset, \quad \text{for all } c < b_i.$$

Let n be the number of all MSO_{m+2} -theories with $|\bar{P}| + 1$ monadic parameters and fix an index $i < \delta$ such that the interval $[b_o, b_i)$ intersects more than n different unbounded \approx -classes, say, C_o, \dots, C_n . By choice of n , there are indices $k < l$ such that

$$\langle \mathfrak{A}^+, C_k \rangle \equiv_{\text{MSO}}^{m+2} \langle \mathfrak{A}^+, C_l \rangle.$$

Setting

$$D := (C_k \cap \downarrow b_i) \cup (C_l \cap [b_i, b_{i+1})) \cup (C_k \cap \uparrow b_{i+1})$$

it follows that

$$\begin{aligned}
& \text{Th}_{\text{MSO}}^{m+2}(\mathfrak{A}^+, D) \\
&= \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_k \rangle|_{\downarrow b_i}) + \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_l \rangle|_{[b_i, b_{i+1}]}) \\
&\quad + \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_k \rangle|_{\uparrow b_{i+1}}) \\
&= \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_k \rangle|_{\downarrow b_i}) + \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_k \rangle|_{[b_i, b_{i+1}]}) \\
&\quad + \text{Th}_{\text{MSO}}^{m+2}(\langle \mathfrak{A}^+, C_k \rangle|_{\uparrow b_{i+1}}) \\
&= \text{Th}_{\text{MSO}}^{m+2}(\mathfrak{A}^+, C_k).
\end{aligned}$$

But C_k is an \approx -class while D is not. We obtain a contradiction since there exists an MSO_{m+2} -formula

$$\vartheta(X) := \exists x \forall z [Xz \leftrightarrow x \approx z]$$

expressing that the set X forms an \approx -class. \square

Using the lexicographic ordering to turn a tree into a linear order as in the previous section, we can derive the following partition theorem for trees.

Proposition 2.4. *Let $\mathfrak{T} = \langle T, \leq_{\text{pf}}, \leq_{\text{so}} \rangle$ be a non-empty successor-ordered tree of height at most $n < \omega$, and let λ be an additive labelling of the order $\langle T, \leq_{\text{lex}} \rangle$. There exists a non-empty prefix $P \subseteq T$ with the following properties.*

- For $u \leq_{\text{lex}} v$ in P , the colour $\lambda(u, v)$ only depends on $|u|$, $|v|$, and $|u \sqcap_{\text{pf}} v|$.
- If $v \in P$ is a leaf of P , it is also a leaf of T .
- If $v \in P$ has infinitely many successors in T , it has infinitely many ones in P .

Proof. We prove the claim by induction on \mathfrak{T} . Hence, we may assume by inductive hypothesis that, for every $u \in \text{Suc}(\langle \rangle)$, we have already found some prefix $P_u \subseteq T|_u$ of the subtree $\mathfrak{T}|_u$ with the above properties.

If $\text{Suc}(\langle \rangle)$ is finite, we pick one $u \in \text{Suc}(\langle \rangle)$ and set $P := \{\langle \rangle\} \cup P_u$. Suppose otherwise. For every $u \in \text{Suc}(\langle \rangle)$, we choose a subset $P_u^o \subseteq P_u$ as follows. If $v \in P_u$ is a vertex with finitely many successors (in P_u), we remove all but one of them (and the attached subtrees). If $v \in P_u$ has infinitely many successors (in P_u), we remove the least and the greatest of them (if they exist; again together with the attached subtrees). Let P_u^o be the remaining set of vertices.

By the Pigeon Hole Principle, there exists an infinite subset $U \subseteq \text{Suc}(\langle \rangle)$ such that the colours associated with P_u are the same for every $u \in U$. That is, there exist colours θ_{ijk} , for $o < i, j, k \leq n$, such that

$$\lambda(v, v') = \theta_{|v|, |v'|, |v \sqcap_{\text{pf}} v'|}, \quad \text{for all } v <_{\text{lex}} v' \text{ in } P_u \text{ with } u \in U.$$

For every $v \in P_u^o$ that is not a leaf, we fix a ‘distinguished successor’ $s(v) \in \text{Suc}(v) \cap P_u^o$, and we define an additive labelling of U by

$$\mu(u, u') := (\lambda(s^i(u), s^j(u'))_{i, j < n}.$$

(If $s^i(u)$ or $s^j(u')$ is not defined, we set $\lambda(s^i(u), s^j(u')) := \perp$.) By the Theorem of Ramsey, there exists an infinite subset $U_\circ \subseteq U$ and colours σ_{ij} such that

$$\mu(u, u') = (\sigma_{ij})_{ij}, \quad \text{for all } u <_{\text{lex}} u' \text{ in } U_\circ.$$

We set

$$P := \{\langle \rangle\} \cup \bigcup_{u \in U_\circ} P_u^\circ.$$

To show that this set has the desired properties, let $v, v' \in P$ be vertices with $v <_{\text{lex}} v'$ and set $i := |v|$, $j := |v'|$, and $k := |v \cap v'|$. If $k > \circ$, we have $v, v' \in P_u^\circ$, for some $u \in U$, which implies that

$$\lambda(v, v') = \theta_{ijk}.$$

Hence, suppose that $k = \circ$. Then $v \in P_u^\circ$ and $v' \in P_{u'}^\circ$, for some $u <_{\text{so}} u'$ in U_\circ . Let p and q be the maximal numbers such that $s^p(u) \leq_{\text{pf}} v$ and $s^q(u') \leq_{\text{pf}} v'$. Then $s^p(u) = v$ or $s^p(u)$ has infinitely many successors, and similarly for $s^q(u')$. First, suppose that $s^p(u) \neq v$ and $s^q(u') \neq v'$. By choice of P_u° , we can then choose a successor $w \in P_u$ of $s^p(u)$ that is larger than both $s^{p+1}(u)$ and the successor v_\circ with $v_\circ \leq_{\text{pf}} v$. Similarly, $s^q(u')$ has a successor $w' \in P_{u'}$ that is smaller than both $s^{q+1}(u')$ and the successor v'_\circ with $v'_\circ \leq_{\text{pf}} v'$. It follows that

$$\begin{aligned} \lambda(v, v') &= \lambda(v, w) + \lambda(w, v') \\ &= \theta_{i(p+2)(p+1)} + \lambda(w, v') \\ &= \lambda(s^{i-1}(u), w) + \lambda(w, v') \\ &= \lambda(s^{i-1}(u), v') \\ &= \lambda(s^{i-1}(u), w') + \lambda(w', v') \\ &= \lambda(s^{i-1}(u), w') + \theta_{(q+2)j(q+1)} \\ &= \lambda(s^{i-1}(u), w') + \lambda(w', s^{j-1}(u')) \\ &= \lambda(s^{i-1}(u), s^{j-1}(u')) \\ &= \sigma_{(i-1)(j-1)}. \end{aligned}$$

If $s^p(u) = v$ and $s^q(u') \neq v'$, we similarly obtain

$$\begin{aligned} \lambda(v, v') &= \lambda(v, w') + \lambda(w', v') \\ &= \lambda(v, w') + \theta_{(q+2)j(q+1)} \\ &= \lambda(v, w') + \lambda(w', s^{j-1}(u')) \\ &= \lambda(v, s^{j-1}(u')) \\ &= \lambda(s^{i-1}(u), s^{j-1}(u')) = \sigma_{(i-1)(j-1)} \end{aligned}$$

(where w' is chosen as above). In the other two cases, it similarly follows that

$$\lambda(v, v') = \sigma_{(i-1)(j-1)}. \quad \square$$

As an application we consider scattered orders, where we obtain homogeneous subsets of order type \mathbb{Z} .

Theorem 2.5. *Let \mathfrak{A} be a linear order and λ an additive labelling with less than n colours of \mathfrak{A} . If $\text{HR}(\mathfrak{A}) > n$, then there exists a subset $Z \subseteq A$ of order type \mathbb{Z} such that*

$$\lambda(u, v) = \lambda(u', v'), \quad \text{for all } u < v \text{ and } u' < v' \text{ in } Z.$$

Proof. Suppose that $\text{HR}(\mathfrak{A}) > n$. Then we can use Lemma 1.22 to find an embedding $\langle T, \leq_{\text{lex}} \rangle \rightarrow \mathfrak{A}$ where $\mathfrak{T} = \langle T, \leq_{\text{pf}} \rangle$ is the tree of height n with domain

$$T := \{w \in \mathbb{Z}^* \mid |w| \leq n, w(i) \geq \circ \text{ for } i \text{ even}, w(i) \leq \circ \text{ for } i \text{ odd}\}.$$

Let $L \subseteq T$ be the set of leaves of \mathfrak{T} and let $P \subseteq T$ be the prefix obtained by applying Proposition 2.4 to \mathfrak{T} and λ . It follows that, for every $i \leq n$, there is some theory θ_i such that

$$\lambda(u, v) = \theta_{|u \cap_{\text{pf}} v|}, \quad \text{for all } u <_{\text{lex}} v \text{ in } P \cap L.$$

By assumption on n , there are indices $\circ < i < j \leq n$ such that $\theta_i = \theta_j$. First, suppose that $j = i + 1$. Pick some vertex $w \in P$ of length $|w| = i$. If

$\text{Suc}(w)$ has order type ω , let u_o be the first element of $\text{Suc}(w) \cap P$ and set

$$\begin{aligned} U_o &:= \text{Suc}(u_o) \cap P, \\ V_o &:= \{s(u) \mid u \in \text{Suc}(w) \cap P \text{ and } u > u_o\}, \\ U &:= \{s^{n-(j+1)}(u) \mid u \in U_o\}, \\ V &:= \{s^{n-(j+1)}(v) \mid v \in V_o\}. \end{aligned}$$

Then $Z := U + V$ has order type \mathbb{Z} and, by choice of P , we have

$$\begin{aligned} \lambda(u, v) &= \theta_j, \quad \text{for all } u <_{\text{lex}} v \text{ with } u, v \in U, \\ \lambda(u, v) &= \theta_i, \quad \text{for all } u <_{\text{lex}} v \text{ with } u \in U + V \text{ and } v \in V. \end{aligned}$$

Since $\theta_i = \theta_j$, this implies that

$$\lambda(u, v) = \lambda(u', v'), \quad \text{for all } u <_{\text{lex}} v \text{ and } u' <_{\text{lex}} v' \text{ in } Z,$$

as desired.

Similarly, if $\text{Suc}(w)$ has order type ω^{op} , we take for u_o the last element of $\text{Suc}(w) \cap P$ and we set

$$\begin{aligned} U_o &:= \{s(u) \mid u \in \text{Suc}(w) \cap P \text{ and } u < u_o\}, \\ V_o &:= \text{Suc}(u_o) \cap P, \\ U &:= \{s^{n-(j+1)}(u) \mid u \in U_o\}, \\ V &:= \{s^{n-(j+1)}(v) \mid v \in V_o\}. \end{aligned}$$

Again, it follows that $Z := U + V$ has order type \mathbb{Z} and

$$\begin{aligned} \lambda(u, v) &= \theta_i, \quad \text{for all } u <_{\text{lex}} v \text{ with } u \in U \text{ and } v \in U + V, \\ \lambda(u, v) &= \theta_j, \quad \text{for all } u <_{\text{lex}} v \text{ with } u, v \in V. \end{aligned}$$

As above, it follows that the set Z is homogeneous.

Finally, suppose that $j > i + 1$. Let $u_o <_{\text{lex}} u_1 <_{\text{lex}} u_2$ be vertices of $P \cap L$ such that $|u_o \sqcap u_2| = i$ and $|u_o \sqcap u_1| = j - 1$. Then

$$\theta_i = \lambda(u_o, u_2) = \lambda(u_o, u_1) + \lambda(u_1, u_2) = \theta_{j-1} + \theta_i,$$

which implies that $\theta_j = \theta_{j-1} + \theta_j$. Similarly, considering vertices $u_o <_{\text{lex}} u_1 <_{\text{lex}} u_2$ of length $|u_i| = n$ such that $|u_o \sqcap u_2| = j - 1$ and $|u_1 \sqcap u_2| = j$, we obtain

$$\theta_{j-1} = \lambda(u_o, u_2) = \lambda(u_o, u_1) + \lambda(u_1, u_2) = \theta_{j-1} + \theta_j.$$

Combining these two equations yields $\theta_j = \theta_{j-1} + \theta_j = \theta_{j-1}$. Consequently, the claim follows as in the case $j = i + 1$ above. \square

Finally, we take a look at well-orders where the colouring is invariant under translations.

Proposition 2.6. *Let λ be an additive colouring of a well-order \mathfrak{A} such that the value $\lambda(i, j)$ only depends on the order type of the interval $[i, j)$. There exists a colour $e \in S$ such that*

$$\lambda(i, j) = e, \quad \text{for all } i < j \text{ such that } [i, j) \text{ has order type } \omega^k \text{ with } k \geq |S|.$$

Furthermore, e is idempotent, provided that the order type of \mathfrak{A} is at least $\omega^{|S|+1}$.

Proof. Let μ be the function such that

$$\lambda(i, j) = \mu(\alpha), \quad \text{for } [i, j) \text{ of order type } \alpha.$$

Set $n := |S|$. If the order type of \mathfrak{A} is less than ω^n , we can choose an arbitrary idempotent $e \in S$. If the order type is in $[\omega^n, \omega^{n+1})$, we can set $e := \mu(\omega^n)$. Finally, suppose that the order type of \mathfrak{A} is at least ω^{n+1} . Then there are exponents $k < l \leq n$ with $\mu(\omega^k) = \mu(\omega^l)$. If $l = k + 1 = n$ and the order type of \mathfrak{A} is less than ω^{n+2} , we can set $e := \mu(\omega^k)$. Then it follows by additivity of λ that

$$e + e = \mu(\omega^k) + \mu(\omega^{k+1}) = \mu(\omega^k + \omega^{k+1}) = \mu(\omega^{k+1}) = e.$$

Otherwise, $\omega^{k+2} = \sum_{i < \omega} (\omega^{k+1} + \omega^k)$ implies by additivity of λ that

$$\begin{aligned} \mu(k+2) &= \mu\left(\sum_{i < \omega} (\omega^{k+1} + \omega^k)\right) \\ &= \sum_{i < \omega} [\mu(\omega^{k+1}) + \mu(\omega^k)] \\ &= \sum_{i < \omega} [\mu(\omega^{k+1}) + \mu(\omega^l)] \\ &= \sum_{i < \omega} \mu(\omega^{k+1} + \omega^l) \\ &= \sum_{i < \omega} \mu(\omega^l) \\ &= \sum_{i < \omega} \mu(\omega^k) = \mu\left(\sum_{i < \omega} \omega^k\right) = \mu(\omega^{k+1}). \end{aligned}$$

By induction on $j > 0$, it therefore follows that

$$\begin{aligned} \mu(k+j+1) &= \mu\left(\sum_{i < \omega} (\omega^{k+j+1} + \omega^{k+j})\right) \\ &= \sum_{i < \omega} [\mu(\omega^{k+j+1}) + \mu(\omega^{k+j})] \\ &= \sum_{i < \omega} [\mu(\omega^{k+j}) + \mu(\omega^{k+j})] \\ &= \sum_{i < \omega} \mu(\omega^{k+j}) \\ &= \mu\left(\sum_{i < \omega} \omega^{k+j}\right) = \mu(\omega^{k+j+1}) = \mu(\omega^{k+1}). \end{aligned}$$

Thus, we can set $e := \mu(\omega^{k+1})$. To see that it is idempotent, note that

$$e + e = \mu(\omega^k) + \mu(\omega^{k+1}) = \mu(\omega^k + \omega^{k+1}) = \mu(\omega^{k+1}) = e. \quad \square$$

3 Interpretations

Let us collect a few results about interpretations between linear orders and interpretations of linear orders in trees. We start with the following obser-

vation about how to encode a convex equivalence relation by a single unary predicate.

Lemma 3.1. *Let $\mathfrak{A} = \langle A, \leq \rangle$ be a linear order and \sim an equivalence relation on A where every \sim -class is convex. There exists a set $P \subseteq A$ such that*

$$a \sim b \quad \text{iff} \quad a \in P \Leftrightarrow b \in P \quad \text{for all } a \leq b.$$

Proof. We consider the equivalence relation \approx on A/\sim defined by

$$C \approx D \quad \text{iff} \quad \text{there are only finitely many } \sim\text{-classes between } C \text{ and } D.$$

We use this relation to define a colouring $\lambda : A/\sim \rightarrow [2]$ such that the set

$$P := \bigcup \{ C \in A/\sim \mid \lambda(C) = 1 \}$$

has the desired properties. For every \approx -class E of A/\sim , we fix some representative $C_E \in E$, and we set, for $D \in E$,

$$\lambda(D) := n \bmod 2, \quad \text{where } n \text{ is the number of } \sim\text{-classes between } C_E \text{ and } D \text{ (including } C_E \text{ and } D).$$

Then it follows that, for all $C < D$ in A/\sim , there is some $C < E \leq D$ with

$$\lambda(C) \neq \lambda(E). \quad \square$$

Next, we consider interpretations of linear orders in the binary tree.

Lemma 3.2. *There exists an FO-interpretation τ mapping the infinite binary tree $\mathfrak{T}_{\text{bin}} = \langle \{0, 1\}^*, \leq_{\text{pf}}, \text{suc}_0, \text{suc}_1 \rangle$ to the order of the rationals $\langle \mathbb{Q}, \leq \rangle$.*

Proof. We set $\tau := \langle \delta, \varphi_{\leq} \rangle$ with

$$\begin{aligned} \delta(x) &:= \exists y [\text{suc}_1(y, x)], \\ \varphi(x, y) &:= x \leq_{\text{pf}} y \vee \exists z \exists u \exists v [\text{suc}_0(z, u) \wedge \text{suc}_1(z, v) \\ &\quad \wedge u \leq_{\text{pf}} x \wedge v \leq_{\text{pf}} y]. \quad \square \end{aligned}$$

Corollary 3.3. $\langle \mathbb{Q}, \leq \rangle$ has a decidable MSO-theory.

Corollary 3.4. There exists an FO-interpretation τ with the following property. For every countable linear order \mathcal{A} , there exists a set $P \subseteq \{0, 1\}^*$ such that

$$\mathcal{A} \cong \tau(\langle \mathfrak{T}_{\text{bin}}, P \rangle).$$

Proof. Let $\sigma = \langle \delta, \varphi_{\leq} \rangle$ be the interpretation from Lemma 3.2. We have seen in Lemma 1.3 that every countable linear order can be embedded into \mathbb{Q} . Consequently, we can obtain the desired interpretation τ by replacing the formula δ by

$$\delta'(x) := \delta(x) \wedge Px. \quad \square$$

Corollary 3.5. Given an MSO-formula φ , it is decidable whether or not there exists a countable C -coloured linear order \mathcal{A} with $\mathcal{A} \models \varphi$.

Proof. Let τ be the FO-interpretation from Corollary 3.4. Modifying τ by using additional unary predicates $\tilde{Q} = (Q_c)_{c \in C}$ for the colours, we can interpret every countable C -coloured linear order \mathcal{A} in $\langle \mathfrak{T}_{\text{bin}}, P, \tilde{Q} \rangle$. Consequently, φ is satisfied by a countable C -countable linear order if, and only if,

$$\langle \mathfrak{T}_{\text{bin}}, P, \tilde{Q} \rangle \models \varphi^\tau,$$

or, equivalently, if

$$\mathfrak{T}_{\text{bin}} \models \exists P \exists \tilde{Q} \varphi^\tau.$$

This property is decidable by Corollary V.3.9. \square

There exists a normal form for interpretations of a linear order in the infinite binary tree, which is sometimes useful.

Lemma 3.6. Let $\mathfrak{T}_{\text{bin}}$ be the infinite binary tree. For each MSO-interpretation τ such that $\tau(\mathfrak{T}_{\text{bin}})$ is a coloured linear order, there exists an MSO-interpretation σ such that

- ♦ $\sigma(\mathfrak{T}_{\text{bin}}) \cong \tau(\mathfrak{T}_{\text{bin}})$,
- ♦ the universe of $\sigma(\mathfrak{T}_{\text{bin}})$ forms an antichain with respect to the tree-order \leq_{pf} of $\mathfrak{T}_{\text{bin}}$, and
- ♦ the order of $\sigma(\mathfrak{T}_{\text{bin}})$ is included in the lexicographic order \leq_{lex} on $\mathfrak{T}_{\text{bin}}$.

Proof. Suppose that $\tau = \langle \delta, \varphi_{\leq} \rangle$, let m be the quantifier-rank of φ_{\leq} , and let Θ be the set of all MSO $_m$ -theories (with one first-order parameter). We use the function

$$h : \{0, 1\}^* \rightarrow (\{0, 1\} \times \Theta + \{\$\})^*$$

mapping a word $w = c_0 \dots c_{n-1}$ to the word

$$\langle c_0, \theta_0 \rangle \dots \langle c_{n-1}, \theta_{n-1} \rangle \$ \quad \text{where} \quad \theta_i := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{c_0 \dots c_{i-1}}, w).$$

Let A be the set defined by δ and let B be its image under h . Then B forms an antichain with respect to \leq_{pf} . If we could find an ordering on the alphabet $\Sigma := \{0, 1\} \times \Theta + \{\$\}$ such that \leq becomes a subset of the corresponding lexicographic ordering, we would be done. Unfortunately, this usually does not work. Instead, we have to choose a different ordering on Σ at each vertex w of the tree.

First note that, given vertices $w, u, u' \in \{0, 1\}^*$ and directions $k, k' \in \{0, 1\}$, the theories

$$\theta := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk}, u) \quad \text{and} \quad \theta' := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk'}, u')$$

determine whether or not $\mathfrak{T}_{\text{bin}} \models \varphi_{\leq}(wku, wk'u')$. We define an ordering \leq_w on Θ by setting $\theta <_w \theta'$ if $\theta \neq \theta'$ and

$$\text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk}, u) = \theta \quad \text{and} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk'}, u') = \theta'$$

implies

$$\mathfrak{T}_{\text{bin}} \models \varphi_{\leq}(wku, wk'u').$$

We define a successor-ordering \leq_{so} on B as follows. For a vertex $w \in \Sigma^*$ and directions $\langle k, \theta \rangle, \langle k', \theta' \rangle \in \Sigma$, we set

$$w\langle k, \theta \rangle \leq_{\text{so}} w\langle k', \theta' \rangle \quad \text{iff} \quad k < k' \quad \text{or} \quad k = k' \text{ and } \theta \leq_w \theta'.$$

It remains to determine the relative order of the vertex $w\$$. Let u be the vertex of $\mathfrak{T}_{\text{bin}}$ with $h(u) = w\$$. We set

$$w\$ \leq_{\text{so}} w\langle k, \theta \rangle \quad \text{iff} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk}, v) = \theta \text{ implies } u \leq wkv.$$

Using a binary encoding of Σ in $\{0, 1\}^m$, for some m , (which depends on the vertex w) we obtain a function

$$g : \{0, 1\}^* \rightarrow \{0, 1\}^*$$

such that the image $C := g[A]$ forms an antichain and is MSO-definable. Furthermore, the image $g[P]$ of each colour predicate P is also MSO-definable.

Let \leq_{lex} be the lexicographic ordering on C induced by this successor-ordering \leq_{so} . To conclude the proof, it is sufficient to show that

$$g(x) \leq_{\text{lex}} g(y) \quad \text{implies} \quad \mathfrak{T} \models \varphi_{\leq}(x, y).$$

Hence, let $x, y \in A$ with $x \neq y$. We distinguish three cases. If $x <_{\text{pf}} y$, we have

$$h(x) = w\$ \quad \text{and} \quad h(y) = w\langle k, \theta \rangle u, \quad \text{for some } w, u, k, \theta.$$

By definition of the ordering on the successors of w , it follows that $h(x) <_{\text{so}} w\langle k, \theta \rangle$. By definition of $<_{\text{so}}$, this implies that $\mathfrak{T}_{\text{bin}} \models \varphi_{\leq}(x, y)$.

If $y <_{\text{pf}} x$, the argument is analogous. Hence, it remains to consider the case where x and y are incomparable. Set $w := x \sqcap y$, let u, u', k, k' be such that $x = wku$ and $y = wku'$, and set

$$\theta := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk}, u) \quad \text{and} \quad \theta' := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{wk'}, u').$$

By definition of \leq_{so} it follows that $\theta \leq_w \theta'$. Consequently,

$$\mathfrak{T}_{\text{bin}} \models \varphi_{\leq}(wku, wku'). \quad \square$$

Finally, let us consider structures interpretable in a linear order. The following result is a variant of Theorem V.1.13 for linear orders.

Proposition 3.7. *Every structure MSO-interpretable in a coloured linear order is also FO-interpretable in some coloured linear order.*

Proof. Let τ be an MSO-interpretation, \mathfrak{A} a coloured linear order, and $\mathfrak{B} := \tau(\mathfrak{A})$ its image. Let m be the quantifier-rank of τ and let λ be the additive labelling on \mathfrak{A} defined by

$$\lambda(a, b) := \text{Th}_{\text{MSO}}^m(\mathfrak{A}|_{[a, b)}, a), \quad \text{for } a < b.$$

By Theorem II.3.7, there exists a Ramseyan split σ for λ . Analogous to the proof of Lemma V.1.12, one can show that λ is FO-definable using suitable monadic parameters. Since the MSO_m -theory of a tuple \bar{a} is determined by its order type and the labels $\lambda(a_i, a_j)$ (both of which are FO-definable), it therefore follows that we can translate every MSO-formula from τ into an FO-formula. \square

4 Regular Linear Orders

We have shown in Theorem V.2.17 that every MSO-definable set of trees is uniquely determined by the regular trees it contains. This result can be generalised to definable sets of countable linear orders. We start by defining what we mean by a regular linear order.

Definition 4.1. (a) Let $\mathfrak{A}_0, \dots, \mathfrak{A}_{n-1}$ be coloured linear orders. The *shuffle* of $\mathfrak{A}_0, \dots, \mathfrak{A}_{n-1}$ is the ordered sum

$$\sum_{q \in \mathbb{Q}} \mathfrak{A}_{\chi(q)},$$

where $\chi : \mathbb{Q} \rightarrow [n]$ is a dense colouring of \mathbb{Q} .

(b) Let C be a fixed finite set of colours. A *Läuchli-Leonard operation* over C is an operation on C -coloured linear orders that takes one of the following forms:

- ♦ constants \circ and c , for $c \in C$,
- ♦ a binary operation $+$,
- ♦ two unary operations $- \times \omega$ and $- \times \omega^{\text{op}}$, and
- ♦ an n -ary operation $- \sqcup \cdots \sqcup -$, for every $\circ < n < \omega$.

The semantics of these operations is as follows. \circ is the empty order; $c \in C$ denotes the singleton order whose element is coloured c ; $\mathcal{A} + \mathcal{B}$ is the ordered sum of \mathcal{A} and \mathcal{B} ; $\mathcal{A} \times \omega$ and $\mathcal{A} \times \omega^{\text{op}}$ denote the corresponding ordered products; and \sqcup is the shuffle of the given orders.

(c) A C -coloured linear order \mathcal{A} is *regular* if it is the value of a finite term of L\"auchli-Leonard operations. \downarrow

Examples. (a) The orders $\langle \mathbb{Z}, \leq \rangle$ and $\langle \mathbb{Q}, \leq \rangle$ are regular
 (b) Every ordinal of the form

$$\omega^{k_0} n_0 + \cdots + \omega^{k_{m-1}} n_{m-1} \quad \text{with} \quad k_0, \dots, k_{m-1}, n_0, \dots, n_{m-1} < \omega$$

is regular. \downarrow

Exercise 4.1. Let \mathcal{A} be a regular linear order. Prove that $\text{gFC}(\mathcal{A}) < \omega$. \downarrow

Below we will prove that every countable linear order has the same MSO_m -theory as a regular one. The proof makes use of a decomposition theorem for linear orders. We start with a general version that also holds for uncountable orders before presenting a special version for the countable case.

Definition 4.2. Let \mathcal{L} be a class of coloured linear orders that is closed under convex suborders, S a set, and $\lambda : \mathcal{L} \rightarrow S$ a function.

(a) The function λ is \mathcal{L} -*additive* if, given two families $(\mathcal{A}_i)_{i \in I}$ and $(\mathcal{B}_i)_{i \in I}$ of orders in \mathcal{L} with $\sum_{i \in I} \mathcal{A}_i \in \mathcal{L}$ and $\sum_{i \in I} \mathcal{B}_i \in \mathcal{L}$, we have

$$\lambda(\mathcal{A}_i) = \lambda(\mathcal{B}_i), \text{ for all } i \in I, \quad \text{implies} \quad \lambda\left(\sum_i \mathcal{A}_i\right) = \lambda\left(\sum_i \mathcal{B}_i\right).$$

(b) The *bicofinality* of a linear order \mathcal{A} is the pair $\langle \alpha, \beta \rangle$ where α is the least ordinal such that A contains a strictly decreasing sequence of length α that is not bounded from below in A , and β is the least ordinal such that

A contains a strictly increasing sequence of length β that is not bounded from above in A . (If A has a minimal element, we set $\alpha := 1$; if it has a maximal element, we set $\beta := 1$.)

(c) A linear order \mathcal{A} is λ -*homogeneous* if

$$\lambda(\mathcal{A}|_C) = \lambda(\mathcal{A}|_D), \quad \text{for all convex } C, D \subseteq A \text{ without end-points} \\ \text{that have the same bicofinality.} \quad \downarrow$$

Before proving the decomposition theorem, let us show how to find λ -homogeneous subsets.

Lemma 4.3. Let $\mathcal{A} = \langle A, \leq, \bar{P} \rangle$ be a dense linear order and λ -an \mathcal{L} -additive labelling where \mathcal{L} is the class of all convex suborders of \mathcal{A} . There exists a non-empty open interval $I \subseteq A$ such that $\mathcal{A}|_I$ is λ -homogeneous.

Proof. Let λ' be the additive labelling of A defined by

$$\lambda'(a, b) := \lambda(\mathcal{A}|_{[a, b)}), \quad \text{for } a < b.$$

As \mathcal{A} is densely ordered, we can apply Theorem 2.2 (b) to λ' to find a non-empty open interval I , a subset $H \subseteq I$, and a colour θ such that H is dense in I and, for all pairs $a < b$ in H , the interval $[a, b)$ has colour θ . We claim that I is the desired λ -homogeneous interval.

For the proof, let $C \subseteq I$ be non-empty, convex, and without end-points, and let $\langle \alpha, \beta \rangle$ be its bicofinality. It follows that

$$\lambda(\mathcal{A}|_C) = (\alpha^{\text{op}} + \beta) \times \theta,$$

which does not depend on C . Consequently, we have

$$\mathcal{A}|_C \equiv_{\text{MSO}}^{\bar{P}} \mathcal{A}|_D, \quad \text{for all non-empty convex sets } C, D \subseteq I \text{ without} \\ \text{end-points with bicofinality } \langle \alpha, \beta \rangle. \quad \square$$

The general decomposition theorem for linear orders can be seen as a variant of Simon's Lemma.

Theorem 4.4. Let C, S, T be finite sets of colours, \mathcal{L} a class of C -coloured linear orders that is closed under convex suborders, the corresponding class of uncoloured linear orders, and let \mathcal{L}_d be the class of all S -coloured dense linear orders whose underlying order belongs to \mathcal{L}_o . Let $\lambda : \mathcal{L} \rightarrow S$ be an \mathcal{L} -additive labelling and $\mu : \mathcal{L}_d \rightarrow T$ an additive one.

Let \mathcal{C} be the smallest class satisfying the following conditions.

- (I) $\mathfrak{A} \in \mathcal{C}$, for all orders \mathfrak{A} with $|A| = 1$.
- (II) $\mathfrak{A}, \mathfrak{B} \in \mathcal{C}$ and $\mathfrak{A} + \mathfrak{B} \in \mathcal{L}$ implies $\mathfrak{A} + \mathfrak{B} \in \mathcal{C}$.
- (III) If α is an ordinal and $(\mathfrak{A}_i)_{i < \alpha}$ a sequence in \mathcal{C} such that

$$\sum_{i \in \alpha} \mathfrak{A}_i \in \mathcal{L} \quad \text{and} \quad \lambda(\mathfrak{A}_i) = \lambda(\mathfrak{A}_j), \quad \text{for all } i, j,$$

then $\sum_{i \in \alpha} \mathfrak{A}_i \in \mathcal{C}$.

- (IV) If α is an ordinal and $(\mathfrak{A}_i)_{i < \alpha}$ a sequence in \mathcal{C} such that

$$\sum_{i \in \alpha^{\text{op}}} \mathfrak{A}_i \in \mathcal{L} \quad \text{and} \quad \lambda(\mathfrak{A}_i) = \lambda(\mathfrak{A}_j), \quad \text{for all } i, j,$$

then $\sum_{i \in \alpha^{\text{op}}} \mathfrak{A}_i \in \mathcal{C}$.

- (v) If $\mathfrak{A} \in \mathcal{L}$ and \approx is a congruence of \mathfrak{A} such that $\mathfrak{A}|_{[a]} \in \mathcal{C}$, for all $[a] \in A/\approx$, and the order $\langle A/\approx, \leq, (P_c)_{c \in S} \rangle$ with

$$P_c := \{ [a] \in A/\approx \mid \lambda([a]) = c \}, \quad \text{for } c \in S,$$

is μ -homogeneous and densely ordered without end-points, then $\mathfrak{A} \in \mathcal{C}$.

Then $\mathcal{C} = \mathcal{L}$.

Proof. Fix $\mathfrak{A} \in \mathcal{L}$. We call a convex subset $D \subseteq A$ decomposable if

$$\mathfrak{A}|_C \in \mathcal{C}, \quad \text{for every non-empty convex } C \subseteq D.$$

We will show that A is decomposable. Consider the following relation on A .

$$x \sim y \quad : \text{iff} \quad x = y \text{ or the interval } (x, y] \text{ or } (y, x] \text{ is decomposable.}$$

This relation is clearly reflexive and symmetric. For transitivity, suppose that $x \sim y \sim z$ where without loss of generality $x < y < z$. Every convex subset J of the interval $[x, z]$ can be factorised as $J = D + E$ where $D \subseteq [x, y]$ and $E \subseteq [y, z]$. By assumption, $\mathfrak{A}|_D, \mathfrak{A}|_E \in \mathcal{C}$ (if they are non-empty). By (II), it follows that $\mathfrak{A}|_J = \mathfrak{A}|_D + \mathfrak{A}|_E \in \mathcal{C}$.

Thus \sim is an equivalence relation. Next, let us show that every \sim -class H is decomposable. If $H = \emptyset$, this is trivial. Otherwise, fix some index $x \in H$ and set

$$H_- := \{ y \in H \mid y \leq x \} \quad \text{and} \quad H_+ := \{ y \in H \mid y > x \}.$$

As we have already proved above that decomposable sets are closed under finite sums, it is sufficient to show that H_- and H_+ are decomposable.

If H_+ has a greatest element z , then $x \sim z$ implies that H_+ is decomposable. Otherwise, fix an increasing sequence $(z_i)_{i < \alpha}$ of elements $z_i \in H$ starting with $z_0 = x$ such that $(z_i)_i$ is not bounded in H . Then $z_i \sim z_j$ implies that every interval $J_{ij} := (z_i, z_j]$ with $i < j$ is decomposable. By Theorem 2.2, there exist a cofinal set $I \subseteq \alpha$ such that

$$\lambda(\mathfrak{A}|_{J_{ij}}) = \lambda(\mathfrak{A}|_{J_{i'j'}}), \quad \text{for all } i < j \text{ and } i' < j' \text{ in } I.$$

Let $K \subseteq H_+$ be convex and non-empty. If K is bounded in H_+ , we have $K \subseteq (x, z_i]$, for some $i \in I$, and $x \sim z_i$ implies that $\mathfrak{A}|_K \in \mathcal{C}$. Suppose otherwise. Then the set $I_o := \{ i \in I \mid z_i \in K \}$ is cofinal in I . Let $(y_i)_{i < \beta}$ be an increasing enumeration of $(z_i)_{i \in I_o}$, for some $\beta \leq \alpha$. It follows that we can write

$$K := C + \sum_{i < \beta} D_i, \quad \text{where} \quad C := K \cap (x, y_o] \text{ and } D_i := (y_i, y_{i+1}].$$

Since $x \sim y_o$ and $y_i \sim y_{i+1}$, it follows that

$$\mathfrak{A}|_C \in \mathcal{C} \quad \text{and} \quad \mathfrak{A}|_{D_i} \in \mathcal{C}.$$

Furthermore, $\lambda(\mathfrak{A}|_{D_i}) = \lambda(\mathfrak{A}|_{D_j})$, for all i, j . Consequently, we can use (II) and (III) to show that

$$\mathfrak{A}|_K \in \mathcal{C}.$$

The proof for H_- is analogous.

To conclude the proof, set $\mathfrak{B} := \langle A/\sim, \leq, (P_c)_{c \in S} \rangle$. If \mathfrak{B} is a singleton, the set A is decomposable and we are done. Hence, suppose otherwise.

Note that $\mathfrak{A}|_{[a]} \in \mathcal{C}$ since each \sim -class $[a]$ is composable. If we can show that \mathfrak{B} is densely ordered, we can use Lemma 4.3 to find an open interval $I \subseteq B$ that is μ -homogeneous. Consequently, it follows by (v) that the union of all \sim -classes in I is composable. This means that I consists of a single \sim -class. A contradiction.

To prove density, fix two \sim -classes $[a]_\sim < [b]_\sim$ in A/\sim . We have shown above that composable sets are closed under finite sums. Hence, if there is no class $[a]_\sim < [c]_\sim < [b]_\sim$, the sum $[a]_\sim + [b]_\sim$ would be composable. In particular, $a \sim b$. A contradiction. \square

Corollary 4.5. *Let C, S, T be finite sets of colours, \mathcal{L} a class of C -coloured linear orders that is closed under convex suborders, \mathcal{U} be the corresponding class of unlabelled linear orders, and let \mathcal{L}_d be the class of all S -coloured dense linear orders whose underlying order belongs to \mathcal{L}_o . Let $\lambda : \mathcal{L} \rightarrow S$ be an \mathcal{L} -additive labelling and $\mu : \mathcal{L}_d \rightarrow T$ an additive one.*

Let $U \subseteq S$ be a set satisfying the following conditions.

- (I) $\lambda(\mathfrak{A}) \in U$, for all $\mathfrak{A} \in \mathcal{L}$ with $|A| = 1$.
- (II) $\lambda(\mathfrak{A}), \lambda(\mathfrak{B}) \in U$ and $\mathfrak{A} + \mathfrak{B} \in \mathcal{L}$ implies $\lambda(\mathfrak{A} + \mathfrak{B}) \in U$.
- (III) If α is an ordinal and $(\mathfrak{A}_i)_{i < \alpha}$ a sequence such that

$$\sum_{i \in \alpha} \mathfrak{A}_i \in \mathcal{L} \quad \text{and} \quad \lambda(\mathfrak{A}_i) = \lambda(\mathfrak{A}_j) \in U, \quad \text{for all } i, j,$$

then $\lambda(\sum_{i \in \alpha} \mathfrak{A}_i) \in U$.

- (IV) If α is an ordinal and $(\mathfrak{A}_i)_{i < \alpha}$ a sequence such that

$$\sum_{i \in \alpha^{op}} \mathfrak{A}_i \in \mathcal{L} \quad \text{and} \quad \lambda(\mathfrak{A}_i) = \lambda(\mathfrak{A}_j) \in U, \quad \text{for all } i, j,$$

then $\lambda(\sum_{i \in \alpha^{op}} \mathfrak{A}_i) \in U$.

- (v) If $\mathfrak{A} \in \mathcal{L}$ and \approx is a congruence of \mathfrak{A} such that the order $\langle A/\approx, \leq, (P_c)_{c \in S} \rangle$ with

$$P_c := \{ [a] \in A/\approx \mid \lambda([a]) = c \}, \quad \text{for } c \in S,$$

is μ -homogeneous and densely ordered without end-points, and if $P_c \neq \emptyset$ implies $c \in U$, then $\lambda(\mathfrak{A}) \in U$.

Then $\lambda(\mathfrak{A}) \in U$, for every $\mathfrak{A} \in \mathcal{L}$.

Proof. Setting $\mathcal{C} := \lambda^{-1}[U] \cap \mathcal{L}$, it follows by Theorem 4.4 that $\mathcal{C} = \mathcal{L}$. Hence, $\lambda[\mathcal{L}] \subseteq U$. \square

The most common case is the one where \mathcal{L} is the class of all countable linear orders and λ maps a linear ordering to its MSO_m -theory. In this case, Theorem 4.4 simplifies to the following statement.

Corollary 4.6. *Let C be a finite set of colours and $m < \omega$ a constant. The class of all countable C -coloured linear orders is the least class \mathcal{L} satisfying the following conditions.*

- Every order with at most one element belongs to \mathcal{L} .
- $\mathfrak{A}, \mathfrak{B} \in \mathcal{L}$ implies $\mathfrak{A} + \mathfrak{B} \in \mathcal{L}$.
- If $\mathfrak{A}_0, \mathfrak{A}_1, \dots \in \mathcal{L}$ are orders with $\text{Th}_{\text{MSO}}^m(\mathfrak{A}_0) = \text{Th}_{\text{MSO}}^m(\mathfrak{A}_1) = \dots$, then $\sum_{i \in \omega} \mathfrak{A}_i \in \mathcal{L}$ and $\sum_{i \in \omega^{op}} \mathfrak{A}_i \in \mathcal{L}$.
- If $(\mathfrak{A}_i)_{i \in I}$ is a family of orders from \mathcal{L} such that the colouring of I mapping each $i \in I$ to $\text{Th}_{\text{MSO}}^m(\mathfrak{A}_i)$ is dense, then $\sum_{i \in I} \mathfrak{A}_i \in \mathcal{L}$.

Theorem 4.7. *Let C be a finite set of colours and $m < \omega$ a constant. For every countable C -coloured linear order \mathfrak{A} , there exists a regular linear order \mathfrak{B} with $\mathfrak{B} \equiv_{\text{MSO}}^m \mathfrak{A}$.*

Proof. We have shown that we can construct \mathfrak{A} in finitely many steps using the operations of Corollary 4.6. We prove the claim by induction on the number of steps.

If \mathfrak{A} has at most one element, it is regular. If $\mathfrak{A} = \mathfrak{B} + \mathfrak{B}'$, we can use the inductive hypothesis to find regular orders \mathfrak{C} and \mathfrak{C}' with

$$\mathfrak{C} \equiv_{\text{MSO}}^m \mathfrak{B} \quad \text{and} \quad \mathfrak{C}' \equiv_{\text{MSO}}^m \mathfrak{B}'.$$

Hence, $\mathfrak{A} \equiv_{\text{MSO}}^m \mathfrak{C} + \mathfrak{C}'$.

If $\mathfrak{A} = \sum_{i \in \omega} \mathfrak{B}_i$, we obtain by inductive hypothesis regular orders

$$\mathfrak{C}_i \equiv_{\text{MSO}}^m \mathfrak{B}_i, \quad \text{for } i < \omega.$$

Consequently, $\mathfrak{A} \equiv_{\text{MSO}}^m \sum_{i \in \omega} \mathfrak{C}_i$.

The proof for the case where $\mathfrak{A} = \sum_{i \in \omega^{\text{op}}} \mathfrak{B}_i$ is analogous.

Finally, suppose that $\mathfrak{A} = \sum_{i \in I} \mathfrak{B}_i$ and that the colouring

$$\mu(i) := \text{Th}_{\text{MSO}}^m(\mathfrak{B}_i)$$

is dense. For every $\theta \in \text{rng } \mu$, we can use the inductive hypothesis to find a regular order \mathfrak{C}_θ with

$$\mathfrak{C}_\theta \equiv_{\text{MSO}}^m \mathfrak{B}_i, \quad \text{for } i \in \mu^{-1}(\theta).$$

Consequently, it follows by the composition theorem for ordered sums that $\mathfrak{A} \equiv_{\text{MSO}}^m \mathfrak{C}_{\theta_0} \sqcup \dots \sqcup \mathfrak{C}_{\theta_{n-1}}$, where $\theta_0, \dots, \theta_{n-1}$ is an enumeration of $\text{rng } \mu$. \square

One of the reasons why Leonard and L\"auchli chose that particular set of operations is that those are what is needed to solve systems of equations.

Definition 4.8. Let C be a finite set of colours, $n < \omega$ a number, and

$$E = \{x_0 = t_0, \dots, x_{n-1} = t_{n-1}\}$$

a finite set of n equations where t_0, \dots, t_{n-1} are terms using binary ordered sums, the variables x_0, \dots, x_{n-1} , and constants for each colour.

A tuple $\langle \mathfrak{A}_0, \dots, \mathfrak{A}_{n-1} \rangle$ is a *solution* of E if it satisfies each equation in E when we substitute \mathfrak{A}_i for the variable x_i and we interpret the constants $c \in C$ as 1-element orders whose single element has colour c . By abuse of terminology, we also say that \mathfrak{A}_i is a solution of E omitting the other components.

The tuple is the *least* solution of E if, furthermore, the orders $\mathfrak{A}_0, \dots, \mathfrak{A}_{n-1}$ can be embedded into every other solution of E . \downarrow

Exercise 4.2. Prove that every finite system of equations has a least solution. \downarrow

Lemma 4.9. Let E be a finite system of equations. The least solution of E is regular.

Proof. Let $E = \{x_0 = t_0, \dots, x_{n-1} = t_{n-1}\}$ be a system of equations. By introducing new variables if necessary, we can transform this system such that every term t_i is of the form $x_k + x_l$ or c , for variables x_k, x_l and a constant c .

We construct a graph with vertices x_0, \dots, x_{n-1} as follows. For each equation $x_i = x_k + x_l$, we add the edges $x_i \rightarrow x_k$ and $x_i \rightarrow x_l$.

We prove the claim by induction on the number of strongly connected components of this graph. Let $X := \{x_0, \dots, x_{m-1}\}$ be one such component and let Y be the set of all variables reachable from the component X . For each variable $y \in Y$, we can use the inductive hypothesis to construct a Leonard-L\"auchli term denoting the value of y . To construct terms for the variables in X , we distinguish two cases.

(I) First, suppose that all equations are of the form $x = y + x'$ or $x = x' + y$ with $x' \in X$ and $y \in Y$. By repeatedly replacing the variables $x' \in X$ by their definitions, we finally obtain equations of the form

$$x = w + x + w'$$

where w and w' are sums (possibly empty) of variables in Y . This equation can be solved by the term

$$x = w \times \omega + w' \times \omega^{\text{op}}.$$

(II) Otherwise there is at least one equations of the form $x_i = x_k + x_l$. Again we eliminate all equations with only one variable from X on the right-hand side by replacing this variable by its definition. Then all equations are of the form

$$x_i = w + x_k + w' + x_l + w''$$

where $x_i, x_k, x_l \in X$ and w, w', w'' are sums of variables in Y . Introducing a new variable y , we can rewrite this equation as

$$x_i = w + x_k + y \quad y = w' + x_l + w''.$$

Replacing x_l in the latter equation by its definition we obtain a system of equations of the form

$$x_i = w + x_k + x_l + w',$$

where $x_i, x_k, x_l \in X$ and w, w' are sums of variables in Y . For each equation of this form, we define orders λ_i, μ_i, ρ_i by

$$\lambda_i = w + \lambda_k,$$

$$\rho_i = \rho_l + w',$$

and $x_i = \lambda_i + \mu_i + \rho_i$.

The equations for λ_i and ρ_i can be solved as above by terms of the form

$$\lambda_i = (w_o + \dots + w_r) \times \omega \quad \text{and} \quad \rho_i = (w'_o + \dots + w'_s) \times \omega^{\text{op}}.$$

The orders μ_i can equivalently be defined by

$$\mu_i = \mu_k + \rho_k + \lambda_l + \mu_l.$$

Let ξ_i be the ordering obtained from μ_i by replacing the sum $\rho_k + \lambda_l$ in each such equation by a constant c_{kl} . Then μ_i can be obtained from ξ_i by substituting the c_{kl} by $\rho_k + \lambda_l$. Since the ξ_i are dense orders without end-points each μ_i denotes the shuffle of the orders $\rho_k + \lambda_l$. \square

Regular linear orders can be characterised in many different ways. (We have proved an analogous characterisation for regular trees in Theorem ??.)

Theorem 4.10. *Let \mathfrak{A} be a linear order. The following statements are equivalent.*

- (1) \mathfrak{A} is regular.
- (2) \mathfrak{A} is the least solution of a finite system of equations.

(3) $\mathfrak{A} \cong \tau(\mathfrak{T}_{\text{bin}})$, for some MSO-interpretation τ and the infinite binary tree $\mathfrak{T}_{\text{bin}}$.

(4) $\mathfrak{A} \cong \langle K, \leq_{\text{lex}}, \bar{P} \rangle$, for some regular language $K \subseteq \{o, i\}^*$ and regular sets $P_c \subseteq K$.

(5) \mathfrak{A} is the unique countable model of a suitable MSO-formula.

Proof. (2) \Rightarrow (1) has already been proved in Lemma 4.9.

(1) \Rightarrow (3) It is sufficient to show that the class of all linear orders interpretable in the infinite binary tree is closed under all L\"auchli-Leonard operations. For a formula $\varphi(\bar{x})$, a vertex z , and a direction $k \in \{o, i\}$, we denote by $\varphi^{(z,k)}(\bar{x})$ the relativisation of φ to the subtree attached to the k -successor of z . Clearly, every order with at least one element can be interpreted in $\mathfrak{T}_{\text{bin}}$. Let $\tau_i = \langle \delta^i, \varphi_{\leq}^i, \varphi_{P_j}^i \rangle$ be an MSO-interpretation of some linear order \mathfrak{A}_i in $\mathfrak{T}_{\text{bin}}$. Then the formulae

$$\delta(x) := (\delta^o)^{(\langle \rangle, o)}(x) \vee (\delta^i)^{(\langle \rangle, i)}(x)$$

$$\varphi_{\leq}(x, y) := (\varphi_{\leq}^o)^{(\langle \rangle, o)}(x, y) \vee (\varphi_{\leq}^i)^{(\langle \rangle, i)}(x, y) \\ \vee [(\delta^o)^{(\langle \rangle, o)}(x) \wedge (\delta^i)^{(\langle \rangle, i)}(y)]$$

$$\varphi_{P_j}(x) := (\varphi_{P_j}^o)^{(\langle \rangle, o)}(x) \vee (\varphi_{P_j}^i)^{(\langle \rangle, i)}(x)$$

provide an interpretation of $\mathfrak{A}_o + \mathfrak{A}_i$. Similarly, if $\psi(x)$ is a formula stating that $x = o^n$, for some n , we obtain interpretations

$$\delta(x) := \exists z [\psi(z) \wedge (\delta^o)^{(z, i)}(x)]$$

$$\varphi_{\leq}(x, y) := \exists z [\psi(z) \wedge (\varphi_{\leq}^o)^{(z, i)}(x, y)]$$

$$\vee \exists u \exists v [\psi(v) \wedge u \leq v \wedge (\delta^o)^{(u, i)}(x) \wedge (\delta^o)^{(v, i)}(y)]$$

$$\varphi_{P_j}(x) := \exists z [\psi(z) \wedge (\varphi_{P_j}^o)^{(z, i)}(x)]$$

and

$$\begin{aligned}\delta(x) &:= \exists z [\psi(z) \wedge (\delta^\circ)^{(z,1)}(x)] \\ \varphi_\leq(x, y) &:= \exists z [\psi(z) \wedge (\varphi_\leq^\circ)^{(z,1)}(x, y)] \\ &\quad \vee \exists u \exists v [\psi(v) \wedge u \leq v \wedge (\delta^\circ)^{(u,1)}(y) \wedge (\delta^\circ)^{(v,1)}(x)] \\ \varphi_{P_j}(x) &:= \exists z [\psi(z) \wedge (\varphi_{P_j}^\circ)^{(z,1)}(x)]\end{aligned}$$

of $\mathfrak{A}_0 \times \omega$ and $\mathfrak{A}_0 \times \omega^{\text{op}}$. Finally, for $\mathfrak{A}_0 \sqcup \dots \sqcup \mathfrak{A}_{n-1}$, we obtain the desired interpretation as follows. Let $\psi_i(w, x)$ be the formula stating that $w \in (10^+)^*$ and $x = w10^{i+1}1$. Then we can use the formulae

$$\begin{aligned}\delta(x) &:= \exists w \exists z \bigvee_{i < n} [\psi_i(w, z) \wedge (\delta^i)^{(z,1)}(x)] \\ \varphi_\leq(x, y) &:= \exists w \exists z \bigvee_{i < n} [\psi_i(w, z) \wedge (\varphi_\leq^i)^{(z,1)}(x, y)] \\ &\quad \vee \exists w \exists w' \exists z \exists z' \bigvee_{i, j < n} [\psi_i(w, z) \wedge \psi_j(w', z') \wedge z \leq_{\text{lex}} z' \\ &\quad \wedge (\delta^i)^{(z,1)}(x) \wedge (\delta^j)^{(z',1)}(y)] \\ \varphi_{P_j}(x) &:= \exists w \exists z \bigvee_{i < n} [\psi_i(w, z) \wedge (\varphi_{P_j}^i)^{(z,1)}(x)].\end{aligned}$$

(To see that this works, note that the order $\langle Q, \leq_{\text{lex}} \rangle$ is dense where Q is the set of vertices satisfying the formula $\exists w \bigvee_{i < n} \psi_i(w, x)$.)

(3) \Rightarrow (4) Let $\sigma = \langle \delta, \varphi_\leq, (\varphi_{P_i})_i \rangle$ be the interpretation from Lemma 3.6. The set $K \subseteq \{0, 1\}^*$ defined by δ , and $Q_i \subseteq \{0, 1\}^*$ the one defined by φ_{P_i} . Then K and Q_i are regular and the order $\langle K, \leq_{\text{lex}}, \tilde{Q} \rangle$ is isomorphic to \mathfrak{A} .

(4) \Rightarrow (2) We can combine the automata recognising the universe K and the predicates P_i into a single deterministic automaton

$$\mathcal{A} = \langle Q, \{0, 1\}, \delta, q_0, F, \tilde{G} \rangle$$

with several sets of final states F, \tilde{G} such that, when using F , it recognises K and, when using \tilde{G}_i , it recognises P_i . (Note that $F = \bigcup_i \tilde{G}_i$.) Let A_q be the language recognised by \mathcal{A} when starting in state q . The orders $\mathfrak{A}_q := \mathfrak{A}|_{A_q}$,

$q \in Q$, satisfy the equations

$$\begin{aligned}\mathfrak{A}_q &= c + \mathfrak{A}_{\delta(q,0)} + \mathfrak{A}_{\delta(q,1)}, & \text{if } q \in F \cap G_c, \\ \mathfrak{A}_q &= \mathfrak{A}_{\delta(q,0)} + \mathfrak{A}_{\delta(q,1)}, & \text{if } q \notin F.\end{aligned}$$

Let E be the system of equations obtained from these equations by replacing \mathfrak{A}_q by a variable x_q , for $q \in Q$. Then $\mathfrak{A} = \mathfrak{A}_{q_0}$ is the least solution of E (for the variable x_{q_0}).

(5) \Rightarrow (2) Fix an MSO-formula φ defining \mathfrak{A} . By Theorem 4.7, φ has a model that is regular. Since \mathfrak{A} is the only countable model of φ , it follows that this model is \mathfrak{A} .

(1) \Rightarrow (5) Let t be a Leonard-Läuchli term denoting \mathfrak{A} . We prove the claim by induction on t . If t is a single constant, \mathfrak{A} is trivially definable. If $t = s + s'$, let φ and φ' be the formulae defining the values of s and s' , respectively. Then we obtain a formula ψ defining \mathfrak{A} by stating that there exists a downwards-closed set Z such that the suborder induced by Z satisfies φ and the complement of Z satisfies φ' .

Suppose that $t = s \times \omega$ and let φ be the formula defining s . We obtain a formula defining t by stating that there exists a set Z such that

- ♦ Z and its complement are both unbounded and
- ♦ every maximal convex set contained in either Z or its complement satisfies φ .

The construction for $s \times \omega^{\text{op}}$ is analogous.

Finally, suppose that $t = s_0 \sqcup \dots \sqcup s_{n-1}$ and let φ_i be a formula defining the value of s_i . The formula defining the value of t states that there exists sets Z_0, \dots, Z_{n-1} such that

- ♦ Z_0, \dots, Z_{n-1} form a partition of the universe,
- ♦ every maximal convex subset of Z_i satisfies φ_i ,
- ♦ for all elements $x < y$, we have either $[x, y] \subseteq Z_i$, for some $i < n$, or $[x, y] \cap Z_k \neq \emptyset$, for all $k < n$. \square

Corollary 4.11. $\text{Th}_{\text{MSO}}(\mathfrak{A})$ is decidable, for every regular linear order \mathfrak{A} .

Proof. We have shown in Corollary V.3.9 that the infinite binary tree $\mathfrak{T}_{\text{bin}}$ has a decidable MSO-theory. Hence, the claim follows from the fact that \mathfrak{A} can be interpreted in $\mathfrak{T}_{\text{bin}}$. \square

Corollary 4.12. *An ordinal α is regular if, and only if, $\alpha < \omega^\omega$.*

Proof. (\Leftarrow) Clearly, every ordinal $\alpha < \omega^\omega$ can be expressed using Leonard-Läuchli operations.

(\Rightarrow) follows by induction on the Leonard-Läuchli term expressing α . Clearly, every singleton is less than ω^ω . If $\alpha = \beta + \gamma$ or $\alpha = \beta \times \omega$, then β and γ must be ordinals and it follows by inductive hypothesis that $\beta, \gamma < \omega^\omega$. Hence, $\alpha < \omega^\omega$. Finally, if $\alpha = \beta \times \omega^{\text{op}}$ or $\alpha = \beta_0 \sqcup \dots \sqcup \beta_{n-1}$, then α is not an ordinal. \square

We conclude this section with a few exercises about the WMSO-theories of ordinals.

Exercise 4.3. Let $1 < m < \omega$ and let $\beta > 0$ be an ordinal. Show that

$$\omega^m \equiv_{\text{FO}}^m \omega^m \beta \quad \text{and} \quad \omega^m \not\equiv_{\text{FO}}^{2m+1} \omega^m + \beta. \quad ,$$

Exercise 4.4. Let $0 < m < \omega$ and let $\alpha < \beta$ be ordinals with $\alpha < \omega^m$. Show that

$$\omega^m \equiv_{\text{WMSO}}^m \omega^m \beta \quad \text{and} \quad \alpha \not\equiv_{\text{WMSO}}^m \beta. \quad ,$$

Exercise 4.5. Let $\alpha, \beta, \gamma, \delta$ be ordinals with $\alpha, \beta < \omega^\omega$. Show that the following statements are equivalent.

- (1) $\omega^\omega \gamma + \alpha \equiv_{\text{FO}} \omega^\omega \delta + \beta$
- (2) $\omega^\omega \gamma + \alpha \equiv_{\text{WMSO}} \omega^\omega \delta + \beta$
- (3) $\alpha = \beta$ and either $\gamma = 0 = \delta$, or $\gamma, \delta > 0$. $\quad ,$

Exercise 4.6. (a) Find an MSO-formula that holds in some uncountable ordinal, but not in any countable one.

(b) Find two countable ordinals α and β with $\alpha \equiv_{\text{MSO}} \beta$. $\quad ,$

5 Modest Linear Orders

We have seen in the preceding section that the MSO-theories of countable linear orders are rather well understood. In particular, we have proved in Corollary 3.3 that the MSO-theory of the order of rational numbers is decidable. It turns out that this is not the case for real numbers. The aim of this section is to find the reason for this difference between the rationals and the reals. We start by introducing some terminology. To keep matters manageable, we will only consider linear orders of the following kind.

Definition 5.1. Let \mathfrak{A} be a linear order.

- (a) \mathfrak{A} is *short* if no subset $X \subseteq A$ has order type ω_1 or ω_1^{op} .
- (b) \mathfrak{A} is *separable* if there exists a countable set D such that every open interval contains an element of D . $\quad ,$

Remark. Every separable linear order is short. $\quad ,$

Exercise 5.1. Prove that every short linear order \mathfrak{A} has size $|A| \leq 2^{\aleph_0}$. $\quad ,$

Remark. A further reason for the restriction to short orders is that one encounters set-theoretic issues when studying decidability problems for general linear orders. For instance, one can show that the decidability of the MSO-theory of $\langle \omega_2, \leq \rangle$ is independent of ZFC [94]. $\quad ,$

We also need generalisations of the notions of a dense set and a perfect one.

Definition 5.2. Let \mathfrak{A} be a linear order.

- (a) We use the usual notation for intervals like

$$(a, b), \quad [a, b), \quad (-\infty, b], \quad (a, \infty), \dots$$
- (b) The *two-sided closure* of $X \subseteq A$ is the set

$$\text{cl}_2(X) := \{ a \in A \mid (b, a) \cap X \neq \emptyset \text{ and } (a, c) \cap X \neq \emptyset, \\ \text{for all } b < a < c \}.$$
- (c) We call a set X *replete* if $X = \text{cl}_2(X) \neq \emptyset$, and *coherent* if it satisfies $\emptyset \neq X \subseteq \text{cl}_2(X)$. $\quad ,$

Example. Every dense subset without end-points is coherent. In particular, if \mathfrak{A} is densely ordered without end-points, the set A is replete. \square

Exercise 5.2. Prove that $\text{cl}_2(\text{cl}_2(X)) \subseteq \text{cl}_2(X)$, for all sets X . The converse holds if X is coherent. \square

Lemma 5.3. Let \mathfrak{A} be a linear order, $D \subseteq A$ coherent, and $X \subseteq D$.

- (a) If X is dense in D , then X is coherent.
- (b) $\text{cl}_2(D)$ is replete.
- (c) $\text{cl}_2(X/D) = \text{cl}_2(X/A) \cap D$, where $\text{cl}_2(X/Y)$ denotes the two-sided closure of X computed in the linear order $\langle Y, \leq \rangle$.
- (d) X is coherent in the linear order $\langle D, \leq \rangle$, if and only if, it is coherent in \mathfrak{A} .
- (e) If D is replete, then X is replete in the linear order $\langle D, \leq \rangle$ if and only if, it is replete in \mathfrak{A} .

Proof. (a) Fix an element $c \in X$. Since $c \in D$, there exists a strictly increasing sequence $(a_i)_i$ of elements $a_i \in D$ with $a_i > c$ and $\sup_i a_i = c$. As X is dense in D , there is some element $x_i \in X$ with $a_i < x_i < c$. It follows that $\sup_i x_i = c$. In the same way we can find a strictly decreasing sequence in X with $\inf_i x_i = c$. Hence, $c \in \text{cl}_2(X)$.

(b) As D is coherent, we have $D \subseteq \text{cl}_2(D)$. This implies that $\text{cl}_2(D) \subseteq \text{cl}_2(\text{cl}_2(D))$. For the converse inclusion, let $x \in \text{cl}_2(\text{cl}_2(D))$ and consider an element $a < x$. Since $x \in \text{cl}_2(\text{cl}_2(D))$, there is some $b \in (a, x) \cap \text{cl}_2(D)$. Similarly, $b \in \text{cl}_2(D)$ implies that there is some $c \in (b, x) \cap D$. We have shown that, for every $a < x$, there is some $c \in D$ with $a < c < x$. Similarly, it follows that, for every $a > x$, there is some $c \in D$ with $x < c < a$. Consequently, $x \in \text{cl}_2(D)$.

(c) Let $(x_i)_i$ be a strictly increasing sequence of elements $x_i \in X$. It is sufficient to show that an element $c \in D$ is the supremum of $(x_i)_i$ in $\langle D, \leq \rangle$ if, and only if, it is the supremum in \mathfrak{A} . (The corresponding statement for strictly decreasing sequences then follows by symmetry.)

(\Leftarrow) is trivial since the supremum computed in \mathfrak{A} is less than or equal to the supremum computed in $\langle D, \leq \rangle$.

(\Rightarrow) For a contradiction, suppose that c is the supremum of $(x_i)_i$ in $\langle D, \leq \rangle$ but not in \mathfrak{A} . Then there exists an element $a \in A$ with $x_i < a < c$, for all i . Since D is coherent and $c \in D \subseteq \text{cl}_2(D)$, there is some element $d \in (a, c) \cap D$. Consequently, c is not the supremum of $(x_i)_i$ in $\langle D, \leq \rangle$. A contradiction.

(d) Let X be coherent in $\langle D, \leq \rangle$. Then it follows by (c) that

$$X \subseteq \text{cl}_2(X/D) = \text{cl}_2(X/A) \cap D \subseteq \text{cl}_2(X/A).$$

Conversely, let X be coherent in \mathfrak{A} . Then it follows by (c) that

$$X \subseteq \text{cl}_2(X/A) \cap D = \text{cl}_2(X/D).$$

(e) Let X be replete in $\langle D, \leq \rangle$. Then it follows by (c) that

$$\begin{aligned} X &= \text{cl}_2(X/D) = \text{cl}_2(X/A) \cap D \\ &= \text{cl}_2(X/A) \cap \text{cl}_2(D/A) = \text{cl}_2(X/A). \end{aligned}$$

Conversely, let X be replete in \mathfrak{A} . Then it follows by (c) that

$$\begin{aligned} \text{cl}_2(X/D) &= \text{cl}_2(X/A) \cap D \\ &= \text{cl}_2(X/A) \cap \text{cl}_2(D/A) = \text{cl}_2(X/A) = X. \end{aligned} \quad \square$$

It turns out that the dividing line between simple and complex linear orders is given by the following property.

Definition 5.4. Let \mathfrak{A} be a linear order and $p < \omega$.

(a) We say that a sequence $\Gamma_0, \dots, \Gamma_{p-1} \subseteq A$ is a *guard* for a set $D \subseteq A$ if

- the sets $\Gamma_0, \dots, \Gamma_{p-1} \subseteq D$ are disjoint and coherent in the linear order $\langle D, \leq \rangle$,
- every Γ_i is dense in D ,
- there is no set $P \subseteq \Gamma_0 \cup \dots \cup \Gamma_{p-1}$ such that each Γ_i is dense in P and P is replete in the linear order $\langle D, \leq \rangle$.

(b) A set X is *guarded* by a guard \bar{I} for D if $X \subseteq D$ and every Γ_i is dense in X .

(c) \mathcal{A} is *p-modest* if it is short and there exists no subset $D \subseteq A$ such that D is densely ordered without end-points and there exists a guard $\Gamma_0, \dots, \Gamma_{p-1}$ of length p for D . \mathcal{A} is *modest* if it is *p-modest* for all $p < \omega$.

Examples. (a) It follows by Corollary 5.26 below that every countable order is modest. (The proof is non-trivial.) In particular, $\langle \mathbb{Q}, \leq \rangle$ is modest.

(b) We will show in Lemma 5.19 below that $\langle \mathbb{R}, \leq \rangle$ is not 1-modest.

Exercise 5.3. Prove that every *p-modest* linear order \mathcal{A} is also *q-modest*, for all $q \leq p$.

Lemma 5.5. *The class of p-modest linear orders is closed under suborders and quotients by congruences.*

Proof. Let \mathcal{A} be *p-modest*. For closure under suborders, consider a subset $B \subseteq A$. For a contradiction, suppose that $\langle B, \leq \rangle$ is not *p-modest*. Let $D \subseteq B$ be the corresponding subset with a guard. Then D is also a subset of A and witnesses that \mathcal{A} is not *p-modest*.

For closure under quotients, consider a congruence \sim on A and let $q : A \rightarrow A/\sim$ be the corresponding projection. Fix any function $e : A/\sim \rightarrow A$ with $q \circ e = \text{id}$. Then the set $B := \text{rng } e$ induces a suborder of \mathcal{A} that is isomorphic to \mathcal{A}/\sim . By the first part of the proof, it follows that B is *p-modest*. Hence, so is \mathcal{A}/\sim . \square

Lemma 5.6. *Let \mathcal{A} be a linear order and \sim a congruence. Then \mathcal{A} is p-modest if, and only if, \mathcal{A}/\sim is p-modest and $\mathcal{A}|_{[a]}$ is p-modest, for every $[a] \in A/\sim$.*

Proof. (\Rightarrow) follows immediately from Lemma 5.5. For (\Leftarrow) , consider a subset $D \subseteq A$ that is densely ordered without end-points. Fix sets $\Gamma_0, \dots, \Gamma_{p-1} \subseteq D$ that are disjoint and coherent in $\langle D, \leq \rangle$ and such that each Γ_i is dense in D . We have to show that \bar{I} is no guard for D .

First, suppose that $a \sim b$, for all $a, b \in D$. Then $\mathcal{A}|_D$ is isomorphic to a suborder of \mathcal{A}/\sim and it follows by Lemma 5.5 that there is no guard for D .

It remains to consider the case where there are $a < b$ in D with $a \sim b$. Then D has a non-trivial interval $I := (a, b) \cap D$ that is contained in some

\sim -class. Set $\Gamma'_i := \Gamma_i \cap I$. Since each \sim -class is *p-modest*, \bar{I}' is no guard for I . Hence, we can find some set $P \subseteq \Gamma'_0 \cup \dots \cup \Gamma'_{p-1}$ that is guarded by \bar{I}' and replete in $\langle I, \leq \rangle$. It follows that Γ_i is dense in P . Furthermore, as I is coherent in $\langle D, \leq \rangle$, Lemma 5.3 (c) implies that

$$P = \text{cl}_2(P/I) = \text{cl}_2(P/D) \cap I = \text{cl}_2(P/D).$$

Hence, P is replete in $\langle D, \leq \rangle$ and it follows that \bar{I} is not a guard for D . \square

Let us give an overview of the properties of modest and non-modest orders. The proofs will be deferred to the end of this section.

Theorem 5.7. *Let \mathcal{A} be a short separable linear order. Then*

$$\mathcal{A} \equiv_{\text{MSO}} \langle \mathbb{Q}, \leq \rangle \quad \text{iff} \quad \mathcal{A} \text{ is modest and densely ordered without end-points.}$$

Corollary 5.8. *Let \mathcal{A} be a short linear order that is separable and dense without end-points. The MSO-theory \mathcal{A} is decidable if, and only if, \mathcal{A} is modest.*

Corollary 5.9. *The MSO-theory of $\langle \mathbb{R}, \leq \rangle$ is undecidable.*

For modest orders that are not necessarily dense, we obtain the following results.

Lemma 5.10. *Every modest short linear order can be embedded into one that is densely ordered and has no end-points.*

Proof. Every linear order \mathcal{A} can be embedded in $\mathbb{Q} \times \mathcal{A}$ via the map $a \mapsto \langle 0, a \rangle$. Applying Lemma 5.6 to the congruence associated with the projection $\mathbb{Q} \times A \rightarrow A$, it follows that $\mathbb{Q} \times \mathcal{A}$ is modest. \square

Theorem 5.11. *Let \mathcal{A} be a separable linear order. The following statements are equivalent.*

(I) \mathcal{A} is modest.

(II) For every $m < \omega$, there is some countable linear order \mathcal{B} with $\mathcal{A} \equiv_{\text{MSO}}^m \mathcal{B}$.

(III) For every $m < \omega$, there is some regular linear order \mathfrak{B} with $\mathfrak{A} \equiv_{\text{MSO}}^m \mathfrak{B}$.

Proof. (II) \Rightarrow (III) follows by Theorem 4.7.

(III) \Rightarrow (I) For each $p < \omega$, there exists an MSO-formula ψ_p expressing that a linear order is p -modest. Let m_p be the quantifier-rank of ψ_p . By assumption, there exist regular linear orders \mathfrak{B}_p with

$$\mathfrak{B}_p \equiv_{\text{MSO}}^{m_p} \mathfrak{A}.$$

Since \mathfrak{B}_p can be embedded in \mathbb{Q} , it is modest by Lemma 5.5. Hence,

$$\mathfrak{B}_p \models \psi_p \quad \text{implies} \quad \mathfrak{A} \models \psi_p.$$

Consequently, \mathfrak{A} is modest.

(I) \Rightarrow (II) Suppose that \mathfrak{A} is modest. Then we can use Lemma 5.10 to embed \mathfrak{A} in some dense modest order \mathfrak{B} without end-points. Given $m < \omega$, set

$$\varphi := \exists X \theta^{(X)}, \quad \text{where} \quad \theta := \bigwedge \text{Th}_{\text{MSO}}^m(\mathfrak{A})$$

and $\theta^{(X)}$ denotes the relativisation of θ to the set X . By Theorem 5.7, we have $\mathfrak{B} \equiv_{\text{MSO}} \langle \mathbb{Q}, \leq \rangle$. Hence,

$$\mathfrak{B} \models \varphi \quad \text{implies} \quad \langle \mathbb{Q}, \leq \rangle \models \varphi,$$

and there exists some set $X_m \subseteq \mathbb{Q}$ such that

$$\mathfrak{A} \equiv_{\text{MSO}}^m \langle X_m, \leq \rangle. \quad \square$$

Corollary 5.12. *The MSO-theory of the class of all separable modest linear orders is decidable.*

Proof. Let T be the MSO-theory of all separable modest linear orders and let \mathcal{C} be the class of all countable linear orders. For an MSO-formula φ we denote by $\varphi^{(X)}$ the relativisation of φ to the set X . Given $\varphi \in \text{MSO}_m$, it follows by Theorem 5.11 and Corollary 5.26 that

$$\begin{aligned} \varphi \in T \quad & \text{iff} \quad \mathfrak{A} \models \varphi, \quad \text{for all countable orders } \mathfrak{A}, \\ & \text{iff} \quad \langle \mathbb{Q}, \leq \rangle \models \forall X \varphi^{(X)}. \end{aligned}$$

The latter is decidable by Corollary 3.3. \square

Open Question. *Do Theorems 5.7 and 5.11 hold without the assumption of separability?*

Coherent Sets and Frames

It remains to provide the proofs we have omitted above because they require some preparation. We start with some tools to work with coherent sets. It will be convenient to use the language of topology.

Definition 5.13. Let \mathfrak{A} be a linear order.

(a) The *order topology* on A is generated by the open basis consisting of all open intervals

$$(-\infty, b), \quad (a, \infty), \quad (a, b), \quad \text{for } a < b \text{ in } A.$$

(b) We denote the *topological closure* of a set $X \subseteq A$ by $\text{cl}(X)$.

(c) \mathfrak{A} is *zero-dimensional* if the order topology has a basis consisting of clopen sets. \lrcorner

Remark. Let \mathfrak{A} be a dense linear order. A convex set $C \subseteq A$ is clopen in the order topology if C has no infimum and no supremum. \lrcorner

Exercise 5.4. Prove that the order topology is always regular, that is, for every closed set C and every element $a \notin C$, there are two disjoint open sets U, V with $a \in U$ and $C \subseteq V$. \lrcorner

Exercise 5.5. Prove that a linear order \mathfrak{A} is separable as defined above if, and only if, its order topology is separable in the usual sense, i.e., if it has a countable dense subset. \lrcorner

Our first observation is that, similar to perfect sets, we can characterise coherent sets via embeddings of certain trees. Instead of binary trees, we use trees of the form $(\omega + \omega^{\text{op}})^*$ although, in order to denote the vertices of such a tree more easily, we will work instead with \mathbb{Z}^* equipped with the following non-standard ordering.

Definition 5.14. (a) We denote by \trianglelefteq the linear ordering on \mathbb{Z}^* defined as follows. For $w, u, v \in \mathbb{Z}^*$ and $d, d' \in \mathbb{Z}$, we set

$$\begin{aligned} wdu \trianglelefteq wd'v & \quad \text{iff} \quad [d < d' \text{ and } (d, d' < 0 \text{ or } d, d' \geq 0)] \\ & \quad \text{or} \quad d \geq 0 > d', \\ w \trianglelefteq wd & \quad \text{iff} \quad d < 0. \end{aligned}$$

$$w0 < w1 < w2 < \dots < w < \dots < w(-3) < w(-2) < w(-1)$$

(b) Let \mathfrak{A} be a linear order. A *frame* for a set $D \subseteq A$ is a bijective function $\varphi : \mathbb{Z}^* \rightarrow D$ with the following properties.

- ♦ φ maps \trianglelefteq to \leq .
- ♦ $\varphi(w) = \inf_{d < 0} \varphi(wd) = \sup_{d \geq 0} \varphi(wd)$, for all $w \in \mathbb{Z}^*$.

Let us collect some basic facts about how the image of a frame is embedded in the surrounding linear order.

Lemma 5.15. Let φ be a frame for D . For $w \in \mathbb{Z}^*$, we denote by $\llbracket w \rrbracket$ the convex closure of $\varphi[\uparrow w]$.

- (I) $\varphi(w) \in \llbracket w \rrbracket$
- (II) $\llbracket w \rrbracket$ is open.
- (III) $\text{cl}(\llbracket wd \rrbracket) \subseteq \llbracket w \rrbracket$, for $d \neq 0, -1$.
- (IV) $\text{cl}(\llbracket wd \rrbracket) \cap \text{cl}(\llbracket wd' \rrbracket) = \emptyset$, for $|d - d'| > 1$.
- (V) Let $(w_n)_{n < \omega}$ be a sequence in \mathbb{Z}^* such that $a := \lim_{n < \omega} \varphi(w_n)$ exists. Then

$$a \in \text{rng } \varphi \quad \text{or} \quad a \in \bigcap_{v \in \beta} \text{cl}(\llbracket v \rrbracket), \quad \text{for some branch } \beta \text{ of } \mathbb{Z}^*.$$

Proof. (I) $\varphi(w) \in \varphi[\uparrow w] \subseteq \llbracket w \rrbracket$

(II) We have

$$\llbracket w \rrbracket = \bigcup \{ (\varphi(wu), \varphi(wv)) \mid u \trianglelefteq v \},$$

which is open.

(III) Let U be the open interval between $\varphi(w(d-1))$ and $\varphi(w(d+1))$. Then

$$\llbracket wd \rrbracket \subseteq U \quad \text{implies} \quad \text{cl}(\llbracket wd \rrbracket) \subseteq \text{cl}(U) \subseteq \llbracket w \rrbracket.$$

(IV) By symmetry, we may assume that $d' > d+1 > 0$. Then

$$\begin{aligned} \text{cl}(\llbracket wd \rrbracket) & \subseteq (-\infty, \varphi(w(d+1)0)], \\ \text{cl}(\llbracket wd' \rrbracket) & \subseteq [\varphi(w(d+1)(-1)), \infty). \end{aligned}$$

In particular, these sets are disjoint.

(V) Set

$$W := \{w_n \mid n < \omega\} \quad \text{and} \quad L_m := \downarrow W \cap \mathbb{Z}^m, \quad \text{for } m < \omega.$$

We distinguish two cases. First, suppose that L_m is infinite for some m . Let m_0 be the least such index. Then there is some $v \in L_{m-1}$ such that the set

$$I := \{d \in \mathbb{Z} \mid v d \in L_m\} \text{ is finite.}$$

Fix a strictly increasing or strictly decreasing sequence d_0, d_1, \dots of elements from I . Then

$$a = \lim_{n < \omega} w_n = \lim_{n < \omega} \varphi(v d_n) = \varphi(v) \in \text{rng } \varphi.$$

It remains to consider the case where all L_m are finite. Then we can use König's Lemma to find an infinite branch $v_0 <_{\text{pf}} v_1 <_{\text{pf}} \dots$ with $v_n \in L_n$, for all $n < \omega$. It follows that

$$a = \lim_{n < \omega} w_n = \lim_{n < \omega} \varphi(v_n) \in \bigcap_{n < \omega} \text{cl}(\llbracket v_n \rrbracket).$$

□

We can use frames to characterise choherency. One of the directions holds without further assumption.

Lemma 5.16. Let φ be a frame for D . Then D is coherent.

Proof. Let $\varphi : \mathbb{Z}^* \rightarrow D$ be a frame for D . To show that D is coherent, fix an element $a \in D$. Then $a = \varphi(w)$, for some $w \in \mathbb{Z}^*$, and it follows that

$$a = \varphi(w) = \inf_{d < 0} \varphi(wd) = \sup_{d \geq 0} \varphi(wd).$$

Hence, $a \in \text{cl}_2(D)$. \square

But without additional assumptions on the linear order, we can only prove the following weaker form of the converse statement.

Lemma 5.17. *Let \mathfrak{A} be a linear order and $D \subseteq A$ a coherent set. There exists a frame φ with $\text{rng } \varphi \subseteq D$.*

Proof. We define $\varphi(w)$ by induction on $|w|$. For $\varphi(\langle \rangle)$ we choose an arbitrary element of D . For the inductive step, suppose that we have already defined $\varphi(w)$ for all $w \in \mathbb{Z}^{\leq n}$. Let $w_0 \triangleleft \dots \triangleleft w_{m-1}$ be an enumeration of all vertices w of length $|w| = n$. As D is coherent and, therefore, densely ordered, we can choose, for every $i < m$, some interval I_i containing $\varphi(w_i)$ such that

$$I_i \cap I_j = \emptyset, \quad \text{for } i < j.$$

Given $i < m$, we can use coherence to choose a strictly increasing sequence $(x_j)_{j < \omega}$ and a strictly decreasing one $(y_j)_{j < \omega}$ such that

$$x_j, y_j \in I_i \cap D \quad \text{and} \quad \sup_{j < \omega} x_j = \varphi(w_i) = \inf_{j < \omega} y_j.$$

For $k \in \mathbb{Z}$, we define

$$\varphi(w; k) := \begin{cases} x_k & \text{if } k \geq 0, \\ y_{-k+1} & \text{if } k < 0. \end{cases} \quad \square$$

With additional assumptions, we obtain a precise characterisation. We start with a simple version of the construction. Below we will present two more elaborate versions with additional requirements on the constructed frame.

Lemma 5.18. *Let \mathfrak{A} be a short zero-dimensional linear order. A countable subset $D \subseteq A$ has a frame if, and only if, D is non-empty and coherent.*

Proof. (\Rightarrow) has already been proved in Lemma 5.16. For (\Leftarrow) , let $D \subseteq A$ be non-empty and coherent. Fix some ordering Ξ on D of order type ω .

In order to define φ , we construct a family $(C_w)_{w \in \mathbb{Z}^*}$ of clopen subsets $C_w \subseteq A$ with $D \cap C_w \neq \emptyset$ and then choose for $\varphi(w)$ the Ξ -minimal element of $D \cap C_w$. We start with $C_{\langle \rangle} := A$. For the inductive step, suppose that we have already defined C_w . As \mathfrak{A} is zero-dimensional and short, there exists a strictly decreasing sequence $C_w = U_0 \supset U_1 \supset \dots$ of clopen sets U_i , $i < \omega$, such that

$$\bigcap_{i < \omega} U_i = \{\varphi(w)\}$$

and each difference $U_i \setminus U_{i+1}$ contains elements strictly above $\varphi(w)$ and ones strictly below $\varphi(w)$. Since D is coherent, there are infinitely many indices $i < \omega$ such that $U_i \setminus U_{i+1}$ contains some element $a \in D$ with $a > \varphi(w)$. Replacing $(U_i)_{i < \omega}$ by a suitable subsequence we may therefore assume that, for every $i < \omega$, there is some $a \in D \cap (U_i \setminus U_{i+1})$ with $a > \varphi(w)$. In the same way we may assume that, for every $i < \omega$, there is some $a \in D \cap (U_i \setminus U_{i+1})$ with $a < \varphi(w)$.

We set

$$C_{wd} := \begin{cases} (\varphi(w), \infty) \cap (U_{-d-1} \setminus U_{-d}) & \text{if } d < 0, \\ (-\infty, \varphi(w)) \cap (U_d \setminus U_{d+1}) & \text{if } d \geq 0, \end{cases}$$

It follows that

$$D \cap C_{wd} \neq \emptyset \quad \text{and} \quad D \cap C_w = \{\varphi(w)\} \cup \bigcup_{d \in \mathbb{Z}} (D \cap C_{wd}).$$

We claim that φ is the desired frame for D . Since

$$D \cap C_w = \{\varphi(w)\} \cup \bigcup_{d \in \mathbb{Z}} (D \cap C_{wd})$$

and since $\varphi(w)$ is the \sqsubseteq -minimal suitable element of D , it follows that every element of D is eventually added to the range of φ . Hence, φ is surjective. Injectivity follows since $\varphi(w) \notin C_u$, for $u >_{\text{pf}} w$.

Furthermore, we have

$$\begin{aligned} \varphi(wd) \leq \varphi(wd') & \quad \text{iff} \quad d < d' < \circ \text{ or } \circ \leq d < d' \text{ or } d \geq \circ > d', \\ \varphi(w) > \varphi(wd), & \quad \text{for } d \geq \circ, \\ \varphi(w) < \varphi(wd), & \quad \text{for } d < \circ. \end{aligned}$$

This implies that φ maps \trianglelefteq to \leq .

To show that $\varphi(w) = \sup_{d \geq \circ} \varphi(wd)$, note that $\bigcap_i U_i = \{\varphi(w)\}$ implies that there is no element a with

$$\varphi(wd) \leq a < \varphi(w), \quad \text{for all } d \geq \circ.$$

The fact that $\varphi(w) = \inf_{d < \circ} \varphi(wd)$ follows in the same way. \square

As an application, let us show that the real numbers are not modest.

Lemma 5.19. $\langle \mathbb{R}, \leq \rangle$ is not 1-modest.

Proof. The set $\Gamma := \mathbb{Q}$ is coherent and dense in \mathbb{R} . We claim that it is a guard for \mathbb{R} . For a contradiction, suppose that there exists a guarded replete set $P \subseteq \mathbb{Q}$. Since \mathbb{Q} is zero-dimensional, we can use Lemma 5.18 to construct a frame $\varphi_\circ : \mathbb{Z}^* \rightarrow \mathbb{Q}$ for P . As the embedding $e : \mathbb{Q} \rightarrow \mathbb{R}$ preserves infima and suprema, it follows that $\varphi := e \circ \varphi_\circ : \mathbb{Z}^* \rightarrow \mathbb{R}$ is also a frame for P . Let β be a branch of \mathbb{Z}^* . As \mathbb{R} is Dedekind-complete, there exists some $x \in \bigcap_{v \in \beta} \text{cl}(\llbracket v \rrbracket)$. Note that $x \in \text{cl}_2(P)$ since, for all elements $a < x < b$, there is some $v \in \beta$ with $\text{cl}(\llbracket v \rrbracket) \subseteq (a, b)$. This implies that there are $u, w \in \mathbb{Z}^*$ with $\varphi(u) \in (a, x)$ and $\varphi(w) \in (x, b)$. Since $\text{cl}_2(P) = P$, it follows that $x = \varphi(w)$, for some $w \in \mathbb{Z}^*$. This implies that $w \in \beta$. Fix a vertex $v \in \beta$ with $w <_{\text{pf}} v$. Then $x \in \text{cl}(\llbracket v \rrbracket)$, but $\varphi(w) \notin \text{cl}(\llbracket v \rrbracket)$. A contradiction. \square

For the undecidability proof below, we will reduce arithmetic to the monadic theory of a non-modest linear order. To do so, we construct a frame

and colour the levels of \mathbb{Z}^* using one colour for each natural number. Then we can show that addition and multiplication of colours is definable.

Definition 5.20. (a) For $P \subseteq \mathbb{Z}^*$, we set

$$\text{len}(P) := \{ |v| \mid v \in P \}.$$

(b) A *stratification* of \mathbb{Z}^* is an infinite partition $(S_n)_{n < \omega}$ of \mathbb{Z}^* such that the sets $\text{len}(S_n)$, for $n < \omega$, are infinite and disjoint. \square

In the following, more complicated construction of a frame, we additionally control the colouring of elements.

Lemma 5.21. Let \mathfrak{A} be a short zero-dimensional linear order containing disjoint sets $\Gamma_0, \dots, \Gamma_{p-1} \subseteq A$ such that each Γ_i is countable and dense in A . Let $D \subseteq A$ be a countable set containing each Γ_i such that the difference $D \setminus (\Gamma_0 \cup \dots \cup \Gamma_{p-1})$ is dense in A . Let \mathcal{P} be a countably infinite partition of D such that each $P \in \mathcal{P}$ is dense in D and $\Gamma_0 \cup \dots \cup \Gamma_{p-1}$ can be written as a union of guarded sets in \mathcal{P} . Then D has a frame φ such that

- ♦ each set $\llbracket w \rrbracket$ is clopen,
- ♦ the sets $(\varphi^{-1}[P])_{P \in \mathcal{P}}$ form a stratification of \mathbb{Z}^* , and
- ♦ $\text{len}(\varphi^{-1}[\Gamma_i]) = \{ n(p+1) + i \mid n < \omega \}$.

Proof. We use a similar construction as in the proof of Lemma 5.18. Fix an enumeration $(P_n)_{n < \omega}$ of \mathcal{P} . Let $\mathcal{P}_g \subseteq \mathcal{P}$ be the set of all sets in \mathcal{P} that are guarded by $\tilde{\Gamma}$ and let $\mathcal{P}_{ng} := \mathcal{P} \setminus \mathcal{P}_g$ be the remaining ones. Note that these two subsets are both non-empty since $D \setminus (\Gamma_0 \cup \dots \cup \Gamma_{p-1})$ is dense in A and $\Gamma_0 \cup \dots \cup \Gamma_{p-1}$ can be written as a union of sets in \mathcal{P}_g . We choose a partition $(L_n)_{n < \omega}$ of ω such that

- ♦ each class L_n is infinite,
- ♦ if $P_n \in \mathcal{P}_g$, then L_n contains no element k with $k \equiv p \pmod{p+1}$ but, for every $q < p$, there are infinitely many $k \in L_n$ with $k \equiv q \pmod{p+1}$,
- ♦ if $P_n \in \mathcal{P}_{ng}$, then $k \equiv p \pmod{p+1}$, for all $k \in L_n$.

To see that such a partition exists, let \mathcal{Q}_k , for $k < p+1$, be a partition of the set

$$\{i(p+1) + k \mid i < \omega\}$$

such that each $\mathcal{Q} \in \mathcal{Q}^k$ is infinite, $|\mathcal{Q}^p| = |\mathcal{P}_{ng}|$, and $|\mathcal{Q}^k| = |\mathcal{P}_g|$, for $k < p$. Then we can set

$$L_n := \mathcal{Q}_0 \cup \dots \cup \mathcal{Q}_{p-1} \quad \text{or} \quad L_n := \mathcal{Q}_p, \quad \text{for a suitable } \mathcal{Q}_k \in \mathcal{Q}_k.$$

Fix some ordering Ξ on D of order type ω .

In order to define φ , we construct a family $(C_w)_{w \in \mathbb{Z}^*}$ of clopen subsets $C_w \subseteq A$ with $C_w \cap D \neq \emptyset$ starting with $C_{\langle \rangle} := A$. For the inductive step, suppose that we have already defined C_w . Let $n < \omega$ be the index such that $|w| \in L_n$. For $\varphi(w)$ we choose the Ξ -minimal element $a \in D$ satisfying

- ♦ $a \in C_w \cap P_n$ and
- ♦ $P_n \in \mathcal{P}_g$ implies $a \in I_{|w| \bmod (p+1)}$.

To see that this is possible note that coherence of D implies that C_w contains infinitely many elements of D . Since P_n is dense in D , it follows that $C_w \cap P_n \neq \emptyset$. Furthermore, if $P_n \in \mathcal{P}_g$, then each I_i is dense in P_n , which implies that I_i is dense in D . As above, it follows that $C_w \cap P_n \cap I_i \neq \emptyset$.

As in the proof of Lemma 5.18, we construct a strictly decreasing sequence $C_w = U_0 \supset U_1 \supset \dots$ of clopen sets U_i , $i < \omega$, such that

$$\bigcap_{i < \omega} U_i = \{\varphi(w)\}$$

and each difference $U_i \setminus U_{i+1}$ contains elements of D that are strictly above $\varphi(w)$ and ones that are strictly below $\varphi(w)$.

$$C_{wd} := \begin{cases} (\varphi(w), \infty) \cap (U_{-d-1} \setminus U_{-d}) & \text{if } d < 0, \\ (-\infty, \varphi(w)) \cap (U_d \setminus U_{d+1}) & \text{if } d \geq 0, \end{cases}$$

It follows that

$$D \cap C_{wd} \neq \emptyset \quad \text{and} \quad D \cap C_w = \{\varphi(w)\} \cup \bigcup_{d \in \mathbb{Z}} (D \cap C_{wd}).$$

We claim that φ is the desired frame for D . The fact that φ is a frame for D follows as in the proof of Lemma 5.18. Hence, it remains to prove the two additional properties.

By choice of $\varphi(w)$, we have

$$\varphi(w) \in P_n \quad \text{iff} \quad |w| \in L_n.$$

This implies that

$$\text{len}(\varphi^{-1}[P_n]) = L_n.$$

Hence, the sets $\varphi^{-1}[P_n]$ form a stratification of \mathbb{Z}^* .

Furthermore, for $q < p$, we have

$$|w| \equiv q \pmod{p+1} \quad \text{iff} \quad |w| \in L_n \text{ for some } n \text{ with } P_n \in \mathcal{P}_g.$$

By choice of $\varphi(w)$, this implies that $\varphi(w) \in I_q$. Conversely, suppose that $\varphi(w) \in I_q$ and let n be the index such that $\varphi(w) \in P_n$. As $I_0 \cup \dots \cup I_{p-1}$ can be written as a union of sets in \mathcal{P}_g , the set P_n must belong to \mathcal{P}_g . By definition of $\varphi(w)$, this implies that $|w| \equiv q \pmod{p+1}$.

Finally, by choice of $\varphi(w)$ and C_{wd} , we have $\llbracket w \rrbracket \subseteq C_w$ and $\llbracket w \rrbracket$ is dense in C_w . Since $\llbracket w \rrbracket$ and C_w are both open it, follows that $\llbracket w \rrbracket = C_w$. In particular, $\llbracket w \rrbracket$ is clopen. \square

Subframes and Frame Embeddings

Once we have constructed a frame, we can use the tree structure it imposes for further constructions. In particular, we can find subtrees with additional properties.

Definition 5.22. A *frame embedding* is an injective function $\mu : \mathbb{Z}^* \rightarrow \mathbb{Z}^*$ with the following two properties.

- ♦ $\mu(u \sqcap_{\text{pf}} v) = \mu(u) \sqcap_{\text{pf}} \mu(v)$, for all $u, v \in \mathbb{Z}^*$.
- ♦ $\mu(w)d \leq_{\text{pf}} \mu(wd')$ implies $d \geq 0 \Leftrightarrow d' \geq 0$, for all $w \in \mathbb{Z}^*$ and $d, d' \in \mathbb{Z}$.

Lemma 5.23. *Let φ be a frame for D and μ a frame embedding. Then $\varphi \circ \mu$ is a frame for some subset of D .*

Proof. Set $E := \text{rng}(\varphi \circ \mu)$. Then $\varphi \circ \mu$ is a bijection $\mathbb{Z}^* \rightarrow E$. Furthermore, $\varphi \circ \mu$ maps \leq to \leq since $\mu(v)d \leq_{\text{pf}} \mu(vd')$ implies $d \geq \circ \Leftrightarrow d' \geq \circ$.

Finally, note that the set

$$I := \{d \in \mathbb{Z} \mid d \geq \circ, \mu(v)d \leq_{\text{pf}} \varphi(vd') \text{ for some } d' \geq \circ\}$$

is infinite. Consequently, we have

$$\sup_{d \geq \circ} \varphi(\mu(vd)) = \sup_{d \in I} \varphi(\mu(v)d) = \sup_{d \geq \circ} \varphi(\mu(v)d) = \varphi(\mu(v)).$$

The lower bound follows analogously. \square

In the same way we can divide every perfect set into countably many pairwise disjoint perfect sets, we can divide a coherent set into continuum many ‘nearly disjoint’ coherent sets.

Proposition 5.24. *Let \mathfrak{A} be a short linear order and $\mathcal{G} \subseteq \wp(A)$ a countable family of sets whose union $\Gamma_* := \bigcup \mathcal{G}$ is coherent and such that each $\Gamma \in \mathcal{G}$ is dense in A . There exist a countable coherent set $D \subseteq \bigcup \mathcal{G}$ and a family $\mathcal{S} \subseteq \wp(D)$ with the following properties.*

- ♦ $|\mathcal{S}| = 2^{\aleph_0}$
- ♦ Every $S \in \mathcal{S}$ is coherent, closed in D , and nowhere dense in D .
- ♦ Every $\Gamma \in \mathcal{G}$ is dense in each $S \in \mathcal{S}$.
- ♦ $\text{cl}(S) \cap \text{cl}(T) \subseteq D$ is scattered, for all distinct $S, T \in \mathcal{S}$.

Proof. Fix a surjective function $\chi : \mathbb{Z} \rightarrow \mathcal{G}$. We start by constructing a frame $\varphi : \mathbb{Z}^* \rightarrow \Gamma_*$ and a family $(U_w)_{w \in \mathbb{Z}^*}$ of open intervals with the following properties.

- (I) $\varphi(w) \in U_w$
- (II) $U_v \subseteq U_w$, for $w \leq_{\text{pf}} v$.
- (III) $U_{wd} \cap U_{wd'} = \emptyset$, for $d \neq d'$.

- (IV) $\varphi(w(5d+i)) \in \chi(d)$, for all $d \in \mathbb{Z}$ and $\circ \leq i < 5$.

We construct $\varphi(w)$ and U_w by induction on w . For $w = \langle \rangle$, we choose an arbitrary open interval $U_{\langle \rangle}$ and an arbitrary element $\varphi(\langle \rangle) \in U_{\langle \rangle} \cap \Gamma_*$. For the inductive step, suppose that we have already defined $\varphi(w)$ and U_w . Since \mathfrak{A} is short, we can choose a strictly increasing sequence $a_\circ < a_1 < \dots$ of elements in U_w with $\sup_{n < \omega} a_n = \varphi(w)$. Similarly, U_w contains a strictly decreasing sequence $b_\circ > b_1 > \dots$ with $\inf_{n < \omega} b_n = \varphi(w)$. For $d \in \mathbb{Z}$, fix an open interval

$$U_{wd} \subseteq \begin{cases} (a_{d+1}, a_d) & \text{if } d \geq \circ, \\ (b_{-d}, b_{-d-1}) & \text{if } d < \circ. \end{cases}$$

Since $\chi(d)$ is dense, we can choose elements

$$\varphi(w(5d+i)) \in U_{w(5d+i)} \cap \chi(d), \quad \text{for } d \in \mathbb{Z} \text{ and } \circ \leq i < 5.$$

Clearly, the function φ and the sets $(U_w)_w$ satisfy properties (I)–(IV). Furthermore, we have

$$\varphi(w) = \inf_{d < \circ} \varphi(wd) = \sup_{d \geq \circ} \varphi(wd),$$

so φ is indeed a frame.

Having defined φ , we set $D := \text{rng } \varphi$. Then φ is a frame for D , and it follows by Lemma 5.16 that D is coherent. For each $\alpha \in 2^\omega$, we define a frame embedding $\mu_\alpha : \mathbb{Z}^* \rightarrow \mathbb{Z}^*$ as follows. Let $\gamma_\circ, \gamma_1 : \mathbb{Z} \rightarrow \mathbb{Z}$ be the two functions with

$$\gamma_i(d) := 5d + 2i + 1.$$

We set

$$\mu_\alpha(\langle \rangle) := \langle \rangle \quad \text{and} \quad \mu_\alpha(wd) := \mu_\alpha(w) \gamma_{\alpha(|w|)}(d).$$

We claim that the family

$$\mathcal{S} := \{S_\alpha \mid \alpha \in 2^\omega\} \quad \text{with} \quad S_\alpha := \text{rng}(\varphi \circ \mu_\alpha)$$

has the desired properties. Since μ_α is a frame embedding, it follows by Lemmas 5.23 and 5.16 that each S_α is coherent. Furthermore, (iv) implies that every $I \in \mathcal{G}$ is dense in S_α .

To see that S_α is closed in D , suppose that $\varphi(w) \notin S_\alpha$. Then (ii) and (iii) implies that $U_w \cap S_\alpha = \emptyset$. Hence, $\varphi(w) \notin \text{cl}(S_\alpha)$.

To check that S_α is nowhere dense, let V be non-empty and open. If V is disjoint from S_α , we are done. Otherwise, fix some $w \in (\varphi \circ \mu_\alpha)^{-1}[V]$ and set $v := \mu_\alpha(w)$. Since V is open, there is some d such that

$$U_{vd} \cap V \neq \emptyset \quad \text{and} \quad d \equiv 2 \pmod{5}.$$

This implies that $U_{vd} \cap S_\alpha = \emptyset$.

Finally, consider an element $a \in \text{cl}(S_\alpha) \cap \text{cl}(S_{\alpha'})$ where $\alpha \neq \alpha'$. Let $m < \omega$ be the least number with $\alpha(m) \neq \alpha'(m)$. We start by showing that $a \in D$. For a contradiction, suppose that it is not. By Lemma 5.15 (iv), we can then find a branch β of ω^* with $a \in \bigcap_{v \in \beta} \text{cl}(U_v)$. Let $v \in \beta$ be the vertex with $|v| = m + 1$. Then $\text{cl}(U_v)$ is disjoint from either S_α or $S_{\alpha'}$. Consequently, $a \notin S_\alpha \cap S_{\alpha'}$. A contradiction.

It remains to show that $\text{cl}(S_\alpha) \cap \text{cl}(S_{\alpha'})$ is scattered. Let $w := \varphi^{-1}(a)$. Note that $\varphi(w) \in \text{cl}(S_\alpha) \cap \text{cl}(S_{\alpha'})$ implies that $m \leq |w|$. Consequently, we have

$$\text{cl}(S_\alpha) \cap \text{cl}(S_{\alpha'}) \subseteq \{ \varphi(w) \mid |w| \leq m \},$$

which is scattered. \square

Corollary 5.25. *Let \mathfrak{A} be a short linear order and \tilde{I} a guard for A . Every guarded replete set has size 2^{\aleph_0} .*

Proof. For a contradiction, suppose that $P \subseteq A$ is guarded and replete, but $|P| < 2^{\aleph_0}$. Note that, by Lemma 5.3 (c) and (d), a set $X \subseteq P$ is coherent or replete in $\langle P, \leq \rangle$ if, and only if, it is coherent or replete in \mathfrak{A} . As each set I_i is dense in P , it further follows by Lemma 5.3 (a) that $P \cap (I_0 \cup \dots \cup I_{p-1})$ is coherent. Hence, we can apply Proposition 5.24 to the linear order $\langle P, \leq \rangle$ and the family $\mathcal{G} = \{I_0 \cap P, \dots, I_{p-1} \cap P\}$. Let $D \subseteq P \cap (I_0 \cup \dots \cup I_{p-1})$

be the resulting countable coherent set and $\mathcal{S} \subseteq \wp(D)$ the resulting family of coherent sets. Then we have

$$\text{cl}(S) \cap \text{cl}(T) \subseteq D, \quad \text{for all } S \neq T \text{ in } \mathcal{S},$$

(we compute the closure $\text{cl}(S)$ in the order $\langle P, \leq \rangle$), which implies that the family

$$\{ \text{cl}(S) \setminus D \mid S \in \mathcal{S} \}$$

is disjoint. Since $|P| < 2^{\aleph_0} = |\mathcal{S}|$, we can find some $S \in \mathcal{S}$ with $\text{cl}(S) \setminus D = \emptyset$. Consequently,

$$\text{cl}(S) \subseteq D \subseteq P \cap (I_0 \cup \dots \cup I_{p-1}).$$

Note that, by choice of \mathcal{S} , each set I_i is dense in S . If we can show that the set S is replete, we obtain a contradiction to the fact that \tilde{I} is a guard. For the proof, note that coherency implies that $S \subseteq \text{cl}_2(S)$, while the above inclusion and the fact that S is closed in D imply that

$$\text{cl}_2(S) \subseteq \text{cl}(S) \subseteq \text{cl}(S) \cap D \subseteq S. \quad \square$$

Corollary 5.26. *Every short linear order \mathfrak{A} of size $|A| < 2^{\aleph_0}$ is modest.*

Proof. For a contradiction, suppose that there exists a set $D \subseteq A$ with a guard I_0, \dots, I_{p-1} such that D is densely ordered without end-points. Applying Proposition 5.24 to the order $\mathfrak{D} := \langle D, \leq \rangle$, we obtain a countable coherent set $C \subseteq I_0 \cup \dots \cup I_{p-1}$ and a family $\mathcal{S} \subseteq \wp(C)$ of coherent sets of size $|\mathcal{S}| = 2^{\aleph_0}$ such that

- I_i is dense in S , for all $S \in \mathcal{S}$ and $i < p$, and
- $\text{cl}(S) \cap \text{cl}(T) \subseteq C$, for all $S \neq T$ in \mathcal{S} .

(We compute the closure $\text{cl}(S)$ in \mathfrak{D} .) This implies that the family

$$\{ \text{cl}(S) \setminus C \mid S \in \mathcal{S} \}$$

is disjoint. As C is countable, we can therefore find some $S \in \mathcal{S}$ with $\text{cl}(S) \subseteq C$. Consequently,

$$\text{cl}(S) \subseteq C \subseteq \Gamma_0 \cup \dots \cup \Gamma_{p-1}.$$

By Lemma 5.3, it follows that the set $P := \text{cl}_2(S)$ is replete in \mathfrak{D} . Furthermore, since Γ_i is dense in S , Γ_i is dense in P . Hence, P is guarded and replete, which contradicts the definition of a guard. \square

Encoding Numbers

For the undecidability proof, we have to encode natural numbers in the levels of a frame. This will be done in Theorem 5.32 below. Before doing so, let us observe that we can restrict our attention to orders that are zero-dimensional and separable.

Lemma 5.27. *Let $0 < p < \omega$. Every short linear order \mathfrak{A} that is not p -modest contains a set $D \subseteq A$ such that the order $\langle D, \leq \rangle$ is zero-dimensional, separable, and not p -modest.*

Proof. As \mathfrak{A} is not p -modest, there exists a set $X^\circ \subseteq A$ with a guard $\Gamma_0^\circ, \dots, \Gamma_{p-1}^\circ$ such that X° is densely ordered without end-points. Replacing \mathfrak{A} by the order $\langle X^\circ, \leq \rangle$ we may assume without loss of generality that $X^\circ = A$. By Proposition 5.24, we can find a (in fact uncountably many) countable coherent set $C \subseteq \Gamma_*^\circ := \Gamma_0^\circ \cup \dots \cup \Gamma_{p-1}^\circ$ such that every set Γ_i° is dense in C . Set

$$D_* := \text{cl}_2(C), \quad \Gamma_i := \Gamma_i^\circ \cap C, \quad \text{and} \quad \Gamma_* := \Gamma_0 \cup \dots \cup \Gamma_{p-1}.$$

For each $i < p$, we choose a subset $Z_i \subseteq \Gamma_i$ such that Z_i and $\Gamma_i \setminus Z_i$ are dense in Γ_i . Setting $Z_* := Z_0 \cup \dots \cup Z_{p-1}$ it follows that

- ♦ Z_* and $\Gamma_* \setminus Z_*$ are dense in Γ_* , and
- ♦ each $\Gamma_i \cap Z_* = Z_i$ is dense in Z_* and $\Gamma_i \setminus Z_* = \Gamma_i \setminus Z_i$ is dense in $\Gamma_* \setminus Z_*$.

We claim that the set $D := D_* \setminus Z_*$ has the desired properties.

To check that $\langle D, \leq \rangle$ is separable, note that the set $C \setminus Z_*$ is dense in $D_* \setminus Z_*$ since $\Gamma_* \setminus Z_*$ is dense in Γ_* , Γ_* is dense in C , and C is dense in D_* .

To see that $\langle D, \leq \rangle$ is zero-dimensional, note that

$$\{ (a, b) \cap (D_* \setminus Z_*) \mid a, b \in Z_*, a < b \}$$

is an open basis for the topology of D because Z_* is dense in Γ_* and, hence, dense in D_* . Furthermore, the complement

$$(D_* \setminus Z_*) \setminus (a, b) = [(-\infty, a) \cup (b, \infty)] \cap (D_* \setminus Z_*)$$

of each base set is also open.

It remains to prove that $\langle D, \leq \rangle$ is not p -modest. By Lemma 5.3 (b), the set $D_* := \text{cl}_2(C)$ is replete. In particular, it is densely ordered and it has no end-points. Thus, it is sufficient to show that the sequence $\Gamma_0 \setminus Z_*, \dots, \Gamma_{p-1} \setminus Z_*$ forms a guard for D .

Since $\Gamma_i \setminus Z_*$ is dense in Γ_i , it follows by Lemma 5.3 (a) that $\Gamma_i \setminus Z_*$ is coherent. Furthermore, the facts that $\Gamma_i \setminus Z_*$ is dense in Γ_i , that Γ_i is dense in C , and that C is dense in D_* imply that $\Gamma_i \setminus Z_*$ is dense in $D = D_* \setminus Z_*$.

For a contradiction suppose that there exists a guarded set $P \subseteq \Gamma_* \setminus Z_*$ that is replete in $\langle D_* \setminus Z_*, \leq \rangle$. By Lemma 5.3, the set

$$P' := \text{cl}_2(P) \cap D_*$$

is replete in $\langle D_*, \leq \rangle$ and, therefore, also in \mathfrak{A} . Furthermore, Γ_i° is dense in P' since $\Gamma_i \setminus Z_*$ is dense in P and P is dense in $\text{cl}_2(P) \cap D_* = P'$. Finally, since P is replete in $\langle D_* \setminus Z_*, \leq \rangle$, we have

$$\text{cl}_2(P) \cap (D_* \setminus Z_*) = P \subseteq \Gamma_*.$$

Hence, $Z_* \subseteq \Gamma_*$ implies that

$$P' = \text{cl}_2(P) \cap D_* = [\text{cl}_2(P) \cap (D_* \setminus Z_*)] \cup [\text{cl}_2(P) \cap Z_*] \subseteq \Gamma_* . \square$$

As mentioned above we will colour the levels of a frame. These colours can then be used to construct subframes with additional properties.

Definition 5.28. Let $X \subseteq \mathbb{Z}^*$

(a) We say that X *varies* at level $l < \omega$ if there are $u, v \in X$ with $|u|, |v| > l$ and $u(l) \neq v(l)$.

(b) Given a set $S \subseteq \omega$, we say that X has *colour* S if it only varies at levels $l \in S$.

(c) Given a family $\mathcal{S} \subseteq \mathcal{P}(\omega)$, we call X *entirely \mathcal{S} -colourless* if the only subsets $C \subseteq X$ with a colour in \mathcal{S} are chains.

(d) Given a tuple \bar{I} of disjoint infinite subsets $I_i \subseteq \omega$, we say that X is \bar{I} -*embedded* if

- ♦ X is non-empty,
- ♦ for every $x \in X$ there are infinitely many $d < \circ$ with $\uparrow x d \cap X \neq \emptyset$ and there are infinitely many $d \geq \circ$ with $\uparrow x d \cap X \neq \emptyset$,
- ♦ for every $x \in X$ and every index i , there is some $u \in X$ with $u >_{\text{pf}} x$ and $|u| \in I_i$.

Remark. Note that every chain $C \subseteq \mathbb{Z}^*$ varies at no levels and has therefore every colour.

In the following two technical lemmas we construct subframes with certain colourings. They will be used in the proof of Theorem 5.32 below.

Lemma 5.29. Let φ be a frame for D , and let $S, \Gamma_0, \dots, \Gamma_{p-1} \subseteq \omega$ and $X, V \subseteq \mathbb{Z}^*$ be sets such that

- ♦ V is upwards closed,
- ♦ $\Gamma_0, \dots, \Gamma_{p-1}$ are infinite and disjoint,
- ♦ $\text{len}(V \cap X) \subseteq S \subseteq \Gamma_0 \cup \dots \cup \Gamma_{p-1}$,
- ♦ for all $u \in V$ and $q < p$, there is some $v \in X$ with $v >_{\text{pf}} u$ and $|v| \in \Gamma_q$.

There exists a frame embedding $\mu : \mathbb{Z}^* \rightarrow V$ such that

- ♦ $\text{rng } \mu$ has colour S and
- ♦ $\varphi[X \cap \text{len}^{-1}(\Gamma_q) \cap \text{rng } \mu]$ is dense in $\varphi[\text{rng } \mu]$, for all $q < p$.

Proof. Choose an enumeration $(w_n)_{n < \omega}$ of \mathbb{Z}^* such that $|w_n| \leq n$ and, for every $u \in \mathbb{Z}^*$ and $q < p$, there is some index

$$n \equiv q \pmod{p} \quad \text{with} \quad w_n = u.$$

To see that such an enumeration exists, fix a bijection $g : \omega \rightarrow \mathbb{Z}^*$ satisfying $|g(n)| \leq n$, for all n . Then we can set

$$w_{np+q} := g(n), \quad \text{for } q < p \text{ and } n < \omega.$$

Using $(w_n)_n$ we now define a frame embedding $\mu : \mathbb{Z}^* \rightarrow V$ and a function $f : \omega \rightarrow \omega$ such that

- ♦ $|\mu(v)| = f(|v|)$, for all $v \in \mathbb{Z}^*$,
- ♦ $\text{rng } \mu$ varies at l iff $l \in \text{rng } f$,
- ♦ for all $n < \omega$ and $q < p$, there is some $v \in X \cap \text{rng } \mu$ with $v >_{\text{pf}} \mu(w_n)$ and $|v| \in \Gamma_q$.

We start with choosing an element $\mu(\langle \rangle) \in V \cap X$ such that $\uparrow \mu(\langle \rangle) \subseteq V$ and we set $f(0) := |\mu(\langle \rangle)|$. For the inductive step, suppose that we have already defined $\mu \upharpoonright \mathbb{Z}^{\leq n}$ and $f(n)$. Let $q := n \bmod p$ and choose a vertex $u \in \mathbb{Z}^n$ with $u >_{\text{pf}} w_n$. By assumption on X and V , there exists an element

$$v \in X \quad \text{with} \quad v >_{\text{pf}} \mu(v) \quad \text{and} \quad |v| \in \Gamma_q.$$

We set

$$f(n+1) := |v|$$

and, for $s \in \mathbb{Z}^n$ and $d \in \mathbb{Z}$,

$$\mu(sd) := t \quad \text{where} \quad t(i) := \begin{cases} \mu(s)(i) & \text{if } i < f(n), \\ d & \text{if } i = f(n), \\ v(i) & \text{if } f(n) < i < f(n+1). \end{cases}$$

Then $\uparrow \mu(s) \subseteq V$ implies that $\mu(sd) \in V$. Furthermore,

- ♦ $|\mu(sd)| = f(n+1)$,
- ♦ $\text{rng } \mu \upharpoonright \mathbb{Z}^{\leq (n+1)}$ varies only at $f(0), \dots, f(n+1)$,
- ♦ $v >_{\text{pf}} \mu(w_n)$, $v \in X \cap \text{rng } (\mu \upharpoonright \mathbb{Z}^{n+1})$, and $|v| \in \Gamma_q$.

We claim that the frame embedding μ defined in this may has the desired properties. First, note that $\text{rng } \mu$ has colour S because it varies only at levels in

$$\text{rng } f \subseteq \text{len}(V \cap X) \subseteq S.$$

It therefore remains to prove that every set $\varphi[X \cap \text{len}^{-1}(\Gamma_q) \cap \text{rng } \mu]$ is dense in $\varphi[Y]$. It is sufficient to show that, for every $u \in \mathbb{Z}^*$, there is some $v \in X \cap \text{len}^{-1}(\Gamma_q) \cap \text{rng } \mu$ with $v \geq_{\text{pf}} \mu(u)$. Hence, fix $u \in \mathbb{Z}^*$. By choice of $(w_n)_n$, there exists an index $n < \omega$ such that

$$w_n = u \quad \text{and} \quad n \equiv q \pmod{p}.$$

By definition of μ , we can further find an element $v >_{\text{pf}} \mu(w_n)$ with

$$v \in X \cap \text{rng}(\mu \upharpoonright \mathbb{Z}^{n+1}) \quad \text{and} \quad |v| \in \Gamma_q.$$

Consequently, $v \in X \cap \text{len}^{-1}(\Gamma_q) \cap \text{rng } \mu$, as desired. \square

Lemma 5.30. *Let $\Gamma_0, \dots, \Gamma_{p-1} \subseteq \omega$ a finite sequence of infinite, disjoint sets, and let $\mathcal{S} \subseteq \mathcal{P}(\omega)$ be a family of colours such that*

$$S \cap S' \cap \Gamma_q = \emptyset, \quad \text{for all } S \neq S' \text{ in } \mathcal{S} \text{ and all } q < p.$$

Given a $\bar{\Gamma}$ -embedded set $X \subseteq \mathbb{Z}^$ such that no set of the form $X \cap \upharpoonright u$ with $u \in \mathbb{Z}^*$ has a colour $S \in \mathcal{S}$, there exists a subset $Y \subseteq X$ that is $\bar{\Gamma}$ -embedded and entirely \mathcal{S} -colourless.*

Proof. Below we will construct a frame embedding $\mu : \mathbb{Z}^* \rightarrow X$ with the following properties.

- ♦ $|\mu(v)| \in \Gamma_{|v| \bmod p}$
- ♦ If $|\mu(v)| \in S \in \mathcal{S}$, then each set of the form $\{\mu(vd), \mu(vd')\}$ with $d \neq d'$ varies at some $l \notin S$.

Before doing so, let us show that the resulting set $Y := \text{rng } \mu$ is $\bar{\Gamma}$ -embedded and entirely \mathcal{S} -colourless.

We start with the former. Clearly, $Y \neq \emptyset$. For the second condition, fix $y \in Y$. Let $u := \mu^{-1}(y)$ and, for $d \in \mathbb{Z}$, let $k_d \in \mathbb{Z}$ be the direction such that $\mu(ud) \geq \mu(u)k_d$. It follows that

$$\{d \geq 0 \mid \upharpoonright yd \cap Y \neq \emptyset\} = \{k_d \mid d \geq 0\}$$

is infinite. In the same way, we can show that there are infinitely many $k < 0$ with $\upharpoonright yk \cap Y \neq \emptyset$.

For the third condition, fix $y \in Y$ and $q < p$. Choose some $u >_{\text{pf}} \mu^{-1}(y)$ with $|u| \equiv q \pmod{p}$. Then $\mu(u) \in Y \cap \upharpoonright y \setminus \{y\}$ and $|\mu(u)| \in \Gamma_q$.

We also have to show that Y is entirely \mathcal{S} -colourless. For a contradiction, suppose that some subset $\{\mu(u), \mu(v)\}$ has a colour $S \in \mathcal{S}$. Set $w := u \sqcap_{\text{pf}} v$. Then $\{\mu(u), \mu(v)\}$ varies at $|w|$. As $\{\mu(u), \mu(v)\}$ has colour S , this implies that $\mu(w) \in S$. By definition of μ , it therefore follows that $\{\mu(u), \mu(v)\}$ also varies at some level $l \notin S$. Consequently, S cannot be the colour of the set $\{\mu(u), \mu(v)\}$. A contradiction.

It remains to construct the embedding μ . We start by picking an arbitrary element $\mu(\langle \rangle) \in X \cap \text{len}^{-1}[\Gamma_0]$. For the inductive step, suppose that we have already defined $\mu(v)$. Set

$$M_o := \{m \in \mathbb{Z} \mid \upharpoonright \mu(v)m \cap X \neq \emptyset\}.$$

Since X is $\bar{\Gamma}$ -embedded, M_o contains infinitely many non-negative numbers and infinitely many negative ones.

If $|\mu(v)| \notin \bigcup \mathcal{S}$, we choose elements

$$\mu(vi) \in \upharpoonright \mu(v)m_i \cap X \cap \text{len}^{-1}[\Gamma_q], \quad \text{where} \quad q \equiv |v| + 1 \pmod{p}$$

and m_i is the i -th non-negative element of M_o , if $i \geq 0$, and the i -th negative element of M_o , if $i < 0$. Then $|\mu(vi)| \in \Gamma_q$ and the second condition holds vacuously.

It remains to consider the case where $\mu(v) \in S \in \mathcal{S}$. Since $|\mu(v)| \in \Gamma_q$ and $S \cap S' \cap \Gamma_q = \emptyset$, for all distinct $S, S' \in \mathcal{S}$, it follows that this set S is unique. Fix an enumeration $(d_i)_{i < \omega}$ of \mathbb{Z} . By induction on i , we choose elements $\mu(vd_i)$, disjoint finite sets $L_i \subseteq \omega \setminus S$, and a decreasing sequence

$M_0 \supset M_1 \supset \dots$ of infinite sets $M_{i+1} \subseteq \mathbb{Z}$ with the following properties. (Intuitively, M_i is the set of directions we can still choose from and L_i is the set of levels where the successors vary.)

- ♦ $\mu(v) \leq_{\text{pf}} \mu(vd_i)$
- ♦ M_i contains infinitely many non-negative elements and infinitely many negative ones.
- ♦ $m \in M_i$ implies that $\mu(v)m \not\leq_{\text{pf}} \mu(vd_j)$, for all $j < i$.
- ♦ For every $m \in M_i$, there is some $z \in \uparrow \mu(v)m \cap X$ such that $z \upharpoonright L_j$ is defined and different from $\mu(vd_j) \upharpoonright L_j$, for all $j < i$.

(Note that the last condition implies that $L_j \subseteq \uparrow |\mu(v)|$.) Hence, suppose that we have already defined $\mu(vd_j)$, L_j , and M_{j+1} , for all $j < i$. Set

$$n := \begin{cases} \min \{ m \in M_i \mid m \geq 0 \} & \text{if } d_i \geq 0, \\ \max \{ m \in M_i \mid m < 0 \} & \text{if } d_i < 0. \end{cases}$$

By inductive hypothesis, we can fix some $u \in \uparrow \mu(v)n \cap X$ such that $u \upharpoonright L_j$ is defined and different from $\mu(vd_j) \upharpoonright L_j$, for all $j < i$. By assumption on X , we can find elements $u_0, u_1 \geq_{\text{pf}} u$ and a set $L_i \subseteq \omega \setminus S$ such that $u_0, u_1 \in X$ and the subsequences $u_0 \upharpoonright L_i$ and $u_1 \upharpoonright L_i$ are defined and different. For $s < 2$, set

$$N_s := \{ m \in M_i \mid \text{there is some } z \geq_{\text{pf}} vm \text{ such that} \\ z \upharpoonright L_j \text{ is defined and different from } \mu(vd_j) \upharpoonright L_j \text{ for } j < i \\ \text{and } z \upharpoonright L_i \text{ is defined and different from } u_s \upharpoonright L_i \}.$$

Since $|\mu(v)| \notin L_i$ and there is at most one index $j < i$ with $|\mu(v)| \in L_j$, it follows that $|M_i \setminus N_s| \leq 1$. Consequently, there is some $s < 2$ such that

$$N_s \cap [0, \infty) \text{ is infinite and } N_s \cap (-\infty, 0) \text{ is infinite.}$$

We set $M_{i+1} := N_s \setminus \{n\}$. Then

- ♦ M_{i+1} contains infinitely many non-negative numbers and infinitely many negative ones,

- ♦ $m \in M_{i+1}$ implies that $\mu(v)m \not\leq_{\text{pf}} \mu(vd_j)$, for all $j < i+1$, and
- ♦ for every $m \in M_{i+1}$, there is some $z \in \uparrow \mu(v)m \cap X$ such that $z \upharpoonright L_j$ is defined and different from $\mu(vd_j) \upharpoonright L_j$, for $j < i$, and also different from $z_s \upharpoonright L_i$.

Finally, since X is \tilde{T} -embedded, we can choose an element

$$\mu(vd_i) \in \uparrow u_s \cap X \cap \text{len}^{-1}[\Gamma_q] \quad \text{where } q := |v| + 1 \bmod p.$$

Then $|\mu(vd_i)| \in \Gamma_q$ and each set of the form $\{\mu(vd_i), \mu(vd_j)\}$ with $j < i$ varies at some level $l \in L_j \subseteq \omega \setminus S$. \square

While frames are usually not definable, it is possible to find a set W encoding a sufficient amount of local information about a given frame.

Definition 5.31. Let \mathcal{A} be a linear order with a guard $\Gamma_0, \dots, \Gamma_{p-1}$, let $D \subseteq A$ and $S \subseteq \wp(D)$.

- (a) We say that a set $X \subseteq D$ is an \mathcal{S} -stratum if, for every non-empty open set U , there are $S \in \mathcal{S}$ and a non-empty open set $V \subseteq U$ with $V \cap X = V \cap S$.
- (b) \mathcal{S} is a *skeleton* for D if there exists a set $\mathcal{P} \subseteq \wp(D)$ with the following properties.

- ♦ \mathcal{P} is a countably infinite partition of D such that every $P \in \mathcal{P}$ is dense in D .
- ♦ $\Gamma_0 \cup \dots \cup \Gamma_{p-1}$ can be written as a union of guarded sets in \mathcal{P} .
- ♦ Every $S \in \mathcal{S}$ is a union of classes in \mathcal{P} , each $S \in \mathcal{S}$ meets some Γ_i , and the family $\{S \cap \Gamma_* \mid S \in \mathcal{S}\}$ is disjoint.
- ♦ $\llbracket S = T \rrbracket = \emptyset$, for all $S \neq T$ in \mathcal{S} .
- ♦ For every \mathcal{S} -stratum $X \subseteq A$, the set $\bigcup_{S \in \mathcal{S}} \llbracket S = X \rrbracket$ is dense in A .

Theorem 5.32. There exists an MSO-formula $\text{Code}(X; D, W, \tilde{T})$ with the following properties. Suppose we are given

- ♦ a short zero-dimensional linear order \mathcal{A} ,
- ♦ a guard \tilde{T} for A such that each Γ_i is countable,
- ♦ a countable set $D \subseteq A$ containing $\Gamma_* := \Gamma_0 \cup \dots \cup \Gamma_{p-1}$ such that $D \setminus \Gamma_*$ is dense in A ,

- ♦ a skeleton $\mathcal{S} \subseteq \wp(D)$ for D .

Then there exists a set $W \subseteq A \setminus D$ such that, for every $X \subseteq D$,

$$\mathfrak{A} \models \text{Code}(X; D, W, \tilde{r}) \quad \text{iff} \quad X \subseteq D \text{ is dense in } D, \text{ guarded, and an } \mathcal{S}\text{-stratum.}$$

Proof. By Lemma 5.21, D has a frame φ such that

- ♦ the sets $(\varphi^{-1}[P])_{P \in \mathcal{P}}$ form a stratification of \mathbb{Z}^* and
- ♦ $\text{len}(\varphi^{-1}[I_i]) = \{n(p+1) + i \mid n < \omega\}$.

Let

$$\llbracket w \rrbracket := \{a \in A \mid \varphi(u) \leq a \leq \varphi(v), u, v \geq_{\text{pf}} w\}$$

be the set from Lemma 5.15.

We say that a set $X \subseteq D$ has the *colour* $S \in \mathcal{S}$ if $\varphi^{-1}[X]$ has colour $\text{len}(S)$. Similarly, we call X *entirely colourless* if $\varphi^{-1}[X]$ is entirely \mathcal{S}' -colourless where $\mathcal{S}' := \{\text{len}(\varphi^{-1}[S]) \mid S \in \mathcal{S}\}$.

Below we will prove the following claims.

- (I) A set $X \subseteq D$ is coherent and guarded if, and only if, $\varphi^{-1}[X]$ is \tilde{L} -embedded where $L_q := \text{len}(\varphi^{-1}[I_q])$.
- (II) There exists a set $W \subseteq A \setminus D$ such that, for every coherent, guarded set $X \subseteq D$,
 - ♦ $|\text{cl}_2(X) \cap W| \leq 1$, if X has some colour in \mathcal{S} , and
 - ♦ $\text{cl}_2(X) \cap W \neq \emptyset$, if X is entirely colourless.
- (III) Let $X \subseteq D$ be a guarded set and $U_o \subseteq A$ a non-empty open set such that X is dense in U_o , and suppose that every non-empty open subset $U' \subseteq U_o$ meets every set of the form $X \setminus S$, for $S \in \mathcal{S}$. Then there exist a non-empty open subset $V \subseteq U_o$ and sets $X_i \subseteq I_i \cap X$, for $i < p$, and $X_p \subseteq X$ such that X_i is dense in V , for all $i < p+1$, and every $S \in \mathcal{S}$ is disjoint from one of the sets $V \cap X_i$, $i < p+1$.
- (IV) There exists an MSO-formula $\text{Code}_o(X, U; D, W, \tilde{r})$ (that does not dependent on \mathfrak{A}) with the following properties. Given \mathfrak{A} , D , \tilde{r} , \mathcal{P} ,

and \mathcal{S} as above, there exists a set $W \subseteq A \setminus D$ such that, for non-empty open $U \subseteq A$ and every guarded $X \subseteq D$ such that X is dense in U ,

$$\mathfrak{A} \models \text{Code}_o(X, U; D, W, \tilde{r})$$

if, and only if, for every non-empty open subset $U_o \subseteq U$, there are $S \in \mathcal{S}$ and a non-empty open set $V \subseteq U_o$ with $V \cap X \subseteq S$.

Before giving the proofs, let us show how the statement of the theorem follows. Let $\text{Code}(X; D, W, \tilde{r})$ be the MSO-formula stating that

“ $X \subseteq D$ is dense in D and guarded, the formula $\text{Code}_o(X, A; D, W, \tilde{r})$ holds, and there is no non-empty open set U_o and no $Y \subseteq U_o \setminus X$ such that Y is dense in U_o and $\text{Code}_o(X \cup Y, U_o; D, W, \tilde{r})$ holds.”

where $\text{Code}_o(X, U; D, W, \tilde{r})$ is the formula from (IV). We claim that $\text{Code}(X; D, W, \tilde{r})$ has the desired properties. Let W be the set from (IV).

(\Leftarrow) Let $X \subseteq D$ be guarded, dense in D , and an \mathcal{S} -stratum. Then $\text{Code}_o(X, A; D, W, \tilde{r})$ holds. Let U_o be open and $Y \neq \emptyset$ dense in U_o and disjoint from X . Since X is an \mathcal{S} -stratum, we can find an open set $V \subseteq U_o$ and some $S \in \mathcal{S}$ such that $X \cap V = S \cap V$. As Y is dense in U_o , it follows that

$$(X \cup Y) \cap V_o \supset X \cap V_o = S \cap V_o, \quad \text{for all } V_o \subseteq V.$$

Consequently, $\text{Code}_o(X \cup Y, U_o; D, W, \tilde{r})$ does not hold.

(\Rightarrow) Suppose that $\text{Code}(X; D, W, \tilde{r})$ holds. Then $X \subseteq D$ is dense in D and guarded. To show that it is also an \mathcal{S} -stratum, fix an open set $U \subseteq A$. Since $\text{Code}_o(X, A; D, W, \tilde{r})$ holds, we can find a set $S \in \mathcal{S}$ and an open subset $V \subseteq U$ such that $V \cap X \subseteq S$. If there exists an open subset $V_o \subseteq V$ such that $V_o \cap X = V_o \cap S$, we are done. For a contradiction, suppose otherwise. Then the set $Y := (S \setminus X) \cap V$ is dense in V , and the formula $\text{Code}_o(X \cup Y, V; D, W, \tilde{r})$ holds. A contradiction.

To conclude the proof it remains to establish the four claims above.

(I) (\Leftarrow) Let $\varphi^{-1}[X]$ be \tilde{L} -embedded. For coherence, fix $x \in X$. Since there are infinitely many $d \geq 0$ with $\uparrow x d \cap \varphi^{-1}[X] \neq \emptyset$, we have

$$x = \inf \{z \in X \mid z \geq x\}.$$

The other bound follows analogously.

To show that X is guarded, fix an index $i < p$ and elements $u, v \in \varphi^{-1}[X]$ with $\varphi(u) < \varphi(v)$. Then $u \triangleleft v$. As $\varphi^{-1}[X]$ is \bar{L} -embedded, we can find some $w \in \varphi^{-1}[X]$ and some $d < 0$ with $w \geq_{\text{pf}} ud$ and $|w| \in L_i$. This implies that $u \triangleleft w \triangleleft v$. Hence, $\varphi(u) < \varphi(w) < \varphi(v)$ and $\varphi(w) \in \Gamma_i$.

(\Rightarrow) Let X be coherent and guarded. Then $\varphi^{-1}[X]$ is non-empty. For the second condition, fix $x \in \varphi^{-1}[X]$ and $d \geq 0$. Since

$$\varphi(x) = \inf \{ \varphi(xk) \mid k \geq 0 \},$$

it follows by coherence that there is some $k \geq d$ with $\uparrow xk \cap \varphi^{-1}[X] \neq \emptyset$. By symmetry, there are also infinitely many $k < 0$ with $\uparrow xk \cap \varphi^{-1}[X] \neq \emptyset$.

For the third condition, fix $w \in \varphi^{-1}[X]$ and $i < p$. Since X is guarded and we have seen in Lemma 5.15 that the set $\llbracket w \rrbracket$ is open, we can find some element $a \in \llbracket w \rrbracket \cap \Gamma_i \cap X$ with $a \neq \varphi(w)$. Set $u := \varphi^{-1}(a)$. Then $u \geq_{\text{pf}} w$ and $|u| \in L_i$.

(II) Let $(X_i)_{i < 2^{\aleph_0}}$ be an enumeration of all coherent, guarded, entirely colourless subsets of D . By induction on i , we construct a sequence $(w_i)_{i < 2^{\aleph_0}}$ of points $w_i \in \text{cl}_2(X_i) \setminus D$ such that the set

$$W := \{ w_i \mid i < 2^{\aleph_0} \}$$

has the desired properties.

Thus, let $i < 2^{\aleph_0}$ and suppose that we have already chosen w_j , for all $j < i$. Given $a \in \text{cl}_2(D) \setminus D$ and $n < \omega$, we denote by $a \upharpoonright n$ the unique vertex $v \in \mathbb{Z}^*$ such that $a \in \llbracket v \rrbracket$ and $|v| = n$. To see that this is well-defined, let us show by induction on n , that the set

$$\{ \llbracket v \rrbracket \mid |v| = n \}$$

forms a partition of $\text{cl}_2(D) \setminus \varphi[\mathbb{Z}^{<n}]$. For $n = 0$, we have $\text{cl}_2(D) = \llbracket \langle \rangle \rrbracket$ since D is dense in $\text{cl}_2(D)$. For the inductive step, it is sufficient to note that, since each set $\llbracket vd \rrbracket$ is closed by choice of φ , there are no elements a with $\llbracket vd \rrbracket < a < \llbracket v(d+1) \rrbracket$ (for $d \neq -1$). Consequently, we have

$$\llbracket v \rrbracket = \{ \varphi(v) \} \cup \bigcup_{d \in \mathbb{Z}} \llbracket vd \rrbracket, \quad \text{for all } v \in \mathbb{Z}^n,$$

and the claim follows by inductive hypothesis.

To choose the next element w_i , we first construct a frame embedding $\mu : \mathbb{Z}^* \rightarrow \mathbb{Z}^*$ with such that

$$\varphi(\mu(w)) \in X_i \cap \Gamma_{|w| \bmod p}, \quad \text{for all } w \in \mathbb{Z}^*.$$

We start with $\mu(\langle \rangle) := v$, for an arbitrary $v \in \varphi^{-1}[X_i \cap \Gamma_0]$. For the inductive step, suppose that we have already defined $\mu(w)$. As X_i is coherent and $\Gamma_{(|w|+1) \bmod p}$ is dense in X_i , we can find directions

$$\dots < k_{-2} < k_{-1} < 0 \leq k_0 < k_1 < \dots$$

and vertices $v_d \geq_{\text{pf}} \mu(w)k_d$ such that

$$\varphi(v_d) \in X_i \cap \Gamma_{(|w|+1) \bmod p}.$$

We set $\mu(wd) := v_d$.

Having defined μ , let $\psi := \varphi \circ \mu$. Below we will prove that

$$|\text{rng } \psi \cap Y_j(S)| \leq 1, \quad \text{for all } j < i \text{ and } S \in \mathcal{S},$$

where

$$Y_j(S) := \left\{ y \in \text{cl}_2(D) \setminus D \mid \text{for every } n < \omega, \{w_j \upharpoonright n, y \upharpoonright n\} \text{ has colour } S \right\}.$$

Setting $Y_j := \bigcup_{S \in \mathcal{S}} Y_j(S)$ we can then choose

$$\begin{aligned} w_i &\in \text{cl}_2(X_i) \setminus \left(D \cup \bigcup_{j < i} Y_j \right) \\ &\supseteq \text{cl}_2(\text{rng } \psi) \setminus \left(D \cup \bigcup_{j < i} Y_j \right) \\ &= \text{cl}_2(\text{rng } \psi) \setminus \left(D \cup \bigcup_{j < i} (Y_j \cap \text{rng } \psi) \right) \end{aligned}$$

since $D \cup \bigcup_{j < i} (Y_j \cap \text{rng } \psi)$ has size at most $\aleph_0 + |i| \cdot |\mathcal{S}| < 2^{\aleph_0}$, while we have $|\text{cl}_2(\text{rng } \psi)| = 2^{\aleph_0}$ by Corollary 5.25. (Note that \mathcal{S} is countable since, by definition of a skeleton, there exists a surjective map $\Gamma_* \rightarrow \mathcal{S}$.)

To prove that the resulting set $W := \{w_i \mid i < 2^{\aleph_0}\}$ has the desired properties note that, by choice of w_i , we have $\text{cl}_2(X_i) \cap W \neq \emptyset$, for all i . To show that $|\text{cl}_2(X) \cap W| \leq 1$, for every set X with a colour $S \in \mathcal{S}$, suppose for a contradiction that there are indices $j < i$ with $w_i, w_j \in \text{cl}_2(X)$. Then $w_i \notin Y_j$ implies that there is some $n < \omega$ such that the set $\{w_i \upharpoonright n, w_j \upharpoonright n\}$ has no colour. In particular,

$$\{k < n \mid (w_i \upharpoonright n)(k) \neq (w_j \upharpoonright n)(k)\} \notin S.$$

Furthermore, $w_i, w_j \in \text{cl}_2(X)$ implies that $w_i \upharpoonright k, w_j \upharpoonright l \in X$, for some $k, l \geq n$. Hence, X does not have colour S . A contradiction.

It remains to prove the above claim. For a contradiction, suppose that there are two distinct elements $y, y' \in \text{cl}_2(\text{rng } \psi) \cap Y_j(S)$. Let $n < \omega$ be the minimal length such that $y \upharpoonright n \neq y' \upharpoonright n$. Since $y \in \text{cl}_2(\text{rng } \psi)$, there is some $v \in \mathbb{Z}^*$ with $\psi(v) \in \llbracket y \upharpoonright n \rrbracket$, which implies that $\mu(v) \geq_{\text{pf}} y \upharpoonright n$. In the same way we find some $v' \in \mathbb{Z}^*$ with $\mu(v') \geq_{\text{pf}} y' \upharpoonright n$. Since $\psi(v), \psi(v') \in X_i$, the vertices $\mu(v)$ and $\mu(v')$ are incomparable, and X_i is entirely colourless, it follows that $\{\psi(v), \psi(v')\}$ has no colour. In particular, $n - 1 \notin \text{len}(S)$. But $y, y' \in Y_j(S)$ implies that the sets

$$\{w_j \upharpoonright n, y \upharpoonright n\} \quad \text{and} \quad \{w_j \upharpoonright n, y' \upharpoonright n\}$$

have colour S . Since at least one of these two sets varies at $n - 1$, it follows that $n - 1 \in \text{len}(S)$. A contradiction.

(III) First, suppose that there exist sets $S \in \mathcal{S}$, $V \subseteq U_0$, and an index $q < p$ such that U_0 is open and $S \cap \Gamma_q \cap X$ is dense in V . We set

$$X_p := X \setminus S, \quad X_q := S \cap \Gamma_q \cap X, \quad X_i := \Gamma_i \cap X, \quad \text{for } i \neq q, p.$$

Note that, for $i \neq q, p$, the set X_i is dense in V since X is guarded and X is dense in V . Furthermore, $X_p \cap V$ is dense in V because every non-empty open $V \subseteq V$ contains some element of $X \setminus S$. Finally, S is disjoint from $V \cap X_p$, while all other elements of \mathcal{S} are disjoint from X_q .

It remains to consider the case where every set of the form $S \cap U_0 \cap \Gamma_i \cap X$ with $S \in \mathcal{S}$ and $i < p$ is nowhere dense in U_0 . We start by constructing a sequence $(x_n)_{n < \omega}$ of elements of $U_0 \cap X$ with the following properties.

- ♦ For every $i < p + 1$, the set $\{x_n \mid n \equiv i \pmod{p+1}\}$ is dense in U_0 .
- ♦ $n \equiv i \pmod{p+1}$ implies $x_n \in \Gamma_i$, for $i < p$.
- ♦ $x_i \in S$ implies $x_j \notin S$, for $S \in \mathcal{S}$ and $i < j < \omega$.

Fix a countable basis $(B_l)_{l < \omega}$ of the topology of A and let $I \subseteq \omega$ be the set of all indices l with $B_l \subseteq U_0$. Suppose we have already chosen x_0, \dots, x_{n-1} and let $S_0, \dots, S_{m-1} \in \mathcal{S}$ be the sets containing some x_i . (Since the restrictions of the sets in \mathcal{S} to each Γ_i are disjoint, there are at most n such sets.) To define x_n , set $i := n \bmod (p+1)$. If $i < p$, we choose some element $x_n \in B_n \cap \Gamma_i \cap X \setminus (S_0 \cup \dots \cup S_{m-1})$ (which is possible since the sets $S_j \cap U_0 \cap \Gamma_i \cap X$ are nowhere dense in U_0). If $i = p$, we choose some $x_n \in B_n \cap X \setminus (S_0 \cup \dots \cup S_{m-1})$.

Having defined $(x_n)_n$, we claim that the sets

$$V := U_0 \quad \text{and} \quad X_i := \{x_n \mid n \equiv i \pmod{p+1}\}, \quad \text{for } i \leq p,$$

have the desired properties. By construction, X_i is dense in $U_0 = V$. Furthermore, given $S \in \mathcal{S}$, there is at most one index $n < \omega$ with $x_n \in S$. Hence, S is disjoint from X_i , for all $i \neq n \pmod{p+1}$.

(IV) Let $\text{Code}_0(X, U; D, W, \bar{I})$ be an MSO-formula stating that

“For every non-empty open set $V \subseteq U$ and all sets $X_i \subseteq \Gamma_i \cap X$, $i < p$, and $X_p \subseteq X$ such that X_p is dense in V , there exists a coherent set $Y \subseteq D \cap V$ such that X_0, \dots, X_p are dense in Y and W contains at most one element of $\text{cl}_2(Y)$.”

We start by proving that, if W is the set from (II), $U \subseteq A$ is non-empty and open, and $X \subseteq D$ is a guarded set such that X is dense in U , then

$$\mathfrak{A} \models \text{Code}_0(X, U; D, W, \bar{I})$$

if, and only if, for every non-empty open subset $U_0 \subseteq U$, there are $S \in \mathcal{S}$ and a non-empty open set $V \subseteq U_0$ with $V \cap X \subseteq S$.

(\Leftarrow) Let V, X_0, \dots, X_p be sets as in the formula $\text{Code}_o(X, U; D, W, \tilde{F})$. We start by showing that there exists some $w \in \mathbb{Z}^*$ such that $\llbracket w \rrbracket \subseteq V$ and $\llbracket w \rrbracket \cap X \subseteq S$, for some $S \in \mathcal{S}$. By assumption, we can find a non-empty open set $V_o \subseteq V$ with $V_o \cap X \subseteq S$, for some $S \in \mathcal{S}$. Since D is dense, there is some $w \in \mathbb{Z}^*$ with $\varphi(w) \in V_o$. As

$$\varphi(w) = \sup_{d \geq 0} \varphi(wd),$$

we can find some $d_o \in \mathbb{Z}$, $d \geq 0$, with $\varphi(wd) \in U$, for all $d \geq d_o$. In particular, $\llbracket w(d_o + 1) \rrbracket \subseteq V_o$.

Thus, replacing V by $\llbracket w \rrbracket$ we may assume that $V = \llbracket w \rrbracket$, for some $w \in \mathbb{Z}^*$, and that $V \cap X \subseteq S$, for some $S \in \mathcal{S}$. Let $\mu : \mathbb{Z}^* \rightarrow \varphi^{-1}[V \cap D]$ be the frame embedding from Lemma 5.29. To see that we can apply this lemma note that

- $\varphi^{-1}[V] = \uparrow w$ is upwards closed,
- the sets $L_q := \{n(p+1) + q \mid n < \omega\}$, $q < p+1$, are infinite and disjoint,
- $\text{len}(\varphi^{-1}[V \cap X]) \subseteq \text{len}(S) \subseteq \omega = L_0 \cup \dots \cup L_p$,
- for all $u \in \varphi^{-1}[V]$ and $q < p$, there is some $v \in \varphi^{-1}[X]$ with $v >_{\text{pf}} u$ and $|v| \equiv q \pmod{p+1}$.

(For the last property, note that X is guarded and X is dense in V . Hence, each X_q is dense in X and X_q is dense in V .)

We claim that the image $Y := \varphi[\text{rng } \mu]$ is the desired set. Coherence of Y follows by Lemmas 5.23 and 5.16. Furthermore, for $q < p$, $X_q \cap Y = X \cap \Gamma_q \cap Y$ is dense in Y . Similarly, $X_p \cap Y \supseteq X \cap \Gamma_p \cap Y$ is dense in Y . Finally, we have $|\text{cl}_2(X) \cap W| \leq 1$ by (II) and the fact that Y has colour S .

(\Rightarrow) Suppose that there is some open $U_o \subseteq U$ such that every non-empty open subset $V \subseteq U_o$ meets every set of the form $X \setminus S$, for $S \in \mathcal{S}$. Let V and X_0, \dots, X_p be sets from (III) and let $Y \subseteq D \cap V$ be a coherent set such that each X_i is dense in Y . To prove that $\text{Code}_o(X, U; D, W, \tilde{F})$ does not hold, it is sufficient to show that $|\text{cl}_2(Y) \cap W| \geq 2$.

We call a set Y *nowhere coloured* if there is no open set U such that $Y \cap U$ is non-empty and has a colour $S \in \mathcal{S}$. We start by proving that Y is nowhere coloured. Let $S \in \mathcal{S}$ and let G be an open set with $G \cap Y \neq \emptyset$. By choice of

V and X_i , there is some $i < p+1$ such that $V \cap X_i \cap S = \emptyset$. Since $Y \subseteq V$ is coherent, G is open, and X_i is dense in V , we have $G \cap X_i \cap Y \neq \emptyset$. Fix some $w \in \varphi^{-1}[G \cap X_i \cap Y]$. As Y is coherent, we have

$$\varphi(w) = \sup \{ \varphi(wd) \mid d \geq 0 \}.$$

Consequently, there are arbitrarily large $d \geq 0$ with $\uparrow xd \cap \varphi^{-1}[G \cap Y] \neq \emptyset$. By symmetry, there are also infinitely many $d < 0$ with $\uparrow xd \cap \varphi^{-1}[G \cap Y] \neq \emptyset$. In particular, the set $\varphi^{-1}[G \cap Y]$ varies at level $|w|$. But $\varphi(w) \in V \cap X_i$ and $V \cap X_i \cap S = \emptyset$ implies that $|w| \notin \text{len}(S)$. Hence, $G \cap Y$ does not have colour S .

Note that Y is guarded since $X_i \subseteq \Gamma_i$ is dense in Y , for every $i < p$. It therefore follows by (I) that the preimage $\varphi^{-1}[Y]$ is \tilde{L} -embedded. Furthermore, as Y is nowhere coloured, no set of the form $\varphi^{-1}[Y] \cap \uparrow u$ with $u \in \mathbb{Z}^*$ has a colour $S \in \mathcal{S}'$ where

$$\mathcal{S}' := \{ \text{len}(\varphi^{-1}[S]) \mid S \in \mathcal{S} \}.$$

It follows that we can use Lemma 5.30 to find a subset $Z \subseteq \varphi^{-1}[Y]$ that is \tilde{L} -embedded and entirely \mathcal{S}' -colourless. Fix two \leq_{pf} -incomparable elements $z_o, z_1 \in Z$ and set $Z_i := \uparrow z_i \cap Z$. Then Z_o and Z_1 are disjoint, \tilde{L} -embedded, and entirely \mathcal{S}' -colourless and $\text{cl}_2(\varphi[Z_o])$ and $\text{cl}_2(\varphi[Z_1])$ are disjoint. Consequently, (I) and (II) imply that

$$\text{cl}_2(\varphi[Z_i]) \cap W \neq \emptyset.$$

It follows that

$$\begin{aligned} |\text{cl}_2(Y) \cap W| &\geq |\text{cl}_2(\varphi[Z]) \cap W| \\ &\geq |\text{cl}_2(\varphi[Z_o]) \cap W| + |\text{cl}_2(\varphi[Z_1]) \cap W| \geq 2. \quad \square \end{aligned}$$

Arithmetic

In order to use the formula Code from the preceding section to encode arithmetic in a non-modest order, we have to deal with the fact that we can check for equality with a set in \mathcal{S} only locally in some open interval.

Definition 5.33. Let \mathfrak{A} be a linear order.

(a) Given a formula $\varphi(\tilde{X})$ and values \tilde{P} for the variables \tilde{X} , we say that $\varphi(\tilde{P})$ holds locally at an element $a \in A$ if there exists an open neighbourhood U of a such that

$$\mathfrak{A} \models \varphi^{(U)}(\tilde{P}),$$

where $\varphi^{(U)}$ denotes the relativisation of φ to U .

(b) We write

$$\llbracket \varphi(\tilde{P}) \rrbracket := \{ a \in A \mid \varphi \text{ holds locally at } a \},$$

$$\mathfrak{A} \models_{\text{loc}} \varphi(\tilde{P}) \quad \text{iff} \quad \llbracket \varphi(\tilde{P}) \rrbracket \text{ is dense in } A.$$

Definition 5.34. Let \mathfrak{A} be a linear order and $W, D, \Gamma_0, \dots, \Gamma_{p-1} \subseteq A$ parameters, and $\mathcal{S} \subseteq \wp(D)$. We set $\Gamma_* := \Gamma_0 \cup \dots \cup \Gamma_{p-1}$.

A tuple $\langle D, C^\leq, C^+, C^\times \rangle$ of sets is a *code* if it satisfies the following conditions.

- ♦ $\Gamma_* \subseteq D$, the complement $D \setminus \Gamma_*$ is dense in A , and $W \cap D = \emptyset$.
- ♦ $C^\leq, C^+, C^\times \subseteq D \setminus \Gamma_*$ are dense in A .
- ♦ For every pair of \mathcal{S} -strata S and T ,

$$\mathfrak{A} \models_{\text{loc}} (S = T) \vee (S \cap T \cap \Gamma_* = \emptyset),$$

$$\mathfrak{A} \models_{\text{loc}} (S = T) \vee (S \cap C^\leq \subset Y) \vee (T \cap C^\leq \subset S).$$

The main idea of our encoding is contained in the following lemma.

Lemma 5.35. Let \mathfrak{A} be a short zero-dimensional separable linear order and let $\Gamma_0, \dots, \Gamma_{p-1}$ be a guard for A . There exists a code $\langle D, C^\leq, C^+, C^\times \rangle$ and a set $W \subseteq A \setminus D$ encoding a skeleton $\mathcal{S} = \{ S_n \mid n < \omega \}$ such that, for all $i < j < k < \omega$,

$$\begin{aligned} i + j = k & \quad \text{iff} \quad S_i \cap S_j \cap S_k \cap C^+ \text{ is dense in } A \\ & \quad \text{iff} \quad S_i \cap S_j \cap S_k \cap C^+ \neq \emptyset, \end{aligned}$$

$$\begin{aligned} i \cdot j = k & \quad \text{iff} \quad S_i \cap S_j \cap S_k \cap C^\times \text{ is dense in } A \\ & \quad \text{iff} \quad S_i \cap S_j \cap S_k \cap C^\times \neq \emptyset. \end{aligned}$$

Furthermore, for $i, j < \omega$, we have

$$\begin{aligned} i \leq j & \quad \text{iff} \quad \llbracket S_i \cap C^\leq \subseteq S_j \rrbracket \text{ is dense in } A \\ & \quad \text{iff} \quad \llbracket S_i \cap C^\leq \subseteq S_j \rrbracket \neq \emptyset. \end{aligned}$$

Proof. Set $\Gamma_* := \Gamma_0 \cup \dots \cup \Gamma_{p-1}$. As \mathfrak{A} is separable, there exists a countable set $D' \subseteq A \setminus \Gamma_*$ that is dense in A . Let $D := \Gamma_* \cup D'$.

For every $i < p$, we can use Lemma 1.8 to choose a partition $(P_n^i)_{n < \omega}$ of Γ_i into countably infinitely many classes. Set

$$\mathcal{P}^0 := \{ P_n^0 \cup \dots \cup P_n^{p-1} \mid n < \omega \}.$$

We claim that each $P \in \mathcal{P}^0$ is guarded and dense in A . For denseness, note that P is dense in Γ_0 and Γ_0 is dense in A . For guardedness, note that $\Gamma_i \cap P = P_n^i$, for some n . Furthermore, P_n^i is dense in Γ_i , which is dense in A . Hence, P_n^i is also dense in P .

Similarly, let \mathcal{P}' be a partition of D' into countably infinitely many classes that each are dense in D' . We choose a partition $\mathcal{P}^\leq \cup \mathcal{P}^+ \cup \mathcal{P}^\times$ of \mathcal{P}' into three infinite parts. Let

$$(C_n^0)_{n < \omega}, \quad (C_n^\leq)_{n < \omega}, \quad (C_{ij}^+)_{i, j < \omega}, \quad (C_{ij}^\times)_{i, j < \omega}$$

be enumerations (without repetitions) of, respectively, \mathcal{P}^0 , \mathcal{P}^\leq , \mathcal{P}^+ , and \mathcal{P}^\times , set

$$C^\leq := \bigcup_{n < \omega} C_n^\leq, \quad C^+ := \bigcup_{i, j < \omega} C_{ij}^+, \quad C^\times := \bigcup_{i, j < \omega} C_{ij}^\times,$$

and let $\mathcal{S} := \{ S_n \mid n < \omega \}$ where

$$\begin{aligned} S_n := C_n^0 \cup \bigcup_{i < n} C_i^\leq \cup \bigcup \{ C_{ij}^+ \mid i = n \text{ or } j = n \text{ or } i + j = n \} \\ \cup \bigcup \{ C_{ij}^\times \mid i = n \text{ or } j = n \text{ or } i \cdot j = n \}. \end{aligned}$$

For $i < j < k$, it follows that

$$S_i \cap S_j \cap S_k \cap C^+ = C_{ij}^+ \cap C_{ik}^+ \cap C_{jk}^+,$$

which is either dense or empty, depending on whether or not $k = i + j$. An analogous statement holds for $S_i \cap S_j \cap S_k \cap C^\times$ and for $S_i \cap C^\leq \subset S_j$.

Set $\mathcal{P} := \{C_n^\circ, C_n^\leq, C_{ij}^+, C_{ij}^\times \mid i, j, n < \omega\}$ and let $W \subseteq A \setminus D$ be the set from Theorem 5.32.

To prove that \mathcal{S} is a skeleton, we have to check several conditions. The following properties follow immediately from the construction.

- ◆ Every $S \in \mathcal{S}$ is trivially an \mathcal{S} -stratum.
- ◆ \mathcal{P} is a countably infinite partition of D .
- ◆ Every $P \in \mathcal{P}$ is dense in D .
- ◆ Γ_* can be written as a union of guarded sets in \mathcal{P} .
- ◆ Every $S \in \mathcal{S}$ is a union of classes in \mathcal{P} .
- ◆ Each $S \in \mathcal{S}$ meets some Γ_i .
- ◆ The family $\{S \cap \Gamma_* \mid S \in \mathcal{S}\}$ is disjoint.

Let us give more detailed proofs of the two remaining conditions.

Suppose that $a \in \llbracket S = T \rrbracket$. Then there exists a non-empty open set $U \subseteq A$ with $S \cap U = T \cap U$. Since each $P \in \mathcal{P}$ is dense in A , it follows that $S \cap P = T \cap P$, which implies that $S = T$.

Fix an \mathcal{S} -stratum X . For a contradiction, suppose that there exists a non-empty open set U disjoint from $\bigcup_{S \in \mathcal{S}} \llbracket S = X \rrbracket$. Since X is an \mathcal{S} -stratum, there exists a non-empty open set $V \subseteq U$ and some $S \in \mathcal{S}$ such that

$$X \cap V = S \cap V.$$

This implies that $V \subseteq \bigcup_{S \in \mathcal{S}} \llbracket S = X \rrbracket$. A contradiction.

We also have to show that $\langle D, C^\leq, C^+, C^\times \rangle$ is a code. By choice of the sets D, C^\leq, C^+, C^\times , we have $\Gamma_* \subseteq D$, $D \setminus \Gamma_*$ is dense in A , and $C^\leq, C^+, C^\times \subseteq D \setminus \Gamma_*$ are dense in A . Furthermore, $W \cap D = \emptyset$ holds by Theorem 5.32.

It therefore remains to show that

$$\begin{aligned} \mathfrak{A} \models_{\text{loc}} (S = T) \vee (S \cap T \cap \Gamma_* = \emptyset), \\ \mathfrak{A} \models_{\text{loc}} (S = T) \vee (S \cap C^\leq \subset T) \vee (T \cap C^\leq \subset S), \end{aligned}$$

for every pair of \mathcal{S} -strata S and T . Fix two \mathcal{S} -strata S, T and let $a \in A$. We claim that

$$\begin{aligned} a \in \llbracket (S = T) \vee (S \cap T \cap \Gamma_* = \emptyset) \rrbracket, \\ a \in \llbracket (S = T) \vee (S \cap C^\leq \subset T) \vee (T \cap C^\leq \subset S) \rrbracket. \end{aligned}$$

Since S and T are \mathcal{S} -strata, there exists an open neighbourhood U of a such that $S \cap U = S_m \cap U$ and $T \cap U = S_n \cap U$, for some $m, n < \omega$. If $m = n$, we have

$$\begin{aligned} U \subseteq \llbracket (S = T) \vee (S \cap T \cap \Gamma_* = \emptyset) \rrbracket, \\ U \subseteq \llbracket (S = T) \vee (S \cap C^\leq \subset T) \vee (T \cap C^\leq \subset S) \rrbracket. \end{aligned}$$

If $m \neq n$, say, $m < n$, then

$$\begin{aligned} U \cap (S \cap T \cap \Gamma_*) &= C_m^\circ \cap C_n^\circ = \emptyset, \\ U \cap (S \cap C^\leq) &= C_0^\leq \cup \dots \cup C_{m-1}^\leq \subseteq T, \end{aligned}$$

which also implies that

$$\begin{aligned} U \subseteq \llbracket (S = T) \vee (S \cap T \cap \Gamma_* = \emptyset) \rrbracket, \\ U \subseteq \llbracket (S = T) \vee (S \cap C^\leq \subset T) \vee (T \cap C^\leq \subset S) \rrbracket. \quad \square \end{aligned}$$

Corollary 5.36. *Let \mathfrak{A} be a short zero-dimensional separable linear order and let $\tilde{\Gamma}$ be a guard for A . For every FO-formula $\varphi(\bar{x})$ over the signature $\{+, \cdot, <\}$, there exists an MSO-formula $\varphi^*(\bar{X})$ (with parameters $D, \bar{C}, W, \tilde{\Gamma}$) such that*

$$\langle \mathbb{N}, +, \cdot, < \rangle \models \varphi(k_0, \dots, k_{n-1}) \quad \text{iff} \quad \mathfrak{A} \models_{\text{loc}} \varphi^*(S_{k_0}, \dots, S_{k_{n-1}}),$$

where $(S_n)_n$ are the sets from Lemma 5.35.

Proof. Given an FO-formula $\varphi(\bar{x})$, we will inductively construct an MSO-definable set $\varphi^*(\bar{X}) \subseteq A$ with the stronger property that

$$\begin{aligned} \langle \mathbb{N}, +, \cdot \rangle \models \varphi(\bar{k}) \quad \text{implies} \quad \varphi^*(\bar{S}) \text{ is dense in } A, \\ \langle \mathbb{N}, +, \cdot \rangle \not\models \varphi(\bar{k}) \quad \text{implies} \quad \varphi^*(\bar{S}) = \emptyset. \end{aligned}$$

W.l.o.g. we may assume that all occurrences of addition and multiplication in φ are in a subformula of the form

$$x < y < z \wedge x + y = z \quad \text{or} \quad x < y < z \wedge x \cdot y = z.$$

Furthermore, we can replace every equality $x = y$ between variables, by $x \leq y \wedge y \leq x$. Then we can set

$$\begin{aligned} (x \leq y)^* &:= \llbracket X \cap C^{\leq} \subseteq Y \rrbracket, \\ (x < y < z \wedge x + y = z)^* &:= (x < y)^* \cap (y < z)^* \\ &\quad \cap \llbracket X \cap Y \cap Z \cap C^+ \rrbracket, \\ (x < y < z \wedge x \cdot y = z)^* &:= (x < y)^* \cap (y < z)^* \\ &\quad \cap \llbracket X \cap Y \cap Z \cap C^{\times} \rrbracket, \\ (\varphi \vee \psi)^* &:= \varphi^* \cup \psi^*, \\ (\neg \varphi)^* &:= A \setminus \text{cl}(\varphi^*), \\ (\exists x \varphi)^* &:= \bigcup \{ \varphi^*(X) \mid \text{Code}(X; D, W, \bar{I}) \}. \end{aligned}$$

We prove the correctness of our translation by induction on the formula. For formulae of the form $x \leq y$, the claim follows immediately by Lemma 5.35. Suppose that φ is of the form $x < y < z \wedge x + y = z$. By Lemma 5.35, it is sufficient to show that $\varphi^*(S_k, S_l, S_m)$ is dense in A if, and only if, the three sets

$$\begin{aligned} X &:= \llbracket S_k \cap C^{\leq} \subseteq S_l \rrbracket, \\ Y &:= \llbracket S_l \cap C^{\leq} \subseteq S_m \rrbracket, \\ Z &:= \llbracket X \cap Y \cap Z \cap C^+ \rrbracket \end{aligned}$$

are all dense in A . For the non-trivial direction, suppose that X , Y , and Z are dense in A . Let $U \subseteq A$ be non-empty and open. To show that

$$\varphi^*(S_k, S_l, S_m) \cap U \neq \emptyset,$$

we fix an element $x \in X \cap U$. By definition of X , there exists an open neighbourhood V_x of x with $V_x \subseteq X$. Replacing V_x by $V_x \cap U$, we may

assume that $V_x \subseteq U$. Fix an element $y \in Y \cap V_x$. Then there exists an open neighbourhood V_y of y with $V_y \subseteq Y$. Again we may assume that $V_y \subseteq V_x$. Finally, let $z \in Z \cap V_y$. Repeating the same argument one more time we obtain an open neighbourhood V_z of z with $V_z \subseteq Z \cap V_y$. It follows that $z \in X \cap Y \cap Z \cap U$.

The proof for formulae of the form $x < y < z \wedge x \cdot y = z$ is analogous.

The inductive step for boolean connectives is straightforward. It therefore remains to consider the existential quantifier. First, suppose that

$$\langle \mathbb{N}, +, \cdot \rangle \models \exists x \varphi(x).$$

By inductive hypothesis, it follows that there is some $n < \omega$ such that $\varphi^*(S_n)$ is dense in A . Hence, so is the larger set

$$(\exists x \varphi)^* = \bigcup \{ \varphi^*(X) \mid \text{Code}(X; D, W, \bar{I}) \}.$$

Conversely, suppose that $(\exists x \varphi)^*$ is dense in A . This means that, for every $a \in (\exists x \varphi)^*$, there is some \mathcal{S} -stratum T with $a \in \varphi^*(T)$. By definition of a stratum, we can find an open neighbourhood U of a such that $T \cap U = S_n \cap U$, for some $n < \omega$. Hence, $a \in \varphi^*(S_n)$, which implies that $\varphi^*(S_n) \neq \emptyset$. By inductive hypothesis, it follows that $\langle \mathbb{N}, +, \cdot \rangle \models \varphi(n)$. \square

Our undecidability proof rests on the following standard result about the undecidability of arithmetic.

Theorem 5.37. *There exists an FO-formula φ over the signature $\{+, \cdot\}$ such that $\langle \mathbb{N}, +, \cdot \rangle \models \varphi$ and every consistent set Φ of FO-formulae containing φ and closed under entailment is undecidable.*

After these preparations we can finally state and prove our undecidability result.

Theorem 5.38. *Let $0 < p < \omega$ and let \mathcal{C} be a class of short linear orders that are not p -modest. The MSO-theory of \mathcal{C} is undecidable.*

Proof. Given an FO-formula ψ , we denote by ψ^* the MSO-formula associated to ψ via Corollary 5.36. Let φ be the formula from Theorem 5.37 and

consider the theory

$$T := \{ \psi \in \text{FO} \mid \text{for all } \mathfrak{A} \in \mathcal{C}, \mathfrak{B} \subseteq \mathfrak{A}, \text{ and all choices of } D, \tilde{C}, W, \tilde{I} \text{ such that } \mathfrak{B} \text{ is zero-dimensional and separable, and } \mathfrak{B} \models_{\text{loc}} \varphi^*, \text{ we have } \mathfrak{B} \models_{\text{loc}} \psi^* \}.$$

By Lemma 5.27 and Corollary 5.36, we have

$$T \subseteq \text{Th}_{\text{FO}}(\langle \mathbb{N}, +, \cdot, < \rangle) \quad \text{and} \quad \varphi \in T.$$

In particular, T is consistent. Furthermore, T is closed under entailment. By choice of φ it follows that T is undecidable. Since we can compute T from the MSO-theory of \mathcal{C} , the latter is also undecidable. \square

Modest Orders

Having shown that non-modest orders have undecidable theories, it remains to prove the decidability statement for modest ones. Recall the notion of a quantifier structure from Definition I.4.22.

Definition 5.39. (a) In this section we will use a variant of the logic $\text{MSO}_{\tilde{n}}$, for some quantifier structure \tilde{n} , where the quantifier-free formulae can use atomic formulae of the form

$$t = \emptyset,$$

where t is a finite boolean combination of variables X_i and unary predicates P_i .

(b) Given a linear order \mathfrak{A} , we define the sets

$$\begin{aligned} \text{First} &:= \begin{cases} \{a\} & \text{if } \mathfrak{A} \text{ has a minimal element } a, \\ \emptyset & \text{otherwise,} \end{cases} \\ \text{Last} &:= \begin{cases} \{a\} & \text{if } \mathfrak{A} \text{ has a maximal element } a, \\ \emptyset & \text{otherwise.} \end{cases} \end{aligned}$$

We will prove decidability using a composition argument where the key pieces are uniform linear orders of the following form.

Definition 5.40. Let \tilde{n} be a quantifier structure. A coloured linear order $\mathfrak{A} = \langle A, \leq, \tilde{P} \rangle$ is \tilde{n} -uniform if \mathfrak{A} is densely ordered without end-points and

$$\mathfrak{A} \equiv_{\text{MSO}}^{\tilde{n}} \mathfrak{A}|_I, \quad \text{for all non-empty convex sets } I \subseteq A \text{ without end-points.}$$

Exercise 5.6. Prove that a linear order \mathfrak{A} with l parameters \tilde{P} is $\langle \rangle$ -uniform if, and only if, \mathfrak{A} is densely ordered without end-points and each of the sets

$$Q_s := \{ a \in A \mid a \in P_i \Leftrightarrow i \in s \}, \quad \text{for } s \subseteq [l],$$

is dense in A .

Using a Ramsey argument we can show that every dense linear order has a uniform suborder.

Lemma 5.41. Let \tilde{n} be a quantifier structure and $\mathfrak{A} = \langle A, \leq, \tilde{P} \rangle$ a short dense linear order. Then there exists a non-empty open interval $I \subseteq A$ such that $\mathfrak{A}|_I$ is \tilde{n} -uniform.

Proof. Let λ be the additive labelling of A defined by

$$\lambda(C) := \text{Th}_{\text{MSO}}^{\tilde{n}}(\mathfrak{A}|_C), \quad \text{for every convex set } C.$$

By Lemma 4.3, there exists an open interval $I \subseteq A$ that is λ -homogeneous. As \mathfrak{A} is short, every convex set without end-points has bicofinality $\langle \omega, \omega \rangle$. This implies that the interval I is \tilde{n} -uniform. \square

The central argument of the decidability proof of the MSO-theory of $\langle \omega, \leq \rangle$ in Section III.2 was as follows. When we quantify over a subset $P \subseteq \omega$, we can use the Theorem of Ramsey to cut the resulting order $\langle \omega, \leq, P \rangle$ into pieces that all have the same $\text{MSO}_{\tilde{n}}$ -theory. Using the composition theory for generalised sums we can then reduce the theory of $\langle \omega, \leq, P \rangle$ to that of $\langle \omega, \leq, Q \rangle$ where $Q = \emptyset$ or $Q = \omega$. Since the set Q does not provide any

information, we can drop it, thereby obtaining a reduction to $\langle \omega, \leq \rangle$ (with 1 less quantifier). In this way, every quantifier can be eliminated.

For modest orders \mathfrak{A} , a similar argument works. Each set $P \subseteq A$ induces an equivalence relation \sim on A as in Lemma 3.1. We can therefore use a composition theorem to reduce the theory of $\langle \mathfrak{A}, P \rangle$ to that of $\langle \mathfrak{A}/\sim, Q \rangle$. If we can show that this structure is $\langle \rangle$ -uniform, we can eliminate Q and use induction to compute the theory.

Definition 5.42. Let $\mathfrak{A} = \langle A, \leq, \bar{P} \rangle$ be a coloured linear order with l parameters P_0, \dots, P_{l-1} and let \bar{n} be a quantifier structure.

(a) A *congruence* is an equivalence relation \sim such that each \sim -class is convex.

(b) A *condensation* of \mathfrak{A} is a structure of the form

$$\mathfrak{C}^{\bar{n}}(I, \sim) := \langle I/\sim, \leq, (T_\theta)_{\theta \in \Theta} \rangle,$$

where I is a non-empty convex subset of A , \sim a congruence, Θ is the set of all $\text{MSO}_{\bar{n}}$ -theories, and

$$T_\theta := \{ [a] \in I/\sim \mid \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_{[a]}, \text{First}, \text{Last}) = \theta \}.$$

The following observation follows immediately from Theorem I.4.24.

Lemma 5.43. *There exists a function $q : \omega \times \omega^* \rightarrow \omega^*$ with the following properties. Given linear orders \mathfrak{A} and \mathfrak{B} with l monadic parameters, convex subsets $I \subseteq A$ and $J \subseteq B$, and congruences \sim and \approx of, respectively, I and J ,*

$$\mathfrak{C}^{\bar{n}}(I, \sim) \equiv_{\text{MSO}}^{q(l, \bar{n})} \mathfrak{C}^{\bar{n}}(J, \approx) \quad \text{implies} \quad \mathfrak{A}|_I \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{B}|_J.$$

Definition 5.44. Let $\mathfrak{A} = \langle A, \leq, \bar{P} \rangle$ be a coloured linear order with l parameters P_0, \dots, P_{l-1} , let \bar{n} be a quantifier structure, and let $q : \omega \times \omega^* \rightarrow \omega^*$ be the function from Lemma 5.43.

(a) The set $\text{UC}^{\bar{n}}(\mathfrak{A})$ of *uniform condensations* of \mathfrak{A} consists of all $q(l, \bar{n})$ -uniform linear orders of the form $\mathfrak{C}^{\bar{n}}(I, \sim)$ where I is a non-empty convex subset $I \subseteq A$ and \sim is a congruence on I .

(b) The *uniform theory* of \mathfrak{A} is the set defined by

$$\begin{aligned} \text{U}^{(\cdot)}(\mathfrak{A}) &:= \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{A}, \text{First}, \text{Last}), \\ \text{U}^{\bar{m}k}(\mathfrak{A}) &:= \{ \text{U}^{q(l+k, \bar{m})}(\mathfrak{C}) \mid \mathfrak{C} \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \}. \end{aligned}$$

We write

$$\mathfrak{A} \equiv_{\text{U}}^{\bar{n}} \mathfrak{B} \quad \text{iff} \quad \text{U}^{\bar{n}}(\mathfrak{A}) = \text{U}^{\bar{n}}(\mathfrak{B}).$$

(c) We define the (quantifier-free) *type* of an element $a \in A$ by

$$\text{tp}(a) := \{ i < l \mid a \in P_i \}.$$

For a non-empty convex subset $C \subseteq A$, we set

$$\text{tp}(C) := \langle s, t, b \rangle \in (\wp([l]) + \{\perp\}) \times (\wp([l]) + \{\perp\}) \times \{1, \infty\},$$

where

$$\begin{aligned} s &:= \begin{cases} \text{tp}(\min C) & \text{if } \min C \text{ exists,} \\ \perp & \text{otherwise,} \end{cases} \\ t &:= \begin{cases} \text{tp}(\max C) & \text{if } \max C \text{ exists,} \\ \perp & \text{otherwise,} \end{cases} \\ b &:= \begin{cases} \infty & \text{if } |C| > 1, \\ 1 & \text{if } |C| = 1. \end{cases} \end{aligned}$$

Lemma 5.45. *Let \mathfrak{A} and \mathfrak{B} be $\langle \rangle$ -uniform linear orders. Then*

$$\mathfrak{A} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{B} \quad \text{and} \quad \text{tp}(C) = \text{tp}(D) \quad \text{implies} \quad \mathfrak{A}|_C \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{B}|_D,$$

for all non-empty convex $C \subseteq A$ and $D \subseteq B$.

Proof. Since \mathfrak{A} is $\langle \rangle$ -uniform, we have

$$\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{A}|_C) = \begin{cases} \theta & \text{if } \text{tp}(C) = \langle \perp, \perp, \infty \rangle, \\ \sigma_s + \theta & \text{if } \text{tp}(C) = \langle s, \perp, \infty \rangle, \\ \theta + \sigma_t & \text{if } \text{tp}(C) = \langle \perp, t, \infty \rangle, \\ \sigma_s + \theta + \sigma_t & \text{if } \text{tp}(C) = \langle s, t, \infty \rangle, \\ \sigma_s & \text{if } \text{tp}(C) = \langle s, s, \perp \rangle, \end{cases}$$

where $\theta := \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{A})$ and σ_s denotes the $\text{MSO}_{(\cdot)}$ -theory of the 1-element structure whose unique element a has type $\text{tp}(a) = s$. \square

We start with some remarks about cuts in modest linear orders.

Definition 5.46. Let \mathfrak{A} be a linear order.

(a) A cut $\langle C, D \rangle$ of \mathfrak{A} is a partition $A = C + D$ such that C is downwards closed and D is upwards closed.

(b) A cut $\langle C, D \rangle$ is *principal* if C has a maximal element or D has a minimal one.

(c) \mathfrak{A} is *Dedekind-complete* if every cut $\langle C, D \rangle$ with $C \neq \emptyset$ and $D \neq \emptyset$ is principal.

(d) Let $X \subseteq A$ and let $\langle C, D \rangle$ be a cut of X . An element $a \in A$ *realises* the cut $\langle C, D \rangle$ if $C < a < D$. If such an element exists we say that $\langle C, D \rangle$ is *realised* in A . \downarrow

Exercise 5.7. Let \mathfrak{A} be a linear order and $X \subseteq A$ coherent. Prove that no principal cut of X is realised in A . \downarrow

Lemma 5.47. No p -modest linear order contains a convex set that is densely ordered and Dedekind-complete.

Proof. Suppose that \mathfrak{A} contains a convex set $I \subseteq A$ that is densely ordered and Dedekind-complete. Since I is densely ordered, there exists an embedding $\mathbb{Q} \rightarrow I$. As it is Dedekind complete, we can extend this embedding to one $\mathbb{R} \rightarrow I$. We have shown in Lemma 5.19 that $\langle \mathbb{R}, < \rangle$ is not p -modest. Consequently, it follows by Lemma 5.5 that neither is \mathfrak{A} . \square

Lemma 5.48. Let \mathfrak{A} be a separable dense linear order, $D \subseteq A$, and let φ be a frame for D . Then $\text{cl}_2(D)$ has only countably many non-principal cuts that are realised in A .

Proof. Since \mathfrak{A} is separable and dense, there exists a countable subset $B \subseteq A$ that is dense in A . It is sufficient to prove that every non-principal cut $\langle C, C' \rangle$ of $\text{cl}_2(D)$ that is realised in A is realised by some element of B . Hence, fix such a cut and let I be the set of all elements of A realising it. It follows by Lemma 5.3 that the set $\text{cl}_2(D)$ is replete. Consequently, I is infinite. In particular, it contains an open interval. This implies that $I \cap B \neq \emptyset$. \square

Our key argument is based on the following technical property of modest orders.

Lemma 5.49. Let \mathfrak{A} be a linear order with $l \leq \log_2 p$ parameters that is p -modest, separable, and $\langle \rangle$ -uniform, and set

$$P_s := \{ a \in A \mid \text{tp}(a) = s \}, \quad \text{for } s \subseteq [l],$$

$$H := \{ \perp \} \cup \{ s \subseteq [l] \mid P_s \neq \emptyset \}.$$

For every set $S \subseteq H \times H \times \{1, \infty\}$ such that either

- ♦ $S = \{ \langle s, s, 1 \rangle \mid P_s \neq \emptyset \}$ or
- ♦ S contains at least one element of the form $\langle s, t, \infty \rangle$ with $s, t \in H$,

there exists a congruence \sim on A such that

$$\mathfrak{C}^{(\cdot)}(A, \sim) \in \text{UC}^{(\cdot)}(\mathfrak{A}) \quad \text{and} \quad S = \{ \text{tp}([a]) \mid [a] \in A/\sim \}.$$

Proof. If S does not contain an element of the form $\langle s, t, \infty \rangle$, we can use the identity for \sim . Hence, suppose otherwise. Let

$$K := \{ s \subseteq [l] \mid \langle s, t, b \rangle \in w \text{ for some } w \in S \} \\ \cup \{ t \subseteq [l] \mid \langle s, t, b \rangle \in w \text{ for some } w \in S \}.$$

Since \mathfrak{A} is short and densely ordered, there exists a family $(I_k)_{k \in \mathbb{Z}}$ of convex subsets of A such that

$$\cdots < I_{-2} < I_{-1} < I_0 < I_1 < I_2 < \cdots$$

and $\bigcup_{k \in \mathbb{Z}} I_k$ is cofinal in both directions. Since \mathcal{A} is p modest and $|K| \leq 2^l \leq p$, we know that the sets $(P_s \cap I_k)_{s \in K}$ do not form a guard for I_k . As each $P_s \cap I_k$ is dense in I_k (and therefore coherent), this implies that we can find some replete set $C_k \subseteq I_k \cap \bigcup_{s \in K} P_s$ that is guarded by $(P_s \cap I_k)_{s \in K}$. We can use Lemma 5.17 to construct a frame φ_k with $\text{rng } \varphi \subseteq C_k$. Set

$$X := \text{cl}_2\left(\bigcup_{k \in \mathbb{Z}} \text{rng } \varphi_k\right).$$

By Lemma 5.3 (b), X is replete. Furthermore, since each C_k is replete, so is $\bigcup_k C_k$, and it follows that

$$X = \text{cl}_2\left(\bigcup_{k \in \mathbb{Z}} \text{rng } \varphi_k\right) \subseteq \text{cl}_2\left(\bigcup_{k \in \mathbb{Z}} C_k\right) = \bigcup_{k \in \mathbb{Z}} \text{cl}_2(C_k) = \bigcup_{k \in \mathbb{Z}} C_k.$$

Hence,

$$X \subseteq \bigcup_{k \in \mathbb{Z}} C_k \subseteq \bigcup_{s \in K} P_s,$$

X is cofinal in A in both directions, and each set P_s with $s \in K$ is dense in X . (P_s is dense in $\bigcup_k C_k$, and the latter is dense in X .) Let $(x_i)_{i < \omega}$ be an enumeration of X . We have shown in Lemma 5.48 that there are only countably many non-principal cuts of X that are realised in A . For each such cut, we fix one element of A realising it. Let $(a_i)_{i < \omega}$ be an enumeration of these elements. Below we will construct an increasing sequence $\emptyset = F_0 \subseteq F_1 \subseteq \dots$ of finite families of convex subsets of A with the following properties.

- (III.a) The sets in F_i are disjoint and each $D \in F_i$ is convex.
- (III.b) $x_i, a_i \in \bigcup F_{i+1}$
- (III.c) If $D_0 < \dots < D_{m-1}$ is an enumeration of F_i , then there are sets $D'_0, \dots, D'_m \in F_{i+1}$ such that

$$D'_0 < D_0 < D'_1 < D_1 < \dots < D_{m-2} < D'_{m-1} < D_{m-1} < D'_m,$$

and $\text{tp}(D'_j) = w_{i \bmod n}$ for all j .

- (III.d) For all $D < D'$ in F_i there are infinitely many elements $x, y, z \in X$ with $x < D < y < D' < z$.

(III.e) If $D \in F_i$ has a minimal or maximal element, it belongs to X .

(III.f) Let $D \in F_i$ and let E_-, E_+ be the sets such that $A = E_- + D + E_+$. If D has no minimal element, the set $E_- \cap X$ is not bounded from above in E_- . If D has no maximal element, the set $E_+ \cap X$ is not bounded from below in E_+ .

Let $F := \bigcup_{i < \omega} F_i$ be the limit. By (III.b), we have $X \subseteq \bigcup F$. Furthermore, it follows by (III.b), (III.e), and (III.f) that $A \setminus X \subseteq \bigcup F$ as well. Hence, $\bigcup F = A$ and F forms a partition of A . Let \sim be the corresponding congruence. Then, (III.c) implies that

$$S = \{ \text{tp}([a]) \mid [a] \in A/\sim \},$$

where \sim is the equivalence relation corresponding to F . It therefore remains to show that $\mathfrak{C}^{(\cdot)}(A, \sim)$ is $\langle \rangle$ -uniform. By (III.c), the classes $[a] \in A/\sim$ with $\text{tp}([a]) = w_i$, are dense in A/\sim . Since the value $\text{tp}([a])$ determines which predicate T_θ the element $[a]$ belongs to, it follows that every convex subset $J \subseteq A/\sim$ without end-points has the same $\text{MSO}_{\langle \rangle}$ -theory.

To conclude the proof, it remains to construct the sequence $(F_i)_i$. We do so by induction on i starting with $F_0 := \emptyset$. For the inductive step, suppose that we have already defined F_i . Let $D_0 < \dots < D_{m-1}$ be an enumeration of F_i and set $w := w_{i \bmod n}$.

We start by finding a set F' satisfying all conditions except for (III.b). To do so, we distinguish two cases. First, suppose that $w = \langle s, s, 1 \rangle$. Since P_s is dense in X , we can pick elements $y_0, \dots, y_m \in X \cap P_s$ with

$$y_0 < D_0 < y_1 < D_1 < \dots < y_{m-1} < D_{m-1} < y_m.$$

The set

$$F' := F_i \cup \{ \{y_0\}, \dots, \{y_m\} \}.$$

satisfies (III.a), (III.c)–(III.f).

It remains to consider the case where $w = \langle s, t, \infty \rangle$. First, assume that $s, t \neq \perp$. By (III.d), and since P_s and P_t are dense in X , there are elements $D_{j-1} < y < z < D_j$ with $y \in X \cap P_s$ and $z \in X \cap P_t$. We set $D'_{j-1} := [y, z]$.

Next suppose that $s = \perp$ and $t = \perp$. Since X is dense, it has uncountably many non-principal cuts. By Lemma 5.48, only countably many of these are realised in A . Consequently, we can find two non-principal cuts $\langle Y_0, Y_1 \rangle$ and $\langle Z_0, Z_1 \rangle$ of X that are not realised in A and such that

$$D_{j-1} \subset Y_0 \subset Z_0 \quad \text{and} \quad Y_1 \supset Z_1 \supset D_j.$$

We choose for D'_{j-1} the set

$$D'_{j-1} := \{ a \in A \mid Y_0 < a < Z_1 \}.$$

In the other two cases, we proceed analogously (choosing one element of X of the correct type and one non-principal cut of X). The resulting set

$$F' := F_i \cup \{ D'_j \mid j \leq m \}$$

satisfies (III.a), (III.c)–(III.f).

To satisfy the remaining condition (III.b), we now set

$$F_{i+1} := F' \cup \{ D', D'' \},$$

where D' and D'' are sets with $x_i \in D'$ and $a_i \in D''$ that are chosen as follows. If $x_i \in \bigcup F'$, we can find a suitable class $D' \in F'$. Hence, suppose otherwise. By assumption, there is at least one $w \in S$ of the form $w = \langle s, t, \infty \rangle$. If $s, t \neq \perp$, we pick $y \in P_s \cap X$ and $z \in P_t \cap X$ with $y \leq x_i \leq z$ such that $D' := [y, z]$ is disjoint from all classes in F' . If $s = \perp$ or $t = \perp$, we use an interval $[y, z)$, $(y, z]$, or (y, z) instead. The construction of D'' is analogous. \square

Our inductive decidability proof is split into the following three lemmas.

Lemma 5.50. *Let \mathcal{A} and \mathcal{B} be two linear orders with $l \leq \log_2 p$ parameters that are p -modest, separable, and $\langle \rangle$ -uniform. Then*

$$\mathcal{A} \equiv_{\text{MSO}}^{(\perp)} \mathcal{B} \quad \text{implies} \quad \mathcal{A} \equiv_{\text{U}}^{(\circ)} \mathcal{B}.$$

Proof. To simplify notation, we will write

$$P_s := \{ a \in A \mid \text{tp}(a) = s \}, \quad \text{for } s \subseteq [l].$$

We will prove the following claims.

(I) For each $\mathfrak{C} \in \text{UC}^{(\perp)}(\mathcal{A})$, we can compute $\text{Th}_{\text{MSO}}^{(\perp)}(\mathfrak{C})$ from

$$\text{tp}(\sim) := \{ \text{tp}([a]) \mid [a] \in I/\sim \} \quad \text{and} \quad \text{Th}_{\text{MSO}}^{(\perp)}(\mathcal{A}).$$

(II) $\text{U}^{(\circ)}(\mathcal{A})$ is uniquely determined by $\text{Th}_{\text{MSO}}^{(\perp)}(\mathcal{A})$ and

$$T := \{ \text{tp}(\sim) \mid \mathfrak{C}^{(\perp)}(I, \sim) \in \text{UC}^{(\perp)}(\mathcal{A}) \}.$$

(III) We can compute T from $\text{Th}_{\text{MSO}}^{(\perp)}(\mathcal{A})$.

Then the lemma follows from (II) and (III).

(I) Let $\mathfrak{C}^{(\perp)}(I, \sim) \in \text{UC}^{(\perp)}(\mathcal{A})$. Since $\mathfrak{C}^{(\perp)}(I, \sim)$ is $\langle \rangle$ -uniform, each predicate T_θ is dense in I/\sim . This implies that $\text{Th}_{\text{MSO}}^{(\perp)}(\mathfrak{C}^{(\perp)}(I, \sim))$ only depends on the set

$$H := \{ \theta \mid T_\theta \neq \emptyset \}.$$

By Lemma 5.45, we can use $\text{tp}(\sim)$ to determine, for each theory θ , whether or not there is some $[a] \in I/\sim$ with $\text{Th}_{\text{MSO}}^{(\perp)}(\mathcal{A}|_{[a]}) = \theta$. Consequently, we can compute H from $\text{tp}(\sim)$.

(II) Since

$$\begin{aligned} \text{U}^{(\circ)}(\mathcal{A}) &= \{ \text{U}^{q(l+\circ, \langle \rangle)}(\mathfrak{C}) \mid \mathfrak{C} \in \text{UC}^{(\perp)}(\mathcal{A}) \} \\ &= \{ \text{Th}_{\text{MSO}}^{(\perp)}(\mathfrak{C}) \mid \mathfrak{C} \in \text{UC}^{(\perp)}(\mathcal{A}) \}, \end{aligned}$$

the claim follows by (I).

(III) Set

$$\begin{aligned} H_* &:= \{ \perp \} \cup \{ s \subseteq [l] \mid P_s \neq \emptyset \}, \\ T_* &:= \{ S \subseteq H_* \times H_* \times \{1, \infty\} \mid \langle s, t, \infty \rangle \in S \text{ for some } s, t \} \\ &\quad \cup \{ \{ \langle s, s, 1 \rangle \mid s \in H_* \} \}. \end{aligned}$$

As H_* and, therefore, T_* are computable from $\text{Th}_{\text{MSO}}^{(\cdot)}(\mathcal{A})$, it is sufficient to prove that $T = T_*$.

(\supseteq) Given $S \in T_*$, let \sim be the congruence from Lemma 5.49. Then $S = \text{tp}(\sim) \in T$.

(\subseteq) Let $S = \text{tp}(\sim) \in T$. If \sim is the identity, we have $S = \{ \langle s, s, 1 \rangle \mid s \in H_* \} \in T_*$. Hence, suppose otherwise. Then there is at least one non-trivial \sim -class and S has an element of the form $\langle s, t, \infty \rangle$. Consequently, $S \in T_*$. \square

Lemma 5.51. *Let \mathcal{A} and \mathcal{B} be \bar{n} -uniform linear orders. Then*

$$\mathcal{A} \equiv_{\text{U}}^{\bar{n}} \mathcal{B} \quad \text{implies} \quad \mathcal{A} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{B}.$$

Proof. Let $\mathcal{A} = \langle A, \leq, \bar{P} \rangle$ be a linear orders with l parameters. We proceed by induction on $|\bar{n}|$. Since $\text{Th}_{\text{MSO}}^{(\cdot)}(\mathcal{A}) = \text{U}^{(\cdot)}(\mathcal{A})$, the claim is trivial for $\bar{n} = \langle \rangle$. For the inductive step, suppose that $\bar{n} = \bar{m}k$ and let

$$\begin{aligned} U &:= \{ \text{Th}_{\text{MSO}}^{q(l+k, \bar{m})}(\mathcal{C}) \mid \mathcal{C} \in \text{UC}^{\bar{m}}(\mathcal{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \}, \\ S &:= \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{A}|_I, \bar{P}') \mid I \text{ convex}, \bar{P}' \in \wp(I)^k \}. \end{aligned}$$

It is sufficient to prove the following claims.

- (I) U is uniquely determined by $\text{U}^{\bar{n}}(\mathcal{A})$.
 - (II) S is uniquely determined by U .
 - (III) $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A})$ is uniquely determined by S .
- (i) We prove the claim by induction on $|\bar{m}|$. For $\bar{m} = \langle \rangle$, the claim is trivial. For $|\bar{m}| > 0$, it follows by inductive hypothesis that

$$\text{U}^{\bar{m}k}(\mathcal{A}) = \{ \text{U}^{q(l+k, \bar{m})}(\mathcal{C}) \mid \mathcal{C} \in \text{UC}^{\bar{m}}(\mathcal{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \}$$

determines U .

(iii) Since \mathcal{A} is $\bar{m}k$ -uniform, we have

$$\mathcal{A} \equiv_{\text{MSO}}^{\bar{m}k} \mathcal{A}|_I, \quad \text{for every convex set } I.$$

Consequently,

$$\begin{aligned} S &= \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{A}|_I, \bar{P}') \mid I \text{ convex}, \bar{P}' \in \wp(I)^k \} \\ &= \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{A}, \bar{P}') \mid \bar{P}' \in \wp(A)^k \}, \end{aligned}$$

which determines $\text{Th}_{\text{MSO}}^{\bar{m}k}(\mathcal{A})$.

(ii) Set

$$\begin{aligned} \mathcal{K} &:= \{ \langle \mathcal{A}|_I, \bar{P}' \rangle \mid I \text{ convex}, \bar{P}' \in \wp(I)^k \}, \\ S_o &:= \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}) \mid \mathcal{C} \in \mathcal{K}, |\mathcal{C}| = 1 \} \subseteq S, \end{aligned}$$

and let S_* be the least set of theories with the following properties.

- ♦ $\text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}) \in S_*$, for all $\mathcal{C} \in \mathcal{K}$ with $|\mathcal{C}| = 1$.
- ♦ $\text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}), \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{D}) \in S_*$, $\mathcal{C} + \mathcal{D} \in \mathcal{K}$ implies $\text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C} + \mathcal{D}) \in S_*$.
- ♦ If $(\mathcal{C}_i)_{i < \omega}$ is a sequence such that

$$\sum_{i \in \omega} \mathcal{C}_i \in \mathcal{K} \quad \text{and} \quad \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}_i) = \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}_j) \in S_*, \quad \text{for all } i, j,$$

then $\text{Th}_{\text{MSO}}^{\bar{m}}(\sum_{i \in \omega} \mathcal{C}_i) \in S_*$.

- ♦ If $(\mathcal{C}_i)_{i < \omega}$ is a sequence such that

$$\sum_{i \in \omega^{\text{op}}} \mathcal{C}_i \in \mathcal{K} \quad \text{and} \quad \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}_i) = \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}_j) \in S_*, \quad \text{for all } i, j,$$

then $\text{Th}_{\text{MSO}}^{\bar{m}}(\sum_{i \in \omega^{\text{op}}} \mathcal{C}_i) \in S_*$.

- ♦ If $\mathcal{C} \in \mathcal{K}$ and \approx is a congruence of \mathcal{C} such that the order $\langle C/\approx, \leq, (P_c)_{c \in S} \rangle$ with

$$P_c := \{ [a] \in C/\approx \mid \text{Th}_{\text{MSO}}^{\bar{m}}([a]) = c \}, \quad \text{for } c \in S,$$

is $q(l, \bar{m})$ -uniform and densely ordered without end-points, and if $P_c \neq \emptyset$ implies $c \in S_*$, then $\text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{C}) \in S_*$.

Note that by definition, we have

$$S_* \subseteq \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{A}) \mid \mathcal{A} \in \mathcal{K} \} = S.$$

We obtain the converse inclusion by applying Corollary 4.5 to the class $\mathcal{L} := \mathcal{K}$, the set $U := S_*$, and the labellings $\lambda := \text{Th}_{\text{MSO}}^{\bar{m}}$ and $\mu := \text{Th}^{q(l, \bar{m})}$. Consequently, $S_* = S$. Hence, it remains to prove that we can compute S_* from U . To do so, we will make use of the following claims.

(II.a) S_0 is computable from U .

(II.b) $\langle \mathcal{A}|_C, \text{First}, \text{Last} \rangle \equiv_{\text{MSO}}^{(\cdot)} \langle \mathcal{A}|_D, \text{First}, \text{Last} \rangle$ implies $\mathcal{A}|_C \equiv_{\text{MSO}}^{\bar{m}k} \mathcal{A}|_D$,
for all convex $C, D \subseteq A$.

(II.c) Given $\theta, \theta' \in S_*$, there exist $\mathfrak{C}, \mathfrak{D} \in \mathcal{K}$ with

$$\text{Th}^{\bar{m}}(\mathfrak{C}) = \theta, \quad \text{Th}^{\bar{m}}(\mathfrak{D}) = \theta', \quad \text{and} \quad \mathfrak{C} + \mathfrak{D} \in \mathcal{K}$$

if, and only if, at least one of the following four cases holds.

- ♦ $\theta \models \text{Last} = \emptyset$ and \mathcal{A} is incomplete.
- ♦ $\theta' \models \text{First} = \emptyset$ and \mathcal{A} is incomplete.
- ♦ $\theta \models \text{Last} = \emptyset$ and $\theta' \models \text{First} \neq \emptyset$.
- ♦ $\theta \models \text{Last} \neq \emptyset$ and $\theta' \models \text{First} = \emptyset$.

(II.d) For $\theta \in S$,

$$\omega \times \theta \in S \quad \text{iff} \quad \omega^{\text{op}} \times \theta \in S \quad \text{iff} \quad \theta + \theta \in S.$$

(II.e) Given a theory $\sigma \in U$ such that

$$\sigma \models T_\theta \neq \emptyset \quad \text{implies} \quad \theta \in S_*,$$

there exists $\mathfrak{C} \in \mathcal{K}$ and a congruence \approx of \mathfrak{C} such that the order $\langle C/\approx, \leq, (Q_c)_{c \in S} \rangle$ with

$$Q_c := \{ [a] \in C/\approx \mid \text{Th}_{\text{MSO}}^{\bar{m}}([a]) = c \}, \quad \text{for } c \in S,$$

is $q(l, \bar{m})$ -uniform and densely ordered without end-points, and we have

$$Q_c \neq \emptyset \quad \text{iff} \quad \sigma \models T_\theta \neq \emptyset.$$

Before giving the proofs, let us show how (II) follows. We have to check that all operations that generate S_* are determined by U . We can then compute S_* by iterating all operations until no new theories are generated anymore. By (II.d), we can compute S_0 . Furthermore, given theories θ, θ' , there exist unique theories $\theta + \theta'$, $\omega \times \theta$, and $\omega^{\text{op}} \times \theta$ by the composition theorem. Finally, we can compute the set of theories added by the last operation since it follows by Lemma 5.43 that

$$\text{Th}_{\text{MSO}}^{\bar{m}}(\mathcal{A}|_I) \quad \text{is determined by} \quad \text{Th}_{\text{MSO}}^{q(l+k, \bar{m})}(\mathfrak{C}^{\bar{m}}(I, \sim)),$$

and the latter theories can be obtained from U .

It therefore remains to show, given $\theta, \theta' \in S_*$, how to determine whether or not $\theta + \theta'$, $\omega \times \theta$, and $\omega^{\text{op}} \times \theta$ also belong to S_* . For binary sums, we can decide whether $\theta + \theta' \in S_*$ by (II.c). The cases $\omega \times \theta$ and $\omega^{\text{op}} \times \theta$ now follow by (II.d).

To conclude the proof, it remains to prove the above claims.

(II.a) By Lemma 5.41, we have $U \neq \emptyset$. Choose an arbitrary theory $\theta \in U$. For $\sigma \subseteq [l]$, let

$$P_\sigma := \bigcap_{i \in \sigma} P_i \setminus \bigcup_{i \notin \sigma} P_i.$$

Since \mathcal{A} is \bar{n} -uniform, we can use θ to compute the set

$$H := \{ \sigma \subseteq [l] \mid P_\sigma \neq \emptyset \}.$$

Let 1 be some 1-element set. Setting

$$Q_i^\sigma := \begin{cases} 1 & \text{if } i \in \sigma, \\ \emptyset & \text{if } i \notin \sigma, \end{cases} \quad \text{for } \sigma \subseteq [l],$$

it follows that

$$S_0 = \{ \text{Th}_{\text{MSO}}^{\bar{m}}(\langle 1, \bar{Q}^\sigma, \bar{Q}' \rangle) \mid \sigma \in H, \bar{Q}' \in \wp(1)^k \},$$

which can be computed from H .

(11.b) Let $C, D \subseteq A$ be $\langle \rangle$ -equivalent convex subsets. If C and D do not have end-points, the claim follows immediately by \bar{n} -uniformity. If C and D have a minimal element but no maximal one, we can write $C = \{c\} + C_o$ and $D = \{d\} + D_o$. As above, it follows by \bar{n} -uniformity that

$$\mathcal{A}|_{C_o} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{D_o}.$$

Furthermore, we have $\mathcal{A}|_{\{c\}} \cong \mathcal{A}|_{\{d\}}$ because of the predicate First. Consequently, the composition theorem implies that

$$\mathcal{A}|_C \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_D.$$

The other cases follow analogously.

(11.c) (\Leftarrow) Fix $\mathfrak{C}, \mathfrak{D} \in \mathcal{K}$ with theories θ and θ' , respectively, and suppose that the conditions on the right-hand side are not satisfied. To show that $\mathfrak{C} + \mathfrak{D} \notin \mathcal{K}$, we distinguish two cases.

If $\theta \models \text{Last} \neq \emptyset$ and $\theta' \models \text{First} \neq \emptyset$, then $\mathfrak{C} + \mathfrak{D}$ is not densely ordered. Hence, it is not isomorphic to a convex suborder of $\langle \mathcal{A}, \bar{P}' \rangle$.

Similarly, if \mathcal{A} is complete, $\theta \models \text{Last} = \emptyset$, and $\theta' \models \text{First} = \emptyset$, then $\mathfrak{C} + \mathfrak{D}$ is not complete and, therefore, not isomorphic to a convex suborder of $\langle \mathcal{A}, \bar{P}' \rangle$.

(\Rightarrow) Fix $\mathfrak{C}, \mathfrak{D} \in \mathcal{K}$ with theories θ and θ' , respectively. There are four cases to consider.

- ♦ \mathfrak{C} has no maximal element and \mathcal{A} is incomplete.
- ♦ \mathfrak{D} has no minimal element and \mathcal{A} is incomplete.
- ♦ \mathfrak{C} has no maximal element and \mathfrak{D} has a minimal element.
- ♦ \mathfrak{C} has a maximal element and \mathfrak{D} has no minimal element.

In each of these cases, we can choose two disjoint convex subsets $C', D' \subseteq A$ such that $C' + D'$ is convex,

$$\mathcal{A}|_{C', \text{First}, \text{Last}} \equiv_{\bar{n}\text{MSO}}^{\langle \rangle} \mathcal{A}|_{C, \text{First}, \text{Last}},$$

$$\mathcal{A}|_{D', \text{First}, \text{Last}} \equiv_{\bar{n}\text{MSO}}^{\langle \rangle} \mathcal{A}|_{D, \text{First}, \text{Last}}.$$

By (11.b), it follows that

$$\mathcal{A}|_{C'} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_C \quad \text{and} \quad \mathcal{A}|_{D'} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_D.$$

Consequently, there are sets $\bar{P}' \in \wp(C)^k, \bar{Q}' \in \wp(D)^k$ with

$$\mathfrak{C}' := \langle \mathcal{A}|_{C'}, \bar{P}' \rangle \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathfrak{C} \quad \text{and} \quad \mathfrak{D}' := \langle \mathcal{A}|_{D'}, \bar{Q}' \rangle \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathfrak{D}.$$

It follows that $\text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathfrak{C}') = \theta, \text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathfrak{D}') = \theta'$, and $\mathfrak{C}' + \mathfrak{D}' \in \mathcal{K}$.

(11.d) We only show that $\omega \times \theta \in S$ if, and only if, $\theta + \theta \in S$. The other proof is analogous.

(\Leftarrow) If there is some convex subset of the form $\sum_{i < \omega} C_i$ with

$$\text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C_i}, \bar{P}'_i) = \theta, \quad \text{for all } i < \omega,$$

then $\text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C_o + C_i}, \bar{P}'_o + \bar{P}'_i) = \theta + \theta$, which implies that $\theta + \theta \in S_*$.

(\Rightarrow) Suppose that there are disjoint subsets $C_o, D_o \subseteq A$ such that $C_o + D_o$ is convex and

$$\text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C_o}, \bar{P}'_o) = \theta \quad \text{and} \quad \text{Th}_{\bar{n}\text{MSO}}^{\bar{n}}(\mathcal{A}|_{D_o}, \bar{Q}'_o) = \theta.$$

We claim that there exists a partition $D_o = C_i + D_i$ such that

$$\mathcal{A}|_{C_i} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{C_o} \quad \text{and} \quad \mathcal{A}|_{D_i} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{D_o}.$$

For the proof, we distinguish several cases. If C_o and D_o have no end-points, the claim follows immediately by \bar{n} -uniformity of \mathcal{A} .

If C_o and D_o have a minimal element, but no maximal one, we can write $C_o = \{c\} + C'_o$ and $D_o = \{d\} + D'_o$. By \bar{n} -uniformity of \mathcal{A} , it follows that there exists a partition $D'_o = C'_i + D_i$ such that

$$\mathcal{A}|_{C'_i} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{C'_o} \quad \text{and} \quad \mathcal{A}|_{D_i} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{D_o}.$$

Furthermore,

$$\mathcal{A}|_{C_o} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{D_o} \quad \text{implies} \quad \mathcal{A}|_{\{c\}} \equiv_{\bar{n}\text{MSO}}^{\bar{n}} \mathcal{A}|_{\{d\}}.$$

Consequently,

$$\mathfrak{A}|_{\{c\}+C'_o} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_{\{d\}+C'_i}.$$

Hence, we can set $C_i := \{d\} + C'_i$.

If C_o and D_o have no minimal element, but a maximal one, we can write $C_o = C'_o + \{c\}$ and $D_o = D'_o + \{d\}$. By \bar{n} -uniformity of \mathfrak{A} , it follows that there exists a partition $D'_o = C_i + D'_i$ such that

$$\mathfrak{A}|_{C_i} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_{C_o} \quad \text{and} \quad \mathfrak{A}|_{D'_i} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_{D'_o}.$$

Consequently,

$$\mathfrak{A}|_{D'_o+\{d\}} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_{D'_i+\{d\}},$$

and we can set $D_i := D'_i + \{d\}$.

Finally, it is impossible that C_o and D_o have both a minimal element and a maximal one, since $C_o + D_o$ is densely ordered.

Having proved the claim, we can use it to find parameters \bar{P}'_i, \bar{Q}'_i such that

$$\mathfrak{A}|_{C_i, \bar{P}'_i} \equiv_{\text{MSO}}^{\bar{m}} \mathfrak{A}|_{C_o, \bar{P}'_o} \quad \text{and} \quad \mathfrak{A}|_{D_i, \bar{Q}'_i} \equiv_{\text{MSO}}^{\bar{m}} \mathfrak{A}|_{D_o, \bar{Q}'_o}.$$

Repeating this step for D_i and \bar{Q}'_i , we obtain a convex subset of the form $\sum_{i < \omega} C_i$ and parameters $\bar{P}'_i \in \mathcal{P}(C_i)^k$ such that

$$\text{Th}_{\text{MSO}}^{\bar{m}}(\mathfrak{A}|_{C_i, \bar{P}'_i}) = \theta, \quad \text{for all } i < \omega.$$

Consequently, $\omega \times \theta \in S_*$. The proof for $\omega^{\text{op}} \times \theta$ is analogous.

(11.e) Fix $\sigma \in U$ such that

$$\sigma \models T_\theta \neq \emptyset \quad \text{implies} \quad \theta \in S_*.$$

Then there exists an order $\mathfrak{C} \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}')$ with $\sigma = \text{Th}_{\text{MSO}}^{q(l+k, \bar{m})}(\mathfrak{C})$. Suppose that $\mathfrak{C} = \mathfrak{C}^{\bar{m}}(I, \approx)$. Note that $[a]_\approx \in T_\theta$ implies

$$\theta \in S_* \quad \text{and} \quad \text{Th}^{\bar{m}}(\mathfrak{A}|_{[a]_\approx, \bar{Q}_a}) = \theta, \quad \text{for some } \bar{Q}_a.$$

Let \bar{Q} be the tuple consisting of the concatenations of the respective components of \bar{Q}_a , for $[a]_\approx \in I/\approx$. Then $\langle \mathfrak{A}|_I, \bar{Q} \rangle \in \mathcal{K}$, \mathfrak{C} is $q(l+k, \bar{m})$ -uniform, and

$$Q_c \neq \emptyset \quad \text{iff} \quad \sigma \models T_\theta \neq \emptyset. \quad \square$$

Lemma 5.52. *There exists a function $L : \omega^* \rightarrow \omega$ with the following properties.*

(a) *Let \mathfrak{A} and \mathfrak{B} be two linear orders with $l \leq \log_2 L(\bar{n})$ parameters that are p -modest, separable, and $\langle \rangle$ -uniform. Then*

$$\mathfrak{A} \equiv_{\text{MSO}}^{\langle \rangle} \mathfrak{B} \quad \text{implies} \quad \mathfrak{A} \equiv_{\text{U}}^{\bar{n}} \mathfrak{B}.$$

(b) *Every linear order with at most $\log_2 L(\bar{n}o)$ parameters that is $L(\bar{n}o)$ -modest, separable, and $\langle \rangle$ -uniform, is \bar{n} -uniform.*

Proof. Let $L : \omega \times \omega^* \rightarrow \omega$ be the function from Lemma 5.50. In the following, we will work with the linear order \mathfrak{A} and with $\langle \mathfrak{A}, \bar{P}' \rangle$ at the same time. To avoid confusion, we will use the notation

$$\text{tp}(a), \quad \text{tp}(C), \quad \mathfrak{C}^{\bar{n}}(I, \sim)$$

when referring to \mathfrak{A} and the notation

$$\text{tp}(a; \bar{P}'), \quad \text{tp}(C; \bar{P}'), \quad \mathfrak{C}^{\bar{n}}(I, \sim; \bar{P}')$$

when referring to $\langle \mathfrak{A}, \bar{P}' \rangle$.

(a) We proceed by induction on $|\bar{n}|$. For $\bar{n} = \langle \rangle$, the claim follows by Lemma 5.50. Hence, suppose that $\bar{n} = \bar{m}k$. By inductive hypothesis, there exists a function g such that

$$g(\text{Th}_{\text{MSO}}^{\langle \rangle}(\mathfrak{B})) = \text{U}^{q(l+k, \bar{m})}(\mathfrak{B}),$$

for all $\langle \rangle$ -uniform $L(q(l+k, \bar{m}))$ -modest \mathfrak{B} with at most $\log_2 L(q(l+k, \bar{m}))$ parameters.

Let r be the function mapping $\mathfrak{C}^{\bar{m}}(I, \sim; \bar{P}')$ to $\mathfrak{C}^{(\cdot)}(I, \sim; \bar{P}')$. Given $u \in U^{(\circ)}(\mathfrak{A})$, set

$$S_u := \left\{ \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}) \mid \mathfrak{C} \text{ a linear order such that } r(\mathfrak{C}) \in \text{UC}^{(\cdot)}(\mathfrak{A}, \bar{P}') \right. \\ \left. \text{and } \text{Th}_{\text{MSO}}^{(\cdot)}(r(\mathfrak{C})) \models u \right\}.$$

We will establish the following two claims.

- (I) S_u is uniquely determined by u and $\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{A})$.
- (II) $\left\{ g(\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C})) \mid \mathfrak{C} \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \right\}$
 $= \left\{ g(\theta) \mid \theta \in S_u, u \in U^{(\circ)}(\mathfrak{A}) \right\}.$

Then it follows that

$$U^{\bar{m}k}(\mathfrak{A}) = \left\{ \bigcup^{q(l+k, \bar{m})}(\mathfrak{C}) \mid \mathfrak{C} \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \right\} \\ = \left\{ g(\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C})) \mid \mathfrak{C} \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}'), \bar{P}' \in \wp(A)^k \right\} \\ = \left\{ g(\theta) \mid \theta \in S_u, u \in U^{(\circ)}(\mathfrak{A}) \right\}.$$

To see that the latter expression is uniquely determined by $\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{A})$, note that we have shown in Lemma 5.50 that

$$\mathfrak{A} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{B} \text{ implies } U^{(\circ)}(\mathfrak{A}) = U^{(\circ)}(\mathfrak{B}),$$

for all $\langle \cdot \rangle$ -uniform $L(\bar{n})$ -modest orders \mathfrak{A} and \mathfrak{B} . Hence, it remains to prove the two claims above.

(I) Suppose that $\mathfrak{A} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{B}$ are $\langle \cdot \rangle$ -uniform, fix u , and let \bar{P}' be parameters and \mathfrak{C} a linear order such that

$$r(\mathfrak{C}) \in \text{UC}^{(\cdot)}(\mathfrak{A}, \bar{P}') \quad \text{and} \quad \text{Th}_{\text{MSO}}^{(\cdot)}(r(\mathfrak{C})) \models u.$$

It follows that $r(\mathfrak{C}) = \mathfrak{C}^{(\cdot)}(I, \sim; \bar{P}')$, for some convex set $I \subseteq A$ and some congruence \sim . By Lemma 5.50, we have $\mathfrak{A} \equiv_U^{(\circ)} \mathfrak{B}$. Consequently, there exists some order $\mathfrak{D}_0 = \mathfrak{C}^{(\cdot)}(J, \approx) \in \text{UC}^{(\cdot)}(\mathfrak{B})$ with

$$\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{D}_0) = u.$$

For each \approx -class $[b] \in J/\approx$, we choose predicates $\bar{Q}_b \in \wp([b])^k$ as follows. Given a type $\sigma := \text{tp}([a])$ with $[a] \in I/\sim$, set

$$H_\sigma := \left\{ \text{tp}([a']; \bar{P}') \mid [a'] \in I/\sim, \text{tp}([a']) = \sigma \right\}.$$

We choose a partition $(K_w)_{w \in H_\sigma}$ of the set

$$K := \left\{ [b] \in B/\approx \mid \text{tp}([b]) = \sigma \right\}$$

such that each set K_w is dense in K . Let $\bar{Q}_b \in \wp([b])^k$ be sets such that

$$\text{tp}([b]; \bar{Q}_b) = w, \quad \text{for } [b] \in K_w.$$

Then we set $Q'_j := \bigcup_{[b] \in B/\approx} (Q_b)_j$, for $j < k$.

It follows that the resulting order $\mathfrak{D}_1 := \mathfrak{C}^{\bar{m}}(J, \approx; \bar{Q}')$ is $\langle \cdot \rangle$ -uniform. Furthermore, we have $r(\mathfrak{C}) \equiv_{\text{MSO}}^{(\cdot)} r(\mathfrak{D}_1)$ since, for every $[a] \in I/\sim$, there is some $[b] \in B/\approx$ with

$$\langle \mathfrak{A}|_{[a]}, \text{First}, \text{Last} \rangle \equiv_{\text{MSO}}^{(\cdot)} \langle \mathfrak{B}|_{[b]}, \text{First}, \text{Last} \rangle.$$

This implies that $\text{Th}_{\text{MSO}}^{(\cdot)}(r(\mathfrak{D}_1)) \models u$. Consequently, it remains to find some $\mathfrak{D} \in r^{-1}(\mathfrak{D}_1)$ with $\mathfrak{C} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{D}$.

By construction, every predicate T_σ of \mathfrak{D}_1 is either empty or infinite. For an $\text{MSO}_{\langle \cdot \rangle}$ -theory σ , we denote by H_σ the set of all $\text{MSO}_{\bar{m}}$ -theories τ such that $\tau \models \sigma$ and $T_\tau \neq \emptyset$ in \mathfrak{C} . For each σ , we choose a partition $(T'_\tau)_{\tau \in H_\sigma}$ of T_σ such that every T'_τ is infinite. For theories $\tau \notin \bigcup_\sigma H_\sigma$, we set $T'_\tau := \emptyset$. The resulting order $\mathfrak{D} := \langle B/\approx, \bar{T}' \rangle$ satisfies $r(\mathfrak{D}) = \mathfrak{D}_1$ and $\mathfrak{C} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{D}$.

(II) (\subseteq) Fix $\bar{P}' \in \wp(A)^k$ and let $\mathfrak{C}^{\bar{m}}(I, \sim; \bar{P}') \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}')$. Since $\mathfrak{C}^{\bar{m}}(I, \sim; \bar{P}')$ is $q(l+k, \bar{m})$ -uniform, its reduct $\mathfrak{C}^{(\cdot)}(I, \sim; \bar{P}')$ is $\langle \cdot \rangle$ -uniform. Consequently, we have

$$\mathfrak{C}^{(\cdot)}(I, \sim; \bar{P}') \in \text{UC}^{(\cdot)}(\mathfrak{A}, \bar{P}').$$

Set

$$\theta := \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}^{\bar{m}}(I, \sim; \bar{P}')) \quad \text{and} \quad u := \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}^{\bar{m}}(I, \sim))$$

Then $\theta \in S_u$ and $u \in U^{(o)}(\mathfrak{A})$.

(\supseteq) Let $u \in U^{(o)}(\mathfrak{A})$ and $\theta \in S_u$. By definition of S_u , there exists some order \mathfrak{C} with

$$\text{Th}_{\text{MSO}}^{(\cdot)}(r(\mathfrak{C})) \models u \quad \text{and} \quad \text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}) = \theta.$$

We can use Lemma 5.41 to find a $\langle \cdot \rangle$ -uniform convex subset I of C . Let $\bar{T} = (T_\tau)_\tau$ be the predicates from \mathfrak{C} . For every $\text{MSO}_{\langle \cdot \rangle}$ -theory σ over the signature without the predicates \bar{P}' , we set

$$H_\sigma := \{ \tau \mid \tau_o = \sigma, T_\tau \cap I \neq \emptyset \},$$

where $\tau_o \subseteq \tau$ denotes the theory without the predicates \bar{P}' that corresponds to τ . Let \mathfrak{C}' be the order obtained from \mathfrak{C} as follows. For each theory σ , we choose a partition $(T'_\tau)_{\tau \in H_\sigma}$ of

$$\bigcup_{\tau \in H_\sigma} T_\tau$$

such that each T'_τ is dense in $\bigcup_{\tau \in H_\sigma} T_\tau$. For $\tau \notin H_{\tau_o}$, we choose $T'_\tau := \emptyset$. Setting $\mathfrak{C}' := \langle C, \leq, (T'_\tau)_\tau \rangle$ it follows that \mathfrak{C}' is \emptyset -uniform and

$$\mathfrak{C}' \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{C}.$$

We can use Lemma 5.41 to find a $q(k+l, \bar{m})$ -uniform convex subset J of \mathfrak{C}' . Then $\mathfrak{C}'|_J \in \text{UC}^{\bar{m}}(\mathfrak{A}, \bar{P}'')$, for some \bar{P}'' , and it follows by $\langle \cdot \rangle$ -uniformity of \mathfrak{C}' that

$$\bigcup^{q(k+l, \bar{m})} (\mathfrak{C}'|_J) = g(\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}'|_J)) = g(\text{Th}_{\text{MSO}}^{(\cdot)}(\mathfrak{C}')) = g(\theta).$$

Furthermore, $\text{Th}_{\text{MSO}}^{(\cdot)}(r(\mathfrak{C}')) \models u$.

(b) Let \mathcal{L} be the set of all convex suborders of \mathfrak{A} and set $S := S_o \cup S_i$ where

$$S_o := \{ \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_C) \mid C \subseteq A, |C| = 1 \},$$

$$S_i := \{ \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_C) \mid C \subseteq A \text{ convex with } \bar{n}\text{-uniform interior} \}.$$

We claim that S satisfies the following conditions.

- $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}) \in S$, for all $\mathfrak{C} \in \mathcal{L}$ with $|C| = 1$.
- $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}), \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{D}) \in S, \mathfrak{C} + \mathfrak{D} \in \mathcal{L}$ implies $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C} + \mathfrak{D}) \in S$.
- If $(\mathfrak{C}_i)_{i < \omega}$ is a sequence such that

$$\sum_{i \in \omega} \mathfrak{C}_i \in \mathcal{L} \quad \text{and} \quad \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}_i) = \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}_j) \in S, \quad \text{for all } i, j,$$

then $\text{Th}_{\text{MSO}}^{\bar{n}}(\sum_{i \in \omega} \mathfrak{C}_i) \in S$.

- If $(\mathfrak{C}_i)_{i < \omega}$ is a sequence such that

$$\sum_{i \in \omega^{\text{op}}} \mathfrak{C}_i \in \mathcal{L} \quad \text{and} \quad \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}_i) = \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}_j) \in S, \quad \text{for all } i, j,$$

then $\text{Th}_{\text{MSO}}^{\bar{n}}(\sum_{i \in \omega^{\text{op}}} \mathfrak{C}_i) \in S$.

- If $\mathfrak{C} \in \mathcal{L}$ and \approx is a congruence of \mathfrak{C} such that the order $\langle C/\approx, \leq, (P_c)_{c \in S} \rangle$ with

$$P_c := \{ [a] \in C/\approx \mid \text{Th}_{\text{MSO}}^{\bar{n}}([a]) = c \}, \quad \text{for } c \in S,$$

is $q(l, \bar{n})$ -uniform and densely ordered without end-points, and if $P_c \neq \emptyset$ implies $c \in S$, then $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{C}) \in S$.

Then it follows by Corollary 4.5 that

$$S = \{ \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_C) \mid C \subseteq A \text{ convex} \}.$$

To show that \mathfrak{A} is \bar{n} -uniform, consider a convex set $C \subseteq A$. Since

$$\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}), \text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_C) \in S,$$

we can find \bar{n} -uniform subsets $I, J \subseteq A$ such that

$$\mathfrak{A} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_I \quad \text{and} \quad \mathfrak{A}|_C \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_J.$$

As \mathfrak{A} is $\langle \cdot \rangle$ -uniform, we have $\mathfrak{A}|_I \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{A}|_J$, which, by (a), implies that $\mathfrak{A}|_I \equiv_{\bar{n}}^{\bar{n}} \mathfrak{A}|_J$. Consequently, it follows by Lemma 5.51 that

$$\mathfrak{A} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_I \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_J \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_C.$$

Hence, it remains to check the above closure properties of S . The first one is trivial since $S_0 \subseteq S$. For the second one, let $C, D \subseteq A$ be disjoint sets such that $C + D$ is convex and $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_C), \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_D) \in S$. Then there are convex sets C' and D' with \bar{n} -uniform interior such that

$$\mathcal{A}|_C \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C'} \quad \text{and} \quad \mathcal{A}|_D \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{D'}.$$

Let C° be the interior of C' and D° the one of D' .

Let us prove that the set X obtained from $C' + D'$ by removing the minimal and maximal elements (if they exist) is also \bar{n} -uniform. Hence, consider a convex set $I \subseteq X$ without end-points. Set $J := I \cap C'$ and $K := I \cap D'$. Since the interiors of C' and D' are \bar{n} -uniform, it follows that

$$\mathcal{A}|_{J^\circ} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C^\circ} \quad \text{and} \quad \mathcal{A}|_{K^\circ} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{D^\circ},$$

where J° and K° are the interiors of J and K , respectively. Let Z be the set such that $X = C^\circ + Z + D^\circ$. (Hence, $Z = \emptyset$ or $|Z| = 1$.) If $J^\circ \neq \emptyset$ and $K^\circ \neq \emptyset$, we have

$$\mathcal{A}|_I = \mathcal{A}|_{J^\circ} + \mathcal{A}|_Z + \mathcal{A}|_{K^\circ} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C^\circ} + \mathcal{A}|_Z + \mathcal{A}|_{D^\circ} = \mathcal{A}|_X,$$

as desired.

Hence suppose otherwise. By symmetry, we may assume that $K^\circ = \emptyset$. Then $I \subseteq C^\circ$, which implies by \bar{n} -uniformity that

$$\mathcal{A}|_I \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C^\circ}.$$

As \mathcal{A} is \emptyset -uniform, there exists a partition $I = Y_0 + Y_1 + Y_2$ of I into convex sets such that $Y_0 \neq \emptyset$, $Y_2 \neq \emptyset$, and $\mathcal{A}|_{Y_1} \equiv_{\text{MSO}}^{\emptyset} \mathcal{A}|_Z$. It follows that

$$\mathcal{A}|_{Y_0} \equiv_{\text{MSO}}^{\emptyset} \mathcal{A}|_{C^\circ}, \quad \mathcal{A}|_{Y_1} \equiv_{\text{MSO}}^{\emptyset} \mathcal{A}|_Z, \quad \mathcal{A}|_{Y_2} \equiv_{\text{MSO}}^{\emptyset} \mathcal{A}|_{D^\circ}.$$

Since C° is \bar{n} -uniform, we therefore have

$$\mathcal{A}|_{Y_0} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C^\circ}.$$

Furthermore, since Y_2 and D° are \bar{n} -uniform, we can use (a) and Lemma 5.51 to show that

$$\mathcal{A}|_{Y_2} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{D^\circ}.$$

Hence,

$$\mathcal{A}|_I = \mathcal{A}|_{Y_0 + Y_1 + Y_2} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C^\circ + Z + D^\circ} = \mathcal{A}|_X,$$

as desired.

Since \mathcal{A} is $\langle \rangle$ -uniform, there exists a convex subset $I \subseteq C'$ such that

$$\mathcal{A}|_I, \text{First}, \text{Last} \equiv_{\text{MSO}}^{\langle \rangle} \mathcal{A}|_{C' + D'}, \text{First}, \text{Last}.$$

By (a), we have

$$\mathcal{A}|_I \equiv_{\text{U}}^{\bar{n}} \mathcal{A}|_{C' + D'}.$$

Since $I \subseteq C'$, the interior of I is \bar{n} -uniform. Furthermore, we have shown above that so is the interior of $C' + D'$. Consequently, Lemma 5.51 implies that

$$\mathcal{A}|_I \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C' + D'}.$$

Hence,

$$\begin{aligned} \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_C) + \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_D) &= \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C' + D'}) \\ &= \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_I) \in S. \end{aligned}$$

For the third condition, let C_0, C_1, \dots be a sequence of disjoint sets such that $\sum_{i < \omega} C_i$ is convex and $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C_i}) \in S$, for all i . Then there are convex sets C'_i with \bar{n} -uniform interior such that

$$\mathcal{A}|_{C_i} \equiv_{\text{MSO}}^{\bar{n}} \mathcal{A}|_{C'_i}.$$

Analogously to above we obtain

$$\sum_{i < \omega} \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_{C_i}) = \text{Th}_{\text{MSO}}^{\bar{n}}(\mathcal{A}|_{\sum_{i < \omega} C'_i}) \in S.$$

The proof of the fourth property is analogous.

Finally, let $C \subseteq A$ be convex and let \sim be a congruence on C such that $\mathfrak{C}^{\bar{n}}(C, \sim) = \langle C/\sim, \leq, (T_\theta)_\theta \rangle$ is $q(l, \bar{n})$ -uniform and $T_\theta \neq \emptyset$ implies $\theta \in S$. We have to show that $\text{Th}_{\text{MSO}}^{\bar{n}}(\mathfrak{A}|_C) \in S$. By definition,

$$u := \bigcup^{q(l, \bar{n})} (\mathfrak{C}^{\bar{n}}(C, \sim)) \in \bigcup^{\bar{n}o} (\mathfrak{A}).$$

Fix a $q(l, \bar{n})$ -uniform convex set $I \subseteq A$ without end-points. Since $l \leq \log_2 L(\bar{n}o)$ and \mathfrak{A} is $L(\bar{n}o)$ -modest, it follows by (a) that $\mathfrak{A}|_I \equiv_{\bar{n}o}^{\bar{n}} \mathfrak{A}$. Hence,

$$u = \bigcup^{q(l, \bar{n})} (\mathfrak{D}), \quad \text{for some } \mathfrak{D} \in \text{UC}^{\bar{n}}(\mathfrak{A}|_I).$$

Fix J and \approx with $\mathfrak{D} = \mathfrak{C}^{\bar{n}}(J, \approx)$. Then

$$\mathfrak{C}^{\bar{n}}(C, \sim) \equiv_{\bar{U}}^{q(l, \bar{n})} \mathfrak{C}^{\bar{n}}(J, \approx)$$

implies by Lemma 5.51 that

$$\mathfrak{C}^{\bar{n}}(C, \sim) \equiv_{\text{MSO}}^{q(l, \bar{n})} \mathfrak{C}^{\bar{n}}(J, \approx).$$

Consequently, we can use Lemma 5.43 to show that

$$\mathfrak{A}|_C \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{A}|_J.$$

Hence, $\text{Th}^{\bar{n}}(\mathfrak{A}|_C) = \text{Th}^{\bar{n}}(\mathfrak{A}|_J) \in S$. \square

Proposition 5.53. *For every $m < \omega$, there exists a constant $p < \omega$ such that*

$$\mathfrak{A} \equiv_{\text{MSO}}^m \mathfrak{B}, \quad \text{for all } p\text{-modest separable dense linear orders } \mathfrak{A}, \mathfrak{B} \\ \text{without end-points.}$$

Proof. Let L be the function from Lemma 5.52, fix a quantifier structure \bar{n} , and let \mathfrak{A} and \mathfrak{B} be $L(\bar{n})$ -modest and densely ordered without end-points. Since \mathfrak{A} and \mathfrak{B} have no parameters, we have $\mathfrak{A} \equiv_{\text{MSO}}^{(\cdot)} \mathfrak{B}$. Hence, we can use Lemma 5.52 (a) to show that $\mathfrak{A} \equiv_{\bar{U}}^{\bar{n}} \mathfrak{B}$. Furthermore, it follows by Lemma 5.52 (b) that \mathfrak{A} and \mathfrak{B} are \bar{n} -uniform. Consequently, Lemma 5.51 implies that $\mathfrak{A} \equiv_{\text{MSO}}^{\bar{n}} \mathfrak{B}$. \square

We are finally able to give the missing proofs.

Proof of Theorem 5.7. (\Leftarrow) We have seen in Corollary 5.26 that $\langle \mathbb{Q}, \leq \rangle$ is modest. Hence, the claim follows by Proposition 5.53.

(\Rightarrow) For each $p < \omega$, there exists an MSO-formula φ_p expressing that a given linear order is p -modest and densely ordered without end-points. By assumption,

$$\langle \mathbb{Q}, \leq \rangle \models \varphi_p \quad \text{implies} \quad \mathfrak{A} \models \varphi_p. \quad \square$$

Proof of Corollary 5.8. (\Leftarrow) follows from Theorem 5.7 and Corollary 3.3, while (\Rightarrow) follows from Theorem 5.38. \square

Proof of Corollary 5.9. We have shown in Lemma 5.19 that $\langle \mathbb{R}, \leq \rangle$ is not modest. Hence, the claim follows by Theorem 5.38. \square

6 Choice Functions

In this section we study the existence of MSO-definable choice functions for linear orders and trees.

Definition 6.1. An MSO-formula $\psi(x, X)$ defines a *choice function* for a structure \mathfrak{A} if

- ♦ $\mathfrak{A} \models \psi(a, U)$ implies $a \in U$, and
- ♦ for every non-empty set U , there is exactly one element $a \in A$ with $\mathfrak{A} \models \psi(a, U)$.

We say that \mathfrak{A} has *strong MSO-choice* if some MSO-formula defines a choice function for \mathfrak{A} . It has (*weak*) *MSO-choice* if some expansion $\langle \mathfrak{A}, \bar{P} \rangle$ by monadic predicates \bar{P} has strong MSO-choice. \lrcorner

It turns out that not very many trees have MSO-choice. We can characterise them as follows. We will be working with generalised trees so we can treat trees and linear orders at the same time.

Definition 6.2. Let \mathfrak{T} be a generalised tree.

(a) The *branching degree* of \mathfrak{T} at a vertex $v \in T$ is the number of classes of the following equivalence relation on $T|_v$.

$$u \sim u' \quad : \text{iff} \quad u \sqcap u' > v.$$

If this number is at most n , we say that the *branching degree* of \mathfrak{T} is *bounded* by n . If the number is finite, we say that \mathfrak{T} is *finitely branching*.

(b) We call \mathfrak{T} *tame* if there are constants $k, n < \omega$ such that

- ♦ \mathfrak{T} is thin,
- ♦ its branching degree is bounded by n , and
- ♦ $\text{HR}(\beta) \leq k$, for every branch β of \mathfrak{T} .

Theorem 6.3. Let \mathfrak{T} be a generalised tree. The following statements are equivalent.

- (1) \mathfrak{T} has weak MSO-choice.
- (2) There exists an MSO-formula $\varphi(x, y)$ (with monadic parameters) defining a well-ordering on \mathfrak{T} .
- (3) \mathfrak{T} is tame.

The proof constitutes most of the rest of this section. But first let us mention that, for ordinary trees and linear orders, this characterisation simplifies as follows.

Corollary 6.4. Let \mathfrak{T} be an order-tree. The following statements are equivalent.

- (1) \mathfrak{T} has weak MSO-choice.
- (2) There exists an MSO-formula $\varphi(x, y)$ (with monadic parameters) defining a well-ordering on \mathfrak{T} .
- (3) \mathfrak{T} is thin and there is some number $n < \omega$, such that every vertex has at most n successors.

Corollary 6.5. Let \mathfrak{A} be a linear order. The following statements are equivalent.

- (1) \mathfrak{A} has weak MSO-choice.

(2) There exists an MSO-formula $\varphi(x, y)$ (with monadic parameters) defining a well-ordering on \mathfrak{A} .

(3) $\text{HR}(\mathfrak{A}) < \omega$.

Concerning strong choice, let us just mention the following result, which immediately follows from Corollary V.5.17.

Proposition 6.6. Every finitely branching successor-ordered tree \mathfrak{T} with finite Cantor-Bendixson rank has strong MSO-choice.

In general, the formula obtained from Corollary V.5.17 depends on the Cantor-Bendixson rank of the tree. It is still open whether we can find a single formula that works for all thin trees.

Open Question. Does there exist an MSO-formula $\varphi(x, X)$ that defines a choice function over every thin, finitely branching, successor-ordered tree \mathfrak{T} ?

Non-Tame Trees

We start the proof of Theorem 6.3 with the negative results. Our main counterexample is the infinite binary tree.

Theorem 6.7. The infinite binary tree $\mathfrak{T}_{\text{bin}} = \langle \{0, 1\}^*, \leq_{\text{pf}}, \text{suc}_0, \text{suc}_1 \rangle$ does not have MSO-choice.

Proof. For a contradiction, suppose that there exists an MSO_r-formula $\varphi(x, X; \bar{P})$ that defines a choice function for $\mathfrak{T}_{\text{bin}}$. First note that the statement

“Using the parameters \bar{P} , $\varphi(x, X; \bar{P})$ defines a choice function.”

can be expressed in MSO. Consequently, if there exists a choice function definable using parameters \bar{P} , we can use Theorem V.2.17 to replace \bar{P} by a tuple of regular parameters \bar{P}' . As regular sets $P' \subseteq \{0, 1\}^*$ are MSO-definable, it is therefore possible to eliminate the parameters \bar{P}' from the formula φ . Consequently, we can assume without loss of generality that $\varphi(x, X)$ is a formula without parameters.

Let m be the number of all MSO_r -theories of structures of the form $\langle \mathfrak{T}_{\text{bin}}, Q, Q' \rangle$ with 2 parameters $Q, Q' \subseteq \{o, 1\}^*$. Set

$$W_i := (o^m o^* 1)^i, \quad V_i := (o + 1)^* W_i, \quad \text{and} \quad U := V_m (o + 1)^*.$$

and $t := \langle \mathfrak{T}_{\text{bin}}, U \rangle$. Let $t_i := t|_{x_i}$ be the subtree rooted at the vertex

$$x_i := (o^m 1)^i \in V_i.$$

Note that $t_o = t$ and that every subtree of t attached at some vertex in V_i is isomorphic to t_i . (Since $w, w' \in W_i$ implies $v w u \in U \Leftrightarrow v' w' u \in U$, for all u, v, v' .)

By choice of m , there exist indices $o \leq i < j \leq m$ such that

$$t_i, \{x_i\} \equiv_{\text{MSO}}^r t_j, \{x_j\}.$$

Note that we can obtain t_{j-i} from t_o by replacing all subtrees attached to a vertex in V_i (which are isomorphic to t_i) by t_j . Consequently, it follows by Proposition V.1.7 that

$$t_l, \{x_l\} \equiv_{\text{MSO}}^r t_o, \{x_o\}, \quad \text{where} \quad l := j - i.$$

To conclude the proof, it is now sufficient to show that, for every $u \in U$, there is some $u' \neq u$ with

$$\langle \mathfrak{T}_{\text{bin}}, U, \{u\} \rangle \equiv_{\text{MSO}}^r \langle \mathfrak{T}_{\text{bin}}, U, \{u'\} \rangle.$$

Then it follows that

$$\mathfrak{T}_{\text{bin}} \models \varphi(u, U) \quad \text{iff} \quad \mathfrak{T}_{\text{bin}} \models \varphi(u', U),$$

which is the desired contradiction.

Hence, consider an element $u \in U$. Let u_m be the shortest prefix of u with $u_m \in U$ and, for $i < m$, let u_i be the longest prefix of u_m with $u_i \in V_i \setminus V_{i+1}$. Set

$$v_i := u_{l+1} o^i, \quad \text{where } l \text{ is the index from above.}$$

By choice of m , there are indices $i < j < m$ such that

$$t|_{v_i}, \{u\} \equiv_{\text{MSO}}^r t|_{v_j}, \{u\}.$$

Hence, it follows by Proposition V.1.7 that

$$t|_{v_m}, \{u\} \equiv_{\text{MSO}}^r t|_{v_{m+i-j}}, \{u\} \cong t|_{v_m}, \{u'\}, \quad \text{for some } u',$$

which implies that $\langle \mathfrak{T}_{\text{bin}}, U, \{u\} \rangle \equiv_{\text{MSO}}^r \langle \mathfrak{T}_{\text{bin}}, U, \{u'\} \rangle$. \square

As an application of the theorem (and the construction used in its proof), we present a game where the winning strategies are not definable.

Theorem 6.8. *There exists a game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega, W \rangle$ with a regular winning condition $W \subseteq C^*$ (where $\Omega : V \rightarrow C$) such that Player \square has a winning strategy but not an MSO -definable one.*

Proof. The game graph $\langle V_\diamond, V_\square, E \rangle$ is the infinite binary tree with partition

$$V_\diamond := 1^* \quad \text{and} \quad V_\square := 1^* o 1^*.$$

To define the winning condition, let

$$U_n := (o + 1)^* (o^n o^* 1)^n (o + 1)^*$$

be the language from the proof of Theorem 6.7. We define the priority function $\Omega : V \rightarrow \{1, 2, 3\}$ by

$$\Omega(1^n) := 3 \quad \text{and} \quad \Omega(1^n o u) := \begin{cases} 1 & \text{if } u \in U_n, \\ 2 & \text{otherwise,} \end{cases}$$

and we say that a play $(u_n)_{n < \omega}$ is winning for Player \diamond if

$$\inf_{n < \omega} \Omega(u_n) \text{ is even.}$$

(Note that this differs from a parity condition in that we are not taking the \liminf , that is, we do not look at the priorities that occur infinitely often, but at those that occur at least once.)

First, note that Player \square has a winning strategy in this game. Player \diamond cannot chose to stay on the right-most branch forever since this branch has priority 3. Hence, she must choose to move away from this branch at some point. In the rest of the game, Player \square can choose to move to a position of the form $i^n o u$ with $u \in U_n$. The minimal priority of the resulting play is 1. Hence, Player \square wins.

We claim that he does not have a definable winning strategy. For a contradiction, suppose otherwise. Then there exists a formula that, given some position of the form $i^n o$, produces an infinite branch β_n that starts with $i^n o$ and that contains some position $i^n o u$ with $u \in U_n$. Using this formula we can construct an MSO-formula $\varphi(x, y, X)$ such that

$$\mathfrak{T}_{\text{bin}} \models \varphi(i^n o, i^n o u, i^n o U_n) \quad \text{iff} \quad u \text{ is the minimal word such that } i^n o u <_{\text{pf}} \beta_n \text{ and } u \in U_n.$$

Let \mathfrak{T}_n be the graph obtained from $\mathfrak{T}_{\text{bin}}$ by removing the subtree rooted at $i^n o$. For every MSO_r-theory θ , there exists a formula $\psi_\theta(x, X)$ such that

$$\mathfrak{T}_{\text{bin}} \models \varphi(i^n o, i^n o u, i^n o U_n) \quad \Leftrightarrow \quad \mathfrak{T}_{\text{bin}} \models \psi_\theta(u, U_n),$$

where $\theta := \text{Th}_{\text{MSO}}^r(\mathfrak{T}_n)$. Fix a linear ordering on all MSO_r-theories and let $\vartheta(x, X)$ be the formula stating that $\psi_\theta(x, X)$ holds where θ is the least theory such that $\psi_\theta(x, X)$ holds for exactly one vertex x . Then $\vartheta(x, X)$ chooses a unique element from every set U_n . By (the proof of) Theorem 6.7 this is not possible. \square

Remark. Note that, while the *winning strategy* of this game is not definable, the *winning regions* (the sets \emptyset and V) clearly are. \lrcorner

It follows that all trees with choice are thin. The proof that they are even tame is split into the following two lemmas.

Lemma 6.9. *Let $\mathfrak{T} = \langle T, \leq, \sqcap, \bar{P} \rangle$ be a coloured generalised tree, $\beta \subseteq T$ a branch, and $Z \subseteq \beta$ a chain of order type \mathbb{Z} such that*

$$\mathfrak{T}_\beta[u, v], u \equiv_{\text{MSO}}^m \mathfrak{T}_\beta[u', v'], u', \quad \text{for all } u < v \text{ and } u' < v' \text{ in } Z,$$

where

$$\mathfrak{T}_\beta[u, v] := \mathfrak{T}_\beta[I] \quad \text{with} \quad I := \{x \in \downarrow Z \mid u \leq z < v\}$$

are the structures from Proposition V.1.9. Then no MSO_m-formula $\psi(x, X)$ (without parameters) defines a choice function for \mathfrak{T} .

Proof. For a contradiction, suppose that there exists an MSO_m-formula $\psi(x, X)$ defining a choice function. Let $w \in Z$ be the element such that $\mathfrak{T} \models \psi(w, Z)$. Setting

$$[u, v] := \uparrow u \setminus \uparrow v, \quad I := T \setminus \bigcup_{u \in Z} \uparrow u \quad \text{and} \quad J := \bigcap_{u \in Z} \uparrow u,$$

we can decompose \mathfrak{T} as in Proposition V.1.9 into the pieces

$$I, \quad \dots, \quad [w - 1, w), \quad [w, w + 1), \quad [w + 1, w + 2), \quad \dots, \quad J,$$

where $w + 1$ denotes the successor of w in Z . The corresponding theories are

$$\begin{aligned} \sigma &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_\beta[I], Z, w), \\ \tau &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_\beta[J], Z, w), \\ \theta_* &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_\beta[w + 1, w + 2), Z, w), \\ \theta_u &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_\beta[u, u + 1), Z, u), \quad \text{for } u \in Z. \end{aligned}$$

Since $\mathfrak{T}[u, u + 1], u \equiv_{\text{MSO}}^m \mathfrak{T}[w, w + 1], w$ implies

$$\mathfrak{T}[u, u + 1], \{u\} \equiv_{\text{MSO}}^m \mathfrak{T}[w, w + 1], \{w\} = \theta_*$$

and $\theta_u = \mathfrak{T}[u, u + 1], \{u\}, u \equiv_{\text{MSO}}^m \mathfrak{T}[w, w + 1], \{w\}, w = \theta_w$,

it follows that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{T}, \beta, Z, u) = \sigma + \sum_{v < u} \theta_* + \theta_u + \sum_{v > u} \theta_* + \tau$$

(if we define the sum of these types as in Proposition V.1.9), and that this expression does not depend on u . Consequently, we have

$$\mathfrak{T}, Z, u \equiv_{\text{MSO}}^m \mathfrak{T}, Z, v, \quad \text{for all } u, v \in Z.$$

A contradiction to the fact that $\mathfrak{T} \models \psi(w, Z) \wedge \neg \psi(w + 1, Z)$. \square

Lemma 6.10. *Every generalised tree \mathfrak{T} with weak MSO-choice is tame.*

Proof. Suppose that \mathfrak{T} is not tame. We distinguish three cases depending on which of the three conditions fails.

(i) For a contradiction, suppose that \mathfrak{T} has unbounded branching degree but there exists an MSO_m -formula $\varphi(x, X; \bar{P})$ with k parameters \bar{P} defining a choice function for \mathfrak{T} . To simplify notation, we again set $\mathfrak{T}^+ := \langle \mathfrak{T}, \bar{P} \rangle$. Let n be the number of MSO_m -theories of structures of the form $\langle S, \leq_{\text{pf}}, \sqcap, \bar{Q}, v \rangle$ with k unary predicates \bar{Q} and a constant v . By assumption, there is some vertex $w \in T$ whose branching degree is larger than n . Let \sim be the equivalence relation used to define the branching degree of v . Then there are two \sim -classes S and S' and vertices $s \in S$ and $s' \in S'$ such that

$$\langle \mathfrak{T}^+|_S, s \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}^+|_{S'}, s' \rangle.$$

Set $C := \{s, s'\}$. By Proposition V.1.7, it follows that

$$\mathfrak{T} \models \varphi(s, C; \bar{P}) \quad \text{iff} \quad \mathfrak{T} \models \varphi(s', C; \bar{P}).$$

Hence, φ does not define a choice function.

(ii) Suppose that, for every $k < \omega$, there is some branch β with $\text{HR}(\beta) > k$. Let $\varphi(x, X; \bar{P})$ be an MSO_m -formula with n monadic parameters \bar{P} . We have to show that φ does not define a choice function for \mathfrak{T} . To simplify notation, set $\mathfrak{T}^+ := \langle \mathfrak{T}, \bar{P} \rangle$.

We choose $k < \omega$ larger than the number of MSO_m -theories of generalised trees of the form $\langle T, \leq, \sqcap, \bar{P}, \beta, p, a \rangle$ with a distinguished branch β , a unary function p , k unary predicates \bar{P} , and a constant a . By assumption, there is some branch β with $\text{HR}(\beta) > k$. We define an additive labelling λ on β by setting

$$\lambda(u, v) := \text{Th}_{\text{MSO}}^m(\mathfrak{T}^+[u, v], u), \quad \text{for } u < v \text{ in } \beta.$$

Since λ uses less than k colours, we can use Theorem 2.5 to find a subset $Z \subseteq \beta$ of order type \mathbb{Z} such that

$$\lambda(u, v) = \lambda(u', v'), \quad \text{for all } u < v \text{ and } u' < v' \text{ in } Z.$$

By Lemma 6.9, it follows that φ does not define a choice function for \mathfrak{T}^+ .

(iii) Finally, suppose for a contradiction that \mathfrak{T} is not thin but that there exists a formula $\varphi(x, X; \bar{P})$ with parameters \bar{P} defining a choice function for \mathfrak{T} . Without loss of generality, we may assume that φ only uses the tree order \leq , but not the infimum operation \sqcap . We will derive a contradiction to Theorem 6.7 by defining a choice function for the binary tree $\mathfrak{S}_{\text{bin}}$.

As \mathfrak{T} is not thin, there exists a subset $S \subseteq T$ such that $\mathfrak{T}|_S$ is isomorphic to $\mathfrak{S}_{\text{bin}}$. We denote the root of S by s_o and the two successors of a vertex $v \in S$ by $\text{suc}_o(v)$ and $\text{suc}_i(v)$. Let \mathfrak{T}^+ be the structure obtained from \mathfrak{T} by adding the predicates \bar{P} and removing the infimum operation \sqcap . We decompose \mathfrak{T}^+ into the sets

$$U_* := T \setminus \uparrow s_o \quad \text{and} \quad U_v := \uparrow v \setminus \uparrow \{\text{suc}_o(v), \text{suc}_i(v)\}, \quad \text{for } v \in S.$$

Then we can express \mathfrak{T}^+ as a generalised sum followed by a quantifier-free interpretation τ as

$$\mathfrak{T}^+ \cong \tau\left(\sum_{i \in \mathfrak{S}} \mathfrak{U}_i\right),$$

where $\mathfrak{S} := \{\{*\} + \{o, i\}^*, \leq_{\text{pf}}, \text{suc}_o, \text{suc}_i\}$ and

$$\mathfrak{U}_v := \langle \mathfrak{T}^+|_{U_v}, I_o, I_i, v \rangle \quad \text{with} \quad I_i := U_v \cap \downarrow \text{suc}_i(v),$$

for $v \in \{*\} + \{o, i\}^*$. The interpretation τ is given by

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= [x \sim y \wedge x \leq_U y] \\ &\vee \left[x \not\sim y \wedge \bigvee_{i < 2} (\text{suc}_i(x) \leq_I y \wedge I_i x) \right], \end{aligned}$$

where \sim is the equivalence relation induced by the partition, \leq_U is the ordering of the components \mathfrak{U}_v , \leq_I the ordering of the index structure \mathfrak{S} , and I_i are the additional predicates added to \mathfrak{U}_v . By the composition theorem for generalised sums, we obtain formulae ψ_o and $\chi_o, \dots, \chi_{n-1}$ such that

$$\mathfrak{T} \models \varphi(a, B; \bar{P}) \quad \text{iff} \quad \mathfrak{S} \models \psi_o(\llbracket \chi_o(a, B) \rrbracket, \dots, \llbracket \chi_{n-1}(a, B) \rrbracket).$$

For a vertex v of \mathfrak{I} , set

$$\theta_v(a, B) := \{ i < n \mid v \in \llbracket \chi_i(a, B) \rrbracket \}.$$

Note that, if $a \in S$ and $B \subseteq S$, then, for every fixed v , there are only 4 possible values of $\theta_v(a, B)$, depending on

whether or not $v = a$ and whether or not $v \in B$.

We denote these values by $\theta_v^0, \theta_v^1, \theta_v^2, \theta_v^3$. Furthermore, for $v = *$, the value $\theta_*(a, B)$ does not depend on a and B . We denote it by θ_* . Let \mathfrak{K} be the structure obtained from \mathfrak{I} by labelling the vertex $*$ by θ_* and every other vertex v by the four values $\theta_v^0, \theta_v^1, \theta_v^2, \theta_v^3$. Note that, given a and B , a formula can use this labelling to decide to which sets of the form

$$\llbracket \chi_i(a, B) \rrbracket$$

a given vertex belongs. Consequently, there exists a formula $\psi_I(x, X)$ such that

$$\mathfrak{K} \models \psi_I(a, B) \quad \text{iff} \quad \mathfrak{I} \models \psi_o(\llbracket \chi_0(a, B) \rrbracket, \dots, \llbracket \chi_{n-1}(a, B) \rrbracket).$$

Finally, note that we can write

$$\mathfrak{K} \cong \sigma(1 \oplus \langle \mathfrak{T}_{\text{bin}}, \bar{Q} \rangle),$$

for some interpretation σ and some unary predicates \bar{Q} . Hence, there is some formula $\psi(x, X; \bar{Q})$ with parameters \bar{Q} such that

$$\mathfrak{T}_{\text{bin}} \models \psi(a, B; \bar{Q}) \quad \text{iff} \quad \mathfrak{K} \models \psi_I(a, B).$$

For all a and B in S , it therefore follows that

$$\mathfrak{T}_{\text{bin}} \models \psi(a, B; \bar{Q}) \quad \text{iff} \quad \mathfrak{T} \models \varphi(a, B; \bar{P}).$$

Hence, the fact that φ defines a choice function implies that so does ψ . Consequently, $\mathfrak{T}_{\text{bin}}$ has weak MSO-choice. A contradiction to Theorem 6.7. \square

Tame Trees

It remains to prove the converse, that tame trees have choice. We start by defining a well-ordering for linear orders of bounded Hausdorff rank.

Lemma 6.11. *Let $n < \omega$.*

(a) *There exists an MSO-formula $\varphi(x, y; \bar{Z})$ with the following property. For every linear order \mathfrak{A} with $\text{HR}(\mathfrak{A}) \leq n$, there are monadic parameters \bar{P} such that $\varphi(x, y; \bar{P})$ defines a well-ordering \sqsubseteq on A .*

(b) *There exists an MSO-formula $\psi(x, y; \bar{Z})$ with the following property. For every linear order \mathfrak{A} with $\text{HR}(\mathfrak{A}) \leq n$, there are monadic parameters \bar{Q} such that, in the well-order $\langle A, \sqsubseteq \rangle$ defined by $\varphi(x, y; \bar{P})$, the formula $\psi(x, y; \bar{Q})$ defines the original ordering \leq .*

Proof. We construct both formulae by induction on n . First, note that the formula

$$\text{wo} := \exists x \forall y (x \leq y) \wedge \forall x \exists y [x < y \wedge \neg \exists z (x < z < y)]$$

holds in every well-orders, but it fails for the opposite of every infinite well-order.

If $n \leq 1$, then one of \mathfrak{A} and \mathfrak{A}^{op} is a well-order. Hence, we can take the formula

$$\varphi(x, y) := [\text{wo} \wedge x \leq y] \vee [\neg \text{wo} \wedge x \geq y]$$

to well-order A . Conversely, using one monadic parameter $Q = A$ or $Q = \emptyset$, we can set

$$\psi(x, y; Q) := [Qx \wedge x \leq y] \vee [\neg Qx \wedge x \geq y].$$

For the inductive step, suppose that $\text{HR}(\mathfrak{A}) = n + 1$. By definition, this means we can decompose \mathfrak{A} as

$$\mathfrak{A} = \sum_{i \in I} \mathfrak{B}_i$$

where I is a well-order or the opposite of a well-order and each factor \mathfrak{B}_i has Hausdorff rank at most n . By inductive hypothesis, there exist formulae $\varphi_o(x, y; \bar{Z})$ and $\psi_o(x, y; \bar{Z}')$ and parameters \bar{P}^i and \bar{Q}^i in \mathfrak{B}_i such that $\varphi_o(x, y; \bar{P}^i)$ well-orders \mathfrak{B}_i and $\psi_o(x, y; \bar{Q}^i)$ defines the original order. Let \sim be the equivalence relation whose classes are the sets B_i . By Lemma 3.1, we can define \sim using a single monadic parameter R_* . We define a well-order \sqsubseteq on A by

$$\begin{aligned} \varphi(x, y; \bar{R}R_*) := & [\text{wo}^\sim \wedge x \not\sim y \wedge x \leq y] \\ & \vee [\neg \text{wo}^\sim \wedge x \not\sim y \wedge x \geq y] \\ & \vee [x \sim y \wedge \varphi_o^{[x]^\sim}(x, y; \bar{R})], \end{aligned}$$

where

$$R_j := \bigcup_{i \in I} P_j^i, \quad \text{for } j < n - 1,$$

$\varphi_o^{[x]^\sim}$ denotes the relativisation of φ_o to the class $[x]^\sim$, and wo^\sim is the variant of the formula wo talking about $A/\sim \cong I$ (that is, the formula obtained from wo by replacing every atom of the form $x < y$ by $x \leq y \wedge x \not\sim y$).

To define ψ , we use one additional monadic parameter $S' = A$ or $S' = \emptyset$ and we set

$$\begin{aligned} \psi(x, y; R_* \bar{S}S') := & [x \sim y \wedge \psi_o^{[x]^\sim}(x, y)] \vee [x \not\sim y \wedge S'x \wedge x < y] \\ & \vee [x \not\sim y \wedge \neg S'x \wedge x > y], \end{aligned}$$

where

$$S_j := \bigcup_{i \in I} Q_j^i, \quad \text{for } j < 2n - 1. \quad \square$$

Given a tame tree, we can use this lemma to define a well-ordering on every branch. Then we can combine these well-orderings to one of the whole tree by fixing a skeleton and well-ordering its classes. To do so, we have to show that the classes of a skeleton are definable.

Lemma 6.12. *Let \mathfrak{T} be a generalised tree whose branching degree is bounded by $n < \omega$. There exists an MSO-formula $\psi(x, y; \bar{Z})$ such that, for every skeleton S , there are monadic parameters \bar{P} such that*

$$\mathfrak{T} \models \psi(u, v; \bar{P}) \quad \text{iff} \quad u, v \in \beta, \quad \text{for some } \beta \in S.$$

Proof. Note that every skeleton S induces a well-founded tree $\langle S, \leq \rangle$ with order

$$\beta \leq \gamma \quad : \text{iff} \quad \beta \cap \Downarrow \gamma \neq \emptyset.$$

Since the branching degree of \mathfrak{T} is bounded by n , we can choose a colouring $\lambda_o : S \rightarrow [n + 1]$ such that

- ♦ $\lambda_o(\gamma) \neq \lambda_o(\beta)$, for all $\gamma \in \text{Suc}(\beta)$, and
- ♦ $\lambda_o(\gamma) \neq \lambda_o(\gamma')$, for all $\gamma, \gamma' \in \text{Suc}(\beta)$ that are attached to the same vertex $v \in \beta$, that is, that satisfy $\beta \cap \Downarrow \gamma = \Downarrow v = \beta \cap \Downarrow \gamma$.

This colouring induces a corresponding colouring $\lambda : T \rightarrow [n + 1]$. It follows that $u, v \in T$ belong to the same $\beta \in S$ if, and only if,

- ♦ u and v are comparable,
- ♦ $\lambda(u) = \lambda(v)$, and
- ♦ $\lambda(u) = \lambda(w)$, for every w between u and v .

These conditions can be expressed in MSO with the help of the relations $P_i := \lambda^{-1}(i)$. \square

As explained above, it remains to well-order the classes of a skeleton. This can be done as follows.

Lemma 6.13. *Let $\mathfrak{T} = \langle T, \leq_{\text{pf}}, \leq_{\text{so}} \rangle$ be a well-founded successor-ordered tree where the successor-ordering \leq_{so} well-orders every set $\text{Suc}(v)$, $v \in T$. Then the lexicographic ordering*

$$u \leq_{\text{lex}} v \quad : \text{iff} \quad u \leq_{\text{pf}} v \quad \text{or} \quad u_o \sqsubseteq_{\text{so}} v_o, \quad \text{where } u_o \text{ and } v_o \text{ are the successors of } u \sqcap v \text{ with } u_o \leq_{\text{pf}} u \text{ and } v_o \leq_{\text{pf}} v,$$

forms a well-ordering on T .

Proof. For a contradiction, suppose that there exists a strictly decreasing sequence $u_0 >_{\text{lex}} u_1 >_{\text{lex}} \dots$. Set $w_{ij} := u_i \sqcap u_j$. Since the tree-order \leq_{pf} has no strictly decreasing chains, we can use the Theorem of Ramsey to find an infinite subsequence $(u_i)_{i \in I}$ such that

$$v_{ij} \sqsubseteq_{\text{so}} v'_{ij}, \quad \text{where } v_{ij} \text{ and } v'_{ij} \text{ are the successors of } w_{ij} \text{ with} \\ v_{ij} \leq_{\text{pf}} u_i \text{ and } v'_{ij} \leq_{\text{pf}} u_j,$$

for all $i < j$ in I . Applying the Theorem of Ramsey one more time, we obtain an infinite subsequence $(u_i)_{i \in J}$ such that one of the following conditions is satisfied

$$\begin{aligned} w_{ij} &\leq_{\text{pf}} w_{ik}, & \text{for all } i < j < k \text{ in } J, \\ w_{ij} &= w_{ik}, & \text{for all } i < j < k \text{ in } J, \\ w_{ij} &>_{\text{pf}} w_{ik}, & \text{for all } i < j < k \text{ in } J. \end{aligned}$$

The first case is impossible since \leq_{pf} has no infinite descending chains, and the third case is impossible since \mathfrak{T} has no infinite branches. Note that, for $i < j < k$ in J ,

$$w_{ij} = w_{ik} \quad \text{implies} \quad w_{ij} \leq_{\text{pf}} w_{jk}.$$

Fix an increasing enumeration $j_0 < j_1 < \dots$ of J . We obtain an increasing sequence

$$w_{j_0 j_1} \leq_{\text{pf}} w_{j_1 j_2} \leq_{\text{pf}} w_{j_2 j_3} \leq_{\text{pf}} \dots$$

Since \mathfrak{T} has no infinite branches, this sequence must stabilise at some point. Thus, there exists a constant $k < \omega$ such that

$$w_{j_k j_{k+1}} = w_{j_{k+2} j_{k+3}} = \dots$$

Let v_i be the successor of $w := w_{j_k j_{k+1}}$ with $v_i \leq_{\text{pf}} u_{j_{k+i}}$. These successors form an infinite descending sequence $v_0 >_{\text{so}} v_1 >_{\text{so}} \dots$. A contradiction to the fact that \leq_{so} well-orders the successor set $\text{Suc}(w)$. \square

Lemma 6.14. *For every tame generalised tree \mathfrak{T} , there exists an MSO-formula $\varphi(x, y)$ with monadic parameters that defines a well-ordering on T .*

Proof. Let \mathfrak{T} be tame and let $k, n < \omega$ be the corresponding bounds. Then we can use Proposition V.5.9 to find a skeleton S of \mathfrak{T} . Let $\psi(x, y; \bar{P})$ be the formula from Lemma 6.12 defining S . Note that, by the definition of a skeleton, the set S forms a well-founded tree $\mathfrak{S} = \langle S, \leq \rangle$ with ordering

$$\beta \leq \gamma \quad : \text{iff} \quad \beta \cap \Downarrow \gamma \neq \emptyset,$$

and that the ordering \leq is MSO-definable using the formula $\psi(x, y; \bar{P})$.

Furthermore, since the Hausdorff rank of every $\beta \in S$ is bounded by k , we can use Lemma 6.11 (a) to find an MSO-formula $\chi(x, y; \bar{Q})$ with monadic parameters \bar{Q} defining a well-ordering \trianglelefteq_β on each $\beta \in S$.

Below we will construct a formula $\vartheta(X, Y, Z; \bar{R})$ with monadic parameters \bar{R} such that, for every $\beta \in S$, the formula $\vartheta(X, Y, \beta; \bar{R})$ defines a well-ordering $\sqsubseteq_{\text{so}}^\beta$ on $\text{Suc}(\beta)$. Let \leq_{lex} be the lexicographic ordering on the corresponding successor-ordered tree $\langle S, \leq, \sqsubseteq_{\text{so}} \rangle$. Since S is well-founded, it follows by Lemma 6.13 that \leq_{lex} is a well-ordering.

We obtain the desired well-ordering \sqsubseteq on T by setting

$$\begin{aligned} x \sqsubseteq y \quad : \text{iff} \quad & (\exists \beta \in S)[x, y \in \beta \wedge x \trianglelefteq_\beta y] \\ & \vee (\exists \beta, \gamma \in S)[x \in \beta \wedge y \in \gamma \wedge \beta <_{\text{lex}} \gamma]. \end{aligned}$$

This condition can be expressed in MSO using the parameters $\bar{P}\bar{Q}\bar{R}$. Furthermore, \sqsubseteq is a well-ordering since the orderings \trianglelefteq_β and \leq_{lex} are well-orderings.

It remains to find the formulae ϑ . Note that, if γ is a successor of β in \mathfrak{S} , there is some vertex $v \in \beta$ with

$$\beta \cap \Downarrow \gamma = \Downarrow v.$$

We say that γ is *attached* at v . Since the branching degree of \mathfrak{T} is bounded by n , there are at most n successors $\gamma \in \text{Suc}(\beta)$ attached at each vertex v . Hence, we can choose a colouring $\lambda_o : S \rightarrow [n]$ such that

$$\lambda_o(\gamma) \neq \lambda_o(\gamma'),$$

for all $\gamma, \gamma' \in \text{Suc}(\beta)$ that are attached to the same vertex $v \in \beta$. This colouring induces a corresponding colouring $\lambda : T \rightarrow [n]$. We can now define the desired successor-ordering by

$$\gamma \sqsubseteq_{\beta}^{\beta} \gamma' \quad : \text{iff} \quad \begin{array}{l} \gamma \text{ is attached at } v, \gamma' \text{ is attached at } v', \text{ and} \\ v \triangleleft_{\beta} v' \text{ or } [v = v' \text{ and } \lambda_o(\gamma) \leq \lambda_o(\gamma')] \end{array}.$$

This can clearly be expressed in MSO using the parameters \tilde{Q} and

$$R_i := \lambda^{-1}(i). \quad \square$$

Proof of Theorem 6.3. (2) \Rightarrow (1) is trivial since we can define a choice function by returning the minimal element of the given set X , while (1) \Rightarrow (3) \Rightarrow (2) have been proved in Lemmas 6.10 and 6.14, respectively. \square

Well-Ordered Trees

In light of Theorem 6.3 it is interesting to study trees that can be well-ordered. The canonical well-ordering on a tree is the so-called *length-lexicographic* one.

Definition 6.15. Let D be a set. The *length-lexicographic order* \leq_{lex} on D^* is defined by

$$u \leq_{\text{lex}} v \quad : \text{iff} \quad \begin{array}{l} |u| < |v|, \text{ or} \\ |u| = |v| \text{ and } u = wcx, v = wdy, \text{ for some} \\ c, d \in D \text{ and } w, x, y \in D^* \text{ with } c \leq d. \end{array}$$

Proposition 6.16. *There exists an MSO-formula χ such that*

$$\mathfrak{A} \models \chi \quad \text{iff} \quad \mathfrak{A} \cong \langle \{o, i\}^*, \text{suc}_o, \text{suc}_i, \leq_{\text{lex}} \rangle,$$

for all structures $\mathfrak{A} = \langle A, S_o, S_i, \leq \rangle$.

Proof. Let \leq_{pf} be the transitive closure of $S_o \cup S_i$ and let E be the relation consisting of all pairs $\langle u, v \rangle$ such that $u \leq v$ and the function

$$f(u') = \min \{ v' \leq_{\text{pf}} v \mid u' \leq v' \} \quad (\text{minimum with respect to } \leq_{\text{pf}})$$

is a bijection between $\Downarrow u$ and $\Downarrow v$ (again with respect to \leq_{pf}). Since \leq_{pf} and E are MSO-definable, we can write down an MSO-formula expressing the following conditions.

- ♦ Every vertex has exactly one S_o -successor and exactly one S_i -successor, and these are different.
- ♦ For every vertex v , the set $\Downarrow v$ is linearly ordered by \leq_{pf} .
- ♦ There exists a least element with respect to \leq_{pf} .
- ♦ E is a partial order that is a union of infinitely many finite chains.
- ♦ If u_o is the S_o -successor of v and u_i is its S_i -successor, then $\langle u_o, u_i \rangle \in E$.
- ♦ If u' is a successor of u and v' is one of v , then $\langle u, v \rangle \in E \Leftrightarrow \langle u', v' \rangle \in E$.
- ♦ We have $u \leq v$ if, and only if,
 - $\langle u, v \rangle \in E$, or
 - there is some $v' < v$ with $\langle u, v' \rangle \in E$, or
 - there is some $v' < v$ with $\langle v', u \rangle \in E$.

Note that the first three points ensure that the relations S_o and S_i are the edge relations of the infinite binary tree and that \leq_{pf} is the corresponding prefix relation. The next three points imply that E the left-to-right ordering on each level of the tree and the last point guarantees that \leq is the length-lexicographic order.

Concerning our claim about E , note that, by definition of f , we know that $\langle u, v \rangle \in E$ implies $|u| = |v|$. Conversely, one can use the above axioms to show by induction on $|u|$ that $|u| = |v|$ implies $\langle u, v \rangle \in E$ or $\langle v, u \rangle \in E$. \square

Our next goal is to show that every binary tree equipped with a well-ordering is as least as complicated as $\langle \mathbb{T}_{\text{bin}}, \leq_{\text{lex}} \rangle$.

Definition 6.17. (a) A *well-ordered tree* is a structure of the form

$$\mathfrak{T} = \langle T, \text{suc}_o, \text{suc}_i, \leq \rangle,$$

where $\langle T, \text{suc}_o, \text{suc}_i \rangle$ is isomorphic to the infinite binary tree and \leq is a well-order on T .

(b) Let \mathfrak{S} and \mathfrak{T} be well-ordered trees. An *embedding* $\varphi : \mathfrak{S} \rightarrow \mathfrak{T}$ is an injective function satisfying

- ♦ $\text{suc}_d(\varphi(v)) \leq_{\text{pf}} \varphi(\text{suc}_d(v))$, for all vertices v .
- ♦ $u \leq v$ implies $\varphi(u) \leq \varphi(v)$.

We will show that $\langle \mathbb{T}_{\text{bin}}, \leq_{\text{lex}} \rangle$ can be embedded into every well-ordered tree.

Lemma 6.18. *For every MSO-formula φ , there exists an MSO-formula $\varphi^*(Z)$ such that, for every well-ordered tree \mathfrak{T} and every set $U \subseteq T$,*

$$\mathfrak{T} \models \varphi^*(U) \quad \text{iff} \quad U \text{ is the range of some embedding } \varphi : \mathfrak{S} \rightarrow \mathfrak{T} \text{ such that } \mathfrak{S} \text{ satisfying } \varphi.$$

Proof. There exists a formula $\vartheta(Z)$ expressing that Z induces a well-ordered tree. Consequently, we can set $\varphi^*(Z) := \vartheta(Z) \wedge \varphi^{(Z)}$, where $\varphi^{(Z)}$ denotes the relativisation of φ to Z . \square

Definition 6.19. Let \mathfrak{T} be a well-ordered tree. We say that a vertex $w \in T$ has *enough successors* if, for every pair of vertices $u, v \geq_{\text{pf}} w$, there is some $v' \geq_{\text{pf}} v$ with $u < v'$.

Lemma 6.20. *Every well-ordered tree has a vertex with enough successors.*

Proof. For a contradiction, suppose that there exists a well-ordered tree \mathfrak{T} where no vertex has enough successors. By induction on i , we construct vertices u_i and v_i such that

$$u_i > u_{i+1} \quad \text{and} \quad u_i \geq w, \quad \text{for all } w \geq_{\text{pf}} v_i.$$

Since $u_0 > u_1 > \dots$ is an infinite descending chain this will contradict the fact that \leq is a well-order.

For $i = 0$, note that the root of \mathfrak{T} does not have enough successors. Hence, there exists a pair of vertices u and v such that $u \geq v'$, for all $v' \geq_{\text{pf}} v$. We set $u_0 := u$ and we choose for v_0 some vertex $v_0 \geq_{\text{pf}} v$ with $v_0 \not\leq_{\text{pf}} u$. By choice of u and v it follows that $u_0 > v'$, for all $v' \geq_{\text{pf}} v_0$.

For the inductive step, suppose that we have already defined u_i and v_i . Since v_i does not have enough successors, there exist vertices $u, v \geq_{\text{pf}} v_i$ such that $u \geq v'$, for all $v' \geq_{\text{pf}} v$. As above we set $u_{i+1} := u$ and we choose for v_{i+1} some vertex $v_{i+1} \geq_{\text{pf}} v$ with $v_{i+1} \not\leq_{\text{pf}} u$. Then $u_{i+1} \geq_{\text{pf}} v_i$ implies that $u_i > u_{i+1}$. \square

Lemma 6.21. *Let α be an ordinal.*

- (a) $\langle 2^\alpha, \leq_{\text{lex}} \rangle$ forms a complete linear order.
- (b) $u = \inf H$ implies that, for every prefix u_0 of u , there is some $w \in H$ with $u_0 <_{\text{pf}} w$.

Proof. Clearly, \leq_{lex} is linear. Hence, it remains to prove completeness. Fix a set $H \subseteq 2^\alpha$. We have to show that H has an infimum. We construct an increasing sequence $u_0 <_{\text{pf}} u_1 <_{\text{pf}} \dots$ of length α such that $u_i \in 2^i$, for $i \leq \alpha$, and

$$u_i \circ^{\gamma_i} \leq_{\text{lex}} w, \quad \text{for all } w \in H \text{ and } i < \alpha$$

(where γ_i is the ordinal such that $i + \gamma_i = \alpha$). We start with $u_0 := \langle \rangle$. For the successor step, suppose that we have already defined u_i . If there is some $w \in H$ with $u_i \circ <_{\text{pf}} w$, we set $u_{i+1} := u_i \circ$. Otherwise, $u_{i+1} := u_i \mathbf{I}$. For the limit step, let δ be a limit ordinal and suppose that we have already defined u_i , for all $i < \delta$. We take for u_δ the limit of all u_i .

Let $u := u_\alpha$ be the limit of the sequence constructed in this way. Then u satisfies (b) and it remains to prove that $u = \inf H$. By construction, we have $u \leq_{\text{lex}} w$, for all $w \in H$. Conversely, let v be a lower bound of H . To prove that $v \leq_{\text{lex}} u$ it is sufficient to show that

$$v \leq_{\text{lex}} u_i \mathbf{I}^{\gamma_i}, \quad \text{for all } i \leq \alpha.$$

We do so by induction on i . Clearly, $v \leq_{\text{lex}} \mathbf{I}^\alpha = u_0 \mathbf{I}^{\gamma_0}$. For the inductive step, suppose that $v \leq_{\text{lex}} u_i \mathbf{I}^{\gamma_i}$. If $u_{i+1} = u_i \mathbf{I}$, we obtain $v \leq_{\text{lex}} u_{i+1} \mathbf{I}^{\gamma_{i+1}}$. Hence, suppose otherwise. Then there is some $w \in H$ with $u_{i+1} = u_i \circ <_{\text{pf}} w$. If u_i is not a prefix of v , it follows that $v <_{\text{lex}} u_i \circ^{\gamma_{i+1}}$ and we are done. Otherwise, $v \leq_{\text{lex}} w$ implies that $u_{i+1} = u_i \circ$ is also a prefix of v . Hence, $v \leq_{\text{lex}} u_{i+1} \mathbf{I}^{\gamma_{i+1}}$. \square

Lemma 6.22. Let $\mathfrak{T}_{\text{lex}} := \langle \{0, 1\}^*, \text{succ}_0, \text{succ}_1, \leq_{\text{lex}} \rangle$. For every well-ordered tree \mathfrak{T} , there exists an embedding $\varphi : \mathfrak{T}_{\text{lex}} \rightarrow \mathfrak{T}$ such that $\text{rng } \varphi$ forms an MSO-definable subset of T .

Proof. We define $\varphi(u)$ by induction on u with respect to \leq_{lex} . We start by using Lemma 6.20 to find a vertex z with enough successors and setting

$$\varphi(\langle \rangle) := z.$$

For the inductive step, suppose that we have already defined $\varphi(w)$. Let w' be the \leq_{lex} -successor of w and suppose that $w' = ud$ for $d \in \{0, 1\}$ and $u \in \{0, 1\}^*$. Clearly, $u \leq_{\text{lex}} w$. Since z has enough successors, there exists a vertex $v \geq_{\text{pf}} \varphi(u)d$ with $\varphi(w) < v$. We set

$$\varphi(w') := v.$$

We have found an embedding $\varphi : \mathfrak{T}_{\text{lex}} \rightarrow \mathfrak{T}$. It remains to show that we can choose φ such that $U := \text{rng } \varphi$ is MSO-definable. We choose for U the lexicographically minimal such set, where the lexicographic ordering on subsets is defined by

$$U \leq_{\text{lex}} V \quad : \text{iff} \quad U = V \text{ or the } \leq_{\text{lex}}\text{-minimal element belonging to one of } U \text{ and } V, \text{ but not to both, belongs to } V.$$

By Lemma 6.18 and Proposition 6.16, there exists a formula $\vartheta(Z)$ stating that Z is the range of some embedding of $\mathfrak{T}_{\text{lex}}$. Since the lexicographic order \leq_{lex} on sets is also MSO-definable, it follows that there exists an MSO-formula $\psi(Z)$ saying that Z is the \leq_{lex} -minimal set satisfying $\vartheta(Z)$. Clearly, this set Z is unique.

Hence, it remains to prove that such a set Z exists. Since

$$\langle \wp(\{0, 1\}^*), \leq_{\text{lex}} \rangle \cong \langle 2^\omega, \leq_{\text{lex}} \rangle$$

it follows by Lemma 6.21 (a) that every family $H \subseteq \wp(\{0, 1\}^*)$ has a \leq_{lex} -infimum. Hence, we only need to prove that, for the family $H := \vartheta^\mathfrak{T}$ of all sets Z satisfying ϑ , the infimum $U := \bigcap \vartheta^\mathfrak{T}$ also satisfies ϑ . Let $u_0 <_{\text{lex}}$

$u_1 <_{\text{lex}} \dots$ be an increasing enumeration of U and let $w_0 <_{\text{lex}} w_1 <_{\text{lex}} \dots$ be one of $\{0, 1\}^*$. We have to show that the function mapping w_i to u_i is an embedding. Clearly, it is monotone with respect to \leq_{lex} . Hence, we only need to prove that it respects the two successor relations. For a contradiction, suppose that $w_k = \text{succ}_d(w_i)$, for $d \in \{0, 1\}$, but $u_k \not\geq_{\text{pf}} \text{succ}_d(u_i)$. Note that $w_k = \text{succ}_d(w_i)$ implies $i < k$. By Lemma 6.21 (b), there exists a set $V \in \vartheta^\mathfrak{T}$ such that the first $k + 1$ elements of V are u_0, \dots, u_k . But this is not possible since $u_k \not\geq_{\text{pf}} \text{succ}_d(u_i)$ implies that no embedding $\mathfrak{T}_{\text{lex}} \rightarrow \mathfrak{T}$ can have range V . \square

Theorem 6.23. There exists an MSO-interpretation τ that maps every well-ordered infinite binary tree \mathfrak{T} to $\langle \{0, 1\}^*, \text{succ}_0, \text{succ}_1, \leq_{\text{lex}} \rangle$.

Proof. By Lemma 6.18 and Proposition 6.16, there exists an embedding $\mathfrak{T}_{\text{lex}} \rightarrow \mathfrak{T}$ with an MSO-definable range U . As we can define the relations of $\mathfrak{T}_{\text{lex}}$ from U , the claim follows. \square

Corollary 6.24. Let $\mathfrak{T}_{\text{bin}}$ be the infinite binary tree and \leq a well-ordering on T . Then $\langle \mathfrak{T}_{\text{bin}}, \leq \rangle$ has an undecidable MSO-theory.

Proof. By the preceding theorem it is sufficient to show that the theory of $\mathfrak{T}_{\text{lex}} := \langle \mathfrak{T}, \leq_{\text{lex}} \rangle$ is undecidable. We do so by defining arbitrarily large finite grids inside of $\mathfrak{T}_{\text{lex}}$. Given a size $n < \omega$, choose a length k such that $2^k \geq n$ and fix a set $P \subseteq \{0, 1\}^k$ of size n . Let $p \in P$ be the \leq_{lex} -minimal element of P and set

$$\begin{aligned} Q &:= \{ \text{succ}_0^i(p) \mid i < n \}, \\ V &:= \{ \text{succ}_0^i(q) \mid i < n, q \in P \}, \\ E_0 &:= \{ \langle u, v \rangle \in V^2 \mid v = \text{succ}_0(u) \}, \\ E_1 &:= \{ \langle u, v \rangle \in V^2 \mid u <_{\text{lex}} v, v \notin Q, \text{ and there is no } w \text{ with} \\ &\quad u <_{\text{lex}} w <_{\text{lex}} v \}. \end{aligned}$$

Then $\langle V, E_0, E_1 \rangle$ forms a directed grid of size $n \times n$ which can be defined using the monadic parameters P, Q , and V . Hence, there exists an MSO-interpretation τ mapping $\langle \mathfrak{T}_{\text{lex}}, P, Q, V \rangle$ to $\langle V, E_0, E_1 \rangle$.

Furthermore, there exists an MSO-formula ψ stating that the given graph forms a finite grid. It follows that an MSO-formula φ is satisfied by all finite grids if, and only if,

$$\mathfrak{I}_{\text{lex}} \models \exists P \exists Q \exists V [\psi^T \wedge \varphi^T].$$

As the theory of all finite grids is undecidable, so is therefore the theory of $\mathfrak{I}_{\text{lex}}$. \square

7 Uniformisation

A choice functions picks a unique element out of a given set. In this section we consider the more general problem of picking a set out of a definable family of sets. More precisely, we look for definable Skolem functions, that is, given a definable relation $R \subseteq A \times B$, we want to construct a definable function $f : A \rightarrow B$ with $f \subseteq R$.

Definition 7.1. An MSO-formula $\varphi_o(\bar{X}, \bar{Y})$ is a *uniformisation* of another MSO-formula $\varphi(\bar{X}, \bar{Y})$ over a structure \mathfrak{A} if, for all parameters $\bar{P}, \bar{Q}, \bar{Q}'$,

$$\begin{aligned} \mathfrak{A} \models \varphi_o(\bar{P}, \bar{Q}) & \quad \text{implies} \quad \mathfrak{A} \models \varphi(\bar{P}, \bar{Q}), \\ \mathfrak{A} \models \exists \bar{Y} \varphi(\bar{P}, \bar{Y}) & \quad \text{implies} \quad \mathfrak{A} \models \exists \bar{Y} \varphi_o(\bar{P}, \bar{Y}), \\ \mathfrak{A} \models \varphi_o(\bar{P}, \bar{Q}) \wedge \varphi_o(\bar{P}, \bar{Q}') & \quad \text{implies} \quad \bar{Q} = \bar{Q}'. \end{aligned}$$

We say that \mathfrak{A} has *strong MSO-uniformisation* if every MSO-formula φ has an uniformisation (without parameters). It has *(weak) MSO-uniformisation* if, for every MSO-formula φ , there exists an MSO-formula φ_o with monadic parameters that is an uniformisation of φ . Finally, we say that \mathfrak{A} has *effective strong/weak MSO-uniformisation* if the formula φ_o can be computed from φ . \downarrow

Clearly, uniformisation implies choice.

Lemma 7.2. *Weak MSO-uniformisation implies weak MSO-choice, and strong MSO-uniformisation implies strong MSO-choice.*

Proof. Let $\varphi_o(X, Y)$ be a uniformisation of the formula

$$\varphi(X, Y) := Y \subseteq X \wedge [X \neq \emptyset \rightarrow |Y| = 1].$$

Then the formula $\psi(x, X) := \varphi_o(X, \{x\})$ defines a choice function. \square

Exercise 7.1. Prove that a structure \mathfrak{A} has [weak] MSO-uniformisation if, and only if, every formula $\varphi(\bar{X}, Y)$ with a single variable Y has a uniformisation [with monadic parameters]. \downarrow

Definable Well-Orderings

We start by investigating which well-orderings can be defined within a given linear order. To do so, we consider the following kind of reduction between ordinals.

Definition 7.3. (a) Given two ordinals α, β , we define the relation

$$\alpha \leq \beta \quad \text{iff} \quad \text{there exists an MSO-formula } \varphi(x, y; \bar{P}) \text{ with monadic parameters } \bar{P} \text{ that defines a well-order of order type } \alpha \text{ in } \langle \beta, \leq \rangle.$$

(b) For a labelling $\lambda : \beta \rightarrow [k]$, we write $\alpha \leq_\lambda \beta$ if the order on β defined by

$$a \sqsubseteq b \quad \text{iff} \quad \lambda(a) < \lambda(b) \quad \text{or} \quad [\lambda(a) = \lambda(b) \text{ and } a \leq b]$$

has order type α . \downarrow

Note that in the above reductions we do *not* allow removing elements from β . We will prove that $\alpha \leq \beta$ if, and only if, $\alpha \leq_\lambda \beta$, for some λ . In the proof, we make use of the following basic properties of the relation \leq .

Lemma 7.4. *Let α and β be ordinals.*

- (a) \leq is reflexive and transitive.
- (b) $\alpha \leq \alpha'$ and $\beta \leq \beta'$ implies $\alpha + \beta \leq \alpha' + \beta'$ and $\alpha\beta \leq \alpha'\beta'$.

- (c) $\alpha + \beta \leq \beta + \alpha$.
- (d) $\alpha n + k \leq \alpha$, for all $n, k < \omega$ and $\alpha \geq \omega$.
- (e) $\beta \leq \alpha$ and $\alpha\omega \leq \beta$ implies $\alpha\omega \leq \alpha$.

Proof. (a) Reflexivity is trivial. For transitivity, suppose that φ is a formula defining β inside α , and ψ a formula defining γ inside β . Then we obtain a formula defining γ inside α by replacing in ψ every atomic subformula of the form $x \leq y$ by $\varphi(x, y)$.

(b) Let φ be a formula defining α in α' and ψ a formula defining β in β' . To prove the first claim, let λ be the labelling of $\alpha' + \beta'$ mapping the elements of α' to 0 and those of β' to 1. We can define $\alpha + \beta$ in $\alpha' + \beta'$ by applying φ to $\lambda^{-1}(0)$ and ψ to $\lambda^{-1}(1)$.

For the second claim, let $P \subseteq \alpha' \beta'$ be the set containing the first element of every copy of α' . Then P has order type β' . Let Ξ_1 be the well-order defined on P by ψ . Since each copy of α' is definable using P , we can use the formula φ to define an ordering Ξ_2 on each copy. Finally, let μ be the function mapping each element i to the maximal element $\mu(i) \in P$ with $\mu(i) \leq i$. Using these two orderings we obtain the desired well-order of $\alpha' \beta'$ by

$$i \Xi j \quad \text{iff} \quad \mu(i) \sqsubset_1 \mu(j) \quad \text{or} \quad [\mu(i) = \mu(j) \text{ and } i \Xi_2 j].$$

(c) Let λ be the labelling of $\beta + \alpha$ mapping the first β elements to 1 and the remaining ones to 0. Then $\alpha + \beta \leq_\lambda \beta + \alpha$.

(d) It is sufficient to prove that $\alpha n \leq \alpha$ and $\alpha + 1 \leq \alpha$. Then it follows by (a) that

$$\alpha n + k \leq \alpha n \leq \alpha.$$

For the second statement, note that (c) implies

$$\alpha + 1 \leq 1 + \alpha = \alpha.$$

For the first statement, we start with the case where α is a limit ordinal. Let $\lambda : \alpha \rightarrow [n]$ be the labelling defined by

$$\lambda(\omega\beta + m) := m \bmod n, \quad \text{for } m < \omega.$$

Then every preimage $\lambda^{-1}(i)$ has order type α . Hence, $\alpha n \leq_\lambda \alpha$.

For the general case, suppose that $\alpha = \delta + m$, for some limit ordinal δ and $m < \omega$. Then (b) implies that

$$\alpha n = (\delta + m)(n - 1) + \delta + m = \delta(n - 1) + \delta + m \leq \delta + m = \alpha.$$

(e) For a contradiction, suppose that $\beta \leq \alpha$, for some $\beta \geq \alpha\omega$. Let φ be a formula defining a well-ordering Ξ of order type β . By assumption, there is some set $P \subseteq \alpha$ such the restriction of Ξ to P has order type $\alpha\omega$. Let $\lambda : \alpha \rightarrow [2]$ be the labelling with $\lambda^{-1}(1) = P$. Since $\alpha \setminus P \subseteq \alpha$ has order type $\gamma \leq \alpha$, we have $\alpha\omega = \gamma + \alpha\omega \leq_\lambda \alpha$. \square

Next let us prove an upper bound, which is tight by Lemma 7.4 (d).

Lemma 7.5. $\beta \leq \alpha$ implies $\beta < \alpha\omega$.

Proof. For a contradiction, suppose that $\beta \leq \alpha$, for some $\beta \geq \alpha\omega$. Then Lemma 7.4 (e) implies that $\alpha\omega \leq \alpha$. In particular, there exists a smallest ordinal α_0 with $\alpha_0\omega \leq \alpha_0$. Let $\varphi(x, y; \bar{P})$ be the formula defining within α_0 an order Ξ of order type $\alpha_0\omega$, let $\sigma : \alpha_0 \rightarrow \alpha_0\omega$ be the corresponding isomorphism, and let m be the quantifier-rank of φ . We start by establishing the following claims.

(I) $\beta < \alpha_0$ implies $\beta\omega \leq \alpha_0$.

(II) $\beta < \alpha_0 \leq \gamma$ implies $\gamma \not\leq \beta$.

(I) Let $\beta < \alpha_0$. For a contradiction, suppose that $\beta\omega > \alpha_0$. Then there exists some $n < \omega$ such that $\beta n \leq \alpha_0 < \beta(n + 1)$. Let γ be the ordinal such that $\alpha_0 + \gamma = \beta(n + 1)$. Then $\gamma \leq \beta$ and $\beta\omega = \alpha_0\omega$. By Lemma 7.4, it follows that

$$\beta\omega = \gamma + \beta\omega \leq \beta\omega + \gamma = \alpha_0\omega + \gamma \leq \alpha_0 + \gamma = \beta(n + 1) \leq \beta.$$

A contradiction to the minimality of α_0 .

(II) Let $\beta < \alpha_0 \leq \gamma$. If $\gamma \leq \beta$, it follows by Lemma 7.4 (e) and (I) that

$$\omega\beta \leq \alpha_0 \leq \gamma \quad \text{implies} \quad \omega\beta \leq \beta.$$

A contradiction to the minimality of α_o .

Having proved the above claims, we conclude the proof as follows. Let $\mu : \alpha_o \rightarrow \omega$ be the function defined by

$$\mu(i) := k \quad : \text{iff} \quad \alpha_o k \leq \sigma(i) < \alpha_o(k+1).$$

The equivalence relation

$$i \approx j \quad : \text{iff} \quad \mu(i) = \mu(j)$$

is convex on $\langle \alpha_o, \Xi \rangle$. By Lemma 3.1, \approx is MSO-definable with the help of one monadic parameter Q . Since \approx has infinitely many classes, we can use Lemma 2.3 to find a bounded (with respect to \leq) \approx -class $C = \mu^{-1}(k)$. Suppose that $C \subseteq \beta$ with $\beta < \alpha$. Let γ be the order type of $\beta \setminus C$. Since the order type of $\langle C, \Xi \rangle$ is α_o , it follows what we can use the restriction of the formula φ to β to define in β a well-order of type $\alpha_o + \gamma$. Hence, $\alpha_o + \gamma \leq \beta$, which contradicts (11). \square

Finally, let us provide a corresponding lower bound. We start with a technical lemma.

Lemma 7.6. *Let α be a limit ordinal such that $\gamma < \alpha$ implies $\gamma \not\leq \alpha$, and let $k < \omega$. Assume that $k = 1$ or that $\alpha = \omega^n$, for some n . Then*

$$\alpha k \not\leq \beta, \quad \text{for all } \beta \geq \alpha(k+1).$$

Proof. For a contradiction suppose that $\alpha k \leq \beta$ and $\beta \geq \alpha(k+1)$. Let $\varphi(x, y; \tilde{P})$ be an MSO_m -formula defining in β a well-order Ξ of order type αk and let $h : \langle \beta, \Xi \rangle \rightarrow \langle \alpha k, \leq \rangle$ be the corresponding isomorphism. Set $\mathfrak{B} := \langle \beta, \leq, \tilde{P} \rangle$. For $s \leq k$ and $\alpha s \leq i, j < \alpha(s+1)$, we have

$$\begin{aligned} \text{Th}_{\text{MSO}}^m(\mathfrak{B}, i, j) &= \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[0, \alpha s)}) + \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha s, \alpha(s+1))}, i, j) \\ &\quad + \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha(s+1), \beta)}). \end{aligned}$$

Consequently, the restriction of Ξ to $[\alpha s, \alpha(s+1))$ is MSO_m -definable in $[\alpha s, \alpha(s+1)) \cong \alpha$. Let $C_s := h[[\alpha s, \alpha(s+1))]$ be its image in αk and let γ be

its order type. Then $\gamma \leq \alpha$, which implies by assumption that $\gamma = \alpha$. Since C_s and $[\alpha s, \alpha(s+1))$ both have the same cofinality $\delta := \text{cf } \alpha$, there exists a sequence $(c_i^s)_{i < \delta}$ in $[\alpha s, \alpha(s+1))$ that is cofinal both in $\langle [\alpha s, \alpha(s+1)), \leq \rangle$ and in $\langle C_s, \Xi \rangle$.

By the Pigeon Hole Principle, there exists a cofinal subset $I_s \subseteq \delta$ and a theory σ_s such that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha s, \alpha(s+1))}, c_i^s) = \sigma_s, \quad \text{for all } i \in I_s.$$

Replacing $(c_i^s)_{i < \delta}$ by the subsequence $(c_i^s)_{i \in I_s}$, we may assume that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha s, \alpha(s+1))}, c_i^s) = \sigma_s, \quad \text{for all } i < \delta.$$

Consequently,

$$\begin{aligned} \text{Th}_{\text{MSO}}^m(\mathfrak{B}, c_i^s, c_j^t) &:= \\ &\quad \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[0, \alpha s)}) + \text{Th}^m(\mathfrak{B}|_{[\alpha s, \alpha(s+1))}, c_i^s) \\ &\quad + \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha(s+1), \alpha t)}) + \text{Th}^m(\mathfrak{B}|_{[\alpha t, \alpha(t+1))}, c_j^t) \\ &\quad + \text{Th}_{\text{MSO}}^m(\mathfrak{B}|_{[\alpha(t+1), \beta)}), \end{aligned}$$

which, by choice of I_s and I_t , depends only on s and t , but not on i and j .

Since Ξ is a linear order, there exists an enumeration s_0, \dots, s_k of $[k+1]$ such that $c_1^{s_0} \sqsubset \dots \sqsubset c_1^{s_k}$. As $\text{Th}^m(\mathfrak{B}, c_i^s, c_j^t)$ does not depend on i and j , it follows that

$$c_i^{s_l} \sqsubset c_j^{s_{l+1}}, \quad \text{for all } 0 < i, j < \delta \text{ and all } l < k.$$

The sequence $(c_i^s)_i$ being cofinal in C_s , this implies that

$$c \sqsubset c_j^{s_l}, \quad \text{for all } 0 < j < \delta \text{ and } c \in C_{s_l} \text{ with } l < k.$$

If $k = 1$, it follows that $C_{s_0} \cup \{c_1^{s_1}\}$ has order type $\alpha + 1$. Hence, the order type of Ξ is strictly greater than α . A contradiction.

It remains to consider the case where $k > 1$. By assumption, this implies that $\alpha = \omega^n$, for some n . Each set C_{s_l} with $l > 0$ can be partitioned into two

parts $C_{s_l} = J_l + K_l$ where

$$J_l := \{ a \in C_{s_l} \mid \uparrow a \cap (C_{s_0} \cup \dots \cup C_{s_{l-1}}) \neq \emptyset \},$$

$$K_l := \{ a \in C_{s_l} \mid C_{s_0} \cup \dots \cup C_{s_{l-1}} \subseteq \downarrow a \}.$$

Let γ_l and γ'_l be the order types of J_l and K_l . Since C_{s_l} has order type $\alpha = \omega^n$, it follows that

$$\omega^n = \gamma_l + \gamma'_l.$$

The only sums with value ω^n are $\omega^n + 0$ and $\gamma + \omega^n$ with $\gamma < \omega^n$. The first case is not possible since $c_1^{s_l} \in K_l$ implies $\gamma'_l > 0$. Consequently, K_l has order type $\omega^n = \alpha$ and it follows that $C_{s_0} + K_1 + \dots + K_{k-1} + \{c_1^{s_k}\}$ has order type $\alpha + \alpha(k-1) + 1 = \alpha k + 1$. Again a contradiction to the fact that the order type of Ξ is αk . \square

Lemma 7.7. For ordinals $\alpha < \beta$, we have

$$\alpha \leq \beta \quad \text{iff} \quad \alpha = \gamma + \delta \text{ and } \beta = \delta + \gamma, \quad \text{for some } \gamma, \delta \leq \beta.$$

Proof. (\Leftarrow) follows by Lemma 7.4. For (\Rightarrow), let us call a pair $\alpha < \beta$ of ordinals *bad* if

$$\alpha \leq \beta, \quad \text{but there are no } \gamma, \delta \leq \beta \text{ with } \alpha = \gamma + \delta \text{ and } \beta = \delta + \gamma.$$

We choose a bad pair $\langle \alpha, \beta \rangle$ such that α minimal and β minimal for this α . We start by establishing the following claims.

- (I) α is a limit ordinal.
- (II) $\gamma \not\leq \alpha$, for all $\gamma < \alpha$.
- (III) If $\gamma < \alpha$ implies $\gamma\omega < \alpha$, then $\eta \geq \alpha$ where η is the ordinal such that $\beta = \alpha + \eta$.
- (1) For a contradiction suppose that $\alpha = \alpha_0 + 1$ is a successor. Then $\alpha_0 \leq \alpha_0 + 1 \leq \beta$ implies, by minimality of α , that

$$\alpha_0 = \gamma + \delta \text{ and } \beta = \delta + \gamma, \quad \text{for some } \gamma, \delta \leq \beta.$$

If $0 < \gamma < \omega$, we obtain

$$\alpha = \gamma + \delta + 1 = (\gamma - 1) + \delta + 1 \quad \text{and} \quad \beta = \delta + \gamma = (\delta + 1) + (\gamma - 1).$$

A contradiction. Hence, $\gamma \geq \omega$ and

$$\alpha = \alpha_0 + 1 = \gamma + \delta + 1 \quad \text{and} \quad \beta = \delta + \gamma = \delta + 1 + \gamma.$$

Again, a contradiction.

(II) For a contradiction, suppose that there is some $\gamma < \alpha$ with $\gamma \leq \alpha$. Then $\gamma \leq \alpha \leq \beta$ implies, by minimality of α , that

$$\gamma = \delta + \varepsilon \text{ and } \beta = \varepsilon + \delta, \quad \text{for some } \delta, \varepsilon \leq \beta.$$

Suppose that $\alpha = \gamma + \xi$, $\beta = \alpha + \eta$, and that

$$\delta = \omega^k d + \delta_0 \quad \text{and} \quad \varepsilon = \omega^l e + \varepsilon_0 \quad \text{with } \delta_0 < \omega^k \text{ and } \varepsilon_0 < \omega^l.$$

We distinguish three cases.

If $k > l$, we have $\beta = \varepsilon + \delta = \delta \leq \gamma$. A contradiction.

If $k < l$, we have $\gamma = \delta + \varepsilon = \varepsilon$, which implies that $\beta = \gamma + \delta$ and $\delta = \xi + \eta$. Hence, $\xi + \eta + \gamma = \delta + \varepsilon = \gamma$ and it follows that

$$\beta = (\gamma + \xi) + \eta \quad \text{and} \quad \alpha = \gamma + \xi = \eta + (\gamma + \xi).$$

A contradiction to our choice of α and β .

Finally, suppose that $k = l$. Then

$$\gamma = \omega^k(d + e) + \varepsilon_0 \quad \text{and} \quad \beta = \omega^k(d + e) + \delta_0.$$

Hence,

$$\alpha = \gamma + \xi = \omega^k(d + e) + \varepsilon_0 + \xi < \beta,$$

which implies that $\varepsilon_0 + \xi < \delta_0$. It follows that $\xi + \delta = \delta$ and

$$\alpha = \gamma + \xi = \delta + \varepsilon + \xi \quad \text{and} \quad \beta = \varepsilon + \delta = \varepsilon + \xi + \delta.$$

Again a contradiction to our choice of α and β .

(III) For a contradiction, suppose that $\eta < \alpha$. We write

$$\eta = \omega^k m + \eta_0 \quad \text{and} \quad \alpha = \omega^l n + \alpha_0,$$

where $\eta_0 < \omega^k$, $\alpha_0 < \omega^l$, and $0 < m, n < \omega$. Then

$$\eta\omega = \omega^{k+1} + \eta_0 \leq \alpha \quad \text{implies} \quad l \geq k+1.$$

Hence, we have

$$\alpha = \eta + \alpha \quad \text{and} \quad \beta = \alpha + \eta,$$

which contradicts our choice of α and β .

Having established the above claims, we conclude the proof as follows.

Suppose that

$$\alpha = \omega^n k + \gamma, \quad \text{for } \gamma < \omega^n \text{ and } 0 < k < \omega.$$

If $\gamma > 0$, we would have

$$\omega^n k = \omega^n(k-1) + \gamma + \omega^n \leq \omega^n(k-1) + \omega^n + \gamma = \alpha$$

and $\omega^n k < \alpha$. A contradiction to (II).

Consequently, $\gamma = 0$. If $k = 1$, that is, if $\alpha = \omega^n$, we have

$$\gamma\omega \leq \alpha, \quad \text{for all } \gamma < \alpha.$$

Consequently, it follows by (III) that $\eta \geq \alpha$. But $\beta = \alpha + \eta \geq \alpha + \alpha$ together with (II) contradicts Lemma 7.6 for $k = 1$.

It follows that $\alpha = \omega^n k$ with $k > 1$. Suppose that $\beta = \alpha + \eta$. If $\eta < \omega^n$, we obtain

$$\alpha = \eta + \alpha \quad \text{and} \quad \beta = \alpha + \eta,$$

which contradicts our choice of α and β . Consequently, $\eta \geq \omega^n$ and $\beta \geq \omega^n(k+1)$. By Lemma 7.6 it follows that there is some $\gamma < \omega^n$ with $\gamma \leq \omega^n$. Furthermore, $\gamma \leq \omega^n \leq \beta$ implies, by minimality of α , that

$$\gamma = \delta + \varepsilon \text{ and } \beta = \varepsilon + \delta, \quad \text{for some } \delta, \varepsilon \leq \beta.$$

But $\delta, \varepsilon \leq \gamma < \omega^n$ implies that $\beta = \varepsilon + \delta < \omega^n 2 \leq \beta$. A contradiction. \square

Combining our two bounds we obtain the following characterisation of the relation \leq .

Proposition 7.8. $\alpha \leq \beta$ if, and only if, $\alpha \leq_\lambda \beta$, for some $\lambda : \beta \rightarrow [k]$ with $k < \omega$.

Proof. (\Leftarrow) is trivial. For (\Rightarrow), let $\alpha \leq \beta$. If $\alpha = \beta$, the claim is trivial and, if $\alpha < \beta$, it follows by Lemma 7.7. Hence, suppose that $\alpha > \beta$. Then α and β are infinite. Furthermore, we have $\alpha < \beta\omega$ by Lemma 7.5. Hence, the claim follows by Lemma 7.4 (d). \square

Remark. (a) For this description it follows that $\alpha \leq \beta$ implies that we can define α in β using a *quantifier-free* formula (with monadic parameters).

(b) An alternative proof of the above proposition is based on Proposition IX.5.1. By that proposition, if we can define a linear order Ξ inside another linear order (A, \leq) , then every \leq -interval can be partitioned into a bounded number of Ξ -intervals. \lrcorner

For ordinals that are not too large, we obtain the following explicit description.

Theorem 7.9. Let α and β be ordinals.

(a) If $\alpha < \omega$, then

$$\beta \leq \alpha \quad \text{iff} \quad \beta = \alpha.$$

(b) If $\omega^k n \leq \alpha < \omega^k(n+1)$, for $0 < n, k < \omega$, then

$$\beta \leq \alpha \quad \text{iff} \quad \omega^k n \leq \beta \leq \omega^{k+1}.$$

(c) If $\alpha \geq \omega^\omega$, then

$$\beta \leq \alpha \quad \text{implies} \quad \beta \geq \omega^\omega.$$

Proof. (a) Note that $\beta \leq \alpha$ implies $|\beta| = |\alpha|$.

(b) (\Rightarrow) By Lemma 7.5,

$$\beta \leq \alpha \quad \text{implies} \quad \beta < \alpha\omega \leq \omega^k(n+1)\omega = \omega^{k+1}.$$

Furthermore, if $\beta < \alpha$ and $\beta \leq \alpha$, then Lemma 7.7 implies that $\beta = \gamma + \delta$ and $\alpha = \delta + \gamma$, for some $\gamma, \delta \leq \alpha$. It follows that $\gamma \geq \omega^k i$ and $\delta \geq \omega^k j$ where $i + j = n$. Consequently,

$$\beta = \gamma + \delta \geq \omega^k i + \omega^k j = \omega^k n.$$

(\Leftarrow) Suppose that $\alpha = \omega^k n + \gamma$ and $\beta = \omega^k(n + i) + \delta$ with $\gamma, \delta < \omega^k$ and $i < \omega$. Then

$$\begin{aligned} \beta &= \omega^k(n + i) + \delta \\ &\leq \delta + \omega^k(n + i) \\ &= \omega^k(n + i) \\ &= \omega^k(n - 1) + \omega^k(i + 1) \\ &\leq \omega^k(n - 1) + \omega^k \\ &= \omega^k n \\ &= \gamma + \omega^k n \\ &\leq \omega^k n + \gamma = \alpha. \end{aligned}$$

(c) For a contradiction suppose that $\beta \leq \alpha$ with $\beta < \omega^\omega$. By Lemma 7.7, there are $\gamma, \delta \leq \alpha$ with $\beta = \gamma + \delta$ and $\alpha = \delta + \gamma$. Since $\delta + \gamma = \alpha \geq \omega^\omega$, we have $\gamma \geq \omega^\omega$ or $\delta \geq \omega^\omega$. In both cases it follows that $\beta = \gamma + \delta \geq \omega^\omega$. \square

Proposition 7.10. *Let \mathcal{A} be a linear order. Exactly one of the following two conditions holds.*

- (I) *Every MSO-definable well-ordering on A has order type at least ω^ω .*
- (II) *There is some $k < \omega$, such that every MSO-definable well-ordering on A has an order type in the interval $[\omega^k, \omega^{k+1})$.*

Proof. Let Ξ and Ξ' be two MSO-definable well-orderings of order types α and α' respectively. Then we can use Lemma 6.11 (b) to construct an MSO-formula φ defining an order \leq on α that is isomorphic to α' . Consequently, $\alpha' \leq \alpha$ and the claim follows by Theorem 7.9. \square

According to the preceding proposition, we can associate the following invariant with every linear ordering.

Definition 7.11. The *well-ordering index* $\text{WO}(\mathcal{A})$ of a structure \mathcal{A} is the least number $k < \omega$ such that every MSO-definable (with monadic parameters) well-ordering on A has an order type less than ω^{k+1} (and at least one such well-ordering exists). If such a number k does not exist, we set $\text{WO}(\mathcal{A}) := \infty$. \square

Uniformisation For Well-Orderings

We can use the above results to give a description of which well-orders (and, more generally, which linear orders) have weak MSO-uniformisation. We start by showing that ω^ω does not have uniformisation. The proof rests on the following technical lemma.

Lemma 7.12. *Let λ be an additive labelling of ω^n with at most $n - 2$ colours such that $\lambda(\alpha, \beta)$ depends only on the order type of $[\alpha, \beta)$. For every cofinal increasing sequence $(\alpha_i)_{i < \omega}$, there exists a cofinal increasing sequence $(\alpha'_i)_{i < \omega}$ different from $(\alpha_i)_{i < \omega}$ such that*

$$\alpha'_0 = \alpha_0 \quad \text{and} \quad \lambda(\alpha'_i, \alpha'_{i+1}) = \lambda(\alpha_i, \alpha_{i+1}), \quad \text{for all } i < \omega.$$

Proof. By assumption on λ , there exists a function μ such that

$$\lambda(\alpha, \alpha + \gamma) = \mu(\gamma), \quad \text{for all } \alpha, \gamma < \omega^n.$$

Furthermore, we can use Proposition 2.6 to find an idempotent colour θ such that

$$\mu(\omega^k) = \theta, \quad \text{for all } k \geq n - 2.$$

Let β_i be the ordinal such that $\alpha_{i+1} = \alpha_i + \beta_i$. Since $\sum_{i < \omega} \beta_i = \omega^n$, there are infinitely many indices i such that $\beta_i \geq \omega^{n-1}$. In particular, there are indices $i < j < \omega$ such that $\beta_i, \beta_j \geq \omega^{n-1}$ and $\beta_k < \omega^{n-1}$, for $i < k < j$. We write

$$\beta_i = \omega^{n-1}\gamma + \delta \quad \text{and} \quad \beta_j = \omega^{n-1}\gamma' + \delta' \quad \text{with} \quad \delta, \delta' < \omega^{n-1}.$$

Set $l := j - i$ and

$$\begin{aligned}\xi_0 &:= \alpha_i, \\ \xi_1 &:= \alpha_i + \omega^{n-1}\gamma + \omega^{n-2} + \delta, \\ \xi_{k+1} &:= \xi_k + \beta_{i+k}, & \text{for } 0 < k < l, \\ \xi_{l+1} &:= \alpha_{i+l+1}.\end{aligned}$$

Then $\omega^{n-2} + \delta + \beta_{i+1} + \dots + \beta_{i+l-1} < \omega^{n-1}$ implies that

$$\begin{aligned}\xi_l + \beta_j &= \alpha_i + (\omega^{n-1}\gamma + \omega^{n-2} + \delta) + \beta_{i+1} + \dots + \beta_{i+l-1} + \beta_j \\ &= \alpha_i + \omega^{n-1}\gamma + \omega^{n-2} + \delta + \beta_{i+1} + \dots + \beta_{i+l-1} + \omega^{n-1}\gamma' + \delta' \\ &= \alpha_i + \omega^{n-1}\gamma + \omega^{n-1}\gamma' + \delta' \\ &= \alpha_i + \omega^{n-1}\gamma + \delta + \beta_{i+1} + \dots + \beta_{i+l-1} + \omega^{n-1}\gamma' + \delta' \\ &= \alpha_i + \beta_i + \dots + \beta_{i+l} \\ &= \xi_{l+1}.\end{aligned}$$

By assumption on λ , we have

$$\lambda(\xi_k, \xi_{k+1}) = \mu(\beta_{i+k+1}) = \lambda(\alpha_{i+k}, \alpha_{i+k+1}), \quad \text{for } 0 < k \leq l.$$

Let m be the largest number such that $\gamma = \omega^m \gamma_0$, for some γ_0 . Then $\gamma_0 = \gamma_1 + 1$ is a successor ordinal. Hence, $\mu(\omega^{n-1}) = \theta = \mu(\omega^{n-2})$ and the fact that θ is idempotent implies that

$$\begin{aligned}\mu(\omega^{n-1}\gamma) + \mu(\omega^{n-2}) &= \mu(\omega^{n-1+k}\gamma_0) + \mu(\omega^{n-2}) \\ &= \theta \cdot \gamma_0 + \theta \\ &= \theta \cdot \gamma_1 + \theta + \theta \\ &= \theta \cdot \gamma_1 + \theta \\ &= \mu(\omega^{n-1+k}\gamma_0) = \mu(\omega^{n-1}\gamma).\end{aligned}$$

Hence,

$$\begin{aligned}\lambda(\xi_0, \xi_l) &= \mu(\omega^{n-1}\gamma) + \mu(\omega^{n-2}) + \mu(\delta) \\ &= \mu(\omega^{n-1}\gamma) + \mu(\delta) \\ &= \lambda(\alpha_i, \alpha_{i+l}).\end{aligned}$$

Consequently, the sequence $(\alpha'_k)_{k < \omega}$ obtained from $(\alpha_k)_{k < \omega}$ by replacing the subsequence $\alpha_{i+1}, \dots, \alpha_j$ by ξ_1, \dots, ξ_l has the desired properties. \square

Our counterexample to uniformisation is the formula stating the existence of a cofinal sequence of length ω .

Lemma 7.13. *Let $\varphi(X, Y)$ be the formula stating that*

“The set X has no maximal element, $Y \subseteq X$ is unbounded in X , and its order type is ω .”

and, for $n < \omega$, let $\psi_n(X, Y; \bar{P}_n)$ be a uniformisation of φ over the structure $\langle \omega^n, \leq \rangle$. Suppose that ψ_n has quantifier-rank m_n and $l_n := |\bar{P}_n|$ parameters. Then at least one of the sequences $(m_n)_n$ and $(l_n)_n$ is unbounded.

Proof. For a contradiction, suppose that the two sequences are bounded by some number m . Since there are only finitely many different formulae of quantifier-rank m with m parameters, we can then find a formula $\psi(X, Y; \bar{Z})$, an unbounded set $I \subseteq \omega$, and parameters \bar{P}_n , for $n \in I$, such that, for every $n \in I$, $\psi(X, Y; \bar{P}_n)$ is a uniformisation of φ over $\langle \omega^n, \leq \rangle$. Let M be the number of MSO_{m+1} -theories with $m+1$ monadic parameters, and choose $n \in I$ with $n \geq M+2$. Set $\mathfrak{D} := \langle \omega^n, \leq, \bar{P}_n \rangle$. We obtain the desired contradiction by proving that ψ is not a uniformisation of φ over \mathfrak{D} .

Let λ be the additive labelling on ω^n defined by

$$\lambda(\alpha, \beta) := \text{Th}_{\text{MSO}}^{m+1}(\mathfrak{D}|_{[\alpha, \beta)}, \alpha).$$

We induce a tree structure on ω^n using the bijection $\mu : \omega^{<n} \rightarrow \omega^n$ given by

$$\mu(\langle k_0, \dots, k_{i-1} \rangle) := \omega^{n-1}k_0 + \dots + \omega^{n-i+1}k_{i-2} + \omega^{n-i}(k_{i-1} + 1).$$

Note that the successors of every vertex form an order of order type ω and that the lexicographic ordering on $\omega^{<n}$ is mapped via μ to the linear ordering of ω^n . Applying Proposition 2.4 to the tree $T := \omega^{<n}$ and the labelling

$$\lambda'(u, v) := \lambda(\mu(u), \mu(v))$$

we obtain a prefix $P \subseteq T$ and theories θ_{ijk} such that

$$\lambda(\mu(u), \mu(v)) = \theta_{|u|, |v|, |u \sqcap v|}, \quad \text{for all } u < v \text{ in } P.$$

Note that, for leaves $u, v \in P$, the order type of the interval $[\mu(u), \mu(v))$ determines the value of $|u \sqcap v|$ since

$$\beta = \alpha + \omega^{s_o} k_o + \dots + \omega^{s_l} k_l \quad \text{implies} \quad |u \sqcap v| = n - s_o - 1.$$

Let $H \subseteq \omega^n$ be the image of the leaves of P under μ . Then H has order type ω^n and, for $\alpha < \beta$ in H , the colour

$$\lambda(\alpha, \beta) \text{ only depends on the order type of } [\alpha, \beta) \cap H.$$

Let θ_o be the theory such that

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha, \beta)}) = \theta_o, \quad \text{for } \alpha \in H \text{ with successor } \beta \text{ in } H.$$

(Note that θ_o can be computed from $\lambda(\alpha, \beta)$.) Since we have

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha, \beta)}) = \sum_{i < \gamma} \theta_o,$$

for $\alpha < \beta$ in H such that $[\alpha, \beta) \cap H$ has order type γ , it follows that the theory

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha, \beta)})$$

also only depends on the order type of $[\alpha, \beta) \cap H$.

Since $\psi(X, Y; \bar{P}_n)$ is an uniformisation of $\varphi(X, Y)$, there exists some set A such that

$$\langle \omega^n, \leq \rangle \models \psi(H, A; \bar{P}_n).$$

To obtain the desired contradiction, we will show that there exists a second such set. Let $(\alpha_i)_{i < \omega}$ be an enumeration of A . By Lemma 7.12, we obtain a second cofinal sequence $(\alpha'_i)_{i < \omega}$ such that $\alpha'_o = \alpha_o$ and

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha'_i, \alpha'_{i+1})}) = \text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})}),$$

for all $i < \omega$. We claim that

$$\langle \omega^n, \leq \rangle \models \psi(H, A'; \bar{P}_n), \quad \text{where } A' := \{ \alpha'_i \mid i < \omega \}.$$

A contradiction to the uniqueness of A .

Note that the theory

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})})$$

uniquely determines

$$\text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})}, \{ \alpha_i \}),$$

since a formula ϑ belongs to the latter if, and only if, the formula

$$\exists z [\forall u (z \leq u) \wedge \vartheta[x \in \{ \alpha_i \} \mapsto x = z]]$$

belongs to the former, where $\vartheta[x \in \{ \alpha_i \} \mapsto x = z]$ denotes the formula obtained from ϑ by replacing every atomic subformula of the form $x \in \{ \alpha_i \}$ by the formula $x = z$. (We may assume without loss of generality that $m \geq 1$.) Consequently,

$$\text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha'_i, \alpha'_{i+1})}) = \text{Th}_{\text{MSO}}^{m+1}(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})})$$

implies that

$$\text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[\alpha'_i, \alpha'_{i+1})}, \{ \alpha'_i \}) = \text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})}, \{ \alpha_i \}).$$

It follows that

$$\begin{aligned} & \text{Th}_{\text{MSO}}^m(\mathfrak{D}, H, A') \\ &= \text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[o, \alpha'_o)}, \emptyset) + \sum_{i < \omega} \text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[\alpha'_i, \alpha'_{i+1})}, \{ \alpha'_i \}) \\ &= \text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[o, \alpha_o)}, \emptyset) + \sum_{i < \omega} \text{Th}_{\text{MSO}}^m(\langle \mathfrak{D}, H \rangle|_{[\alpha_i, \alpha_{i+1})}, \{ \alpha_i \}) \\ &= \text{Th}_{\text{MSO}}^m(\mathfrak{D}, H, A), \end{aligned}$$

which proves our claim. \square

Proposition 7.14. $\langle \alpha, \leq \rangle$ does not have MSO-uniformisation, for $\alpha \geq \omega^\omega$.

Proof. We claim that the formula $\varphi(X, Y)$ from the preceding lemma expressing that

“The set X has no maximal element, $Y \subseteq X$ is unbounded in X , and its order type is ω .”

has no MSO-uniformisation. For a contradiction, suppose that there is an MSO-formula $\psi(X, Y; \bar{P})$ uniformising φ and let m be its quantifier-rank. For $n < \omega$ and $A \subseteq B \subseteq \omega^n$, we have

$$\begin{aligned} & \text{Th}^m(\langle \alpha, \leq, A, B, \bar{P} \rangle) \\ &= \text{Th}^m(\langle \alpha, \leq, A, B, \bar{P} \rangle|_{[\omega, \omega^n]}) + \text{Th}^m(\langle \alpha, \leq, A, B, \bar{P} \rangle|_{[\omega^n, \alpha]}) \\ &= \text{Th}^m(\langle \alpha, \leq, A, B, \bar{P} \rangle|_{[\omega, \omega^n]}) + \text{Th}^m(\langle \alpha, \leq, \emptyset, \emptyset, \bar{P} \rangle|_{[\omega^n, \alpha]}) . \end{aligned}$$

As the second theory does not depend on A and B it follows that, for sets $B \subseteq A \subseteq \omega^n$, the truth value of $\psi(A, B; \bar{P})$ only depends on

$$\text{Th}^m(\langle \alpha, \leq, A, B, \bar{P} \rangle|_{[\omega, \omega^n]}) .$$

Therefore we can find, for every $n < \omega$, an MSO_m -formula $\psi_n(X, Y; \bar{P}_n)$ with $|\bar{P}_n| = |\bar{P}|$ parameters that uniformises φ over $\langle \omega^n, \leq \rangle$. A contradiction to Lemma 7.13. \square

It remains to show that all smaller ordinals do have uniformisation. Since each such ordinal can be constructed from 1 and ω using addition and multiplication, it is sufficient to prove the following lemmas.

Lemma 7.15. If α and β are ordinals with effective weak MSO-uniformisation, then so are $\alpha + \beta$ and $\alpha\beta$.

Proof. Let $\varphi(\bar{X}, \bar{Y}) \in \text{MSO}_m$. We start with the sum $\alpha + \beta$. Let A be the prefix of order type α and B the suffix of order type β . By the composition theorem for ordered sums, we know that there exists a set Θ of pairs of

MSO_m -theories such that

$$\begin{aligned} \langle \alpha + \beta, \leq \rangle \models \varphi(\bar{P}, \bar{Q}) \quad & \text{iff} \quad \text{there is some } \langle \sigma, \tau \rangle \in \Theta \text{ such that} \\ & \text{Th}_{\text{MSO}}^m(\langle \alpha + \beta, \leq, \bar{P}, \bar{Q} \rangle|_A) = \sigma \\ & \text{and } \text{Th}_{\text{MSO}}^m(\langle \alpha + \beta, \leq, \bar{P}, \bar{Q} \rangle|_B) = \tau . \end{aligned}$$

By assumption we can choose, for each pair $\langle \sigma, \tau \rangle$, a uniformisation ψ_σ of $\wedge \sigma$ over α and one ϑ_τ of $\wedge \tau$ over β . Then φ has a uniformisation stating that

“There is some $\langle \sigma, \tau \rangle \in \Theta$ such that the restriction to A satisfies ψ_σ and the restriction to B satisfies ϑ_τ .”

This statement can be expressed in MSO using the parameters A, B and those used by the formulae ψ_σ and ϑ_τ .

For $\alpha\beta$ we proceed analogously. Note that we can write $\alpha\beta$ as an ordered sum of β many copies of α . By Lemma 3.1 we can define the relation \sim of being in the same copy using some MSO-formula with a single monadic predicate. Furthermore, by the composition theorem for generalised sums, there exists some MSO-formula χ such that

$$\langle \alpha\beta, \leq \rangle \models \varphi(\bar{P}, \bar{Q})$$

if, and only if, there exists a labelling λ of β by MSO_m -theories satisfying

$$\langle \beta, \leq, \lambda \rangle \models \chi .$$

Again we fix, for every theory θ , a uniformisation ψ_θ of $\wedge \theta$ over α . Since the relation \sim is definable, we can express in MSO the statement

“There exists a labelling λ of the \sim -classes satisfying χ such that every restriction of $\alpha\beta$ to a single \sim -class C satisfies $\psi_{\lambda(C)}$.” \square

We can generalise the proof of the preceding lemma as follows.

Exercise 7.2. Let \mathfrak{S} be a structure with weak MSO-uniformisation, let $m, n < \omega$, and let $(\mathfrak{A}_i)_{i \in I}$ be a family of structures such that, for every

$i \in I$ and for every MSO_n -formula $\varphi(\bar{X}, \bar{Y}; \bar{P})$ with at most n parameters, there exists an MSO_m -formula $\psi(\bar{X}, \bar{Y}; \bar{Q})$ with at most m parameters that is an uniformisation of φ over \mathfrak{A}_i . Prove that the generalised sum $\sum_{i \in \mathfrak{S}} \mathfrak{A}_i$ has weak MSO-uniformisation. \square

Proposition 7.16. $\langle \omega, \leq \rangle$ has effective weak MSO-uniformisation.

Proof. Let $\varphi(\bar{X}, \bar{Y})$ be an MSO-formula (possibly with parameters). Replacing φ by the formula

$$\varphi'(\bar{X}, \bar{Y}) := \varphi(\bar{X}, \bar{Y}) \vee \left[\neg \exists \bar{Y}' \varphi(\bar{X}, \bar{Y}') \wedge \bigwedge_i Y_i = \emptyset \right]$$

we may assume without loss of generality that, for every \bar{X} , there is some \bar{Y} satisfying $\varphi(\bar{X}, \bar{Y})$. Note that we can encode every tuple of sets \bar{X} by an ω -word over the alphabet $\Sigma := \wp(\bar{X})$. Similarly, we can encode a tuple $\bar{X}\bar{Y}$ by an ω -word over the alphabet $\Sigma \times \Gamma$, where $\Gamma := \wp(\bar{Y})$. According to Theorem III.4.7, we can use this encoding to translate φ into a deterministic ω -automaton $\mathcal{A} = \langle Q, \Sigma \times \Gamma, \Delta, q_0, \Omega \rangle$.

To construct the desired uniformisation, we define a parity game \mathcal{G} simulating \mathcal{A} on an unspecified input word. The positions are

$$V_\diamond := \omega \times Q \times \Sigma \quad \text{and} \quad V_\square := \omega \times Q.$$

The initial position is $\langle 0, q_0 \rangle \in V_\square$. In the position $\langle k, p \rangle \in V_\square$, Player \square chooses some letter $a \in \Sigma$ and the game proceeds to position $\langle k, p, a \rangle$. Then Player \diamond chooses a transition $\langle p, c, q \rangle \in \Delta$ such that $c = \langle a, b \rangle$, for some $b \in \Gamma$, and the game continues in position $\langle k + 1, q \rangle$. Player \diamond wins a play

$$\langle 0, q_0 \rangle, \langle 0, q_0, a_0 \rangle, \langle 1, q_1 \rangle, \langle 1, q_1, a_1 \rangle, \dots$$

of this game if the corresponding sequence q_0, q_1, \dots of states satisfies the parity condition.

By construction it follows that Player \diamond wins this game from every position and that every play corresponds to a pair of tuples \bar{X}, \bar{Y} satisfying φ .

As parity games are positionally determined, we can fix a memory-less winning strategy σ for Player \diamond . Furthermore, being memory-less, we can encode σ by a tuple \bar{S} of monadic predicates. Let $\psi(\bar{X}, \bar{Y}; \bar{S})$ be an MSO-formula stating that \bar{Y} are the sets determined by the strategy σ in the game \mathcal{G} when Player \square uses (an ω -word encoding) the sets \bar{X} to make his choices. This formula is a uniformisation of φ . \square

In summary, we have obtained the following characterisation.

Theorem 7.17. Let \mathfrak{A} be a well-order of order type α . The following statements are equivalent.

- (1) \mathfrak{A} has weak MSO-uniformisation.
- (2) \mathfrak{A} has effective weak MSO-uniformisation.
- (3) $\alpha < \omega^\omega$.
- (4) $\text{WO}(\mathfrak{A}) < \infty$.

Proof. (3) \Leftrightarrow (4) follows by Theorem 7.9; (2) \Rightarrow (1) is trivial; and (1) \Rightarrow (3) follows by Proposition 7.14. For the remaining direction (3) \Rightarrow (2), note that the 1-element order 1 trivially has effective weak MSO-uniformisation. Furthermore, we have shown in Proposition 7.16 that the same holds for ω . As every ordinal $\alpha < \omega^\omega$ can be obtained from ω and 1 using addition and multiplication, it follows by Lemma 7.15 that every ordinal less than ω^ω has effective weak MSO-uniformisation. \square

For arbitrary linear orders, we obtain the following result.

Corollary 7.18. Let \mathfrak{A} be a linear order. The following statements are equivalent.

- (1) \mathfrak{A} has weak MSO-uniformisation.
- (2) \mathfrak{A} has effective weak MSO-uniformisation.
- (3) $\text{WO}(\mathfrak{A}) < \infty$.

Proof. We start by proving that, if Ξ is a well-order definable in \mathfrak{A} , then $\langle A, \leq \rangle$ has [effective] weak MSO-uniformisation if, and only if, $\langle A, \Xi \rangle$ has. By symmetry, it is sufficient to prove one direction. Hence, assume that

$\langle A, \sqsubseteq \rangle$ has weak MSO-uniformisation. To show that so does $\langle A, \leq \rangle$, consider a formula $\varphi(\bar{X}, \bar{Y}) \in \text{MSO}$. Let $\vartheta(x, y; \bar{P})$ be the formula defining \sqsubseteq . By Lemma 6.11 (b), there exists a formula defining the original ordering \leq in terms of \sqsubseteq . Let φ' be the formula obtained from φ by replacing every occurrence of the relation \leq by χ . By assumption, there exists a uniformisation $\psi'(\bar{X}, \bar{Y}; \bar{S})$ of φ' over $\langle A, \sqsubseteq \rangle$. Let ψ be the formula obtained from ψ' by replacing every occurrence of the relation \sqsubseteq by ϑ . Then ψ is a uniformisation of φ . Furthermore, if ψ' is computable from φ' , we can also compute ψ from φ .

It remains to prove the above equivalences. (2) \Rightarrow (1) is trivial.

(1) \Rightarrow (3) Suppose that \mathfrak{A} has weak MSO-uniformisation. By the above claim every ordinal definable inside of \mathfrak{A} also has MSO-uniformisation. By Theorem 7.17, it follows that every such ordinal is less than ω^ω . Furthermore, since \mathfrak{A} has weak MSO-choice, it follows by Theorem 6.3 that at least one such ordinal exist. Hence, $\text{WO}(\mathfrak{A}) < \infty$.

(3) \Rightarrow (2) If $\text{WO}(\mathfrak{A}) < \infty$, we can define a well-ordering \sqsubseteq on A of order type $\alpha < \omega^\omega$. By Theorem 7.17, $\langle A, \sqsubseteq \rangle$ has effective weak MSO-uniformisation. By the above claim, so does $\langle A, \leq \rangle$. \square

Uniformisation For Trees

The above results can now be generalised to trees. Recall the notion of a tame tree from Definition 6.2 which we introduced above to characterise when a generalised tree has weak MSO-choice. We use the following stronger notion to characterise weak MSO-uniformisation.

Definition 7.19. A generalised tree \mathfrak{T} is *very tame* if it is tame and there exists a constant $l < \omega$ such that $\text{WO}(\beta) \leq l$, for all branches β .

Theorem 7.20. Let \mathfrak{T} be a generalised tree. The following statements are equivalent.

- (1) \mathfrak{T} has weak MSO-uniformisation.
- (2) \mathfrak{T} has effective weak MSO-uniformisation.
- (3) \mathfrak{T} is very tame.

Before giving the proof, let us note the following composition lemma.

Lemma 7.21. Given an MSO-formula $\varphi(\bar{X}; \bar{Y})$, we can compute an MSO-formula $\psi(\bar{X}; \bar{Z})$ with the following property. For every generalised tree \mathfrak{T} , every branch β , and all sets $A_i \subseteq \beta$ and $P_i \subseteq T$,

$$\mathfrak{T} \models \varphi(\bar{A}, \bar{P}) \quad \text{iff} \quad \langle \beta, \leq \rangle \models \psi(\bar{A}; \bar{Q}),$$

where the predicates $\bar{Q} = (Q_\theta)_\theta$ are defined by

$$Q_\theta := \{ v \in \beta \mid \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}, \bar{P} \rangle|_{B_v}) = \theta \},$$

$$B_v := \{ w \in T \mid w > v \text{ and } w \not\geq z \text{ for all } z \in \beta \text{ with } z > v \}.$$

Proof. Fix a formula $\varphi(\bar{X}, \bar{Y})$. Given a branch β , we can write \mathfrak{T} as a generalised sum followed by a quantifier-free interpretation τ :

$$\langle \mathfrak{T}, \bar{P} \rangle = \tau \left(\sum_{v \in \beta} \langle \mathfrak{T}, \bar{P}, \{v\} \rangle|_{B_v \cup \{v\}} \right).$$

Consequently, there exist formulae $\psi_o, \chi_o, \dots, \chi_{n-1}$ such that

$$\mathfrak{T} \models \varphi(\bar{A}, \bar{P}) \quad \text{iff} \quad \langle \beta, \leq \rangle \models \psi_o(\llbracket \chi_o(\bar{A}, \bar{P}) \rrbracket, \dots), \quad \text{for all } \bar{A}, \bar{P}.$$

Furthermore, there exists a quantifier-free interpretation σ such that

$$\langle \mathfrak{T}, \bar{A}, \bar{P} \rangle|_{B_v \cup \{v\}} = \sigma(\langle \mathfrak{T}, \bar{A}, \bar{P} \rangle|_{\{v\}} \oplus \langle \mathfrak{T}, \emptyset, \bar{P} \rangle|_{B_v}),$$

for $v \in \beta$, $A_i \subseteq \beta$, and $P_i \subseteq T$. We can therefore construct formulae $\vartheta_i(x; \bar{Z})$ such that

$$\langle \mathfrak{T}, \bar{A}, \bar{P} \rangle|_\beta \models \vartheta_i(v; \bar{Q}) \quad \text{iff} \quad \langle \mathfrak{T}, \bar{A}, \bar{P} \rangle|_{B_v \cup \{v\}} \models \chi_i,$$

for v, \bar{A}, \bar{P} as above. Substituting these formulae into ψ_o , we obtain an MSO-formula ψ such that

$$\langle \beta, \leq \rangle \models \psi(\bar{A}, \bar{P}, \bar{Q}) \quad \text{iff} \quad \langle \beta, \leq \rangle \models \psi_o(\llbracket \chi_o(\bar{A}, \bar{P}) \rrbracket, \dots)$$

$$\text{iff} \quad \mathfrak{T} \models \varphi(\bar{A}; \bar{P}). \quad \square$$

Proof of Theorem 7.20. (2) \Rightarrow (1) is trivial.

(1) \Rightarrow (3) Suppose that \mathfrak{T} has weak MSO-uniformisation. Then it in particular has weak MSO-choice, which implies by Theorem 6.3 that \mathfrak{T} is tame. Let $\varphi(X, Y)$ be the formula stating that

“ X is a chain with no maximal element, $Y \subseteq X$ is cofinal, and the order type of Y is ω .”

By assumption, $\varphi(X, Y)$ has a uniformisation $\psi(X, Y; \bar{P})$. We can use Lemma 7.21 to construct formulae $\vartheta, \chi_0, \dots, \chi_{n-1}$ such that

$$\mathfrak{T} \models \varphi(X, Y; \bar{P}) \quad \text{iff} \quad \langle \beta, \leq \rangle \models \vartheta(X, Y; \bar{Q}),$$

for every branch β and all sets $X, Y \subseteq \beta$. It follows that ϑ is a uniformisation of φ over $\langle \beta, \leq \rangle$. Since the formula ϑ does not depend on β , it follows by Lemma 7.13 that $\text{WO}(\beta)$ is bounded.

(3) \Rightarrow (2) Suppose that \mathfrak{T} is very tame. By definition, this means that \mathfrak{T} is thin and there exist numbers $k, n, l < \omega$ such that the branching degree of \mathfrak{T} is bounded by n and

$$\text{HR}(\beta) \leq k \quad \text{and} \quad \text{WO}(\beta) \leq l, \quad \text{for every branch } \beta.$$

Furthermore, it follows by Proposition V.5.9 that \mathfrak{T} has a skeleton S , and we have seen in Lemma 6.12 that S is definable. Note that S forms a well-founded tree $\mathfrak{S} = \langle S, \leq \rangle$ with tree-order

$$\beta \leq \gamma \quad : \text{iff} \quad u \leq_{\text{pf}} v, \text{ for some } u \in \beta \text{ and } v \in \gamma.$$

Let U be a set containing one vertex from each subbranch $\beta \in S$. Note that U with the order induced by that of S forms a tree $\mathfrak{U} = \langle U, \leq \rangle$ isomorphic to \mathfrak{S} . For $u \in U$, we denote by β_u the subbranch $\beta_u \in S$ with $u \in \beta_u$. Then the map $u \mapsto \beta_u$ is an isomorphism $\mathfrak{U} \rightarrow \mathfrak{S}$. Finally, for $u \in U$ and $v \in T$, we set

$$\begin{aligned} A_u &:= \{ v \in T \mid v \geq_{\text{pf}} w, \text{ for some } w \in \beta_u \}, \\ E_v &:= \{ w \in \text{Suc}(u) \mid v \in \beta_u, \Downarrow w \cap \Downarrow \beta_u = \Downarrow v \}. \end{aligned}$$

We start by proving the following claims.

- (I) The tree-order on U which is induced by that of S is definable.
- (II) Given $u \in U$, we can define the following sets: the subbranch $\beta_u \in S$; the successor set $\text{Suc}(u)$; and the subtree A_u . Similarly, given $v \in T$, we can define E_v .
- (III) $\text{Suc}(u)$ is a set of representatives for the partition $(A_v)_{v \in \text{Suc}(u)}$.
- (IV) $|E_v| \leq n$ and $\text{Suc}(u) = \sum_{v \in \beta_u} E_v$.
- (V) The tree-order on β_u induces a (definable) ordering on $\text{Suc}(u)$ with

$$\text{WO}(\text{Suc}(u)) = \text{WO}(\beta_u) \leq l.$$

- (VI) For every $m < \omega$, there exists a constant $m_+ < \omega$ such that, for all vertices $u \in U$ and all predicates \bar{P} , the theory

$$\text{Th}_{\text{MSO}}^{m_+}(\langle \mathfrak{T}, \bar{P} \rangle|_{\beta_u}, \bar{Q}_u) \quad \text{determines} \quad \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}, \bar{P} \rangle|_{A_u}),$$

where the predicates $\bar{Q}_u = (Q_{u, \theta})_\theta$ are defined by

$$\begin{aligned} Q_{u, \theta} &:= \{ v \in \beta_u \mid \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}, \bar{P} \rangle|_{B_v}) = \theta \}, \\ B_v &:= \bigcup_{w \in E_v} A_w. \end{aligned}$$

- (I) and (II) follow by definability of S (cf. Lemma 6.12).
- (III) By definition, we have $v \in A_v$. Conversely, if v and v' are distinct successors of u , then $A_v \cap A_{v'} = \emptyset$.
- (IV) We have $|E_v| \leq n$ since the branching degree of \mathfrak{T} is bounded n . Clearly, $E_v \cap E_w = \emptyset$, for $v \neq w$, and every $w \in \text{Suc}(u)$ belongs to some set E_v with $v \in \beta_u$.
- (V) By (IV), $\text{Suc}(u) = \text{Suc}(\beta_u)$ is obtained from β_u by replacing each vertex by a finite chain of length at most n . Let γ_u be the branch of \mathfrak{T} containing β_u . Then

$$\text{WO}(\text{Suc}(u)) = \text{WO}(\beta_u) \leq \text{WO}(\gamma_u) \leq l.$$

Furthermore, choosing n monadic predicates \bar{W} such that $|W_i \cap E_v| \leq 1$, for all i, v , we can define the ordering on $\text{Suc}(u)$ with the help of the ordering \leq_{pf} on β_u and the parameters \bar{W} .

(v1) By Lemma 7.21, exist sets \bar{Q} such that, for every MSO_m -theory θ , there is some formula ψ_θ with

$$\langle \mathfrak{T}, \bar{P} \rangle|_{A_u} \models \wedge \theta \quad \text{iff} \quad \langle \mathfrak{T}, \bar{P} \rangle|_{\beta_u} \models \psi_\theta(\bar{Q}).$$

We can take for m_+ the maximum of the quantifier-ranks of these formulae ψ_θ .

Having established the above claims, we now construct the desired uniformisation for a formula $\varphi(\bar{X}, \bar{Y})$ similarly as in Lemma 7.15. Given sets \bar{X} , we consider its restriction to each subbranch $\beta \in S$. Since $\text{WO}(\beta) \leq l$, β has uniformisation, we can define unique subsets \bar{Y}_β in β for $\bar{X}|_\beta$. The union of all these will be the desired sets \bar{Y} . To make this work, our choices \bar{Y}_β have to be consistent with each other. With each $u \in U$ we associate an instruction of how to choose \bar{Y}_u in β_u . This instruction must be consistent with the already chosen subset of A_u . As \mathbb{U} is well-founded, we can then ensure consistency by induction on u .

The details are as follows. Let \bar{P} be all the parameters used in φ and in the definitions of U and S , and set $\mathfrak{T}^* := \langle T, \leq_{\text{pf}}, \bar{P} \rangle$. Fix sets \bar{X} in T such that

$$\mathfrak{T}^* \models \exists \bar{Y} \varphi(\bar{X}, \bar{Y}).$$

Let m be the quantifier rank of φ and m_+ the constant from (v1), set $s := |P| + |X|$ and $t := |Y|$, let r be the number of predicates \bar{Q}_u from (v1), and let Θ_j^i be the set of all MSO_i -theories with j monadic parameters.

By induction on the distance of $u \in U$ from the root, we will construct

- two MSO -formulae $\psi'_u(\bar{X}, \bar{Y}; \bar{C}_u)$ and $\psi''_u(\bar{X}, \bar{Y}; \bar{M}_u^1, \bar{D}_u)$ with monadic parameters \bar{C}_u, \bar{M}_u^1 in $\text{Suc}(u)$ and \bar{D}_u in β_u , respectively,
- two labellings $\mu_u^1 : \text{Suc}(u) \rightarrow \Theta_{s+t}^m$ and $\mu_u^2 : \text{Suc}(u) \rightarrow \Theta_{s+t+r}^{m_+}$,
- sets \bar{Y}_u in β_u , and
- a theory θ

such that the union of the sets \bar{Y}_u will be the desired set \bar{Y} and the formulae ψ'_u and ψ''_u can be used to determine μ_u^1, μ_u^2 , and \bar{Y}_u .

First, suppose that $u \in U$ is the root of \mathbb{U} . We say that a pair of labellings $\mu_u^1 : \text{Suc}(u) \rightarrow \Theta_{s+t}^m$ and $\mu_u^2 : \text{Suc}(u) \rightarrow \Theta_{s+t+r}^{m_+}$ is θ -consistent, for some

theory θ , if there exist sets \bar{Y} in A_u such that

$$\begin{aligned} \mu_u^1(v) &= \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{A_v}), & \text{for all } v \in \text{Suc}(u), \\ \mu_u^2(v) &= \text{Th}_{\text{MSO}}^{m_+}(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_v}, \bar{Q}_v'), & \text{for all } v \in \text{Suc}(u), \end{aligned}$$

where

$$\bar{Q}_{v,\sigma}' := \{ w \in \beta_v \mid \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{B_w}) = \sigma \},$$

and, for all \bar{Y} ,

$$\text{Th}_{\text{MSO}}^{m_+}(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_u}, \bar{Q}_u) = \theta \quad \text{implies} \quad \mathfrak{T}^* \models \varphi(\bar{X}, \bar{Y}),$$

where

$$\bar{Q}_{u,\sigma} := \{ v \in \beta_u \mid \bigoplus_{w \in E_v} \mu_u^1(w_o) = \sigma \}$$

are the sets from (v1) (where \bigoplus is the operation on theories corresponding to taking the disjoint union of the underlying structures). Note that $\mathfrak{T}^* \models \exists \bar{Y} \varphi(\bar{X}, \bar{Y})$ implies that there exist a theory θ and labellings μ_u^1, μ_u^2 that are θ -consistent. We choose θ minimal (in some fixed enumeration of all theories) with this property. Let \bar{M}_u^i be unary predicates encoding μ_u^i . By the composition theorem for finite disjoint unions, the sets \bar{Q}_u can be defined in terms of the labelling μ_u^1 . Consequently, there exists an MSO -formula φ' such that

$$\langle \mathfrak{T}^*, \bar{X} \rangle|_{A_u \setminus \beta_u} \models \varphi'(\bar{M}_u^1, \bar{M}_u^2) \quad \text{iff} \quad \mu_u^1, \mu_u^2 \text{ are } \theta\text{-consistent}.$$

By (v1), we obtain a formula φ'' and monadic predicates \bar{R} in $\text{Suc}(u)$ (which are definable in $\langle \mathfrak{T}^*, \bar{X} \rangle$) such that

$$\mathfrak{T}^*|_{\text{Suc}(u)} \models \varphi''(\bar{R}, \bar{M}_u^1, \bar{M}_u^2) \quad \text{iff} \quad \mu_u^1, \mu_u^2 \text{ are } \theta\text{-consistent}.$$

By (v) and Corollary 7.18, it follows that φ'' has a uniformisation ψ'_u (with monadic parameters \bar{C}_u) over $\mathfrak{T}^*|_{\text{Suc}(u)}$ (where we consider \bar{R} as the given

variables \bar{X} , and $\bar{M}^1 \bar{M}^2$ as the unique ones \bar{Y} we want to compute). Similarly, the statement

$$\text{Th}_{\text{MSO}}^+(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_u}, \bar{Q}_u) = \theta$$

has a uniformisation $\psi_u''(\bar{X}, \bar{Y}; \bar{M}_u^1, \bar{D}_u)$ with additional parameters \bar{D}_u . We choose for \bar{M}_u^1 , \bar{M}_u^2 , and \bar{Y}_u the sets determined by ψ_u' and ψ_u'' .

For the inductive step, suppose that we have already defined ψ_w' , ψ_w'' , \bar{C}_w , \bar{D}_w , μ_w^1 , μ_w^2 , and \bar{Y}_w , and let $u \in \text{Suc}(w)$ be a successor of w . By consistency of μ_w^1 , μ_w^2 we can find sets \bar{Y}_u' in A_w such that

$$\begin{aligned} \mu_w^1(u) &= \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}^*, \bar{X}, \bar{Y}_u' \rangle|_{A_u}), \\ \mu_w^2(u) &= \text{Th}_{\text{MSO}}^{m+}(\langle \mathfrak{T}^*, \bar{X}, \bar{Y}_u', \bar{Q}_v' \rangle|_{\beta_u}). \end{aligned}$$

Similarly to above, we say that a pair of labellings $\mu_u^1 : \text{Suc}(u) \rightarrow \mathcal{O}_{s+t}^m$ and $\mu_u^2 : \text{Suc}(u) \rightarrow \mathcal{O}_{s+t+r}^{m+}$ is *consistent* if

$$\begin{aligned} \mu_u^1(v) &= \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{A_v}), \quad \text{for all } v \in \text{Suc}(u), \\ \mu_u^2(v) &= \text{Th}_{\text{MSO}}^{m+}(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_v}, \bar{Q}_v'), \quad \text{for all } v \in \text{Suc}(u), \end{aligned}$$

and $\text{Th}_{\text{MSO}}^+(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_u}, \bar{Q}_u) = \mu_w^2(u)$

$$\text{implies } \text{Th}_{\text{MSO}}^m(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{A_u}) = \mu_w^1(u), \quad \text{for all } \bar{Y},$$

where \bar{Q}_u and \bar{Q}_v' are the predicates from (vi) as above. Let \bar{M}_u^i be unary predicates encoding μ_u^i . As above, we obtain a formula φ'' and monadic predicates \bar{R} in $\text{Suc}(u)$ such that

$$\mathfrak{T}^*|_{\text{Suc}(u)} \models \varphi''(\bar{R}, \bar{M}_u^1, \bar{M}_u^2) \quad \text{iff} \quad \mu_u^1, \mu_u^2 \text{ are consistent.}$$

Then φ'' has a uniformisation ψ_u' (with monadic parameters \bar{C}_u) over $\mathfrak{T}^*|_{\text{Suc}(u)}$. Similarly, the statement

$$\text{Th}_{\text{MSO}}^+(\langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_u}, \bar{Q}_u) = \mu_w^2(u)$$

has a uniformisation $\psi_u''(\bar{X}, \bar{Y}; \bar{M}_u^1, \bar{D}_u)$ with additional parameters \bar{D}_u . We choose for \bar{M}_u^1 , \bar{M}_u^2 , and \bar{Y}_u the sets determined by ψ_u' and ψ_u'' .

This concludes the inductive definition of ψ_u' , ψ_u'' , μ_u^1 , μ_u^2 , \bar{C}_u , \bar{D}_u , \bar{Y}_u , and θ . Note that both the quantifier-ranks and the number of monadic parameters of the formulae ψ_u' and ψ_u'' are bounded. Hence, there are only finitely many possible choices for these formulae. Let \bar{C} and \bar{D} be the unions of the sets $(\bar{C}_u)_u$ and $(\bar{D}_u)_u$, respectively. The desired uniformisation of φ is the formula stating that there exists predicates \bar{M}^1 , \bar{M}^2 such that

$$\begin{aligned} \langle \mathfrak{T}^*, \bar{R}, \bar{M}^1 \bar{M}^2 \rangle|_{\text{Suc}(u)} &\models \psi_u', \quad \text{for all } u \in U, \\ \langle \mathfrak{T}^*, \bar{X}, \bar{Y} \rangle|_{\beta_u} &\models \psi_u''(\bar{X}, \bar{Y}; \bar{M}^1|_{\text{Suc}(u)}, \bar{D}_u), \quad \text{for all } u \in U. \end{aligned}$$

To see that we can express this in MSO, note that the formulae ψ_u' and ψ_u'' can both be determined from the labellings \bar{M}^1 , \bar{M}^2 . Furthermore, all steps of the above construction are effective. Hence, we can compute the resulting formula from φ . \square

8 First-Order Logic

We have already shown in Theorem II.5.1 that the logics FO and LTL are equivalent of over finite words. The aim of this section is to generalise this statement to all Dedekind-complete linear orders. To do so, we need to introduce a dual version of the modal operator U that looks backwards. But first, let us define the notion of completeness we use.

Definition 8.1. A linear order \mathcal{Q} is *Dedekind-complete* if, for every decomposition $A = I + K$ with $I, K \neq \emptyset$, the supremum $\sup I$ and the infimum $\inf K$ exist. \dashv

Thus, Dedekind-completeness is weaker than completeness since unbounded sets are not required to have a supremum or an infimum, that is, we do not require the order to have a least element or a greatest one.

Definition 8.2. Let Σ be an alphabet.

(a) The formulae of *bidirectional linear temporal logic* BLTL are built up from atomic formulae of the form P_a with $a \in \Sigma$ using boolean operations

and two modal operators of the form

$$\varphi \text{ U } \psi \quad \text{and} \quad \varphi \text{ S } \psi, \quad \text{for } \varphi, \psi \in \text{BLTL}.$$

We read these formulae as, respectively, ‘ φ until ψ ’ and ‘ φ since ψ ’. The semantics is defined as follows. Given a Σ -labelled linear order \mathfrak{A} and an element $s \in A$, we set

$$\begin{aligned} \mathfrak{A}, s \models P_a & : \text{iff} \quad s \in P_a^{\mathfrak{A}}, \\ \mathfrak{A}, s \models \varphi \text{ U } \psi & : \text{iff} \quad \text{there is some } t > s \text{ such that } \mathfrak{A}, t \models \psi \\ & \quad \text{and } \mathfrak{A}, u \models \varphi \text{ for all } s < u < t, \\ \mathfrak{A}, s \models \varphi \text{ S } \psi & : \text{iff} \quad \text{there is some } t < s \text{ such that } \mathfrak{A}, t \models \psi \\ & \quad \text{and } \mathfrak{A}, u \models \varphi \text{ for all } t < u < s. \end{aligned}$$

Boolean operations are interpreted in the usual way.

(b) (*Forward*) *linear temporal logic* LTL is the fragment of BLTL without the operator S.

(c) It is common to use the following abbreviations:

$$\begin{aligned} X\varphi &:= \text{false U } \varphi && (\text{‘next } \varphi\text{’}), \\ Y\varphi &:= \text{false S } \varphi && (\text{‘yesterday } \varphi\text{’}), \\ F\varphi &:= \text{true U } \varphi && (\text{‘eventually } \varphi\text{’}), \\ P\varphi &:= \text{true S } \varphi && (\text{‘past } \varphi\text{’}), \\ G\varphi &:= \neg F\neg\varphi && (\text{‘always } \varphi\text{’}), \\ H\varphi &:= \neg P\neg\varphi && (\text{‘hitherto } \varphi\text{’}). \end{aligned}$$

Sometimes it is also convenient to introduce starred versions of these operators that include the current position:

$$\varphi \text{ U}^* \psi := \varphi \vee \varphi \text{ U } \psi \quad \text{and} \quad \varphi \text{ S}^* \psi := \varphi \vee \varphi \text{ S } \psi,$$

and analogously for F^* , G^* , P^* , and H^* .

Examples. (a) The formula

$$\text{GFtrue}$$

states that there is no maximal element.

(b) The formula

$$\neg X\text{true} \equiv \neg(\text{false U true})$$

states that the given element does not have an immediate successor.

(c) Similarly,

$$\neg(\neg P \text{ U true})$$

states that, above the given element, the elements in P are not bounded from below. J

Theorem 8.3 (Kamp). *FO and BLTL are equivalent over the class of all coloured Dedekind-complete linear orders.*

‘Equivalent’ here means that there are translations (in both directions) between FO-formulae $\varphi(x)$ with one free variable and BLTL-formulae φ' such that

$$\mathfrak{A} \models \varphi(a) \quad \text{iff} \quad \mathfrak{A}, a \models \varphi',$$

for all coloured Dedekind-complete linear orders \mathfrak{A} and all $a \in A$.

The remainder of this section is devoted to a proof of this theorem. We start with the trivial direction: the translation of BLTL into FO.

Lemma 8.4. *For every BLTL-formula φ , there exists an FO-formula $\varphi^*(x)$ such that*

$$\mathfrak{A} \models \varphi(s) \quad \text{iff} \quad \mathfrak{A}, s \models \varphi^*,$$

for every linear order \mathfrak{A} and all $s \in A$.

Proof. We construct φ^* by induction on φ .

$$\begin{aligned} P_a^*(x) &:= P_a x, \\ (\varphi \wedge \psi)^*(x) &:= \varphi^*(x) \wedge \psi^*(x), \\ (\neg\varphi)^*(x) &:= \neg\varphi^*(x), \\ (\varphi \cup \psi)^*(x) &:= \exists y[y > x \wedge \varphi^*(y) \wedge \forall z[x < z < y \rightarrow \psi^*(z)]], \\ (\varphi S \psi)^*(x) &:= \exists y[y < x \wedge \varphi^*(y) \wedge \forall z[y < z < x \rightarrow \psi^*(z)]]. \quad \square \end{aligned}$$

The other direction is more involved. We introduce the following normal form for FO-formulae which can then easily be translated into BLTL.

Definition 8.5. (a) A *basic temporal formula* is a first-order formula of the form

$$\begin{aligned} \varphi(\bar{x}) = \exists \bar{z} \Big[& \bigwedge_{i < m} x_i = z_{\sigma(i)} \wedge \bigwedge_{i < n} z_i < z_{i+1} \wedge \bigwedge_{i < n} \alpha_i(z_i) \\ & \wedge \bigwedge_{i < n} \forall y[z_i < y < z_{i+1} \rightarrow \beta_{i+1}(y)] \\ & \wedge \forall y[y > z_{n-1} \rightarrow \beta_n(y)] \wedge \forall y[y < z_0 \rightarrow \beta_0(y)] \Big] \end{aligned}$$

where $\bar{x} = \langle x_0, \dots, x_{m-1} \rangle$, $\bar{z} = \langle z_0, \dots, z_n \rangle$, $\sigma : [m] \rightarrow [n]$ is an arbitrary function, and α_i and β_i are (translations of) BLTL-formulae. To simplify notation, we will write such a formula more suggestively as

$$\langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_\sigma(\bar{x}).$$

(b) A *standard temporal formula* is a disjunction of basic temporal formulae. J

We will prove below that every first-order formula can be translated into a standard temporal formula. Before doing so, let us note how to translate standard temporal formulae into BLTL.

Lemma 8.6. *Over coloured linear orders, every standard temporal formula with at most one free variable is equivalent to a BLTL-formula.*

Proof. First, suppose that

$$\varphi(x_o) = \langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_\sigma(x_o)$$

has exactly one free variable x_o . Let $k := \sigma(o)$. Then φ is equivalent to the BLTL-formula

$$\begin{aligned} \psi &:= \alpha_k \wedge [\beta_{k+1} \cup [\alpha_{k+1} \wedge [\beta_{k+2} \cup \dots \cup [\alpha_{n-1} \wedge G\beta_n] \dots]]] \\ &\wedge \alpha_k \wedge [\beta_k S [\alpha_{k-1} \wedge [\beta_{k-1} S \dots S [\alpha_o \wedge H\beta_o] \dots]]]. \end{aligned}$$

It remains to consider the case where

$$\varphi = \langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle$$

has no free variables. Then

$$\varphi \equiv \exists x_o \langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_\sigma(x_o)$$

with $\sigma(o) = o$. Let $\psi \in \text{BLTL}$ be the translation of

$$\langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_\sigma(x_o).$$

Then φ is equivalent to the BLTL-formula $\text{PF}\psi$. □

It remains to show that every FO-formula is equivalent to a standard temporal formula. To do so we prove that standard temporal formulae are closed under all operations of first-order logic. But first note that we only need to consider formulae with at most two free variables.

Lemma 8.7. *Every basic temporal formula is equivalent to a conjunction of such formulae each of which have at most two free variables.*

Proof. Consider a basic temporal formula

$$\varphi(\bar{x}) = \langle\langle \beta_0[\alpha_0]\beta_1[\alpha_1] \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_\sigma(\bar{x}),$$

with $m := |\bar{x}| > 2$ free variables. We prove the claim by induction on m .

First, consider the case where $\sigma(i) = \sigma(j)$ for some $i \neq j$. Let \bar{x}' be all variables except for x_j and let $\sigma' : [m-1] \rightarrow [n]$ be the restriction of σ to $[m] \setminus \{j\}$. Then

$$\begin{aligned} \varphi(\bar{x}) &\equiv \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma'}(\bar{x}') \\ &\quad \wedge \langle\langle \text{true}[\text{true}]\text{true} \rangle\rangle_{\tau}(x_i, x_j), \end{aligned}$$

where τ is the unique function $[2] \rightarrow [1]$. Hence, the claim follow by inductive hypothesis.

It remains to consider the case where σ is injective. Pick an index $k < m$ such that

$$\sigma(i) < \sigma(k) < \sigma(j), \quad \text{for some } i, j.$$

Let \bar{x}' be the set of all variables x_i with $\sigma(i) < \sigma(k)$ and \bar{x}'' those with $\sigma(i) > \sigma(k)$. Let $\mu' : [|\bar{x}'|] \rightarrow [|\bar{x}|]$ and $\mu'' : [|\bar{x}''|] \rightarrow [|\bar{x}|]$ be the functions such that

$$x'_i = x_{\mu'(i)} \quad \text{and} \quad x''_i = x_{\mu''(i)}.$$

Then

$$\begin{aligned} \varphi(\bar{x}) &\equiv \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_k[\alpha_k]\text{true} \rangle\rangle_{\sigma'}(\bar{x}'x_k) \\ &\quad \wedge \langle\langle \text{true}[\alpha_k]\beta_{k+1} \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma''}(\bar{x}''x_k), \end{aligned}$$

where $\sigma' : [|\bar{x}'|+1] \rightarrow [k+1]$ and $\sigma'' : [|\bar{x}''|+1] \rightarrow [n-k]$ are the functions with

$$\begin{aligned} \sigma'(i) &:= \begin{cases} \sigma(\mu'(i)) & \text{if } i < |\bar{x}'|, \\ k & \text{if } i = |\bar{x}'|, \end{cases} \\ \sigma''(i) &:= \begin{cases} \sigma(\mu''(i)) - k & \text{if } i < |\bar{x}''|, \\ o & \text{if } i = |\bar{x}''|. \end{cases} \end{aligned}$$

Hence, the claim follows by inductive hypothesis. \square

The closure properties of standard temporal formulae are relatively straightforward to establish, except for negation, which we will treat separately below.

Lemma 8.8. *The set of standard temporal formulae is closed under disjunctions, conjunctions, and existential quantifiers.*

Proof. Closure under disjunctions and existential quantifiers is trivial. Before taking a look at conjunction let us show that, if $\varphi(\bar{x})$ is a basic temporal formula with $m := |\bar{x}|$ free variables, then $\varphi(\bar{x}y)$ (considered as a formula with one additional unused variable y) is equivalent to a standard temporal formula. If

$$\varphi(\bar{x}) = \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma}(\bar{x}),$$

there are $2n+1$ possible intervals where the element y can be located. This results in the definition

$$\varphi(\bar{x}) \equiv \bigvee_{0 \leq i < n} \varphi'_i(\bar{x}y) \vee \bigvee_{-1 \leq i < n} \varphi''_i(\bar{x}y),$$

where

$$\begin{aligned} \varphi'_i(\bar{x}y) &:= \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma'_i}(\bar{x}y), \\ \varphi''_i(\bar{x}y) &:= \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_i[\beta_i]\beta_i \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma''_i}(\bar{x}y), \end{aligned}$$

and

$$\begin{aligned} \sigma'_i(k) &:= \begin{cases} \sigma(k) & \text{if } k < m, \\ i & \text{if } k = m, \end{cases} \\ \sigma''_i(k) &:= \begin{cases} \sigma(k) & \text{if } k < m \text{ and } \sigma(k) \leq i, \\ \sigma(k) + 1 & \text{if } k < m \text{ and } \sigma(k) > i, \\ i + 1 & \text{if } k = m. \end{cases} \end{aligned}$$

For closure under conjunctions, let

$$\begin{aligned} \varphi(\bar{x}) &= \langle\langle \beta_o[\alpha_o]\beta_o \dots \beta_{m-1}[\alpha_{m-1}]\beta_m \rangle\rangle_{\sigma}(\bar{x}) \\ \psi(\bar{x}) &= \langle\langle \delta_o[\gamma_o]\delta_o \dots \delta_{n-1}[\gamma_{n-1}]\delta_n \rangle\rangle_{\tau}(\bar{x}) \end{aligned}$$


be two basic temporal formulae. (By the above remark, we can assume that the free variables \bar{x} are the same in both formulae.) The formula φ guesses some increasing m -tuple \bar{a} and ψ guesses some n -tuple \bar{b} . The union $\bar{a} \cup \bar{b}$ is then an increasing l -tuple where $l \leq m+n$ depends on how many elements a_i coincide with some b_j . In particular, if a_i and b_j correspond to the same free variable x_k , i.e., if $i = \sigma(k)$ and $j = \tau(k)$, then $a_i = b_j$. Consequently, we can describe the embeddings of \bar{a} and \bar{b} in $\bar{a} \cup \bar{b}$ by two injective monotone functions $\xi : [m] \rightarrow [l]$ and $\zeta : [n] \rightarrow [l]$ such that

$$\text{rng } \xi \cup \text{rng } \zeta = [l] \quad \text{and} \quad \xi \circ \sigma = \zeta \circ \tau.$$

Then

$$\varphi \wedge \psi \equiv \bigvee_{\xi, \zeta} \vartheta_{\xi\zeta},$$

where, for each pair ξ, ζ as above, the basic temporal formula $\vartheta_{\xi\zeta}$ describes those tuples $\bar{a} \cup \bar{b}$ where \bar{a} arises from a guess made by φ and \bar{b} from a guess made by ψ . Thus, $\vartheta_{\xi\zeta}$ guesses an l -tuple and checks that each interval satisfies the correct formulae. The general definition turns out to be rather technical and not very enlightening. Therefore, we omit it and just illustrate the general idea using an example. Suppose that $\xi : [5] \rightarrow [8]$ and $\zeta : [4] \rightarrow [8]$ are the functions with ranges $\text{rng } \xi = \{0, 2, 3, 4, 6\}$ and $\text{rng } \zeta = \{1, 4, 5, 7\}$, and that $\sigma : [1] \rightarrow [5]$ and $\tau : [1] \rightarrow [4]$ are given by $\sigma(0) = 3$ and $\tau(0) = 1$.



Then we obtain

$$\begin{aligned} \vartheta_{\xi\zeta}(\bar{x}) := & \langle\langle \beta_0 \wedge \delta_0[\alpha_0 \wedge \delta_0] \\ & \beta_1 \wedge \delta_0[\beta_1 \wedge \gamma_0] \\ & \beta_1 \wedge \delta_1[\alpha_1 \wedge \delta_0] \\ & \beta_2 \wedge \delta_1[\alpha_2 \wedge \delta_0] \\ & \beta_3 \wedge \delta_1[\alpha_3 \wedge \gamma_1] \\ & \beta_4 \wedge \delta_2[\beta_3 \wedge \gamma_2] \\ & \beta_4 \wedge \delta_3[\alpha_4 \wedge \delta_2] \\ & \beta_5 \wedge \delta_4[\beta_5 \wedge \gamma_3] \\ & \beta_5 \wedge \delta_4 \rangle\rangle_v(\bar{x}) \end{aligned}$$

with $v(o) = 4$. □

It remains to prove closure under negations. We will do so in several steps, starting with simple formulae and working our way up to the general case. To simplicity, we will omit the subscript σ from all basic temporal formulae with exactly two free variables

$$\langle\langle \beta_0[\alpha_0] \beta_1 \dots \beta_{n-1}[\alpha_{n-1}] \beta_n \rangle\rangle(x_0, x_1)$$

with the convention that σ maps x_0 to the first position and x_1 to the last one.

Lemma 8.9. *Let $\varphi(x_0, x_1)$ be a basic temporal formula of the form*

$$\varphi(x_0, x_1) = \langle\langle \beta_0[\text{true}] \beta_1[\alpha_1] \dots [\alpha_{n-2}] \beta_{n-1}[\text{true}] \beta_n \rangle\rangle_\sigma(x_0, x_1)$$

where $n \geq 3$, $\sigma(o) = o$, $\sigma(1) = n-1$, and $\beta_i = \text{true}$, for all i . Over the class of Dedekind-complete linear orders, the negation $\neg\varphi$ is equivalent to a standard temporal formulae.

Proof. We prove the claim by induction on n . If $n = 3$, we have

$$\begin{aligned} & \neg \langle\langle \text{true}[\text{true}] \text{true}[\alpha_1] \text{true}[\text{true}] \text{true} \rangle\rangle_\sigma(x_0, x_1) \\ & \equiv \langle\langle \text{true}[\text{true}] \neg \alpha_1[\text{true}] \text{true} \rangle\rangle_\sigma(x_0, x_1). \end{aligned}$$

For the inductive step, suppose that $n > 3$ and that we have already some standard temporal formula $\psi(x_o, x_i)$ that is equivalent to

$$\neg \langle\langle \text{true}[\text{true}] \text{true}[\alpha_2] \dots [\alpha_{n-2}] \text{true}[\text{true}] \text{true} \rangle\rangle(x_o, x_i).$$

Further, let

$$\vartheta(x) := (\forall y > x) \exists z [x < z < y \wedge \alpha_1(z)]$$

be the formula stating that, above x , the elements satisfying α_1 are not bounded from below. We claim that

$$\begin{aligned} & \neg \langle\langle \text{true}[\text{true}] \text{true}[\alpha_1] \dots [\alpha_{n-2}] \text{true}[\text{true}] \text{true} \rangle\rangle(x_o, x_i) \\ & \equiv \langle\langle \text{true}[\text{true}] \neg \alpha_1 [\text{true}] \text{true} \rangle\rangle(x_o, x_i) \\ & \vee [\vartheta(x_o) \wedge \psi(x_o, x_i)] \\ & \vee \exists y [x_o < y < x_i \wedge \langle\langle \text{true}[\text{true}] \neg \alpha_1 [\alpha_1 \vee \vartheta] \text{true} \rangle\rangle(x_o, y) \\ & \quad \wedge \psi(y, x_i)]. \end{aligned}$$

Before proving this equivalence, let us check that the above formula can be written as a standard temporal formula. The first term of the above disjunction is already a basic temporal formula. For the second term, note that ϑ is equivalent to the BLTL-formula $\neg(\neg \alpha_1 \cup \text{true})$ and that

$$\vartheta(x_o) \equiv \langle\langle \text{true}[\vartheta] \text{true} \rangle\rangle(x_o).$$

Furthermore, it follows by inductive hypothesis that ψ is equivalent to a standard temporal formula, and we have shown in Lemma 8.8 that standard temporal formulae are closed under conjunctions.

Finally, concerning the third term note that

$$x_o < y < x_i \equiv \langle\langle \text{true}[\text{true}] \text{true}[\text{true}] \text{true}[\text{true}] \text{true} \rangle\rangle(x_o, y, x_i),$$

and that we have shown in Lemma 8.8 that standard temporal formulae are closed under conjunctions and existential quantifications.

It remains to verify that the above formula is correct. We distinguish several cases. If there is no element $y \in (x_o, x_i)$ satisfying α_1 , the formula φ is false and the first term of the above disjunction is true. Hence, suppose otherwise. Then the infimum

$$y := \inf(\alpha_1^{\mathfrak{A}} \cap (x_o, x_i))$$

exists. It follows that the formula φ holds if, and only if,

$$\langle\langle \text{true}[\text{true}] \text{true}[\alpha_2] \dots [\alpha_{n-2}] \text{true}[\text{true}] \text{true} \rangle\rangle(y, x_i)$$

is true. If $y = x_o$, this is equivalent to the second term of the above disjunction, otherwise to the third one. \square

Lemma 8.10. *Let $\varphi(x_o, x_i)$ be a basic temporal formula of the form*

$$\varphi(x_o, x_i) = \langle\langle \text{true}[\alpha_o] \beta_1 \dots \beta_{n-1} [\alpha_{n-1}] \text{true} \rangle\rangle_{\sigma}(x_o, x_i)$$

where $n \geq 2$, $\sigma(o) = o$ and $\sigma(i) = n-1$. Over the class of Dedekind-complete linear orders, the negations of the following two formulae are equivalent to a standard temporal formula.

- (a) $\psi(x_o, x_i) := \exists x' [x_o < x' < x_i \wedge \varphi(x_o, x')]$
- (b) $\psi'(x_o, x_i) := \exists x' [x_o < x' < x_i \wedge \varphi(x', x_i)]$

Proof. (a) We consider the formulae

$$\vartheta_{n-1} := \alpha_{n-1} \quad \text{and} \quad \vartheta_i := \alpha_i \wedge (\beta_{i+1} \cup \vartheta_{i+1}), \quad \text{for } i < n-1.$$

We claim that $\psi(x_o, x_i)$ holds if, and only if, there exist elements $x_o = z_o < \dots < z_{n-1} < x_i$ such that each z_i satisfies the formula ϑ_i . Then

$$\neg \psi(x_o, x_i) \equiv \neg \vartheta(x_o) \vee \langle\langle \text{true}[\vartheta_1] \text{true} \dots \text{true}[\vartheta_{n-1}] \text{true} \rangle\rangle(x_o, x_i).$$

By Lemma 8.9, the latter formula is equivalent to a standard temporal formula. Hence, it remains to prove the claim.

(\Rightarrow) Fix an element $x' \in (x_o, x_i)$ satisfying

$$\langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x_o, x').$$

Then there are elements $x_o = z_o < \dots < z_{n-1} = x'$ such that z_i satisfies α_i and every element in (z_i, z_{i+1}) satisfies β_{i+1} . By induction on i (starting with $i = n-1$), it therefore follows that z_i satisfies ϑ_i .

(\Leftarrow) We proceed by induction on n . For $n = 2$, suppose that there are elements $x_o = z_o < z_1 < x_1$ such that z_o satisfies α_1 and z_1 satisfies $\alpha_o \wedge (\beta_1 \cup \alpha_1)$. Then the formula

$$\langle\langle \text{true}[\alpha_o]\beta_1[\alpha_2]\text{true} \rangle\rangle_\sigma(x_o, z_1)$$

holds. Hence, we can set $x' := z_1$.

For the inductive step, suppose that there are elements $x_o = z_o < \dots < z_{n-1} < x_1$ as above. By inductive hypothesis, we can find an element $y \in (x_o, z_{n-1})$ satisfying the formula

$$\langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_{n-3}[\alpha_{n-3}]\beta_{n-2}[\vartheta_{n-2}]\text{true} \rangle\rangle_\sigma(x_o, y).$$

It follows that y satisfies the formula $\vartheta_{n-2} = \alpha_{n-2} \wedge (\beta_{n-1} \cup \alpha_{n-1})$. Hence, there is some $y' > y$ satisfying α_{n-1} such that every element in the interval (y, y') satisfies β_{n-1} . If $y' < x_1$, we can take $x' := y'$. Hence, suppose otherwise. Then $z_{n-1} \in (y, y')$ implies that β_{n-1} holds in the interval (y, z_{n-1}) . Furthermore, we know by assumption that z_{n-1} satisfies $\vartheta_{n-1} = \alpha_{n-1}$. Consequently, we can set $x' := z_{n-1}$.

(b) follows as in (a) after reversing the order. \square

Lemma 8.II. Let $\varphi(x_o, x_1)$ be a basic temporal formula of the form

$$\varphi(x_o, x_1) = \langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x_o, x_1)$$

where $n \geq 2$, $\sigma(o) = o$, and $\sigma(1) = n-1$. Over the class of Dedekind-complete linear orders, the negation $\neg\varphi$ is equivalent to a standard temporal formula.

Proof. We distinguish four different cases, in each of which we will produce a standard temporal formula that is equivalent to $\neg\varphi$ if we are in that case,

and that is false if we are not. The disjunction of these four formulae will then be the desired formula equivalent to $\neg\varphi$.

(I) Suppose that the interval (x_o, x_1) is empty, that is, $x_1 \leq x_o$. Then $\neg\varphi$ is trivially true. Hence, we can use the formula

$$\langle\langle \text{true}[\text{true}]\text{true} \rangle\rangle_\sigma(x_o, x_1) \vee \langle\langle \text{true}[\text{true}]\text{true}[\text{true}]\text{true} \rangle\rangle_\sigma(x_1, x_o).$$

(II) x_o does not satisfy α_o or the elements satisfying $\neg\beta_1$ are unbounded above x_o . Then $\neg\varphi$ is again trivially true and we can use the formula

$$\langle\langle \text{true}[\neg\alpha_o \vee \neg(\beta_1 \cup \text{true})]\text{true} \rangle\rangle_\sigma(x_o).$$

(III) x_o satisfies α_o and every element in (x_o, x_1) satisfies β_1 . Then $\neg\varphi$ is equivalent to the formula

$$\begin{aligned} &\langle\langle \text{true}[\alpha_o]\beta_1[\text{true}]\text{true} \rangle\rangle_\sigma(x_o, x_1) \\ &\wedge \neg\exists x' [x_o < x' < x_1 \wedge \langle\langle \text{true}[\alpha_1]\beta_2 \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x_o, x_1)]. \end{aligned}$$

By Lemma 8.I0 the second part of this formula is equivalent to a disjunction of basic temporal formulae. Furthermore, we have shown in Lemma 8.8 that such disjunctions are closed under conjunction.

(IV) We are not in case (I) and there is some $y \in (x_o, x_1)$ that does not satisfy β_1 . The fact that we are not in case (I) can be expressed by the formula

$$\chi(x_o) := \langle\langle \text{true}[\neg\alpha_o \vee \neg(\beta_1 \cup \text{true})]\text{true} \rangle\rangle_\sigma(x_o).$$

from (II). Furthermore, the infimum

$$z' := \inf(-\beta_1^{\text{ql}} \cap (x_o, x_1))$$

can be defined by the formula

$$\begin{aligned} &\vartheta(x_o, z', x_1) := \\ &\langle\langle \text{true}[\text{true}]\beta_1[\neg\beta_1 \vee \neg(\beta_1 \cup \text{true})]\text{true}[\text{true}]\text{true} \rangle\rangle_\sigma(x_o, z', x_1). \end{aligned}$$

Consequently, the fact that we are in case (iv) can be expressed by

$$\psi(x_o, x_1) := \chi(x_o) \wedge \exists z' [x_o < z' < x_1 \wedge \vartheta(x_o, z', x_1)],$$

which is equivalent to the basic temporal formula

$$\langle\langle \text{true}[\chi]\beta_1[\neg\beta_1 \vee \neg(\beta_1 \cup \text{true})]\text{true}[\text{true}]\text{true} \rangle\rangle(x_o, x_1).$$

Consequently, we have to prove that the formula

$$\psi(x_o, x_1) \wedge \neg\varphi(x_o, x_1)$$

is equivalent to a standard temporal formula. We do so by induction on n .

If $n = 2$, then $\varphi(x_o, x_1)$ does not hold if at least one of the following conditions is satisfied.

- ◆ x_o does not satisfy α_o .
- ◆ x_1 does not satisfy α_1 .
- ◆ There is some $z \in (x_o, x_1)$ that does not satisfy β_1 .

Consequently,

$$\begin{aligned} \neg\varphi(x_o, x_1) &\equiv \langle\langle \text{true}[\neg\alpha_o]\text{true} \rangle\rangle(x_o) \\ &\vee \langle\langle \text{true}[\neg\alpha_1]\text{true} \rangle\rangle(x_1) \\ &\vee \exists x' \langle\langle \text{true}[\alpha_o]\text{true}[\neg\beta_1]\text{true}[\alpha_1]\text{true} \rangle\rangle(x_o, x', x_1). \end{aligned}$$

For the inductive step, suppose that $n > 2$. Note that, if

$$\varphi(x_o, x_1) = \langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x_o, x_1)$$

holds, then every element $x' \in (x_o, x_1)$ must satisfy one of the formulae α_i or β_j with $1 \leq i \leq n-2$ and $1 \leq j \leq n-1$. Consequently, we have

$$\varphi(x_o, x_1) \equiv (\forall x'. x_o < x' < x_1) \left[\bigvee_{1 \leq i < n-1} \xi_i \vee \bigvee_{1 \leq i < n} \zeta_i \right],$$

where

$$\begin{aligned} \xi_i(x_o, x', x_1) &:= \xi_i^-(x_o, x') \wedge \xi_i^+(x', x_1), \\ \zeta_i(x_o, x', x_1) &:= \zeta_i^-(x_o, x') \wedge \zeta_i^+(x', x_1), \\ \xi_i^-(x_o, x') &:= \langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_i[\alpha_i]\text{true} \rangle\rangle_\sigma(x_o, x'), \\ \xi_i^+(x', x_1) &:= \langle\langle \text{true}[\alpha_i]\beta_{i+1} \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x', x_1), \\ \zeta_i^-(x_o, x') &:= \langle\langle \text{true}[\alpha_o]\beta_1 \dots \beta_{i-1}[\alpha_{i-1}]\beta_i[\beta_i]\text{true} \rangle\rangle_\sigma(x_o, x'), \\ \zeta_i^+(x', x_1) &:= \langle\langle \text{true}[\beta_i]\beta_i[\alpha_i]\beta_{i+1} \dots \beta_{n-1}[\alpha_{n-1}]\text{true} \rangle\rangle_\sigma(x', x_1). \end{aligned}$$

If $(x_o, x_1) \neq \emptyset$, we also obtain the dual equivalence

$$\varphi(x_o, x_1) \equiv (\exists x'. x_o < x' < x_1) \left[\bigvee_{1 \leq i < n-1} \xi_i \vee \bigvee_{1 \leq i < n} \zeta_i \right].$$

It follows that the formula $\psi \wedge \neg\varphi$ is equivalent to

$$(\exists x'. x_o < x' < x_1) \left[\vartheta(x_o, x', x_1) \wedge \bigwedge_{1 \leq i < n-1} \neg\xi_i \wedge \bigwedge_{1 \leq i < n} \neg\zeta_i \right].$$

It therefore remains to prove that all formulae appearing above can be written as standard temporal formulae. Then the claim follows by closure under conjunction and existential quantification.

First, note that we can use the inductive hypothesis to translate the formulae $\neg\xi_i$ and $\neg\zeta_j$, for $1 \leq i < n-1$ and $1 < j < n-1$, to standard temporal formulae. The same is true for the formulae $\neg\zeta_1^-$ and $\neg\zeta_{n-1}^+$. Finally, we have

$$\begin{aligned} \vartheta \wedge \neg\zeta_1^+ &\equiv \vartheta, \\ \vartheta \wedge \neg\zeta_{n-1}^- &\equiv \vartheta \wedge \neg\zeta_{n-1}^- \wedge \forall u [x_o < u < x' \rightarrow \beta_1(u)], \end{aligned}$$

by definition of ϑ and choice of x' . We have already seen above that ϑ is (equivalent to) a standard temporal formula. Furthermore, the formula $\neg\zeta_{n-1}^- \wedge \forall u [x_o < u < x' \rightarrow \beta_1(u)]$ can be translated into a standard temporal formula using the construction from case (iii). Consequently, the claim follows by the closure properties of standard temporal formulae established in Lemma 8.8. \square

Finally, we can prove the general case.

Lemma 8.12. *Over the class of coloured Dedekind-complete linear orders, standard temporal formulae are closed under negation.*

Proof. Let $\varphi = \varphi_o \vee \dots \vee \varphi_{m-1}$ be a disjunction of basic temporal formulae φ_i . By Lemmas 8.7 and 8.8, we may assume that each φ_i has at most two free variables. Since

$$\neg\varphi = \neg\varphi_o \wedge \dots \wedge \neg\varphi_{m-1}$$

and we have shown in Lemma 8.8 that standard temporal formulae are closed under conjunctions, it is therefore sufficient to prove that the negation of a basic temporal formula

$$\varphi(\tilde{x}) = \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma}(\tilde{x})$$

with at most two free variables is equivalent to a standard temporal formula. We distinguish several cases.

(i) First suppose that φ has exactly only one free variable x_o . Then it follows by Lemma 8.6 that $\varphi(x_o)$ is equivalent to some BLTL-formula ψ . Hence,

$$\neg\varphi(x_o) \equiv \langle\langle \text{true}[\neg\psi]\text{true} \rangle\rangle(x_o).$$

(ii) Next, suppose that φ has no free variables. Then we can again use Lemma 8.6 to translate φ to some BLTL-formula ψ . Since ψ does not depend on the position it is evaluated at, it follows that $\varphi \equiv \forall x\psi(x)$ and, therefore,

$$\neg\varphi \equiv \exists x \langle\langle \text{true}[\neg\psi]\text{true} \rangle\rangle(x),$$

(iii) Suppose that φ has two free variables x_o, x_1 and that $\sigma(o) = \sigma(1)$. Then

$$\varphi(x_o, x_1) \equiv x_o = x_1 \wedge \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma'}(x_o),$$

where $\sigma'(o) := \sigma(o)$. Hence,

$$\begin{aligned} \neg\varphi(x_o, x_1) &\equiv \langle\langle \text{true}[\text{true}]\text{true}[\text{true}]\text{true} \rangle\rangle(x_o, x_1) \\ &\vee \langle\langle \text{true}[\text{true}]\text{true}[\text{true}]\text{true} \rangle\rangle(x_1, x_o) \\ &\vee \neg \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle_{\sigma'}(x_o), \end{aligned}$$

and the last formula above is equivalent to a standard temporal formula as in case (i).

(iv) Finally, suppose that φ has two free variables x_o, x_1 and that $\sigma(o) \neq \sigma(1)$. Renaming the variables if necessary, we may assume that $\sigma(o) < \sigma(1)$. Setting $k := \sigma(o)$ and $l := \sigma(1)$, it follows that

$$\begin{aligned} \varphi(x_o, x_1) &\equiv \langle\langle \beta_o[\alpha_o]\beta_1 \dots \beta_k[\alpha_k]\text{true} \rangle\rangle(x_o) \\ &\wedge \langle\langle \text{true}[\alpha_k]\beta_{k+1} \dots \beta_l[\alpha_l]\text{true} \rangle\rangle(x_o, x_1) \\ &\wedge \langle\langle \text{true}[\alpha_l]\beta_1 \dots \beta_{n-1}[\alpha_{n-1}]\beta_n \rangle\rangle(x_o). \end{aligned}$$

The negations of the first and the last term of the above disjunction can be computed as in case (i) above, while the negation of the second term can be obtained from Lemma 8.11. \square

After these preparations, we can finally conclude the proof of Kamp's Theorem.

Proof of Theorem 8.3. First let us note that, over the class of all coloured Dedekind-complete linear orders, every FO-formula $\varphi(x)$ can be translated to a standard temporal formula: this is obvious for atomic formulae, while the inductive step follows by Lemmas 8.8 and 8.12. Furthermore, since $\varphi(x)$ has one free variable, we can use Lemma 8.6 to translate every basic temporal formula in the resulting disjunction to a BLTL-formula. \square

Exercise 8.1. The above proof is very similar in structure to the original proof of the Theorem of Büchi. We consider languages $K \subseteq \Sigma^\omega$ over some alphabet Σ . Let us call such a language *basic ω -regular* if it is of the form

$$K = \bigcup_{i < m} U_i V_i^\omega, \quad \text{for regular languages } U_i, V_i \subseteq \Sigma^+$$

(compare with Exercise III.3.2). As usual we say that an MSO-formula $\varphi(\vec{X})$ with free variables \vec{X} defines the language

$$\{ w \in (\Sigma \times \wp(\bar{X}))^\omega \mid w_0 \models \varphi(\bar{P}) \},$$

where w_0 is the projection of w to the alphabet Σ and P_i is the set of all positions n such that $w(n) = \langle c, s \rangle$ with $X_i \in s$.

Show that every MSO-definable language $K \subseteq (\Sigma \times \wp(\bar{X}))^\omega$ is basic ω -regular by proving the following statements.

- (a) Every language defined by an atomic formula is basic ω -regular.
- (b) The class of basic ω -regular languages is closed under finite unions.
- (c) The class of basic ω -regular languages is closed under complement (see Exercise III.3.3).
- (d) The class of basic ω -regular languages is closed under projections.

Notes

A comprehensive account on the theory of linear orders from a logician's perspective is given by [184].

Theorem 2.2 is from [193] and Theorem 2.5 is implicit in [133]. Theorem 4.4 is based on a result in [93].

Regular linear orders were introduced in [130]. Our exposition is based on [129, 193, 15, 43, 26].

The sections on choice and uniformisation are based on [97, 133, 134, 39, 40]. Proposition 7.16 is from [36]. Strong MSO-choice for thin trees has been investigated in [13, 196].

Our proof of Kamp's Theorem follows the exposition in [158]. Kamp's Theorem can be generalised to the class of all linear orders by adding two new modal operators, see [159] for details.

The undecidability of $\text{Th}_{\text{MSO}}(\langle \mathbb{R}, \leq \rangle)$ is due to [193]. The original proof assumed a weak form of the continuum hypothesis. A proof in ZFC was provided by [96]. The material on modest linear orders is from [93, 95]. Note that there are several problems with these papers: the proof of Claim 3.3

in [96] (Claim (11) in the proof of Theorem 5.32 above) is wrong and that of Lemma 3.1 in [95] (Lemma 5.49 above) has a gap that I was not able to close. The corrected version in Section 5 above is based on ideas of Sven Manthe (personal communication).

One can show that, for each ordinal $\alpha < \omega_2$, the MSO-theory of $\langle \alpha, \leq \rangle$ is decidable [193], while the decidability of $\text{Th}_{\text{MSO}}(\langle \omega_2, \leq \rangle)$ is independent of ZFC [94].

VII Sparse Structures

1 Spanning Forests

THE LARGER EXPRESSIVE POWER of GSO over MSO stems from the fact that GSO can quantify over sets of hyperedges (guarded tuples). Thus, the more hyperedges a structure has, the higher the expressive power. Below we will make this intuition precise by showing that, if the number of hyperedges is linear in the size of the structure, we can replace every quantification over hyperedges by a set quantifier and the expressive power of GSO collapses to that of MSO.

As an example, let us consider the case of undirected trees. Every tree \mathfrak{T} can be oriented by fixing some vertex $v \in T$ as the *root* and orienting every edge such that it points away from v . Having chosen such an orientation, we can represent an edge by the vertex it points to. In this way, we can encode every set of edges by a corresponding set of vertices and every GSO-quantifier can be replaced by an MSO-quantifier. The same idea works in a much more general setting. For simplicity, we will work with hypergraphs in this section instead of relational structures.

Definition 1.1. (a) A *hypergraph* is a triple $\mathfrak{H} = \langle V, E, \text{in} \rangle$ consisting of a set V of *vertices*, a set E of *hyperedges* (or simply *edges*), and an *incidence relation* $\text{in} \subseteq V \times E$. Usually we will identify a hyperedge $e \in E$ with the set

$$\{ v \in V \mid \langle v, e \rangle \in \text{in} \}$$

of its vertices and we will write $v \in e$ instead of $\langle v, e \rangle \in \text{in}$. (But note that there might be several edges with the same underlying set.) To simplify notation further, we will also usually denote \mathfrak{H} by the pair $\langle V, E \rangle$ omitting the relation in .

(b) A hypergraph $\mathfrak{H} = \langle V, E \rangle$ has *rank* at most m if

$$|e| \leq m, \quad \text{for all } e \in E.$$

(c) We can encode a hypergraph $\mathfrak{H} = \langle V, E \rangle$ as a structure in two ways. For the *incidence representation* $\mathfrak{H}_{\text{in}} = \langle V, E, \text{in} \rangle$ we consider \mathfrak{H} as a two-sorted structure with domains V and E and a binary relation $\text{in} \subseteq V \times E$.

If every hyperedge has only finitely many vertices, we can also use the *adjacency representation* $\mathfrak{H}_{\text{adj}} = \langle V, (E_m)_{m < \omega} \rangle$ whose elements are only the vertices and we add adjacency relations

$$E_m := \left\{ \langle v_0, \dots, v_{m-1} \rangle \mid \text{there is some } e \in E \text{ with } |e| = m \text{ and } e = \{v_0, \dots, v_{m-1}\} \right\},$$

for every $m < \omega$. ,

Most notions from graph theory can straightforwardly be generalised to hypergraphs. Below we will need the following ones.

Definition 1.2. Let $\mathfrak{H} = \langle V, E \rangle$ be a hypergraph.

(a) $\mathfrak{G} = \langle U, F \rangle$ is a *subhypergraph* of \mathfrak{H} if $U \subseteq V$ and $F \subseteq E$. We denote this fact by $\mathfrak{G} \subseteq \mathfrak{H}$.

(b) A *path* is a finite sequence $(e_i)_i$ of edges such that

$$e_i \cap e_j \neq \emptyset \quad \text{iff} \quad |i - j| \leq 1, \quad \text{for all } i, j.$$

(c) A subset $C \subseteq E$ is *connected* if, for all $e, f \in C$, there exists a path p between e and f with $p \subseteq C$. Similarly, a set $C \subseteq V$ is *connected* if, for all $u, v \in C$, there exists a path p between u and v with $\bigcup p \subseteq C$. A *connected component* of \mathfrak{H} is a maximal connected set of vertices. ,

We also need a slightly different notion of a subhypergraph where we are allowed to remove vertices from edges.

Definition 1.3. Let $\mathfrak{H} = \langle V, E \rangle$ be a hypergraph

(a) For $U \subseteq V$, we denote by $\mathfrak{H}|_U$ the hypergraph with vertices U and edges

$$E|_U := \{ e \cap U \mid e \in E, e \cap U \neq \emptyset \}.$$

(b) A *weak subhypergraph* of \mathfrak{H} is a hypergraph \mathfrak{G} of the form $\mathfrak{G} = \mathfrak{H}_0|_U$, for some U and some $\mathfrak{H}_0 \subseteq \mathfrak{H}$. We will usually represent such a subhypergraph by the pair $\langle U, F \rangle$ where $F \subseteq E$ is the set of edges of \mathfrak{H}_0 .

(c) Let $\mathfrak{G} = \langle U, F \rangle$ be a weak subhypergraph of \mathfrak{H} . A *border hyperedge* of \mathfrak{G} is an edge $e \in F$ such that $e \setminus U \neq \emptyset$. The other edges are called *internal*.

(d) Two weak subhypergraphs $\mathfrak{F} = \langle U, F \rangle$ and $\mathfrak{G} = \langle W, G \rangle$ of \mathfrak{H} are *disjoint* if $U \cap W = \emptyset$ and $F \cap G = \emptyset$. ,

The translation of GSO into MSO depends on a result of independent interest which we will prove in this section: every hypergraph where the size of the hyperedges is bounded can be ‘oriented’ by a GSO-formula, i.e., there exists an GSO-formula (with parameters) that defines a linear order on each hyperedge.

The construction starts by choosing a suitable spanning forest of the given hypergraph. One technicality we have to deal with is the fact that, for uncountable hypergraphs, such forests only exist if we allow them to have arbitrary *ordinal* height.


Definition 1.4. A *forest of ordinal height* is a partial order $\mathfrak{F} = \langle F, \leq \rangle$ such that, for every vertex $v \in F$, the set of all vertices $u \leq v$ forms a well-order. The *height* of \mathfrak{F} is the least ordinal α such that, for every $v \in F$, the order type of this well-order is less than α . ,

Example. Consider the complete graph \mathfrak{K}_κ , for some uncountable cardinal κ . We can enumerate the vertices of \mathfrak{K}_κ as $(v_\alpha)_{\alpha < \kappa}$ where the index α ranges over all ordinals less than κ . Then we obtain a spanning forest $\mathfrak{F} = \langle V, \leq \rangle$ with ordering

$$v_i \leq v_j \quad : \text{iff} \quad i \leq j.$$

Thus, \mathfrak{F} consists of a single path of length κ . ,

Unfortunately, we cannot in general hope to have a spanning forest that is a *subgraph* of the given hypergraph, since the partial order \leq requires too many edges. Therefore, we will replace the full order \leq by a smaller relation $R \subseteq \leq$ that contains sufficiently many edges to reconstruct \leq from it. The idea is as follows. For a vertex v with an immediate predecessor u , we only keep the edge $u \rightarrow v$. For a vertex v without such a predecessor, we need to know the infinite path it is the limit of. We choose an increasing sequence $(u_i)_{i < \alpha}$ of vertices with limit v and, for every $i < \alpha$, a path connecting u_i to v . These paths can be combined into a tree with root v whose leaves are given by the vertices u_i .



Hence, every vertex v of \mathfrak{F} is attached to its predecessors via some auxiliary graph U_v that is either a single edge or a tree with root v whose leaves form an increasing sequence of predecessors of v with limit v .

Example. Considering again a complete graph \mathfrak{K}_κ , for an uncountable cardinal κ , and an enumeration of $(v_\alpha)_{\alpha < \kappa}$ the vertices. We can encode the spanning forest $\mathfrak{F} = \langle V, \leq \rangle$ defined above by the relation


$$R := \{ \langle v_i, v_{i+1} \rangle \mid i < \kappa \} \cup \{ \langle v_i, v_\delta \rangle \mid \delta \text{ a limit ordinal and } i < \delta \}.$$

The first part of R consists of the successor edges, whereas the second part contains the auxiliary graphs U_{v_δ} attaching a limit vertex v_δ to its predecessors.

Pseudo-Trees

Unfortunately, for hypergraphs this construction is complicated by the fact that, in general, a set of paths cannot be combined into a tree, since there

might be ‘accidental’ intersections between the hyperedges. A typical example is the hypergraph



What we will use instead are hypergraphs that are sufficiently tree-like for our purposes. These come in 4 different kinds.

Definition 1.5. Let $\mathfrak{T} = \langle T, F, v \rangle$ be a hypergraph where $L \subseteq F$ and $v \in T$.

- (a) \mathfrak{T} is a *sunflower* if $L = F$ and $v \in \bigcap F = e \cap f$, for all distinct $e, f \in F$.
- (b) \mathfrak{T} is a *hypertree* if there is some edge $e \in F$ with $v \in e$ such that, for every hyperedge $f \in F$, there exists a unique path from e to f . With each hypertree \mathfrak{T} we associate an order on F by

$$f \leq g \quad : \text{iff} \quad \text{the unique path from } e \text{ to } g \text{ contains } f.$$

We require that L is the set of maximal edges with respect to this ordering and that, for every $f \in F$, there is some $g \in L$ with $g \geq f$.

- (c) \mathfrak{T} is a *star* if the set F can be partitioned into a sunflower $S \subseteq F$ with root v and a family $(p_s)_{s \in S}$ of finite paths such that

- ♦ the first edge of p_s intersects $s \in S$,
- ♦ no edge of p_s intersects any edge of p_t , for $s \neq t$,
- ♦ if $s \neq t$, then s does not intersect any edge of p_t or t does not intersect any edge of p_s ,
- ♦ L is the set of last edges of the paths p_s .

The paths p_s are the *rays* of \mathfrak{T} and S is its *core*.

(d) \mathfrak{T} is a *pseudo-tree* if it is (i) a sunflower, (ii) a star, (iii) a hypertree, or (iv) an infinite path where $L = \emptyset$ and v belongs to the first edge. The vertex v is called the *root* of \mathfrak{T} and L is its set of *leaf-edges*.

Proposition 1.6. Let $\mathfrak{H} = \langle V, E \rangle$ be a connected hypergraph of rank at most m and let $B \subseteq E$ be an infinite set of hyperedges that is equipped with some well-order \sqsubseteq . Then there exists a pseudo-tree $\mathfrak{T} = \langle T, F, L, v \rangle$ such that $\langle T, F \rangle$ is a subhypergraph of \mathfrak{H} and L a cofinal subset of B .

Proof. Replacing B by a suitable subset, we may assume that $\kappa := |B|$ is a regular cardinal. It follows that a subset $L \subseteq B$ is cofinal if, and only if, $|L| = \kappa$.

We start by constructing an order-tree $\mathfrak{S} = \langle S, \leq \rangle$ with vertices $S \subseteq E$ as follows. Pick some edge $e_o \in E$ and, for every $e \in B$, fix a path p_e connecting e_o with e . We construct \mathfrak{S} by induction on $e \in B$ such that every $f \in S$ belongs to some p_e . Suppose that we have already constructed a priority tree \mathfrak{S}_e containing all edges $e' \sqsubset e$. Let $q_e = h_o \dots h_n$ be the shortest suffix of the path p_e that meets S . Let $e' \sqsubset e$ be the \sqsubset -minimal edge such that the path $q_{e'}$ contains some edge f' with $f' \cap h_o \neq \emptyset$, and let f be the \leq -minimal such edge belonging to $q_{e'}$. We add the path q_e to the tree by making h_o a successor of f . Let \mathfrak{S} be the tree constructed in this way. To define the desired priority tree \mathfrak{T} , we distinguish two cases.

(i) First, suppose that there is some $g \in S$ with κ different successors. Let U be the set of these successors. By construction of \mathfrak{S} , every $f \in U$ belongs to some path $q_{e(f)}$. Let r_f be the suffix of $q_{e(f)}$ starting at f . By the Pigeon Hole Principle, we can find a subset $W \subseteq U$ such that

- ♦ either all paths r_f with $f \in W$ consist of a single edge, or all such paths consist of at least two edges,
- ♦ either the first edge of every path r_f belongs to B , or none of these edges does, and
- ♦ $f \cap g = f' \cap g$, for all $f, f' \in W$.

Let \mathfrak{T} be the union of all the paths r_f with $f \in W$. For the root vertex we pick an arbitrary vertex $v \in f \cap g$. If the paths r_f have length 1, \mathfrak{T} forms a sunflower. Otherwise, it is a star.

(ii) It remains to consider the case where every $g \in S$ has less than κ successors. We start by showing that this implies that $\kappa = \aleph_o$. For $n < \omega$, let S_n be the set of all $g \in S$ whose distance from the root e_o is exactly n . By induction on n , it follows that $|S_n| < \kappa$ since

$$|S_{n+1}| = \sum_{g \in S_n} \lambda_g, \quad \text{where } \lambda_g < \kappa \text{ is the number of successors of } g,$$

which is a sum of $|S_n| < \kappa$ cardinals of size $\lambda_g < \kappa$. As κ is regular, it follows that $|S_{n+1}| < \kappa$. If $\kappa > \aleph_o$, this would imply that

$$|S| = \sum_{n < \omega} |S_n| < \kappa.$$

A contradiction to the fact that $B \subseteq S$.

It follows that \mathfrak{S} is a finitely branching infinite tree. By the Lemma of Kőnig, it has an infinite branch β . If we can choose β such that it contains infinitely many edges from B , this branch is the desired pseudo-tree. Hence, suppose otherwise. By construction of \mathfrak{S} , each edge f of β belongs to some path $q_{e(f)}$. Let $f_o < f_1 < \dots$ be an enumeration of β and set $e_i := e(f_i)$. By construction of \mathfrak{S} , we have $e_i \sqsubseteq e_{i+1}$, for all i . Furthermore, the only intersections between edges of q_{e_i} and q_{e_j} , for $i < j$ with $e_i \neq e_j$, are between the first edge of q_{e_j} and some edge of q_{e_i} with $j = i + 1$. Let e'_o, e'_1, \dots be an enumeration of e_o, e_1, \dots without repetitions and let r_i be the prefix of $q_{e'_i}$ whose last edge is the first hyperedge of $q_{e'_i}$ belonging to B . We construct the hypergraph \mathfrak{T} by taking all edges of β together with the edges of r_{2i} , for $i < \omega$. By the above remarks it follows that two hyperedges of \mathfrak{T} intersect if, and only if, one is the immediate successor of the other one. Hence, \mathfrak{T} forms a hypertree and every leaf of \mathfrak{T} belongs to B . \square

For our applications below, we will consider a family of pseudo-trees embedded as weak subhypergraphs as in the preceding proposition, and we will have to be able to reconstruct each of them by a GSO-formula. This can be done using the following two lemmas.

The first one states that the vertices of a pseudo-trees can be linearly ordered in GSO.

Lemma 1.7. Let $m < \omega$ and let $\vartheta(x, y; \bar{Z})$ be an MSO-formula. There exists an MSO-formula $\varphi(x, y; \bar{Z}, \bar{Z}')$ with the following property. Given

- a hypergraph $\mathfrak{H} = \langle V, E \rangle$ of rank at most m ,
- a weak subhypergraph $\mathfrak{T} = \langle T, F, L, v \rangle$ of \mathfrak{H} that forms a pseudo-tree,
- monadic parameters \bar{P} such that $\vartheta(x, y; \bar{P})$ defines in \mathfrak{H}_{in} a well-order on the leaf-edges L of \mathfrak{T} , and

there are monadic parameters \bar{P}' with $P'_i \subseteq T \cup F$ such that $\varphi(x, y; \bar{P}, \bar{P}')$ defines in \mathfrak{H}_{in} a linear order on T .

Proof. Below we will construct an MSO-definable partial order \sqsubseteq on F such that, for every vertex $u \in T$, the set $\{e \in F \mid u \in e\}$ is well-ordered by \sqsubseteq . Then the relation

$$u \leq_o w \quad : \text{iff} \quad \text{the } \sqsubseteq\text{-least } e \in L \text{ containing } u \text{ is } \sqsubseteq\text{-smaller than the } \sqsubseteq\text{-least } f \in L \text{ containing } w$$

defines a linear preorder on T where each class has at most m elements. Hence, there is some colouring $\lambda : T \rightarrow [m]$ assigning distinct colours to elements of the same \leq_o -class. The desired linear ordering of T is now given by

$$u \leq w \quad : \text{iff} \quad \lambda(u) < \lambda(w), \quad \text{or} \quad \lambda(u) = \lambda(w) \text{ and } u \leq_o w.$$

This relation is MSO-definable using the parameters for the definition of \sqsubseteq and m additional predicates $Q_k := \lambda^{-1}(k)$.

To find this order \sqsubseteq , we distinguish four cases depending on the type of \mathfrak{T} .

- (i) If \mathfrak{T} is a sunflower, we can take the well-ordering \sqsubseteq on L defined by ϑ .
- (ii) If \mathfrak{T} is an infinite path, we choose for \sqsubseteq the well-order on F given by

$$e \sqsubseteq f \quad : \text{iff} \quad e \text{ belongs to every prefix of } \mathfrak{T} \text{ containing } f.$$

(iii) Suppose that \mathfrak{T} is a hypertree. Let \leq_{pf} be the tree-order on F associated with \mathfrak{T} , and let \leq_{lr} be the left-to-right ordering on F given by

$$e \leq_{\text{lr}} f \quad : \text{iff} \quad \begin{aligned} &\text{the } \leq_{\text{pf}}\text{-least leaf-edge } e' \text{ with } e \leq_{\text{pf}} e' \text{ is } \vartheta\text{-less} \\ &\text{or equal to the } \leq_{\text{pf}}\text{-least leaf-edge } f' \text{ with} \\ &e \leq_{\text{pf}} f'. \end{aligned}$$

We can define the lexicographic ordering on F by

$$e \leq_{\text{lex}} f \quad : \text{iff} \quad e \leq_{\text{pf}} f, \text{ or } e \text{ and } f \text{ have the same } \leq_{\text{pf}}\text{-predecessor and } e \leq_{\text{lr}} f.$$

Then $\sqsubseteq := \leq_{\text{lex}}$ has the desired properties.

(iv) Finally, suppose that \mathfrak{T} is a star. Let S be its core and $(p_s)_{s \in S}$ the family of rays. We can linearly order the edges of each path p_s such that the smaller edges are those closer to s . Let \leq_{path} be the union of the corresponding orders. Furthermore, we define the partial order

$$e \leq_o f \quad : \text{iff} \quad e \in p_i \text{ and } f \in p_j \text{ and the leaf-edge of } p_i \text{ is } \vartheta\text{-smaller than the leaf-edge of } p_j$$

on the edges of the paths p_i . Similarly, we define a well-order order on S by

$$s \leq_t t \quad : \text{iff} \quad \text{the } \leq_o\text{-least edge } e \text{ intersecting } s \text{ is } \leq_o\text{-smaller than the } \leq_o\text{-least edge } f \text{ intersecting } t.$$

Then we obtain the desired order \sqsubseteq by setting

$$e \sqsubseteq f \quad : \text{iff} \quad e \leq_t f, \quad \text{or} \quad e \leq_o f, \quad \text{or} \quad e \in S \text{ and } f \notin S. \quad \square$$

For the second lemma, which states that we can encode a family of pseudo-trees by finitely many monadic predicates, we need the following construction.

Definition 1.8. A weak subhypergraph $\mathfrak{G} = \langle U, F \rangle$ has the k -colouring property if, every function $\lambda_o : \bigcup F \setminus U \rightarrow [k]$ can be extended to a function $\lambda : \bigcup F \rightarrow [k]$ such that, for every $e \in F$,

$$\lambda(u) \neq \lambda(v), \quad \text{for all } u \in e \cap U \text{ and } v \in e \setminus U,$$

Lemma 1.9. Let \mathfrak{H} be a hypergraph of rank at most m , $\mathfrak{T} = \langle T, F \rangle$ a weak subhypergraph of \mathfrak{H} , and $B \subseteq F$ the set of its border hyperedges. If \mathfrak{T} is a hypertree or a union of disjoint paths, then the weak subhypergraph $\langle T, B \rangle$ has the $3m$ -colouring property.

Proof. Fix a map $\lambda_o : \bigcup B \setminus T \rightarrow [2m]$. Given $v \in T$, let

$$U := \{ e \in B \mid v \in e \}.$$

If \mathfrak{T} is a hypertree, we have $|U| \leq 3$. If \mathfrak{T} is a union of paths, $|U| \leq 2$. Thus, $|\bigcup U| \leq 3m$. Since $v \in T$, it follows that the set

$$C := \lambda_o[\bigcup U \setminus T]$$

has less than $3m$ elements. We pick some $c \in [3m] \setminus C$ and set $\lambda(v) := c$. \square

Lemma 1.10. *Let $m < \omega$ and let $\gamma(x, y, z; \bar{Z})$ be an MSO-formula. There exists an MSO-formula $\varphi(x, X, Y; \bar{Z}')$ with the following property. Suppose we are given a hypergraph $\mathfrak{H} = \langle V, E \rangle$ of rank at most m and a family $\mathfrak{T}_i = \langle T_i, F_i \rangle$, $i \in I$, of weak subhypergraphs of \mathfrak{H} such that*

- ♦ *each \mathfrak{T}_i is a pseudo-tree with root v_i ,*
- ♦ *\mathfrak{T}_i and \mathfrak{T}_j are disjoint, for $i \neq j$, and*
- ♦ *there exist parameters \bar{P}' such that*

$$\begin{aligned} \mathfrak{H}_{\text{in}} \models \gamma(u, w, e; \bar{P}') \quad \text{iff} \quad & \text{there is some } i \in I \text{ such that} \\ & u = v_i, \quad e \in F_i, \quad u, w \in e \cap T_i, \\ & \text{and } \mathfrak{T}_i \text{ is a star or a sunflower.} \end{aligned}$$

Then there exist monadic parameters \bar{P} such that

$$\mathfrak{H}_{\text{in}} \models \varphi(u, A, U; \bar{P}) \quad \text{iff} \quad u = v_i, \quad A = T_i, \quad \text{and} \quad U = F_i, \\ \text{for some } i \in I.$$

Proof. We can partition the index set I into four sets I_0, I_1, I_2, I_3 such that all pseudo-trees in I_t have the same type (sunflower, star, hypertree, or path). It is sufficient to construct a separate formula φ_t for each subfamily $(\mathfrak{T}_i)_{i \in I_t}$. Hence, suppose that all hypergraphs \mathfrak{T}_i have the same type. Set

$$T := \bigcup_{i \in I} T_i, \quad F := \bigcup_{i \in I} F_i, \quad R := \{ v_i \mid i \in I \}.$$

We distinguish three cases.

(i) If every \mathfrak{T}_i is a sunflower, we can use the formula

$$\begin{aligned} \varphi(x, X, Y) := & Rx \wedge \forall e [e \in Y \leftrightarrow \gamma(x, x, e; \bar{P}')] \\ & \wedge \forall z [z \in X \leftrightarrow (\exists e \in Y) \gamma(x, z, e; \bar{P}')] . \end{aligned}$$

(ii) Suppose that each \mathfrak{T}_i is a path or a hypertree. Fix a well-ordering \leq on I and let

$$B_i := \{ e \in F_i \mid e \text{ is a border edge of } \mathfrak{T}_i \} \quad \text{and} \quad B := \bigcup_{i \in I} B_i.$$

We introduce two colourings $\lambda : \bigcup F \rightarrow [3m]$ and $\mu : F \rightarrow \wp([3m])$ such that the edge of \mathfrak{T}_i represented by $e \in F_i$ is

$$e \cap \lambda^{-1}[\mu(e)].$$

We define these colourings by induction on $i \in I$. Let $W := \bigcup F \setminus T$. We start with an arbitrary colouring $\lambda : W \rightarrow [3m]$. For the inductive step, suppose that we have already defined $\lambda(v)$ and $\mu(e)$ for all $v \in W \cup \bigcup_{j < i} T_j$ and $e \in \bigcup_{j < i} F_j$. By Lemma 1.9, the subhypergraph $\langle T_i, B_i \rangle$ has the $3m$ -colouring property. We use this fact to extend λ to the set $W \cup \bigcup_{j < i} T_j \cup B_i$. Then we can extend λ to all of $W \cup \bigcup_{j < i} T_j \cup T_i$ by using an arbitrary colour for the vertices in $T_i \setminus \bigcup B_i$. Having defined λ , we can set

$$\mu(e) := \lambda[e \cap T_i], \quad \text{for } e \in F_i.$$

We can use these two colourings to construct the desired formula as follows. Using the parameters T, F, R , and

$$P_k := \lambda^{-1}(k), \quad Q_s := \mu^{-1}(s), \quad \text{for } k < m \text{ and } s \leq m,$$

we can define the relation $F|_T$ by

$$\psi(X) := (\exists e \in F) [X = e \cap \lambda^{-1}[\mu(e)]] .$$

Using this formula there exists a formula $\vartheta(Y; \bar{Z})$ such that

$$\mathfrak{H}_{\text{in}} \models \vartheta(C; \bar{P}\bar{Q}) \quad \text{iff} \quad C \subseteq F \text{ corresponds to a connected component of } \langle T, F|_T \rangle.$$

Note that the sets F_i are the connected components of the hypergraph $\langle T, F|_T \rangle$. Hence, we can set

$$\begin{aligned} \varphi(x, X, Y) &:= x \in X \cap R \wedge \vartheta(Y) \\ &\quad \wedge \forall z [z \in X \leftrightarrow (\exists e \in Y)[z \in e \wedge \lambda(z) \in \mu(e)]] . \end{aligned}$$

(III) It remains to consider the case where every \mathfrak{Z}_i is a star. We can use the formula from (I) to define the core of \mathfrak{Z}_i and the formula from (II) to define its rays. \square

Spanning Forests

With these preparations out of the way, we are finally able to define what we mean by a spanning forest. We use a forest of ordinal height where limit vertices are attached using a pseudo-tree. In addition, we require two further properties of such a forest: (I) every edge of the given hypergraph should contain some vertex of the forest and (II) if an edge contains several such vertices, these must be comparable in the forest order. (The latter roughly corresponds to a forest that can be obtained by a depth-first traversal of the hypergraph.)

Definition 1.11. Let $\mathfrak{H} = \langle V, E \rangle$ be a hypergraph of rank at most m , and let

$$\mathfrak{F} = \langle F, \leq, (U_v)_{v \in F} \rangle$$

be a structure where $F \subseteq V$ is a subset of the vertices, $\langle F, \leq \rangle$ forms a forest of ordinal height, and each $U_v \subseteq E$ is a set of edges, for $v \in F$.

(a) The set of *auxiliary nodes* associated to a vertex $v \in F$ is

$$A_v := \{v\} \cup \bigcup_{u < v} U_v \setminus \bigcup_{u < v} A_u .$$

The *attachment set* of a set $X \subseteq V$ and its *principle point of attachment* are

$$\begin{aligned} B(X/\mathfrak{F}) &:= \{v \in F \mid X \cap A_v \neq \emptyset\} , \\ \beta(X/\mathfrak{F}) &:= \max B(X/\mathfrak{F}) . \end{aligned}$$

(If $B(X/\mathfrak{F})$ does not have a unique maximal element, we let $\beta(X/\mathfrak{F})$ be undefined.)

(b) We call \mathfrak{F} an \mathfrak{H} -*forest* if it satisfies the following conditions.

(I) For all $u \neq v$,

$$U_u \cap U_v = \emptyset, \quad A_u \cap A_v = \emptyset, \quad \text{and} \quad A_u \cap F = \{u\} .$$

(II) $\bigcup U_u \cap \bigcup U_v \neq \emptyset$ implies $u \leq v$ or $v \leq u$.

(III) Each vertex $v \in F$ is assigned one of the following *types*: (R) the *roots*; $(s_l)_{l < m}$ the *successors*; and $(\tau_l)_{l < m}$ the *limits*; where the successor and limit vertices are further subdivided into m subclasses. This assignment is subject to the following conditions:

(R) If v has type R, it is a root of F and $U_v = \emptyset$.

(s_l) If v has type s_l, it is the (immediate) successor of some vertex $u \in F$ and there exists an edge $e \in E$ such that

$$v \in e, \quad U_v = \{e\}, \quad \beta(e \setminus \{v\}/\mathfrak{F}) = u ,$$

$B(e/\mathfrak{F})$ is linearly ordered by \leq , and v is the only vertex in $B(e/\mathfrak{F})$ of type s_l.

(τ_l) If v has type τ_l, it has no immediate predecessor and the weak sub-hypergraph $\mathfrak{Z}_v := \langle A_v, U_v, v \rangle$ of \mathfrak{H} forms a pseudo-tree satisfying the following conditions:

- $\{\beta(e \setminus \{v\}/\mathfrak{F}) \mid e \text{ a border edge of } U_v\}$ is a cofinal subset of $\{u \in T \mid u < v\}$,
- if \mathfrak{Z}_v is a sunflower, then v is the only vertex in $B(\bigcup U_v/\mathfrak{F})$ of type τ_l,
- if \mathfrak{Z}_v is a star, then v is the only vertex in $B(\bigcup C/\mathfrak{F})$ of type τ_l, where $C := \{e \in U_v \mid v \in e\}$ is the core of \mathfrak{Z}_v . \lrcorner

We are interested in \mathfrak{H} -forests that span the whole hypergraph and that allow us to orient every of its hyperedges.

Definition 1.12. Let $\mathfrak{H} = \langle V, E \rangle$ be a hypergraph of rank at most m and $\mathfrak{F} = \langle F, \leq, (U_v)_{v \in F} \rangle$ an \mathfrak{H} -forest.

- (a) \mathfrak{F} is *spanning* if $B(e/\mathfrak{F}) \neq \emptyset$, for every non-empty $e \in E$.
- (b) \mathfrak{F} is *depth-first* if, $B(e/\mathfrak{F})$ is linearly ordered by \leq , for every $e \in E$.

Let us start by showing that such \mathfrak{H} -forests always exist.

Theorem 1.13. Every hypergraph \mathfrak{H} has a depth-first spanning \mathfrak{H} -forest.

Proof. We construct an increasing sequence

$$\mathfrak{F}_\alpha = \langle F_\alpha, \leq, (U_v)_{v \in F_\alpha} \rangle, \quad \alpha < \kappa,$$

of depth-first \mathfrak{H} -forests with the property that, for every connected component C of $W_\alpha := V \setminus \bigcup_{v \in F_\alpha} A_v$, the set

$$N_\alpha(C) := \bigcup \{ B(e/\mathfrak{F}_\alpha) \mid e \in E \text{ with } e \cap C \neq \emptyset \}$$

is linearly ordered by \leq . The limit of this sequence will then be the desired depth-first spanning \mathfrak{H} -forest.

We start with the empty tree $F_0 := \emptyset$. For limit ordinals δ , we take the limit $\mathfrak{F}_\delta := \bigcup_{\alpha < \delta} \mathfrak{F}_\alpha$. Clearly, \mathfrak{F}_δ is a depth-first \mathfrak{H} -forest and $N_\delta(C)$ is linearly ordered for every C . For the successor step, suppose that we have already defined \mathfrak{F}_α . Fix some connected component C of W_α . We distinguish three cases.

(a) If $N_\alpha(C) = \emptyset$, we pick some vertex $v \in C$ and we add it to \mathfrak{F}_α as a new root with $U_v := \emptyset$ and type (r) . Clearly, the resulting structure $\mathfrak{F}_{\alpha+1}$ is an \mathfrak{H} -forest.

(b) If $N_\alpha(C)$ has a maximal element u , we choose some edge e with $e \cap A_u \neq \emptyset$ and $e \cap C \neq \emptyset$, and we pick some vertex $v \in e \cap C$. We add v to \mathfrak{F}_α as an immediate successor of u and we set $U_v := \{e\}$. It follows that $A_v = e \cap W_\alpha$. Since $B(e/\mathfrak{F}_\alpha)$ contains at most $|e \setminus \{v\}| < m$ vertices, there is some number $l < m$ such that $B(e/\mathfrak{F}_\alpha)$ contains no vertex of type s_l .

In the new forest $\mathfrak{F}_{\alpha+1}$, we assign the type s_l to v . To show that $\mathfrak{F}_{\alpha+1}$ is an \mathfrak{H} -forest, we have to check three conditions.

(i) As $e \cap W_\alpha \neq \emptyset$ while $W_\alpha \cap \bigcup U_u = \emptyset$, for every $u \in F_\alpha$, we have $e \not\subseteq U_u$. Hence, $U_v \cap U_u = \emptyset$. Furthermore, the fact that $B(e/\mathfrak{F}_\alpha)$ is linearly ordered implies that e only intersects sets A_u with $u \leq v$. By definition of A_v it therefore follows that $A_v \cap A_u = \emptyset$, for $u \neq v$. Finally, $A_v \cap F_{\alpha+1} = (e \setminus F_\alpha) \cap F_{\alpha+1} = \{v\}$.

(ii) Suppose that $e \cap \bigcup U_u \neq \emptyset$, for $u \neq v$. Since $B(e/\mathfrak{F}_\alpha)$ is linearly ordered and bounded by v , it follows that $u \leq v$.

(iii) holds by construction of $\mathfrak{F}_{\alpha+1}$ and by choice of the type s_l .

(c) Suppose that $N_\alpha(C)$ has no maximal element. Set $\mathfrak{C} := \mathfrak{H}|_C$ and let B be the set of border hyperedges of \mathfrak{C} . Note that B is preordered by the relation

$$e \leq f \quad : \text{iff} \quad \beta(e/\mathfrak{F}_\alpha) \leq \beta(f/\mathfrak{F}_\alpha).$$

Let $B_o \subseteq B$ be a set of representatives of the corresponding equivalence classes. Then B_o is well-ordered and we can use Proposition 1.6 to find a pseudo-tree $\mathfrak{S} = \langle S, H, L, v \rangle$ such that L is a cofinal subset of B_o . We simplify \mathfrak{S} as follows.

If \mathfrak{S} is a sunflower, we can find a subset $H_o \subseteq H$ and a number $l < m$ such that

- ♦ $N_\alpha(H_o)$ is unbounded and,
- ♦ for every $e \in H_o$, the set $B(e/\mathfrak{F}_\alpha)$ contains no vertex of type τ_l .

We replace \mathfrak{S} by the sunflower $\mathfrak{S}_o := \langle S \cap \bigcup H_o, H_o, H_o, v \rangle \subseteq \mathfrak{S}$.

Similarly, if \mathfrak{S} is a star, we can find a subhypergraph $\mathfrak{S}_o \subseteq \mathfrak{S}$ and a number $l < m$ such that \mathfrak{S}_o forms a star and, for no edge e in the core of \mathfrak{S}_o , there is a vertex of type τ_l in the set $B(e/\mathfrak{F}_\alpha)$. Again, we replace \mathfrak{S} by \mathfrak{S}_o .

Finally, if \mathfrak{S} is a path or a hypertree, we fix an arbitrary number $l < m$ and we leave \mathfrak{S} unchanged.

We define $F_{\alpha+1} := F_\alpha \cup \{v\}$ where v is the limit of $N_\alpha(C)$, we assign the type τ_l to v , and we set $U_v := H$. It follows that $A_v = S$.

It remains to show that $\mathfrak{F}_{\alpha+1}$ is an \mathfrak{H} -forest. We have to check three conditions.

(i) Since every edge $e \in U_v$ intersects W_α while $W_\alpha \cap \bigcup U_\alpha = \emptyset$, for $u \in F_\alpha$, we have $U_v \cap U_u = \emptyset$. Furthermore, the fact that $B(\bigcup U_v / \mathfrak{F}_\alpha) \subseteq N_\alpha(C)$ is linearly ordered, implies that $\bigcup U_v$ only intersects sets A_u with $u \leq v$. By definition of A_v it therefore follows that $A_v \cap A_u = \emptyset$, for $u \neq v$. Finally, $A_v \cap F_{\alpha+1} = (\bigcup U_v \setminus F_\alpha) \cap F_{\alpha+1} = \{v\}$.

(ii) Suppose that $\bigcup U_v \cap \bigcup U_u \neq \emptyset$, for $u \neq v$. Since $B(\bigcup U_v / \mathfrak{F}_\alpha) \subseteq N_\alpha(C)$ is linearly ordered and bounded by v , it follows that $u \leq v$.

(iii) We have

$$\{\beta(e/\mathfrak{F}_{\alpha+1}) \mid e \text{ a border edge of } U_v\} \subseteq N_\alpha(C) \cup \{v\}$$

which is linearly ordered. Furthermore, by choice of the type τ_I , v is the only vertex of this type in $B(\bigcup U_v / \mathfrak{F}_{\alpha+1})$. Finally, every internal edge $e \in U_v$ is contained in $C \subseteq W_\alpha$. This implies that $e \cap A_u = \emptyset$, for all such edges e and all $u \neq v$.

This concludes the construction of $\mathfrak{F}_{\alpha+1}$. It remains to show that $\mathfrak{F}_{\alpha+1}$ is of the correct form. We start by showing that $\mathfrak{F}_{\alpha+1}$ is depth-first. Let v be the new vertex added to \mathfrak{F}_α . If $e \cap A_v = \emptyset$, then $B(e/\mathfrak{F}_{\alpha+1}) = B(e/\mathfrak{F}_\alpha)$ and we are done. Otherwise, we have $B(e/\mathfrak{F}_{\alpha+1}) = B(e/\mathfrak{F}_{\alpha+1}) \cup \{v\}$. Note that $A_v \subseteq C$ implies $e \cap C \neq \emptyset$. Therefore, we have $B(e/\mathfrak{F}_\alpha) \subseteq N_\alpha(C)$. As v is larger than every element in $N_\alpha(C)$ the claim follows.

Finally, let D be a connected component of $W_{\alpha+1} = V \setminus \bigcup_{x \in F_{\alpha+1}} A_x$. We have to show that $N_{\alpha+1}(D)$ is linearly ordered. Since $W_{\alpha+1} \subseteq W_\alpha$ there is some connected component D' of W_α containing D . If $D' \neq C$, then $W_\alpha \setminus W_{\alpha+1} \subseteq C$ implies that $D = D'$ and the set

$$N_{\alpha+1}(D) = N_\alpha(D')$$

is linearly ordered. If, on the other hand, $D \subseteq C$ then we have

$$N_{\alpha+1}(D) \subseteq N_\alpha(C) \cup \{v\}$$

and the latter set is linearly ordered since v is greater than every element of $N_\alpha(C)$. \square

We will use depth-first spanning forests to encode orientations of a hypergraph. First, let us show how to encode such forest by finitely many parameters.

Lemma 1.14. *Let $m < \omega$. There exist MSO-formulae $\varphi(X; \bar{Z})$, $\psi(x, Y; \bar{Z})$, $\vartheta(x, Y; \bar{Z})$, and $\chi(x, y; \bar{Z})$ such that, for every hypergraph \mathfrak{H} of rank at most m and each depth-first spanning \mathfrak{H} -forests $\mathfrak{F} = \langle F, \leq, (U_v)_v \rangle$, there are parameters \bar{S} such that*

$$\begin{aligned} \mathfrak{H}_{\text{in}} &\models \psi(v, P; \bar{S}) && \text{iff} && v \in F \text{ and } P = U_v, \\ \mathfrak{H}_{\text{in}} &\models \vartheta(v, P; \bar{S}) && \text{iff} && v \in F \text{ and } P = A_v, \\ \mathfrak{H}_{\text{in}} &\models \varphi(P; \bar{S}) && \text{iff} && P \subseteq F \text{ is downwards closed w.r.t. } \leq, \\ \mathfrak{H}_{\text{in}} &\models \chi(u, v; \bar{S}) && \text{iff} && u, v \in F \text{ and } u \leq v. \end{aligned}$$

Proof. For every type τ , we use the following parameters:

- a unary predicate F_τ containing all vertices $v \in F$ of type τ and
- the sets $A^\tau := \bigcup \{A_v \mid v \in F_\tau\}$ and $U^\tau := \bigcup \{U_v \mid v \in F_\tau\}$.

First, we construct the formulae ψ and ϑ . To simplify our task we define separate versions $\psi_\tau(x, Y)$ and $\vartheta_\tau(x, Y)$ for each type τ such that

$$\begin{aligned} \mathfrak{H} &\models \psi_\tau(v, P) && \text{iff} && v \in F_\tau \text{ and } P = U_v, \\ \mathfrak{H} &\models \vartheta_\tau(v, P) && \text{iff} && v \in F_\tau \text{ and } P = A_v. \end{aligned}$$

Then the disjunctions of these formulae yield the desired formula ψ and ϑ .

If v has type \mathbf{r} , then $A_v = \{v\}$, $U_v = \emptyset$, and we can set

$$\begin{aligned} \psi_{\mathbf{r}}(x, Y) &:= F_{\mathbf{r}}x \wedge Y = \emptyset, \\ \vartheta_{\mathbf{r}}(x, Y) &:= F_{\mathbf{r}}x \wedge Y = \{v\}. \end{aligned}$$

If the type of v is s_I , then $A_v = e \cap A^{s_I}$ and $U_v = \{e\}$, where e is the unique edge in U^{s_I} containing v . Hence, we can define

$$\begin{aligned} \psi_{s_I}(x, Y) &:= F_{s_I}x \wedge (\exists e \in U^{s_I})[x \in e \wedge Y = \{e\}], \\ \vartheta_{s_I}(x, Y) &:= F_{s_I}x \wedge (\exists e \in U^{s_I})[x \in e \wedge Y = e \cap A^{s_I}]. \end{aligned}$$

Finally, suppose that v has type τ_l . We construct the desired formulae ψ_{τ_l} and ϑ_{τ_l} with the help of Lemma 1.10. To do so, we need a formula $\gamma(x, y, z)$ (with parameters) such that

$$\begin{aligned} \mathfrak{H}_{\text{in}} \models \gamma(u, w, e) \quad \text{iff} \quad & u = v_i, \quad e \in F_i, \quad \text{and} \quad u, w \in e \cap A_i, \\ & \text{for some } i \in I \text{ such that } \mathfrak{T}_i \text{ is a star or} \\ & \text{a sunflower.} \end{aligned}$$

This formula is given by

$$\gamma(x, y, z) := x \in F_{\tau_l} \wedge z \in S \wedge x, w \in z \cap A^{\tau_l},$$

where

$$S := \{ e \mid \mathfrak{T}_i \text{ is a sunflower or a star, } v_i \in e \in F_i, \text{ and } v_i \text{ has type } \tau_l \}.$$

(Note that, by definition of an \mathfrak{H} -forest, each edge $e \in S$ has a unique vertex of type τ_l .) Let $\varphi'(x, X, Y)$ be the formula obtained from Lemma 1.10. Then we can set

$$\psi_{\tau_l}(x, Y) := \exists X \varphi'(x, X, Y) \quad \text{and} \quad \vartheta_{\tau_l}(x, Y) := \exists X \varphi'(x, Y, X).$$

For the final two formulae note that, using ψ and ϑ , we can construct a formula $\alpha(x, Y)$ such that

$$\mathfrak{H} \models \alpha(v, Q) \quad \text{iff} \quad v \in F \text{ and } Q = B(\bigcup U_v / \mathfrak{F}).$$

With the help of α we can then set

$$\begin{aligned} \varphi(X) &:= X \subseteq F \wedge \forall x \forall Y [Xx \wedge \alpha(x, Y) \rightarrow Y \subseteq X], \\ \chi(x, y) &:= Fx \wedge Fy \wedge \forall X [\varphi(X) \rightarrow (Xy \rightarrow Xx)]. \end{aligned} \quad \square$$

Orientations

We can use depth-first spanning forests to define an orientation (i.e., a linear ordering) of every hyperedge in a given hypergraph.

Definition 1.15. Let $\mathfrak{H} = \langle V, E \rangle$ be a hypergraph. A formula $\varphi(x, y, z)$ defines an *edge ordering* of \mathfrak{H} if, for every edge $e \in E$, the formula $\varphi(x, y, e)$ defines (in \mathfrak{H}_{in}) a linear ordering on the vertices of e . \lrcorner

Theorem 1.16. Let $m < \omega$. There exists an MSO-formula $\varphi(x, y, z, \bar{U})$ such that, for every hypergraph \mathfrak{H} of rank at most m , there are parameters \bar{P} such that $\varphi(x, y, z, \bar{P})$ defines (in the structure \mathfrak{H}_{in}) an edge ordering of \mathfrak{H} .

Proof. Fix a depth-first spanning \mathfrak{H} -forest \mathfrak{F} . For each type τ , we will construct a formula $\eta_\tau(x, y)$ such that, for every vertex v of type τ , η_τ defines a linear preorder on the set A_v whose equivalence classes have size at most m . Using the formulae $\chi(x, y)$ and $\vartheta(x, Y)$ from Lemma 1.14 (we omit the parameters \bar{S} from the notation), we can set

$$\begin{aligned} \psi_o(x, y) &:= (\exists u \in F)(\exists v \in F)[x \in A_u \wedge y \in A_v \wedge u < v] \\ &\vee (\exists v \in F)[x, y \in A_v \wedge "v \text{ has type } \tau" \wedge \eta_\tau(x, y)]. \end{aligned}$$

This formula defines a preorder Ξ_o such that

- the restriction of Ξ_o to F coincides with \leq ,
- Ξ_o linearly preorders every set $X \subseteq V$ such that $B(X/\mathfrak{F})$ is linearly ordered by \leq , and
- each Ξ_o -class has at most m elements.

In particular, Ξ_o linearly preorders every hyperedge of \mathfrak{H} . Using m additional unary predicates P_0, \dots, P_{m-1} we can assign a unique colour to each element of a Ξ_o -class, and these colours can then be used to linear order each class. The resulting partial order Ξ is MSO-definable and linearly orders every hyperedge of \mathfrak{H} .

Hence, it remains to construct the formulae $\eta_\tau(x, y)$. Let F_τ be the set of all vertices $v \in F$ of type τ . If v is of type τ or s_l , the set A_v has at most m elements and we can set

$$\eta_\tau(x, y) := F_\tau x \wedge F_\tau y.$$

For vertices of type τ_l , U_v is the set of edges of a priority tree

$$\mathfrak{T} = \langle \bigcup U_v, U_v, L_v, \lambda_v, \mu_v, v \rangle.$$

Since the leafs of \mathfrak{T} are well-ordered by the (definable) order \leq of \mathfrak{F} , we can use the formula φ from Lemma 1.7 to order the set $A_v \subseteq \bigcup U_v$. To make this work, we have to deal with one technicality. Let \tilde{S}_v be the parameters used by $\varphi(x, y; \tilde{S}_v)$ to order A_v . We have to encode the parameters for the various vertices v into a single tuple \tilde{S} . Since $(S_v)_i \subseteq A_v \cup U_v$, these sets are disjoint for different v . Hence, we can use the union

$$S_i := \bigcup_v (S_v)_i$$

and the formula

$$\eta_\tau(x, y; \tilde{S}) := (\exists v \in F_\tau) [x, y \in A_v \wedge \varphi(x, y; \tilde{S} \upharpoonright A_v \cup U_v)],$$

which is definable by Lemma 1.14. \square

As an application, let us show how to order the predecessors of each vertex in a directed graphs.

Definition 1.17. A formula $\varphi(x, y; z)$ defines an *predecessor ordering* of a directed graph $\mathfrak{G} = \langle V, E \rangle$ if, for every vertex $v \in V$, the formula $\varphi(x, y; v)$ defines a linear ordering on the set $p(v) := \{u \in V \mid \langle u, v \rangle \in E\}$.

Corollary 1.18. Let $m < \omega$. There exists an MSO-formula $\varphi(x, y; z, \bar{U})$ with the following property: for every directed graph \mathfrak{G} of indegree at most m , there are MSO-parameters \bar{P} such that the formula $\varphi(x, y; z, \bar{P})$ defines (in \mathfrak{G}) a predecessor ordering of \mathfrak{G} .

Proof. Let $\mathfrak{H} := \langle V, F \rangle$ be the hypergraph where the hyperedges are of the form

$$F := \{p(v) \mid v \in V\} \quad \text{with} \quad p(v) := \{u \in V \mid \langle u, v \rangle \in E\}.$$

Note that we can interpret the incidence representation \mathfrak{H}_{in} in \mathfrak{G} by encoding each edge $p(v)$ by the corresponding vertex v . More precisely, we choose for every hyperedge e some vertex v with $p(v) = e$. Let Q be the set of

these vertices. Then we obtain an interpretation $\tau = \langle \delta_V, \delta_E, \varphi_{\text{in}} \rangle$ of \mathfrak{H}_{in} in $\langle \mathfrak{G}, Q \rangle$ by setting

$$\delta_V(v) := \text{true}, \quad \delta_E(e) := Qe, \quad \varphi_{\text{in}}(v, e) := Eve.$$

Let $\psi(x, y; z, \bar{U})$ be the MSO-formula from Theorem 1.16 defining an edge ordering of \mathfrak{H} and let \bar{P} be the corresponding parameters. Using the interpretation τ we can evaluate $\psi(x, y; z, \bar{P})$ on $\langle \mathfrak{G}, Q \rangle$. Let $\psi'(x, y; z, \bar{P}', Q)$ be the resulting formula. Since τ encodes each hyperedge $e = p(v)$ by the set of predecessors of v , this edge ordering induces a predecessor ordering on \mathfrak{G} . Consequently, we can use the formula

$$\begin{aligned} \varphi(x, y; z, \bar{P}', Q) := & \exists z' [Qz' \wedge \forall u [Euz \leftrightarrow Euz']] \\ & \wedge \psi'(x, y; z', \bar{P}', Q)]. \end{aligned}$$

to define a predecessor ordering on \mathfrak{G} . (We have to replace the given vertex z by the representative of the hyperedge $p(z)$, i.e., by the unique vertex $z' \in Q$ with $p(z) = p(z')$.) \square

2 Sparse Hypergraphs

It is finally time to define the class of structures where the translation of GSO into MSO is possible. As already mentioned above, this only works if there are not too many hyperedges since, otherwise, we have not enough vertices to encode them by.

Definition 2.1. Let $k < \omega$.

(a) A hypergraph $\mathfrak{H} = \langle V, E \rangle$ is *k-sparse* if

$$|E \upharpoonright X| \leq k \cdot |X|, \quad \text{for all finite } X \subseteq V,$$

where $E \upharpoonright X := \{e \in E \mid e \subseteq X\}$.

(b) A relational Σ -structure \mathfrak{A} is *k-sparse* if

$$|R \upharpoonright X| \leq k \cdot |X|, \quad \text{for all finite } X \subseteq A \text{ and all } R \in \Sigma.$$

Remark. These two definitions do not quite agree for adjacency representations of hypergraphs. A k -sparse hypergraph \mathfrak{H} of rank m has an adjacency representation $\mathfrak{H}_{\text{adj}}$ that is only $km!$ -sparse. \lrcorner

Examples. (a) Trees are 1-sparse.

(b) Planar graphs are 3-sparse since every planar graph with n vertices has at most $3n - 6$ edges.

(c) If an undirected graph $\mathfrak{G} = \langle V, E \rangle$ has degree at most $2k$, it is k -sparse since, for every finite subset $X \subseteq V$, we have

$$2 \cdot |E \upharpoonright X| = \sum_{v \in X} \deg(v) \leq 2k \cdot |X|.$$

(d) Cliques with n vertices are k -sparse for $k \geq \frac{n(n-1)/2}{n} = \frac{1}{2}n$, but not for smaller k .

(e) Complete bipartite graphs $\mathfrak{K}_{n,n}$ are k -sparse for $k \geq \frac{n^2}{2n} = \frac{1}{2}n$, but not for smaller k .

(f) The class of k -sparse graphs is closed under subdivisions. Furthermore, for every graph \mathfrak{G} , we can construct a 2-sparse subdivision by replacing each edge by a path of length 2.

(g) There are k -sparse Σ -structures \mathfrak{A} whose Gaifman graph is not k -sparse. For instance, let $\mathfrak{A} = \langle A, R \rangle$ by the structure with universe

$$A := C + (C \times C), \quad \text{for some set } C \text{ of size } |C| = n,$$

and a ternary relation

$$R := \{ \langle a, b, \langle a, b \rangle \rangle \mid a, b \in C \}.$$

Then \mathfrak{A} is 1-sparse but its Gaifman graph contains a clique of size n . \lrcorner

Clearly, a structure must be sparse if we want to encode every guarded tuple by a single element. Surprisingly, this necessary condition turns out to be also sufficient. Before giving the proof in the next section, we collect a few results about MSO-definable choice functions, i.e., functions selecting a unique vertex from every hyperedge.

Definition 2.2. A *choice function* for a hypergraph $\mathfrak{H} = \langle V, E \rangle$ is a function $\varepsilon : E \rightarrow V$ such that $\varepsilon(e) \in e$, for all $e \in E$. The *indegree* of a choice function ε is the cardinal

$$\sup_{v \in V} |\varepsilon^{-1}(v)|.$$

Theoretically, we could obtain such a function using the linear order of Theorem 1.16 and map each $e \in E$ to the minimal vertex $v \in e$. But below we need an MSO-definable choice function while Theorem 1.16 only produces GSO-definable ones. Furthermore, we want our choice function to be of bounded indegree and there is no guarantee that the orientation from Theorem 1.16 produces such a function. For sparse hypergraphs, there fortunately exists a direct argument (not based on orientations) proving the existence of a choice function of bounded indegree. It follows that there is an MSO-definable way of representing hyperedges by vertices such that every vertex encodes at most k hyperedges, for some constant k . This means we can encode these hyperedges by using at most k copies of v . We start with proving existence (ignoring the issue of definability).

Proposition 2.3. A hypergraph \mathfrak{H} is k -sparse if, and only if, it has a choice function ε if indegree at most k .

Proof. For (\Leftarrow) , let $X \subseteq V$ be finite. Then

$$|E \upharpoonright X| \leq \sum_{v \in X} |\varepsilon^{-1}(v)| \leq k \cdot |X|.$$

(\Rightarrow) First, let us consider the case where \mathfrak{H} is finite. If ε is an arbitrary choice function of \mathfrak{H} then

$$\sum_{v \in V} |\varepsilon^{-1}(v)| = |E| \leq k \cdot |V|.$$

Hence, if there is some element $v \in V$ with $|\varepsilon^{-1}(v)| > k$ then there must be some other element $u \in V$ with $|\varepsilon^{-1}(u)| < k$. Let us define the *weight* of ε by

$$w(\varepsilon) := \sum \{ |\varepsilon^{-1}(v)| - k \mid v \in V, |\varepsilon^{-1}(v)| > k \}.$$

We have to construct a choice function of weight 0. To do so we transform a given choice function ε with $w(\varepsilon) > 0$ into one with smaller weight. Given ε , fix an element $v \in V$ with $|\varepsilon^{-1}(v)| > k$. Let $F \subseteq E$ be the smallest subset of E such that v belongs to the set $U := \bigcup F$ and we have $\varepsilon^{-1}(u) \subseteq F$, for every $u \in U$. The subhypergraph $\mathfrak{H}|_U$ induced by U is k -sparse. Hence, there exists some element $u \in U$ with $|\varepsilon^{-1}(u)| < k$. By choice of F we can find a sequence of edges $e_0, \dots, e_n \in F$ with

$$u \in e_0, \quad \varepsilon(e_i) \in e_{i+1}, \quad \text{and} \quad \varepsilon(e_n) = v.$$

We define a new choice function ι by setting

$$\iota(e) := \begin{cases} u & \text{if } e = e_0, \\ \varepsilon(e_{i-1}) & \text{if } e = e_i, i > 0, \\ \varepsilon(e) & \text{otherwise.} \end{cases}$$

It follows that

$$|\iota^{-1}(x)| = \begin{cases} |\varepsilon^{-1}(v)| - 1 & \text{if } x = v, \\ |\varepsilon^{-1}(u)| + 1 & \text{if } x = u, \\ |\varepsilon^{-1}(x)| & \text{otherwise.} \end{cases}$$

Hence, $w(\iota) < w(\varepsilon)$. Repeating this construction we obtain the desired choice function ε with $w(\varepsilon) = 0$.

The general case where \mathfrak{H} may be infinite can be proved using the Compactness Theorem for first-order logic. Let Δ be the elementary diagram of \mathfrak{H}_{in} . We can write down a formula φ stating that $\varepsilon : E \rightarrow V$ is a function such that

$$\varepsilon(e) \in e \quad \text{and} \quad |\varepsilon^{-1}(v)| \leq k, \quad \text{for all } e \in E \text{ and } v \in V.$$

By assumption and the first part of the proof, every finite subset of $\Delta \cup \{\varphi\}$ is satisfiable. Therefore we can use the Compactness Theorem to find a model \mathfrak{H}^+ of $\Delta \cup \{\varphi\}$. As \mathfrak{H}^+ satisfies Δ , there exists an elementary embedding $h : \mathfrak{H}_{\text{in}} \rightarrow \mathfrak{H}^+$. Furthermore, since every edge of \mathfrak{H} has only finitely

many vertices, it follows that

$$v \in h(e) \quad \text{implies} \quad v = h(u), \quad \text{for some } u \in e.$$

Hence, we can obtain the desired choice function of \mathfrak{H} by restricting that of \mathfrak{H}^+ . \square

Exercise 2.1. Find a hypergraph \mathfrak{H} that has a choice function φ of indegree at most k such that the hypergraph obtained from \mathfrak{H} by removing from every hyperedge e the vertex $\varepsilon(e)$ has no choice function of indegree k .

It remains to show that the choice function obtained via the preceding proposition is MSO-definable. The following sequence of lemmas shows how we can encode such a function by a finite set of unary predicates. We start simply, with the case of graphs.

Theorem 2.4. Let \mathcal{S}_k be the class of all directed graphs whose edge relation is irreflexive, antisymmetric, and has indegree at most k .

- (a) Every $\mathfrak{G} \in \mathcal{S}_k$ is $(2k+1)$ -colourable.
- (b) There exists a finite graph $\mathfrak{T}_k \in \mathcal{S}_k$ such that every $\mathfrak{G} \in \mathcal{S}_k$ admits a homomorphism $\mathfrak{G} \rightarrow \mathfrak{T}_k$.

Proof. (a) If $\mathfrak{G} \in \mathcal{S}_k$ is finite, we can prove the claim by induction on the number of vertices. Note that \mathfrak{G} is k -sparse. Since the sum of all outdegrees is equal to the number of edges, which is bounded by $k \cdot |V|$, there must be some vertex v of outdegree at most k . By inductive hypothesis, the subgraph $\mathfrak{G} - v$ has a $(2k+1)$ -colouring χ . Since v has at most $2k$ neighbours, we can pick some colour c that is different from the colours of these neighbours. Consequently, we can extend χ to a colouring of \mathfrak{G} by setting $\chi(v) := c$.

It remains to consider the case where $\mathfrak{G} = \langle V, E \rangle$ is infinite. The set

$$\Phi := \bigcup_{i < 2k+1} \left\{ Euv \wedge \neg(P_i u \wedge P_i v) \mid \langle u, v \rangle \in E \right\} \\ \cup \left\{ \forall x \bigvee_i P_i x \wedge \forall x \bigwedge_{i \neq j} \neg(P_i x \wedge P_j x) \right\}$$

states that the family $(P_i)_{i < 2k+1}$ encodes a $(2k+1)$ -colouring of \mathfrak{G} . As every finite subgraph of \mathfrak{G} is $(2k+1)$ -colourable, all finite subsets of Φ are

satisfiable. By the Compactness Theorem, so is therefore the whole set Φ . Let $\mathfrak{G}_+ = \langle V_+, E_+, \bar{P} \rangle$ be a model. Then \mathfrak{G} is a subgraph of $\langle V_+, E_+ \rangle$ and the restriction of the predicates \bar{P} to V induce a $(2k+1)$ -colouring of \mathfrak{G} .

(b) The graph \mathfrak{T}_k has the vertex set

$$T := [2]^{<p-1} \quad \text{where} \quad p := 2k(k+1) + 1.$$

Given $\bar{a}, \bar{b} \in T$ with $|\bar{b}| = i < |\bar{a}|$, we add an edge $\bar{a} \rightarrow \bar{b}$ if $a_i = 1$, and an edge $\bar{b} \rightarrow \bar{a}$ if $a_i = 0$. If $|\bar{a}| = |\bar{b}|$ there is no edge between \bar{a} and \bar{b} .

We claim that the graph \mathfrak{T}_k defined in this way has the desired property. Fix $\mathfrak{G} \in \mathcal{S}_k$. Let \mathfrak{G}' be the graph obtained from \mathfrak{G} by adding all edges $u \rightarrow v$ such that, in \mathfrak{G} , there is a path of length 2 from u to v . Then \mathfrak{G}' has indegree at most $k(k+1)$ and we can use (a) to construct a p -colouring χ of \mathfrak{G}' . We use χ to define the desired homomorphism $h : \mathfrak{G} \rightarrow \mathfrak{T}_k$ as follows. Given a vertex $v \in V$ with colour $c := \chi(v)$, we set

$$h(v) := \langle a_0, \dots, a_{c-1} \rangle,$$

where

$$a_i = 0 \quad : \text{iff} \quad \text{there is some edge } u \rightarrow v \text{ with } u \in \chi^{-1}(i).$$

To show that this defines a homomorphism, consider two vertices u, v connected by an edge (in either direction). By symmetry, we may assume that $i := \chi(u) < \chi(v) =: c$. Let $\bar{a} := h(v)$ and $\bar{b} := h(u)$. If the edge is directed from u to v , we have $a_i = 0$, which implies that $\bar{a} \rightarrow \bar{b}$ is an edge of \mathfrak{T}_k . If the edge has the other direction, then the fact that χ is a colouring of \mathfrak{G}' implies that there is no edge $w \rightarrow v$ with $\chi(w) = i = \chi(u)$. Consequently, $a_i = 1$ and it follows that $\bar{b} \rightarrow \bar{a}$ is an edge of \mathfrak{T}_k . \square

The graphs \mathfrak{T}_k from the preceding theorem can be used to encode choice functions of hypergraphs. We start with a bit of terminology.

Definition 2.5. Let $\mathfrak{T} = \langle V, E \rangle$ be a directed graph and $\mathfrak{H} = \langle V, E \rangle$ a hypergraph.

(a) Given a choice function ε of \mathfrak{H} , we define a directed graph $\mathcal{O}_\varepsilon(\mathfrak{H}) := \langle V, F \rangle$ with edge relation

$$F := \{ \langle u, v \rangle \mid u \neq v \text{ and there is some edge } e \in E \text{ with } u \in e \text{ and } \varepsilon(e) = v \}.$$

(b) A \mathfrak{T} -colouring of \mathfrak{H} consists of a pair $\langle \varepsilon, h \rangle$ where ε is a choice function of \mathfrak{H} and h is a homomorphism $\mathcal{O}_\varepsilon(\mathfrak{H}) \rightarrow \mathfrak{T}$. The *indegree* of such a \mathfrak{T} -colouring $\langle \varepsilon, h \rangle$ is the indegree of ε .

(c) A family $(P_v)_{v \in V}$ of unary predicates *encodes* an \mathfrak{T} -colouring $\langle \varepsilon, h \rangle$ of \mathfrak{H} if $P_v = h^{-1}(v)$, for all $v \in V$.

(d) We say that a formula $\varphi(x, Y)$ *defines a choice function* ε for \mathfrak{H} if

$$\mathfrak{H}_{\text{adj}} \models \varphi(v, e) \quad \text{iff} \quad e \in E \text{ and } v = \varepsilon(e).$$

Proposition 2.6. Let $\mathfrak{H} = \langle V, E \rangle$ be a k -sparse hypergraph of rank m where $0 < k < \omega$ and $1 < m < \omega$. Then \mathfrak{H} has a choice function ε of indegree at most mk^2 such that the edge relation of $\mathcal{O}_\varepsilon(\mathfrak{H})$ is antisymmetric.

Proof. First, we consider the case where \mathfrak{H} is finite. We call a vertex $u \in V$ *bad* for a choice function ε of \mathfrak{H} if there is some vertex $v \in V$ such that $\mathcal{O}_\varepsilon(\mathfrak{H})$ contains both edges $\langle u, v \rangle$ and $\langle v, u \rangle$. Note that this implies that the vertex v is also bad.

We construct a sequence of choice functions $(\varepsilon_n)_n$ such that

$$|\varepsilon_n^{-1}(v)| \leq \begin{cases} k & \text{if } v \text{ is bad for } \varepsilon_n, \\ mk^2 & \text{otherwise,} \end{cases}$$

and the number of bad elements decreases at every step. We start with an arbitrary choice function ε_0 bounded by k .

Given a choice function ε_n with the above properties we construct a new choice function ε_{n+1} with fewer bad elements as follows. Let v be a bad element, set $X := \varepsilon_n^{-1}(v)$, and let

$$Y := \{ e \mid v \in e \text{ and } \varepsilon_n(e) \in \bigcup X \setminus \{v\} \}.$$

Since v is bad we have

$$|X| \leq k \quad \text{and} \quad |\bigcup X| \leq k(m-1).$$

Note that every element of the form $u := \varepsilon_n(e)$ with $e \in Y$ is also bad since, by definition of X , there is an edge $e' \in X$ with

$$u \in e' \quad \text{and} \quad \varepsilon_n(e') = v.$$

Consequently, $\mathcal{O}_{\varepsilon_n}(\mathfrak{H})$ contains the edges $\langle u, v \rangle$ (since $\varepsilon_n(e') = v$) and $\langle v, u \rangle$ (since $\varepsilon_n(e) = u$). It follows that

$$|Y| \leq k \cdot |\bigcup X \setminus \{v\}| \leq k^2(m-1).$$

We define the new choice function ε_{n+1} by

$$\varepsilon_{n+1}(e) := \begin{cases} v & \text{if } e \in Y, \\ \varepsilon_n(e) & \text{otherwise.} \end{cases}$$

Then we have

$$|\varepsilon_{n+1}(x)^{-1}| \leq \begin{cases} k + k^2(m-1) & \text{if } x = v, \\ |\varepsilon_n(x)^{-1}| & \text{otherwise.} \end{cases}$$

In particular, ε_{n+1} is bounded by mk^2 . By construction, the element v is not bad for ε_{n+1} . Furthermore, if $\langle u, w \rangle$ is an edge in $\mathcal{O}_{\varepsilon_{n+1}}(\mathfrak{H})$ with $u, w \neq v$ then this edge is induced by an edge e in \mathfrak{H} with $e \notin X \cup Y$. Hence, $\langle u, w \rangle$ is also an edge of $\mathcal{O}_{\varepsilon_n}(\mathfrak{H})$. Therefore, every element that is bad for ε_{n+1} is also bad for ε_n .

It remains to prove the claim for infinite hypergraphs \mathfrak{H} . Let Φ be the union of the elementary diagram of \mathfrak{H} and formulae stating that ε is a choice function of \mathfrak{H} that is bounded by mk^2 and that $\mathcal{O}_\varepsilon(\mathfrak{H})$ has an antisymmetric edge relation. If \mathfrak{M} is a model of Φ then there exists an embedding $h : \mathfrak{H} \rightarrow \mathfrak{M}$ and the desired choice function of \mathfrak{H} can be obtained via h from that of \mathfrak{M} . Hence, it is sufficient to show that Φ is satisfiable. Note that every finite subset $\Phi_o \subseteq \Phi$ is satisfiable since every finite substructure of \mathfrak{H} has a choice function of the desired form. By the Compactness Theorem, it therefore follows that whole set Φ is also satisfiable. \square

Corollary 2.7. *Let $m, k < \omega$.*

- (a) *Every k -sparse hypergraph \mathfrak{H} of rank at most m has a \mathfrak{T}_{mk^2} -colouring of indegree at most mk^2 .*
- (b) *There exists an FO-formula $\eta_k(\bar{X})$ such that, for every undirected hypergraph \mathfrak{H} of rank at most m ,*

$$\mathfrak{H}_{\text{adj}} \models \eta_k(\bar{P}) \quad \text{iff} \quad \bar{P} \text{ encodes a } \mathfrak{T}_k\text{-colouring of } \mathfrak{H} \text{ of indegree at most } k.$$

- (c) *There exists an FO-formula $\varphi(x, Y; \bar{Z})$ such that, for every k -sparse hypergraph \mathfrak{H} of rank at most m , there are monadic parameters \bar{P} such that $\varphi(x, Y; \bar{P})$ defines a choice function ε for \mathfrak{H} of indegree at most mk^2 .*

Proof. (a) In Proposition 2.6, we have shown that such a graph $\mathfrak{H} = \langle V, E \rangle$ has a choice function $\varepsilon : E \rightarrow V$ of indegree at most mk^2 such that the edge relation of $\mathcal{O}_\varepsilon(\mathfrak{H})$ is antisymmetric. Then $\mathcal{O}_\varepsilon(\mathfrak{H})$ has indegree at most mk^2 as well, and we can use Theorem 2.4 (b) to find a homomorphism $h : \mathcal{O}_\varepsilon(\mathfrak{H}) \rightarrow \mathfrak{T}_{mk^2}$. Thus, $\langle \varepsilon, h \rangle$ is the desired \mathfrak{T}_{mk^2} -colouring.

(b) Suppose that $\mathfrak{T}_k = \langle T, F \rangle$. The formula $\eta_k(\bar{Z})$ has to express the following three conditions.

- (i) The sets Z_t form a partition of the vertices (some Z_t may be empty).

$$\bigwedge_{s \neq t} Z_s \cap Z_t = \emptyset \wedge \forall x \bigvee_{t \in T} Z_t x$$

- (ii) The function $\mathcal{O}_\varepsilon(\mathfrak{H}) \rightarrow \mathfrak{T}_k$ encoded in \bar{Z} is a homomorphism.

$$\forall \bar{x} \left[E_m \bar{x} \rightarrow \bigvee_{j < m} \left\{ \bigwedge_{i < m} Z_{t_i} x_i \mid \langle t_i, t_j \rangle \in F \text{ for all } i \neq j \right\} \right]$$

- (iii) The indegree is at most mk^2 .

$$\forall x \exists y_0 \dots \exists y_{k-1} \forall z \left[\psi(z, x) \rightarrow \bigvee_{i < k} z = y_i \right]$$

where

$$\psi(x, y) := \exists \bar{z} \left[E_m \bar{z} \wedge \bigvee_{i \neq j} [x = z_i \wedge y = z_j] \wedge \bigvee_{\langle s, t \rangle \in F} [Z_s x \wedge Z_t y] \right].$$

states that $\langle x, y \rangle$ is an edge of $\mathcal{O}_\varepsilon(\mathfrak{H})$ that is directed from x to y by the colouring.

(c) We set

$$\begin{aligned} \varphi(x, Y; \bar{Z}) := & Yx \wedge \exists \bar{z} \left[E_m \bar{z} \wedge \forall y \left[Yy \leftrightarrow \bigvee_{i < m} y = z_i \right] \right] \\ & \wedge \forall y [Yx \wedge y \neq x \rightarrow \psi(x, y)], \end{aligned}$$

where $\psi(x, y)$ is the formula from (b) (III). Given parameters \bar{P} encoding a \mathfrak{T}_{mk^2} -colouring of \mathfrak{H} of indegree at most mk^2 , it follows that $\varphi(x, Y; \bar{P})$ defines a choice function of at most that indegree. \square

Exercise 2.2. Prove that the class of all k -sparse undirected graphs is finitely MSO-axiomatisable. \urcorner

3 Translating GSO Into MSO

We are finally able to prove that one can translate every GSO-formula φ into an MSO-formula ψ that is equivalent to φ on all k -sparse hypergraphs. The general idea of the proof is as follows. When we want to represent sets of hyperedges by sets of vertices, we can fix a choice function ε of bounded indegree k . Then we can encode each hyperedge e by a pair $\langle i, v \rangle$ where $v := \varepsilon(e)$ and $i < k$ is some number. To know which number i to use, we can use the results of Section 1 to define a linear ordering on each hyperedge.

Exercise 3.1. Prove that, for every GSO-formula φ , one can construct an MSO-formula φ' such that

$$\mathfrak{G} \models \varphi \quad \text{iff} \quad \mathfrak{G} \models \varphi', \quad \text{for every finite grid } \mathfrak{G}. \quad \urcorner$$

It will be convenient to phrase the proof below in terms of the following operation, which is a combination of an interpretation and the copying operation. Given a structure \mathfrak{A} it first creates a structure of the form $\mathfrak{A} \oplus \mathfrak{B}$, where \mathfrak{B} is definable in $\text{copy}_k(\mathfrak{A})$, for some k , and then it adds definable relations between A and B . To simplify the definition, we will work with many-sorted structures in this section.

Definition 3.1. Let $\Xi \subseteq \Xi_+$ be sets of sorts and Σ and Γ two signatures where Σ is Ξ -sorted and Γ is Ξ_+ -sorted. An MSO-definable expansion, or an MSO-expansion for short, is an operation of the form

$$\tau = \tau_o \circ \text{copy}_k, \quad \text{for some } k < \omega,$$

where $\tau_o = \langle (\delta_\xi(x))_{\xi \in \Xi_+}, (\varphi_R(\bar{x}))_{R \in \Gamma} \rangle$ is an MSO-interpretation from Σ to Γ such that

$$\begin{aligned} \delta_\xi(x) &= H_o x, & \text{for } \xi \in \Xi, \\ \delta_\xi(x) &\text{ implies } \neg H_o x, & \text{for } \xi \in \Xi_+ \setminus \Xi, \\ \varphi_R(\bar{x}) &= R\bar{x} \wedge \bigwedge_i H_o x_i, & \text{for } R \in \Sigma \cap \Gamma, \end{aligned}$$

where H_o, \dots, H_{k-1} are the colour predicates introduced by copy_k . \urcorner

Remarks. (a) Note that there is a canonical inclusion $\mathfrak{A} \rightarrow \tau(\mathfrak{A})$ for every MSO-expansion τ and every input structure \mathfrak{A} .

(b) An MSO-expansion is a special form of an operation called an MSO-transduction. We will study these in Chapter X. \urcorner

We obtain a composition theorem for MSO-expansions immediately from the corresponding theorems for interpretations and the copy operation.

Lemma 3.2. Let τ be an MSO-expansion and $\varphi(\bar{x}, \bar{Y}; \bar{Z})$ an MSO-formula. There exists an MSO-formula $\varphi^\tau(\bar{x}, \bar{Y}; \bar{Z}')$ such that, for every structure \mathfrak{A} and all MSO-parameters \bar{Q} in $\tau(\mathfrak{A})$, there are MSO-parameters \bar{Q}' in \mathfrak{A} satisfying

$$\tau(\mathfrak{A}) \models \varphi(\bar{a}, \bar{P}; \bar{Q}) \quad \text{iff} \quad \mathfrak{A} \models \varphi^\tau(\bar{a}, \bar{P}; \bar{Q}'),$$

for all parameters \bar{a}, \bar{P} in \mathfrak{A} .

Slightly less trivial is the fact that MSO-expansions are closed under composition.

Lemma 3.3. *If σ and τ are MSO-expansions, so is $\tau \circ \sigma$.*

Proof. Suppose that $\sigma = \sigma_o \circ \text{copy}_k$ and $\tau = \tau_o \circ \text{copy}_m$. By Lemmas I.4.20 and I.4.7, there exist MSO-interpretations σ'_o and ρ such that

$$\begin{aligned} \tau \circ \sigma &= \tau_o \circ \text{copy}_m \circ \sigma_o \circ \text{copy}_k \\ &= \tau_o \circ \sigma'_o \circ \text{copy}_m \circ \text{copy}_k = \rho \circ \text{copy}_{mk}. \end{aligned}$$

It remains to check that ρ is of the correct form. Let H_i and H'_j be the colour predicates introduced by, respectively, copy_k and copy_m . Then the colour predicates of copy_{mk} are of the form $H_{i,j} = H_i \times H'_j$ (up to a renaming of the indices). By assumption on σ_o and τ_o , there exists a bijection $A \rightarrow H_{o,o}$ that preserves all relations. Consequently, we can choose the formulae of ρ such that they also have the required form. \square

Definition 3.4. Let \mathfrak{B} and \mathfrak{C} be Γ -structures and \mathfrak{A} a Σ -structure with $\mathfrak{A}|_{\Sigma \cap \Gamma} \subseteq \mathfrak{B}|_{\Sigma \cap \Gamma}$ and $\mathfrak{A}|_{\Sigma \cap \Gamma} \subseteq \mathfrak{C}|_{\Sigma \cap \Gamma}$. We denote by $\mathfrak{B} \cup_{\mathfrak{A}} \mathfrak{C}$ the Γ -structure obtained from the disjoint union $\mathfrak{B} \oplus \mathfrak{C}$ by identifying the two copies of \mathfrak{A} .

Lemma 3.5. *Let σ and τ be two MSO-expansions. There exists an MSO-expansion $\sigma \cup \tau$ such that*

$$(\sigma \cup \tau)(\mathfrak{A}) = \sigma(\mathfrak{A}) \cup_{\mathfrak{A}} \tau(\mathfrak{A}), \quad \text{for all } \mathfrak{A}.$$

Proof. Suppose that $\sigma = \sigma_o \circ \text{copy}_k$ and $\tau = \tau_o \circ \text{copy}_m$ with

$$\begin{aligned} \sigma_o &= \langle (\delta_\xi(x))_{\xi \in \Xi_+}, (\varphi_R(\tilde{x}))_{R \in \Gamma} \rangle, \\ \tau_o &= \langle (\gamma_\xi(x))_{\xi \in \Xi_+}, (\psi_R(\tilde{x}))_{R \in \Gamma} \rangle. \end{aligned}$$

We set

$$\sigma \cup \tau := \rho \circ \text{copy}_{k+m-1},$$

where the interpretation ρ is defined as follows. Intuitively, if H_o, \dots, H_{k-1} and H'_o, \dots, H'_{m-1} are the partitions of $\sigma(\mathfrak{A})$ and $\tau(\mathfrak{A})$ induced by the copy-operations, the structure $\sigma(\mathfrak{A}) \cup_{\mathfrak{A}} \tau(\mathfrak{A})$ has a partition

$$H_o + H_1 + \dots + H_{k-1} + H'_1 + \dots + H'_{m-1},$$

which the interpretation ρ recreates. To do so, it has to re-index the sets H'_i from $\tau(\mathfrak{A})$. For a formula φ , we denote by φ' the formula obtained from φ by replacing every atom of the form $H_i z$ with $i > o$ by $H_{k+i-1} z$. Furthermore, let

$$\alpha(x) := \bigvee_{o < i < k} H_i x \quad \text{and} \quad \beta(x) := \bigvee_{o < i < m} H_{k+i-1} x.$$

We set $\rho := \langle (\eta_\xi(x))_{\xi \in \Xi_+}, (\vartheta_R(\tilde{x}))_{R \in \Gamma} \rangle$ where

$$\begin{aligned} \eta_\xi(x) &:= \begin{cases} H_o x & \text{for } \xi \in \Xi, \\ [\alpha(x) \wedge \delta_\xi(x)] \vee [\beta(x) \wedge \gamma'_\xi(x)] & \text{for } \xi \in \Xi_+ \setminus \Xi, \end{cases} \\ \vartheta_R(\tilde{x}) &:= \left[\bigwedge_i [H_o x_i \vee \alpha(x_i)] \wedge \varphi_R(\tilde{x}) \right] \\ &\quad \vee \left[\bigwedge_i [H_o x_i \vee \beta(x_i)] \wedge \psi_R(\tilde{x}) \right]. \end{aligned} \quad \square$$

With these preparations out of the way we can finally state and proof the theorem we have worked towards the preceding sections.

Theorem 3.6 (Courcelle). *Let $k, m < \omega$. There exists an MSO-expansion τ such that, for every non-empty k -sparse hypergraph \mathfrak{H} of rank at most m , there are monadic parameters \tilde{P} such that*

$$\tau(\langle \mathfrak{H}_{\text{adj}}, \tilde{P} \rangle) = \mathfrak{H}_{\text{in}}.$$

Proof. Note that $\mathfrak{H}_{\text{in}} \cong \mathfrak{H}_{\text{in}}^o \cup_{\mathfrak{H}_{\text{adj}}} \dots \cup_{\mathfrak{H}_{\text{adj}}} \mathfrak{H}_{\text{in}}^m$, where

$$\mathfrak{H}^i := \langle V, E_i \rangle \quad \text{with} \quad E_i := \{ e \in E \mid |e| = i \}.$$

It is therefore sufficient to construct MSO-expansions τ_i mapping $\mathfrak{H}_{\text{adj}}^i$ to $\mathfrak{H}_{\text{in}}^i$. Then $\tau := \tau_0 \cup \dots \cup \tau_m$ is the desired MSO-expansion mapping $\mathfrak{H}_{\text{adj}}$ to \mathfrak{H}_{in} . Replacing \mathfrak{H} by some \mathfrak{H}^i , we may therefore assume that all edges of \mathfrak{H} have exactly m vertices. We construct the desired MSO-expansion τ by induction on m .

($m = 0$) Since \mathfrak{H} is k -sparse, it has at most k hyperedges of size 0. Consequently, we can interpret \mathfrak{H}_{in} in $\text{copy}_{k+1}(\mathfrak{H}_{\text{adj}})$.

($m = 1$) As \mathfrak{H} is k -sparse there are, for every vertex $v \in V$, at most k hyperedges e with vertex v . Again, we can interpret \mathfrak{H}_{in} in $\text{copy}_{k+1}(\mathfrak{H}_{\text{adj}})$.

($m = 2$) If all hyperedges have two vertices, the hypergraph is an undirected graph. By Corollary 2.7 (c), there exists an FO-formula $\varphi(x, Y; \bar{P})$ defining a choice function ε for \mathfrak{H} of indegree at most $2k^2$. We use this function to turn \mathfrak{H} into a directed graph by orienting each edge e towards $\varepsilon(e)$. By Corollary 1.18, there exists an MSO-formula $\psi(x, y; z, \bar{Q})$ defining a predecessor ordering on the resulting graph \mathfrak{G} . For a vertex v , we denote by $p_i(v)$ the i -th predecessor with respect to this ordering. We encode \mathfrak{H}_{in} in $\text{copy}_{k+1}(\mathfrak{H}, \bar{P}, \bar{Q})$ as follows. Vertices v of \mathfrak{H} are represented by the $(k+1)$ -th copy $\langle k, v \rangle$, while an edge $p_i(v) \rightarrow v$ is represented by the i -th copy $\langle i, v \rangle$. Clearly, the incidence relation in is definable using this encoding. Hence, there exists an MSO-interpretation τ_0 mapping $\text{copy}_{k+1}(\mathfrak{H}, \bar{P}, \bar{Q})$ to \mathfrak{H}_{in} . The composition $\tau_0 \circ \text{copy}_{k+1}$ is the desired MSO-expansion mapping $\langle \mathfrak{H}, \bar{P} \rangle$ to \mathfrak{H}_{in} .

($m > 2$) The basic idea is simple: given a hypergraph of rank m , we choose a choice function ε and we remove from every hyperedge e the vertex $\varepsilon(e)$. This results in a hypergraph of rank $m-1$, to which we apply the inductive hypothesis. Unfortunately, this plan does not work: the new hypergraph might not be sparse. Hence, we need the following more complicated construction.

Given a k -sparse hypergraph $\mathfrak{H} = \langle V, E \rangle$ of rank at most m , we use Corollary 2.7 (c) to construct an FO-formula $\varphi(x, Y; \bar{P})$ defining a choice function ε of indegree at most mk^2 . Let $\mathfrak{D} := \langle D, R \rangle$ be the hypergraph

with

$$\begin{aligned} D &:= \{ \langle v, \varepsilon(e) \rangle \mid e \in E, v \in e \setminus \{\varepsilon(e)\} \}, \\ R &:= \{ \{ \langle v_0, u \rangle, \dots, \langle v_{m-2}, u \rangle \} \mid \\ &\quad e = \{u, v_0, \dots, v_{m-2}\} \in E, u = \varepsilon(e) \}. \end{aligned}$$

Note that \mathfrak{D} has rank $m-1$ and that its vertices are the edges of $\mathcal{O}_\varepsilon(\mathfrak{H})$. Let $s, t : D \rightarrow V$ be the functions

$$s(\langle u, v \rangle) := u \quad \text{and} \quad t(\langle u, v \rangle) := v.$$

We start by proving that \mathfrak{D} is $2k$ -sparse. Fix a finite set $X \subseteq D$. Note that

$$\{ \langle v_0, u \rangle, \dots, \langle v_{m-2}, u \rangle \} \in R \upharpoonright X$$

implies that

$$\{u, v_0, \dots, v_{m-2}\} \in E \upharpoonright Y, \quad \text{where } Y := s[X] \cup t[X].$$

Since \mathfrak{H} is k -sparse, it therefore follows that

$$|R \upharpoonright X| \leq |E \upharpoonright Y| \leq k \cdot |Y| \leq k \cdot (|s[X]| + |t[X]|) \leq 2k \cdot |X|,$$

as desired.

Using the formula $\varphi(x, Y; \bar{P})$ from above, we can interpret the graph $\langle V, D, s, t \rangle$ (where we consider D as a binary relation) in $\langle \mathfrak{H}, \bar{P} \rangle$. As every vertex of $\langle V, D \rangle$ has indegree at most $mk^2(m-1)$, this graph is $mk^2(m-1)$ -sparse. By inductive hypothesis, there therefore exists an MSO-expansion mapping $\langle V, D, s, t \rangle$ to its incidence structure $\langle V, D, \text{in} \rangle$ (where D is now considered as one of the domains). A trivial modification of this MSO-expansion yields the structure $\langle V, D, E, s, t \rangle$ instead (where we consider E as a relation of arity m). Since R (as a relation of arity $m-1$) is definable using φ , there therefore exists an MSO-expansion mapping $\langle \mathfrak{H}, \bar{P} \rangle$ to $\langle V, D, E, R, s, t \rangle$. As $\langle D, R \rangle$ is $2k$ -sparse of rank $m-1$, we can use the inductive hypothesis to find an MSO-expansion mapping $\langle V, D, E, R, s, t \rangle$

to the structure $\langle V, D, R, E, s, t, \text{in}^\mathfrak{D} \rangle$ where R is now considered as one of the domains and $\text{in}^\mathfrak{D} \subseteq D \times R$ is the incidence relation of \mathfrak{D} . Finally, we can define $\mathfrak{H}_{\text{in}} = \langle V, E, \text{in}^\mathfrak{H} \rangle$ in this latter structure since

$$\langle v, e \rangle \in \text{in}^\mathfrak{H} \quad \text{iff} \quad v = s(d) \text{ or } v = t(d),$$

where $d \in R$ is the hyperedge corresponding to e .

Hence,

$$\mathfrak{H}_{\text{in}} \cong \langle V, R, F \rangle \quad \text{where} \quad F := \{ \langle s(d), d \rangle, \langle t(d), d \rangle \mid d \in R \}. \quad \square$$

Theorem 3.7. *Let Σ be a finite relational signature and $k < \omega$. There exists an MSO-expansion τ such that, for every non-empty k -sparse Σ -structure \mathfrak{A} , there are monadic parameters \bar{P} such that*

$$\tau(\langle \mathfrak{A}, \bar{P} \rangle) = \mathfrak{A}_{\text{in}}.$$

Proof. Recall that $\mathfrak{A}_{\text{in}} = \langle A, E, (P_{R,\sigma})_{R,\sigma}, (\text{in}_i)_i \rangle$. For each $R \in \Sigma$ and every map σ , we will construct an MSO-expansion $\tau_{R,\sigma}$ mapping \mathfrak{A} to the structure $\mathfrak{A}_{R,\sigma} := \langle A, P_{R,\sigma}, (\text{in}_i)_i \rangle$, where $P_{R,\sigma}$ is considered as one of the domains. Then we can obtain \mathfrak{A}_{in} from the union

$$\bigcup_{\mathfrak{A}} \{ \mathfrak{A}_{R,\sigma} \mid R, \sigma \}$$

by (i) merging elements of $P_{R,\sigma}$ and $P_{R',\sigma'}$ that represent the same tuple and (ii) redefining some of the relations. Consequently, we can obtain the desired MSO-expansion τ from $\bigcup_{R,\sigma} \tau_{R,\sigma}$ by a straightforward modification.

It therefore remains to construct the MSO-expansions $\tau_{R,\sigma}$. Fix a relation $R \in \Sigma$ of arity n and an injective map $\sigma : [k] \rightarrow [n]$. By Theorem 3.6, there exists an MSO-expansion σ mapping the hypergraph $\langle A, P_{R,\sigma} \rangle$ to its incidence representation $\langle A, P_{R,\sigma}, \text{in} \rangle$. We have to define the relations in_i in this structure. We can use Theorem 1.16 to define an edge ordering of $\langle A, P_{R,\sigma}, \text{in} \rangle$. For every bijection $\rho : [k] \rightarrow [k]$, let Q_ρ be the set of all tuples $\bar{c} \in P_{R,\sigma}$ such that the edge ordering defined on \bar{c} is $c_{\rho(0)} < \dots < c_{\rho(k-1)}$.

Then

$$\langle v, e \rangle \in \text{in}_i \quad \text{iff} \quad \langle v, e \rangle \in \text{in} \text{ and there is some } \rho \text{ such that } e \in Q_\rho \text{ and } v \text{ is the } \rho(i)\text{-th element of } e \text{ in its edge ordering.}$$

Clearly, this can be expressed in MSO. \square

Corollary 3.8. *Let Σ be a finite relational signature and $k < \omega$. Every GSO-definable property is MSO-definable over the class of all k -sparse Σ -structures.*

4 Sparse Distributions

The results so far concern ways to encode edges by vertices. In this section we consider a more general problem. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph. We denote by $\wp_{\text{fin}}(V)$ the set of all finite subsets of V . We would like to encode a given subset $F \subseteq \wp_{\text{fin}}(V)$ by a set of vertices, that is, we would like to find a definable function $h : F \rightarrow V$ that is injective. For $F = E$ this reduces to the problem considered in the preceding sections. For an arbitrary F , such a function h does not always exist. But we will show that sometimes we can transform a given function $h_0 : F \rightarrow V$ into an injective one.

Suppose we are given a function $h_0 : F \rightarrow V$ that we want to transform into an injective function $h : F \rightarrow V$. Let $\delta(v) := |h_0^{-1}(v)|$. The first step in the construction of h consists in finding a definable function $g : V \rightarrow V$ such that $|g^{-1}(v)| = \delta(v)$, for all v . Of course, this is not always possible. For instance, if the graph is finite and we have $\delta(v) > 1$, for all vertices v . Therefore, we consider only functions δ that are sparse in the sense of the following definition.

Definition 4.1. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) The *border* of a subset $Z \subseteq V$ is the set

$$B_{\mathfrak{G}}(Z) := E \cap Z \times (V \setminus Z)$$

of all edges connecting a vertex in Z with a vertex outside of Z .

(b) A *distribution* of \mathfrak{G} is a map $\delta : V \rightarrow \omega$. For a set $Z \subseteq V$, we set

$$\delta(Z) := \sum_{v \in Z} \delta(v).$$

(c) A distribution δ is *k-sparse* if

$$\delta(Z) \leq |Z| + k \cdot |B_{\mathfrak{G}}(Z)|, \quad \text{for every finite } Z \subseteq V.$$

Given a *k-sparse* distribution δ we will construct the desired function $g : V \rightarrow V$ by solving a network flow problem.

Definition 4.2. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) A *flow* of \mathfrak{G} is a function $f : V \times V \rightarrow \mathbb{Z}$ such that, for all $u, v \in V$,

- ♦ $f(u, v) = -f(v, u)$ and
- ♦ $f(u, v) \neq 0$ implies $\langle u, v \rangle \in E$.

(b) A flow f is *acyclic* if there is no cycle u_0, \dots, u_m of \mathfrak{G} with

$$f(u_i, u_{i+1}) > 0, \quad \text{for all } i \leq m \quad (\text{index arithmetic modulo } m).$$

(c) The *defect* of a flow f is the function

$$d_f(v) := \sum_{u \in V} f(v, u).$$

(d) Let $\delta : V \rightarrow \omega$ be a distribution. A flow f is a δ -*flow* if, for every $v \in V$, either

$$d_f(v) = \delta(v) - 1, \quad \text{or} \quad \delta(v) = 0 \text{ and } d_f(v) = 0.$$

(e) A flow f is *edge-bounded* by k if $|f(u, v)| \leq k$, for all $u, v \in V$. We call f *vertex-bounded* by k if

$$\sum_{u \in V} |f(u, v)| \leq k, \quad \text{for all } v \in V.$$

Our aim is to show that, for every *k-sparse* distribution δ there is a bounded δ -flow f and a function $g : V \rightarrow V$ inducing δ . Furthermore, if δ is definable then g should also be definable.

Definition 4.3. Let L be a logic.

(a) A distribution δ is *L-definable* if there exist formulae $\varphi_i(x) \in L$, $i < k$, such that

$$\mathfrak{G} \models \varphi_i(v) \quad \text{iff} \quad \delta(v) = i.$$

(b) Similarly, a flow f is *L-definable* if there exist formulae $\varphi_i(x, y) \in L$ such that

$$\mathfrak{G} \models \varphi_i(u, v) \quad \text{iff} \quad f(u, v) = i.$$

Remark. Note that every edge-bounded flow f can be encoded with the help of the GSO-parameters

$$S_i := \{ \langle u, v \rangle \in E \mid f(u, v) = i \}.$$

We start with a few lemmas about bounded flows. The first one follows immediately from the definitions.

Lemma 4.4. Let \mathfrak{G} be an undirected graph.

- (a) Every flow that is vertex-bounded by k is also edge-bounded by k .
- (b) If \mathfrak{G} has maximal degree d , every flow that is edge-bounded by k is vertex-bounded by dk .

Lemma 4.5. Let \mathfrak{G} be an undirected graph, δ a distribution, and f a flow such that

$$d_f(v) \geq \delta(v) - 1, \quad \text{for all vertices } v.$$

Then there exists a distribution $\delta' \geq \delta$ such that f is a δ' -flow.

Proof. We set

$$\delta'(v) := d_f(v) + 1.$$

Then $\delta'(v) \geq d_f(v) + 1 \geq \delta(v) \geq 0$ implies that δ' is a distribution and that $\delta' \geq \delta$. Furthermore, f is a δ' -flow since $d_f(v) = \delta'(v) - 1$, for all v . \square

Lemma 4.6. For every δ -flow f there exists an acyclic δ -flow f' such that, if f is edge-bounded by k or vertex-bounded by k , then so is f' .

Proof. We repeat the following construction until the flow is acyclic. Select a cycle u_o, \dots, u_m such that

$$c := \min \{ f(u_i, u_{i+1}) \mid i \leq m \} > 0$$

(with index arithmetic modulo m). We define f' by

$$f'(x, y) := \begin{cases} f(x, y) - c & \text{if } x = u_i \text{ and } y = u_{i+1}, \text{ for some } i, \\ f(x, y) + c & \text{if } x = u_{i+1} \text{ and } y = u_i, \text{ for some } i, \\ f(x, y) & \text{otherwise.} \end{cases} \quad \square$$

Our first main theorem is the fact that every distribution induces a flow.

Proposition 4.7. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph and δ a k -sparse distribution. Then \mathfrak{G} has a δ -flow f that is edge-bounded by k .

Proof. First, we assume that \mathfrak{G} is finite. In this case we can reduce the task to a network flow problem. Let \mathfrak{H} be the graph obtained from \mathfrak{G} by adding two new vertices s and t that are connected to every vertex of \mathfrak{G} . We define the capacity $c(e)$ of an edge e of \mathfrak{H} as follows. If e is an edge of \mathfrak{G} , we set $c(e) := k$. If $e = \langle s, v \rangle$ with $v \in V$ we set $c(e) := \max\{0, \delta(v) - 1\}$. Finally, if $e = \langle v, t \rangle$ with $v \in V$ we define

$$c(e) := \begin{cases} 0 & \text{if } \delta(v) > 0, \\ 1 & \text{otherwise.} \end{cases}$$

Let f be a maximal flow from s to t with respect to c . We claim that its restriction to the edges of \mathfrak{G} is the desired flow.

According to the Max-Flow Min-Cut Theorem, there is a set X of vertices containing s but not t such that the maximal flow m from s to t equals

$$m = \sum_{e \in B_{\mathfrak{H}}(X)} c(e).$$

Let $X_o := X \setminus \{s\} \subseteq V$ and $Y := \delta^{-1}(0)$. Since

$$B_{\mathfrak{H}}(X) = B_{\mathfrak{G}}(X_o) \cup \{ \langle v, t \rangle \mid v \in X_o \} \cup \{ \langle s, v \rangle \mid v \in V \setminus X_o \},$$

we have

$$\begin{aligned} m &= \sum_{e \in B_{\mathfrak{H}}(X)} c(e) \\ &= k \cdot |B_{\mathfrak{G}}(X_o)| + |X_o \cap Y| + (\delta(V \setminus X_o) - |(V \setminus X_o) \setminus Y|) \\ &= k \cdot |B_{\mathfrak{G}}(X_o)| + |X_o| + \delta(V \setminus X_o) \\ &\quad - |(V \setminus X_o) \setminus Y| - |X_o \setminus Y| \\ &\geq \delta(X_o) + \delta(V \setminus X_o) - |V \setminus Y| \\ &= \delta(V) - |V \setminus Y|. \end{aligned}$$

On the other hand, for the set $X = \{s\}$, we have

$$m \leq \sum_{e \in B_{\mathfrak{H}}(X)} c(e) = \sum_{v \in V} \max\{0, \delta(v) - 1\} = \delta(V) - |V \setminus Y|.$$

Consequently, the maximal flow m from s to t equals

$$m = \delta(V) - |V \setminus Y|.$$

This implies that

$$f(s, v) = \max\{0, \delta(v) - 1\}, \quad \text{for every } v \in V.$$

For each $v \in V$, we therefore have

$$0 = \sum_{u \in V \cup \{s, t\}} f(u, v) = \max\{0, \delta(v) - 1\} + f(t, v) + \sum_{u \in V} f(u, v).$$

If $\delta(v) > 0$ this implies

$$\delta(v) - 1 - \sum_{u \in V} f(v, u) = 0, \quad \text{that is} \quad d_f(v) = \delta(v) - 1,$$

while, for $\delta(v) = 0$, we have

$$-f(v, t) - \sum_{u \in V} f(v, u) = 0.$$

Hence, either $d_f(v) = -1 = \delta(v) - 1$ or $d_f(v) = 0$.

It remains to prove the lemma for infinite graphs. Let $\Phi(\mathfrak{G})$ consist of the elementary diagram of \mathfrak{G} together with first-order formulae stating that f is a δ -flow on \mathfrak{G} that is edge-bounded by k . We will use the compactness theorem to show that $\Phi(\mathfrak{G})$ is satisfiable.

Let $\Phi_0 \subseteq \Phi(\mathfrak{G})$ be finite. There exists a finite induced subgraph $\mathfrak{G}_0 = \langle V_0, E_0 \rangle$ of \mathfrak{G} such that $\Phi_0 \subseteq \Phi(\mathfrak{G}_0)$. Let $\langle u_0, v_0 \rangle, \dots, \langle u_{m-1}, v_{m-1} \rangle$ be an enumeration (without repetitions) of all edges $\langle u, v \rangle$ with $u \in V_0$ and $v \in V \setminus V_0$. We construct a new graph $\mathfrak{G}'_0 = \langle V'_0, E'_0 \rangle$ by attaching to each vertex u_i a path P_i of length k . Let δ' be the distribution on \mathfrak{G}'_0 with $\delta'(v) = \delta(v)$, for $v \in V_0$, and $\delta'(v) = 0$, for $v \in V'_0 \setminus V_0$. In order to show that Φ_0 is satisfiable, it is sufficient to prove that \mathfrak{G}'_0 has a flow of the desired form. Consider an arbitrary set $X \subseteq V'_0$ of vertices. Let

$$I := \{i \mid u_i \in X\} \quad \text{and} \quad J := \{i \mid u_i \in X \text{ and } P_i \subseteq X\}.$$

It follows that

$$\begin{aligned} \delta'(X) &= \delta(X \cap V_0) \leq |X \cap V_0| + k \cdot |B_{\mathfrak{G}}(X \cap V_0)| \\ &\leq |X| - k \cdot |J| + k \cdot |B_{\mathfrak{G}_0}(X \cap V_0)| \\ &\leq |X| - k \cdot |J| + k \cdot (|B_{\mathfrak{G}_0}(X)| + |J|) \\ &= |X| + k \cdot |B_{\mathfrak{G}'_0}(X)|. \end{aligned}$$

By the first part of the proof it follows that \mathfrak{G}'_0 has a flow of the desired form. \square

For the application below, we need flows that are WMSO-definable. But to prove their existence we have to make a few concessions: the proof only works for trees and the flow we obtain is only a δ' -flow, for some $\delta' \geq \delta$. We start with the case of finite trees.

Lemma 4.8. *Let $\mathfrak{T} = \langle T, (\text{succ}_i)_{i < d} \rangle$ be a finite successor tree where every vertex has at most d successors and let δ be a k -sparse distribution on \mathfrak{T} . Then \mathfrak{T} has a δ' -flow f , for some $\delta' \geq \delta$, such that f is WMSO-definable (without parameters) and edge-bounded by $dk + 1$. Furthermore, for every edge $\langle u, v \rangle$ of \mathfrak{T} , there is some connected set $Z \subseteq T$ with minimal element v such that*

$$f(u, v) = |Z| + k \cdot |B_{\mathfrak{T}}(Z)| - \delta(Z) - k.$$

Proof. The flow f we define below has the property that, for every $v \in T$ with predecessor u , there is some connected set $Z_v \subseteq T$ with minimal element v such that

$$\begin{aligned} f(u, v) &= |Z_v| + k \cdot |B_{\mathfrak{T}}(Z_v)| - \delta(Z_v) - k, \\ f(x, y) &\geq k, \quad \text{for all } \langle x, y \rangle \in B_{\mathfrak{T}}(Z_v) \setminus \{\langle v, u \rangle\}. \end{aligned}$$

We proceed by induction on the size of the subtree attached at v . Hence, let $\langle u, v \rangle$ be an edge such that f is already defined for all edges of the subtree attached at v , let w_0, \dots, w_{n-1} be the successors of v , and set $I := \{i < n \mid f(v, w_i) < k\}$. We set

$$Z_v := \{v\} \cup \bigcup_{i \in I} Z_{w_i}.$$

Then $f(u, v)$ is given by the formula above.

This concludes the definition of f . To see that f has the desired properties, we start by computing its defect. Let v be a vertex a predecessor u and

successors w_0, \dots, w_{n-1} , and set $I := \{i < n \mid f(v, w_i) < k\}$. Then

$$\begin{aligned}
 d_f(v) &= \sum_{i < n} f(v, w_i) - f(u, v) \\
 &= \sum_{i < n} f(v, w_i) - [|Z_v| + k \cdot |B_{\mathfrak{Z}}(Z_v)| - \delta(Z_v) - k] \\
 &= \sum_{i < n} f(v, w_i) - \left[1 + \sum_{i \in I} |Z_{w_i}| + \left[\delta(v) + \sum_{i \in I} \delta(Z_{w_i})\right] + k \right. \\
 &\quad \left. - k \cdot \left[1 + \sum_{i \in I} [|B_{\mathfrak{Z}}(Z_{w_i})| - 1] + (n - |I|)\right] \right] \\
 &= \sum_{i < n} f(v, w_i) - 1 + \delta(v) + k - k - k(n - |I|) \\
 &\quad - \sum_{i \in I} [|Z_{w_i}| - \delta(Z_{w_i}) + k \cdot |B_{\mathfrak{Z}}(Z_{w_i})| - k] \\
 &= \sum_{i < n} f(v, w_i) - 1 + \delta(v) + k - k - k(n - |I|) - \sum_{i \in I} f(v, w_i) \\
 &= \delta(v) - 1 - k(n - |I|) + \sum_{i \notin I} f(v, w_i).
 \end{aligned}$$

To see that f is edge-bounded by $dk + 1$ note that, by sparseness of δ , we have

$$f(u, v) = |Z_v| + k \cdot |B_{\mathfrak{Z}}(Z_v)| - \delta(Z_v) - k \geq -k.$$

Conversely, we have

$$\begin{aligned}
 f(u, v) &= \sum_{i < n} f(v, w_i) - d_f(v) \\
 &= \sum_{i < n} f(v, w_i) - \delta(v) + 1 + k(n - |I|) - \sum_{i \notin I} f(v, w_i) \\
 &= \sum_{i \in I} f(v, w_i) - \delta(v) + 1 + k(n - |I|) \\
 &\leq (k - 1) \cdot |I| - \delta(v) + 1 + k(n - |I|) \\
 &= kn + 1 + (k - 1) \cdot |I| - \delta(v) - k \cdot |I| \\
 &\leq kd + 1.
 \end{aligned}$$

To find the desired distribution δ' , we can use Lemma 4.5. To do so we have to show that $d_f(v) \geq \delta(v) - 1$. If v is different from the root, we have

$$\begin{aligned}
 d_f(v) + 1 &= \delta(v) - 1 - k(n - |I|) + \sum_{i \notin I} f(v, w_i) + 1 \\
 &= \delta(v) + \sum_{i \notin I} [f(v, w_i) - k] \\
 &\geq \delta(v) + 0.
 \end{aligned}$$

For the root v , let

$$Z' := \{v\} \cup \bigcup_{i < n} Z_{w_i},$$

where w_0, \dots, w_{n-1} are the successors of v . Then sparsity of δ implies that

$$\begin{aligned}
 d_f(v) + 1 &= \sum_{i < n} f(v, w_i) + 1 \\
 &= \sum_{i < n} [|Z_{w_i}| + k \cdot |B_{\mathfrak{Z}}(Z_{w_i})| - \delta(Z_{w_i}) - k] + 1 \\
 &= |Z'| - 1 + k \cdot [|B_{\mathfrak{Z}}(Z')| + n] - \delta(Z') + \delta(v) - nk + 1 \\
 &= |Z'| + k \cdot |B_{\mathfrak{Z}}(Z')| - \delta(Z') - 1 + nk + \delta(v) - nk + 1 \\
 &= |Z'| + k \cdot |B_{\mathfrak{Z}}(Z')| - \delta(Z') + \delta(v) \\
 &\geq \delta(v).
 \end{aligned}$$

It remains to show that f is WMSO-definable. Since f is edge-bounded, an WMSO-formula can guess the values $f(u, v)$, for every edge $\langle u, v \rangle$, and then check that the guessed values are correct using the relation

$$f(u, v) = \sum_{i \in I} f(v, w_i) - \delta(v) + 1 + k(n - |I|)$$

from above. \square

We can extend this result to countable trees with a slightly worse bound.

Proposition 4.9. Let $\mathfrak{T} = \langle T, (\text{succ}_i)_{i < d} \rangle$ be a successor tree where every vertex has at most d successors and let δ be a k -sparse distribution on \mathfrak{T} . Then \mathfrak{T} has a δ' -flow f , for some $\delta' \geq \delta$, such that f is WMSO-definable (without parameters) and edge-bounded by $(dk + 1)(d + 2)$.

Proof. We start to define the desired flow f on all finite subtrees of \mathfrak{T} . We would like to Lemma 4.8 for this. Unfortunately, the restriction of δ to a subtree of \mathfrak{T} does not need to be k -sparse. Given a finite subtree \mathfrak{S} of \mathfrak{T} , we therefore first add a path of length k to the root of \mathfrak{S} and we set $\delta(v) := 0$, for every of the new vertices v . The resulting distribution is k -sparse. Let f' be the flow obtained by applying Lemma 4.8 to it. For every edge $\langle u, v \rangle$ of \mathfrak{S} (including the edge connecting the root \mathfrak{S} to the rest of \mathfrak{T}), we set $f(u, v) := f'(u, v)$. Note that this part of f is edge-bounded by $dk + 1$.

Hence, it remains to define $f(u, v)$ for edges $\langle u, v \rangle$ belonging to some infinite branch of \mathfrak{T} . Let B be the union of all infinite branches of \mathfrak{T} . Hence, it remains to define $f(u, v)$ for $u, v \in B$. We do so by induction on the distance of u from the root. Hence, let $v \in B$ a vertex with successors w_0, \dots, w_{n-1} and set

$$I := \{ i < n \mid w_i \in B \}.$$

If v has a predecessor u , we may assume by inductive hypothesis that $f(u, v)$ is already defined. Let

$$s := \delta(v) - 1 - \sum_{i \notin I} f(v, w_i) + \begin{cases} f(u, v) & \text{if } u \text{ exists,} \\ 0 & \text{otherwise,} \end{cases}$$

set

$$c := \left\lfloor \frac{\max\{s, 0\}}{|I|} \right\rfloor \quad \text{and} \quad c' := \left\lceil \frac{\max\{s, 0\}}{|I|} \right\rceil,$$

and let $0 \leq m < |I|$ be the constant such that $\max\{s, 0\} = c \cdot |I| + m$. For $i \in I$, we set

$$f(v, w_i) := \begin{cases} c & \text{if there are less than } m \text{ indices } j < i \text{ in } I, \\ c' & \text{otherwise.} \end{cases}$$

Then

$$\begin{aligned} d_f(v) &= (\delta(v) + 1) - s + \sum_{i \in I} f(v, w_i) \\ &= \delta(v) + 1 - s + c \cdot |I| + m \\ &\geq \delta(v) + 1. \end{aligned}$$

In particular, it follows by Lemma 4.5 that f is a δ' -flow for some $\delta' \geq \delta$. Furthermore, f is WMSO-definable since we can express in WMSO that the subtree attached to a given vertex is finite, the definition for finite subtrees is WMSO-definable, and the values of f for vertices in B can be computed inductively starting at the root.

Since $f(v, w_i) \geq 0$ it therefore remains to show that

$$f(v, w_i) \leq (dk + 1)(d + 2).$$

Hence, consider an edge $\langle u, v \rangle$ in B . If u has $m \geq 2$ successors in B , we have (with the notation from above)

$$\begin{aligned} f(u, v) &\leq c' \leq \frac{1}{m} \cdot \max\{s, 0\} + 1 \\ &\leq \frac{1}{2} \cdot \max\{s, 0\} + 1 \\ &\leq \frac{1}{2} [\delta(v) - 1 + (dk + 1)(n - m) + (dk + 1)(d + 2)] \\ &\leq \frac{1}{2} [(d + 1)k + 1 - 1 + (dk + 1)(n - m) \\ &\quad + (dk + 1)(d + 2)] \\ &\leq \frac{1}{2} [(d + 1)k + (dk + 1)(n - m) + (dk + 1)(d + 2)] \\ &\leq \frac{1}{2} [(d + 1)k + (dk + 1)(d - 2) + (dk + 1)(d + 2)] \\ &= \frac{1}{2} [(d + 1)k + (dk + 1)2d] \end{aligned}$$

$$\begin{aligned}
&= (dk+1)d + \frac{1}{2}(d+1)k \\
&= (dk+1)(d+2) - 2(dk+1) + \frac{1}{2}(d+1)k \\
&= (dk+1)(d+2) - \left[2dk+2 - \frac{1}{2}dk - \frac{1}{2}k\right] \\
&= (dk+1)(d+2) - \left[\frac{3d-1}{2}k+2\right] \\
&\leq (dk+1)(d+2) - \left[\frac{5}{2}k+2\right] \\
&\leq (dk+1)(d+2) - 2k.
\end{aligned}$$

Hence, it remains to consider the case where v is the only successor of u that belongs to B . Let u_0, \dots, u_m be the longest suffix of the path from the root to u such that $f(u_j, u_{j+1}) > (dk+1)(d+2) - 2k$, for all i . Let $w_0^j, \dots, w_{n_j-1}^j$ be the successors of u_j different from u_{j+1} . By the case above, it follows that u_{j+1} is the unique successor of u_j that belongs to B . For $i < n_j$, let Z_i^j be the set such that

$$\begin{aligned}
f(u_j, w_i^j) &= |Z_i^j| + k \cdot |B_{\mathfrak{Z}}(Z_i^j)| - \delta(Z_i^j) - k, \\
f(x, y) &\geq k, \quad \text{for all } \langle x, y \rangle \in B_{\mathfrak{Z}}(Z_i^j) \setminus \{\langle w_i^j, u_j \rangle\}.
\end{aligned}$$

Let

$$Z := \{u_0, \dots, u_m\} \cup \bigcup_{\substack{j \leq m \\ i \in I_j}} Z_i^j.$$

Let u_{-1} be the predecessor of u_0 , if it exists, and set $f(u_{-1}, u_0) := 0$, in case

it does not. By definition of f , we have

$$\begin{aligned}
f(u_m, v) &= \max\{s, 0\} \\
&= \delta(u_m) - 1 - \sum_{i < n_m} f(u_m, w_i^m) + f(u_{m-1}, u_m) \\
&= \dots \\
&= \sum_{j=0}^m \left[\delta(u_j) - 1 - \sum_{i < n_j} f(u_j, w_i^j) \right] + f(u_{-1}, u_0) \\
&= \sum_{j=0}^m \left[\delta(u_j) - 1 - \sum_{i < n_j} \left[|Z_i^j| + k \cdot |B_{\mathfrak{Z}}(Z_i^j)| - \delta(Z_i^j) - k \right] \right] \\
&\quad + f(u_{-1}, u_0) \\
&= \delta(Z) - |Z| - \sum_{j=0}^m \sum_{i < n_j} k \cdot |B_{\mathfrak{Z}}(Z_i^j)| + k(n_0 + \dots + n_m) \\
&\quad + f(u_{-1}, u_0) \\
&= \delta(Z) - |Z| - k \cdot [|B_{\mathfrak{Z}}(Z)| - 2] + f(u_{-1}, u_0) \\
&\leq |Z| + k \cdot |B_{\mathfrak{Z}}(Z)| - |Z| - k \cdot [|B_{\mathfrak{Z}}(Z)| - 2] + f(u_{-1}, u_0) \\
&= f(u_{-1}, u_0) + 2k \\
&\leq (dk+1)(d+2) - 2k + 2k = (dk+1)(d+2). \quad \square
\end{aligned}$$

It remains to show how we can use the δ -flow f we have just constructed to define the desired function $g : V \rightarrow V$. We start by selecting a certain family of definable paths. Note that we allow paths of length 0. Such paths are uniquely determined by the vertex they start (and end) at.

Lemma 4.10. *Let \mathfrak{G} be a countable undirected graph and f an acyclic δ -flow of \mathfrak{G} . There exists a set \mathcal{P} of finite directed paths through \mathfrak{G} satisfying the following conditions:*

- (I) *For every $v \in V$, there are exactly $\delta(v)$ paths in \mathcal{P} starting at v .*
- (II) *For every $v \in V$ there is at most one path in \mathcal{P} ending at v .*

(III) For every pair $u, v \in V$ of vertices there are at most $\max\{0, f(u, v)\}$ paths in \mathcal{P} containing the edge $\langle u, v \rangle$ (in this direction).

Proof. Fix an enumeration $\langle v_n, k_n \rangle_{n < \omega}$ of the set

$$\{ \langle v, k \rangle \mid v \in V, 0 \leq k < \delta(v) \}.$$

For $n < \omega$, we construct paths π_n with the following properties:

- ♦ π_n starts at v_n .
- ♦ The endpoints of π_m and π_n are different for $m \neq n$.
- ♦ For every edge $\langle u, v \rangle$ there are at most $f(u, v)$ paths π_n containing the edge $\langle u, v \rangle$.

By induction, suppose that we have already defined π_i , for $i < n$. Let

$$\begin{aligned} \alpha(v) &:= \left| \{ i < n \mid \pi_i \text{ starts at } v \} \right|, \\ \beta(v) &:= \left| \{ i < n \mid \pi_i \text{ ends at } v \} \right|, \\ \mu(u, v) &:= \left| \{ i < n \mid \pi_i \text{ contains the edge } \langle u, v \rangle \} \right|. \end{aligned}$$

We construct a path $u_0 \dots u_m$ inductively starting with $u_0 := v_n$. For the inductive step, suppose that we have already defined u_0, \dots, u_i . If $\beta(u_i) = 0$, we stop and set $\pi_n := u_0 \dots u_i$. Otherwise, we claim that there is some neighbour w of u_i with $f(u_i, w) > \mu(u_i, w)$. Hence, we can set $u_{i+1} := w$.

To prove the claim, we distinguish two cases. If $i = 0$, then $\alpha(u_0) < \delta(u_0)$ implies that

$$\begin{aligned} \sum_{x \in V} \mu(u_0, x) &= \alpha(u_0) - \beta(u_0) + \sum_{x \in V} \mu(x, u_0) \\ &\leq \alpha(u_0) - 1 + \sum \{ f(x, u_0) \mid f(x, u_0) \geq 0 \} \\ &= \alpha(u_0) - 1 - d_f(u_0) + \sum \{ f(u_0, x) \mid f(u_0, x) \geq 0 \} \\ &= \alpha(u_0) - 1 - (\delta(u_0) - 1) \\ &\quad + \sum \{ f(u_0, x) \mid f(u_0, x) \geq 0 \} \\ &< \sum \{ f(u_0, x) \mid f(u_0, x) \geq 0 \}, \end{aligned}$$

as desired. Similarly, if $i > 0$ then $\mu(u_{i-1}, u_i) < f(u_{i-1}, u_i)$ implies that

$$\begin{aligned} \sum_{x \in V} \mu(u_i, x) &= \alpha(u_i) - \beta(u_i) + \sum_{x \in V} \mu(x, u_i) \\ &< \alpha(u_i) - 1 + \sum \{ f(x, u_i) \mid f(x, u_i) \geq 0 \} \\ &= \alpha(u_i) - 1 - d_f(u_i) + \sum \{ f(u_i, x) \mid f(u_i, x) \geq 0 \} \\ &= \alpha(u_i) - 1 - (\delta(u_i) - 1) \\ &\quad + \sum \{ f(u_i, x) \mid f(u_i, x) \geq 0 \} \\ &\leq \sum \{ f(u_i, x) \mid f(u_i, x) \geq 0 \}. \end{aligned}$$

Note that the construction of π_n must terminate after at most $n + 1$ steps since the flow f is acyclic and there are only n vertices u with $\beta(u) = 1$. \square

Lemma 4.11. Let $m < \omega$. There exists an MSO-formula $\varphi(X; \bar{Z})$ with the following property. For every graph \mathfrak{G} and each set \mathcal{P} of finite paths such that every vertex and every edge of \mathfrak{G} is contained in at most m paths of \mathcal{P} , there exists a tuple \bar{S} of monadic parameters such that

$$\mathfrak{G}_{\text{in}} \models \varphi(P; \bar{S}) \quad \text{iff} \quad P \text{ is (the set of edges of) a non-empty path in } \mathcal{P}.$$

Proof. For every edge $e = \langle u, v \rangle$ of \mathfrak{G} , we fix a bijection $\mu(u, v) : [n] \rightarrow \mathcal{P}_e$ where $\mathcal{P}_e \subseteq \mathcal{P}$ is the set of all paths containing the edge $\langle u, v \rangle$ (in either direction) and $n := |\mathcal{P}_e|$. We assume that $\mu(u, v) = \mu(v, u)$.

Let S be the set of all edges of \mathfrak{G} contained in some path of \mathcal{P} . By Corollary 1.18, there exists an MSO-formula $\chi(x, y, z; \bar{S}')$ with parameters \bar{S}' such that, for every $v \in V$, the formula $\chi(x, y, v; \bar{S}')$ linearly orders the set of all vertices that are connected to v via an edge in S .

Finally, we define unary predicates Q_{ji}^{ik} containing all vertices v such that there exists a path $\pi \in \mathcal{P}$ containing edges $\langle u, v \rangle$ and $\langle v, w \rangle$ where

- ♦ $\mu(u, v)(k) = \pi = \mu(v, w)(l)$,
- ♦ u is the i -th neighbour of v (in the order defined by χ),
- ♦ w is the j -th neighbour of v .

It follows that a non-empty set $P \subseteq E$ of edges is a path in \mathcal{P} if, and only if, P is a minimal non-empty subset of E satisfying the following condition:

P can be written as a union $P = P_0 \cup \dots \cup P_{m-1}$ such that, for all vertices u, v, w such that $v \in Q_{jl}^{ik}$ and u and w are, respectively, the i -th and j -th neighbour of v , we have $\langle u, v \rangle \in P_k \Leftrightarrow \langle v, w \rangle \in P_l$.

This condition can be expressed in GSO. \square

Remark. Note that the set of empty paths in \mathcal{P} is trivially definable with the help of the parameter

$$Q := \{v \in V \mid \mathcal{P} \text{ contains an empty path from } v \text{ to } v\}.$$

Using the family \mathcal{P} we can construct a formula φ defining the function g .

Theorem 4.12. *Let $m < \omega$. There exists an MSO-formula $\varphi(x, y; \bar{Z})$ with the following property. For every graph $\mathfrak{G} = \langle V, E \rangle$ and each acyclic δ -flow f of \mathfrak{G} that is vertex-bounded by m , there exist monadic parameters \bar{S} such that $\varphi(x, y; \bar{S})$ defines in \mathfrak{G}_{in} a partial function $g : V \rightarrow V$ with*

$$|g^{-1}(v)| = \delta(v), \quad \text{for all } v \in V.$$

Proof. Let \mathfrak{G}' be the graph obtained from \mathfrak{G} by removing every edge $\langle u, v \rangle$ with $f(u, v) = 0$. Note that f is also a δ -flow of \mathfrak{G}' . Since f is vertex-bounded by m it follows that every vertex of \mathfrak{G}' has degree at most $m < \omega$. Consequently, each connected component \mathfrak{G}_o of \mathfrak{G} is countable. For each such component \mathfrak{G}_o , let \mathcal{P}_o be the set of paths obtained by applying Lemma 4.10 to the restriction of f to \mathfrak{G}_o . By Lemma 4.11, there exists a formula $\psi(X; \bar{Z})$ and a set \bar{S} of guarded relations such that

$$\mathfrak{G} \models \psi(P; \bar{S}) \quad \text{iff} \quad P \text{ is a non-empty path in } \mathcal{P},$$

where \mathcal{P} is the union of all sets \mathcal{P}_o corresponding to the connected components of \mathfrak{G}' . With the help of ψ we can define a partial function $g : V \rightarrow V$ such that

$$g(v) = u \quad : \text{iff} \quad \mathcal{P} \text{ contains a path from } u \text{ to } v.$$

By construction of \mathcal{P} , we have $|g^{-1}(v)| = \delta(v)$, for every $v \in V$. \square

Again, for the application below, we need a version of these results where the set of paths is definable without parameters.

Theorem 4.13. *Let $d, k < \omega$ and let $(\psi_i)_{i < m}$ be a family of WMSO-formulae. There exists an WMSO-formula $\varphi(x, y)$ with the following property. Given a successor tree $\mathfrak{T} = \langle T, (\text{suc}_i)_{i < d} \rangle$ such that*

- ♦ every vertex has at most d successors and,
- ♦ the formulae $(\psi_i)_i$ define a k -sparse distribution δ on \mathfrak{T} ,

the formula $\varphi(x, y)$ defines in \mathfrak{T} a partial function $g : T \rightarrow T$ with

$$|g^{-1}(v)| = \delta(v), \quad \text{for all } v \in V.$$

Proof. Let f be the WMSO-definable δ' -flow from Proposition 4.9. We proceed as in the above proofs: we use f to construct a set of paths \mathcal{P} , we show that \mathcal{P} is definable, and we take for g the function mapping the end-point of each path to its starting point.

To simplify the construction of \mathcal{P} we modify the tree \mathfrak{T} as follows. We delete every edge $\langle u, v \rangle$ with $f(u, v) = 0$ and we direct the remaining edges by

$$u \rightarrow v \quad \text{iff} \quad f(u, v) > 0.$$

Finally, we duplicate each of the resulting edges $f(u, v)$ times. Let \mathfrak{S} be the resulting directed graph. Note that the structure \mathfrak{S}_{in} is WMSO-interpretable in $\text{copy}_m(\mathfrak{T})$ using the formulae defining f , where $m := (dk+1)(d+2)$ is the edge-bound of f . Consequently, it is sufficient to find an WMSO-definable set of finite paths \mathcal{P} in \mathfrak{S} such that

- ♦ every edge of \mathfrak{S} belongs to exactly one path in \mathcal{P} ,
- ♦ for every vertex v , there are exactly $\delta'(v)$ paths in \mathcal{P} starting at v , and
- ♦ for every vertex v , there is at most 1 path in \mathcal{P} ending at v .

We order the set of neighbours of each vertex v of \mathfrak{S} by starting with its predecessor followed by its successors $\text{suc}_0(v), \text{suc}_1(v), \dots$. This ordering induces, for each vertex of \mathfrak{S} , an ordering of the incoming edges and an

ordering of the outgoing ones. To define \mathcal{P} it is now sufficient to specify, for each incoming edge e , whether the path containing e stops at the current vertices, or which of the outgoing edges it continues with. We have to do the same with the $\delta'(v)$ paths starting at v .

We proceed as follows.

- ◆ The first incoming path stops at the given vertex.
- ◆ The $(i + 1)$ -th incoming path continues with the i -th outgoing edge.
- ◆ The i -th path starting at v leaves with the i -th not-yet-used outgoing edge.

These conditions can obviously be expressed in WMSO. It thus remains to check that this definition results in a set of finite paths. First note that \mathfrak{G} is acyclic since it originates from a tree. Thus, \mathcal{P} is indeed a set of paths. To see that each of them is finite, note that every path P in \mathfrak{G} starts with a (possibly empty) prefix of edges directed towards the root followed by a (empty, finite, or infinite) suffix directed away from the root. Let u_o, u_1, \dots be this suffix and suppose that P uses the i_n -th edge from u_n to u_{n+1} . Then $i_{n+1} = i_n - 1$, which implies that P stops at u_{i_o+1} . \square

Let us show how we can use these results to encode sets by single elements.

Definition 4.14. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph and $h : X \rightarrow V$ a function. The distribution $\delta : V \rightarrow \omega$ induced by h is given by

$$\delta(v) := |h^{-1}(v)|, \quad \text{for } v \in V.$$

Again we present a general version for arbitrary graphs and definability with parameters, and a special version for trees where we do not need parameters.

Theorem 4.15. Let $d, k < \omega$ and let $\varphi(X, y; \bar{Z})$ be an MSO-formula. There exist MSO-formulae $\psi(x, y; \bar{Z}')$ and $\vartheta(X, y; \bar{Z}')$ with the following property. Given a connected graph $\mathfrak{G} = \langle V, E \rangle$ of degree at most d and monadic parameters \bar{S} such that

- ◆ $\varphi(X, y; \bar{S})$ defines (in \mathfrak{G}_{in}) a partial function $h : \wp_{\text{fin}}(V) \rightarrow V$, and

- ◆ the distribution δ induced by h is k -sparse,

there exist monadic parameters \bar{S}' such that the formulae $\psi(x, y; \bar{S}')$ and $\vartheta(X, y; \bar{S}')$ define (in \mathfrak{G}) partial functions $g : V \rightarrow V$ and $h_o : \wp_{\text{fin}}(V) \rightarrow V$ such that

$$h = g \circ h_o \quad \text{and} \quad h_o \text{ is injective.}$$

Proof. We start by constructing an MSO-formula $\chi(X, Y; z)$ (with parameters) such that $\chi(X, Y; v)$ linearly orders the set $h^{-1}(v)$, for $v \in V$. Let $T_o \subseteq E$ be a spanning tree of \mathfrak{G} (in the usual graph-theoretic sense) and let $r \in V$ be some vertex acting as the root. Using the parameters T_o and r we can define the forest ordering on V by

$$u \leq v \quad : \text{iff} \quad \text{the unique path in } T_o \text{ between } r \text{ and } v \text{ contains } u.$$

Let $T \subseteq V \times V$ be the set obtained from T_o by orienting the edges according to this ordering. Then T is a directed successor tree. Furthermore, since the degree of \mathfrak{G} is bounded by d , we can use Corollary 1.18 (applied to the inverse of T) to linearly order the successors of every vertex in T . We use these two orderings to define the lexicographic ordering \leq_{lex} on T . Finally, we obtain the desired ordering on $\wp_{\text{fin}}(V)$ by setting

$$X < Y \quad : \text{iff} \quad \begin{array}{l} \text{the } \leq_{\text{lex}}\text{-minimal element of } (X \setminus Y) \cup (Y \setminus X) \\ \text{belongs to } Y. \end{array}$$

Each of these definitions can be expressed in MSO.

Having defined χ , we now use Proposition 4.7 to obtain a δ -flow f that is edge-bounded by k . Since \mathfrak{G} has degree at most d , it follows that f is vertex-bounded by dk . Hence, Theorem 4.12 produces a definable function $g : V \rightarrow V$ with $|g^{-1}(v)| = \delta(v) = |h^{-1}(v)|$. Fix unary predicates P_o, \dots, P_{k-1} such that we have $i \neq j$ whenever $u \in P_i$ and $v \in P_j$ are distinct vertices with $g(u) = g(v)$. Using these predicate we can define partial functions $g_o, \dots, g_{k-1} : V \rightarrow V$ such that $g_i(v)$ is the unique element of $g^{-1}(v) \cap P_i$. We define $h_o : \wp(V) \rightarrow V$ by $h_o(X) := (g_i \circ h)(X)$, where the index i is chosen such that X is the i -th element of $h^{-1}(h(X))$ (in the order defined

by χ). It follows that $h(X) = g(h_o(X))$ and h_o is injective. Furthermore, the function h_o is clearly MSO-definable in \mathfrak{G}_{in} . As the graph \mathfrak{G} has degree at most d , it is d -sparse. Therefore, we can translate the corresponding formulae from \mathfrak{G}_{in} to \mathfrak{G} . \square

The version without parameters looks as follows.

Theorem 4.16. *Let $d, k < \omega$ and let $\varphi(X, y)$ be a WMSO-formula. There exist WMSO-formulae $\psi(x, y)$ and $\vartheta(X, y)$ with the following property. Given a successor tree $\mathfrak{T} = \langle T, (\text{succ}_i)_{i < d} \rangle$ such that*

- ♦ *every vertex has at most d successors,*
- ♦ *$\varphi(X, y)$ defines a partial function $h : \wp_{\text{fin}}(V) \rightarrow V$, and*
- ♦ *the distribution δ induced by h is k -sparse,*

the formulae $\psi(x, y)$ and $\vartheta(X, y)$ define partial functions $g : V \rightarrow V$ and $h_o : \wp_{\text{fin}}(V) \rightarrow V$ such that

$$h = g \circ h_o \quad \text{and} \quad h_o \text{ is injective.}$$

Proof. We can use Theorem 4.13 to construct a WMSO-definable function $g : V \rightarrow V$ with $|g^{-1}(v)| = \delta(v) = |h^{-1}(v)|$. Let \leq_{lex} be the lexicographic ordering on T . Since this ordering is FO-definable, we can define partial functions $g'_0, \dots, g'_{k-1} : V \rightarrow V$ such that $g'_i(v)$ is the i -th element of $g^{-1}(v)$ with respect to \leq_{lex} .

Furthermore, we can use \leq_{lex} to define a linear order on $\wp_{\text{fin}}(T)$ by setting

$$X <_{\text{set}} Y \quad : \text{iff} \quad \text{the } \leq_{\text{lex}}\text{-minimal element of } (X \setminus Y) \cup (Y \setminus X) \text{ belongs to } Y.$$

We define $h_o : \wp(V) \rightarrow V$ by $h_o(X) := (g'_i \circ h)(X)$, where the index i is chosen such that X is the i -th element of $h^{-1}(h(X))$ (in the order defined by \leq_{set}). It follows that $h(X) = g(h_o(X))$ and h_o is injective. Furthermore, the function h_o is WMSO-definable. \square

The Finite Power-Set Operation

As an application let us prove that the finite power-set operation \wp_{fin} on structures commutes with interpretations in the following sense.

Theorem 4.17. *For every FO-interpretation τ , there exists an WMSO-interpretation σ such that*

$$\tau(\wp_{\text{fin}}(\mathfrak{T})) \cong \wp_{\text{fin}}(\mathfrak{A}) \quad \text{implies} \quad \sigma(\mathfrak{T}) \cong \mathfrak{A},$$

for every tree \mathfrak{T} and every structure \mathfrak{A} .

$$\begin{array}{ccc} \wp_{\text{fin}}(\mathfrak{T}) & \xrightarrow{\tau} & \wp_{\text{fin}}(\mathfrak{A}) \\ \wp_{\text{fin}} \uparrow & & \uparrow \wp_{\text{fin}} \\ \mathfrak{T} & \xrightarrow{\sigma} & \mathfrak{A} \end{array}$$

We split the proof into several steps. Let us start with a bit of terminology and some conventions. First of all, note that the FO-formulae the interpretation τ evaluates in $\wp_{\text{fin}}(\mathfrak{T})$ can be translated into WMSO-formulae that can be evaluated directly in \mathfrak{T} . Consequently, we will be interested in structures of the form $\wp_{\text{fin}}(\mathfrak{A})$ that can be defined inside \mathfrak{T} using WMSO-formulae (encoding each element of $\wp_{\text{fin}}(\mathfrak{A})$ by a *finite subset* of T). In the following, whenever we write $\mathfrak{T} \models \varphi(P)$ or $P \in \wp^{\mathfrak{T}}$, for some WMSO-formula $\varphi(X)$, we always tacitly assume that the set P is finite.

Definition 4.18. Let \mathfrak{T} be a structure and $\text{sing}(X)$, $\text{in}(X, Y)$ two WMSO-formulae. We say that the pair $\langle \text{sing}, \text{in} \rangle$ *encodes a finite power-set* on \mathfrak{T} if

- ♦ $\mathfrak{T} \models \text{in}(P, Q)$ implies $\mathfrak{T} \models \text{sing}(P)$
- ♦ for every finite $H \subseteq \text{sing}^{\mathfrak{T}}$, there exists exactly one finite set $Q \subseteq T$ with

$$\mathfrak{T} \models \text{in}(P, Q) \quad \text{iff} \quad P \in H.$$

Given a tree \mathfrak{T} and a vertex v , recall that $\mathfrak{T}[v]$ denotes the tree obtained from \mathfrak{T} by removing the subtree rooted at v and $\mathfrak{T}|_v$ denotes this subtree. Finally, $\uparrow v$ is the set of vertices of $\mathfrak{T}|_v$.

Given a tree \mathfrak{T} and a vertex $v \in T$, we are interested in which subsets of the part of \mathfrak{T} ‘above’ v can be completed to some set satisfying a given formula $\text{sing}(X)$, and which subsets in the part ‘below’ v . In addition, we will have to distinguish these subsets by their theory. This leads to the following definition.

Definition 4.19. Let \mathfrak{T} be a tree, $v \in T$ a vertex with predecessor v' , $\text{sing}(X) \in \text{WMSO}_m$ a formula, and σ, τ two WMSO_m -theories.

(a) We set

$$\begin{aligned} \mathcal{A}(v, \sigma) &:= \{ P \setminus \uparrow v \mid \mathfrak{T} \models \text{sing}(P) \text{ and } \\ &\quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}[v], P \setminus \uparrow v, v') = \sigma \}, \\ \mathcal{B}(v, \tau) &:= \{ P \cap \uparrow v \mid \mathfrak{T} \models \text{sing}(P) \text{ and } \\ &\quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_v, P \cap \uparrow v, v) = \tau \}. \end{aligned}$$

(b) We say that the pair $\langle \sigma, \tau \rangle$ *implies* $\text{sing}(X)$ if

$$\mathfrak{T} \models \text{sing}(P \cup Q), \quad \text{for } P \in \mathcal{A}(v, \sigma) \text{ and } Q \in \mathcal{B}(v, \tau).$$

We denote this fact by $\langle \sigma, \tau \rangle \models \text{sing}$.

Remark. Note that the theories

$$\text{Th}_{\text{WMSO}}^m(\mathfrak{T}[v], P \setminus \uparrow v, v') \quad \text{and} \quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_v, P \cap \uparrow v, v)$$

uniquely determine $\text{Th}_{\text{WMSO}}^m(\mathfrak{T}, P)$ since there exists a quantifier-free interpretation τ such that

$$\langle \mathfrak{T}, P \rangle = \tau(\langle \mathfrak{T}[v], P \setminus \uparrow v, v' \rangle \oplus \langle \mathfrak{T}|_v, P \cap \uparrow v, v \rangle).$$

In particular, the definition of the relation $\langle \sigma, \tau \rangle \models \text{sing}(X)$ does not depend on \mathfrak{T} or v .

Our main combinatorial lemma consists in the observation that, for every decomposition of the tree, the above set or the below set has to be small. This is a direct consequence of the properties of the formula $\text{in}(X, Y)$.

Lemma 4.20. Let $\text{sing}, \text{in} \in \text{WMSO}_m$. There exists a constant k_o with the following property. Given a tree \mathfrak{T} such that $\langle \text{sing}, \text{in} \rangle$ encodes a finite power-set on \mathfrak{T} , we have

$$\langle \sigma, \tau \rangle \models \text{sing} \quad \text{implies} \quad |\mathcal{A}(v, \sigma)| < k_o \quad \text{or} \quad |\mathcal{B}(v, \tau)| < k_o,$$

for all $v \in T$ and all WMSO_m -theories σ and τ .

Proof. Let Θ be the set of all WMSO_m -theories with two set variables and one first-order variable, $s := |\Theta|$, and set $k_o := 2s + 1$. For a contradiction suppose that $|\mathcal{A}(v, \sigma)| \geq k_o$ and $|\mathcal{B}(v, \tau)| \geq k_o$. Fix two subsets $\mathcal{A}_o \subseteq \mathcal{A}(v, \sigma)$ and $\mathcal{B}_o \subseteq \mathcal{B}(v, \tau)$ of size $|\mathcal{A}_o| = 2s + 1$ and $|\mathcal{B}_o| = 2s + 1$, and let S be the set of all sets $W \subseteq T$ encoding some union of these singletons $A \cup B \in \text{sing}^{\mathfrak{T}}$. That is, $W \in S$ if, and only if,

$$\mathfrak{T} \models \text{in}(U, W) \quad \text{implies} \quad U = A \cup B, \quad \text{for some } A \in \mathcal{A}_o, B \in \mathcal{B}_o.$$

With each $W \in S$ we associate the two functions $f_W : \mathcal{A}_o \rightarrow \Theta$ and $g_W : \mathcal{B}_o \rightarrow \Theta$ defined by

$$\begin{aligned} f_W(A) &:= \text{Th}_{\text{WMSO}}^m(\mathfrak{T}[v], A, W \setminus \uparrow v, v'), \\ g_W(B) &:= \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_v, B, W \cap \uparrow v, v). \end{aligned}$$

For two sets $V, W \in S$ it follows that

$$\begin{aligned} \langle f_V, g_V \rangle &= \langle f_W, g_W \rangle \\ \Rightarrow \quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}, A \cup B, V) &= \text{Th}_{\text{WMSO}}^m(\mathfrak{T}, A \cup B, W), \\ &\quad \text{for all } A \in \mathcal{A}_o \text{ and } B \in \mathcal{B}_o, \\ \Rightarrow \quad \mathfrak{T} \models \text{in}(A \cup B, V) &\leftrightarrow \text{in}(A \cup B, W), \\ &\quad \text{for all } A \in \mathcal{A}_o \text{ and } B \in \mathcal{B}_o, \\ \Rightarrow \quad V &= W. \end{aligned}$$

Since there are only $|\Theta|^{|\mathcal{A}_o|} \cdot |\Theta|^{|\mathcal{B}_o|} = s^{2(2s+1)}$ pairs $\langle f, g \rangle$ of such functions, this implies that

$$2^{(2s+1)^2} = |S| \leq s^{2(2s+1)} = 2^{(2s+1) \cdot 2 \log s} < 2^{(2s+1)^2}.$$

A contradiction. \square

When we want to encode the sets $P \in \text{sing}^{\mathfrak{T}}$ by single vertices, we can mostly ignore subtrees $\mathfrak{T}|_v$ of \mathfrak{T} such that there are few choices for the values of $P \cap T|_v$ since most information about P is already contained in the value $P \setminus T|_v$. In the following we will therefore concentrate on the prefix of \mathfrak{T} obtained by removing all these subtrees.

Definition 4.21. Suppose that $\langle \text{sing}, \text{in} \rangle$ encodes a finite power-set on a tree \mathfrak{T} and let $k < \omega$. We say that, for a set $P \in \text{sing}^{\mathfrak{T}}$, *there is little choice below a vertex $v \in T$ if*

$$|\{Q \cap \uparrow v \mid \mathfrak{T} \models \text{sing}(Q), Q \setminus \uparrow v = P \setminus \uparrow v\}| \leq k.$$

Otherwise, we say that there is *a lot of choice below v* . We denote by $\text{Ch}_k(P)$ the set of all vertices below which there is a lot of choice for P . \lrcorner

Remark. (a) Each set of the form $\text{Ch}_k(P)$ is prefix-closed.

(b) $k \leq l$ implies $\text{Ch}_k(P) \supseteq \text{Ch}_l(P)$. \lrcorner

Lemma 4.22. Let $\text{sing}, \text{in} \in \text{WMSO}_m$. There exists a constant k such that, given a tree \mathfrak{T} such that $\langle \text{sing}, \text{in} \rangle$ encodes a finite power-set on \mathfrak{T} , the set

$$\mathcal{P}_k(v) := \{P \setminus \uparrow v \mid \mathfrak{T} \models \text{sing}(P) \text{ and } v \in \text{Ch}_k(P)\}$$

has size less than k , for every $v \in T$.

Proof. We set $k := s^2 k_o$ where k_o the constant from Lemma 4.20 and s is the number of WMSO $_m$ -theories with one free set variable and one free first-order variable. For a contradiction, suppose that there is some $v \in T$ and k sets $P_0, \dots, P_{k-1} \in \text{sing}^{\mathfrak{T}}$ such that

$$v \in \text{Ch}_k(P_i) \quad \text{and} \quad P_i \setminus \uparrow v \neq P_j \setminus \uparrow v, \quad \text{for } i \neq j.$$

By definition, $v \in \text{Ch}_k(P_i)$ implies that there are more than k sets $Q \subseteq \uparrow v$ such that

$$Q \cup (P_i \setminus \uparrow v) \in \text{sing}^{\mathfrak{T}}.$$

Hence, we can find a theory $\tau_i(X)$ such that more than $k/s = sk_o$ of these sets Q satisfy

$$\text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_v, Q, v) = \tau_i.$$

This implies that $|\mathcal{B}(v, \tau_i)| > k/s \geq k_o$. There is a subset $H \subseteq [k]$ of size $|H| \geq k/s = sk_o$ such that $\tau_i = \tau_j$, for all $i, j \in H$. Let τ be this common theory.

Finally, there exists a theory σ and subset $J \subseteq H$ of size $|J| \geq |H|/s \geq k_o$ such that

$$\text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_v, P_i \setminus \uparrow v, v') = \sigma, \quad \text{for all } i \in J,$$

where v' is the predecessor of v . Then $P_i \setminus \uparrow v \in \mathcal{A}(v, \sigma)$, for $i \in J$, which implies that $|\mathcal{A}(v, \sigma)| \geq |J| \geq k_o$. Since $\langle \sigma, \tau \rangle \models \text{sing}$, this contradicts Lemma 4.20. \square

Let us also remark that the above property is monotone in k .

Lemma 4.23. $k \leq l$ implies $\mathcal{P}_l(v) \subseteq \mathcal{P}_k(v)$.

Proof. Let $k \leq l$. To show that $\mathcal{P}_l(v) \subseteq \mathcal{P}_k(v)$ it is sufficient to note that $v \in \text{Ch}_l(P)$ implies $v \in \text{Ch}_k(P)$. \square

After these combinatorial results we can now start to explain how to encode each set $P \in \text{sing}^{\mathfrak{T}}$ by a single vertex v of \mathfrak{T} . By Theorem 4.16, it is sufficient to find a WMSO-definable function $h : \text{sing}^{\mathfrak{T}} \rightarrow T$ such that the induced distribution is sparse. We choose $h(P)$ based on $\text{Ch}_k(P)$, but we have to distinguish two cases. The complicated one is where $\text{Ch}_k(P)$ forms an infinite branch of \mathfrak{T} . We will deal with it separately below. For all other cases, we can use the following simple definition.

Definition 4.24. Let \mathfrak{T} be a tree and $C \subseteq T$ a prefix-closed set that does not form an infinite branch. The *prime branching point* $\text{pbr}(C)$ of C is the maximal vertex $v \in C$ that is comparable to all other vertices in C . \lrcorner

Thus $\text{pbr}(C)$ is either the minimal vertex with more than one successor, or C is a finite path and $\text{pbr}(C)$ is its end-point.

Lemma 4.25. *Let $\text{sing}, \text{in} \in \text{WMSO}_m$. There exist a constant $c < \omega$ with the following property. If $k < \omega$ is sufficiently large and \mathfrak{T} a tree such that $\langle \text{sing}, \text{in} \rangle$ encodes a finite power-set on \mathfrak{T} , then the distribution induced by $\text{pbr} \circ \text{Ch}_k$ is ck -sparse.*

Proof. Let $Z \subseteq T$ be finite and connected. Then

$$Z = \uparrow v \setminus (\uparrow u_0 \cup \dots \cup \uparrow u_{n-1}), \quad \text{for some } v, u_0, \dots, u_{n-1} \in T.$$

Let k_o be the constant from Lemma 4.20, k_i the constant from Lemma 4.22, let s_1 be the number of WMSO_m -theories with one free set variable and one first-order variable, and s_2 the number with two free set variables and one first-order variable. We will prove the claim for all $k \geq \max\{k_i, s_1 k_o\}$. We start by showing that the sets

$$\begin{aligned} \mathcal{C} &:= \{P \setminus \uparrow v \mid \mathfrak{T} \models \text{sing}(P), \text{pbr}(\text{Ch}_k(P)) \in Z\}, \\ \mathcal{D}_i &:= \{P \cap \uparrow u_i \mid \mathfrak{T} \models \text{sing}(P), \text{pbr}(\text{Ch}_k(P)) \in Z\}, \quad \text{for } i < n, \end{aligned}$$

all have size at most $s_1 k$.

For the first bound, note that $\text{pbr}(\text{Ch}_k(P)) \in Z$ implies $v \in \text{Ch}_k(P)$. Hence,

$$\begin{aligned} \mathcal{C} &= \{P \setminus \uparrow v \mid \mathfrak{T} \models \text{sing}(P), \text{pbr}(\text{Ch}_k(P)) \in Z\} \\ &\subseteq \{P \setminus \uparrow v \mid \mathfrak{T} \models \text{sing}(P), v \in \text{Ch}_k(P)\} = \mathcal{P}_k(v). \end{aligned}$$

By Lemmas 4.22 and 4.23, it follows that

$$|\mathcal{C}| \leq |\mathcal{P}_k(v)| \leq |\mathcal{P}_{k_i}(v)| < k_i \leq k \leq s_1 k.$$

For the second bound, it is sufficient to show that

$$|\mathcal{B}(u_i, \tau_i(P))| \leq k, \quad \text{for all } P \in \text{pbr}^{-1}(Z) \text{ and } i < n,$$

where

$$\tau_i(P) := \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_{u_i}, P \cap \uparrow u_i, u_i).$$

Since $\mathcal{D}_i \subseteq \bigcup_{\tau} \mathcal{B}(u_i, \tau)$, we then obtain $|\mathcal{D}_i| \leq s_1 k$, as desired.

If $u_i \notin \text{Ch}_k(P)$, the composition theorem for WMSO_m -theories implies that

$$\begin{aligned} \mathcal{B}(u_i, \tau_i(P)) &= \{Q \cap \uparrow u_i \mid Q \in \text{sing}^{\mathfrak{T}}, \\ &\quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_{u_i}, Q \cap \uparrow u_i, u_i) = \tau_i(P)\} \\ &= \{Q \cap \uparrow u_i \mid Q \in \text{sing}^{\mathfrak{T}}, Q \setminus \uparrow u_i = P \setminus u_i\}. \end{aligned}$$

Hence, $u_i \notin \text{Ch}_k(P)$ implies that $|\mathcal{B}(u_i, \tau_i(P))| \leq k$.

Suppose now that $u_i \in \text{Ch}_k(P)$. Then $\text{pbr}(\text{Ch}_k(P)) \in Z$ implies that $\text{pbr}(\text{Ch}_k(P)) < u_i$. Hence, there exists some $w \in \text{Ch}_k(P)$ such that

$$\text{pbr}(P) = u_i \sqcap w.$$

By definition of $\text{Ch}_k(P)$, there is some theory τ such that $\uparrow w$ contains more than k/s_1 subsets Q with

$$Q \cup (P \setminus \uparrow w) \in \text{sing}^{\mathfrak{T}} \quad \text{and} \quad \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_w, Q, w) = \tau.$$

Hence, $|\mathcal{B}(w, \tau)| \geq k/s_1 \geq k_o$. By Lemma 4.20, it follows that

$$|\mathcal{A}(w, \sigma(P))| < k_o \quad \text{where} \quad \sigma(P) := \text{Th}_{\text{WMSO}}^m(\mathfrak{T}|_w, P \setminus \uparrow w, w'),$$

and w' is the predecessor of w . Since $\uparrow u_i \subseteq T \setminus \uparrow w$ and the types $\sigma(P)$ and $\tau_i(P)$ are both obtained from the same set P , we obtain

$$|\mathcal{B}(u_i, \tau_i(P))| \leq |\mathcal{A}(w, \sigma(P))| < k_o \leq k.$$

Having established the above bounds, we can now conclude the proof by a similar argument as in the proof of Lemma 4.22: let Θ be the set of all WMSO_m -theories with two free variables and let S be the set of all sets $W \subseteq T$ encoding some union of singletons $U \in \text{sing}^{\mathfrak{T}}$ with $\text{pbr}(\text{Ch}_k(U)) \in Z$. That is, $W \in S$ if, and only if,

$$\mathfrak{T} \models \text{in}(U, W) \quad \text{implies} \quad \text{pbr}(\text{Ch}_k(U)) \in Z.$$

Note that every $W \in S$ is uniquely determined by the following data:

- ♦ the piece $W \cap Z$,
- ♦ the function $f_W : \mathcal{C} \rightarrow \Theta$ given by

$$f_W(P) := \text{Th}_{\text{WMSO}}^m(\mathfrak{T}[\nu], P, W \setminus \uparrow\nu, \nu'),$$

- ♦ the functions $g_W^i : \mathcal{D}_i \rightarrow \Theta$, $i < n$, given by

$$g_W^i(P) := \text{Th}_{\text{WMSO}}^m(\mathfrak{T}[\nu], P, W \cap \uparrow\nu, \nu'),$$

where ν' is the predecessor of ν . It follows that

$$|S| \leq 2^{|Z|} \cdot |\Theta|^{|\mathcal{C}|} \cdot \prod_{i < n} |\Theta|^{|\mathcal{D}_i|} \leq 2^{|Z|} \cdot |\Theta|^{s_1 k(n+1)}.$$

We therefore have

$$2^{|(\text{pbr} \circ \text{Ch}_k)^{-1}[Z]|} \leq |S| \leq 2^{|Z| + s_1 k \log s_2 \cdot |B_{\mathfrak{T}}(Z)|},$$

which implies that

$$|(\text{pbr} \circ \text{Ch}_k)^{-1}[Z]| \leq |Z| + s_1 k \log s_2 \cdot |B_{\mathfrak{T}}(Z)|.$$

Consequently, the distribution induced by $\text{pbr} \circ \text{Ch}_k$ is ck -sparse where $c := s_1 \log s_2$. \square

It remains to consider the chase where $C := \text{Ch}_k(P)$ forms a single infinite branch. Note that, since P is finite, there always exists some vertex $\nu \in C$ such that $\uparrow\nu$ does not contain any element from P . We would like to encode P by the smallest such vertex, which we call the *minimal prefix* of C . Unfortunately, the resulting function is not necessarily sparse. Therefore, we have to modify it in a second step which results in the notion of the *prime prefix* of C .

Definition 4.26. Let \mathfrak{T} be a tree and $C \subseteq T$ an infinite branch.

- (a) The *minimal prefix* of C is

$$\text{mpr}_k(C) := \min \left\{ \nu \in C \mid C = \text{Ch}_k(P) \text{ for some } P \in \text{sing}^{\mathfrak{T}} \text{ with } \uparrow\nu \cap P = \emptyset \right\}.$$

- (b) The *prime prefix* of C is

$$\text{ppr}_k(C) := \min \left\{ \nu \in C \mid \text{mpr}_k(C) \leq \nu \text{ and } \nu \notin D \text{ for every infinite branch } D \text{ with } \text{mpr}_k(D) < \text{mpr}_k(C) \right\}.$$

Before proving that ppr_k is sparse, we collect a few basic properties of this function.

Lemma 4.27. Let \mathfrak{T} be a tree and \mathcal{C} the set of all infinite branches C of \mathfrak{T} such that $C = \text{Ch}_k(P)$, for some $P \in \text{sing}^{\mathfrak{T}}$, and let $k < \omega$ be sufficiently large.

- (a) $|\text{Ch}_k^{-1}(C)| < k$, for all $C \in \mathcal{C}$.
 (b) $|\text{mpr}_k^{-1}(\nu)| < k$, for all $\nu \in T$.
 (c) For every $C \in \mathcal{C}$, there are only finitely many $D \in \mathcal{C}$ with

$$\text{mpr}_k(D) \leq \text{mpr}_k(C).$$

- (d) $\text{ppr}_k(C) \leq \text{ppr}_k(D)$ implies $\text{mpr}_k(C) \leq \text{mpr}_k(D)$.
 (e) $|\text{ppr}_k^{-1}(\nu)| < k$, for all $\nu \in T$.
 (f) $\text{ppr}_k(C) < \text{ppr}_k(D)$ implies $\text{ppr}_k(D) \notin \mathcal{C}$, for $C, D \in \mathcal{C}$.

Proof. Let k be the constant from Lemma 4.22.

- (a) For a contradiction, suppose that there are a branch $C \in \mathcal{C}$ and k sets $P_0, \dots, P_{k-1} \in \text{sing}^{\mathfrak{T}}$ with $\text{Ch}_k(P_i) = C$, for all i . Fix a vertex $\nu \in C$ such that

$$P_i \setminus \uparrow\nu \neq P_j \setminus \uparrow\nu, \quad \text{for all } i \neq j.$$

It follows that

$$|\{ P \setminus \uparrow\nu \mid \mathfrak{T} \models \text{sing}(P) \text{ and } \nu \in \text{Ch}_k(P) \}| = |\mathcal{P}_k(\nu)| \geq k.$$

A contradiction to our choice of k .

- (b) Suppose that $\text{mpr}_k(C) = \nu$. Then there exists some set P such that

$$P \in \text{sing}^{\mathfrak{T}}, \quad \text{Ch}_k(P) = C, \quad \nu \in \text{Ch}_k(P), \quad \text{and} \quad P \subseteq T \setminus \uparrow\nu.$$

In particular, $P \in \mathcal{P}_k(v)$. Thus

$$\text{mpr}_k^{-1}(v) \subseteq \text{Ch}_k[\mathcal{P}_k(v)].$$

By choice of k , the latter set has less than k elements.

(c) Fix $C \in \mathcal{C}$. The condition $\text{mpr}_k(D) \leq \text{mpr}_k(C)$ is equivalent to

$$D \in \text{mpr}_k^{-1}(v), \quad \text{for some } v \leq \text{mpr}_k(C).$$

Hence, it follows by (b) that there are only finitely many such D .

(d) Suppose that $\text{ppr}_k(C) \leq \text{ppr}_k(D)$. Then

$$\text{mpr}_k(C) \leq \text{ppr}_k(C) \leq \text{ppr}_k(D) \in D \quad \text{implies} \quad \text{mpr}_k(C) \in D.$$

Since $\text{mpr}_k(D) \in D$ and D is a chain, it follows that $\text{mpr}_k(C)$ and $\text{mpr}_k(D)$ are comparable. To see that $\text{mpr}_k(C) \leq \text{mpr}_k(D)$ note that, by definition of $\text{ppr}_k(C)$, $\text{mpr}_k(D) < \text{mpr}_k(C)$ would imply that $\text{ppr}_k(C) \notin D$, which we have already seen is not the case.

(e) Suppose that $\text{ppr}_k(C) = v$ and set $u := \text{mpr}_k(C)$. By (d), it follows that

$$\text{ppr}_k(D) = v \quad \text{implies} \quad \text{mpr}_k(D) = u, \quad \text{for } D \in \mathcal{C}.$$

Hence, $\text{ppr}_k^{-1}(v) \subseteq \text{mpr}_k^{-1}(u)$. By (b), the latter set has size less than k .

(f) Suppose that $\text{ppr}_k(C) < \text{ppr}_k(D)$, for $C, D \in \mathcal{D}$. Then (d) implies that $\text{mpr}_k(C) \leq \text{mpr}_k(D)$. If $\text{mpr}_k(C) < \text{mpr}_k(D)$, the claim follows immediately by definition of $\text{ppr}_k(D)$. We claim that the remaining case where $\text{mpr}_k(C) = \text{mpr}_k(D)$ is not possible. For a contradiction, suppose otherwise. Since $\text{ppr}_k(C) < \text{ppr}_k(D)$ both lie on the branch D , it then would follow that $\text{ppr}_k(C) = \text{ppr}_k(D)$. A contradiction. \square

Lemma 4.28. *Let $\text{sing, in} \in \text{WMSO}_m$ and let \mathfrak{T} be a tree such that $\langle \text{sing, in} \rangle$ encodes a finite power-set on \mathfrak{T} . The distribution induced by $\text{ppr}_k \circ \text{Ch}_k$ is k^2 -sparse, for every sufficiently large $k < \omega$.*

Proof. Let k be the constant from Lemma 4.22 and let Z be a finite connected set. Then

$$Z = \uparrow v \setminus (\uparrow u_0 \cup \dots \uparrow u_{n-1}), \quad \text{for some } v, u_0, \dots, u_{n-1} \in T.$$

Suppose that $\text{ppr}_k(C), \text{ppr}_k(D) \in Z$, for some $C, D \in \mathcal{C}$. If there is some i with $u_i \in C \cap D$, we have $C \cap Z = D \cap Z$. Hence,

$$\text{ppr}_k(C) \in C \cap Z \quad \text{and} \quad \text{ppr}_k(D) \in D \cap Z$$

implies that $\text{ppr}_k(C) \in D$ and $\text{ppr}_k(D) \in C$. By Lemma 4.27 (f), this implies that $\text{ppr}_k(C) = \text{ppr}_k(D)$. Consequently, there is some $w_i \in Z$ such that

$$\{ C \in \text{ppr}_k^{-1}(Z) \mid u_i \in C \} \subseteq \text{ppr}_k^{-1}(w_i).$$

This implies that

$$\text{ppr}_k^{-1}(Z) \subseteq \bigcup_{i < n} \text{ppr}_k^{-1}(w_i),$$

and it follows by Lemma 4.27 (e) and (a) that

$$\begin{aligned} |(\text{ppr}_k \circ \text{Ch}_k)^{-1}(Z)| &\leq \sum_{i < n} |(\text{ppr}_k \circ \text{Ch}_k)^{-1}(w_i)| \\ &\leq \sum_{i < n} k \cdot |\text{ppr}_k^{-1}(w_i)| \\ &\leq \sum_{i < n} k^2 \\ &\leq k^2 \cdot |B_{\mathfrak{T}}(Z)| \leq |Z| + k^2 \cdot |B_{\mathfrak{T}}(Z)|. \quad \square \end{aligned}$$

Definition 4.29. Let \mathfrak{T} be a tree and $P \in \text{sing}^{\mathfrak{T}}$ and let k be the constant from Lemma 4.22. The *preliminary encoding* of P is the vertex

Combining the two functions pbr and ppr we obtain our desired encoding of sets.

Proof of Theorem 4.17. Suppose that

$$\tau = \langle \delta(x), \varphi_{\subseteq}(x, y), (\varphi_R(\tilde{x}))_R \rangle.$$

Then there exists WMSO-formulae $\delta'(X)$ and $\varphi'_R(\tilde{X})$ such that

$$\begin{aligned} \mathfrak{T} \models \delta'(P) & \quad \text{iff} \quad \wp_{\text{fin}}(\mathfrak{T}) \models \delta(P), \\ \mathfrak{T} \models \varphi'_R(\tilde{P}) & \quad \text{iff} \quad \wp_{\text{fin}}(\mathfrak{T}) \models \varphi_R(\tilde{P}). \end{aligned}$$

We will construct an WMSO-formula $\text{encode}(X, y)$ that defines an injective function mapping every set X representing an atom $\{a\} \in \wp_{\text{fin}}(A)$ to some element $y \in T$. Then we can define $\sigma := \langle \gamma(x), (\psi_R(\tilde{x}))_R \rangle$ by

$$\begin{aligned} \gamma(x) &:= \exists X [\delta'(X) \wedge \text{encode}(X, x)], \\ \psi_R(\tilde{x}) &:= \exists \tilde{X} [\varphi'_R(\tilde{X}) \wedge \bigwedge_i \text{encode}(X_i, x_i)]. \end{aligned}$$

First, note that the set of singletons and the membership relation are WMSO-definable by

$$\begin{aligned} \text{sing}(X) &:= \forall Y [\varphi'_{\subseteq}(Y, X) \wedge X \neq Y \rightarrow \forall Z \varphi'_{\subseteq}(Y, Z)], \\ \text{in}(X, Y) &:= \text{atom}(X) \wedge \varphi'_{\subseteq}(X, Y). \end{aligned}$$

By Lemmas 4.25 and 4.28, we can choose a sufficiently large constant $k < \omega$ such that the function

$$\text{enc}_k(P) := \begin{cases} \text{ppr}_k(\text{Ch}_k(P)) & \text{if } \text{Ch}_k(P) \text{ forms an infinite branch,} \\ \text{pbr}(\text{Ch}_k(P)) & \text{otherwise,} \end{cases}$$

is l -sparse, for some $l < \omega$. In addition, the definitions of ppr_k , pbr , and $\text{Ch}_k(P)$, can be expressed in WMSO. Hence, we can use Theorem 4.16 to turn the function $\text{enc}_k : \text{sing}^{\mathfrak{T}} \rightarrow T$ into an injective one. \square

Notes

The collapse of guarded second-order logic to monadic second-order logic on sparse structures was proved by Courcelle [58]. A corrected and generalised

version was provided in [19]. Our exposition mostly follows Section 9.4 of [60]. Theorem 2.4 originally appeared in [146].

The power-set construction was first systematically investigated by Colcombet and Löding in [52]. Theorem 4.17 is from that article.

VIII Tree-Width and Graph Minors

1 Tree-Decompositions

THE CENTRAL NOTION of this chapter is that of a *tree-decomposition* and the associated notion of *tree-width*. The latter is a complexity measure for graphs (and more generally relational structures) that tells us how far a given graph is from being a tree. The idea is to cover a given graph by small pieces (i.e., induced subgraphs) such that the gluing operations we have to perform to obtain the whole graph can be arranged in the form of a tree. An example of a tree-decomposition is given in Figure 1 where the dashed lines are the gluing instructions indicating which vertices have to be identified.

Definition 1.1. Let \mathfrak{A} be a Σ -structure.

(a) A *tree-decomposition* of \mathfrak{A} is a family $(U_t)_{t \in T}$ of subsets $U_t \subseteq A$ indexed by an undirected tree T such that

- ♦ every element $a \in A$ is contained in some U_t ;
- ♦ every tuple $\bar{c} \in R$ in some relation R of \mathfrak{A} is contained in some U_t ; and
- ♦ for every element $a \in A$, the set $\{t \in T \mid a \in U_t\}$ of components containing a forms a connected subset of T .

We call T the *index tree* of the decomposition and the sets U_t its *components*. The *width* of a tree-decomposition $(U_t)_{t \in T}$ is the cardinal

$$\max_{t \in T} |U_t|.$$

(b) The *tree-width* $\text{twd } \mathfrak{A}$ of \mathfrak{A} is the minimal width of a tree-decomposition of \mathfrak{A} .




Figure 1: A graph (upper left) with a tree-decomposition (bottom) and its index tree (upper right).

(c) A *path-decomposition* is a tree-decomposition where the index tree is a path. The *path-width* $\text{pwd } \mathfrak{A}$ of \mathfrak{A} is the minimal width of a path-decomposition of \mathfrak{A} .

(d) The *height- n tree-width* $\text{twd}_n \mathfrak{A}$ of \mathfrak{A} is the minimal width of a tree-decomposition of \mathfrak{A} whose index tree has height at most n (as a directed tree).

Remark. For historical reasons most authors subtract 1 in the definition of the tree-width. This means most bounds on tree-width derived below differ by 1 from results found in the literature.

Examples. (a) Every tree with at least one edge has tree-width 2 (see Figure 2 for an example). The path-width of a tree depends on its height (see Proposition 1.7 below).

(b) Every cycle has path-width 3.




Figure 2: Tree-decompositions and their index trees of (a) a tree (upper left); (b) a cycle (upper right); and (c) a grid (bottom).

(c) An $m \times n$ grid with $m \leq n$ has tree-width (and path-width) $m + 1$ (see Proposition 3.6 below).

(d) A complete bipartite graph $\mathfrak{K}_{m,n}$ with $m \leq n$ has tree-width $m + 1$.

(e) A complete graph \mathfrak{K}_n has tree-width n .

We can simplify our lives by going from arbitrary relational structures to undirected graphs.

Definition 1.2. Let \mathfrak{A} be a relational structure. The *Gaifman graph* of \mathfrak{A} is

the graph with vertices A and edge relation

$$\{ \langle a, b \rangle \mid a \neq b \text{ and some relation } R \text{ of } \mathfrak{A} \text{ contains a tuple } \bar{c} \text{ with } a, b \in \bar{c} \}.$$

Let us show that the operation of forming the Gaifman graph does not change the tree-width of a structure. The main argument is contained in the following lemma.

Lemma 1.3. *Let $(U_t)_{t \in T}$ be a tree-decomposition of a graph \mathfrak{G} and let C be a clique of \mathfrak{G} . Then $C \subseteq U_t$, for some $t \in T$.*

Proof. Let v_0, \dots, v_n be an enumeration of C . Fixing some vertex of T as the root, we may turn T into an order-tree $\langle T, \leq \rangle$. As $(U_t)_t$ is a tree-decomposition, it follows that, for every $i \leq n$, there exists a unique minimal vertex $t_i \in T$ with $v_i \in U_{t_i}$. Furthermore, for all $i \neq j$, we can find some $s_{ij} \in T$ such that $U_{s_{ij}}$ contains the edge $\langle v_i, v_j \rangle$. By choice of t_i and t_j it follows that $t_i, t_j \leq s_{ij}$. This implies that $t_i \leq t_j$ or $t_j \leq t_i$. Renumbering the v_0, \dots, v_n , we may assume without loss of generality that $t_0 \leq \dots \leq t_n$. Then $t_i \leq t_n \leq s_{in}$ implies that $v_i \in U_{t_n}$. Hence, $C \subseteq U_{t_n}$. \square

Corollary 1.4. *Let \mathfrak{A} be a structure and \mathfrak{G} its Gaifman graph. Every tree-decomposition of \mathfrak{G} is also one of \mathfrak{A} and vice versa.*

Exercise 1.1. Prove that every graph of tree-width k is k -colourable. \square

The following observation is also frequently useful.

Lemma 1.5. *Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph and $(U_t)_{t \in T}$ a tree-decomposition of \mathfrak{G} . For every connected set $C \subseteq V$, the set*

$$\langle C \rangle := \{ t \in T \mid U_t \cap C \neq \emptyset \}$$

forms a connected subset of T .

Proof. We prove the claim by induction on $|C|$. If $C = \{v\}$, then $\langle v \rangle$ is connected by definition of a tree-decomposition. Hence, suppose that $|C| > 1$.

Fix a vertex $v \in C$ and set $C_o := C \setminus \{v\}$. Since C is connected, there exists an edge $e = \langle v, u \rangle$ with $u \in C_o$. By the definition of a tree-decomposition, we can find some component U_t with $e \subseteq U_t$. Then $u, v \in U_t$ implies $t \in \langle C_o \rangle \cap \langle v \rangle$. Thus, $\langle C \rangle = \langle C_o \rangle \cup \langle v \rangle$ is the union of two overlapping connected sets. Hence, it is itself also connected. \square

As an example let us use this lemma to compute the path-width of trees. Similar computations of the tree-width have to be postponed to Section 3 where we have the necessary tools available to prove the required lower bounds.

Definition 1.6. A full k -ary tree is a successor-tree T where each internal vertex has exactly k successors and all the leaves are on the same level. \square

Proposition 1.7. *Let $k \geq 3$. The full k -ary tree of height n has path-width $n + 1$.*

Proof. Let \mathfrak{T} be the full k -ary tree of height n . For the upper bound, fix a left-to-right enumeration v_0, \dots, v_{m-1} of all leaves of \mathfrak{T} and let U_i be the set of vertices of the branch connecting the root to v_i . Then $(U_i)_{i < m}$ is a path-decomposition of \mathfrak{T} of width $n + 1$.

For the lower bound, let \mathfrak{S} be the undirected version of \mathfrak{T} . It is sufficient to prove that $\text{pwd } \mathfrak{S} \geq n + 1$. Consider a path-decomposition $(U_i)_{i < m}$ of \mathfrak{S} . We prove the claim by induction on n . If $n = 0$, then \mathfrak{S} consists of a single vertex v and there is some component U_i with $v \in U_i$. Hence, $|U_i| = 1$.

For the inductive step, suppose that $n > 0$. Removing all empty components from $(U_i)_i$ we may assume that $U_0 \neq \emptyset$ and $U_{m-1} \neq \emptyset$. Fix vertices $u \in U_0$ and $v \in U_{m-1}$ and a path P connecting u and v . By Lemma 1.5, every component U_i contains some vertex from P . The graph $\mathfrak{S} - P$ obtained from \mathfrak{S} by removing all vertices in P (and all incident edges) has the path-decomposition $(U_i \setminus P)_{i \leq m}$. Let $r \in S$ be the vertex of \mathfrak{S} corresponding to the root of \mathfrak{T} . Then r has $k \geq 3$ neighbours. At most 2 of them lie on the path P . Let \mathfrak{S}_o be the subtree of \mathfrak{S} attached to one of the other neighbours of r . Then \mathfrak{S}_o is the undirected version of a full k -ary tree of height $n - 1$. By inductive hypothesis, it follows that $\text{pwd } \mathfrak{S}_o \geq n$. This implies that $|U_i \setminus P| \geq n$, for some i . As U_i contains at least one vertex from P , it follows that $|U_i| \geq n + 1$. \square

For binary trees one can use a trick that reduces the path-width by 1.

Exercise 1.2. Find a path-decomposition of the full binary tree of height n whose width is n . \lrcorner

Exercise 1.3. Prove that the path-width of full binary trees of height n is unbounded as n grows. *Hint.* Use the fact that every full binary tree of height $2n$ has the full 4-ary tree of height n as a minor. \lrcorner

Note that structures of bounded tree-width are sparse. Hence, we can apply the results from Chapter VII.

Lemma 1.8. Let Σ be a signature of size $m := |\Sigma|$ and let n be the maximal arity of a relation in Σ . Every Σ -structure \mathcal{A} of tree-width k is mk^n -sparse.

Proof. Given a finite subset $X \subseteq A$, let $(U_t)_{t \in T}$ be a tree-decomposition of the restriction $\mathcal{A}|_X$ of width at most k such that T is minimal. This implies that $U_s \not\subseteq U_t$, for every edge $\langle s, t \rangle$ of T , as, otherwise, we would obtain a smaller decomposition by removing the component U_s from T . Consequently, $|T| \leq |X|$. Fix a relation $R \in \Sigma$ of arity r . As every tuple $\bar{c} \in R \upharpoonright X$ is contained in some component U_t , it follows that

$$|R \upharpoonright X| \leq \sum_{t \in T} |U_t|^r \leq |T| \cdot k^r \leq k^r |X|.$$

Consequently,

$$\sum_{R \in \Sigma} |R \upharpoonright X| \leq mk^n |X|. \quad \square$$

Corollary 1.9. Over the class of all Σ -structures of tree-width at most k , every GSO-formula is equivalent to some MSO-formula.

Lemma 1.10. For every finite signature Σ , there exist constants $c, d < \omega$ such that

$$\text{twd } \mathcal{A} \leq \text{twd } \mathcal{A}_{\text{in}} \leq c(\text{twd } \mathcal{A})^d, \quad \text{for every } \Sigma\text{-structure } \mathcal{A}.$$

Proof. For the lower bound, fix a tree-decomposition $(U_t)_{t \in T}$ of \mathcal{A}_{in} . Setting $U'_t := U_t \cap A$, we obtain a tree-decomposition $(U'_t)_{t \in T}$ of \mathcal{A} whose width is at most that of $(U_t)_{t \in T}$.

For the upper bound, fix a tree-decomposition $(U_t)_{t \in T}$ of \mathcal{A} of width k . Set

$$U'_t := U_t \cup \{ \bar{c} \mid \bar{c} \in R \cap (U_t)^n \text{ for some } n\text{-ary relation } R \text{ of } \mathcal{A} \}.$$

It follows that $(U'_t)_{t \in T}$ is a tree-decomposition of \mathcal{A}_{in} whose width is at most $k + sk^r \leq (s+1)k^r$, where s is the number of relations in Σ and r their maximal arity. \square

First-Order Properties

Finally, let us relate the tree-width of a structure with its first-order theory. The technical problem we have to deal with is that some constructions will yield tree-decompositions that are not indexed by trees but by non-standard models of the theory of all trees. To preserve sufficiently much of the tree structure when going to a non-standard extension we will work with order-trees and we add an explicit infimum-operation.

Definition 1.11. (a) A *non-standard tree* is a lower semilattice $\langle T, \leq, \sqcap \rangle$ where every set of the form $\downarrow t := \{ s \in T \mid s \leq t \}$ is linearly ordered.

(b) A *generalised path-decomposition* of a graph $\mathfrak{G} = \langle V, E \rangle$ is a sequence $(U_i)_{i \in I}$ of sets indexed by a linear order I such that

- ♦ $\bigcup_{i \in I} U_i = V$
- ♦ For every edge $e \in E$, there is some index $i \in I$ with $e \subseteq U_i$.
- ♦ $v \in U_i \cap U_k$ implies $v \in U_j$, for all $i \leq j \leq k$.

(c) A *generalised tree-decomposition* of a graph $\mathfrak{G} = \langle V, E \rangle$ is a family $(U_t)_{t \in T}$ of sets indexed by a non-standard tree T such that

- ♦ $\bigcup_{t \in T} U_t = V$
- ♦ For every edge $e \in E$, there is some index $t \in T$ with $e \subseteq U_t$.
- ♦ $U_s \cap U_t \subseteq U_u$, for all u with $s \sqcap t \leq u \leq s$ or $s \sqcap t \leq u \leq t$. \lrcorner

Remark. Note that the property of being a non-standard tree is expressible in first-order logic. \square

Before constructing generalised tree-decompositions let us show how to transform them into ordinary ones. We start with path-decompositions.

Lemma 1.12. *For every generalised path-decomposition $(U_i)_{i \in I}$ of finite width, there exists a tree-decomposition $(W_t)_{t \in T}$ of the same width such that*

- ♦ *for every $i \in I$, there is some $t \in T$ with $W_t = U_i$, and*
- ♦ *there is some $t_* \in T$ with $W_{t_*} = \bigcap_{i \in I} U_i$.*

Proof. If there are several indices i with the same set U_i we can remove all duplicates without destroying the property of being a path-decomposition. Therefore, we may assume without loss of generality that $U_i \neq U_j$, for $i \neq j$.

Let n be the width of $(U_i)_{i \in I}$. We prove the claim by induction on $n - |C|$, where

$$C := \bigcap_{i \in I} U_i.$$

For each $i \in I$, we define the set of *shared* elements

$$A_i := \{a \in U_i \mid a \in U_k \text{ for some } k \neq i\}$$

and the relation

$$i \sim j \quad \text{iff} \quad A_k = A_i \quad \text{for all } i \leq k \leq j.$$

Note that \sim is an equivalence relation with convex classes. For a \sim -class w , we write

$$A_w := A_i, \quad \text{for any index } i \in w.$$

Let $J := \{w \in I/\sim \mid A_w \supset C\}$ and $K := \{i \in I \mid A_i = C\}$. For every $w \in J$ we can use the inductive hypothesis to find a tree-decomposition $(Y_r^w)_{r \in R_w}$ corresponding to the generalised path-decomposition $(U_i)_{i \in w}$.

Furthermore, we have $|A_i| < n$ as the sequence $(U_i)_i$ does not contain duplicates. Consequently, $(A_w)_{w \in J}$ forms a generalised path-decomposition of width at most $n - 1$. Since $C \subseteq A_i$, for all i , the value of $n - |C|$ for $(A_w)_w$ is strictly smaller than the one for $(U_i)_i$. Hence, we can use the inductive hypothesis to find a tree-decomposition $(Z_s)_{s \in S}$ corresponding to $(A_w)_{w \in J}$.

We construct the desired tree-decomposition $(W_t)_{t \in T}$ as follows. We take the disjoint union of all decompositions $(Z_s)_{s \in S}$ and $(Y_r^w)_{r \in R_w}$, for $w \in J$. In addition, we introduce a new vertex t_* with component $W_{t_*} := C$ and, for every $k \in K$, a vertex t_k with component $W_{t_k} := U_k$. To turn this into a tree-decomposition, we add edges between

- ♦ t_* and every t_k , for $k \in K$,
- ♦ t_* and some vertex in S , and
- ♦ for every $w \in J$, we choose some vertex $s_w \in S$ with $Z_{s_w} = A_w$, and add an edge between s_w and an arbitrary vertex in R_w . \square

Lemma 1.13. *For every generalised tree-decomposition $(U_s)_{s \in S}$ of finite width, there exists a tree-decomposition $(W_t)_{t \in T}$ of the same width such that*

- ♦ *for every $s \in S$, there is some $t \in T$ with $W_t = U_s$, and*
- ♦ *there is some $t_* \in T$ with $W_{t_*} = \bigcap_{s \in S} U_s$.*

Proof. We say that a tree-decomposition $(W_t)_{t \in T}$ *covers* an index $s \in S$ if $U_s = W_t$, for some $t \in T$. Fix an enumeration $(s_i)_{i < \kappa}$ of S where $\kappa := |S|$. By induction on the ordinal i , we will construct an increasing chain of tree-decompositions $(W_t^{(i)})_{t \in T^{(i)}}$, $i \leq \kappa$, such that $(W_t^{(i)})_{t \in T^{(i)}}$ covers every $u \in \bigcup_{j < i} \downarrow s_j$. Then $(W_t^{(\kappa)})_{t \in T^{(\kappa)}}$ will then be the desired tree-decomposition.

We start with a tree $T^{(0)}$ consisting of a single vertex t_* with component $W_{t_*}^{(0)} = \bigcap_{s \in S} U_s$. For a limit ordinal δ , we take for $(W_t^{(\delta)})_{t \in T^{(\delta)}}$ the limit of $(W_t^{(i)})_{t \in T^{(i)}}$, for $i < \delta$.

For the successor step, suppose that we have already defined $(W_t^{(i)})_{t \in T^{(i)}}$. Let $I \subseteq S$ be the set of all $u \leq s_i$ that are not covered by $(W_t^{(i)})_{t \in T^{(i)}}$. By

inductive hypothesis, I forms a final segment of the chain $\downarrow s_i$. If I is empty, there is nothing to do and we cat set $(W_t^{(i+1)})_{t \in T^{(i+1)}} := (W_t^{(i)})_{t \in T^{(i)}}$. Otherwise, $(S_u)_{u \in I}$ forms a generalised path-decomposition and we can use Lemma 1.12 to transform it into a tree-decomposition $(Z_r)_{r \in R}$. Set

$$C := \bigcup_{u \in I} \bigcap_{v \in I \cap \downarrow u} S_v \quad \text{and} \quad D := \bigcup_{u \in \downarrow s_i \setminus I} \bigcap_{v \in (\downarrow s_i \setminus I) \cap \uparrow u} S_v.$$

As all sets S_v are finite, there exist indices $u \in I$ and $v \in \downarrow s_i \setminus I$ such that

$$C \subseteq S_u \quad \text{and} \quad D \subseteq S_v.$$

By assumption, there is some $r_o \in R$ such that $Z_{r_o} = S_u$ and some $t_o \in T^{(i)}$ with $W_t^{(i)} = S_v$. To construct $(W_t^{(i+1)})_{t \in T^{(i+1)}}$ we attach $(Z_r)_{r \in R}$ to $(W_t^{(i)})_{t \in T^{(i)}}$ by adding an edge between r_o and t_o . Note that the result is indeed a tree-decomposition since every element that is contained both in some component $W_t^{(i)}$ and in some component Z_r must belong to $C \cap D$. \square

After these preparations we can start looking at the first-order theory of a structure with finite tree-width. The main result is the following compactness property which implies that in many cases it is sufficient to consider finite graphs only.

Theorem 1.14. *Let Σ be a finite relational signature. For every Σ -structure \mathcal{A} , we have*

$$\text{twd } \mathcal{A} = \sup \{ \text{twd } \mathcal{B} \mid \mathcal{B} \subseteq \mathcal{A} \text{ finite} \}.$$

Proof. Let \mathcal{A} be a graph and suppose that

$$\sup \{ \text{twd } \mathcal{B} \mid \mathcal{B} \subseteq \mathcal{A} \text{ finite} \} = k.$$

Clearly, each tree-decomposition of \mathcal{A} induces one for every substructure $\mathcal{B} \subseteq \mathcal{A}$. This implies that $k \leq \text{twd } \mathcal{A}$. Hence, it remains to show that $\text{twd } \mathcal{A} \leq k$.

We encode every generalised tree-decomposition $(U_t)_{t \in T}$ of \mathcal{A} by the relation

$$U := \{ \langle t, v \rangle \mid v \in U_t \}$$

on the disjoint union $\mathcal{A} \oplus T$. Thus, we work with structures over the signature $\Sigma \cup \{ \leq, \sqcap, U \}$ where \leq and \sqcap are the relations of \mathfrak{T} . Let us call the resulting $(\Sigma \cup \{ \leq, \sqcap, U \})$ -structure \mathcal{A}_T .

To prove the bound on the tree-width of \mathcal{A} , we will construct suitable relations \leq , \sqcap , and U . Let Δ be the so-called *atomic diagram* of \mathcal{A} , that is, the following set of first-order formulae:

$$\Delta := \{ \alpha(\bar{c}_a) \mid \mathcal{A} \models \alpha(\bar{a}) \} \cup \{ \neg \alpha(\bar{c}_a) \mid \mathcal{A} \not\models \alpha(\bar{a}) \},$$

where $\alpha(\bar{x})$ ranges over all atomic formula (over Σ) and the c_a, c_b, \dots are constant symbols, one for each element $a \in A$. Furthermore, we need a set Φ of formulae stating that \leq , \sqcap , and U encode a generalised tree-decomposition of width at most k . We use the following formulae to state the properties of a tree-decomposition.

- ♦ “There exists a partition $P \cup Q$ of the universe such that $U \subseteq Q \times P$, the elements in Q are related by \leq and \sqcap , while all other relations are over P .”
- ♦ $\langle Q, \leq, \sqcap \rangle$ is a non-standard tree.
- ♦ “Every element is contained in some component.”

$$\forall x [Px \rightarrow \exists t Utx].$$

- ♦ “Every edge is covered by some component.”

$$\forall \bar{x} [R\bar{x} \rightarrow \exists t [Utx_o \wedge \dots \wedge Utx_{n-1}]],$$

for every relation R .

- ♦ “The set of components containing a given element v is connected.”

$$\forall x \forall s \forall t \forall u [s \sqcap t \leq u \leq s \wedge Uxs \wedge Utx \rightarrow Uux].$$

- ♦ “Every component has size at most k .”

$$\forall t \exists y_0 \dots \exists y_{k-1} \forall x \left[U_t x \rightarrow \bigvee_{i < k} x = y_i \right].$$

We start by showing that the set $\Phi \cup \Delta$ is satisfiable. Given a finite subset $\Psi \subseteq \Phi \cup \Delta$, let $\mathfrak{B} \subseteq \mathfrak{A}$ be a finite substructure containing all elements of Δ mentioned in Ψ . By assumption, \mathfrak{B} has a tree-decomposition $(U_t)_{t \in T}$ of width k . Then the encoding \mathfrak{B}_T satisfies Ψ .

By the Compactness Theorem for first-order logic, it follows that there is some model of $\Phi \cup \Delta$. By choice of Φ , it follows that this model is of the form \mathfrak{A}_T^+ . Where $\mathfrak{A}^+ \supseteq \mathfrak{A}$ is an extension of \mathfrak{A} and $(U_t)_{t \in T}$ is a generalised tree-decomposition of \mathfrak{A}^+ of width at most k . Set

$$W_t := U_t \cap A, \quad \text{for } t \in T.$$

Then $(W_t)_{t \in T}$ is a generalised tree-decomposition of \mathfrak{A} of width at most k . Consequently, it follows by Lemma 1.13 that $\text{twd } \mathfrak{A} \leq k$. \square

One noteworthy consequence of this theorem is the fact that having finite tree-width is a first-order property.

Theorem 1.15. *Let Σ be a finite relational signature, \mathfrak{A} and \mathfrak{B} two Σ -structures, and $k < \omega$.*

$$\mathfrak{A} \equiv_{\text{FO}} \mathfrak{B} \quad \text{implies} \quad \text{twd } \mathfrak{A} \leq k \Leftrightarrow \text{twd } \mathfrak{B} \leq k.$$

Proof. Suppose that $\mathfrak{A} \equiv_{\text{FO}} \mathfrak{B}$ and $\text{twd } \mathfrak{A} \leq k$. For every $n < \omega$, there exists a first-order formula φ_n stating every substructure with n elements has tree-width at most k . Then

$$\mathfrak{A} \models \varphi_n \quad \text{implies} \quad \mathfrak{B} \models \varphi_n.$$


Consequently, it follows by Theorem 1.14 that $\text{twd } \mathfrak{B} \leq k$. \square

2 Minors

The notion of tree-width is closely related to graph minor theory. Let us collect a few results relating these two areas.

Definition 2.1. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) The operation of *contracting* an edge $e = \langle u, v \rangle$ of \mathfrak{G} replaces the two end-points u and v by a single new vertex, which is a neighbour of every former neighbour of u and of v .



(b) A *minor* of \mathfrak{G} is a graph \mathfrak{H} obtained from \mathfrak{G} by (i) removing an arbitrary number of vertices and edges and (ii) contracting an arbitrary number of the remaining edges.

(c) We denote the function mapping a graph \mathfrak{G} to the set of its minors by Min . \lrcorner

We can describe a minor of a graph using the following data.

Definition 2.2. Let $\mathfrak{G} = \langle V, E \rangle$ and $\mathfrak{H} = \langle U, F \rangle$ be undirected graphs.

(a) Two sets $A, B \subseteq V$ *touch* if they have some vertex in common or if there exists an edge between a vertex of A and one of B .

(b) A *minor-embedding* of \mathfrak{H} in \mathfrak{G} consists of a family $(C_u)_{u \in U}$ of disjoint subsets $C_u \subseteq V$ such that

- ♦ each C_u is connected and non-empty and
- ♦ for every edge $\langle u, u' \rangle \in F$ of \mathfrak{H} , the associated sets C_u and $C_{u'}$ touch. \lrcorner

Lemma 2.3. *There exists a minor-embedding of \mathfrak{H} in \mathfrak{G} if, and only if, \mathfrak{H} is (isomorphic to) a minor of \mathfrak{G} .*

We begin our study of the relationship between minors and tree-decompositions with the following simple observation.

Lemma 2.4. *Let \mathfrak{H} be a minor of \mathfrak{G} . For every tree-decomposition of \mathfrak{G} of width k there exists a tree-decomposition of \mathfrak{H} of width at most k that has the same index tree.*

Proof. Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{G} of width $\text{twd } \mathfrak{G}$. We construct the desired tree-decomposition $(U'_t)_{t \in T}$ of \mathfrak{H} as follows. Each component U'_t is obtained from U_t by performing the following two operations.

- ♦ We remove from U_t all vertices v that were deleted from \mathfrak{G} when constructing \mathfrak{H} .
- ♦ For every contracted edge $\langle u, v \rangle$, we replace in U_t the vertices u and v with the vertex they were merged into.

Clearly, the new family $(U'_t)_t$ is a valid tree-decomposition and its width is at most that of $(U_t)_t$. \square

Corollary 2.5. $\mathfrak{H} \in \text{Min}(\mathfrak{G})$ implies $\text{twd } \mathfrak{H} \leq \text{twd } \mathfrak{G}$.

The preceding result gives a first indication that there might be a connection between minors and tree-width. We will present three theorems that provide an even tighter relationship between these notions, one for each of the three variants of tree-width we have introduced. The first two, for height- n tree-width and for path-width, follow below; the third version, for general tree-width, is a bit more involved and therefore deferred to Section 4. For the proofs we need the following notion of a spanning tree.

Definition 2.6. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

- (a) A *spanning tree* of \mathfrak{G} is an undirected tree $\mathfrak{T} = \langle T, F \rangle$ such that

$$T = V \quad \text{and} \quad F \subseteq E.$$

- (b) A *normal spanning tree* of \mathfrak{G} is an order-tree $\mathfrak{T} = \langle T, \leq \rangle$ such that the underlying undirected tree is a spanning tree of \mathfrak{G} and we have

$$u \leq v \quad \text{or} \quad v \leq u, \quad \text{for every edge } \langle u, v \rangle \in E \text{ of } \mathfrak{G}.$$

For a finite graph, we can construct a normal spanning tree by a depth-first traversal. For countable graphs, we can do so as well but we need to

be slightly more careful. Uncountable graphs do not need to have normal spanning trees. For instance, no uncountable complete graph has one.

Lemma 2.7. *Every countable graph $\mathfrak{G} = \langle V, E \rangle$ has a normal spanning tree.*

Proof. Fix a well-order \leq of V and let v_0, v_1, \dots be the corresponding increasing enumeration of the vertices. We inductively define an increasing sequence of trees $\mathfrak{T}_0 \subseteq \mathfrak{T}_1 \subseteq \dots$ with the following properties.

- ♦ \mathfrak{T}_n is a normal spanning tree of the subgraph of \mathfrak{G} induced by T_n .
- ♦ \mathfrak{T}_n contains v_0, \dots, v_n .
- ♦ If there exists a path \mathfrak{P} of \mathfrak{G} linking two vertices of \mathfrak{T}_n and such that the internal vertices of \mathfrak{P} are disjoint from \mathfrak{T}_n , then

$$v_i \leq v_k \quad \text{or} \quad v_k \leq v_i, \quad \text{in } \mathfrak{T}_n.$$

We start with the singleton tree \mathfrak{T}_0 that consists just of v_0 . For the inductive step, suppose that we have already defined \mathfrak{T}_{n-1} . If \mathfrak{T}_{n-1} contains v_n , we set $\mathfrak{T}_n := \mathfrak{T}_{n-1}$. Otherwise, let P be the set of all paths of \mathfrak{G} of minimal length connecting v_n to some vertex $u \in T_{n-1}$, and let $C \subseteq T_{n-1}$ be the set of these vertices u . By assumption on \mathfrak{T}_{n-1} , the set C forms a chain in \mathfrak{T}_{n-1} . Let w be its maximal vertex and let \mathfrak{P} be the corresponding path from w to v_n . For \mathfrak{T}_n , we choose the tree obtained from \mathfrak{T}_{n-1} by attaching this path to w . It remains to check that \mathfrak{T}_n has the desired properties. It obviously contains v_0, \dots, v_n . By minimality of \mathfrak{P} , it also follows that \mathfrak{T}_n is a normal spanning tree of the subgraph induced by T_n . Furthermore, if \mathfrak{Q} is any path linking a vertex $u \in P$ to some $v \in T_{n-1}$, then $v \in C$, which implies that $v \leq u$. \square

Our first theorem relates height- n tree-width and path minors.

Theorem 2.8 (Excluded Path Theorem). *Every undirected graph \mathfrak{G} with $\text{twd}_n(\mathfrak{G}) > n$ has a path of length n as a minor.*

Proof. By Theorem 1.14, it is sufficient to prove the claim for finite graphs \mathfrak{G} . Hence, suppose that \mathfrak{G} is finite and that it does not have the path of length n

as a minor. To show that $\text{twd}_n(\mathfrak{G}) \leq n$, we fix a normal spanning tree $\langle T, \leq \rangle$ of \mathfrak{G} and we define a tree-decomposition $(U_v)_{v \in T}$ of \mathfrak{G} by setting

$$U_v := \{u \in T \mid u \leq v\}.$$

As T is depth-first, it follows that every edge $\langle u, v \rangle$ of \mathfrak{G} is contained in some component U_w where for w we can choose the maximum of u and v (with respect to \leq). Consequently, $(U_v)_v$ is in fact a tree-decomposition.

To compute its width, note that the height of the tree T must be less than n since \mathfrak{G} contains no path of length n . Hence, $|U_v| = |v| + 1 \leq n$ and the width of $(U_v)_v$ is at most n . \square

Let us also provide a converse.

Proposition 2.9. *Let \mathfrak{G} be an undirected graph with $\text{twd}_n \mathfrak{G} \leq k$. Every path contained in \mathfrak{G} has length at most $(k-1)(k+1)^n$.*

Proof. Suppose that $\text{twd}_n \mathfrak{G} \leq k$. We prove the claim by induction on n . If $n = 0$, then \mathfrak{G} has at most k vertices and every path contained in \mathfrak{G} has length at most $k-1$.

For the inductive step, suppose that $n > 0$. Fix a tree-decomposition $(U_t)_{t \in T}$ of \mathfrak{G} of width at most k such that T has height at most n . Let r be the root of T and consider some path P in \mathfrak{G} . We can decompose P into segments $P = P_0 \dots P_s$ such that no internal vertex of P_i belongs to U_r . We choose the number s of segments minimal. This implies that the common end-point of P_i and P_{i+1} belongs to U_r . Since each vertex can appear at most once in P , it follows that $s \leq |U_r| \leq k$. By inductive hypothesis, each segment P_i has length at most $(k-1)(k+1)^{n-1}$. Consequently, the length of P is at most $(s+1)(k-1)(k+1)^{n-1} \leq (k-1)(k+1)^n$. \square

For path-width, a similar result holds if we replace paths by trees. For the proof, we need some preparations. First, let us recall a standard result from graph theory about the existence of disjoint paths.

Definition 2.10. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) A *separation* $\langle A, B \rangle$ of \mathfrak{G} consists of two subsets $A, B \subseteq V$ such that $A \cup B = V$ and no edge of \mathfrak{G} connects a vertex in $A \setminus B$ to one in $B \setminus A$. The *order* of a separation $\langle A, B \rangle$ is the number $|A \cap B|$.

(b) For a set $F \subseteq E$ of edges, we denote by $\mathfrak{G} - F$ the graph obtained from \mathfrak{G} by deleting all edges in F (but keeping all the vertices). For $F = \{e\}$, we usually simply write $\mathfrak{G} - e$. \downarrow

Theorem 2.11 (Menger). *Let $\mathfrak{G} = \langle V, E \rangle$ be a finite undirected graph and $X, Y \subseteq V$ two sets of size at least k . There exists a family of k pairwise disjoint paths connecting X and Y if, and only if, \mathfrak{G} has no separation $\langle A, B \rangle$ of order less than k with $X \subseteq A$ and $Y \subseteq B$.*

Proof. (\Rightarrow) Suppose that $\langle A, B \rangle$ is a separation of order $m < k$ with $X \subseteq A$ and $Y \subseteq B$. Since every path between X and Y must contain a vertex in $A \cap B$, at most $m < k$ such paths can be pairwise disjoint.

(\Leftarrow) Suppose that there is no separation $\langle A, B \rangle$ of order less than k with $X \subseteq A$ and $Y \subseteq B$. We construct the desired family of paths by induction on $|E|$.

First suppose that $E = \emptyset$. Then $\langle X, (V \setminus X) \cup (X \cap Y) \rangle$ is a separation of \mathfrak{G} . By assumption, its order is $|X \cap Y| \geq k$. Hence, selecting k vertices in $X \cap Y$ produces a family of k paths of length 0 between X and Y .

For the inductive step, suppose that there is some edge $e = \langle u, v \rangle \in E$. Let \mathfrak{G}/e be the graph obtained from \mathfrak{G} by contracting the edge e and let w_* denote the vertex into which u and v were merged by this. If \mathfrak{G}/e has k disjoint paths between X and Y , then so does \mathfrak{G} . Hence, we may assume that this is not the case. By inductive hypothesis, it follows that \mathfrak{G}/e has a separation $\langle A_0, B_0 \rangle$ of order less than k with $X \subseteq A_0$ and $Y \subseteq B_0$. If $w_* \notin A_0 \cap B_0$, then $\langle A_0, B_0 \rangle$ would induce a separation of \mathfrak{G} of the same order, in contradiction to our assumption. It follows that the sets

$$A := (A_0 \setminus \{w_*\}) \cup \{u, v\} \quad \text{and} \quad B := (B_0 \setminus \{w_*\}) \cup \{u, v\}$$

form a separation $\langle A, B \rangle$ of \mathfrak{G} of order $|A \cap B| \leq k$.

Furthermore, any separation $\langle C, D \rangle$ of $\mathfrak{G} - e$ with $X \subseteq C$ and $Y \subseteq D$ is also a separation of \mathfrak{G} with $Y \subseteq D$. By assumption, this implies that every

such $\langle C, D \rangle$ has order at least k . Hence, we can use the inductive hypothesis to construct a family R_0, \dots, R_{k-1} of disjoint paths between X and S in the graph $\mathfrak{G}|_A$. In the same way, we can obtain a family R'_0, \dots, R'_{k-1} of disjoint paths between S and Y in $\mathfrak{G}|_B$. As $|S| \leq k$, it follows that each path R_i shares an end-point with a unique R'_j . Renumbering these paths if necessary, we may assume that $j = i$. Consequently, we obtain k disjoint paths between X and Y by concatenating R_i and R'_i , for each $i < k$. \square

Suppose we want to construct a path-decomposition of a given graph inductively. Say, we have already defined components U_0, \dots, U_n and we would like to add the next one U_{n+1} . A necessary condition for such a naïve strategy to succeed is that the last component U_n contains every vertex that is connected to some vertex not yet covered by the components we have defined so far. Let us give this property a name.

Definition 2.12. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) The *boundary* of a set $A \subseteq V$ is

$$\partial A := \{v \in A \mid \langle v, u \rangle \in E \text{ for some } u \in V \setminus A\}.$$

(b) A *k-extendible sequence* for a set $C \subseteq V$ is a path-decompositions $(U_i)_{i \leq n}$ of $\mathfrak{G}|_C$ of width at most k such that $\partial C \subseteq U_n$. \lrcorner

We will need the following lemma about restrictions of *k-extendible* sequences.

Lemma 2.13. Let $B \subseteq A$ be sets and let $(P_u)_{u \in \partial B}$ be a family of disjoint paths between ∂B and ∂A where P_u starts at the vertex u . If there exists a *k-extendible* sequence for A , there also exists one for B .

Proof. Given a *k-extendible* sequence $(U_t)_{t \leq m}$ for A , we set

$$U_{<t} := U_0 \cup \dots \cup U_{t-1}$$

$$\text{and } W_t := ((U_t \setminus U_{<t}) \cap B) \cup \partial(U_{<t} \cap B).$$

We claim that $(W_t)_{t \leq m}$ is a *k-extendible* sequence for B .

Let us start by checking the axioms of a path-decomposition. First, consider a vertex $v \in B \subseteq A$. Let t be the minimal index with $v \in U_t$. Then $v \in (U_t \setminus U_{<t}) \cap B \subseteq W_t$.

Next, consider an edge $e = \langle u, v \rangle \in E$ with $u, v \in B$. Let s and t be the minimal indices with $u \in U_s$ and $v \in U_t$. By symmetry, we may assume that $s \leq t$. If $s = t$, then $u, v \in (U_t \setminus U_{t-1}) \cap B \subseteq W_t$ and we are done. Hence, suppose that $s < t$. Then the existence of the edge e implies that $u \in \partial(U_{<t} \cap B) \subseteq W_t$. Since $v \in (U_t \setminus U_{t-1}) \cap B \subseteq W_t$, the claim follows.

Finally, consider indices $s < r < t$ and a vertex $v \in W_s \cap W_t$. Then $W_s \subseteq U_{<r} \subseteq U_{<t}$ implies that $v \in U_{<r}$ and $v \in U_{<t} \cap W_t$. By definition of W_t , it follows that $v \in \partial(U_{<t} \cap B)$. Hence,

$$v \in \partial(U_{<t} \cap B) \cap U_{<r} \subseteq \partial(U_{<r} \cap B) \subseteq W_r.$$

Furthermore, we have $\partial B \subseteq W_m$ since

$$\partial B \cap U_{<(m-1)} \subseteq \partial(U_{<(m-1)} \cap B) \subseteq W_m,$$

$$\partial B \cap (U_m \setminus U_{<(m-1)}) \subseteq (U_m \setminus U_{m-1}) \cap B \subseteq W_m.$$

To conclude the proof it therefore remains to compute the width of $(W_t)_t$. To do so, it is sufficient to show that

$$|\partial(U_{<t} \cap B)| \leq |\partial U_{<t}|,$$

since this implies that

$$\begin{aligned} |W_t| &= |((U_t \setminus U_{<t}) \cap B) \cup \partial(U_{<t} \cap B)| \\ &\leq |U_t \setminus U_{<t}| + |\partial U_{<t}| \\ &\leq |U_t \setminus U_{<t}| + |U_t \cap U_{<t}| = |U_t| \leq k. \end{aligned}$$

For the proof of the above inequality, let $C := U_{<t} \cap B$ and $D := U_{<t}$. We construct an injective map $\alpha : \partial C \setminus \partial D \rightarrow \partial D \setminus \partial C$. Consider a vertex $u \in \partial C \setminus \partial D$. Then there exists an edge $\langle u, v \rangle$ with $v \notin C$. Since $u \notin \partial D$, it follows that $v \in D$. Hence, $v \in D \setminus C = U_{<t} \setminus B$. This implies that $u \in \partial B$. Since P_u is a path from $u \in U_{<t}$ to a vertex $w \in \partial A \subseteq U_m$ it must visit some

vertex $\alpha(u) \in \partial U_{<t} = \partial D$. As u is the only vertex of P_u belonging to B , it follows that $\alpha(u) \in \partial D \setminus B \subseteq \partial D \setminus \partial C$. Furthermore, the map α defined in this way is injective, since that paths P_u are disjoint. \square

These preparations out of the way we can state and prove the desired excluded minor theorem.

Theorem 2.14 (Excluded Tree Theorem). *Let \mathfrak{T} be an undirected tree with n vertices. Every graph \mathfrak{G} with path-width $\text{pwd } \mathfrak{G} \geq n$ has \mathfrak{T} as a minor.*

Proof. Suppose that $\text{pwd } \mathfrak{G} \geq n$ and let \mathfrak{T} be a tree with n vertices. Choosing a root for \mathfrak{T} , we may assume that \mathfrak{T} is directed. We fix an enumeration t_0, \dots, t_{n-1} of T starting with the root t_0 and such that the parent of every other vertex t_i appears before it in this list. Let \mathfrak{T}_m be the subtree of \mathfrak{T} consisting of the vertices t_0, \dots, t_{m-1} and set

$$H_m := \{ C \subseteq V \mid |\partial C| \leq m \text{ and there exists an } (n-1)\text{-extendible sequence for } C \}.$$

Let $m < n$ be the maximal number such that there exists a set $A \subseteq V$ such that

- ♦ A is a maximal element of H_m and
- ♦ A contains a minor-embedding $(C_i)_{i < m}$ of \mathfrak{T}_m with $|C_i \cap \partial A| = 1$, for all i .

Such a number exists since $\emptyset \in H_0$. Hence, we could take $m = 0$ and a maximal element of H_0 .

We can use the set A to construct a minor-embedding of \mathfrak{T} . We start by fixing a vertex $v \in V \setminus A$ as follows. If $m = 0$, we pick an arbitrary element $v \in V \setminus A$. (This is possible since $\text{pwd } \mathfrak{G} \geq n$ implies $V \notin H_0$. Hence, $A \neq V$.) If $m > 0$, let t_j be the parent of t_m in \mathfrak{T} . By choice of our enumeration of \mathfrak{T} , we have $j < m$. Let u be the unique vertex with $C_j \cap \partial A = \{u\}$, and let v be some neighbour of u that belongs to $V \setminus A$.

If $m = n-1$, we set

$$D_i := \begin{cases} C_i & \text{for } i < n-1, \\ \{v\} & \text{for } i = n-1. \end{cases}$$

Then $(D_i)_i$ is a minor-embedding of \mathfrak{T} in \mathfrak{G} and we are done.

For a contradiction, suppose that $m < n-1$. Set $B_0 := A \cup \{v\}$ and let $(U_i)_{i \leq k}$ be an $(n-1)$ -extendible sequence for A . We claim that the family $(U'_i)_{i \leq (k+1)}$ defined by

$$U'_i := \begin{cases} U_i & \text{if } i \leq k, \\ \{v\} \cup \partial A & \text{if } i = k+1. \end{cases}$$

forms an $(n-1)$ -extendible sequence for B_0 . Clearly, $(U'_i)_i$ is a path-decomposition and we have $\partial B_0 \subseteq \{v\} \cup \partial A = U'_{k+1}$. Furthermore,


$$|U'_{k+1}| = |\{v\} \cup \partial A| \leq 1 + m < n$$

implies that $(U'_i)_i$ has width less than n .

Hence, $(U'_i)_i$ is an $(n-1)$ -extendible sequence. By maximality of A , it follows that $B_0 \notin H_m$, that is, $|\partial B_0| > m$. Since $\partial B_0 \subseteq \partial A \cup \{v\}$, it follows that $|\partial B_0| = m+1$ and $\partial B_0 = \partial A \cup \{v\}$. In particular, $B_0 \in H_{m+1}$ and we can find some maximal element $B \in H_{m+1}$ with $B_0 \subseteq B$.

To obtain the desired contradiction, it remains to prove that B contains a minor-embedding of \mathfrak{T}_{m+1} . In the subgraph $\mathfrak{G}|_B$ we fix a maximal family L of pairwise disjoint paths connecting B_0 and ∂B . By Menger's Theorem, there exists a separation $\langle C, D \rangle$ of $\mathfrak{G}|_B$ with $B_0 \subseteq C$, $\partial B \subseteq D$, and $|C \cap D| = |L|$. Set $S := C \cap D$ and let W be the union of S and all connected components of $\mathfrak{G}|_{V \setminus S}$ that contain a vertex of B_0 . Then

$$(*) \quad \partial W \subseteq S \quad \text{and} \quad A \subset B_0 \subseteq W.$$



Let us show that $W \subseteq B$. For a contradiction, suppose that there is some vertex $w \in W \setminus B$. Then $w \notin S$ since $S \subseteq B$. By choice of W , we can therefore find a path P from B_0 to w that is disjoint from S . Since $B_0 \subseteq B$ and $w \notin B$ this path must contain some vertex in ∂B . Let u be the first such vertex and let P_0 be the prefix of P connecting B_0 to u . Then P_0 is a path in B connecting B_0 and ∂B . Hence, P_0 must contain some vertex of S . A contradiction to the fact that we chose P disjoint from S .

Hence, $W \subseteq B \in H_{m+1}$ and we can use Lemma 2.13 to find an $(n-1)$ -extendible sequence for W . By $(*)$, maximality of A , and the fact that every path in L meets ∂B_0 , it follows that

$$m < |\partial W| \leq |S| = |L| \leq |\partial B_0| = m + 1.$$

Hence, the paths in L induce a bijection between ∂B_0 and ∂B . For $i < m$, let $P_i \in L$ be the path containing the unique vertex in $\partial A \cap C_i$, and let $P_m \in L$ be the path containing the vertex v from above. We define $(D_i)_{i < m+1}$ by

$$D_i := \begin{cases} C_i \cup P_i & \text{if } i < m, \\ P_m & \text{if } i = m. \end{cases}$$

To obtain the desired contradiction, we show that the set B together with the family $(D_i)_{i < m+1}$ satisfy the above conditions for $m+1$. Thus, m is not maximal. First, note that the sets D_i are connected and pairwise disjoint.

Furthermore, D_i and D_j touch if t_i is the parent of t_j . Hence, $(D_i)_i$ forms a minor-embedding of \mathfrak{T}_{m+1} . To see that every D_i meets ∂B in a unique point, note that

$$C_i \subseteq A \subseteq B_0 \quad \text{implies} \quad D_i \cap \partial B = P_i \cap \partial B,$$

and each path P_i contains a single vertex of ∂B . \square

Corollary 2.15. (a) A class of finite graphs has bounded path-width if and only if it excludes some tree (equivalently, some forest) as a minor.

(b) A class of finite graphs has bounded height- n tree-width, for some n , if and only if it excludes some path as a minor (equivalently, as a subgraph).

3 Brambles

Upper bounds on the tree-width of a graph are easy to obtain: one just has to find a tree-decomposition of that width. But for a lower bound one has to prove that no such tree-decomposition exists. That is much more difficult. To simplify this task we will introduce a combinatorial configuration called a *bramble* whose existence implies the non-existence of a tree-decomposition.

Definition 3.1. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph and $k < \omega$.

(a) We say that a set $S \subseteq V$ *covers* a family $\mathcal{B} \subseteq \wp(V)$ if $S \cap B \neq \emptyset$, for all $B \in \mathcal{B}$.

(b) A k -*bramble* is a family \mathcal{B} of subsets $B \subseteq V$ with the following properties.

(B1) Every $B \in \mathcal{B}$ is non-empty and connected.

(B2) Every pair of sets $B, B' \in \mathcal{B}$ touch.

(B3) No set $S \subseteq V$ of size $|S| < k$ covers \mathcal{B} . $\quad \text{,}$

Theorem 3.2. A graph \mathfrak{G} has a k -bramble if, and only if, $\text{twd } \mathfrak{G} \geq k$.

Before presenting the proof we need to make a few technical preparations. The following notions will only be used in this section.

Definition 3.3. Let $(U_t)_{t \in T}$ be a tree-decomposition of a graph \mathfrak{G} , and fix a number k .

- (a) A vertex $t \in T$ is called *big* if $|U_t| \geq k$.
- (b) We call $(U_t)_t$ *good* if not every vertex of T is big and if all big vertices are leaves of T .
- (c) A *petal* of $(U_t)_t$ is a set of the form $U_t \setminus U_s$, where t is a big leaf and s is the neighbour of t .
- (d) In a graph \mathfrak{G} , we denote the set of *neighbours* of a set P by

$$N(P) := \{v \in V \setminus P \mid v \text{ is a neighbour of some } u \in P\}.$$

We regard a good tree-decomposition as a *partial* tree-decomposition where the big components still need to be decomposed further. The following, rather technical lemma states that, given two good tree-decompositions satisfying certain compatibility conditions, we can find a common refinement.

Lemma 3.4. Let $(U_t)_{t \in T}$ and $(U'_t)_{t \in T'}$ be two good tree-decompositions of a graph \mathfrak{G} , let P be a petal of $(U_t)_t$ and P' one of $(U'_t)_t$. Suppose that

- ♦ there are vertices $t_o \in T$ and $t'_o \in T'$ such that $U_{t_o} = P \cup N(P)$ and $U'_{t'_o} = P' \cup N(P')$;
- ♦ no petal of $(U_t)_t$ contains P' and no petal of $(U'_t)_t$ contains P ; and
- ♦ P and P' do not touch.

Then there exists a good tree-decomposition $(W_s)_{s \in S}$ such that every petal of $(W_s)_s$ are contained in some petal of $(U_t)_t$ or of $(U'_t)_t$, while P and P' are not petals of $(W_s)_s$.

Proof. Since P and P' do not touch, the set $N(P)$ is disjoint from P' . By definition, it is also disjoint from P . Since P is a petal and $U_{t_o} = P \cup N(P)$, the set $N(P)$ separates P from the rest of the graph. In particular, it separates P from P' . Consequently, there exists a separation $\langle A, B \rangle$ of \mathfrak{G} with $P \subseteq A$ and $P' \subseteq B$. We choose $\langle A, B \rangle$ such that the set $C := A \cap B$ has minimal size. Then $|C| \leq |N(P)| \leq |U_{s_o}| \leq k$, where s_o is the neighbour of t_o . By minimality of C and the Theorem of Menger, it follows that A contains

a family $(Q_v)_{v \in C}$ of disjoint paths connecting $N(P)$ and C . In the same way, we can find a family $(Q'_v)_{v \in C}$ of disjoint paths in B that connect $N(P')$ and C .

Let \mathfrak{H} be the minor of \mathfrak{G} obtained by (i) deleting every vertex in A that does not belong to one of the paths Q_v and (ii) contracting every path Q_v . Note that the resulting graph can also be obtained from $\mathfrak{G}|_B$ by adding some edges to C (those between Q_v and Q_u in the original graph \mathfrak{G}). By Lemma 2.4, $(U_t)_{t \in T}$ induces a tree-decomposition of \mathfrak{H} and, therefore, also one of $\mathfrak{G}|_B$. We denote the latter by $(\tilde{U}_t)_{t \in T}$. It follows that

$$\tilde{U}_t = (U_t \cap B) \cup \{v \in C \mid U_t \cap Q_v \neq \emptyset\}.$$

Since $U_{t_o} = P \cup N(P) \subseteq A$ and $N(P)$ contains vertices of every path Q_v , it follows in particular that $\tilde{U}_{t_o} = C$.

In the same way we obtain a tree-decomposition $(\tilde{U}'_t)_{t \in T'}$ of $\mathfrak{G}|_A$ with $\tilde{U}'_{t'_o} = C$. Let S be the tree obtained from the disjoint union of T and T' by identifying the vertices t_o and t'_o , and denote the resulting vertex by s_* . For $s \in S$, set

$$W_s := \begin{cases} \tilde{U}_s & \text{if } s \in T, \\ \tilde{U}'_s & \text{if } s \in T'. \end{cases}$$

(Note that this is well-defined since $\tilde{U}_{t_o} = C = \tilde{U}'_{t'_o}$.)

We claim that $(W_s)_{s \in S}$ is the desired good tree-decomposition of \mathfrak{G} . It is straightforward to check that it is a tree-decomposition. Furthermore, the internal vertices s of S are those of T and T' plus the new vertex s_* . We have $|W_{s_*}| = |C| \leq k$ and

$$|W_s| = |\tilde{U}_s| \leq |U_s| \leq k, \quad \text{for } s \in T.$$

A similar calculation shows that $|W_s| \leq k$ for $s \in T'$. Hence, none of these is big and $(W_s)_s$ is good.

To conclude the proof, let R be a petal of $(W_s)_s$ located at, say, the vertex $s \in S$. By symmetry, we may assume that $s \in T$. Then $R = W_s \setminus W_t = \tilde{U}_s \setminus \tilde{U}_t$,

where t is the neighbour of s . Hence, $|U_s| \geq |\tilde{U}_s| \geq |W_s| \geq k$ implies that s is also big as a vertex of $(U_s)_s$. Hence, $\hat{R} := U_s \setminus U_t$ is a petal of $(U_t)_t$. It remains to show that $R \subseteq \hat{R}$ and that \hat{R} is different from P and P' .

As $(W_s)_s$ is a tree-decomposition and $s \neq s_*$, we have $W_s \cap W_{s_*} \subseteq W_t \cap W_{s_*}$. This implies that

$$\begin{aligned} R \cap W_{s_*} &= (W_s \setminus W_t) \cap W_{s_*} = (W_s \cap W_{s_*}) \setminus W_t \\ &\subseteq (W_t \cap W_{s_*}) \setminus W_t = \emptyset. \end{aligned}$$

Consequently, $R \subseteq B \setminus A$. Since $W_s \setminus A = \tilde{U}_s \setminus A = U_s \setminus A$ and similarly for W_t , it follows that

$$R = R \setminus A = (W_s \setminus A) \setminus (W_t \setminus A) \subseteq U_s \setminus U_t = \hat{R}.$$

Furthermore, $R \cap A = \emptyset$ implies that $R \neq P'$. And we have $R \neq P$ since, by assumption, no petal of $(U_t)_t$ contains P . \square

Proof of Theorem 3.2. (\Rightarrow) Let \mathcal{B} be a k -bramble and let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{G} . We claim that some component U_t covers \mathcal{B} . By (B3), this implies that $|U_t| \geq k$.

If there exists some edge $e = \langle t_o, t_i \rangle$ of T such that $U_{t_o} \cap U_{t_i}$ covers \mathcal{B} , we are done. Hence, suppose otherwise. Then we can orient each edge $e = \langle t_o, t_i \rangle$ of T as follows. Let C_i be the connected component of $T - e$ containing t_i and set $A_i := \bigcup_{t \in C_i} U_t$. Then $\langle A_o, A_i \rangle$ is a separation of \mathfrak{G} . As $U_{t_o} \cap U_{t_i}$ does not cover \mathcal{B} , the set

$$\mathcal{C}_e := \{ B \in \mathcal{B} \mid B \cap U_{t_o} \cap U_{t_i} = \emptyset \}$$

is non-empty. Since all members of \mathcal{C}_e are connected and touch each other, it follows that they are either all contained in $A_o \setminus A_i$, or in $A_i \setminus A_o$. In the first case, we direct e from t_i to t_o , in the second case, from t_o to t_i .

Having defined an orientation of T , there exists some vertex $t \in T$ such that all edges with end-point t are directed towards t . (Just start at an arbitrary vertex and follow outgoing edges until you arrive at t .) We claim that U_t covers \mathcal{B} .

For a contradiction, suppose otherwise. Then there is some $B \in \mathcal{B}$ with $B \cap U_t = \emptyset$. If t has no neighbours in T , then $U_t = V$ and $B \cap U_t = B \neq \emptyset$. A contradiction. Hence, t has neighbours s_o, \dots, s_{n-1} in T . For each edge $e_i := \langle s_i, t_i \rangle$, $B \cap U_t = \emptyset$ implies that $B \in \mathcal{C}_{e_i}$. Furthermore, as the edge e_i is directed towards t , we have

$$B \subseteq \bigcup_{u \in C_i} U_u,$$

where C_i is the connected component of $T - e_i$ containing t . It follows that

$$B \subseteq \bigcap_{i < n} \bigcup_{u \in C_i} U_u = U_t.$$

A contradiction.

(\Leftarrow) Suppose that $\text{twd } \mathfrak{G} \geq k$. Let \mathcal{B}^+ be a set of minimal size satisfying the following two conditions.

- (1) \mathcal{B}^+ contains a petal of every good tree-decomposition.
- (2) \mathcal{B}^+ is upwards closed in the sense that, if $P \subseteq Q$ are both petals of good tree-decompositions, then $P \in \mathcal{B}$ implies $Q \in \mathcal{B}$.

Note that $\text{twd } \mathfrak{G} \geq k$ implies that every good tree-decomposition has at least one petal. Consequently, the set of *all* petals of good tree-decompositions satisfies (1) and (2), which means that a set like \mathcal{B}^+ exists. Set

$$\mathcal{B} := \{ B \in \mathcal{B}^+ \mid B \text{ is connected} \}.$$

We claim that \mathcal{B} is the desired k -bramble.

By definition, every set in \mathcal{B} is non-empty and connected. This proves (B1). For (B3), we fix a set S with $|S| < k$. Since \mathfrak{G} has tree-width at least k , we have $|V| \geq k$. In particular, $V \setminus S$ is non-empty. Let C_o, \dots, C_{m-1} be the connected components of $V \setminus S$. We obtain a tree-decomposition of \mathfrak{G} with components S and $S \cup C_i$, for $i < m$, where the underlying index tree consists of one internal vertex and m leaves. This tree-decomposition is good. Consequently, one of its petals C_i belongs to \mathcal{B}^+ . As C_i is connected, we have $C_i \in \mathcal{B}$. Then $S \cap C_i = \emptyset$ implies that the set S does not cover \mathcal{B} .

It remains to prove (B2). For a contradiction, suppose that there are two sets $P, P' \in \mathcal{B}$ that do not touch. We can choose them minimal (with respect to inclusion). Then $\mathcal{B}^+ \setminus \{P\}$ and $\mathcal{B}^+ \setminus \{P'\}$ still satisfy (2), and it follows by minimality of \mathcal{B} that they must violate (1). Consequently, we can find two good tree-decompositions $(U_t)_{t \in T}$ and $(U'_t)_{t \in T'}$ such that P is the only petal of $(U_t)_t$ belonging to \mathcal{B}^+ , and P' is only one of $(U'_t)_t$. Let $t_o \in T$ and $t'_o \in T'$ be the vertices corresponding to, respectively, P and P' . Removing every vertex of U_{t_o} not belonging to $P \cup N(P)$ and every vertex of $U'_{t'_o}$ not belonging to $P' \cup N(P')$, we may assume that $U_{t_o} = P \cup N(P)$ and $U'_{t'_o} = P' \cup N(P')$. (Note that $|P \cup N(P)| \geq k$ since otherwise the modified tree-decomposition would have no petal in \mathcal{B}^+ , which contradicts (1). Furthermore, note that every removed vertex belongs to $U_{t_o} \cap U_s$, where s is the neighbour of t_o . Hence, the new tree-decompositions still cover every vertex of \mathfrak{G} .)

By (2), every petal of $(U_t)_t$ containing P' would belong to $\mathcal{B}^+ \setminus \{P\}$. Hence, there is no such petal. In the same way it follows that $(U'_t)_t$ has no petal containing P . Consequently, we can use Lemma 3.4 to obtain a good tree-decomposition $(W_s)_s$ that has no petal in \mathcal{B}^+ . A contradiction. \square

As a prototypical example of graphs with large tree-width we consider grids.

Definition 3.5. Let $m, n > 1$. The $m \times n$ grid is the graph with vertices

$$V := [m] \times [n]$$

and edges

$$E := \{ \langle \langle i, j \rangle, \langle k, l \rangle \rangle \mid |i - k| + |j - l| = 1 \}.$$

Proposition 3.6. Let $m \leq n$. The $m \times n$ grid has tree-width and path-width $m + 1$.

Proof. For the upper bound, we construct a tree-decomposition $(U_u)_{u \in T}$ of the $m \times n$ grid of width $m + 1$ as follows. Consider the enumeration

v_0, \dots, v_{mn-1} of V where

$$v_{mj+i} := \langle i, j \rangle, \quad \text{for } i < m \text{ and } j < n.$$

We set

$$U_k := \{v_k, \dots, v_{k+m}\}, \quad \text{for } k < m(n-1)$$

(see Figure 2 (c)). Then $(U_k)_{k < m(n-1)}$ forms a path-decomposition of width $|U_k| = m + 1$.

It remains to prove the lower bound. By Theorem 3.2 it is sufficient to find an $(m + 1)$ -bramble \mathcal{B} . We set


$$\mathcal{B} := \{C_{ij} \mid i < m-1, j < n-1\} \cup \{P, Q\},$$

where

$$C_{ij} := (\{i\} \times [n-1]) \cup ([m-1] \times \{j\}),$$

$$P := \{m-1\} \times [n],$$

$$Q := [m-1] \times \{n-1\}.$$



To check the bramble properties, note that each set C_{ij}, P, Q is non-empty and connected. Furthermore, $C_{ij} \cap C_{i'j'} \neq \emptyset$ and there are edges between C_{ij} and P , between C_{ij} and Q , and between P and Q . So all members of \mathcal{B} touch.

Finally, consider a set $S \subseteq V$ of size $|S| \leq m$. If $S \cap P = \emptyset$ or $S \cap Q = \emptyset$, we are done. Hence, suppose otherwise. Then S contains a vertex $u \in P$ and a vertex $v \in Q$. Let $I \subseteq [m]$ and $J \subseteq [n]$ be the projections of $S \setminus (P \cup Q)$ to, respectively, the first and the second coordinate. Since $|S \setminus (P \cup Q)| \leq m - 2$, we can find two indices $i \in [m - 1] \setminus I \neq \emptyset$ and $j \in [n - 1] \setminus J \neq \emptyset$. Consequently, $C_{ij} \cap Z = \emptyset$. \square

Exercise 3.1. Prove that every cycle of length $n \geq 3$ has tree-width 3.

Exercise 3.2. Prove that the complete bipartite graph $\mathcal{K}_{n,n}$ has tree-width $n + 1$.

4 The Excluded Grid Theorem

Similarly to the Excluded Path Theorem and the Excluded Tree Theorem, there exists also a version for tree-width. The corresponding minors are grids. The present section is devoted to the proof of this result.

Theorem 4.1 (Excluded Grid Theorem). *There exists a function $g : \omega \rightarrow \omega$ such that every graph of tree-width at least $g(n)$ contains an $n \times n$ grid as a minor.*

Before presenting the proof let us make a few remarks concerning grids, minors, and planar graphs.

Proposition 4.2. *A finite graph \mathcal{G} is planar if, and only if, it is a minor of the $n \times n$ grid, for some n .*

Proof. (\Leftarrow) Clearly, every grid is planar. Furthermore, every subgraph of a planar graph is planar. Finally, every graph obtained from a planar graph by contracting a single edge is also planar.


(\Rightarrow) Fix an embedding of \mathcal{G} in the plane where each line is represented by a curve that is piecewise linear. Let $S_0 \subseteq \mathbb{R}^2$ be the resulting set of points. We choose three real numbers $\gamma > \delta > \varepsilon > 0$ as follows.

(1) First, we fix $\gamma > 0$ such that the distance between any two vertices of \mathcal{G} is greater than 2γ and the distance from each vertex v of \mathcal{G} to every

edge e that is not incident with v is greater than γ . In the embedding S_0 we replace every point representing a vertex of \mathcal{G} by a closed disc of radius γ . Let $S_1 \subseteq \mathbb{R}^2$ be the resulting shape.

(2) Next, we choose $0 < \delta < \gamma$ such that the distance between all curves representing edges is greater than 2δ and the distance between every edge e and every disc representing some vertex is greater than δ . Let $S_2 \subseteq \mathbb{R}^2$ be the set of points containing S_1 and every point whose distance from a curve representing some edge is at most ε .

(3) Finally, choose $0 < \varepsilon < \delta$ such that $\varepsilon < \frac{1}{3}\delta$ and every line segment that is part of a curve representing some edge has length at least 5ε . Fix numbers $a, b \in \mathbb{R}$ and an integer $k > 0$ such that $S_2 \subseteq [a, a + k\varepsilon] \times [b, b + k\varepsilon]$, and fix an embedding of the $(k + 1) \times (k + 1)$ grid in the plane such that the vertex $\langle i, j \rangle$ is represented by the point $\langle a + i\varepsilon, b + j\varepsilon \rangle$ and every edge is represented by a straight line parallel to one of the coordinate axes.



Then \mathcal{G} can be obtained from this grid by (i) deleting every vertex and every edge whose image is not contained in the set S_2 , and (ii) contracting most of the remaining edges. Hence, \mathcal{G} is a minor of the $(k + 1) \times (k + 1)$ grid. \square

Corollary 4.3. *A class of finite graphs has bounded tree-width if, and only if, it excludes some planar graph as a minor.*

Proof. We have shown in Proposition 3.6 that the $n \times n$ grid has tree-width $n + 1$. Consequently it follows by Corollary 2.5 that no finite graph \mathcal{G} of

tree-width n can have the $n \times n$ grid cannot be a minor. Hence, every class with tree-width n omits the $n \times n$ grid as a minor.

Conversely, suppose that a class \mathcal{C} omits some planar graph \mathfrak{H} as a minor. By Proposition 4.2, we can then find some number n such that every graph in \mathcal{C} omits the $n \times n$ grid as a minor. Hence, the Excluded Grid Theorem tells us that the tree-width of graphs in \mathcal{C} is bounded by $g(n)$. \square

Pseudo-Grids

To prove the Excluded Grid Theorem, we simplify a given graph of large tree-width in several steps until we obtain a grid. We present these steps in reverse order from the more specific to the more general configurations, starting with graphs that already closely resemble a grid. Throughout this section all graphs will be implicitly assumed to be finite and undirected.

Definition 4.4. (a) We denote by $\mathfrak{K}_{m,n}$ the complete bipartite graph with $m+n$ vertices, that is, the graph whose vertex set is divided into two classes $A := [m]$ and $B := [n]$ such that each vertex in A is connected to every vertex in B .

(b) An $m \times n$ pseudo-grid is an undirected graph \mathfrak{G} with vertices $V := [m] \times [n]$ such that

- for all $i < m$ and $k < n-1$, there is an edge between $\langle i, k \rangle$ and $\langle i, k+1 \rangle$,
- for all $k < n$, the subgraph induced by $C(k) := [m] \times \{k\}$ is connected.

We call the set $C(k) := [m] \times \{k\}$ the k -th column of \mathfrak{G} and $R(i) := \{i\} \times [n]$ its i -th row.

(c) Let \mathfrak{G} be a graph and $A, B \subseteq V$ sets of vertices. An A - B -linkage is a set L of disjoint paths where each $P \in L$ has one end-point in A and one in B . A linkage is *tight* if $|A| = |B| = |L|$.

(d) We will make use of the following constants. For $g \geq 1$ and $h \geq 3$, define

$$M(g, h) := 2(2g+1)(h-2), \quad \varepsilon(g, h) := \frac{1}{2M(g, h)}.$$

Furthermore, we denote by $\mathcal{G}(g, h)$ the class of all finite undirected graphs that contain neither the $g \times g$ grid, nor the complete bipartite graph $\mathfrak{K}_{h,h}$ as a minor. \lrcorner

We start by showing that, in every pseudo-grid we can find either a grid or a complete bipartite graph, depending on whether or not the column graphs $C(k)$ contain long paths. If sufficiently many $C(k)$ contain a path of sufficient length, we obtain a grid minor; otherwise, we obtain $\mathfrak{K}_{h,h}$.

Proposition 4.5. Let $g \geq 1$ and $h \geq 3$, and set $m := (2g+1)(2h-5) + 2$ and $n := N(g, h)$. No $m \times n$ pseudo-grid \mathfrak{G} belongs to $\mathcal{G}(g, h)$.

The proof is split into several lemmas. We start with two lemmas that allow us to extract from a pseudo-grid either a grid minor or a complete bipartite graph.

Lemma 4.6. Let \mathfrak{G} be an $m \times n$ pseudo-grid, $k < l$ two column indices, and $A \subseteq C(k)$ and $B \subseteq C(l)$ two sets of the same size $s := |A| = |B|$. If $l-k > s$, then there exists a tight A - B -linkage L of size s such that every vertex in $\bigcup L$ belongs to $A \cup B \cup \bigcup_{k < i < l} C(i)$.

Proof. Let \mathfrak{H} be the subgraph of \mathfrak{G} induced by $A \cup B \cup \bigcup_{k < i < l} C(i)$. We have to show that there exist s disjoint paths between A and B in \mathfrak{H} . For a contradiction, suppose otherwise. By Menger's Theorem, there then exists a set S of size $|S| < s$ such that every connected subgraph of \mathfrak{H} intersecting A and B also intersects S . As $|A| > |S|$, we can find some $i < m$ such that the row $R(i)$ intersects A and but not S . Similarly, we can find an index $j < m$ such that $R(j)$ intersects B , but not S . Furthermore $l-k-1 > s-1 \geq |S|$ implies that there is some index $k < h < l$ with $C(h) \cap S = \emptyset$. It follows that the set $R(i) \cup C(h) \cup R(j)$ induces a connected graph that intersects both A and B , but not S . A contradiction to our choice of S . \square

Lemma 4.7. Let $h \geq 1$ and $m \geq h+1$, and let \mathfrak{G} be an $m \times h^2$ pseudo-grid. If each column $C(i(h+1))$ with $i < h$, has a spanning tree with at least h leaves, then \mathfrak{G} has a $\mathfrak{K}_{h,h}$ minor.

Proof. For each $i < h$, we fix a spanning tree T_i of $C(i(h+1))$ with at least h leaves and we choose a set $A_i \subseteq T_i$ of exactly h such leaves. By choice of A_i , each set $B_i := C(i(h+1)) \setminus A_i$ induces a connected subgraph of \mathfrak{G} . Since $h+1 > |A_i| = |A_{i+1}|$, we can use Lemma 4.6 to find, for every $i < h-1$, a tight A_i - A_{i+1} -linkage L_i with

$$\bigcup L_i \subseteq A_i \cup A_{i+1} \cup \bigcup_{i(h+1) < j < (i+1)(h+1)} C(j).$$

The union $\bigcup_i L_i$ therefore consists of h disjoint paths. Let \mathfrak{H} be the minor of \mathfrak{G} obtained by (i) contracting each path in $\bigcup_i L_i$ to a single vertex and (ii) doing the same with every set B_i . As every $v \in A_i$ is connected to some $u \in B_i$, it follows that each path in $\bigcup_i L_i$ touches every set B_j . Consequently, $\mathfrak{H} \cong \mathfrak{K}_{h,h}$. \square

It remains to prove that, in every pseudo-grid, we can apply one of the two preceding lemmas.

Lemma 4.8. *Let $r \geq 0$ and $h \geq 3$, and let \mathfrak{G} be a connected graph with at least*

$$|V| \geq (r+3)(2h-5) + 2$$


vertices. Then either

- ♦ \mathfrak{G} has a spanning tree with at least h leaves, or
- ♦ \mathfrak{G} contains an path of length r where every internal vertex has degree 2 in \mathfrak{G} .

Proof. For a contradiction, suppose that neither of these cases occurs. We choose a spanning tree T of \mathfrak{G} such that the number of leaves of T is maximal. Then T has at most $h-1$ leaves and, therefore, at most $h-3$ vertices with more than 2 neighbours. Thus, T is a subdivision of a tree with $2h-4$ vertices and $2h-5$ edges. It follows that we can decompose T into $2h-5$ induced paths (induced with respect to T). Let v_0, \dots, v_{k-1} be the longest such path. This implies that T has at most $(k-2)(2h-5) + 2h-4$ vertices.

We claim that, in \mathfrak{G} , the sequence $v_2, v_3, \dots, v_{k-4}, v_{k-3}$ forms the desired path. For a contradiction, suppose that some vertex v_i with $3 \leq i \leq k-4$ has a third neighbour u in \mathfrak{G} . Let P be the path in T between u and v_i . By

symmetry, we may assume that v_{i-1} belongs to P . Let T' be the tree obtained from T by replacing the edge $\langle v_{i-2}, v_{i-1} \rangle$ by the edge $\langle v_i, u \rangle$.



To obtain the desired contradiction, we show that T' has at least one more leaf than T . The proof consists of two cases. If $v_{i-2} \neq u$, then v_{i-1} and v_{i-2} are leaves of T' , while the only vertex that could be a leaf of T , but not of T' is u . Similarly, if $v_{i-2} = u$, then v_{i-1} is a leaf of T' and there is no vertex that is a leaf of T but not of T' .

We have found an path $v_2, v_3, \dots, v_{k-4}, v_{k-3}$ with $k-4$ vertices where the internal vertices have exactly 2 neighbours in \mathfrak{G} . By assumption, this implies that $k-5 < r$. As have seen above that T has at most $(k-2)(2h-5) + 2h-4$ vertices, it follows that

$$\begin{aligned} |V| &= (k-2)(2h-5) + 2h-4 = (k-1)(2h-5) + 1 \\ &\leq (r+3)(2h-5) + 1. \end{aligned}$$

A contradiction. \square

After these preparations, we can prove the proposition.

Proof of Proposition 4.5. Let \mathfrak{G} be a $m \times n$ pseudo-grid. We have to show that \mathfrak{G} has a $g \times g$ grid or a $\mathfrak{K}_{h,h}$ as a minor. For $i < n - 2g + 2$, set

$$D_i := C(i) \cup \dots \cup C(i + 2g - 2)$$

and let \mathfrak{H}_i be the minor of $\mathfrak{G}[D_i]$ obtained by contracting every edge in $R(j) \cap D_i$, for $j < m$. Below we will prove the following claim:

- (*) If, for some $i < n - 2g + 2$, the graph \mathfrak{H}_i has no spanning tree with at least h leaves, then \mathfrak{G} has a $g \times g$ grid as a minor.

Then the proof concludes as follows. If \mathfrak{G} has a $g \times g$ grid as a minor, we are done. Otherwise, it follows by $(*)$ that every \mathfrak{H}_i has a spanning tree with at least h leaves. Let \mathfrak{G}_o be graph obtained from \mathfrak{G} by contracting all edges in

$$R(j) \cap D_{di}, \quad \text{for } j < m, \ i < h, \text{ and } d := 2g + h - 1.$$

Note that this is well-defined since the last column in $D_{d(h-1)}$ has index

$$\begin{aligned} d(h-1) + 2g - 2 &= (2g + h - 1)(h - 1) + 2g - 2 \\ &= (2g + h - 2)(h - 1) + (h - 1) + 2g - 2 \\ &= (2g + h - 2)h - 1 \\ &= n - 1. \end{aligned}$$

It follows that \mathfrak{G}_o is a $m \times n'$ pseudo-grid consisting of the columns

$$\begin{aligned} \mathfrak{H}_o, C(2g - 1), C(2g), \dots, C(2g + h - 2), \\ \mathfrak{H}_1, C(d + 2g - 1), C(d + 2g), \dots, C(d + 2g + h - 2), \\ \mathfrak{H}_2, \dots, \\ \mathfrak{H}_{h-2}, C(d(h - 2) + 2g - 1), C(d(h - 2) + 2g), \dots, \\ C(d(h - 2) + 2g + h - 2), \\ \mathfrak{H}_{h-1}. \end{aligned}$$

Since $n' = (h - 1)(h + 1) + 1 = h^2$ and every $(h + 1)$ -th of these columns has a spanning tree with at least h leaves, we can use Lemma 4.7 to show that \mathfrak{G}_o has a $\mathfrak{R}_{h,h}$ minor. Hence, so does \mathfrak{G} .

It remains to prove $(*)$. Fix an index $i < n - 2g + 2$. Since $|H_i| = m = (2g + 1)(2h - 5) + 2$, it follows by Lemma 4.8 that \mathfrak{H}_i contains a path where every internal vertex has degree 2 in \mathfrak{H}_i . Permuting the rows of \mathfrak{G} if necessary, we may assume that this path is of the form u_o, \dots, u_{2g-2} , where $u_j \in H_i$ is the vertex created by contracting the edges in $R(j) \cap D_i$. For $0 < j < 2g - 2$ and $i \leq k \leq i + 2g - 1$, it follows that the vertex $\langle j, k \rangle \in R(j) \cap C(k)$ has no neighbours in $C(k)$, except possibly for $\langle j - 1, k \rangle$ and $\langle j + 1, k \rangle$. As $C(k)$ is connected, there exists a path from the vertex $\langle g - 1, k \rangle$ to $\langle o, k \rangle$. This path

must start with either

$$\langle g - 1, k \rangle, \dots, \langle o, k \rangle \quad \text{or} \quad \langle g - 1, k \rangle, \dots, \langle 2g - 2, k \rangle.$$

By symmetry, we may assume that the former is the case for at least g of the $2g$ indices k between i and $i + 2g - 1$. These g paths together with the paths $R(o) \cap D_i, \dots, R(g - 1) \cap D_i$ (with suitably contracted edges) thus form a $g \times g$ grid. \square

Linkages

In the next step, we collect several results about finding sufficiently large sub-linkages of a given linkage that avoid certain parts of the ambient graph. We start with a way of enumerating the vertices of a linkage.

Lemma 4.9. *Let \mathfrak{G} be a graph of size $s := |V|$. If \mathfrak{G} has a unique tight A - B -linkage L of size m with $\bigcup L = V$, then there exist an enumeration $(u_i)_{i < s-m}$ of $V \setminus A$ and an enumeration $(v_i)_{i < s-m}$ of $V \setminus B$ such that, for every i ,*

$$\langle u_i, v_i \rangle \in E \quad \text{and} \quad \langle C_i, D_i \rangle \text{ is a separation of order } m,$$

where $C_i := A \cup \{u_j \mid j < i\}$ and $D_i := B \cup \{v_j \mid j \geq i\}$.

Proof. We proceed by induction on s . If there exists a separation $\langle C, D \rangle$ of order m with $A \subseteq C$ and $B \subseteq D$, we can apply the inductive hypothesis to (i) the subgraph $\mathfrak{G}|_C$ and the sets A and $C \cap D$; and to (ii) the subgraph $\mathfrak{G}|_D$ and the sets $C \cap D$ and B . (Note that each path of L must intersect $C \cap D$. Consequently, the restriction of L to the respective subset is a linkage between $C \cap D$ and the other set. Furthermore, it is also unique since, if there were another such linkage, we could combine it with the other part of L to obtain a linkage in \mathfrak{G} different from L .) The concatenations of the resulting sequences has the desired properties.

Consequently, we may assume that there is no such separation $\langle C, D \rangle$. If there are no edges in $\bigcup L$, then we have $A = B$. By uniqueness of L , this implies that $A = V = B$ and there is nothing to do. Hence, we may assume that there is some edge e in $\bigcup L$. By uniqueness of L , the graph $\mathfrak{G} - e$ does

not have a tight A - B -linkage of size m . Consequently, we can use Menger's Theorem to find a separation $\langle C, D \rangle$ of $\mathfrak{G} - e$ of order $m - 1$ with $A \subseteq C$ and $B \subseteq D$. As $\langle C, D \rangle$ is not a separation of \mathfrak{G} , we must have $e = \langle u, v \rangle$ where $u \in C \setminus D$ and $v \in D \setminus C$. It follows that $\langle C \cup \{v\}, D \rangle$ and $\langle C, D \cup \{u\} \rangle$ are both separations of \mathfrak{G} of order m . We apply the inductive hypothesis to (i) the subgraph $\mathfrak{G}|_{C \cup \{v\}}$ and the sets A and $(C \cap D) \cup \{v\}$; and to (ii) the subgraph $\mathfrak{G}|_{D \cup \{u\}}$ and the sets $(C \cap D) \cup \{u\}$ and B . The two concatenations of the resulting sequences (inserting u in the middle of the first one, and v in the middle of the second one) has the desired properties. \square

The next proposition is the main result of this section. It tells us that if we have a linkage and a collection of connected graphs, we can find one of these graphs that is avoided by a sufficiently many paths from the linkage.

Proposition 4.10. *Let $g \geq 1$, $h \geq 3$, $m \geq M(g, h)$, and $n := 2mN(g, h)$. Suppose that $\mathfrak{G} = \langle V, E \rangle \in \mathcal{G}(g, h)$ contains*

- ♦ *an A - B -linkage L of size m and*
- ♦ *a set K of disjoint connected subgraphs of \mathfrak{G} of size $|K| \geq n$.*

Then there exist some graph $\mathfrak{R} = \langle U, F \rangle \in K$ and a subset $L_0 \subseteq L$ of size $|L_0| \geq \varepsilon(g, h) \cdot m$ that is an $(A \setminus U)$ -($B \setminus U$)-linkage disjoint from U .

Proof. We prove the claim by induction on $|V| + |E|$. If there is some vertex or edge that does not appear in L or in K , we can delete it from the graph. Similarly, if some edge appears both in some $P \in L$ and in some $R \in K$, we can contract it. In both cases the claim follows by inductive hypothesis. Therefore, we may assume that every edge $e \in E$ appears either in L or in K , but not in both. If there is a non-isolated vertex that does not appear in L , we can contract one of the outgoing edges and again apply the inductive hypothesis. If some isolated vertex does not appear in L , then K contains a graph $\mathfrak{R} = \langle U, F \rangle$ consisting of just this vertex and L is disjoint from U . So in this case, we are done. Consequently, we may assume that $\bigcup L$ contains every vertex of \mathfrak{G} . Finally, if there exists a second A - B -linkage L' of size m , then L' must use some edge from K . Hence, we can contract this edge as above and the claim follows again by inductive hypothesis.

Thus, we may assume that L is the only A - B -linkage of size m and that it contains every vertex of \mathfrak{G} . This means we can apply Lemma 4.9. Let $\langle C_i, D_i \rangle_{i < |V| - m + 1}$ be the sequence of separations obtained in this way. For each $R \in K$, we set

$$I(R) := \{ i \mid R \text{ contains a vertex in } C_i \cap D_i \}.$$

Note that each such set is non-empty, since every $R \in K$ has at least one vertex and every vertex belongs to $C_i \cap D_i$, for some i . Furthermore, the sets $I(R)$ are convex since every $R \in K$ is connected and every path between $C_i \cap D_i$ and $C_j \cap D_j$ must contain a vertex from every set $C_k \cap D_k$ with $i \leq k \leq j$. Since $|C_i \cap D_i| \leq m$ there are at most m graphs $R \in K$ with $i \in I(R)$. This implies that, for every set $J \subseteq [|V| - m + 1]$ there are at most $m \cdot |J|$ graphs $R \in K$ with $I(R) \cap J \neq \emptyset$. Thus, every set $J \subseteq [|V| - m + 1]$ of size $|J| < \frac{1}{m}|K| \leq n/m$ is disjoint from at least one set $I(R)$, $R \in K$. We construct a sequence of graphs $R_0, \dots, R_{n/m-1} \in K$ and a corresponding sequence of indices $j_0, \dots, j_{n/m-1} \in [|V| - m + 1]$ as follows. We choose $R_0 \in K$ such that $\max I(R_0)$ is minimal and we set $j_0 := \max I(R_0)$. For the inductive step, suppose that R_i and j_i are already defined. By the above remark, there is some $R_{i+1} \in K$ with $I(R_{i+1}) \cap \{j_0, \dots, j_i\} = \emptyset$. We choose R_{i+1} such that $\max I(R_{i+1})$ is minimal and we set $j_{i+1} := \max I(R_{i+1})$. We claim that

$$j_0 < \dots < j_{n/m-1} \quad \text{and} \quad I(R_i) \cap I(R_k) = \emptyset, \quad \text{for } i < k.$$

For the first claim, note that $j_i \notin I(R_{i+1})$ by choice of R_{i+1} . This implies that $j_i \neq j_{i+1}$. Furthermore, if $j_i > j_{i+1}$ then, when choosing R_i , we would have chosen R_{i+1} instead.

For the second claim, suppose that $I(R_i) \cap I(R_k) \neq \emptyset$, for some $i < k$. By choice of R_k we have $\max I(R_i) = j_i \notin I(R_k)$. This implies that $j_{i+1} = \max I(R_k) < \max I(R_i)$, in contradiction to the first claim.

Let P_0, \dots, P_{m-1} be an enumeration of L . If some R_k is disjoint from at least $\varepsilon(g, h)m$ of the P_i , then $\mathfrak{G} - R_k$ contains an A - B -linkage of the desired size and we are done.

For a contradiction, suppose otherwise. By assumption, the set

$$J := \{ \langle i, k \rangle \mid P_i \cap R_k = \emptyset \}$$

has size at most

$$|J| \leq \varepsilon(g, h) \cdot m \cdot \frac{n}{m} = \frac{M(g, h)}{2M(g, h)} \cdot \frac{n}{m} = \frac{n}{2m}.$$

Consequently, at least $\frac{n}{2m}$ of the R_k intersect with all P_i . Furthermore, if we traverse P_i from A to B , it follows by choice of the enumeration $R_0, \dots, R_{n/m-1}$ that every vertex in $P_i \cap R_k$ is before every vertex in $P_i \cap R_l$, for $k < l$. Therefore, these $\frac{n}{2m} = N(g, h)$ graphs R_k together with all P_i form an $m \times \frac{n}{2m}$ pseudo-grid (after contracting some edges if necessary). By Proposition 4.5, the resulting minor of \mathfrak{G} does not belong to $\mathcal{G}(g, h)$. As $\mathcal{G}(g, h)$ is closed under minors, this contradicts the fact that $\mathfrak{G} \in \mathcal{G}(g, h)$. \square

Definition 4.11. Let \mathfrak{G} be a graph and $S \subseteq V$.

(a) Path P is S -proper if it has either zero or more than one edge, its end-points both belong to S , but none of its other vertices do.

(b) A linkage L is S -proper if every path in L is S -proper.

Corollary 4.12. Let $g \geq 1$, $h \geq 3$, $m \geq M(g, h)$, and set $n := 2mN(g, h)$. Suppose that $\mathfrak{G} \in \mathcal{G}(g, h)$ has

- an S -proper A - B -linkage L of size m with $A, B \subseteq S$ and
- an S -proper linkage K of size n .

Then, for some $\mathfrak{H} = \langle U, F \rangle \in K$, there exists a $(S \setminus U)$ -proper $(A \setminus U)$ -($B \setminus U$)-linkage L' of size at least $\varepsilon(g, h) \cdot m$ such that $\bigcup L' \subseteq \bigcup L \setminus U$.

Proof. Let \mathfrak{F} be the subgraph of \mathfrak{G} obtained from the union of the paths in L and K . Applying Proposition 4.10 to \mathfrak{F} , we obtain a graph $\mathfrak{H} = \langle U, F \rangle \in K$ and a subset $L_o \subseteq L$ of size $|L_o| \geq \varepsilon(g, h) \cdot m$ that forms an $(A \setminus U)$ -($B \setminus U$)-linkage disjoint from U .

It follows that there exists an $(A \setminus U)$ -($B \setminus U$)-linkage L' of size $|L'| \geq \varepsilon(g, h) \cdot m$ that is disjoint from U and such that $\bigcup L'$ is minimal. We claim that this linkage L' is $(S \setminus U)$ -proper. As no edge of \mathfrak{F} has both of its end-points in S , no path in L' has exactly one edge. For a contradiction, suppose that there exists a path $P \in L'$ that has an internal vertex v which belongs to S . Since L is S -proper, at most one neighbour of v belongs to $\bigcup L'$. By

the same argument, at most one neighbour of v belongs to $\bigcup K$. Thus, v has exactly two neighbours in \mathfrak{F} . As every neighbour of v in \mathfrak{F} belongs to $\bigcup L'$ or to $\bigcup K$ it follows that v is an end-point of some path in L' and of some path in K . This implies that $v \in A \cup B$. Consequently, some proper subpath P_o of P forms a path in \mathfrak{F} connecting $A \setminus U$ and $B \setminus U$. A contradiction to the minimality of $\bigcup L'$. \square

Corollary 4.13. Let $g \geq 1$, $h \geq 3$, and $n \geq 0$, and set $d := 2N(g, h)$ and

$$m_o := \frac{1}{\varepsilon(g, h)^n} \quad \text{and} \quad m_{i+1} := \frac{d(1+d)^i}{\varepsilon(g, h)^n}.$$

Suppose that $\mathfrak{G} \in \mathcal{G}(g, h)$ is a graph such that, for each $i \leq n$, there exists a S -proper A_i - B_i -linkage L_i of size m_i with $A_i, B_i \subseteq S$. Then \mathfrak{G} contains $n+1$ pairwise disjoint S -proper paths P_o, \dots, P_n such that P_i connects A_i and B_i .

Proof. We prove the claim by induction on n . For $n = 0$, we can take any $P_o \in L_o$ since $m_o > 0$. Hence, suppose that $n > 0$. For $i < n$, we denote by $K_i \subseteq L_n$ the set of all paths $P \in L_n$ such that there is no $(S \setminus P)$ -proper $(A_i \setminus P)$ -($B_i \setminus P$)-linkage of size at least $\varepsilon(g, h) \cdot m_i$ that is disjoint from P . Since

$$|L_i| = m_i \geq m_o \geq \varepsilon(g, h)^{-1} = 2M(g, h),$$

we can apply Corollary 4.12 to L_i and K_i and it follows that $|K_i| < dm_i$. Consequently,

$$\begin{aligned} |K_o \cup \dots \cup K_{n-1}| &\leq d(m_o + \dots + m_{n-1}) \\ &= \frac{d}{\varepsilon(g, h)^n} [1 + d + d(1+d) + \dots + d(1+d)^{n-1}] \\ &= \frac{d}{\varepsilon(g, h)^n} \left[1 + d \frac{(1+d)^n - 1}{1+d-1} \right] \\ &= \frac{d(1+d)^n}{\varepsilon(g, h)^n} = m_n. \end{aligned}$$

Hence, there is some path $P_n \in L_n \setminus (K_o \cup \dots \cup K_{n-1})$. We have found a path P_n from A_n to B_n such that, for every $i < n$, there exists some $(S \setminus P_n)$ -proper $(A_i \setminus P_n)$ – $(B_i \setminus P_n)$ -linkage L'_i of size at least $\varepsilon(g, h) \cdot m_i$. We can now apply the inductive hypothesis to $\mathfrak{G} \setminus P_n$ to obtain the remaining paths P_o, \dots, P_{n-1} . \square

Externally Linked Sets

We are finally able to connect the above combinatorial results with the notion of tree-width. In the proposition below we show that every graph of large tree-width must contain a highly connected set.

Definition 4.14. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph.

(a) A set $S \subseteq V$ is *externally linked* in a set $C \subseteq V$ if, for every choice of subsets $A, B \subseteq S$ with $|A| = |B|$, not necessarily disjoint, there exists a S -proper A - B -linkage L of size $|A|$ with $\bigcup L \subseteq C$.

(b) A separation $\langle A, B \rangle$ *left-contains* a minor-embedding $(C_u)_{u \in U}$ if

$$|A \cap B| = |U| \quad \text{and} \quad |C_u \cap B| = 1, \quad \text{for all } u \in U.$$

The following result is a more complicated variant of the construction from the proof of Theorem 2.14.

Proposition 4.15. Let T be a tree with $w \geq 1$ vertices and let \mathfrak{G} be a graph with $\text{twd } \mathfrak{G} \geq \frac{3}{2}w$. Then there exists a separation $\langle A, B \rangle$ with the following properties.

- ♦ $|A \cap B| = w$.
- ♦ The subgraph induced by $B \setminus A$ is connected.
- ♦ Every vertex $v \in A \cap B$ has a neighbour in $B \setminus A$.
- ♦ $A \cap B$ is externally linked in B .
- ♦ $\langle A, B \rangle$ left-contains a minor-embedding of T .

Proof. We fix an enumeration t_o, \dots, t_{w-1} of T such that every vertex t_i (with $i > o$) is adjacent to some t_j with $j < i$. Let T_i be the subtree of T induced by $\{t_o, \dots, t_{i-1}\}$. According to Theorem 3.2, \mathfrak{G} has a $\frac{3}{2}w$ -bramble \mathcal{B} .

We choose \mathcal{B} maximal. This implies that, if a connected set $C \subseteq V$ contains some $B \in \mathcal{B}$, then $C \in \mathcal{B}$. (Otherwise, we could add C to \mathcal{B} without destroying the bramble properties.) By definition of a bramble, for every set $Z \subseteq V$ of size $|Z| < \frac{3}{2}w$, there is some $B \in \mathcal{B}$ disjoint from Z . Since every pair of sets $B, B' \in \mathcal{B}$ touch, it follows by maximality of \mathcal{B} that there is a unique connected component C of $\mathfrak{G} - Z$ with $C \in \mathcal{B}$. We denote it by $\beta(Z) := C$. Let us call a separation $\langle A, B \rangle$ *\mathcal{B} -right-directed* if $\beta(A \cap B) \subseteq B$. We will prove the following sequence of claims. The last three of them will then imply the existence of a separation with the desired properties.

(i) We start by showing that there exists a separation $\langle A, B \rangle$ with the following four properties.

- (1) $\langle A, B \rangle$ has order k , for some $1 \leq k \leq w$.
- (2) $\langle A, B \rangle$ left-contains a minor-embedding of T_k .
- (3) $\langle A, B \rangle$ is \mathcal{B} -right-directed.
- (4) There is no \mathcal{B} -right-directed separation $\langle C, D \rangle$ of order less than k with $C \supseteq A$ and $D \subseteq B$.

To see this, fix a vertex $v \in \beta(\emptyset)$ and set

$$A := (V \setminus \beta(\emptyset)) \cup \{v\} \quad \text{and} \quad B := \beta(\emptyset).$$

Note that, since $\beta(\emptyset)$ and $\beta(\{v\})$ must touch and $\beta(\emptyset)$ is a connected component of \mathfrak{G} , we have $\beta(\{v\}) \subseteq \beta(\emptyset) = B$. Consequently, $\langle A, B \rangle$ is \mathcal{B} -right-directed. Furthermore, the set $\{v\}$ is a minor-embedding of $T_1 = \{t_o\}$ which is left-contained by the separation $\langle A, B \rangle$, as desired.

In the following, let $\langle A, B \rangle$ be a separation satisfying (i)–(4) such that the difference $|A| - |B|$ is maximal.

(ii) We claim that there is no \mathcal{B} -right-directed separation $\langle C, D \rangle$ of order k that is different from $\langle A, B \rangle$ and satisfies $C \supseteq A$ and $D \subseteq B$. For a contradiction, suppose such a separation $\langle C, D \rangle$ exists. By maximality of $|A| - |B|$ it follows that $\langle C, D \rangle$ cannot left-contain a minor-embedding of T_k . Note that the subgraph induced by $B \cap C$ does not contain k disjoint paths between $A \cap B$ and $C \cap D$ since, otherwise, we could add them to the minor-embedding of T_k in A to obtain one in C . Thus, $\langle C, D \rangle$ would left-contain

a minor-embedding of T_k , which we have seen is not possible. By Menger's Theorem it follows that there exists a separation $\langle C', D' \rangle$ of order less than k with $A \subseteq C'$ and $D \subseteq D'$. As $\beta(C' \cap D')$ touches $\beta(C \cap D)$ (both belong to B), and $\beta(C \cap D) \subseteq D \subseteq D'$, it follows that $\beta(C' \cap D') \subseteq D'$. Thus the separation $\langle C', D' \rangle$ contradicts (4).

(iii) Next we show that $B \setminus A$ is connected and every vertex in $A \cap B$ has a neighbour in $B \setminus A$. By definition, $U := \beta(A \cap B)$ is connected, contained in B , and disjoint from $A \cap B$. Hence, U is a connected component of $\mathfrak{G}|_{B \setminus A}$. Set

$$C := V \setminus \beta(A \cap B) \quad \text{and} \quad D := (A \cap B) \cup \beta(A \cap B).$$

Then $\langle C, D \rangle$ is a separation satisfying (i)–(4). By choice of $\langle A, B \rangle$ it follows that $\langle C, D \rangle = \langle A, B \rangle$. Consequently, $B \setminus A = D \setminus C = \beta(A \cap B)$ which, by definition of β , is connected.

For the second statement, suppose that there is some vertex $v \in A \cap B$ without a neighbour in $B \setminus A$. Then $\langle A, B \setminus \{v\} \rangle$ is a separation of order $k - 1$. Furthermore, $\beta(A \cap B \setminus \{v\})$ touches $\beta(A \cap B)$. This implies that $\beta(A \cap B \setminus \{v\}) \subseteq B \setminus \{v\}$, a contradiction to (4).

(iv) We claim that $k = w$. For a contradiction, suppose that $k < w$. Fix a minor-embedding $(C_i)_{i < k}$ of T_k in A such that every C_i contains a unique vertex v_i in $A \cap B$. By choice of the enumeration t_0, \dots, t_{w-1} there is some vertex t_i with $i < k$ that is a neighbour of t_k . By (iii), v_i has some neighbour $v_k \in B \setminus A$. Setting $A' := A \cup \{v_k\}$ we obtain a separation $\langle A', B \rangle$ of order $k + 1 \leq w$ that left-contains a minor-embedding $C_0, \dots, C_{k-1}, \{v_k\}$ of T_{k+1} . Furthermore, since $\beta(A' \cap B)$ touches $\beta(A \cap B)$, we have $\beta(A' \cap B) \subseteq B$. By (i), it follows that $\langle A', B \rangle$ satisfies (4). This contradicts the maximality of $|A| - |B|$.

(v) It remains to prove that $A \cap B$ is externally linked in B . For a contradiction, suppose otherwise. Then we can find two sets $X, Y \subseteq A \cap B$ of the same size $|X| = |Y|$ such that B does not contain a $(A \cap B)$ -proper X - Y -linkage of size $|X|$. Replacing X and Y by, respectively, $X \setminus Y$ and $Y \setminus X$ if necessary, we may assume that $X \cap Y = \emptyset$. Set $S := (A \cap B) \setminus (X \cup Y)$ and let F be the set of edges between X and Y . By Menger's Theorem, there exists a

separation $\langle C', D' \rangle$ of the graph $\mathfrak{G}|_{B \setminus S} - F$ of order $|C' \cap D'| < |X|$ with $X \subseteq C'$ and $Y \subseteq D'$. Setting $C := C' \cup S$ and $D := D' \cup S$ we obtain a separation $\langle C, D \rangle$ of $\mathfrak{G}|_B - F$ of order less than

$$|C \cap D| < |X| + |S| = |X \cup S| \leq |C \cap A \cap B| \leq |C \cap A|.$$

Consequently, $\langle A \cup C, D \rangle$ is a separation of \mathfrak{G} of order

$$\begin{aligned} |C \cap D| + |(A \setminus C) \cap D| &\leq |C \cap D| + |(A \cap B) \setminus C| \\ &= |C \cap D| + |A \cap B| - |A \cap C| \\ &< |A \cap B| = w. \end{aligned}$$

Hence, $\beta((A \cup C) \cap D)$ exists and it follows by (4) that $\beta((A \cup C) \cap D) \not\subseteq D$. As $\beta((A \cup C) \cap D)$ is connected and $\langle A \cup C, D \rangle$ is a separation, this implies that $\beta((A \cup C) \cap D) \subseteq A \cup C$. In the same way one can show that $\beta((A \cup D) \cap C) \subseteq A \cup D$.

Set $X := (A \cap B) \cup (C \cap D)$. As $A \cap B$, $(A \cup C) \cap D$, and $(A \cup D) \cap C$ are all subsets of X and every vertex of X belongs to at least two of these three sets, it follows that

$$\begin{aligned} 2|X| &\leq |A \cap B| + |(A \cup C) \cap D| + |(A \cup D) \cap C| \\ &\leq w + w - 1 + w - 1 = 3w - 2. \end{aligned}$$

That is $|X| \leq \frac{3}{2}w - 1$ and $\beta(X)$ exists. As $\beta(X)$ touches $\beta(A \cap B) \subseteq B$, we have $\beta(X) \subseteq B$. Since $\langle C, D \rangle$ is a separation of $\mathfrak{G}|_B - F$ and $\beta(X)$ is connected, it follows that $\beta(X) \subseteq C$ or $\beta(X) \subseteq D$. But $\beta(X)$ touches $(A \cup D) \cap C$, which implies that $\beta(X) \not\subseteq C$. Similarly, it follows that $\beta(X) \not\subseteq D$. A contradiction. \square

The preceding proposition can be used to find certain minors in graphs with large tree-width.

Proposition 4.16. *Let \mathfrak{H} be a connected graph with h vertices and f edges that is not a tree, and let $g \geq 1$. Every graph \mathfrak{G} that does not contain the $g \times g$ grid or \mathfrak{H} as a minor has tree-width at most*

$$\text{twd } \mathfrak{G} \leq 3[8h(h-2)(2g+h)(2g+1)]^{f-h} + \frac{3}{2}h.$$

Proof. Fix a graph \mathfrak{G} without the $g \times g$ grid and without \mathfrak{H} as a minor. Fix an enumeration t_0, \dots, t_{h-1} of the vertices of \mathfrak{H} , let T_0 be a spanning tree of \mathfrak{H} , and let f_0, \dots, f_m be an enumeration of all edges of \mathfrak{H} that do not belong to T_0 . Suppose that $f_i = \langle t_{p(i)}, t_{q(i)} \rangle$. As \mathfrak{H} is not a tree, we have $h \geq 3$ and $m \geq 0$. Define

$$k_0 := \frac{1}{\varepsilon(g, h)^m} \quad \text{and} \quad k_{i+1} := \frac{d(d+1)^i}{\varepsilon(g, h)^m}, \quad \text{for } i < m,$$

where $d := 2N(g, h)$. Let T be the tree obtained from T_0 by attaching to $t_{p(i)}$ and to $t_{q(i)}$ $k_i - 1$ new leaves each, for $i \leq m$. (If $p(i) = p(j)$ or $p(i) = q(j)$, we attach new leaves several times to the same vertex.) It follows that T has

$$\begin{aligned} w &:= h + 2(k_0 - 1 + \dots + k_m - 1) \\ &= h + 2 \left[\frac{1 + d + d(d+1) + \dots + d(d+1)^{m-1}}{\varepsilon(g, h)^m} \right] - 2(m+1) \\ &= h + 2 \left[\frac{1 + d \cdot \frac{(d+1)^m - 1}{d+1-1}}{\varepsilon(g, h)^m} \right] - 2(m+1) \\ &= h + 2 \left[\frac{d+1}{\varepsilon(g, h)} \right]^m - 2(m+1) \\ &\leq h + 2 \left[\frac{d+1}{\varepsilon(g, h)} \right]^m \end{aligned}$$

vertices. We will show that $\text{twd } \mathfrak{G} < \frac{3}{2}w$. Since

$$\frac{3}{2}w \leq \frac{3}{2}h + 3 \left[\frac{d+1}{\varepsilon(g, h)} \right]^m \leq \frac{3}{2}h + 3[8h(h-2)(2g+h)(2g+1)]^{f-h}$$

the proposition then follows. For a contradiction, suppose that

$$\text{twd } \mathfrak{G} \geq \frac{3}{2}w.$$

Then we can use Proposition 4.15 to find a separation $\langle A, B \rangle$ such that

- ♦ $|A \cap B| = w$,
- ♦ $S := A \cap B$ is externally linked in B , and
- ♦ $\langle A, B \rangle$ left-contains some minor-embedding $(C_i)_{i < w}$ of T .

Let $(C'_i)_{i < h}$ be the minor-embedding of T_0 obtained from $(C_i)_{i < w}$ by merging all components C_i belonging to one of the new leaves with the component of their respective parents. Then $|S \cap C_i| = 1$ implies that

$$|S \cap C'_{p(i)}|, |S \cap C'_{q(i)}| \geq k_i, \quad \text{for all } i \leq m.$$

As S is externally linked in B , we can find, for every $i \leq m$, some S -proper $(S \cap C'_{p(i)})$ – $(S \cap C'_{q(i)})$ -linkage L_i of size k_i . Since \mathfrak{H} is a minor of $\mathfrak{R}_{h,h}$, we have $\mathfrak{G} \in \mathcal{G}(g, h)$ and we can use Corollary 4.13 to find $m+1$ pairwise disjoint S -proper paths P_0, \dots, P_m in B such that P_i connects $S \cap C'_{p(i)}$ with $S \cap C'_{q(i)}$. Contracting every C_i to a single vertex and every P_i to a single edge, we obtain an \mathfrak{H} -minor in \mathfrak{G} . A contradiction. \square

The following corollary concludes the proof of Theorem 4.1.

Corollary 4.17. *Every graph \mathfrak{G} without a $g \times g$ grid as a minor has tree-width at most*

$$\text{twd } \mathfrak{G} \leq g^{8g^2}.$$

Proof. Let \mathfrak{H} be the $g \times g$ grid. If $g \leq 1$, the claim is trivial. Hence, suppose that $g \geq 2$. Then \mathfrak{H} is not a tree and it has $h = g^2$ vertices and $f = 2g(g-1)$ edges. By Proposition 4.16 it follows that every graph with no \mathfrak{H} -minor has tree-width at most

$$\begin{aligned} \text{twd } \mathfrak{G} &\leq 3[8h(h-2)(2g+h)(2g+1)]^{f-h} + \frac{3}{2}h \\ &= 3[8g^2(g^2-2)(2g+g^2)(2g+1)]^{2g(g-1)-g^2} + \frac{3}{2}g^2 \\ &\leq 4[8g^2 \cdot g^2 \cdot 2g^2 \cdot 3g]^{g^2-g} \\ &\leq 4[48g^7]^{g^2-g} \\ &\leq g^{8g^2}. \end{aligned} \quad \square$$

5 Branch-Decompositions and Tangles

There is a dual to the notions of tree-width and brambles that is based on decompositions of the edges instead of the vertices of a graph. While perhaps less intuitive it turns out that these dual versions are slightly better behaved than the vertex-based ones.

Definition 5.1. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph.

(a) A *branch-decomposition* of \mathfrak{G} is a pair $\langle T, \sigma \rangle$ where T is an undirected ternary tree and σ is a bijection between the leaves of T and the edges of \mathfrak{G} . Frequently we will denote a branch-decomposition simply by the function σ .

(b) Let $\langle T, \sigma \rangle$ be a branch-decomposition. Every edge $\langle u, v \rangle$ of T partitions the set E into two classes: the edges labelling leaves on one side of $\langle u, v \rangle$ and the edges labelling leaves on the other side. Formally, we define a function $\hat{\sigma}$ mapping each edge $\langle u, v \rangle$ of T to the set

$$\hat{\sigma}(\langle u, v \rangle) := \left\{ \sigma(w) \mid w \text{ is a leaf of the component of } T - \langle u, v \rangle \text{ that contains } v \right\}.$$

We call $\hat{\sigma}$ the *edge-flow* associated with σ .

(c) The *support* of a set $X \subseteq E$ is the set

$$\text{supp}(X) := \{ v \in V \mid v \text{ is an end-point of some edge in } X \},$$

and its *rank* $\text{rk}(X)$ is the cardinality

$$\text{rk}(X) := |\text{supp}(X) \cap \text{supp}(E \setminus X)|.$$

(d) The *width* of a branch-decomposition $\langle T, \sigma \rangle$ is

$$\sup \{ \text{rk}(\hat{\sigma}(\langle u, v \rangle)) \mid \langle u, v \rangle \text{ an edge of } T \}.$$

The *branch-width* of \mathfrak{G} is the minimal width of a branch-decomposition of \mathfrak{G} .


Examples. (a) If $|E| \leq 1$, then \mathfrak{G} has branch-width 0.

(b) If \mathfrak{G} has an edge e between two non-leaves, its branch-width is at least 2. To see this, consider a branch-decomposition $\langle T, \sigma \rangle$. Then there is some leaf u of T with $\sigma(u) = e$. Let v be the neighbour of u and let x, y be the two end-points of e . As x and y have each at least one other outgoing edge, it follows that $\text{supp}(\{e\}) = \{x, y\} \subseteq \text{supp}(E \setminus \{e\})$. Hence,

$$\text{rk}(\hat{\sigma}(\langle v, u \rangle)) = \text{rk}(\{e\}) = |\{x, y\}| = 2.$$

In particular, the width of σ is at least 2.

(c) Every tree has branch-width 2. For instance, the tree on the left has the branch-decomposition on the right, where we have labelled every edge with the corresponding rank.



Before proving the relationship between branch-width and tree-width, let us collect a few properties of the rank function. The second one is called *submodularity* and it is the main reason why branch-width is better behaved than tree-width.

Lemma 5.2. Let $X, Y \subseteq E$.

(a) $\text{rk}(X) = \text{rk}(E \setminus X)$

(b) $\text{rk}(X \cap Y) + \text{rk}(X \cup Y) \leq \text{rk}(X) + \text{rk}(Y)$

Proof. (a) The definition of $\text{rk}(X)$ is symmetric with respect to X and its complement.

(b) Let

$$\begin{aligned} A &:= \text{supp}(X \cap Y), & C &:= \text{supp}(Y \setminus X), \\ B &:= \text{supp}(X \setminus Y), & D &:= \text{supp}(E \setminus (X \cup Y)). \end{aligned}$$

Then $2 \cdot |B \cap C| - |A \cap B \cap C| - |B \cap C \cap D| \geq 0$ implies, by the inclusion-exclusion principle that

$$\begin{aligned}
& \text{rk}(X \cap Y) + \text{rk}(X \cup Y) \\
&= |A \cap (B \cup C \cup D)| + |(A \cup B \cup C) \cap D| \\
&= [|A \cap B| + |A \cap C| + |A \cap D| - |A \cap B \cap C| - |A \cap B \cap D| \\
&\quad - |A \cap C \cap D| + |A \cap B \cap C \cap D|] \\
&\quad + [|A \cap D| + |B \cap D| + |C \cap D| - |A \cap B \cap D| - |A \cap C \cap D| \\
&\quad - |B \cap C \cap D| + |A \cap B \cap C \cap D|] \\
&= |A \cap B| + |A \cap C| + 2 \cdot |A \cap D| + |B \cap D| + |C \cap D| \\
&\quad - |A \cap B \cap C| - 2 \cdot |A \cap B \cap D| - 2 \cdot |A \cap C \cap D| - |B \cap C \cap D| \\
&\quad + 2 \cdot |A \cap B \cap C \cap D| \\
&\leq [|A \cap B| + |A \cap C| + 2 \cdot |A \cap D| + |B \cap D| + |C \cap D| \\
&\quad - |A \cap B \cap C| - 2 \cdot |A \cap B \cap D| - 2 \cdot |A \cap C \cap D| - |B \cap C \cap D| \\
&\quad + 2 \cdot |A \cap B \cap C \cap D|] \\
&\quad + [2 \cdot |B \cap C| - |A \cap B \cap C| - |B \cap C \cap D|] \\
&= [|A \cap C| + |B \cap C| - |A \cap B \cap C|] \\
&\quad + [|A \cap D| + |B \cap D| - |A \cap B \cap D|] \\
&\quad - [|A \cap C \cap D| + |B \cap C \cap D| - |A \cap B \cap C \cap D|] \\
&\quad + [|A \cap B| + |C \cap B| - |A \cap B \cap C|] \\
&\quad + [|A \cap D| + |C \cap D| - |A \cap C \cap D|] \\
&\quad - [|A \cap B \cap D| + |C \cap B \cap D| - |A \cap C \cap B \cap D|] \\
&= [|A \cap B| + |A \cap C| + |A \cap D| - |A \cap B \cap C| - |A \cap B \cap D| - |A \cap C \cap D| + |A \cap B \cap C \cap D|] \\
&\quad + [|A \cap C| + |B \cap C| - |A \cap B \cap C|] \\
&= |(A \cup B) \cap (C \cup D)| + |(A \cup C) \cap (B \cup D)| \\
&= \text{rk}(X) + \text{rk}(Y).
\end{aligned}$$

□

Let us show that branch-width and tree-width differ by at most a constant factor.

Proposition 5.3. *Let \mathfrak{G} be a graph with branch-width β . Then*

$$\beta \leq \text{twd } \mathfrak{G} \leq \max \left\{ \frac{3}{2}\beta, 2 \right\}.$$

Proof. If \mathfrak{G} has at most one edge, then $\beta = 0$ and $\text{twd } \mathfrak{G} \leq 2$. Hence, we may assume that \mathfrak{G} has at least two edges. As the addition or removal of isolated vertices does not change the branch-width or the tree-width, we may also assume that \mathfrak{G} has no isolated vertices.

For the upper bound, fix a branch-decomposition $\langle T, \sigma \rangle$ of \mathfrak{G} of width β . We define a tree-decomposition $(U_v)_{v \in T}$ with the same underlying tree as follows.

- ♦ If v is a leaf of T , we set $U_v := \text{supp}(\{\sigma(v)\})$.
- ♦ If v is an internal vertex, U_v contains all vertices u such that there are two leaves w, w' of T such that v lies on the path between w and w' and both edges $\sigma(w)$ and $\sigma(w')$ are incident with u .

We claim that $(U_v)_{v \in T}$ is the desired tree-decomposition of \mathfrak{G} . Clearly, every edge $e \in E$ is contained in some component U_v , namely the one at the leaf $v := \sigma^{-1}(e)$. Furthermore, for every vertex $u \in V$, the set $\{v \in T \mid u \in U_v\}$ of components containing u consists of the union of all paths connecting two leaves w and w' of T such that $\sigma(w)$ and $\sigma(w')$ are edges with end-point u . This is a connected set.

Hence, $(U_v)_{v \in T}$ is a tree-decomposition and it remains to bound its width. We have $|U_v| = 2$, for leaves $v \in T$. Suppose that v is an internal vertex of T with neighbours u_o, u_i, u_2 . Set

$$A_i := \text{supp}(\hat{\sigma}(v, u_i)) \quad \text{and} \quad B_i := \text{supp}(\hat{\sigma}(u_i, v)).$$

Then $\text{rk}(\hat{\sigma}(v, u_i)) = |A_i \cap B_i| \leq \beta$. Note that $U_v \cap U_{u_i} = A_i \cap B_i$. Since every vertex $w \in U_v$ lies in all components on some path between two

leaves of T , it follows that every $w \in U_v$ belongs to at least two of the sets $U_{u_0}, U_{u_1}, U_{u_2}$. Consequently,

$$\begin{aligned} 2 \cdot |U_v| &\leq |U_v \cap U_{u_0}| + |U_v \cap U_{u_1}| + |U_v \cap U_{u_2}| \\ &= |A_0 \cap B_0| + |A_1 \cap B_1| + |A_2 \cap B_2| \\ &= \text{rk}(\hat{\sigma}(v, u_0)) + \text{rk}(\hat{\sigma}(v, u_1)) + \text{rk}(\hat{\sigma}(v, u_2)) \leq 3\beta, \end{aligned}$$

that is, $|U_v| \leq \frac{3}{2}\beta$.

It remains to establish the lower bound. Let $(U_v)_{v \in T}$ be a tree-decomposition of \mathfrak{G} of minimal width. We will transform $(U_v)_v$ in several steps until it has the following four properties.

- (I) For every $e \in E$ there is some leaf $u \in T$ with $U_u = \text{supp}(e)$.
- (II) For every leaf $u \in T$, we have $U_u = \text{supp}(e)$, for some $e \in E$.
- (III) $U_u \neq U_v$, for all leaves $u \neq v$.
- (IV) Every vertex of T has at most 3 neighbours.

Once $(U_v)_{v \in T}$ has the above form, we obtain the desired branch-decomposition as follows. Let S be the tree obtained from T by removing every vertex v of degree 2 (and adding a new edge between the former neighbours of v); and let σ be the function mapping each leaf u of S to the unique edge $e \in E$ with $U_u = \text{supp}(e)$. By (i), (ii), and (iii) σ is a bijection. Hence, $\langle S, \sigma \rangle$ is a branch-decomposition of \mathfrak{G} and it remains to compute its width. Consider an edge $\langle u, v \rangle$ of S and set $A := \text{supp}(\hat{\sigma}(v, u))$ and $B := \text{supp}(\hat{\sigma}(u, v))$. Then

$$A \cap B \subseteq U_u \cap U_v \quad \text{implies} \quad \text{rk}(\hat{\sigma}(u, v)) = |A \cap B| \leq |U_u \cap U_v|.$$

Consequently, the width of σ is at most the width of $(U_v)_v$. To conclude the proof it therefore remains to describe the above simplification steps.


(i) For every edge $e = \langle w, w' \rangle \in E$, we pick some component U_v with $w, w' \in U_v$ (which must exist by definition of a tree-decomposition), and we attach a new leaf u to v and set $U_u := \{w, w'\}$.

(ii) If there is some leaf u of T such that U_u contains no edges, we can delete u from T without destroying the property of it being a tree-decomposition. If there is some leaf u such that U_u contains several edges,

then (i) implies that $U_u \subseteq U_v$, where v is the unique neighbour of u . Hence, we can again delete u from T .

(iii) Suppose that there are two leaves u and v with the same component $U_u = U_v$. By deleting v from T and we obtain a smaller tree-decomposition of the same width.

(iv) Suppose that $v \in T$ has $n > 3$ neighbours u_0, \dots, u_{n-1} . We add a new vertex v' to T with edges $\langle v, v' \rangle$ and $\langle v', u_i \rangle$, for $i > 1$. Deleting the edges $\langle v, u_i \rangle$, for $i > 1$, we obtain a new tree T' , where v has now exactly 3 neighbours (u_0, u_1, v') and v' has $n - 1$ neighbours ($v, u_2, u_3, \dots, u_{n-1}$).



Setting $U_{v'} := U_v$, we obtain a new tree-decomposition $(U_v)_{v \in T'}$ of the same width. Repeating this process, we can reduce the degree of every vertex to 3. \square

The aim of the reminder of this section is to derive a condition for the non-existence of branch-decompositions, in the same ways brambles can be used to find lower bounds for tree-width. As an intermediate step, we start by generalising the notion of a branch-decomposition. The resulting object is called k -flow and it is based on the edge-flow map $\hat{\sigma}$. Working with $\hat{\sigma}$ instead of σ has the advantage that we can replace the global requirement (that the edges be distributed over the leaves of the tree) by a local one: that, at each internal vertex, each edge appears in the label of some outgoing edge. Our generalisation concerns two aspects of such a map: (i) we allow leaves to be labelled by several edges and (ii) edges are allowed to be assigned to several leaves. The former is useful when we are dealing with 'partial branch-decompositions' where some parts of the graph are not yet fully decomposed, while the latter makes constructing decompositions easier, since we have weaker requirements to satisfy.

Definition 5.4. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) A k -flow is a pair $\langle T, \tau \rangle$ where T is a ternary tree and τ a function mapping each (directed) edge $\langle u, v \rangle$ of T to some set $\tau(u, v) \subseteq E$ such that the following conditions are satisfied.

- (F1) $\text{rk}(\tau(u, v)) < k$, for every edge $\langle u, v \rangle$.
- (F2) $\tau(u, v) = E \setminus \tau(v, u)$, for every edge $\langle u, v \rangle$.
- (F3) $\tau(v, u_o) \cup \tau(v, u_1) \cup \tau(v, u_2) = E$,
for every internal vertex $v \in T$ with
neighbours u_o, u_1, u_2 .

(b) A k -flow $\langle T, \tau \rangle$ is *exact* if, for every internal vertex $v \in T$ with neighbours u_o, u_1, u_2 , the sets

$$\tau(v, u_o), \tau(v, u_1), \tau(v, u_2)$$

are pairwise disjoint.

(c) Let $\langle T, \tau \rangle$ is be a k -flow. For a leaf $u \in T$, we write

$$\tau^\circ(u) := \tau(v, u), \quad \text{where } v \text{ is the neighbour of } u.$$

A *block* of τ is a set of the form $\tau^\circ(u)$ where u is a leaf. The set of all blocks is denoted

$$B(\tau) := \{ \tau^\circ(u) \mid u \text{ leaf of } T \}.$$

Let us first show how to get rid of the possible overlap in a k -flow.

Lemma 5.5. Let \mathfrak{G} be a finite undirected graph. For every k -flow $\langle T, \tau \rangle$ there exists an exact k -flow $\langle S, \sigma \rangle$ with $B(\sigma) \subseteq \downarrow B(\tau)$.

Proof. Let us introduce some notation. We turn T into a rooted tree by fixing an arbitrary vertex r . The distance of a given vertex v from r is denoted by $d(v)$. Given a k -flow σ , we call a path u, v, u' of length 2 an *inexactness* if

$$\sigma(v, u) \cap \sigma(v, u') \neq \emptyset.$$

Let $I(\sigma)$ be the set of all inexactnesses of σ . The *defect* is

$$D(\sigma) := \sum_{\langle u, v, u' \rangle \in I(\sigma)} 3^{-d(v)}.$$

Thus, the defect measures how far a k -flow is from being exact with precedence given to inexactnesses close to the root.

We choose a k -flow σ on T satisfying the following conditions.

- (I) $\sigma(u, v) \subseteq \tau(u, v)$ for all edges $\langle u, v \rangle$.
- (II) The total rank $\sum_{u, v} \text{rk}(\sigma(u, v))$ is minimal.
- (III) Subject to (ii), the defect $D(\sigma)$ is minimal.

We claim that σ is exact. For a contradiction, suppose otherwise and fix some inexactness $\langle u, v, u' \rangle \in I(\sigma)$. By symmetry, we may assume that $d(u) \geq d(u')$. Set $A := \sigma(v, u)$ and $A' := \sigma(v, u')$, and define

$$\rho(x, y) := \begin{cases} A \setminus A' & \text{if } \langle x, y \rangle = \langle v, u \rangle, \\ E \setminus (A \setminus A') & \text{if } \langle x, y \rangle = \langle u, v \rangle, \\ \sigma(x, y) & \text{otherwise.} \end{cases}$$

We start by showing that $\text{rk}(A \setminus A') \geq \text{rk}(A)$. If $\text{rk}(A \setminus A') \geq k$, this follows immediately from the fact that $\text{rk}(A) < k$. Hence, we may assume that $\text{rk}(A \setminus A') < k$. Then ρ is a k -flow and it follows by (ii) that

$$\begin{aligned} 2 \cdot \text{rk}(A) &= \text{rk}(A) + \text{rk}(E \setminus A) \\ &= \text{rk}(\sigma(v, u)) + \text{rk}(\sigma(u, v)) \\ &\leq \text{rk}(\rho(v, u)) + \text{rk}(\rho(u, v)) \\ &= \text{rk}(A \setminus A') + \text{rk}(E \setminus (A \setminus A')) = 2 \cdot \text{rk}(A \setminus A'). \end{aligned}$$

Using the same argument, it follows that $\text{rk}(A') \leq \text{rk}(A' \setminus A)$. Furthermore, submodularity implies that

$$\begin{aligned} \text{rk}(A \setminus A') + \text{rk}(A' \setminus A) &= \text{rk}(A \cap (E \setminus A')) + \text{rk}(A' \cap (E \setminus A)) \\ &= \text{rk}(A \cap (E \setminus A')) + \text{rk}(A \cup (E \setminus A')) \\ &\leq \text{rk}(A) + \text{rk}(E \setminus A') \\ &= \text{rk}(A) + \text{rk}(A'). \end{aligned}$$

Together with $\text{rk}(A) \leq \text{rk}(A \setminus A')$ and $\text{rk}(A') \leq \text{rk}(A' \setminus A)$ it therefore follows that

$$\text{rk}(A) = \text{rk}(A \setminus A') \quad \text{and} \quad \text{rk}(A') = \text{rk}(A' \setminus A).$$

Consequently, the k -flow ρ defined above satisfies

$$\sum_{x,y} \text{rk}(\rho(x, y)) = \sum_{x,y} \text{rk}(\sigma(x, y)).$$

Let us compute $D(\rho)$. Note that $d(u) \geq d(u')$ implies $d(u) > d(v)$, i.e., $d(u) = d(v) + 1$. By construction, we have

$$\langle u, v, u' \rangle, \langle u', v, u \rangle \in I(\sigma) \setminus I(\rho)$$

while the only paths that might appear in $I(\rho) \setminus I(\sigma)$ are $\langle v, u, w \rangle$, $\langle w, u, v \rangle$, $\langle v, u, w' \rangle$, and $\langle w', u, v \rangle$, where w and w' are the other two neighbours of u . Consequently,

$$\begin{aligned} D(\rho) &\leq D(\sigma) - 2 \cdot 3^{-d(v)} + 4 \cdot 3^{-d(u)} \\ &= D(\sigma) - 2 \cdot 3^{-d(v)} + 4 \cdot 3^{-d(v)-1} \\ &= D(\sigma) - \frac{2}{3} \cdot 3^{-d(v)} \\ &< D(\sigma), \end{aligned}$$

in contradiction to (iii). \square

Corollary 5.6. *Let \mathfrak{G} be a non-empty, finite, undirected graph. The following statements are equivalent.*

- (1) \mathfrak{G} has a branch-decomposition σ of width less than k .
- (2) \mathfrak{G} has a k -flow τ with $|X| \leq 1$, for all $X \in B(\tau)$.

Proof. (1) \Rightarrow (2) Let $\langle T, \sigma \rangle$ be a branch-decomposition of width less than k . Then $\hat{\sigma}$ is a k -flow where each block is a singleton.

(2) \Rightarrow (1) Let $\langle T, \tau \rangle$ be a k -flow as above. By Lemma 5.5, we can assume that it is exact.

Let us start by showing that the blocks of τ are pairwise disjoint. For a contradiction, suppose that there are two leaves $u \neq u'$ of T and an edge $e \in \tau^\circ(u) \cap \tau^\circ(u')$. Let v_0, \dots, v_n be the path from u to u' . (Note that $\tau^\circ(u) \cap \tau^\circ(u') \neq \emptyset$ implies that $n > 1$.) By induction on i and the fact that τ is exact, it follows that $e \in \tau(v_{i+1}, v_i)$ and $e \in \tau(v_{n-i-1}, v_{n-i})$. Consequently,

$$\tau(v_i, v_{i-1}) \cap \tau(v_i, v_{i+1}) \neq \emptyset.$$

A contradiction.

Furthermore, every edge is contained in some block. To see this, fix an arbitrary vertex $v_0 \in T$. For every edge $e \in E$, exactness of τ implies that v_0 has a unique neighbour v_1 with $e \in \tau(v_0, v_1)$. Repeating this argument, we can construct a path v_0, v_1, \dots, v_n such that $e \in \tau(v_i, v_{i+1})$ for all i . As some point this path must reach a leaf v_n . Consequently, e is contained in the block $\tau(v_{n-1}, v_n)$.

If T has leaves u with $\tau^\circ(u) = \emptyset$, we can remove them from T . This can create vertices with exactly two neighbours. To obtain a valid k -flow we can remove each such vertex v and contract one of the incoming edges. In this way, we obtain a new k -flow with $\emptyset \notin B(\tau)$.

Hence, we may assume that, for every leaf u of T , there exists some $e \in E$ with $\tau^\circ(u) = \{e\}$. We define $\sigma(u) := e$. As the blocks of τ are pairwise disjoint, the resulting function is injective. Since the union of all block is equal to E , σ is also surjective. Consequently, $\langle T, \sigma \rangle$ is a branch-decomposition. As $\hat{\sigma} = \tau$, it further follows that the width of σ is less than k . \square

In order to derive lower bounds on the branch-width it would be nice to have some analogue to a bramble, that is, some kind of witness showing that the branch-width is at least k . In the following definition we present one such witness called a *tangle*. Intuitively, tangles are something like a finite analogue of a non-principal ultrafilter, i.e., they can be seen as a family of ‘complicated’ or ‘big’ subsets of the given graph.

Definition 5.7. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph. A k -tangle is a collection $S \subseteq \wp(E)$ such that

- (τ_1) $\text{rk}(X) < k$, for all $X \in S$.
- (τ_2) $\text{rk}(X) < k$ implies $X \in S$ or $E \setminus X \in S$.
- (τ_3) $X, Y, Z \in S$ implies $X \cap Y \cap Z \neq \emptyset$.
- (τ_4) $X \in S$ implies $|X| > 1$.

Examples. (a) Let \mathfrak{K}_n be the complete graph with n vertices and let $X \subseteq E$. The following conditions are equivalent.

- (1) $\text{rk}(X) < n$.
- (2) $|\text{supp}(X)| < n$ or $|\text{supp}(E \setminus X)| < n$.
- (3) There is some vertex v such that X or $E \setminus X$ contains all edges incident with v .

It follows that \mathfrak{K}_n has a unique n -tangle

$$S := \{ X \subseteq E \mid \text{for some } v, X \text{ contains all edges incident with } v \}.$$

(b) The $n \times n$ grid has an n -tangle consisting of all sets $X \subseteq E$ of rank $\text{rk}(X) < n$ such that X contains all the edges of some column or some row.

The analogy between tangles and ultrafilters is explored in the following lemma.

Lemma 5.8. *Let S be a k -tangle and $X, Y \subseteq E$ sets of rank less than k .*

- (a) $X \in S$ iff $E \setminus X \notin S$.
- (b) $X, Y \in S$ and $\text{rk}(X \cap Y) < k$ implies $X \cap Y \in S$.
- (c) $X \subseteq Y$ and $X \in S$ implies $Y \in S$.

Proof. (a) follows by (τ_2) and (τ_3) since $X \cap (E \setminus X) = \emptyset$.

(b) For a contradiction, suppose that there are sets $X, Y \in S$ such that $\text{rk}(X \cap Y) < k$, but $X \cap Y \notin S$. Then $E \setminus (X \cap Y) \in S$, by (τ_2). But

$$X \cap Y \cap E \setminus (X \cap Y) = \emptyset$$

contradicts (τ_3).

(c) For a contradiction, suppose that $Y \notin S$. Then $E \setminus Y \in S$ by (τ_2). But

$$X \cap (E \setminus Y) \subseteq X \cap (E \setminus X) = \emptyset$$

contradicts (τ_3). \square

In particular, 1-tangles correspond to non-principal ultrafilters on the set of connected components of the given graph. Let us use this fact to prove the existence of 1-tangles.

Lemma 5.9. *A finite graph \mathfrak{G} has a 1-tangle if, and only if, it has a connected component with at least 2 edges.*

Proof. (\Leftarrow) Fix a connected component C such that $F := E|_C$ has at least two elements. Then

$$S := \{ X \subseteq E \mid \text{rk}(X) = 0, F \subseteq X \}$$

satisfies (τ_1)–(τ_4).

(\Rightarrow) Let S be a 1-tangle of \mathfrak{G} and fix some set $X \in S$ of minimal size. By (τ_1), X is the set of edges of a union $C_0 \cup \dots \cup C_{n-1}$ of connected components of \mathfrak{G} and (τ_4) implies that X contains at least two edges. Hence, it remains to prove that the above union consists of a single connected component. For a contradiction, suppose otherwise. Let $F_i := E|_{C_i}$ be the sets of edges in the component C_i . By minimality of X , we have $F_i \notin S$ which, by (τ_2), implies that $E \setminus F_i \in S$. Consequently, it follows by Lemma 5.8 (b) that $Z := (E \setminus F_0) \cap \dots \cap (E \setminus F_{n-1}) \in S$. But $Z = E \setminus X$. Hence, $X, Z \in S$ contradicts (τ_3). \square

Let us finally show that all the various combinatorial notions we have introduced in this section are equivalent.

Theorem 5.10. *A finite undirected graph $\mathfrak{G} = \langle V, E \rangle$ has a k -tangle if, and only if, its branch-width is greater than or equal to k .*

Proof. Let us first take a look at a few trivial cases. Every graph \mathfrak{G} has the 0-tangle $S = \emptyset$ (as there are no sets of rank less than 0) and the branch-width is obviously always at least 0.

It follows by (τ_4) (see Lemma 5.9) that a graph \mathfrak{G} has a 1-tangle if, and only if, it contains a path of length 2, i.e., if not every edge is isolated. This is exactly the case if its branch-width is at least 1.

For the remainder of the proof we may therefore assume that $k > 1$ and that \mathfrak{G} has at least two edges.

(\Rightarrow) For a contradiction, suppose that \mathfrak{G} has both a k -tangle S and a branch-decomposition $\langle T, \sigma \rangle$ of width less than k . Note that $|E| \geq 2$ implies that T has at least one edge. We call a (directed) edge $\langle u, v \rangle$ of T *big* if $\hat{\sigma}(u, v) \in S$. It follows by Lemma 5.8 (a) that, for every edge $\langle u, v \rangle$ of T , exactly one of $\hat{\sigma}(u, v)$ and $\hat{\sigma}(v, u)$ is big. Hence, there are exactly as many big pairs $\langle u, v \rangle$ as T has edges. Since a finite tree has one more vertex than edges, it follows that there is some vertex w such that $\hat{\sigma}(w, v)$ is not big, for any neighbour v of w . Thus, $\hat{\sigma}(w, v) \notin S$, for all neighbours v . Since $\text{rk}(\hat{\sigma}(w, v)) < k$, it therefore follows by (τ_2) that $\hat{\sigma}(v, w) \in S$ for all such v . We distinguish two cases.

If w is a leaf with neighbour v , then $\{\sigma(w)\} = \hat{\sigma}(v, w) \in S$ contradicts (τ_4) . And if w is an internal vertex with neighbours v_0, v_1, v_2 , then

$$\hat{\sigma}(w, v_0) \cap \hat{\sigma}(w, v_1) \cap \hat{\sigma}(w, v_2) = \emptyset$$

contradicts (τ_3) .

(\Leftarrow) Suppose that there is no k -tangle for \mathfrak{G} . By Corollary 5.6, it is sufficient to construct a k -flow τ such that $B(\tau)$ contains only singletons.

We will prove the following more general claim: for every downwards-closed set $A \subseteq \wp(E)$ containing all singletons, we will show that, if there is no k -tangle S with $S \cap A = \emptyset$, then \mathfrak{G} has a k -flow τ such that $B(\tau) \subseteq A$. Then the result follows with $A = \{\emptyset\} \cup \{\{e\} \mid e \in E\}$.

Let $A \subseteq \wp(E)$ be a downwards-closed set that contains all singletons and such that no k -tangle is disjoint from A . We construct the desired k -flow τ by induction on the number N of sets $X \subseteq E$ such that $\text{rk}(X) < k$ and $X, E \setminus X \notin A$.

First, suppose that $N = 0$. For $X \subseteq E$, we then have

$$\text{rk}(X) < k \quad \text{implies} \quad X \in A \quad \text{or} \quad E \setminus X \in A.$$

Hence, the set

$$S := \{X \mid E \setminus X \in A, \text{rk}(X) < k, |X| > 1\}$$

satisfies (τ_1) , (τ_2) , and (τ_4) . As S is disjoint from A , it cannot be a k -tangle. Consequently, S does not satisfy (τ_3) and we can find sets $X_0, X_1, X_2 \in S$ with $X_0 \cap X_1 \cap X_2 = \emptyset$. Let T be the tree consisting of a central vertex v with three leaves u_0, u_1, u_2 as neighbours, and let τ be the function with

$$\tau(u_i) := E \setminus X_i.$$

Then $\langle T, \tau \rangle$ is a k -flow for \mathfrak{G} with $B(\tau) = \{E \setminus X_0, E \setminus X_1, E \setminus X_2\} \subseteq A$.

For the inductive step, suppose that $N > 0$. Let $X \subseteq E$ be a set of minimal size such that $\text{rk}(X) < k$ and $X, E \setminus X \notin A$. Set $A_0 := A \cup \downarrow\{X\}$ and $A_1 := A \cup \downarrow\{E \setminus X\}$. Since no k -tangle is disjoint from A , the same is true for A_0 and for A_1 . Furthermore, the number N corresponding to these two sets is smaller than that for A . Consequently, we can use the inductive hypothesis to construct k -flows $\langle T_0, \tau_0 \rangle$ and $\langle T_1, \tau_1 \rangle$ such that $B(\tau_i) \subseteq A_i$. By Lemma 5.5, we may assume that τ_0 and τ_1 are exact. We call a leaf $u \in T_i$ *special* if $\tau_i^\circ(u) \notin A$. It follows that

$$\tau_0^\circ(u) \subseteq X, \quad \text{for every special leaf } u \text{ of } T_0.$$

If there were some special leaf u with $\tau_0^\circ(u) \subset X$, it would follow by minimality of X that

$$\text{rk}(\tau_0^\circ(u)) < k \quad \text{implies} \quad \tau_0^\circ(u) \in A \quad \text{or} \quad E \setminus \tau_0^\circ(u) \in A.$$

The former is not possible, since u is special; and the latter is impossible since

$$E \setminus X \subset E \setminus \tau_0^\circ(u) \in A \quad \text{implies} \quad E \setminus X \in \downarrow A = A.$$

Consequently, we have $\tau_0^\circ(u) = X$, for every special leaf $u \in T_0$. As τ_0 is exact, it follows that T_0 has at most one special leaf. If there are none, then $\langle T_0, \tau_0 \rangle$ is the desired k -flow and we are done. Hence, we may assume that

there is precisely one special leaf $w \in T_o$. Let w' be its neighbour. Then $\tau_o(w, w') = E \setminus X$ and $\tau_o(w', w) = X$. Note that w' cannot be also a leaf since $\tau_o(w, w') = E \setminus X \notin A_o$.

Let u_o, \dots, u_{n-1} be the special leaves of T_1 and let v_i be the neighbour of u_i . Then

$$\tau_1^\circ(u_i) \subseteq E \setminus X, \quad \text{which implies that} \quad \tau_1(u_i, v_i) \supseteq X.$$

Let S_o, \dots, S_{n-1} be n copies of the tree obtained from T_o by removing the leaf w and the edge $\langle w, w' \rangle$; and let T be the tree obtained from the disjoint union of T_1, S_o, \dots, S_{n-1} by identifying $u_i \in T_1$ with the copy of w' in S_i . Finally, let τ be the labelling on T that coincides with τ_o on each edge of S_i , and with τ_1 on the edges of T_1 . To see that τ is a k -flow, we only have to check (F3) for each copy of the vertex w' . (The labellings of the incoming edges at every other internal vertex have not changed.) Let x_i, y_i, v_i be the neighbours of the i -th copy of w' . Then

$$\begin{aligned} & \tau(w', x_i) \cup \tau(w', y_i) \cup \tau(w', v_i) \\ &= \tau_o(w', x_i) \cup \tau_o(w', y_i) \cup \tau_1(u_i, v_i) \\ &\supseteq \tau_o(w', x_i) \cup \tau_o(w', y_i) \cup X \\ &= \tau_o(w', x_i) \cup \tau_o(w', y_i) \cup \tau_o(w', w) = E. \end{aligned} \quad \square$$

For tree-width we obtain the following bounds.

Corollary 5.11. *Let \mathfrak{G} be a finite undirected graph.*

- (a) *If \mathfrak{G} has a k -tangle, then $\text{twd } \mathfrak{G} \geq k$.*
- (b) *If \mathfrak{G} has no k -tangle, then $\text{twd } \mathfrak{G} \leq \max\{\frac{3}{2}(k-1), 2\}$.*

Proof. Let k be the maximal number such that \mathfrak{G} has a k -tangle. Then it follows by Theorem 5.10 that k is equal to the branch-width of \mathfrak{G} . Consequently, Proposition 5.3 implies that $k \leq \text{twd } \mathfrak{G} \leq \max\{\frac{3}{2}k, 2\}$. \square

6 Well-Quasi-Orderings

Path-width and tree-width were introduced (or rather: rediscovered) in a long series of articles by Robertson and Seymour whose main result was a

solution to a conjecture of Wagner about a certain combinatorial property of the graph minor relation: that it forms what is called a *well-quasi-ordering*, i.e., a non-linear version of a well-ordering.

Definition 6.1. A preorder \leq on a class A is a *well-quasi-ordering* if it contains neither an infinite descending chain, nor an infinite antichain. \lrcorner

Lemma 6.2. *Let \leq be a preorder on a set A . The following conditions are equivalent.*

- (1) \leq is a well-quasi-ordering.
- (2) Every infinite sequence a_o, a_1, a_2, \dots in A contains two indices $i < k$ with $a_i \leq a_k$.
- (3) Every infinite sequence a_o, a_1, a_2, \dots in A contains an infinite increasing subsequence.
- (4) Every upwards-closed set $U \subseteq A$ is of the form $\uparrow C$ for some finite set C .
- (5) Every increasing sequence $U_o \subseteq U_1 \subseteq \dots$ of upwards-closed sets is eventually stationary, i.e., $U_n = U_{n+1} = U_{n+2} = \dots$.
- (6) For every set $X \subseteq A$, there is some finite set C with $C \subseteq X \subseteq \uparrow C$.

Proof. (3) \Rightarrow (1) is trivial.

(1) \Rightarrow (4) Let C be the set of all minimal elements of U . Then C is an antichain with $U = \uparrow C$. Furthermore, it follows by (1) that C is finite.

(4) \Rightarrow (5) The union $W := \bigcup_{i < \omega} U_i$ is also upwards closed. By assumption, we can find a finite subset $C \subseteq W$ with $W = \uparrow C$. As C is finite, there exists an index $k < \omega$, with $C \subseteq U_k$. Consequently,

$$W = \uparrow C \subseteq U_k \subseteq U_i \subseteq W, \quad \text{for all } i \geq k.$$

This implies that $U_k = U_i = W$, for $i \geq k$.

(5) \Rightarrow (2) Given a sequence a_o, a_1, \dots , we set $U_i := \uparrow \{a_o, \dots, a_i\}$. Then $U_o \subseteq U_1 \subseteq \dots$. By assumption, there is some index n with $U_n = U_{n+1}$. Consequently, $a_{n+1} \in U_n$, which implies that $a_i \leq a_{n+1}$, for some $i \leq n$.

(2) \Rightarrow (3) Let a_o, a_1, \dots be an infinite sequence. By the Theorem of Ramsey, there exists an infinite set $H \subseteq \omega$ of indices such that one of the following three conditions is satisfied:

- (I) $a_i < a_k$, for all $i < k$ in H .
- (II) $a_i = a_k$, for all $i < k$ in H .
- (III) $a_i > a_k$, for all $i < k$ in H .

By (2), there are two indices $i < k$ in H with $a_i \leq a_k$. This contradicts (iii). Consequently, (i) or (ii) must be true and a_0, a_1, \dots contains the infinite increasing subsequence $(a_i)_{i \in H}$.

(2) \Rightarrow (6) Let C be the set of all minimal elements of X . Then $X \subseteq \uparrow C$. If C is finite, we are done. For a contradiction, suppose otherwise. Then there exists an infinite sequence $(c_i)_{i < \omega}$ of distinct element $c_i \in C$. By (2), we can find two indices $i < k$ with $c_i \leq c_k$. As the c_i are distinct, it follows that $c_i < c_k$. Hence, c_k is not a minimal element of X . A contradiction.

(6) \Rightarrow (4) If U is upwards closed, then $C \subseteq U \subseteq \uparrow C$ implies $U = \uparrow C$. \square

Exercise 6.1. Let A and B be well-quasi-ordered sets.

- (a) Prove that the disjoint union $A \oplus B$ is also well-quasi-ordered.
- (b) Prove that the ordered sum $A + B$ is also well-quasi-ordered.
- (c) Prove that the direct product $A \times B$ is also well-quasi-ordered.

We won't present the proof of Wagner's conjecture as it would take us too far afield and would take too much space. Instead, we will prove two simpler statements about words and trees (which are quite useful on their own) and then just state the full result without proof.

Proposition 6.3 (Higman). *If A is a well-quasi-ordered set, the set A^* equipped with the subword ordering*

$$\begin{aligned} \langle a_0, \dots, a_{m-1} \rangle \leq \langle b_0, \dots, b_{n-1} \rangle \\ \text{ : iff } \text{ there exists an injective function } \varphi : [m] \rightarrow [n] \text{ such that} \\ a_i \leq b_{\varphi(i)}, \text{ for all } i < m, \end{aligned}$$

is also well-quasi-ordered.

Proof. For a contradiction, suppose that A^* contains an infinite sequence $(w_i)_{i < \omega}$ that does not contain a pair $w_i \leq w_j$ with $i < j$. Let us call such a

sequence a *bad sequence*. We inductively construct a sequence $u_0, u_1, \dots \in A^*$ as follows. We choose for u_n the minimal element of A^* such that there exists a bad sequence starting with u_0, \dots, u_n . The resulting sequence u_0, u_1, \dots is clearly bad. In particular, we have $u_i \neq \langle \rangle$ for all i . Hence, suppose that $u_i = a_i v_i$, for $a_i \in A$ and $v_i \in A^*$. As A is well-quasi-ordered, there exists an infinite subsequence u_{k_0}, u_{k_1}, \dots such that the sequence a_{k_0}, a_{k_1}, \dots of first elements is increasing. Since

$$\begin{aligned} u \not\leq bv & \quad \text{implies} \quad u \not\leq v, \\ a \leq b \text{ and } au \not\leq bv & \quad \text{implies} \quad u \not\leq v, \end{aligned}$$

for $a, b \in A$ and $u, v \in A^*$, it follows that the sequence

$$u_0, \dots, u_{k_0-1}, v_{k_0}, v_{k_1}, \dots$$

is also bad. A contradiction to the minimality of u_{k_0} . \square

We can extend this result from words to trees as follows. Let us call a function $\varphi : s \rightarrow t$ between two order-trees a *meet-embedding* if (i) it is injective; (ii) it preserves the tree order and binary meets (with respect to the tree order); and it is monotone on the labels, i.e., $s(v) \leq t(\varphi(v))$, for all vertices v .

Theorem 6.4 (Kruskal). *Let A be a well-quasi-ordered set. The set of all A -labelled finite order-trees equipped with the ordering*

$$s \leq t \quad \text{ : iff } \quad \text{there exists a meet-embedding } s \rightarrow t,$$

is also well-quasi-ordered.

Proof. We proceed as in the proof of Higman's Theorem. For a contradiction, suppose that there exists a bad sequence of trees. Inductively, let t_n be a tree with the minimal number of vertices such that some bad sequence starts with t_0, \dots, t_n . The resulting sequence t_0, t_1, \dots is bad. Let S_i be the set of all subtrees of t_i attached to the root.

We start by showing that the union $S := \bigcup_i S_i$ is well-quasi-ordered. Let s_0, s_1, \dots be an infinite sequence in S . For every $i < \omega$, there is some index

$n(i)$ with $s_i \in S_{n(i)}$. Let i be an index such that $n(i)$ is minimal. As s_i is a proper subtree of $t_{n(i)}$, it follows by choice of $t_{n(i)}$ that the sequence

$$t_0, \dots, t_{n(i)-1}, s_i, s_{i+1}, \dots$$

cannot be bad. Hence, there is an increasing pair $r \leq r'$ somewhere in this sequence. If $r = s_j$ and $r' = s_l$, for some $i \leq j < l$, the sequence $s_0, s_1, \dots, s_i, \dots$ is also not bad and the claim follows.

Hence, suppose otherwise. Then $r = t_j$, for some j and there is some index $l > j$ with $r' = t_l$ or $r' = s_l \leq t_l$. In both cases it follows that $t_j \leq r' \leq t_l$. A contradiction to the fact that the sequence $(t_i)_i$ is bad.

To conclude the proof, let $w_i \in S^*$ be an enumeration of S_i and let a_i be the label of the root of t_i . As A is well-quasi-ordered, there exists an infinite set $H \subseteq \omega$ such that the sequence $(a_i)_{i \in H}$ is increasing. Furthermore, we can use the Theorem of Higman to find indices $i < j$ in H with $w_i \leq w_j$. Consequently, there exist a function $f : S_i \rightarrow S_j$ with $s \leq f(s)$. Let φ be the union of the corresponding meet-embeddings together with the function mapping the root of t_i to the one of t_j . Then φ is a meet-embedding $t_i \rightarrow t_j$. Consequently, $t_i \leq t_j$. A contradiction. \square

Theorem 6.5 (Robertson, Seymour). *The minor relation is a well-quasi-ordering on the class of all finite graphs.*

Notes

Most of the material in this chapter originated in a long series of articles by Robertson and Seymour [160, 162, 161, 166, 163, 164, 165, 168, 167, 169, 170, 171, 172, 173, 174, 176, 175, 177, 178, 179, 180, 182, 181] which together constitute a proof of the Graph Minor Theorem 6.5.

Tree-width and tree-decompositions were originally introduced (under a different name) by Halin [100]. It was later rediscovered by Robertson and Seymour in [161]. Path-width was introduced in [160], and the material in Section 5 is taken mostly from [169]. The section on brambles follows [71].

The Excluded Tree Theorem and the Excluded Grid Theorem (Theorems 2.14 and 4.1) were proved in [160, 163]. Our proofs follow [69, 70] and [132],

respectively. The original bound in the Excluded Grid Theorem was astronomical, but it has been reduced to a polynomial in the meantime. For instance, [47] derive a bound of $cn^9 \log^d n$.

Embeddings of graphs in surfaces is a well-studied area of graph theory. An introduction can be found in [143]. To keep the proof of Proposition 4.2 simple we have been very wasteful. A more precise computation shows that every planar graph with n vertices is a minor of the $2n \times 2n$ grid (see [183] for a proof).

IX Crossing-Width

FOR GRAPHS, THERE EXISTS a generalisation of tree-width called *clique-width*. While tree-width is asymmetric in the sense that it treats edges and non-edges differently, clique-width treats them the same. The main differences are that clique-width is based on decompositions into non-overlapping parts and that it measures not the size of these parts, but the complexity of the connections between them. In this chapter we will not define clique-width itself, which only works for graphs, but a generalisation that can be used for arbitrary relational structures. Since, while equivalent, this more general measure does not coincide with clique-width over graphs, we will call it by a different name.

1 Partitions and Ranks

We start by taking a look at partitions of a given structure. Given a partition $P \cup Q$ of some Σ -structure \mathfrak{A} , we are interested in what information is needed to recover \mathfrak{A} from the two induced substructures $\mathfrak{A}|_P$ and $\mathfrak{A}|_Q$. Clearly, what we need to know is which tuples of P are connected to which tuples of Q via which relations. In other words, we need to know the classes of the following equivalence relation.

For technical reasons, we will always assume in the following that the signature Σ is finite, or at least that the arity of all relations in Σ are bounded. We denote this bound by r .

Definition 1.1. Let \mathfrak{A} be a relational structure and r the maximal arity of a relation of \mathfrak{A} .

(a) A formula $\varphi(\bar{x}; \bar{y})$ where the free variables are divided into two groups \bar{x} and \bar{y} is called *crossing* if at least one of the variables x_i and at least one y_j

actually occur in φ .

(b) Let $U \subseteq A$. The *crossing equivalence* on A^n over U is the relation

$$\begin{aligned} \bar{a} \simeq_U \bar{b} \quad : \text{iff} \quad \mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \Leftrightarrow \mathfrak{A} \models \varphi(\bar{b}; \bar{c}), \\ \text{for all atomic crossing formulae } \varphi(\bar{x}; \bar{y}) \text{ and} \\ \text{all parameters } \bar{c} \text{ in } U. \end{aligned}$$

(c) The *crossing rank* of a partition $P \sqcup Q$ of A is the cardinal

$$\text{crk}(P/\mathfrak{A}) := |P^{r-1}/\simeq_Q|.$$

When the structure \mathfrak{A} is clear from the context, we will write $\text{crk}(P)$ instead of $\text{crk}(P/\mathfrak{A})$. \lrcorner

Examples. (a) Let $\mathfrak{A} = \langle A, \leq \rangle$ be a linear order and $I \subseteq A$ an interval. Let P be the set of elements below I and Q the set of those above. For $a, b \in A$, it follows that

$$\begin{aligned} a \simeq_{P \sqcup Q} b \quad & \text{iff} \quad a = b \notin I \quad \text{or} \quad a, b \in I, \\ a \simeq_I b \quad & \text{iff} \quad a = b \in I, \quad a, b \in P, \quad \text{or} \quad a, b \in Q. \end{aligned}$$

Consequently,

$$\text{crk}(I/\mathfrak{A}) = 1 \quad \text{and} \quad \text{crk}(P \sqcup Q/\mathfrak{A}) = 2.$$

(b) Let $\mathfrak{T} = \langle T, \leq \rangle$ be a tree and $S \subseteq T$ a subtree. Then all elements of S are $\simeq_{T \setminus S}$ -equivalent, while elements of $T \setminus S$ are \simeq_S -equivalent if either both belong to the path from the root of T to the root of S , or both do not belong to this path. Hence,

$$\text{crk}(S/\mathfrak{T}) = 1 \quad \text{and} \quad \text{crk}(T \setminus S/\mathfrak{T}) = 2. \quad \lrcorner$$

Exercise 1.1. Let \bar{a} and \bar{b} be α -tuples, for some index set α . Prove that $\bar{a} \simeq_U \bar{b}$ if, and only if, $\bar{a}|_H \simeq_U \bar{b}|_H$, for every set $H \subseteq \alpha$ of less than r indices. \lrcorner

Let us collect a few basic relations that come in handy when computing crossing ranks.

Lemma 1.2. Let Σ be a finite relational signature and \mathfrak{A} a Σ -structure.

- (a) $\text{crk}(U/(\mathfrak{A}, P)) = \text{crk}(U/\mathfrak{A})$, for all unary predicates $P \subseteq A$.
- (b) $\text{crk}(U \cap C/\mathfrak{A}|_C) \leq \text{crk}(U/\mathfrak{A})$, for all $U, C \subseteq A$.
- (c) $\text{crk}(P \sqcup Q/\mathfrak{A}) \leq 2^{r-1} \cdot \text{crk}(P/\mathfrak{A}) \cdot \text{crk}(Q/\mathfrak{A})$, for all $P, Q \subseteq A$.
- (d) There exists a constant $m < \omega$ (only depending on Σ) such that

$$\text{crk}(P) \leq m^{\text{crk}(Q)}, \quad \text{for every partition } P \sqcup Q \text{ of } A.$$

- (e) There exists a constant $d < \omega$ (only depending on Σ) such that

$$\text{crk}(P \setminus Q) \leq \text{crk}(P)^{d \cdot \text{crk}(Q)}, \quad \text{for all sets } Q \subseteq P \subseteq A.$$

Proof. (a) As no atomic crossing formula can contain unary predicates, it follows that the relations \simeq_U coincide when computed in, respectively, \mathfrak{A} and $\langle \mathfrak{A}, P \rangle$.

- (b) For $\bar{a}, \bar{b} \in C^n$,

$$\bar{a} \simeq_U \bar{b} \quad \text{in the structure } \mathfrak{A} \quad \text{implies} \quad \bar{a} \simeq_{U \cap C} \bar{b} \quad \text{in } \mathfrak{A}|_C.$$

Hence, every \simeq_U -class of \mathfrak{A} is included in some $\simeq_{U \cap C}$ -class of $\mathfrak{A}|_C$. This means there are at least as many of the former as there are of the latter.

(c) Set $U := A \setminus (P \sqcup Q)$. The \simeq_U -class of a tuple $\bar{c} \in (P \sqcup Q)^{r-1}$ is uniquely determined by

- ♦ the set of indices $I := \{i < r-1 \mid c_i \in P\}$,
- ♦ the $\simeq_{A \setminus P}$ -class of the subtuple $\bar{c} \cap P^{r-1}$,
- ♦ the $\simeq_{A \setminus Q}$ -class of the subtuple $\bar{c} \cap Q^{r-1}$.

The number of possibilities for such sets and subtuples is bounded by 2^{r-1} , $\text{crk}(P)$, and $\text{crk}(Q)$, respectively.

(d) Let $\varphi(\bar{x}; \bar{y})$ be an atomic crossing formula with $|\bar{x}| = m$ and $|\bar{y}| = n$. For $\bar{a} \in P^m$, we define the set

$$C_\varphi(\bar{a}) := \{ \bar{c} \in Q^n \mid \mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \}.$$

Then

$$\bar{a} \simeq_Q \bar{b} \quad \text{iff} \quad C_\varphi(\bar{a}) = C_\varphi(\bar{b}), \quad \text{for all } \varphi.$$

Furthermore, by definition of $C_\varphi(\bar{a})$ we have

$$\bar{c} \simeq_P \bar{d} \quad \text{implies} \quad \bar{c} \in C_\varphi(\bar{a}) \Leftrightarrow \bar{d} \in C_\varphi(\bar{a}), \quad \text{for } \bar{c}, \bar{d} \in Q^n.$$

Consequently, there are at most $2^{\text{crk}(Q)}$ possible sets of the form $C_\varphi(\bar{a})$. If k is the number of atomic crossing formulae $\varphi(\bar{x}; \bar{y})$ with $0 < |\bar{x}| < r$ and $0 < |\bar{y}| < r$, it follows that the relation \simeq_Q has at most $(2^{\text{crk}(Q)})^k = (2^k)^{\text{crk}(Q)}$ classes.

(e) Set $U := A \setminus P$ and let $\bar{q}_0, \dots, \bar{q}_{n-1}$ be an enumeration (of representatives) of $Q^k / \simeq_{A \setminus Q}$. For two tuples $\bar{p}, \bar{p}' \in (P \setminus Q)^{r-1}$, it follows that

$$\bar{p} \simeq_{U \cup Q} \bar{p}' \quad \text{iff} \quad \bar{p} \simeq_Q \bar{p}' \quad \text{and} \quad \bar{p} \bar{q}_i \simeq_U \bar{p}' \bar{q}_i, \quad \text{for all } i < n.$$

Consequently,

$$|(P \setminus Q)^{r-1} / \simeq_{U \cup Q}| \leq |P^{r-1} / \simeq_Q| \cdot |P^{2(r-1)} / \simeq_U|.$$

Furthermore, for $\bar{p}, \bar{p}' \in P^{2(r-1)}$, we have

$$\bar{p} \simeq_U \bar{p}' \quad \text{iff} \quad \bar{p}|_I \simeq_U \bar{p}'|_I, \quad \text{for every subset } I \subseteq [2(r-1)] \\ \text{of size } |I| = r-1,$$

which implies that

$$|P^{2(r-1)} / \simeq_U| \leq |P^{r-1} / \simeq_U|^{\binom{2(r-1)}{r-1}}.$$

As $n \leq \text{crk}(Q)$, it follows that

$$\begin{aligned} \text{crk}(P \setminus Q) &\leq \text{crk}(P) \cdot \left(\text{crk}(P)^{\binom{2(r-1)}{r-1}} \right)^{\text{crk}(Q)} \\ &= \text{crk}(P)^{1 + \binom{2(r-1)}{r-1} \cdot \text{crk}(Q)}. \end{aligned} \quad \square$$

Frequently, it is useful to name \simeq_U -classes, i.e., to assign *colours* to them. Formally, we do so by using additional relations for these colours. Hence, we consider structures over a signature $\Sigma \cup \Xi$ that is divided into two parts: the set Σ of *actual* symbols and the set Ξ of *auxiliary* symbols.

Definition 1.3. Let \mathcal{A} be a Σ -structure, $P \cup Q = A$ a partition, and let \mathfrak{P} be an Ξ -expansion of the substructure $\mathcal{A}|_P$ induced by P .

(a) We say that \mathfrak{P} *respects* the crossing equivalence over Q if, for every tuple \bar{a} in P there exists a quantifier-free formula $\psi(\bar{x})$ over the signature Ξ such that

$$\mathfrak{P} \models \psi(\bar{b}) \quad \text{iff} \quad \bar{b} \simeq_Q \bar{a}.$$

(b) By $\mathfrak{A}[P]$ we denote any Ξ -expansion of $\mathcal{A}|_P$ that respects the crossing equivalence over Q . J

We start by describing the relationship between colourings and the crossing rank.

Lemma 1.4. Let Σ be a finite relational signature.

(a) For every constant $k < \omega$, there exists a finite relational signature Ξ with the following property: if \mathcal{A} is a Σ -structure with a partition $P \cup Q = A$ such that $\text{crk}(P/\mathcal{A}) \leq k$, then $\mathcal{A}|_Q$ has a Ξ -expansion respecting $\simeq_{A \setminus P}$.

(b) For every finite relational signature Ξ , there exists a constant $k < \omega$ with the following property: if \mathcal{A} is a Σ -structure with a partition $P \cup Q = A$ such that $\mathcal{A}|_P$ has a Ξ -expansion respecting \simeq_Q , then $\text{crk}(P/\mathcal{A}) \leq k$.

Proof. (a) For each arity $0 < n < r$, we introduce k relations R_0^n, \dots, R_{k-1}^n of arity n . By assumption, we can fix injective functions $g_n : P^n \rightarrow [k]$ such that

$$\bar{a} \simeq_Q \bar{b} \quad \text{iff} \quad g_n(\bar{a}) = g_n(\bar{b}), \quad \text{for all } \bar{a}, \bar{b} \in P^n.$$

Then the expansion of $\mathcal{A}|_P$ by the relations

$$R_i^n := \{ \bar{a} \in P^n \mid g_n(\bar{a}) = i \}, \quad 0 < n < r, \quad i < k,$$

respects \simeq_Q .

(b) Let k be the number of quantifier-free formulae $\psi(\bar{x})$ over the signature Ξ with $0 < |\bar{x}| < r$. The fact that $\mathfrak{A}|_P$ has a Ξ -expansion respecting \simeq_Q implies that $\text{crk}(P/\mathfrak{A}) \leq k$. \square

The main reason why we are interested in crossing equivalences is the fact that these contain exactly the information needed to reassemble a structure \mathfrak{A} from its parts. The simplest case of this situation is contained in the following lemma. We will prove more general statements below.

Lemma 1.5. *Let \mathfrak{A} be a Σ -structure, $P \cup Q = A$ a partition, and let Ξ be a signature such that the Ξ -expansions $\mathfrak{A}[P]$ and $\mathfrak{A}[Q]$ do exist. Then there exists a quantifier-free interpretation τ such that*

$$\mathfrak{A} \cong \tau(\mathfrak{A}[P] \oplus \mathfrak{A}[Q]).$$

Proof. For a tuple $\bar{c} \in A^n$, we denote by \bar{c}^P and \bar{c}^Q the two sub tuples of all components contained in, respectively, P and Q . Let $R \in \Sigma$ be a symbol of arity n . Then we have

$$R^{\mathfrak{A}} = R^{\mathfrak{A}|_P} \cup R^{\mathfrak{A}|_Q} \cup (R^{\mathfrak{A}} \setminus (P^n \cup Q^n)),$$

where the relation $R' := R^{\mathfrak{A}} \setminus (P^n \cup Q^n)$ is closed under the crossing equivalence in the sense that

$$\bar{c}^P \simeq_Q \bar{d}^P \text{ and } \bar{c}^Q \simeq_P \bar{d}^Q \text{ implies } \bar{c} \in R' \Leftrightarrow \bar{d} \in R'.$$

By assumption, all \simeq_P -classes and \simeq_Q -classes are definable by quantifier-free formulae. Consequently, so is the relation R' . \square

We have defined the crossing equivalence in terms of a very simple class of formulae. One consequence of the preceding lemma is that essentially nothing changes if we use more powerful ones.

Proposition 1.6. *Let Σ be a finite relational signature and $k, m, n, M < \omega$ constants. There exists a number N such that*

$$\text{crk}(U) \leq M \text{ implies } |U^n / \approx_k^m| \leq N$$

for every Σ -structure \mathfrak{A} and every subset $U \subseteq A$, where \approx_k^m is the equivalence relation on U^n defined by

$$\bar{a} \approx_k^m \bar{b} \text{ iff } \bar{a}\bar{c} \equiv_{\text{CMSO}}^m \bar{b}\bar{c}, \text{ for all } \bar{c} \in (A \setminus U)^k.$$

Proof. Set $W := A \setminus U$ and fix a signature Ξ such that $\mathfrak{A}[U]$ and $\mathfrak{A}[W]$ exist. Since disjoint unions and quantifier-free interpretations preserve MSO_m -equivalence it follows that, for $\bar{a}, \bar{b} \in U^n$ and $\bar{c} \in (A \setminus U)^k$,

$$\begin{aligned} & \mathfrak{A}, \bar{a}\bar{c} \equiv_{\text{MSO}}^m \mathfrak{A}, \bar{b}\bar{c} \\ \text{iff } & \mathfrak{A}[U], \bar{a} \equiv_{\text{MSO}}^m \mathfrak{A}[U], \bar{b} \text{ and } \mathfrak{A}[W], \bar{c} \equiv_{\text{MSO}}^m \mathfrak{A}[W], \bar{c} \\ \text{iff } & \mathfrak{A}[U], \bar{a} \equiv_{\text{MSO}}^m \mathfrak{A}[U], \bar{b}. \end{aligned}$$

Hence, the index of \approx_k^m is bounded by the number of MSO_m -theories of n -tuples. \square

Corollary 1.7. *For every CMSO-interpretation τ , there exists a function f such that*

$$\text{crk}(U/\tau(\mathfrak{A})) \leq f(\text{crk}(U/\mathfrak{A})), \text{ for all structures } \mathfrak{A} \text{ and all } U \subseteq A.$$

Proof. Let m be the quantifier rank of τ . For every atomic crossing formula $\varphi(\bar{x}; \bar{y})$ it follows that

$$\begin{aligned} \bar{a} \approx_k^m \bar{b} \text{ implies } & \mathfrak{A} \models \varphi^r(\bar{a}; \bar{c}) \Leftrightarrow \mathfrak{A} \models \varphi^r(\bar{b}; \bar{c}), \\ & \text{for all } \bar{c} \in (A \setminus U)^k, \end{aligned}$$

where \approx_k^m is the relation from the preceding proposition. Consequently, $\approx_{r-1}^m \subseteq \simeq_{A \setminus U}$ (where the first relation is defined in the structure \mathfrak{A} and the second one in $\tau(\mathfrak{A})$). We have shown in Proposition 1.6 that the index of \approx_{r-1}^m is bounded by a function of $\text{crk}(U/\mathfrak{A})$. Hence, so is $\text{crk}(U/\tau(\mathfrak{A}))$. \square

We have seen in Lemma 1.5 that we can reconstruct a structure from a partition into two substructures by adding suitable colourings. Let us generalise this to partitions into more than two components. We start by

introducing a suitable operation on $(\Sigma + \Xi)$ -structures consisting of disjoint unions and quantifier-free interpretations. We denote by $\bigotimes_i^\tau \mathfrak{A}_i$ the structure obtained by

- (I) taking the disjoint union of the structures \mathfrak{A}_i ;
- (II) adding tuples to the relations in Σ if they cross several components and their parts have the right colours; and
- (III) updating all the colours.

The formal definition is as follows.

Definition 1.8. Let Σ and Ξ be finite relational signatures.

- (a) An *update specification* is a quantifier-free interpretation τ of the form $\langle \delta(x), (\varphi_R)_{R \in \Sigma + \Xi} \rangle$ where $\delta(x) = \text{true}$ and the formulae φ_R are over the signature $\Xi \cup \{\leq\}$ for some binary relation symbol \leq .
- (b) An update specification is *symmetric* if all formulae φ_R use the relation \leq only in statements of the form $x \leq y \wedge y \leq x$.
- (c) For an update specification τ and a family $(\mathfrak{A}_i)_{i \in I}$ of $(\Sigma + \Xi)$ -structures indexed by some linear order I , we denote by

$$\bigotimes_{i \in I}^\tau \mathfrak{A}_i$$

the following $(\Sigma + \Xi)$ -structure. Let \mathfrak{B} be the disjoint union of the structures \mathfrak{A}_i , expanded by the preorder

$$a \leq b \quad : \text{iff} \quad a \in A_i \text{ and } b \in A_k, \quad \text{for some } i \leq k.$$

We set $\bigotimes_{i \in I}^\tau \mathfrak{A}_i := \tau'(\mathfrak{B})$ where $\tau' = \langle \delta(x), (\varphi'_R)_{R \in \Sigma + \Xi} \rangle$ is the interpretation with

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi'_R(\bar{x}) &:= \begin{cases} R\bar{x} \vee \bigvee_{i,j} [x_i \not\leq x_j \wedge \varphi_R(\bar{x})] & \text{if } R \in \Sigma, \\ \varphi_R(\bar{x}) & \text{if } R \in \Xi. \end{cases} \end{aligned}$$

- (d) Let \mathfrak{A} be a Σ -structure and $(H_i)_{i \in I}$ a partition of A . For a set $P \subseteq I$ of indices and an arity $n < \omega$, we set

$$H^n[P] := \bigcup_{i \in P} H_i^n.$$

Lemma 1.9. For every finite signature Σ and every number $k < \omega$, there exists a finite signature Ξ with the following property: if \mathfrak{A} is a Σ -structure, $U \subseteq A$, and $(H_i)_{i \in I}$ a partition of $A \setminus U$ such that

$$\text{crk}(H[P]/U \cup H[Q]) \leq k, \quad \text{for every partition } P \cup Q \text{ of } I,$$

then there exist a symmetric update specification τ and Ξ -expansions $\mathfrak{A}[H_i]$, for $i \in I$, such that

$$\mathfrak{A}[A \setminus U] \cong \bigotimes_{i \in I}^\tau \mathfrak{A}[H_i].$$

Proof. Let r be the maximal arity of a relation in Σ . Given a set $J \subseteq I$, we say that a subset $P \subseteq J$ provides J -representatives if, for every $i \in J \setminus P$, every arity $0 < n < r$, and every tuple $\bar{c} \in H_i^n$, there is some $\bar{a} \in H^n[P]$ such that

$$\bar{c} \simeq_{U \cup H[J \setminus (P \cup \{i\})]} \bar{a}.$$

We start by proving that, for every set $J \subseteq I$, there exists a subset $P \subseteq J$ of size $|P| \leq 2k$ providing J -representatives. Fix a partition $K \cup L = J$ such that the number

$$|H^{r-1}[K]/\simeq_{A \setminus H[K]}| + |H^{r-1}[L]/\simeq_{A \setminus H[L]}|$$

is maximal. Let $P \subseteq K$ and $Q \subseteq L$ be minimal sets such that $H[P]$ and $H[Q]$ contain representatives of all $\simeq_{A \setminus H[K]}$ -classes and all $\simeq_{A \setminus H[L]}$ -classes, respectively. Since the number of such classes is bounded by k , it follows that $|P|, |Q| \leq k$. We claim that the union $P \cup Q$ has the desired properties. Hence, fix $i \in J \setminus (P \cup Q)$ and $\bar{c} \in H_i^n$. Without loss of generality, we may assume that $n = r - 1$. By symmetry, we may also assume that $i \in K$. For a contradiction, suppose that there is no $\bar{a} \in H^{r-1}[P]$ with

$$\bar{a} \simeq_{U \cup H[J \setminus (P \cup Q \cup \{i\})]} \bar{c}.$$

Fix representatives $\bar{p}_0, \dots, \bar{p}_{s-1}$ and $\bar{q}_0, \dots, \bar{q}_{t-1}$ of every $\simeq_{A \setminus H[K]}$ -class in $H^{r-1}[P]$ and of every $\simeq_{A \setminus H[L]}$ -class in $H^{r-1}[Q]$, respectively. By choice of P and Q , we have

$$|H^{r-1}[K] / \simeq_{A \setminus H[K]}| = s \quad \text{and} \quad |H^{r-1}[L] / \simeq_{A \setminus H[L]}| = t.$$

But by assumption, $\bar{c}, \bar{p}_0, \dots, \bar{p}_{s-1}$ belong to different $\simeq_{A \setminus H[K]}$ -classes, while $\bar{q}_0, \dots, \bar{q}_{t-1}$ belong to different $\simeq_{A \setminus H[L]}$ -classes. Consequently,

$$(s+1) + t = |H^{r-1}[K] / \simeq_{A \setminus H[K]}| + |H^{r-1}[L] / \simeq_{A \setminus H[L]}| = s + t.$$

A contradiction.

Having established the claim, we can now find τ and the expansions $\mathcal{Q}[H_i]$ as follows. First, we construct a sequence $P_0, \dots, P_{r-1} \subseteq I$ of r disjoint sets by the following inductive procedure: if we have already chosen P_0, \dots, P_{j-1} , we use the above claim to find a set $P_j \subseteq I \setminus (P_0 \cup \dots \cup P_{j-1})$ of size $|P_j| \leq 2k$ of $(I \setminus (P_0 \cup \dots \cup P_{j-1}))$ -representatives.

Let $J := I \setminus (P_0 \cup \dots \cup P_{r-1})$ be the complement of these sets. For every arity $n < r$, every index $j < r$, and every $\simeq_{A \setminus H[P_j]}$ -class σ of n -tuples, we fix a representative $\bar{a}_j^\sigma \in P_j^n$ and we set

$$R_{j,i}^\sigma := \begin{cases} \left\{ \bar{c} \in H_i^n \mid \bar{c} \simeq_{A \setminus H[P_0 \cup \dots \cup P_j \cup \{i\}]} \bar{a}_j^\sigma \right\}, & \text{for } i \in J, \\ \left\{ \bar{c} \in H_i^n \mid \bar{c} \simeq_{A \setminus H[P_j]} \bar{a}_j^\sigma \right\}, & \text{for } i \in P_j, \end{cases}$$

$$Q_j := H[P_j].$$

Let

$$\Xi := \{Q_0, \dots, Q_{r-1}\} \cup \{R_j^\sigma \mid j < r, \sigma \text{ a } \simeq_{A \setminus H[P_j]} \text{-class}\}.$$

We claim that the Ξ -expansion of $\mathcal{Q}|_{H_i}$, $i \in I$, given by these relations has the desired properties.

For the proof it is sufficient to show that we can recover the \simeq_U -class of a tuple $\bar{c} \in (A \cup U)^n$ from the information contained in this expansion. We distinguish two cases. If $\bar{c} \in H_i^n$, for some $i \in I$, we already know its \simeq_U -class.

Otherwise, let $i_0, \dots, i_{s-1} \in I$ be the indices i such that $\bar{c} \cap H_i \neq \emptyset$ and set $\bar{c}_j := \bar{c} \cap H_{i_j}$. We choose the enumeration i_0, \dots, i_{s-1} such that

$$i_j \in P_m \quad \text{implies} \quad i_{j+1} \in P_m \cup \dots \cup P_{r-1}, \quad \text{for } j < s-1 \text{ and } m < r,$$

and we construct a monotone function $\mu : [s] \rightarrow [r]$ and tuples $\bar{a}_j \subseteq H[P_{\mu(j)}]$ satisfying

$$\bar{a}_j \simeq_{A \setminus H[P_0 \cup \dots \cup P_{\mu(j)} \cup \{i_j\}]} \bar{c}_j, \quad \text{for all } j < s,$$

as follows. We proceed by induction on j starting with $j = s-1$. If $\bar{c}_j \subseteq H[P_m]$, we set $\mu(j) := m$ and $\bar{a}_j := \bar{c}_j$. (By choice of i_j , this implies that $\mu(j) \leq \mu(j+1)$.) Otherwise, we set $\mu(j) := \mu(j+1) - 1$ and we use the fact that $P_{\mu(j)}$ is a set of $(I \setminus (P_0 \cup \dots \cup P_{\mu(j)-1}))$ -representatives to find a tuple $\bar{a}_j \subseteq H[P_{\mu(j)}]$ with

$$\bar{a}_j \simeq_{A \setminus H[P_0 \cup \dots \cup P_{\mu(j)} \cup \{i_j\}]} \bar{c}_j.$$

It then follows that

$$\bar{c}_j \simeq_{U \cup \bar{c}_0 \dots \bar{c}_{j-1} \bar{a}_{j+1} \dots \bar{a}_{r-1}} \bar{a}_j, \quad \text{for all } j < s.$$

This implies that

$$\begin{aligned} \bar{c}_0 \dots \bar{c}_{s-2} \bar{c}_{s-1} &\simeq_U \bar{c}_0 \dots \bar{c}_{s-2} \bar{a}_{s-1} \\ &\simeq_U \dots \\ &\simeq_U \bar{c}_0 \bar{a}_1 \dots \bar{a}_{s-1} \simeq_U \bar{a}_0 \bar{a}_1 \dots \bar{a}_{s-1}. \end{aligned}$$

Similarly, it follows that

$$\mathcal{A} \models \varphi(\bar{c}_0; \dots; \bar{c}_{s-1}) \quad \text{iff} \quad \mathcal{A} \models \varphi(\bar{a}_0; \dots; \bar{a}_{s-1}),$$

for every atomic formula $\varphi(\bar{x}_0; \dots; \bar{x}_{s-1})$ that contains variables from at least two different tuples \bar{x}_i and \bar{x}_j .

Hence, the \simeq_U -class of \bar{c} can be determined from the the tuples \bar{a}_j (which can be computed using the relations R_{j,i_j}^σ), while its atomic type can be determined from $\bar{a}_0, \dots, \bar{a}_{s-1}$ and from the atomic types of the component tuples $\bar{c}_0, \dots, \bar{c}_{s-1}$. \square

2 Decompositions

We can use the crossing rank to define decompositions for relational structures. It turns out that these are better behaved if we allow as index sets order-trees of arbitrary *ordinal height*. Let us quickly define what we mean by that.

Definition 2.1. Let $\mathfrak{T} = \langle T, \leq \rangle$ be a partial order.

(a) \mathfrak{T} is an *order-tree of ordinal height* if there exists a set D and an ordinal α such that \mathfrak{T} is isomorphic to a partial order of the form $\langle S, \leq \rangle$ where $S \subseteq D^{<\alpha}$ is prefix-closed and \leq is the prefix ordering on $D^{<\alpha}$.

(b) The *height* of \mathfrak{T} is the least such ordinal α .

(c) The maximal elements of T are called the *leaves* of \mathfrak{T} .

The corresponding decompositions now take the following form.

Definition 2.2. Let \mathfrak{A} be a Σ -structure.

(a) A *partition decomposition* of \mathfrak{A} is a family $(U_t)_{t \in T}$ of non-empty subsets $U_t \subseteq A$ indexed by an order-tree T of ordinal height satisfying the following conditions.

- $s \leq t$ implies $U_s \supseteq U_t$,
- $U_s \cap U_t = \emptyset$, if s and t are incomparable,
- for every $a \in A$, there is some $t \in T$ with $U_t = \{a\}$.

(b) The *crossing-width* of a partition decomposition $(U_t)_{t \in T}$ is the number

$$\sup \left\{ \text{crk}(\cup_{i \in P} U_{s_i}) \mid t \in T, P \text{ a subset of the successors of } t \right\}.$$

The *crossing-width* $\text{cwd } \mathfrak{A}$ of \mathfrak{A} is the minimal crossing width of a partition decomposition of \mathfrak{A} .

Remark. For a graph, the crossing-width is not quite the same as the clique-width, but it is 'equivalent' to it in the sense that these two widths differ by a factor of at most 2.

Exercise 2.1. Let $(U_t)_{t \in T}$ be a partition decomposition of some structure \mathfrak{A} . Prove the following properties.

(a) $U_{\langle \rangle} = A$, for the root $\langle \rangle$ of T .

(b) If S is the set of successors of $t \in T$, then $(U_s)_{s \in S}$ forms a partition of U_t .

Examples. (a) Let $\mathfrak{G} = \langle V, E \rangle$ be a complete graph. We define a partition decomposition $(U_t)_{t \in T}$ as follows. For T we take the tree consisting of a root $*$ to which we attach one leaf t_v for every vertex v of \mathfrak{G} . The components are $U_* := V$ and $U_{t_v} = \{v\}$. Since

$$\text{crk}(P) = 1, \quad \text{for every } P \subseteq V,$$

it follows that $\text{crk}(\mathfrak{G}) = 1$.

(b) Let $\mathfrak{A} = \langle A, B \rangle$ where B is a linear betweenness relation, i.e., a relation of the form

$$B = \{ \langle a, b, c \rangle \mid a \leq b \leq c \}, \quad \text{for some linear order } \leq \text{ on } A.$$

We define a partition decomposition $(U_t)_{t \in T}$ as follows. We start with the root component $U_{\langle \rangle} = A$. Inductively, suppose we have already defined U_t and that U_t forms an interval of \mathfrak{A} (with respect to the underlying order \leq). If $|U_t| = 1$, we stop. Otherwise, we partition U_t into two intervals $U_u + U_v$ and we make u and v successors of t .

It follows that $\text{crk}(\mathfrak{A}) \leq 3$ since, for an interval $I \subseteq A$, there are exactly 3 \simeq_U -classes of pairs $\langle a, b \rangle$ in I^2 : those with $a < b$, those with $a = b$, and those with $a > b$.

(c) Let $\mathfrak{T} = \langle T, \leq \rangle$ be an order-tree. We consider the following partition decomposition $(U_s)_{s \in S}$. Let S be the tree obtained from \mathfrak{T} by adding a new leaf to every vertex. For $s \in S$, we set If $s \in S$ is an original vertex of \mathfrak{T} , we set

$$U_s := \begin{cases} \{ v \in T \mid v \geq s \} & \text{if } s \text{ is an original vertex of } \mathfrak{T}, \\ \{ v \} & \text{if } s \text{ is the leaf added to } v \in T. \end{cases}$$

Let s_0, \dots, s_{n-1} be successors of the same vertex $t \in S$. If the new leaf added to t is not among the s_0, \dots, s_{n-1} , or if s_0, \dots, s_{n-1} consists of all successors of t , we have

$$\text{crk}(U_{s_0} \cup \dots \cup U_{s_{n-1}}) = 1.$$

Otherwise,

$$\text{crk}(U_0 \cup \dots \cup U_{s_{n-1}}) = 2.$$

Hence, this decomposition has width 2.

We can improve this decomposition to reduce the width to 1 as follows. Let S be the tree obtained from \mathfrak{T} by (I) inserting an intermediate vertex in every edge; (II) adding a new root; and (III) adding a new leaf to every of the vertices introduced in (I) and (II). The components are as follows.

$$W_s := \begin{cases} \{v \in T \mid v > s\} & \text{if } s \text{ is an original vertex of } \mathfrak{T}, \\ \{v \in T \mid v \geq s\} & \text{if } s \text{ is the predecessor of } v \in T \text{ introduced} \\ & \text{in (I) or (II),} \\ \{v\} & \text{if } s \text{ is the leaf added to the predecessor} \\ & \text{of } v \in T. \end{cases}$$

(d) Let $\mathfrak{G} = ([m] \times [n], E)$ be the $m \times n$ grid with edge relation

$$E = \{ \langle \langle i, k \rangle, \langle j, l \rangle \rangle \mid |i - j| + |k - l| = 1 \}.$$

We construct a partition decomposition of \mathfrak{G} by induction on m . If $m = 1$, \mathfrak{G} is a path of length n and we can use a decomposition splitting \mathfrak{G} into a single vertex and a path of length $n - 1$. If $m > 1$, we split \mathfrak{G} into a path of length n and an $(m - 1) \times n$ grid.

Let $(U_t)_{t \in T}$ be the resulting decomposition. Then

$$\text{crk}(U_t) \leq n + 1$$

(there are up to n vertices with neighbours outside of U_t and one $\simeq_{V \setminus U_t}$ -class for the other vertices in U_t). As the tree T is binary, this implies that $\text{crk}(\mathfrak{G}) \leq n + 1$. (A (rather bad) lower bound can be obtained via Corollary 4.4 or via Proposition 7.2 below; see also the rest of Section 7.)

An important example of structures with small crossing-widths are linear orders and trees. Since we will use decompositions of linear orders in some

of the constructions below, let us state their properties formally. The second part of this lemma is one of the reasons why we work with trees of height larger than ω .

Lemma 2.3. Let $\mathfrak{A} = \langle A, \leq \rangle$ be a linear order.

- (a) \mathfrak{A} has a partition decomposition $(U_t)_{t \in T}$ of width 1 where the index tree T is binary and has height at most $|A|$.
- (b) If \mathfrak{A} is infinite, it does not have a partition decomposition $(U_t)_{t \in T}$ of finite width where the index tree T has height strictly less than $|A|$.

Proof. (a) Set $\kappa := |A|$ and fix an enumeration $(a_i)_{i < \kappa}$ of A . Let us call a family $(U_t)_{t \in T}$ be a *partial* decomposition if it can be obtained from a partition decomposition by removing some subtrees from T . By induction on $\alpha < \kappa$, we will define a sequence $(U_t^\alpha)_{t \in T_\alpha}$ of partial decompositions of \mathfrak{A} such that

- ♦ the sequence is increasing in the sense that

$$T_\alpha \subseteq T_\beta \quad \text{and} \quad U_t^\alpha = U_t^\beta, \quad \text{for all } \alpha < \beta \text{ and } t \in T_\alpha,$$

- ♦ T_α is a binary tree of height strictly less than $3(\alpha + 1)$,
- ♦ every component U_t^α is an interval of \mathfrak{A} ,
- ♦ for every $i < \alpha$, there is some leaf $t \in T_\alpha$ with $U_t^\alpha = \{a_i\}$.

Then the limit $(U_t)_t$ of this sequence will be a partition decomposition of \mathfrak{A} . Furthermore, component U_t being an interval, we have $\text{crk}(U_t) = 1$, for all t .

We start with a tree T_0 consisting only of the root $\langle \rangle$ and with the component $U_{\langle \rangle} := A$. For the limit step, let δ be a limit ordinal and suppose that we have defined $(U_t^\alpha)_{t \in T_\alpha}$ for all $\alpha < \delta$. We choose for $T_\delta := \bigcup_{\alpha < \delta} T_\alpha$ the limit of the trees and for the components the sets

$$U_t^\delta := U_t^\alpha, \quad \text{for some } \alpha \text{ with } t \in T_\alpha.$$

The height of T_δ is at most

$$\sup_{\alpha < \delta} 3(\alpha + 1) \leq 3\delta < \delta + 1 \leq 3(\delta + 1).$$

For the successor step, suppose that we have already defined $(U_t^\alpha)_{t \in T_\alpha}$. We distinguish several cases. If there is some leaf t with $U_t^\alpha = \{a_\alpha\}$, we simply set

$$T_{\alpha+1} := T_\alpha \quad \text{and} \quad U_t^{\alpha+1} := U_t^\alpha.$$

Suppose otherwise. If there is no leaf t with $a_\alpha \in U_t$, the set of vertices $t \in T_\alpha$ with $a_\alpha \in U_t$ forms an infinite branch of T_α . We add a new vertex s to T_α as the limit of this branch, and we set

$$U_s := \bigcap_{t < s} U_t.$$

Hence, we may assume that T_α has a leaf t (necessarily unique) with $a_\alpha \in U_t$. Let I be the subinterval of U_t consisting of all elements smaller than a_α , and let J be the subinterval of the larger elements. Thus, $U_t = I + \{a_\alpha\} + J$. Let $T_{\alpha+1}$ be the tree obtained from T_α by adding two successors s_o and s_i to t and two successors v_o and v_i to s_i . The new components are

$$U_{s_o}^{\alpha+1} := I, \quad U_{s_i}^{\alpha+1} := \{a_i\} \cup J, \quad U_{v_o}^{\alpha+1} := \{a_i\}, \quad U_{v_i}^{\alpha+1} := J.$$

The height of $T_{\alpha+1}$ is strictly less than

$$3(\alpha + 1) + 1 + 2 = 3(\alpha + 2).$$

(The 1 is in case we had to add the limit of an infinite branch.)

(b) Let $(U_t)_{t \in T}$ be a partition decomposition of \mathfrak{A} where the height of T is strictly less than $\kappa := |A|$. Then there must be some vertex $t \in T$ with κ many successors. Let S be the set of successors of t . For each $s \in S$, we fix some element $a_s \in U_s$. Choose an infinite set $I \subseteq S$ such that, for all $u, v \in I$ with $a_u < a_v$, there is some $w \in S \setminus I$ with $a_u < a_w < a_v$. Then it follows that

$$\text{crk}(\bigcup_{s \in I} U_s) \geq |I| \geq \aleph_0. \quad \square$$

Let us collect a few straightforward observations about the way some simple operations affect (or do not affect) the crossing-width.

Lemma 2.4. Let \mathfrak{A} and \mathfrak{B} be Σ -structures.

- (a) $\mathfrak{A} \subseteq \mathfrak{B}$ implies $\text{cld } \mathfrak{A} \leq \text{cld } \mathfrak{B}$.
- (b) $\text{cld } (\mathfrak{A}, P) = \text{cld } \mathfrak{A}$, for all unary predicates $P \subseteq A$.
- (c) $\text{cld}(\mathfrak{A} \oplus \mathfrak{B}) = \max\{\text{cld } \mathfrak{A}, \text{cld } \mathfrak{B}\}$.

Exercise 2.2. Prove this lemma. J

Exercise 2.3. Prove that every structure \mathfrak{A} has a partition decomposition $(U_t)_{t \in T}$ of width $\text{cld } \mathfrak{A}$ where the index tree T is binary. J

3 Terms

The original definition of clique-width was given in terms of certain graph operations. A similar characterisation exists for the crossing-width using the operations \oplus^τ from above.

Definition 3.1. Let Σ and \mathcal{E} be finite relational signatures. A \oplus -term over the signatures Σ and \mathcal{E} is a tree t (possibly of ordinal height) whose internal vertices are labelled by operations of the form \oplus^τ with a symmetric update specification τ and whose leaves are labelled by $(\Sigma + \mathcal{E})$ -structures with a single element. J

The problem we face when trying to define the value of such a \oplus -term is that they are not well-founded: they might have branches of arbitrary ordinal length. For finite \oplus -terms t , we could simply say

$$\text{val}(t) := \begin{cases} \mathfrak{A} & \text{if } t \text{ is a constant with value } \mathfrak{A}, \\ \bigoplus_{i \in I}^\tau \text{val}(s_i) & \text{if } t = \bigoplus_{i \in I}^\tau s_i. \end{cases}$$

If we want to evaluate the \oplus^τ operation at some vertex v inside a tree with infinite branches, we need to know for every tuple \vec{c} of elements to which relations from \mathcal{E} they currently belong. We will compute this information along the paths from the leaves corresponding to \vec{c} to the vertex v .

Definition 3.2. Let t be a \oplus -term over the signatures Σ and Ξ , w a vertex of t , and $n < \omega$ a number.

(a) An n -ary colour trace to w is a sequence $(\theta_v)_{v \leq w}$ of atomic types over the signature Ξ with n variables x_0, \dots, x_{n-1} .

(b) Such a colour trace $(\theta_v)_{v \leq w}$ is *locally consistent* if, for each vertex $v < w$ with predecessor u labelled $t(u) = \oplus^\tau$, the type θ_u is obtained from θ_v as described by the update specification τ . \lrcorner

We would like to define the relations from Ξ in $\text{val}(t)$ by choosing a locally consistent colour trace for every tuple \bar{c} and then reading off the relations from the first type $\theta_{\langle \rangle}$ in each trace. Unfortunately, there might be several different consistent colour traces to a vertex w since some branches of t might contain limit steps, i.e., vertices v without an immediate predecessor. To have unique colour traces, we have to specify what happens at such steps. Note that the precise type θ_v we choose for such a limit vertex v does not really matter, since we are mostly only concerned by the partition induced by the types, not by which type is assigned to which class of the partition. This allows us to use a very simple definition: we just label such a vertex with the infimum of all the types assigned to smaller vertices.

Definition 3.3. Let $n < \omega$ and fix some linear order on the set of all atomic types of arity n .

(a) We lift the order on types to an order on the set of all colour traces of arity n by setting

$$(\sigma_v)_{v \leq w} \leq (\theta_v)_{v \leq w} \quad : \text{iff} \quad \sigma_v \leq \theta_v, \quad \text{for all } v \leq w.$$

(b) A colour trace $(\theta_v)_{v \leq w}$ of arity n is *globally consistent* if, for every vertex $v \leq w$ without immediate predecessor, the set $\{u < v \mid \theta_u = \theta_v\}$ is unbounded below v .

(c) A colour trace is *consistent* if it is locally and globally consistent.

(d) A \oplus -term t is *order-preserving* if, for every vertex w and every arity $n < \omega$, the set of consistent n -ary colour traces to w is linearly ordered by the above order. \lrcorner

Remark. Note that local consistency propagates information from larger vertices to smaller ones, while global consistency propagates it in the opposite direction. \lrcorner

We start by proving that, for order-preserving terms, colour traces are unique.

Lemma 3.4. Let t be an order-preserving \oplus -term. For every vertex w and every atomic type σ of arity $n < \omega$, there exists at most one consistent colour trace $(\theta_v)_{v \leq w}$ to w ending in $\theta_w = \sigma$.

Proof. Let $(\theta_v)_{v \leq w}$ and $(\theta'_v)_{v \leq w}$ be two consistent colour traces to w with $\theta_w = \theta'_w$. By induction on w starting at the root, we prove that $\theta_v = \theta'_v$ for all v . If $w = \langle \rangle$ is the root, there is only one vertex $v \leq w$ and the claim holds by assumption. Next, suppose that w has a predecessor v for which we have already established the claim. By local consistency, $\theta_w = \theta'_w$ implies $\theta_v = \theta'_v$, and the claim follows by inductive hypothesis.

Finally suppose that w is a vertex without predecessors and that we have already prove the claim for all $v < w$. By symmetry and the fact that t is well-formed, we may then assume that $(\theta_v)_{v \leq w} \leq (\theta'_v)_{v \leq w}$. By global consistency of $(\theta'_v)_{v \leq w}$, the set

$$S := \{v < w \mid \theta'_v = \theta'_w\}$$

is unbounded below w . Since there are only finitely many atomic types σ , there exists a minimal one such that the subset

$$S_o := \{v \in S \mid \theta_v = \sigma\}$$

is also unbounded below w . By global consistency of $(\theta_v)_{v \leq w}$, it therefore follows that

$$\theta_w \leq \sigma = \theta_v \leq \theta'_v = \theta'_w = \theta_w, \quad \text{for all } v \in S_o.$$

Consequently, $\theta_v = \theta'_v$, for all $v \in S_o$. As S_o is unbounded it therefore follows by inductive hypothesis that $\theta_v = \theta'_v$, for all $v < w$. \square

Unfortunately, even for order-preserving terms, consistent colour traces do not always exist. To make sure we have all the colour traces we need to evaluate a given term, we require our terms to be well-behaved in the following sense.

Definition 3.5. Let t be a \oplus -term.

(a) Let \bar{c} be an n -tuple of leaves of t . We say that \bar{c} *splits* at the vertices w, s_0, \dots, s_{m-1} if $w := c_0 \sqcap \dots \sqcap c_{n-1}$ is their common prefix and s_0, \dots, s_{m-1} is an enumeration of all successors s of w with $s \leq c_i$, for some i . In this case the *split* of \bar{c} is the sequence $\bar{c}^0, \dots, \bar{c}^{m-1}$ of subtuples $\bar{c}^i := \bar{c} \cap \uparrow s_i$.

(b) Let \bar{c} be an n -tuple of leaves of t . Under certain circumstances, we can associate with \bar{c} a consistent colour trace $\chi(\bar{c})$ by induction on the number of distinct components of \bar{c} . If $c_0 = \dots = c_{n-1}$, let $\chi(\bar{c})$ be the unique colour trace to c_0 starting in the atomic type of $\langle c_0, \dots, c_{n-1} \rangle$ (if it exists). For the inductive step, suppose that \bar{c} splits at w, s_0, \dots, s_{m-1} , let $\bar{c}^0, \dots, \bar{c}^{m-1}$ be its split, and suppose that $t(w) = \oplus^\tau$. If the colour traces $\chi(\bar{c}^0), \dots, \chi(\bar{c}^{m-1})$ exist, say, $\chi(\bar{c}^i) = (\theta_v^i)_v$, we denote by $\chi(\bar{c})$ the unique colour trace $(\sigma_v)_{v \leq w}$ to w (if it exists) where the type θ_w is obtained from the types $\theta_{s_0}^0, \dots, \theta_{s_{m-1}}^{m-1}$ as described by the update specification τ .

(c) We call t *well-formed* if it is order-preserving and every finite tuple \bar{c} of leaves has an associated trace $\chi(\bar{c})$.

(d) If t is well-formed, we define the *value* $\text{val}(t)$ of t as the following $(\Sigma + \Xi)$ -structure \mathfrak{A} . The universe A is the set of leaves of t . To define the relations of \mathfrak{A} , let $\bar{a} \in A^n$ be a tuple and $\chi(\bar{a}) = (\theta_v)_v$ the associated colour trace. We add \bar{a} to a relation $R \in \Xi$ if $R\bar{x} \in \theta_{\langle \rangle}$, and we add it to a relation $R \in \Sigma$ if it satisfies the following condition. Suppose that \bar{a} splits at w, s_0, \dots, s_{m-1} , let $\bar{a}^0, \dots, \bar{a}^{m-1}$ be its split, and let $\chi(\bar{a}^i) = (\sigma_u^i)_u$. Then we add \bar{a} to R if the operation $t(w) = \oplus^\tau$ at w adds all tuples with types $\sigma_{s_0}^0, \dots, \sigma_{s_{m-1}}^{m-1}$ to R . \downarrow

We are finally able to prove that partition decompositions and \oplus -terms are basically the same thing, only the associated notions of width are slightly different.

Proposition 3.6. Let Σ be a finite relational signature.

(a) For every finite relational signature Ξ , there exists a number $k < \omega$ such that, given a \oplus -term s over the signatures Σ and Ξ , the structure $\text{val}(s)$ has a partition decomposition $(U_t)_{t \in T}$ of width k whose index tree T coincides with the underlying tree of s .

(b) For every $k < \omega$, there exists a finite relational signature Ξ such that, given a Σ -structure \mathfrak{A} with a partition decomposition $(U_t)_{t \in T}$ of width k , there exists a well-formed \oplus -term s over the signatures Σ and Ξ such that $\text{val}(s) = \mathfrak{A}$ and the underlying tree of s coincides with T .

Proof. (a) Let T be the underlying tree of s . For $t \in T$, let U_t be the set of all (elements of A corresponding to) leaves v of s with $t \leq v$. Then $(U_t)_{t \in T}$ is a partition decomposition and it follows by Lemma 1.4 (b) that its width is bounded by some constant depending only on Ξ .

(b) Let Ξ_0 be the signature from Lemma 1.9. We define the desired term s in two phases. First we choose, for every vertex $t \in T$, a Ξ_0 -expansion $\mathfrak{A}[U_t]$ such that, if $(s_i)_{i \in I}$ are the successors of t , there is some update specification τ with

$$\mathfrak{A}[U_t] = \bigoplus_{i \in I}^{\tau} \mathfrak{A}[U_{s_i}].$$

Then we use these update specifications to construct the desired term s , but not directly: to ensure global consistency, we have to use a larger set $\Xi \supseteq \Xi_0$ of colours.

For the first phase, we start at the root of T where we use the Ξ_0 -expansion of \mathfrak{A} where all auxiliary relations in Ξ_0 are empty. Inductively suppose that we have already chosen a Ξ_0 -expansion $\mathfrak{A}[U_t]$ for $t \in T$. By Lemma 1.9, there exists a symmetric update specification τ and Ξ_0 -expansions $\mathfrak{A}[U_{s_i}]$ of the successors s_i of t such that

$$\mathfrak{A}[U_t] = \bigoplus_i^{\tau} \mathfrak{A}[U_{s_i}].$$

Finally for the limit step, let $t \in T$ be a vertex without predecessor and suppose that we have defined $\mathfrak{A}[U_s]$ already for all $s < t$. We choose for $\mathfrak{A}[U_t]$ an arbitrary Ξ_0 -expansion of $\mathfrak{A}|_{U_t}$ respecting $\simeq_{A \setminus U_t}$.

This completes the first phase. For every tuple \bar{c} in A and every vertex $t \in T$ with $\bar{c} \subseteq U_t$, let θ_t be the atomic type of \bar{c} in (the \mathcal{E}_o -reduct of) $\mathfrak{A}[U_t]$. We would like to define the term s such that $\chi(\bar{c}) := (\theta_t)_t$ is the colour trace associated with \bar{c} . The problem is that, as defined above, these traces are locally consistent but not necessarily globally consistent. Also they do not need to be linearly ordered. To fix this we use the larger signature $\mathcal{E} := \mathcal{E}_o + \mathcal{E}_1$, where \mathcal{E}_1 is a disjoint copy of \mathcal{E}_o . Note that every atomic \mathcal{E} -type θ is of the form $\theta = \theta_o + \theta_1$ where θ_i is an atomic \mathcal{E}_i -type. (By \mathcal{E} -type we mean a type over the signature \mathcal{E} .) We define the order on such types by

$$\theta_o + \theta_1 \leq \theta'_o + \theta'_1 \quad \text{iff} \quad \theta_o \leq \theta'_o, \quad \text{or} \quad \theta_o = \theta'_o \text{ and } \theta_1 \leq \theta'_1.$$

For each vertex $t \in T$, we will choose some bijection μ_t between atomic \mathcal{E}_o -types and \mathcal{E}_1 -types and then add the new relations from \mathcal{E}_1 to $\mathfrak{A}[U_t]$ according to μ_t . To avoid confusion in the following, we will denote the \mathcal{E}_o -expansion of $\mathfrak{A}[U_t]$ by $\mathfrak{A}^o[U_t]$ and its \mathcal{E} -expansion by $\mathfrak{A}^1[U_t]$. Using the expansions $\mathfrak{A}^1[U_t]$ we can then read off a term s . We will choose the functions μ_t such that

- ♦ the colour traces from s are globally consistent and
- ♦ all consistent colour traces to the same vertex and with the same arity a linearly ordered.

By the definition of $\text{val}(s)$, it then follows that s is well-formed and that its value is $\text{val}(s) = \mathfrak{A}$ (or, rather the \mathcal{E} -expansion $\mathfrak{A}^1[A]$).

We proceed again by induction starting at the root $\langle \rangle$ of T , where we define $\mathfrak{A}^1[U_{\langle \rangle}]$ such that all relations in \mathcal{E}_1 are empty. For the successor step, suppose that we have already chosen the \mathcal{E} -expansion $\mathfrak{A}^1[U_v]$ and let w be a successor of v . Note that specifying a bijection μ_w is the same thing as specifying a second linear ordering \sqsubseteq on the atomic \mathcal{E}_o -types. (Then we can map the n -th \mathcal{E}_o -type with respect to \sqsubseteq to the n -th \mathcal{E}_1 -type with respect to the given ordering \leq .) For each pair θ, θ' of atomic \mathcal{E}_o -types, set $(\theta_u)_{u \leq w} := \chi(\bar{c})$ and $(\theta'_u)_{u \leq w} := \chi(\bar{c}')$ where \bar{c} and \bar{c}' are any tuples in

$\mathfrak{A}^o[U_w]$ of types θ and θ' , respectively. We set

$$\theta \sqsubseteq \theta' \quad \text{iff} \quad \text{either there is some } u < w \text{ with } \mu_u(\theta_u) < \mu_u(\theta'_u), \\ \text{or } \theta_w \leq \theta'_w \text{ and } \theta_u = \theta'_u, \quad \text{for all } u < w.$$

By inductive hypothesis, this ordering is linear. We choose for μ_w the bijection induced by this ordering.

Having done this construction for all successors of w of v , we also have to find an update specification τ such that

$$\mathfrak{A}^1[U_v] = \bigoplus_w^{\tau} \mathfrak{A}^1[U_w].$$

Then we can set $s(v) := \bigoplus^{\tau}$. From the first phase we already have a specification τ_o satisfying

$$\mathfrak{A}^o[U_v] = \bigoplus_w^{\tau_o} \mathfrak{A}^o[U_w].$$

Since the atomic \mathcal{E}_o -type of every tuple in $\mathfrak{A}^1[U_v]$ determines its atomic \mathcal{E} -type, we can simply extend τ_o by formulae defining the new relations in \mathcal{E}_1 .

It remains to consider vertices v without predecessor. By inductive hypothesis, we may assume that we have already chosen \mathcal{E} -expansions for every $u < v$. Given an atomic \mathcal{E}_o -type θ , set $(\theta_u)_{u \leq v} := \chi(\bar{c})$ with \bar{c} realising θ , and let σ be the minimal atomic \mathcal{E} -type such that the set

$$\{ u < v \mid \theta_u + \mu_u(\theta_u) = \sigma \}$$

is unbounded below v . We define $\mathfrak{A}^1[U_v]$ by assigning the type σ to every tuple that has type θ in $\mathfrak{A}^o[U_v]$. (Note that this can change the \mathcal{E}_o -type of the tuple. But for vertices without predecessors the types are arbitrary, so we can permute them as we like.) This ensures that all colour traces to v are linearly ordered and that they satisfy the global consistency condition. \square

4 Tree-width and Crossing-width

Let us compare crossing-width and tree-width. In Section X.2 below we will not only need to know the relationship between the values of these parameters, but also how the respective index trees of the corresponding decompositions are related. Therefore, we have to include this information in the following statements making them a bit unwieldy. First, note that the crossing-width generalises tree-width in the sense that structures with small tree-width also have a small crossing-width.

Proposition 4.1. *For every finite relational signature Σ and every constant $k < \omega$, there exists a number $N < \omega$ with the following property. For every tree-decomposition $(U_t)_{t \in T}$ of a Σ -structure \mathcal{A} with width at most k , there exists a partition decomposition $(W_s)_{s \in S}$ of \mathcal{A} of width at most N whose index tree S is obtained from T by adding at most k new leaves to every vertex.*

Proof. Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathcal{A} of width $k := \text{twd } \mathcal{A}$. By choosing a root of T , we may assume that T is an order-tree. We construct a partition decomposition $(W_s)_{s \in S}$ of \mathcal{A} as follows. Let S be the tree obtained from T by adding to every vertex t as many new leaves as there are elements in the set

$$C_t := U_t \setminus \bigcup \{ U_v \mid v < t \}.$$

Then S has height at most ω . For original vertices $t \in T$, we define the components by

$$W_t := \bigcup \{ C_v \mid v \geq t \},$$

while, for the new leaves s_0, \dots, s_{n-1} attached to a vertex $t \in T$ with associated set $C_t = \{v_0, \dots, v_{n-1}\}$, we set

$$W_{s_i} := \{v_i\}, \quad \text{for } i < n.$$

Then $(W_s)_{s \in S}$ is a partition decomposition of \mathcal{A} and it remains to compute its width.

Let t be a vertex of S and s_0, \dots, s_{n-1} some of its successors. Note that, if $\bar{c} \in R$ is a tuple in some relation containing elements of both $P := \bigcup_i W_{s_i}$ and its complement, then the latter elements must belong to U_t by the connectedness requirement of a tree-decomposition. Together with the tuples that are not in any relation, it follows that

$$\text{crk}(A \setminus (P \setminus C_t)) \leq (|U_t| + 1)^{r-1} \leq (k + 1)^{r-1}.$$

Hence, Lemma 1.2 (c) implies that

$$\begin{aligned} \text{crk}(A \setminus P) &\leq 2^{r-1} \cdot \text{crk}(C_t) \cdot \text{crk}(A \setminus (P \setminus C_t)) \\ &\leq 2^{r-1} \cdot |C_t|^{r-1} \cdot (k + 1)^{r-1} \\ &\leq 2^{r-1} \cdot k^{r-1} \cdot (k + 1)^{r-1} \\ &= (2k^2 + 2k)^{r-1}. \end{aligned}$$

By Lemma 1.2 (d), it follows that there exists a constant m (only depending on the signature) such that

$$\text{crk}(W_{s_0} \cup \dots \cup W_{s_{n-1}}) \leq m^{\text{crk}(A \setminus P)} \leq m^{(2k^2 + 2k)^{r-1}}. \quad \square$$

Corollary 4.2. *For every finite relational signature Σ and every constant $k < \omega$, there exists a number $N < \omega$ such that every Σ -structure \mathcal{A} with $\text{twd } \mathcal{A} \leq k$ has a partition decomposition of width at most N where the index tree has height at most ω .*

The converse is not true in general. We have seen above that cliques have a small crossing-width while their tree-width is unbounded. The obstruction turns out to be large bipartite graphs embedded in the structure. As above, we start with a technical statement containing additional information about the index trees.

Lemma 4.3. *Let Σ be a finite relational signature and $k, \beta < \omega$. There exists a number $N < \omega$ and an MSO-interpretation τ with the following property.*

If s is a \oplus -term of height at most ω such that the Gaifman graph of $\text{val}(s)$ does not contain the complete bipartite graph $\mathbb{K}_{\beta, \beta}$ (as a non-induced subgraph), then $\tau(s)$ is a tree-decomposition of $\text{val}(s)$ of width at most N whose index tree is equal to the underlying tree of s .

(To produce a tree-decomposition by an interpretation, we encode the decompositions as a tree where each vertex is labelled by the isomorphism type of its component and each edge is labelled by information about which elements of the two components should be identified.)

Proof. Let r be the maximal arity of a relation in Σ . As the claim is trivial for $r = 1$, we may assume w.l.o.g. that $r \geq 2$. Let Ξ be the auxiliary signature used by s and let T be the underlying tree. For $t \in T$, let U_t be the set of all leaves v of T with $v \geq t$. Note that $(U_t)_{t \in T}$ forms a partition decomposition of \mathfrak{A} .

Given a tuple $\bar{c} \subseteq U_t$, we denote by $\chi_t(\bar{c})$ the colour (atomic Ξ -type) assigned to \bar{c} by the subterm $s|_t$. Note that this colour uniquely determines the $\simeq_{A \setminus U_t}$ -class of \bar{c} . We set $[\alpha]_t := \chi_t^{-1}(\alpha)$ and, by abuse of notation, $[\bar{c}]_t := [\chi_t(\bar{c})]_t$. We denote by $S(t)$ the set of successors of a vertex $t \in T$. The support of a colour α at t is the set

$$S(t; \alpha) := \{s \in S(t) \mid [\alpha]_s \neq \emptyset\}.$$

For a vertex $t \in T$, we set

$$W_t := B(t) \cup C(t) \cup D(t),$$

where

$$B(t) := \bigcup \{ [\alpha]_s \mid s \in S(t; \alpha), |[\alpha]_s| < \beta, |S(t; \alpha)| < r\beta \},$$

$$C(t) := \bigcup \{ \bar{c} \setminus U_s \mid \text{there are } R \in \Sigma, \bar{c} \in R, s \leq t \text{ such that } \bar{c} \cap U_s = \bar{c} \cap U_t \text{ and } |[\bar{c} \cap U_s]_s| \geq \beta \},$$

$$D(t) := \bigcup \{ \bar{c} \setminus U_s \mid \text{there are } R \in \Sigma, \bar{c} \in R, s \leq t, p \text{ parent of } s, \bar{c} \cap U_s = \bar{c} \cap U_t, |S(p; \chi_s(\bar{c} \cap U_s))| \geq r\beta \}.$$

Note that each of these sets is MSO-definable from t . Consequently, there exists an MSO-transduction mapping the tree T to (a tree encoding) $(W_t)_{t \in T}$. It therefore remains to prove that $(W_t)_{t \in T}$ is a tree-decomposition of \mathfrak{A} of width at most

$$N := r(\beta + 1)(\beta - 1)M < r\beta^2 M,$$

where M is the number of colours (atomic Ξ -types).

We start by showing that $(W_t)_{t \in T}$ is indeed a tree-decomposition. For every $a \in A$, there is some leaf $t \in T$ with $U_t = \{a\}$. This implies that $a \in B(U_t)$. Hence, $\bigcup_t W_t = A$.

To check that all tuples from the relations of \mathfrak{A} are covered by some component W_t , consider an n -tuple $\bar{c} \in R$. Since the set A^n has a single \simeq_{\emptyset} -class, we have $|\bar{c}|_{\langle \rangle} \geq \beta$ or $|A^n| < \beta$. In the latter case, $B(\langle \rangle) = A$ and we are done. Hence, suppose the former. It is sufficient to show that, for every $t \in T$,

$$|[\bar{c} \cap U_t]_t| \geq \beta \quad \text{implies} \quad \bar{c} \subseteq W_s, \quad \text{for some } s \geq t.$$

Then the claim follows for $t = \langle \rangle$. The proof proceeds by induction on the number of vertices $s \geq t$ with $\bar{c} \cap U_s \neq \emptyset$. Hence, fix t as above. If $|[\bar{c} \cap U_s]_s| \geq \beta$ for some $s > t$, we can use the inductive hypothesis to obtain some $u \geq s$ with $\bar{c} \subseteq W_u$. Consequently, we may assume that $|[\bar{c} \cap U_s]_s| < \beta$, for all successors s of t . Let s_0, \dots, s_{m-1} be an enumeration of all $s \in S(t)$ with $\bar{c} \cap U_s \neq \emptyset$, and set $\bar{c}^i := \bar{c} \cap U_{s_i}$ and $\alpha_i := \chi_{s_i}(\bar{c}^i)$. If there is no index i such that $|S(t; \alpha_i)| \geq r\beta$, then $\bar{c} \cap U_t \subseteq B(t)$ and $\bar{c} \setminus U_t \subseteq C(t)$. Hence, $\bar{c} \subseteq W_t$ and we are done. Otherwise, let i be such an index. Then $\bar{c} \setminus U_{s_i} \subseteq D(s_i)$. We prove by induction on the number of vertices $v \geq s_i$ with $\bar{c} \cap U_v \neq \emptyset$ that there is some $u \geq s_i$ with $\bar{c} \subseteq W_u$. If $\bar{c} \cap U_{s_i} \subseteq B(s_i)$, we are done. Otherwise, there is some successor u of s_i with

$$|S(s_i; \chi_u(\bar{c} \cap U_u))| \geq r\beta.$$

Consequently, $\bar{c} \setminus U_u \subseteq D(u)$ and the claim follows by inductive hypothesis.

It remains to check the connectivity condition. Suppose that $a \in W_s \cap W_t$. We have to show that $a \in W_v$, for every vertex v on the path between s and t . For the proof, we distinguish the following cases.

(i) If $a \in B(s)$ and $a \in B(t)$, then s and t are comparable. By symmetry, we may assume that $s \leq t$. Then $|[a]_v| < \beta$, for all $s < v \leq t$, implies that $a \in B(v) \subseteq W_v$, for every such v .

(ii) If $a \in C(s)$ and $a \in B(t)$, then $s \not\leq t$. Let $w := s \sqcap t$. By definition, there is some $u \leq s$ and a guarded tuple \bar{c} with $a \in \bar{c}$, $|[\bar{c} \cap U_u]_u| \geq \beta$,

and $\bar{c} \cap U_u = \bar{c} \cap U_s$. Hence, $a \in C(v)$, for all $u \leq v \leq s$. Furthermore, $|\bar{c} \cap U_u| \geq \beta$ implies that $|\bar{c} \cap U_v| \geq \beta$, for all $w < v < u$, which implies that $a \in C(v)$, for $w < v < u$. Furthermore, for $w < v \leq t$, we must have $|[a]_v| < \beta$ since, otherwise, $[a]_v \times [\bar{c} \cap U_u]_u$ would contain a copy of $\mathfrak{R}_{\beta, \beta}$. This implies that $a \in B(v)$ for $w < v < t$. Finally, we also have $a \in B(w)$ since $|S(w; \chi_v(a))| \geq r\beta$ (where v is the successor of w with $v \leq t$) would lead to a large bipartite graph $\bigcup_{v \in S(w)} [a]_v \times [\bar{c} \cap U_u]_u$. Thus, $a \in W_v$ for all v between s and t .

(iii) If $a \in C(s)$ and $a \in C(t)$, we consider the leaf $w \in T$ with $U_w = \{a\}$. By symmetry, we may assume that $s \sqcap w \leq t \sqcap w$. Since $a \in B(w)$, it follows by case (ii) above that $a \in W_v$, for all $s \sqcap w < v \leq w$, all $s \sqcap w < v \leq s$, and all $s \sqcap t < v \leq t$. In particular, $a \in W_v$ holds for all vertices v on the path from s to t .

(iv) Finally, suppose that $a \in D(t)$ and $a \in W_u$. Choose $s \leq t$, p , \bar{c} as in the definition of $D(t)$. Then $a \in D(v)$ for all $s \leq v \leq t$. Furthermore,

$$|S(p; \chi_s(\bar{c} \cap U_s))| \geq r\beta \quad \text{implies} \quad |[\bar{c} \cap U_p]_p| \geq r\beta.$$

If $a \notin U_p$, we therefore have $a \in C(p)$. Otherwise, we claim that $a \in B(p)$. In both cases it then follows by (i)–(iii) that a belongs to all vertices between t and u . Hence, it remains to prove the claim. Let v be the successor of p with $a \in U_v$. If $|[a]_v| \geq \beta$, we would have a large bipartite graph

$$[a]_v \times (S(p; \chi_s(\bar{c} \cap U_s)) \setminus (P \cup \{v\})),$$

where $P := \{r \in S(p) \mid \bar{c} \cap U_r \neq \emptyset\}$ has size at most $r - 1$. A contradiction. Hence, $|[a]_v| < \beta$. For a contradiction, suppose that $a \notin B(p)$. Then it follows that $|S(p; \chi_v(a))| \geq r\beta$. Consequently, we can choose disjoint sets

$$P \subseteq S(p; \chi_s(\bar{c} \cap U_s)) \quad \text{and} \quad Q \subseteq S(p; \chi_v(a))$$

of size $|P|, |Q| \geq \beta$. These induce the bipartite graph $P \times Q$. A contradiction.

This concludes the proof that $(W_t)_t$ is a tree-decomposition. It remains to show that $|W_t| \leq N$. Clearly,

$$|B(t)| \leq (\beta - 1)(r\beta - 1)M.$$

To compute $|D(t)|$, we set

$$D_\alpha(t) := \bigcup \left\{ \bar{c} \setminus U_s \mid \begin{array}{l} \text{there are } R \in \Sigma, \bar{c} \in R, s \leq t, p \text{ parent of } s, \\ \chi_t(\bar{c} \cap U_t) = \alpha, \bar{c} \cap U_s = \bar{c} \cap U_t \text{ and} \\ |S(p; \chi_s(\bar{c} \cap U_s))| \geq r\beta \end{array} \right\}.$$

We claim that $|D_\alpha(t)| < \beta$ and, hence, $|D(t)| \leq (\beta - 1)M$. For a contradiction, suppose otherwise and choose β distinct elements $a_0, \dots, a_{\beta-1} \in D_\alpha(t)$. By definition, there exist vertices $s_i \leq t$ and tuples \bar{c}_i in some relation R_i such that

$$a_i \in \bar{c}_i \setminus U_{s_i}, \quad \chi_t(\bar{c}_i \cap U_t) = \alpha, \quad \bar{c}_i \cap U_{s_i} = \bar{c}_i \cap U_t, \\ |S(p_i; \chi_{s_i}(\bar{c}_i \cap U_{s_i}))| \geq r\beta,$$

where p_i is the parent of s_i . Set $s := \max \{s_i \mid i < \beta\}$, let p be the parent of s , and set $\alpha' := \chi_s(\bar{c}_i \cap U_{s_i}) = \chi_s(\bar{c}_i \cap U_t)$. Since every tuple \bar{c}_i can intersect at most $r - 1$ components U_u with $u \in S(p; \alpha')$, it follows that we can find a set $P \subseteq S(p; \alpha')$ of size $|P| \geq |S(p; \alpha')| - \beta(r - 1) \geq \beta$ such that

$$\bar{c}_i \cap \bigcup_{u \in P} U_u = \emptyset, \quad \text{for all } i < \beta.$$

For each $i < \beta$ and $u \in P$, the relation R_i contains a tuple \bar{d} with

$$\bar{d} \setminus U_u = \bar{c}_i \setminus U_s \quad \text{and} \quad \bar{d} \cap U_u \in [\alpha']_u.$$

Consequently, the Gaifman graph of \mathfrak{A} contains the subgraph

$$\{a_i\} \times \bigcup_{u \in P} [\alpha']_u.$$

Taking the union, we obtain the graph

$$\{a_i \mid i < \beta\} \times \bigcup_{u \in P} [\alpha']_u,$$

which contains a copy of $\mathfrak{R}_{\beta, \beta}$. A contradiction.

It remains to show that $|C(t)| \leq (r-1)(\beta-1)M$. For a colour α and a vertex $s \leq t$, let

$$C_\alpha(t; s) := \bigcup \left\{ \bar{c} \setminus U_s \mid \text{there are } R \in \Sigma, \bar{c} \in R, \bar{c} \cap U_s = \bar{c} \cap U_t, \right. \\ \left. |[\bar{c} \cap U_s]_s| \geq \beta, \alpha = \chi_s(\bar{c} \cap U_s) \right\}.$$

For each $a \in C_\alpha(t; s)$, we fix a tuple \bar{c}^a witnessing that $a \in C_\alpha(t; s)$. Then the Gaifman graph of \mathfrak{A} contains the bipartite graph

$$[\bar{c}_o^a]_s \times (\bar{c}^a \setminus U_t) = [\alpha]_s \times (\bar{c}^a \setminus U_t).$$

Taking the union, we obtain the graph

$$[\alpha]_s \times H_\alpha(s) \quad \text{where} \quad H_\alpha(s) := \bigcup_{a \in C_\alpha(t; s)} (\bar{c}^a \setminus U_t).$$

Let I be the set of all possible choices for s . We claim that

$$|\bigcup_{s \in I} H_\alpha(s)| < \beta.$$

Suppose otherwise. Then there is some finite subset $I_o \subseteq I$ with

$$|\bigcup_{s \in I_o} H_\alpha(s)| \geq \beta.$$

Let $s \in I_o$ be its maximal element. Then the Gaifman graph contains the bipartite graph

$$\mathfrak{R}_{\beta, \beta} \subseteq [\alpha]_s \times \bigcup_{s \in I_o} H_\alpha(s).$$

A contradiction. Note that $C_\alpha(t; s)$ is just the union of all tuples in $H_\alpha(s)$ each of which has at most $r-1$ components. Consequently, $|C_\alpha(t; s)| \leq (r-1) \cdot |H_\alpha(s)|$ and

$$|C(t)| = \left| \bigcup_{\alpha} \bigcup_{s \in I} C_\alpha(t; s) \right| \leq M \cdot (r-1) \cdot (\beta-1). \quad \square$$

Corollary 4.4. *Let Σ be a finite relational signature and $k, \beta < \omega$. There exists a number $N < \omega$ with the following property.*

Let \mathfrak{A} be a Σ -structure such that the Gaifman graph of \mathfrak{A} does not contain (as a non-induced subgraph) the complete bipartite graph $\mathfrak{R}_{\beta, \beta}$. For every partition decomposition $(U_t)_{t \in T}$ of \mathfrak{A} of width k whose index tree T has height at most ω , there exists a tree-decomposition $(W_t)_{t \in T}$ of \mathfrak{A} with the same index tree T whose width is at most N .

Proof. Fix a partition decomposition $(U_t)_{t \in T}$ of \mathfrak{A} of width k . By Proposition 3.6, there exists a \oplus -term s denoting \mathfrak{A} such that T is the underlying tree of s . Consequently, we can use Lemma 4.3 to construct a tree-decomposition of \mathfrak{A} of width at most N with index tree T . \square

Theorem 4.5. *Let \mathcal{C} be a class of Σ -structures for a finite signature Σ . The following statements are equivalent.*

- (1) $\text{twd } \mathcal{C} < \infty$.
- (2) $\text{cwd } \mathcal{C} < \infty$ and there is some $s < \omega$ such that all structures in \mathcal{C} are s -sparse.
- (3) $\text{cwd } \mathcal{C} < \infty$ and there exists a number $\beta < \omega$ such that the complete bipartite graph $\mathfrak{R}_{\beta, \beta}$ is not a (non-induced) subgraph of any Gaifman graph of a structure in \mathcal{C} .

Proof. (3) \Rightarrow (1) follows by Corollary 4.4; (1) \Rightarrow (2) by Proposition 4.1 and Lemma VIII.1.8; and (2) \Rightarrow (3) holds since $\mathfrak{R}_{\beta, \beta}$ is not $(\beta-2)$ -sparse (as a structure over the signature $\{E\}$; as a graph it is $\frac{1}{2}(\beta-1)$ -sparse but not $\frac{1}{2}(\beta-2)$ -sparse). \square

Our final remark allows us to transfer many results about crossing-width (such as those in the next section) to tree-width. For the application we have in mind, it is also necessary to keep track of the index trees while doing so.

Proposition 4.6. *Let \mathcal{C} be a class of Σ -structures for a finite signature Σ , \mathcal{C}_{in} the corresponding class of incidence structures, let \mathcal{T} be a class of trees (of height at most ω) and let \mathcal{T}_+ be the class of trees obtained from \mathcal{T} by adding to every vertex a bounded number of new leaves. The following statements are equivalent.*

- (1) There exists a constant $k < \omega$ such that every $\mathfrak{A} \in \mathcal{C}$ has a tree-decomposition of width at most k whose index tree belongs to \mathcal{T} .
- (2) There exists a constant $k < \omega$ such that every incidence structure $\mathfrak{A} \in \mathcal{C}_{\text{in}}$ has a tree-decomposition of width at most k whose index tree belongs to \mathcal{T} .
- (3) There exists a constant $k < \omega$ such that every incidence structure $\mathfrak{A} \in \mathcal{C}_{\text{in}}$ has a partition decomposition of width at most k whose index tree belongs to \mathcal{T}_+ .

Proof. (2) \Rightarrow (3) follows by Proposition 4.1.

(3) \Rightarrow (2) By Corollary 4.4, every structure $\mathfrak{A} \in \mathcal{C}_{\text{in}}$ has a tree-decomposition $(U_t)_{t \in T}$ whose width N is bounded in terms of k and whose index tree T belongs to \mathcal{T}_+ . By definition, T is obtained from some tree $T_o \in \mathcal{T}$ by adding at most $m < \omega$ leaves to every vertex. Let $S(t)$ be the set of leaves added to $t \in T_o$. We construct a new tree-decomposition $(W_t)_{t \in T_o}$ with components

$$W_t := U_t \cup \bigcup_{s \in S(t)} U_s.$$

Then $|W_t| \leq (m+1)N$ is bounded.

(1) \Rightarrow (2) Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{A} of width k . For every tuple \bar{c} in some relation of \mathfrak{A} , we pick some index $t(\bar{c}) \in T$ with $\bar{c} \subseteq U_{t(\bar{c})}$ and we set

$$W_t := U_t \cup \{ \bar{c} \mid t(\bar{c}) = t \}.$$

It is straightforward to check that $(W_t)_{t \in T}$ is a tree-decomposition of \mathfrak{A}_{in} . Furthermore, since U_t contains at most $(|U_t| + 1)^r$ tuples of arity at most r , we have

$$|W_t| \leq k + (k+1)^r.$$

(2) \Rightarrow (1) Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{A}_{in} of width k . Let W_t be the set obtained from U_t by replacing each tuple $\bar{c} \in U_t$ by its components, i.e.,

$$W_t := (U_t \cap A) \cup \bigcup \{ \bar{c} \mid \bar{c} \in U_t \setminus A \}.$$

Then $|W_t| \leq kr$ and it is straightforward to check that $(W_t)_{t \in T}$ is a tree-decomposition of \mathfrak{A} . \square

Corollary 4.7. Let \mathcal{C} be a class of Σ -structures and \mathcal{C}_{in} the corresponding class of incidence structures. Then

$$\text{twd } \mathcal{C} < \aleph_o \quad \text{iff} \quad \text{cwd } \mathcal{C}_{\text{in}} < \aleph_o.$$

5 Interpretations

There is a close connection between MSO-interpretations and the decompositions for crossing-width.

Proposition 5.1. Let τ be a CMSO-interpretation. For every $k < \omega$, there exists a number $n < \omega$ such that, if a structure \mathfrak{A} has a partition decomposition $(U_t)_{t \in T}$ of width at most k , then $\tau(\mathfrak{A})$ has a partition decomposition $(W_s)_{s \in S}$ of width at most n where the index tree S can be obtained from T by deleting some subtrees.

Proof. Suppose that $\tau = \langle \delta(x), (\varphi_R(\bar{x}))_R \rangle$ and let $(U_t)_{t \in T}$ be a partition decomposition of \mathfrak{A} of width at most k . We set

$$W_t := U_t \cap \delta^{\mathfrak{A}} \quad \text{and} \quad S := \{ t \in T \mid W_t \neq \emptyset \}.$$

Then $(W_s)_{s \in S}$ is a partition decomposition of $\tau(\mathfrak{A})$. To compute its width, let f be the function from Corollary 1.7. Setting $n := f(k)$ it follows that

$$\text{crk}(W_s/\mathfrak{A}) \leq k \quad \text{implies} \quad \text{crk}(W_s/\tau(\mathfrak{A})) \leq n, \quad \text{for all } s \in S. \quad \square$$

If we require more precise information about the index tree of the decomposition, we have to use a special kind of interpretation.

Definition 5.2. An interpretation $\tau = \langle \delta(x), (\varphi_R(\bar{x}))_R \rangle$ is *leaf-based* if

$$\mathfrak{Z} \models \delta(v) \quad \text{iff} \quad v \text{ is a leaf, for every tree } \mathfrak{Z}.$$

Lemma 5.3. *Let τ be a leaf-based CMSO-interpretation. There exists a number $k < \omega$ such that, for every order tree \mathfrak{T} , the image $\tau(\mathfrak{T})$ has a partition decomposition $(U_t)_{t \in T}$ of width at most k with index tree \mathfrak{T} .*

Proof. Note that every order-tree \mathfrak{T} has a partition decomposition $(U_s)_{s \in S}$ of width 1 whose index tree S can be obtained from \mathfrak{T} by adding a new leaf to every vertex. We can therefore use Proposition 5.1 to find a partition decomposition $(W_s)_{s \in S_0}$ of $\tau(\mathfrak{T})$ whose index tree S_0 is obtained from S by removing some subtrees and whose width is bounded by some constant k only depending on τ . A closer look at the proof of Proposition 5.1 shows that the vertices $s \in S \setminus S_0$ are exactly those such that $U_s \cap \delta^{\mathfrak{T}} = \emptyset$. As τ is leaf-based, it follows that $S \setminus S_0$ contains the new vertices added to an internal vertex of \mathfrak{T} . Hence, S_0 is obtained from \mathfrak{T} by adding a new successor to every leaf. Thus, if $s \in S_0 \setminus T$ is the successor of $t \in T$, then $W_s = W_t = U_t$ is a singleton. It follows that we can remove the component W_s from the decomposition again without violating any of the axioms of a partition decomposition. Consequently, $(W_s)_{s \in T}$ is also a partition decomposition of $\tau(\mathfrak{T})$ of width k . \square

Lemma 5.4. *The function val mapping a \oplus -term to the corresponding structure is a leaf-based MSO-interpretation.*

Proof. Let t be an \oplus -term over the signatures Σ and Ξ . The elements of $\text{val}(t)$ are the leaves of t , which are MSO-definable. Hence, it remains to find formulae φ_R for each relation R .

For a tuple \bar{v} of leaves of t and a vertex $u \leq \inf \bar{v}$, we say that the *colour* of \bar{v} at u is the set

$$C(u; \bar{v}) := \{ R \in \Xi \mid \bar{v} \in R \text{ in the structure } \text{val}(t|_u) \}.$$

Let \bar{v} be a tuple of leaves in t . By induction on $|\bar{v}|$, we construct formulae $\psi_R(x, \bar{y})$, for $R \in \Xi$, such that

$$t \models \psi_R(u, \bar{v}) \quad \text{iff} \quad u \leq \inf \bar{v} \text{ and } R \in C(u; \bar{v}).$$

The formula ψ_R guesses the colours $C(w; \bar{v})$, for every $u \leq w \leq \inf \bar{v}$, and then checks whether $R \in C(u; \bar{v})$. We can verify that the guesses are correct by induction starting at $w := \inf \bar{v}$.

Hence, suppose that $w := \inf \bar{v}$. If $v_0 = \dots = v_{n-1}$, we have $w = v_0$ and we can read off $C(v_0; \bar{v})$ from the label at 0 . Otherwise, \bar{v} splits at some vertices w, s_0, \dots, s_{n-1} . Let $\bar{v}^0, \dots, \bar{v}^{n-1}$ be its split. By inductive hypothesis, we can compute the colours $C(s_i; \bar{v}^i)$. From this data and the operation labelling w , we can now derive $C(w; \bar{v})$.

For the inductive step, suppose that we already know $C(w; \bar{v})$ and let w' be the predecessor of w . Then we can compute $C(w'; \bar{v})$ from $C(w; \bar{v})$ by looking at the operation labelling w . \square

We obtain the following characterisation of classes of finite crossing-width.

Theorem 5.5. *Let \mathcal{T} be a class of order-trees, possibly of ordinal height, and let \mathcal{C} be a class of Σ -structures for a finite relational signature Σ . We denote by $\mathcal{T}^{(k)}$ the class of all coloured trees with k colours whose underlying tree belongs to \mathcal{T} . The following statements are equivalent.*

- (1) *There exists a constant $m < \omega$ such that every structure in \mathcal{C} has a partition decomposition of width at most m whose index tree belongs to \mathcal{T} .*
- (2) *Every structure in \mathcal{C} can be denoted by a \oplus -term whose underlying tree belongs to \mathcal{T} .*
- (3) *There exists a leaf-based MSO-interpretation of \mathcal{C} in some subclass of $\mathcal{T}^{(k)}$, for some finite k .*
- (4) *There exists a leaf-based CMSO-interpretation of \mathcal{C} some subclass of $\mathcal{T}^{(k)}$, for some finite k .*

If the class \mathcal{T} is the class of all trees of height at most α , where α is either finite, a limit ordinal, or $\alpha = \infty$, the following statement is equivalent to the ones above.

- (5) *There exists a leaf-based FO-interpretation of \mathcal{C} in $\mathcal{T}^{(k)}$, for some finite k .*

Proof. The implications (1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (1) follow by, respectively, Proposition 3.6, Lemma 5.4, and Lemma 5.3; and (4) \Rightarrow (3) is trivial.

(3) \Rightarrow (4) If \mathcal{T} is a class of trees whose height is bounded by some finite $n < \omega$, the claim follows by Proposition V.1.17. If \mathcal{T} is the class of trees of height at most α , where α is a limit ordinal, the claim follows by Corollary V.1.16. \square

If we do not care what the class of trees looks like, we obtain the following statement.

Corollary 5.6. *Let \mathcal{C} be a class of Σ -structures for a finite relational signature Σ . The following statements are equivalent.*

- (1) $\text{cwd } \mathcal{C} < \aleph_0$
- (2) Every structure in \mathcal{C} can be denoted by a \oplus -term.
- (3) There exists a leaf-based CMSO-interpretation of \mathcal{C} in some class of trees.
- (4) There exists a CMSO-interpretation of \mathcal{C} in some class of trees.
- (5) There exists a leaf-based MSO-interpretation of \mathcal{C} in some class of trees.
- (6) There exists an MSO-interpretation of \mathcal{C} in some class of trees.
- (7) There exists a leaf-based FO-interpretation of \mathcal{C} in some class of trees.
- (8) There exists an FO-interpretation of \mathcal{C} in some class of trees.

Proof. The equivalences follow from the preceding theorem and the following two additional arguments.

First, note that we can encode every tree with colours $\{c_0, \dots, c_{k-1}\}$ as an uncoloured tree by adding $i + 1$ new leaves to each vertex v with label c_i .

Second, we can turn an arbitrary interpretation (CMSO, MSO, or FO) into a leaf-based one by adding leaves to the input tree. \square

Remark. The same result holds in the important case where all structures in \mathcal{C} are countable and we only consider countable trees of height ω . \square

For a single structure we obtain.

Corollary 5.7. *Let \mathcal{A} be a Σ -structure with a finite relational signature Σ . The following statements are equivalent.*

- (1) $\text{cwd } \mathcal{A} < \aleph_0$
- (2) \mathcal{A} can be denoted by a \oplus -term.

- (3) There exists a leaf-based CMSO-interpretation of \mathcal{A} in some tree.
- (4) There exists a CMSO-interpretation of \mathcal{A} in some tree.
- (5) There exists a leaf-based MSO-interpretation of \mathcal{A} in some tree.
- (6) There exists an MSO-interpretation of \mathcal{A} in some tree.
- (7) There exists a leaf-based FO-interpretation of \mathcal{A} in some tree.
- (8) There exists an FO-interpretation of \mathcal{A} in some tree.

As an application, we obtain the following Löwenheim-Skolem theorem for classes of bounded crossing width.

Proposition 5.8. *Let Σ be a finite signature and \mathcal{C} a non-empty class of Σ -structures with $\text{cwd } \mathcal{C} < \aleph_0$. Then \mathcal{C} contains a structure of size at most κ .*

Proof. By Corollary 5.6, there exists an MSO-interpretation τ and a class \mathcal{T} of trees such that $\tau[\mathcal{T}] = \mathcal{C}$. Furthermore, we can use Corollary V.4.4 to find a countable tree $\mathfrak{T} \in \mathcal{T}$. Its image $\tau(\mathfrak{T}) \in \mathcal{C}$ is the desired countable structure in \mathcal{C} . \square

The analogue of the Excluded Grid Theorem (Theorem VIII.4.1) for crossing-width is still open.

Open Question (Seese's Conjecture). *Let \mathcal{C} be a class of Σ -structures with unbounded crossing-width. Does there exist an MSO-interpretation of the class of all finite grids in \mathcal{C} ?*

Currently, only the following partial result is known.

Theorem 5.9 (Courcelle, Oum). *Let \mathcal{C} be a class of undirected graphs with unbounded crossing-width. Then there exist a CMSO-interpretation mapping \mathcal{C} to the class of all finite grids.*

We conclude this section by taking a look at two other operations that are compatible with monadic second-order logic. The first one is the Muchnik iteration.

Proposition 5.10. $\text{cwd } \mathcal{A}^* \leq 2^{r-1} \cdot \text{cwd } \mathcal{A}$

Proof. Let $(U_t)_{t \in T}$ be a partition decomposition of \mathfrak{A} of width $\text{cwd } \mathfrak{A}$ and let L be the set of leaves of T . Let S be the tree obtained from T by recursively attaching copies of T to every leaf. Formally, we set

$$S := L^*(T + 1),$$

with the following variant of the prefix ordering. We set

$$\langle s_0, \dots, s_{m-1} \rangle \leq \langle t_0, \dots, t_{n-1} \rangle$$

if, and only if, $m \leq n$, $s_i = t_i$, for $i < m$, and either

$$m = n \text{ and } s_{m-1} \leq t_{m-1}, \quad \text{or} \quad m < n \text{ and } s_{m-1} = t_{m-1}.$$

The components are

$$\begin{aligned} W_{\langle t_0, \dots, t_{n-1}, t_n \rangle} &:= \{ a_0 \dots a_{n-1} a_n w \mid a_i \in U_{t_i}, w \in A^* \}, \\ W_{\langle t_0, \dots, t_{n-1}, * \rangle} &:= \{ a_0 \dots a_{n-1} \} \quad \text{where } U_{t_i} = \{ a_i \} \end{aligned}$$

(where $*$ denotes the unique element of 1).

Then $(W_v)_{v \in S}$ is a partition decomposition and it remains to compute its width. Hence, fix a vertex $v = \langle t_0, \dots, t_{n-1} \rangle \in S$ and set P of successors of v . If t_{n-1} is a leaf, then v has exactly two successors with components

$$\begin{aligned} W_{\langle t_0, \dots, t_{n-1}, \langle \rangle \rangle} &:= \{ a_0 \dots a_{n-1} w \mid U_{t_i} = \{ a_i \}, w \in A^* \}, \\ W_{\langle t_0, \dots, t_{n-1}, * \rangle} &:= \{ a_0 \dots a_{n-1} \} \quad \text{where } U_{t_i} = \{ a_i \}. \end{aligned}$$

The corresponding ranks are

$$\text{crk}(W_{\langle t_0, \dots, t_{n-1}, \langle \rangle \rangle}) = 1 \quad \text{and} \quad \text{crk}(W_{\langle t_0, \dots, t_{n-1}, * \rangle}) = 1.$$

Otherwise, we have

$$\bigcup_{s \in P} W_s = wVA^* \quad \text{where } \{w\} = U_{t_0} \times \dots \times U_{t_{n-1}} \text{ and } V := \bigcup_{s \in P} U_s.$$

Setting

$$X := wV, \quad Q := w(A \setminus V), \quad Y := wVA^+, \quad Z := A^* \setminus wA^*$$

we have $\bigcup_{s \in P} W_s = X \cup Y$ and $A \setminus \bigcup_{s \in P} W_s = Q \cup Z$. Furthermore, for $\bar{a}, \bar{a}' \subseteq X$ and $\bar{b}, \bar{b}' \subseteq Y$ we have

$$\bar{a}\bar{b} \simeq_{Q \cup Z} \bar{a}'\bar{b}' \quad \text{iff} \quad \bar{a} \simeq_Q \bar{a}'.$$

Considering the various ways an $(r-1)$ -tuple can be distributed over X and Y it therefore follows that

$$\begin{aligned} \text{crk}(\bigcup_{s \in P} W_s / \mathfrak{A}^*) &= |(X \cup Y)^{r-1} / \simeq_{Q \cup Z}| \\ &\leq \sum_{k \leq r-1} \binom{r-1}{k} |X^k / \simeq_Q| \\ &\leq 2^{r-1} |X^{r-1} / \simeq_Q| \\ &\leq 2^{r-1} \cdot \text{crk}(\bigcup_{s \in P} U_s / \mathfrak{A}), \end{aligned} \quad \square$$

The other operation we are interested in are generalised sums.

Proposition 5.11. *Let Σ and Γ be finite relational signatures, \mathfrak{I} a Γ -structure, and $(\mathfrak{A}_i)_{i \in I}$ a family of Σ -structures. Suppose that there exists a constant $k < \omega$ such that*

$$\text{cwd } \mathfrak{A}_i \leq k, \quad \text{for all } i \in I.$$

Then

$$\text{cwd } \sum_{i \in I} \mathfrak{A}_i \leq \max \{k, \text{cwd } \mathfrak{I}\}.$$

Proof. Fix partition decompositions $(W_s)_{s \in S}$ of \mathfrak{I} and $(U_t)_{t \in T_i}$ of \mathfrak{A}_i of widths at most $\text{cwd } \mathfrak{I}$ and k , respectively. Let $\mu : I \rightarrow S$ be the function mapping an element $i \in I$ to the unique leaf $\mu(i)$ of S with $W_{\mu(i)} = \{i\}$. We construct a partition decomposition $(V_r)_{r \in F}$ of $\sum_i \mathfrak{A}_i$ as follows. The

index tree F is the tree obtained from S by adding the tree $T_{\mu^{-1}(s)}$ to each leaf s . Formally,

$$F = S + \sum_{i \in I} T_i$$

with ordering, for $s, s' \in S$, $i, i' \in I$, $t \in T_i$, $t' \in T_{i'}$,

$$\begin{aligned} s \leq s' & \quad : \text{iff} \quad s \leq s' \text{ in } S, \\ s \leq \langle i, t \rangle & \quad : \text{iff} \quad s \leq \mu(i), \\ \langle i, t \rangle \leq \langle i', t' \rangle & \quad : \text{iff} \quad i = i' \text{ and } t \leq t' \text{ in } T_i. \end{aligned}$$

The components are

$$\begin{aligned} V_s &:= \{ \langle i, a \rangle \mid \mu(i) \geq s \text{ and } a \in A_i \}, \quad \text{for } s \in S, \\ V_{\langle i, t \rangle} &:= \{ \langle i, a \rangle \mid a \in U_t^i \}, \quad \text{for } i \in I \text{ and } t \in T_i. \end{aligned}$$

Then $(V_r)_{r \in F}$ is a partition decomposition and it remains to compute its width.

For a vertex of the form $\langle i, t \rangle$ and a set P of its successors, we have

$$\text{crk}(\bigcup_{s \in P} V_s) = \text{crk}(\bigcup_{s \in P} U_s^i / \mathfrak{A}_i) \leq k.$$

(Where we have identified P with the corresponding set of successors of t in T_i .)

For a vertex of the form $s \in S$ and a set P of its successors, we proceed as follows. Consider two tuples $\bar{c}, \bar{c}' \in V_s^{r-1}$ with components $c_j = \langle i_j, a_j \rangle$ and $c'_j = \langle i'_j, a'_j \rangle$. We claim that

$$\bar{c} \simeq_{\sum_i A_i \setminus V_s} \bar{c}' \quad \text{iff} \quad \bar{i} \simeq_{I \setminus W_s} \bar{i}'.$$

Then it follows that

$$\text{crk}(V_s) \leq \text{crk}(W_s / \mathfrak{S}) \leq \text{cwd } \mathfrak{S}.$$

To prove the claim, note that we cannot distinguish \bar{c} and \bar{c}' using the relation \sim since V_s is a union of \sim -classes. Similarly, we cannot distinguish them using a relation from Σ . Hence, the only way to distinguish \bar{c} and \bar{c}' is via a relation from F . \square

6 Non-Standard Crossing-Width

We have shown in Section VIII.1 that having a given finite tree-width is a first-order property. Here we derive the analogous results for crossing-width. Again the main technical issue is the fact that being a tree is not first-order axiomatisable. Hence, we have to work with decompositions indexed by non-standard trees (which we have introduced in Definition VIII.1.11).

Definition 6.1. (a) Let $\langle T, \leq, \sqcap \rangle$ be a non-standard tree. We call the maximal elements of T its *leaves*. An element $t \in T$ is *meet-irreducible* if it we cannot write t as $s_0 \sqcap s_1$ with $s_0, s_1 > t$. We call t *meet-reducible* if it is neither a leaf nor meet-irreducible.

(b) A *non-standard partition decomposition* of a structure \mathfrak{A} is a family $(U_t)_{t \in T}$ of non-empty subsets $U_t \subseteq A$ indexed by a non-standard tree T satisfying the following conditions.

- ♦ $s \leq t$ implies $U_s \supseteq U_t$,
- ♦ $U_s \cap U_t = \emptyset$, if s and t are incomparable,
- ♦ for every $a \in A$, there is some $t \in T$ with $U_t = \{a\}$.

(c) The *width* of a partition decomposition $(U_t)_{t \in T}$ of \mathfrak{A} is the supremum over all ranks

$$\text{crk}(\bigcup_{t \in I(C)} U_t),$$

where C ranges over all finite non-empty antichains and

$$I(C) := \downarrow C \setminus \downarrow \inf C.$$

The *non-standard crossing-width* $\text{cwd}_{\text{ns}} \mathfrak{A}$ of a structure \mathfrak{A} is the minimal width of a non-standard partition decomposition of \mathfrak{A} . \lrcorner

A new problem we are facing when trying to adapt the material in Section VIII.1 is that a given structure might not contain enough elements to enable us to encode the index tree of a (non-standard) partition decomposition. It only has element for the *leaves* of the tree. To encode the whole tree we have therefore to use *pairs* of elements. This leads to the following ternary relation.

Definition 6.2. Let $\langle T, \leq, \sqcap \rangle$ be a non-standard tree and let $L \subseteq T$ be the set of its maximal elements. The *holographic relation* on L is

$$H := \{ \langle a, b, c \rangle \in L^3 \mid a \sqcap b \leq c \}.$$

Lemma 6.3. Let $\langle T, \leq, \sqcap \rangle$ be a non-standard tree with maximal elements L , and let H be its holographic relation. The 4-ary relation

$$a \sqcap b \leq c \sqcap d$$

is definable in $\langle L, H \rangle$.

Proof. We have

$$a \sqcap b \leq c \sqcap d \quad \text{iff} \quad a \sqcap b \leq c \quad \text{and} \quad a \sqcap b \leq d. \quad \square$$

Thus we can encode the elements of the original tree T by pairs of leaves, provided that we can express every element as an infimum of leaves, i.e., if every non-leaf of T is meet-reducible. Note that, when considering non-standard partition decompositions, this is no loss of generality since we can always remove all meet-irreducible non-leaves (which are the vertices with exactly 1 successor) from its index tree.

Exercise 6.1. Let $(U_t)_{t \in T}$ be a non-standard partition decomposition of a structure \mathfrak{A} of width k and let $S := \{ s \sqcap t \mid s, t \text{ leaves of } T \}$ be the set of its meet-reducible elements and its leaves. Then $(U_t)_{t \in S}$ is also a non-standard partition decomposition of \mathfrak{A} of width k . \square

Using this encoding we can prove the following compactness result.

Proposition 6.4. Let \mathfrak{A} be a Σ -structure. Then

$$\text{cld}_{\text{ns}} \mathfrak{A} = \sup \{ \text{cld } \mathfrak{C} \mid \mathfrak{C} \subseteq \mathfrak{A} \text{ finite} \}.$$

Proof. Let k be the value of the above supremum. Since $\mathfrak{C} \subseteq \mathfrak{A}$ implies $\text{cld } \mathfrak{C} = \text{cld}_{\text{ns}}(\mathfrak{C}) \leq \text{cld}_{\text{ns}}(\mathfrak{A})$, we have $k \leq \text{cld } \mathfrak{A}$. Hence, it is sufficient to prove that $\text{cld}_{\text{ns}}(\mathfrak{A}) \leq k$. We construct a corresponding non-standard partition decomposition as follows.

Let Δ be the atomic diagram of \mathfrak{A} and let Φ be a set of formulae stating the following properties.

- ♦ H is the holographic relation of some non-standard tree T whose set of leaves is A . (We identify T with a quotient of A^2 via the map $\langle a, b \rangle \mapsto a \sqcap b$.)
- ♦ U is a ternary relation such that the family $(U_{s,t})_{s,t \in A}$ defined by $U_{s,t} := \{ a \in A \mid \langle a, s, t \rangle \in U \}$ forms a partition decomposition with $U_{a,a} = \{a\}$, for all $a \in A$.
- ♦ $\text{crk}(\bigcup_{t \in I} U_t) \leq k$, for all I as in the definition of the width of a non-standard partition decomposition.

Since every finite subset of $\Delta \cup \Phi$ is satisfiable there exists a model $\mathfrak{A}_+ \models \Delta \cup \Phi$. Since \mathfrak{A}_+ satisfies Δ , there exists an embedding $\mathfrak{A} \rightarrow \mathfrak{A}_+$. Hence, we may assume that $\mathfrak{A} \subseteq \mathfrak{A}_+$. Let $(U_{s,t})_{s,t}$ be the corresponding family of sets and set $W_{s,t} := U_{s,t} \cap A$. Clearly, $(W_{s,t})_{s,t}$ is a non-standard partition decomposition of \mathfrak{A} . It remains to show that its width is at most k .

Hence, fix a vertex $\langle s, t \rangle$ of the tree and a set $I = I(C)$ of successors associated with a finite antichain C as in the definition of the width. Then

$$\text{crk}(\bigcup_{t \in I} U_t / \mathfrak{A}_+) \leq k \quad \text{implies} \quad \text{crk}(\bigcup_{t \in I} W_t / \mathfrak{A}) \leq k. \quad \square$$

As for tree-width, it follows that the first-order theory of a structure determines its non-standard crossing-width.

Theorem 6.5. Let \mathfrak{A} and \mathfrak{B} be Σ -structures.

$$\mathfrak{A} \equiv_{\text{FO}} \mathfrak{B} \quad \text{implies} \quad \text{cld}_{\text{ns}} \mathfrak{A} = \text{cld}_{\text{ns}} \mathfrak{B}.$$

Proof. Let $k := \text{cld}_{\text{ns}} \mathfrak{A}$. By symmetry and Proposition 6.4, it is sufficient to show that

$$\text{cld } \mathfrak{B}_0 \leq k, \quad \text{for every finite } \mathfrak{B}_0 \subseteq \mathfrak{B}.$$

For every $n < \omega$, let φ_n be an FO-formula stating that every n -tuple of elements induces a substructure whose crossing-width is at most k . Such a formula exists, since there are only finitely many Σ -structures with at most n elements. Then

$$\mathfrak{A} \models \varphi_n \quad \text{implies} \quad \mathfrak{B} \models \varphi_n, \quad \text{for all } n < \omega. \quad \square$$

It remains to determine the consequences of these results for the standard crossing-width. To do so we have to turn non-standard partition decompositions into standard ones.

Theorem 6.6. *For every signature Σ and every constant k , there exists a number n with the following property. Every Σ -structure \mathfrak{A} with $\text{cwd}_{\text{ns}} \mathfrak{A} < k$ has a (standard) partition decomposition of width n whose index tree has height at most $\max\{|A|, \omega\}$.*

Proof. Let $(U_t)_{t \in T}$ be a non-standard partition decomposition of \mathfrak{A} of width at most k and set $\kappa := \max\{|A|, \omega\}$. If A is finite then so is T . Hence, $(U_t)_{t \in T}$ is a normal partition decomposition and there is nothing to do. We may therefore assume that A and T are infinite.

In fact, we may further assume that $|T| = |A|$. To see this, we fix, for every $a \in A$, some vertex $t \in T$ with $U_t = \{a\}$. Let S_o be the set of these vertices and set $S := \{u \sqcap v \mid u, v \in S_o\}$. Then $(U_t)_{t \in S}$ is a non-standard partition decomposition of the same width whose index tree has size $|S| = |S_o| = |A|$.

We will construct the desired standard partition decomposition as the limit of a sequence of *partial* ones, which are decompositions where we drop the requirement that leaf components are singletons. More precisely, we call a family $(W_s)_{s \in S}$ a *partial partition decomposition* if

- each W_s is a non-empty subset of A ,
- S is an order-tree,
- $s \leq t$ implies $W_s \supseteq W_t$,
- $W_s \cap W_t = \emptyset$, if s and t are incomparable,
- for every $a \in A$, there is a leaf $s \in S$ with $a \in W_s$.

We also need the notion of a *subdecomposition* of $(U_t)_{t \in T}$. This is a family of the form $(U_s)_{s \in S}$ where the index set $S \subseteq T$ is a prefix of some subtree of T , that is, S is either equal to T or there is some vertex $t \in T \setminus S$ such that

$$s \in S \text{ implies } u \in S, \text{ for all } t < u < s \text{ and all } u > s.$$

In the latter case we say that the subdecomposition is *below* the vertex t . We also say that $(U_s)_{s \in S}$ is a subdecomposition of some set C if $C = \bigcup_{s \in S} U_s$.

Fix an enumeration $(a_i)_{i < \kappa}$ of A . We will construct a sequence $(W_s)_{s \in S_i}$ of partial partition decompositions of width $f(k)$ with the following properties.

- $S_i \subseteq S_j$ for $i \leq j$.
- Each S_i has height at most κ .
- For every $j < i$, there is some leaf $s \in S_i$ with $W_s = \{a_j\}$.
- For every leaf $s \in S_i$ with $|W_s| > 1$, there exists a subdecomposition of $(U_t)_{t \in W_s}$.

Then the limit $(W_s)_{s \in S}$ with $S = \bigcup_i S_i$ is the desired partition decomposition of \mathfrak{A} .

We construct $(W_s)_{s \in S_i}$ by induction on i starting with the family where S_o consists of a single vertex $\langle \rangle$ and $W_{\langle \rangle} = A$. For the successor step, suppose that we have already defined $(W_s)_{s \in S_i}$. Let $s_i \in S_i$ be the leaf with $a_i \in W_{s_i}$. By inductive hypothesis, there exists a set $T_i \subseteq T$ such that $(U_t)_{t \in T_i}$ is a subdecomposition of the set W_{s_i} . Let $t_i \in T_i$ be the leaf with $U_{t_i} = \{a_i\}$, set

$$I := \{t \in T_i \mid t \leq t_i\},$$

and let $(K_r)_{r \in R}$ be the partition decomposition of the linear order $\langle I, \leq \rangle$ obtained from Lemma 2.3 (a). Let S_{i+1} be the tree obtained from S_i by replacing the leaf s_i with the tree R . Since the heights of S_i and R are at most κ , the same holds for the height of S_{i+1} . We define the new components as follows. Let $r \in R$ and let L_r be the vertices in I that are larger than all vertices in K_r . We set

$$W_r := \bigcup_{t \in K_r} U_t \setminus \bigcup_{t \in L_r} U_t.$$

(If K_r and L_r have least elements u and v , respectively, this definition simplifies to $W_r = U_u \setminus U_v$.) For a leaf $r \in R$ with $K_r = \{t\}$, it follows that

$$\begin{aligned} W_r &= U_t \setminus \bigcup_{s \in L_r} U_s \\ &= \bigcup \{U_s \mid s > t \text{ and there is no } t < u < s \text{ with } u \in I\}. \end{aligned}$$

This implies that $\text{crk}(W_r) \leq k$ and there is a subdecomposition of $(U_t)_{t \in T_i}$ of W_r . For internal vertices $r \in R$, the fact that $(U_t)_t$ has width k implies that

$$\text{crk}(\bigcup_{t \in K_r} U_t) \leq k \quad \text{and} \quad \text{crk}(\bigcup_{t \in L_r} U_t) \leq k.$$

Consequently, it follows by Lemma 1.2 (e) that

$$\text{crk}(W_r) \leq k^{ck},$$

for some constant c only depending on Σ .

Finally, for the limit step let δ be a limit ordinal and suppose that we have already defined $(W_s)_{s \in S_i}$ for all $i < \delta$. We set $S_\delta = \bigcup_{i < \delta} S_i$. Note that, being a union of trees of height at most κ , the tree S_δ also has height at most κ . \square

Corollary 6.7. $\mathfrak{A} \equiv_{\text{FO}} \mathfrak{B}$ and $\text{cwd} \mathfrak{A} < \aleph_0$ implies $\text{cwd} \mathfrak{B} < \aleph_0$.

Corollary 6.8. Let \mathcal{C} be a class of Σ -structures and \mathcal{C}_{fin} the class of all finite substructures of structures in \mathcal{C} . Then

$$\text{cwd} \mathcal{C}_{\text{fin}} < \aleph_0 \quad \text{implies} \quad \text{cwd} \mathcal{C} < \aleph_0.$$

7 The Monadic Independence Property

The aim of this section is to study combinatorial properties of complicated classes of Σ -structures. In particular, these properties can be used to prove that a given class has unbounded crossing-width or unbounded tree-width. We start with a criterion based on the definability of certain functions.

Definition 7.1. We call a function $f : A \times B \rightarrow C$ *cancellative* if

$$\begin{aligned} f(a, b) = f(a', b) & \text{ implies } a = a', \\ f(a, b) = f(a, b') & \text{ implies } b = b'. \end{aligned}$$

Note that, in particular, every bijective function $A \times B \rightarrow C$ is cancellative.

Proposition 7.2. Let C be a set, κ a cardinal, and $f : A \times B \rightarrow C$ (the graph of) a cancellative function where $A, B \subseteq C$ are sets of size $|A|, |B| > 3\kappa$. Then

$$\text{cwd} \langle C, f \rangle \geq \kappa.$$

Proof. For a contradiction, suppose that there exists a partition refinement $(U_t)_{t \in T}$ of $\langle C, f \rangle$ such that

$$\text{crk}(\bigcup_{i \in P} U_{s_i}) < \kappa, \quad \text{for every } t \in T \text{ and every subset } P \text{ of successors of } t.$$

By Exercise 2.3, we may assume that the index tree T is binary. For $t \in T$, set $A_t := A \cap U_t$ and $B_t := B \cap U_t$.

We claim that there exists a vertex $s \in T$ such that

$$|A_s| > \kappa \quad \text{and} \quad |A \setminus A_s| \geq \kappa.$$

To find s , note that the set $S := \{t \in T \mid |A_t| > 2\kappa\}$ forms a chain. Let $(t_i)_{i < \alpha}$ be an increasing enumeration of S .

If κ is finite, then S has a maximal element s . By symmetry, we may assume that $|A_{t_0}| \geq |A_{t_1}|$. Then

$$2|A_{t_0}| \geq |A_{t_0}| + |A_{t_1}| = |A_t| > 2\kappa \quad \text{implies} \quad |A_{t_0}| > \kappa.$$

Hence, we can set $s := t_0$.

If κ is infinite, $|A_t| = |A_{t_0}| + |A_{t_1}|$ implies that α is a limit ordinal. Let

$$X_i := A_{t_i} \setminus A_{t_{i+1}}, \quad Z := \bigcap_{i < \alpha} A_{t_i}, \quad \text{and} \quad I := \{i < \alpha \mid X_i \neq \emptyset\}.$$

Then $A = Z \cup \bigcup_{i < \alpha} X_i$. As $|A| > \kappa$ and $|Z| \leq \kappa$, it follows that

$$\sum_{i \in I} |X_i| = \sum_{i < \alpha} |X_i| = |A \setminus Z| > \kappa.$$

Since $|X_i| \leq \kappa$ and κ^+ is regular we have $|I| > \kappa$. Consequently, there is some $\beta < \alpha$ such that $|I \cap \beta| \geq \kappa$. Setting $s := t_\beta$ it follows that

$$|A_s| > \kappa \quad \text{and} \quad |A \setminus A_s| = \sum_{i < \beta} |X_i| \geq |I \cap \beta| \geq \kappa.$$

Having found a vertex $s \in T$ as above we distinguish two cases. We have $|B_s| \geq \kappa$ or $|B \setminus B_s| \geq \kappa$. We consider both case separately.

(I) First, suppose that $|B_s| \geq \kappa$. We start by showing that there are less than κ elements $b \in B_s$ such that $f(a, b) = c$, for some $a \in A \setminus A_s$ and $c \in C \setminus U_s$. Otherwise, we could find elements $a, a' \in A \setminus A_s, c, c' \in C \setminus U_s$, and $b, b' \in B_s$ with

$$b \neq b', \quad f(a, b) = c, \quad f(a', b') = c', \quad \text{and} \quad b \simeq_{C \setminus U_s} b'.$$

But then $f(a', b) = c'$ implies $f(a', b) = c'$. Hence, $f(a', b) = f(a', b')$. Since f is cancellative it follows that $b = b'$. A contradiction.

Consequently, there is some $b \in B_s$ such that $f(a, b) \in U_s$, for all $a \in A \setminus A_s$. Since $|A \setminus A_s| \geq \kappa$, we can find $a \neq a'$ in $A \setminus A_s$ with $b, f(a, b) \simeq_{C \setminus U_s} b, f(a', b)$. Hence, we have

$$f(a, b) = c \Leftrightarrow f(a', b) = c, \quad \text{for all } a \in A \setminus A_s.$$

Again a contradiction.

(II) It remains to consider the case where $|B \setminus B_s| \geq \kappa$. First, we prove that there are less than κ elements $b \in B \setminus B_s$ such that $f(a, b) = c$, for some $a \in A_s$ and $c \in U_s$. Otherwise, we could find elements $a, a' \in A_s, c, c' \in U_s$, and $b, b' \in B \setminus B_s$ with

$$b \neq b', \quad f(a, b) = c, \quad f(a', b') = c', \quad \text{and} \quad ac \simeq_{C \setminus U_s} a'c'.$$

But then $f(a, b) = c$ implies $f(a', b) = c'$. Hence, $f(a', b) = f(a', b')$. Since f is cancellative it follows that $b = b'$. A contradiction.

Consequently, there is some $b \in B \setminus B_s$ such that $f(a, b) \in C \setminus U_s$, for all $a \in A_s$. Since $|A_s| \geq \kappa$ we can find $a \neq a'$ in A_s with $a \simeq_{C \setminus U_s} a'$. Hence, we have

$$f(a, b) = c \Leftrightarrow f(a', b) = c, \quad \text{for all } c \in C \setminus U_s.$$

Again a contradiction. \square

Corollary 7.3. A group $\mathfrak{G} = \langle G, \cdot \rangle$ has finite crossing-width if, and only if, it is finite.

Remark. In Theorem XI.6.4 below, we will prove that the Cayley graph of a group has finite crossing-width if, and only if, the group is virtually free. \square

Note that every $n \times n$ grid can be used to define (in MSO) a cancellative function $A \times B \rightarrow C$ with $|A| = n = |B|$. But it turns out that we obtain a more natural class if we use first-order logic instead. But in that case grids are no longer sufficient since we cannot compute transitive closures in FO. The following lemma summarises several equivalent configurations that can be used instead of grids.

Lemma 7.4. Let \mathcal{C} be a class of structure. The following statements are equivalent.

- (1) There exist a monadic expansion \mathcal{C}_+ of \mathcal{C} and an FO-formula $\chi(x, y, z)$ such that, for every $n < \omega$, there is some structure $\mathfrak{M} \in \mathcal{C}_+$ in which φ defines a bijection $A \times B \rightarrow C$ with $A, B, C \subseteq M$ and $|A| = n = |B|$.
- (2) There exist a monadic expansion \mathcal{C}_+ of \mathcal{C} and an FO-formula $\vartheta(x, y)$ such that, for every $n < \omega$, there is some structure $\mathfrak{M} \in \mathcal{C}_+$ and a bijection $\sigma : A \times A \rightarrow B$ with $A, B \subseteq M$ and $|A| = n$ such that

$$\mathfrak{M} \models \vartheta(a, b) \quad \text{iff} \quad \sigma(a, c) = b \text{ or } \sigma(c, a) = b, \quad \text{for some } c.$$

- (3) There exist a monadic expansion \mathcal{C}_+ of \mathcal{C} and two FO-formulae $\varphi(x, y)$ and $\psi(x, y)$ such that, for every $n < \omega$, there is some structure $\mathfrak{M} \in \mathcal{C}_+$ in which φ and ψ define equivalence relations with n classes such that each class of the first one intersects every class of the other one.

Proof. (2) \Rightarrow (1) Fix $n < \omega$ and choose a structure $\mathfrak{M} \in \mathcal{C}_+$ in which the formula ϑ defines pairs for a set A of size $2n$. Fixing a partition $A = P + Q$ with $|P| = n = |Q|$, it follows that the formula

$$\chi(x, y, z) := Px \wedge Qy \wedge \vartheta(x, z) \wedge \vartheta(y, z)$$

defines a bijection $P \times A \rightarrow B$.

(1) \Rightarrow (2) Fix $n < \omega$ and choose a structure $\mathfrak{M} \in \mathcal{C}_+$ in which the formula χ defines a bijection $f : A \times B \rightarrow C$ with $|A| = n = |B|$. Fixing some

bijection $\sigma : A \rightarrow B$, we set

$$D := \{ f(a, \sigma(a)) \mid a \in A \}.$$

Then the formula

$$\vartheta(x, y) := \exists u \exists v [f(u, v) = y \wedge (x = u \vee f(x, v) \in D)]$$

has the desired properties.

(1) \Rightarrow (3) Fix $n < \omega$ and choose a structure $\mathfrak{M} \in \mathcal{C}_+$ in which the formula χ defines a bijection $f : A \times B \rightarrow C$ with $|A| = n = |B|$. We can define two equivalence relations on C by setting

$$\varphi(x, y) := \exists u \exists v \exists w [f(u, w) = x \wedge f(v, w) = y],$$

and $\psi(x, y) := \exists u \exists v \exists w [f(w, u) = x \wedge f(w, v) = y]$.

(3) \Rightarrow (1) Fix $n < \omega$ and choose a structure $\mathfrak{M} \in \mathcal{C}_+$ in which the formulae φ and ψ define two equivalence relations \sim and \approx with n classes each. We fix a \approx -class A , a \sim -class B , and a set C of representatives of each class of the relation $\sim \cap \approx$. Then the formula

$$\chi(x, y, z) := Ax \wedge By \wedge Cz \wedge z \sim x \wedge z \approx y$$

defines a bijection $A \times B \rightarrow C$. \square

For our official definition, we will use the second of the above variants.

Definition 7.5. (a) Let $k < \omega$. A formula $\varphi(\bar{x}, \bar{y})$ has *pairing rank* at least k in a structure \mathfrak{A} if there exist families $(\bar{a}_i)_{i < k}$ and $(\bar{b}_{ij})_{0 \leq i < j < k}$ such that

$$\mathfrak{A} \models \varphi(\bar{a}_m, \bar{b}_{ij}) \quad \text{iff} \quad m = i \text{ or } m = j.$$

(b) A class \mathcal{C} has *bounded pairing rank* if there exist a number $k < \omega$ such that no FO-formula $\varphi(\bar{x}, y)$ with a single variable y has pairing rank k in any structure $\mathfrak{A} \in \mathcal{C}$.

(c) Let \mathcal{C} be a class of Σ -structures. A *monadic expansion* of \mathcal{C} is a class of the form

$$\mathcal{C}_+ = \{ \langle \mathfrak{A}, \bar{P} \rangle \mid \mathfrak{A} \in \mathcal{C}, \bar{P} \in \wp(A)^m \}, \quad \text{for some } m < \omega.$$

Exercise 7.1. Let τ be an FO-interpretation and \mathcal{C} a class of structures. Show that, if $\tau(\mathcal{C})$ has unbounded pairing rank, then so does \mathcal{C} . \lrcorner

As our notion of a ‘complicated class’ we will use classes that have a monadic expansion with unbounded pairing rank. But before showing that these are indeed quite complicated, we will prove that the following, seemingly much stronger property, is in fact equivalent to having such a monadic expansion.

Definition 7.6. Let \mathcal{C} be a class of Σ -structures. A formula $\varphi(\bar{x}, \bar{y})$ has the *independence property* over \mathcal{C} if, for every $n < \omega$, there exist a structure $\mathfrak{A} \in \mathcal{C}$ containing two families $(\bar{a}_i)_{i < n}$ and $(\bar{b}_w)_{w \in [n]}$ of tuples such that

$$\mathfrak{A} \models \varphi(\bar{a}_i, \bar{b}_w) \quad \text{iff} \quad i \in w.$$

We say that the class \mathcal{C} has the *independence property* if there exists a first-order formula $\varphi(\bar{x}, \bar{y})$ with the independence property over \mathcal{C} . We say that \mathcal{C} has the *monadic independence property* if some monadic expansion of \mathcal{C} has the independence property. \lrcorner

In the definition of the independence property we may assume w.l.o.g. that the tuples \bar{x} and \bar{y} are singletons. We omit the proof, as it is combinatorially quite involved. It can be found in [6].

Proposition 7.7. Let \mathcal{C} be a class of Σ -structures. If \mathcal{C} has the monadic independence property, there exists an FO-formula $\varphi(x, y)$ with the independence property over some monadic expansion of \mathcal{C} , where x and y are single variables.

It remains to give the proof that the monadic independence property is in fact equivalent to the seemingly weaker condition of having unbounded pairing rank.

Proposition 7.8. Let \mathcal{C} be a class of finite Σ -structures. The following conditions are equivalent.

- (1) \mathcal{C} has the monadic independence property.
- (2) \mathcal{C} has a monadic expansion of unbounded pairing rank.

- (3) For every signature Γ , there exists an FO-interpretation mapping some monadic expansion of \mathcal{C} to the class of all finite Γ -structures.

Proof. (2) \Rightarrow (1) is trivial.

(3) \Rightarrow (2) Fix a monadic expansion \mathcal{C}_+ of \mathcal{C} and an FO-interpretation $\tau = \langle \delta, \varphi_E \rangle$ mapping \mathcal{C}_+ to the class of all $\{E\}$ -structures, where E is a binary relation. For every $n < \omega$, then there exists some $\mathfrak{C} \in \mathcal{C}_+$ such that

$$\tau(\mathfrak{C}) \cong \langle A + A^2, E \rangle,$$

with $|A| = n$ and

$$E := \{ \langle a, \langle a, b \rangle \rangle \mid a, b \in A \} \cup \{ \langle a, \langle b, a \rangle \rangle \mid a, b \in A \}.$$

Then the formula φ_E has pairing rank at least n in the structure \mathfrak{C} .

(1) \Rightarrow (3) Let \mathcal{C}_+ be a monadic expansion of \mathcal{C} with the independence property and let $\varphi(x, y)$ be the corresponding formula. By Theorem 7.7, we may assume that x and y are single variables. Fix a finite Γ -structure \mathfrak{M} of size n and set

$$k := |M| + \sum_{R \in \Gamma} |R^{\mathfrak{M}}|.$$

We choose a structure $\mathfrak{C} \in \mathcal{C}_+$ which contains families $(a_i)_{i < k}$ and $(b_w)_{w \subseteq [k]}$ such that

$$\mathfrak{C} \models \varphi(a_i, b_w) \quad \text{iff} \quad i \in w.$$

We construct the desired interpretation τ as follows. Fix a bijection

$$\sigma : [k] \rightarrow M + \sum_{R \in \Gamma} R^{\mathfrak{M}}$$

and set

$$A := \{ a_i \mid \sigma(i) \in M \}.$$

For a relation symbol $R \in \Gamma$ and a number l that is less than the arity of R , we define

$$P_{R,l} := \{ b_{\{i,j\}} \mid \sigma(i) \in M, \sigma(j) \in R^{\mathfrak{M}}, \text{ and } \sigma(i) \text{ is the } l\text{-th component of the tuple } \sigma(j) \}.$$

We can define the i -th incidence relation by

$$\psi_R(\bar{x}) := \exists y \exists \bar{z} \bigwedge_l [x_l \in A \wedge z_l \in P_{R,l} \wedge \varphi(x_l, z_l) \wedge \varphi(y, z_l)].$$

Setting $\delta(x) := Ax$ it follows that the interpretation $\tau := \langle \delta, (\psi_R)_{R \in \Gamma} \rangle$ maps $\langle \mathfrak{C}, A, \bar{P} \rangle$ to \mathfrak{M} . As τ is independent of \mathfrak{M} , it follows that $\tau[\mathcal{D}]$ is the set of all Γ -structures, where

$$\mathcal{D} := \{ \langle \mathfrak{C}, A, \bar{P} \rangle \mid \mathfrak{C} \in \mathcal{C}_+ \}$$

is a monadic expansion of \mathcal{C} . \square

As an immediate consequence of these equivalences, let us note that classes with the monadic independence property are quite complicated.

Corollary 7.9. *If a class \mathcal{C} has the monadic independence property, its crossing-width is unbounded and its MSO-theory is undecidable.*

Proof. Let τ be an FO-interpretation mapping some monadic expansion \mathcal{C}_+ of \mathcal{C} to the class of all Γ -structures. Then $\tau[\mathcal{C}_+]$ has unbounded crossing-width. By Proposition 5.1, it follows that so does \mathcal{C}_+ . As $\text{cwd } \mathcal{C}_+ = \text{cwd } \mathcal{C}$, the first claim follows.

For undecidability, note that there exist signatures Γ such that validity for FO-formulae over Γ is undecidable. Consequently, the FO-theory of $\tau[\mathcal{C}_+]$ is undecidable. Since this theory reduces to the MSO-theory of \mathcal{C} , it follows that the latter theory is also undecidable. \square

Notes

Clique-width has been introduced in [61, 62]. Our definition of crossing-width is a modified version of the so-called partition-width defined in [15, 17].

The version for uncountable structures can also be found there. Most of this chapter is based on these two sources.

Seese's Conjecture was originally stated in [191]. The partial solution in Theorem 5.9 is from [63].

The relation between tree-width and clique-width is due to [57].

The independence property was introduced by Shelah [194]. Proposition 7.7 is from [6]. The relation of the monadic independence property to crossing width was first observed in [20, 21]. It has recently gained renewed attention in the study of first-order model-checking [33, 72].

X Guarded Second-Order Transductions

1 Transductions

IT IS TIME TO TAKE a look at the connection between tree-decompositions and logic. To do so, we introduce a variant of an interpretation that is better-suited as a transformation between classes of structures, in particular, classes of finite structures.

Definition 1.1. Let Σ and Γ be two signatures.

(a) A *transduction* (from Σ to Γ) is a binary relation τ between Σ -structures and Γ -structures. Usually we use functional notation and write $\mathfrak{B} \in \tau(\mathfrak{A})$ instead of $\langle \mathfrak{B}, \mathfrak{A} \rangle \in \tau$. For a class \mathcal{C} of Σ -structure, we set

$$\tau(\mathcal{C}) := \bigcup \{ \tau(\mathfrak{A}) \mid \mathfrak{A} \in \mathcal{C} \}.$$

(b) The *composition* of two transductions τ and σ is the relation

$$\tau \circ \sigma := \{ \langle \mathfrak{C}, \mathfrak{A} \rangle \mid \langle \mathfrak{C}, \mathfrak{B} \rangle \in \tau \text{ and } \langle \mathfrak{B}, \mathfrak{A} \rangle \in \sigma \}.$$

Frequently, transductions construct the output structure from the elements of the input structure. Sometimes the information of how the elements are encoded in the input structure is useful when reasoning about transductions. In such cases we can add this data to the transduction. (Although we will not make much use of the resulting notion in this book.)

Definition 1.2. Let $\mathbb{O} : \text{Set} \rightarrow \text{Set}$ be a functor. An *transduction with origin information of type \mathbb{O}* , or a *\mathbb{O} -transduction* for short, is a set of triples $\langle \mathfrak{A}, \mathfrak{B}, o \rangle$ where $o : \mathbb{O}A \rightarrow B$ is a partial function called the *origin map*.

Examples. (a) An interpretation τ defines an Id-transduction $\langle \mathcal{A}, \tau(\mathfrak{B}), o \rangle$ where $o : \delta^{\mathcal{A}} \rightarrow B$ maps each element of $\delta^{\mathcal{A}}$ to the element of B it represents.

(b) A k -dimensional interpretation τ is a \mathbb{P}_k -transduction where $\mathbb{P}_k A = A^k$.

(c) The copying operation copy_k defines a \mathbb{C}_k -transduction where $\mathbb{C}_k A = [k] \times A$.

(d) The power-set operations \wp and \wp_{fin} are, respectively, \wp -transductions and \wp_{fin} -transductions. J

Lemma 1.3. *Let σ be a \mathbb{O} -transduction and τ a \mathbb{P} -transduction. Then $\tau \circ \sigma$ is a $(\mathbb{P} \circ \mathbb{O})$ -transduction.*

Proof. If $\langle \mathcal{A}, \mathfrak{B}, o \rangle \in \sigma$ and $\langle \mathfrak{B}, \mathbb{C}, p \rangle \in \tau$, then $\langle \mathcal{A}, \mathbb{C}, p \circ \mathbb{P}o \rangle \in \tau \circ \sigma$ where $p \circ \mathbb{P}o : \mathbb{P}\mathbb{O}A \rightarrow C$. \square

Of course, we are particularly interested in transductions that are definable in some logic.

Definition 1.4. (a) An MSO-transduction is a transduction that can be decomposed into a sequence of the following kinds of basic transductions.

(i) Every MSO-interpretation is an MSO-transduction.

(ii) Every copy operation copy_k is an MSO-transduction.

(iii) An m -expansion is the transduction

$$\exp_m(\mathcal{A}) := \{ \langle \mathcal{A}, \bar{P} \rangle \mid P_0, \dots, P_{m-1} \subseteq A \}.$$

(iv) An MSO-filtering is a transduction of the form

$$\{ \langle \mathcal{A}, \mathcal{A} \rangle \mid \mathcal{A} \models \chi \}, \quad \text{for some MSO-formula } \chi.$$

(b) A GSO-transduction is a transduction of the form

$$\tau = \{ \langle \mathfrak{B}, \mathcal{A} \rangle \mid \langle \mathfrak{B}_{\text{in}}, \mathcal{A}_{\text{in}} \rangle \in \sigma \}, \quad \text{for some MSO-transduction } \sigma. \text{ J}$$

Exercise 1.1. Find MSO-transductions with the following properties.

(i) A transduction mapping the class of all finite paths to the class of all circles.

(ii) A transduction mapping the class of all circles to the class of all finite paths.

(iii) A transduction mapping the class of all finite paths to the class of all complete bipartite graphs.

(iv) For every $n < \omega$, a transduction τ_n mapping the class of all finite paths to the class of all finite trees of height n .

(v) A transduction mapping the class of all finite grids to the class of all finite graphs.

Which of these maps are GSO-transductions? J

Every element of a structure \mathfrak{B} obtained via a transduction from some structure \mathcal{A} is represented by some element of \mathcal{A} , but note that, because of copy-operations, every element of \mathcal{A} can represent several elements of \mathfrak{B} .

Definition 1.5. (a) A transduction containing the copy operations

$$\text{copy}_{k_0}, \dots, \text{copy}_{k_{n-1}}$$

is called k -copying, where $k := k_0 \cdot \dots \cdot k_{n-1}$. A 1-copying transduction is usually called *non-copying*.

(b) For each pair $\langle \mathfrak{B}, \mathcal{A} \rangle$ of structures related by a k -copying transduction τ , we can define an *origin map* $o : B \rightarrow [k] \times A$ that maps each element of \mathfrak{B} to the element of \mathcal{A} representing it and to the number of the copy of \mathcal{A} it belongs to. We start by defining o for basic transductions.

(i) For an MSO-interpretation $\langle \delta(x), (\varphi_R(\bar{x}))_R \rangle$ we set

$$o(b) := \langle o, b \rangle, \quad \text{for } b \in B \subseteq A.$$

(ii) For an expansion \exp_m , we similarly define

$$o(b) := \langle o, b \rangle, \quad \text{for } b \in B = A.$$

(iii) For a copying transduction copy_k , we set

$$o(\langle i, b \rangle) := \langle i, b \rangle, \quad \text{for } b \in B = A \text{ and } i < k.$$

(iv) For an MSO-filtering, we again have

$$o(b) := \langle o, b \rangle, \quad \text{for } b \in B = A.$$

Finally, for a composition $\tau_o \circ \dots \circ \tau_{n-1}$ of basic transductions we compose the respective origin maps o_o, \dots, o_{n-1} as follows. The composition of

$$o : C \rightarrow [l] \times B \quad \text{and} \quad o' : B \rightarrow [k] \times A$$

is the map $oo' : C \rightarrow [lk] \times A$ defined by

$$oo'(c) := \langle jk + i, a \rangle, \quad \text{for } o'(c) = \langle j, b \rangle \text{ and } o(b) = \langle i, a \rangle.$$

Let us give a few examples of what we can do with transductions. We start by showing that we can use transductions to take quotients by definable equivalence relations.

Lemma 1.6. *For every MSO-formula $\varepsilon(x, y)$, there exists a transduction τ mapping a structure \mathcal{A} to the quotient \mathcal{A}/\sim , where \sim is the equivalence relation that is generated by the relation defined by ε .*

Proof. Note that \sim is MSO-definable as the transitive closure of

$$x = y \vee \varepsilon(x, y) \vee \varepsilon(y, x).$$

We use one parameter P that contains representatives for each equivalence class. Thus, we can set $\tau := \tau_o \circ \rho \circ \exp_1$ where ρ is the filtering by the formula stating that every \sim -class contains exactly one element in P and τ_o is the interpretation defined by the formulae

$$\delta(x) := Px \quad \text{and} \quad \varphi_R(\bar{x}) := \exists \bar{y} \left[R\bar{y} \wedge \bigwedge_i y_i \sim x_i \right]. \quad \square$$

Lemma 1.7. *There exists a GSO-transduction mapping a class of hypergraphs to the class of their minors.*

Proof. Each minor of a graph can be described by three parameters: the set R of removed vertices, the set D of deleted edges, and the set C of contracted edges. We can therefore construct all minors by a transduction of the form $\tau = \tau_o \circ \sigma_o \circ \exp_3$, where \exp_3 guesses R , D , and C , τ_o is the transduction forming the quotient by the equivalence relation generated by C (cf. Lemma 1.6), and σ_o is the interpretation removing the vertices and edges in R and D , which can be defined by

$$\delta_V(x) := \neg Rx, \quad \delta_E(x) := \neg Dx, \quad \text{and} \quad \varphi_{\text{in}}(x, y) := \text{in}(x, y). \quad \square$$

The next two examples will be used in Section 4 below.

Lemma 1.8. *For every $n < \omega$, there exists an MSO-transduction mapping the class of all finite paths to the class of all finite trees of height at most n .*

Proof. We can encode a tree T of height n with m vertices as a finite word w of length m over the alphabet $[n]$ as follows (see Figure 1). Let $v_o <_{\text{lex}} \dots <_{\text{lex}} v_{m-1}$ be the enumeration of the vertices of T in lexicographic order, and let l_i be the level of v_i . We encode T by the word $w := l_o \dots l_{m-1}$. A transduction can recover T from w as follows. Each position in w corresponds to a vertex. The predecessor of the i -th vertex v is the maximal vertex to the left of v whose label is less than l_i . Clearly this predecessor relation is definable in monadic second-order logic. Consequently, we obtain a transduction of the form $\chi \circ \tau \circ \exp_n$ where \exp_n guesses the labels l_i , χ is a filtering checking that the graph in question is indeed a tree, and τ is the MSO-interpretation described above. (The reason we need χ is that not every colouring of a path constitutes a valid encoding of a tree. For instance, we could have several vertices with the colour of the root.) \square

Recall that the $m \times n$ grid is the undirected graph $\mathfrak{G} = \langle V, E \rangle$ with vertices $V = [m] \times [n]$ and edge relation

$$E = \{ \langle \langle i, k \rangle, \langle j, l \rangle \rangle \mid |i - j| + |k - l| = 1 \}.$$

Its directed variant is $\langle V, E_o, E_1 \rangle$ where

$$E_o := \{ \langle \langle i, k \rangle, \langle i+1, k \rangle \rangle \mid i < m-1, k < n \},$$




Figure 1: Transforming a path into a tree of bounded height

and $E_1 := \{ \langle \langle i, k \rangle, \langle i, k+1 \rangle \rangle \mid i < m, k < n-1 \}$.

Lemma 1.9. *There exists an MSO-transduction mapping the class of all finite, undirected grids to the class of all finite, directed grids.*

Proof. Given an undirected grid $\mathfrak{G} = \langle V, E \rangle$, we guess six parameters $P_0, P_1, P_2, Q_0, Q_1, Q_2 \subseteq V$ such that

$$P_m := \{ \langle i, k \rangle \mid i \equiv m \pmod{3} \},$$

$$\text{and } Q_m := \{ \langle i, k \rangle \mid k \equiv m \pmod{3} \}.$$

Then we can define

$$E_0 = \{ \langle u, v \rangle \in E \mid u \in P_i \text{ and } v \in P_j \text{ for some } j \equiv i+1 \pmod{3} \},$$

$$\text{and } E_1 = \{ \langle u, v \rangle \in E \mid u \in Q_i \text{ and } v \in Q_j \text{ for some } j \equiv i+1 \pmod{3} \}.$$

To do so, we need a formula checking that the parameters P_m and Q_m are correctly chosen. For such a formula it is sufficient to state that the labelling of each vertex is consistent with those of its neighbours and that the labelling of every cycle of length 4 is consistent (where ‘consistent’ means that the labelling is one of those appearing in a correctly labelled grid). \square

Lemma 1.10. *There exists a GSO-transduction mapping the class of all finite grids to the class of all finite Σ -structures.*

Proof. By Lemma 1.9, it is sufficient to find a transduction mapping directed grids to all graphs. Let \mathfrak{A} be a Σ -structure with

$$\mathfrak{A}_{\text{in}} = \langle A \cup E, (P_R)_R, \text{in}_0, \dots, \text{in}_{r-1} \rangle.$$




Figure 2: Transforming a grid into an arbitrary graph

Fix enumerations a_0, \dots, a_{m-1} of A and e_0, \dots, e_{n-1} of E . We encode \mathfrak{A}_{in} in the directed $m \times n$ grid using the following parameters (see Figure 2).

$$\begin{aligned} A' &:= [m] \times \{0\}, & P'_R &:= \{ \langle 0, k \rangle \mid e_k \in P_R \}, \\ E' &:= \{0\} \times [n], & I'_l &:= \{ \langle i, k \rangle \mid (a_i, e_k) \in \text{in}_l \}. \end{aligned}$$

Then \mathfrak{A}_{in} can be recovered from \mathfrak{G} by an MSO-transduction using these sets as parameters:

$$\delta_A(x) := A'x, \quad \delta_E(x) := E'x, \quad \varphi_{P_R}(x) := P'_R x,$$

$$\begin{aligned} \varphi_{\text{in}_l}(x, y) &:= A'x \wedge E'y \\ &\quad \wedge \exists z [I'_l z \wedge 'z \text{ is the intersection of the row of } x \\ &\quad \text{and the column of } y']. \end{aligned} \quad \square$$

After these examples we turn to the further development of the theory of transductions. We start with a normal form for transductions that frequently comes in handy: we need only one operation of each kind and in a standard order.

Proposition 1.11. *Every MSO-transduction τ can be written in the form*

$$\tau = \tau_0 \circ \rho \circ \text{copy}_k \circ \text{exp}_m,$$

where τ_0 is an MSO-interpretation and ρ an MSO-filtering.

Proof. The proof consists of the following (rather long) sequence of commutation relations.

(a) We start by showing that two operations of the same kind can always be combined into a single one. We can combine two MSO-filterings by just taking the conjunction of the two formulae. Furthermore, for MSO-interpretations note that $\tau \circ \tau' = \tau''$ where τ'' is obtained from τ by replacing every formula ψ by $\psi^{\tau'}$. Finally, we have

$$\exp_m \circ \exp_n = \exp_{mn} \quad \text{and} \quad \text{copy}_k \circ \text{copy}_l = \sigma \circ \text{copy}_{kl},$$

where σ is a simple quantifier-free interpretation that defines the predicates introduced by copy_k and copy_l in terms of those defined by copy_{kl} . More precisely, let \tilde{H}, I, H', I' , and \tilde{H}'', I'' be the predicates introduced by, respectively, $\text{copy}_{kl}, \text{copy}_k$, and copy_l . Then

$$\begin{aligned} H'_i &= \bigcup_{j < l} H_{il+j}, & I' &= I \cap \bigcup_{j < l} (H''_j \times H''_j) \\ H''_j &= \bigcup_{i < k} H_{il+j}, & I'' &= I \cap \bigcup_{i < k} (H'_i \times H'_i), \end{aligned}$$

which are clearly all definable.

(b) If ρ is a filtering with the formula χ and τ is an MSO-interpretations, then $\rho \circ \tau = \tau \circ \rho'$, where ρ' is the filtering with the formula χ^{τ} .

(c) Concerning the copying operation, we have already shown in Lemma I.4.20 that it commutes with interpretations. For a filtering ρ , note that

$$\text{copy}_k \circ \rho = \rho' \circ \text{copy}_k,$$

where ρ' applies the formula from ρ to just one of the copies, i.e., if χ is the formula used by ρ , we take for ρ' the relativisation $\chi^{(H_o)}$ of χ to the set H_o .

(d) Concerning expansions, we have

$$\begin{aligned} \exp_m \circ \tau &= \tau \circ \exp_m, & \text{for every MSO-interpretation } \tau, \\ \exp_m \circ \rho &= \rho \circ \exp_m, & \text{for every MSO-filtering } \rho, \\ \exp_m \circ \text{copy}_k &= \tau' \circ \text{copy}_k \circ \exp_{km}, \end{aligned}$$

where τ' is the quantifier-free interpretation that replaces the parameters P_o, \dots, P_{k-m-1} by P'_o, \dots, P'_{m-1} where

$$P'_i := \bigcup_{j < k} [P_{ik+j} \cap H_j]. \quad \square$$

Remark. Let us note two consequences of this result that will frequently be useful. Suppose that τ is a k -copying transduction in this normal form.

(a) We see that the elements of $\tau(\mathfrak{A})$ are of the form $\langle i, a \rangle$ where $a \in A$ and $i < k$ denotes the copy the element belongs to.

(b) Each formula $\varphi_R(\vec{x})$ from τ defining some relation R can therefore be split into several formulae $\varphi_{\bar{i}}(\vec{x})$ that define R between elements from the copies indicated by \bar{i} . Formally, we can define $\varphi_{\bar{i}}$ as the formula obtained from φ_R by replacing every subformula of the form $H_{ij}x_j$ by true, and all subformulae H_{lx_j} with $l \neq i_j$ by false. Then we have

$$\varphi_R(\vec{x}) \equiv \bigvee_{i_0, \dots, i_{n-1} < k} [H_{i_0}x_0 \wedge \dots \wedge H_{i_{n-1}}x_{n-1} \wedge \varphi_{i_0, \dots, i_{n-1}}(\vec{x})],$$

where n is the arity of R and H_o, \dots, H_{k-1} are the predicates for the various copies. $\quad \dashv$

Lemma I.12. *If τ_o and τ_1 are MSO-transductions, then so is $\tau_o \cup \tau_1$.*

Proof. The rough idea is to use an additional parameter to decide which of the two transductions to apply. Suppose that $\tau_i = \tau'_i \circ \rho_i \circ \text{copy}_{k_i} \circ \exp_{m_i}$. Then

$$\tau_o \cup \tau_1 = \sigma \circ v \circ \text{copy}_k \circ \exp_{m+1}$$

where $k := \max\{k_o, k_1\}$, $m := \max\{m_o, m_1\}$, the filtering v is given by the formula

$$[\exists z Qz \wedge \chi_o] \vee [\neg \exists x Qx \wedge \chi_1],$$

where χ_i is the formula for ρ_i and Q is the additional parameter, and the interpretation σ is $\langle \delta(x), (\varphi_R(\tilde{x}))_R \rangle$ with

$$\begin{aligned}\delta(x) &:= [\exists z Qz \wedge \delta_o(x) \wedge \bigwedge_{i < k_o} H_i x] \\ &\vee [\neg \exists z Qz \wedge \delta_I(x) \wedge \bigwedge_{i < k_I} H_i x], \\ \varphi_R(\tilde{x}) &:= [\exists z Qz \wedge \varphi_R^o(\tilde{x})] \vee [\neg \exists z Qz \wedge \varphi_R^I(\tilde{x})],\end{aligned}$$

where $\tau'_i = \langle \delta_I(x), (\varphi_R^I(\tilde{x}))_R \rangle$ and H_i are the predicates added by copy_k . \square

Exercise 1.2. Given two MSO-transductions σ and τ , construct an MSO-transduction ρ such that

$$\rho(\mathfrak{A} \oplus \mathfrak{B}) = \sigma(\mathfrak{A}) \oplus \tau(\mathfrak{B}), \quad \text{for all structures } \mathfrak{A} \text{ and } \mathfrak{B}.$$

The usefulness of transductions stems from the fact that, similarly to the interpretations they generalise, they provide reductions between theories. The corresponding results for transductions are made more complicated by the fact that they are non-deterministic. Let us introduce a bit of notation to make the rôle of parameters more explicit.

Definition 1.13. Let τ be a transduction of the form $\tau = \tau_o \circ \exp_m$ where τ_o is parameterless. For structures \mathfrak{A} and \mathfrak{B} and an parameters $P_o, \dots, P_{m-1} \subseteq A$ We write

$$\mathfrak{B} = \tau(\mathfrak{A}/\tilde{P}) \quad \text{iff} \quad \mathfrak{B} = \tau_o(\mathfrak{A}, \tilde{P}),$$

that is, if \mathfrak{B} is the structure obtained via τ from \mathfrak{A} when using the parameters \tilde{P} .

When considering formulae with free variables, we need a version of this notation that also takes the parameters of the formula into account. We write

$$\mathfrak{B}, \tilde{b}\tilde{Q} = \tau_i(\mathfrak{A}, \tilde{b}'\tilde{Q}'/\tilde{P})$$

if \tilde{b} is encoded by \tilde{b}' and \tilde{I} , and \tilde{Q} is encoded by \tilde{Q}' . The precise definition is as follows.

Definition 1.14. Let τ be a k -copying MSO-transduction with origin map o , and suppose that $\mathfrak{B} = \tau(\mathfrak{A}/\tilde{P})$. For tuples $\tilde{b} \in B^s$, $\tilde{b}' \in A^s$, and $\tilde{I} \in [k]^s$ and sets $Q_o, \dots, Q_{n-1} \subseteq B$, and $Q'_{ij} \subseteq A$ (for $i < k$, $j < n$), we write

$$\mathfrak{B}, \tilde{b}\tilde{Q} = \tau_i(\mathfrak{A}, \tilde{b}'\tilde{Q}'/\tilde{P})$$

if, and only if,

$$o(b_i) = \langle l_i, b'_i \rangle \quad \text{and} \quad Q_{ij} = \{ a \in A \mid b \in Q_j, o(b) = \langle i, a \rangle \}.$$

Using this notation we can state the desired backwards-translation result for transductions as follows.

Proposition 1.15. Let τ be a k -copying MSO-transduction. For every MSO-formula $\varphi(\tilde{x}, \tilde{Y})$ with $|\tilde{x}| = s$ and every tuple $\tilde{I} \in [k]^s$, there exist an MSO-formula $\varphi_i^r(\tilde{x}', \tilde{Y}'; \tilde{Z})$ such that

$$\mathfrak{B}, \tilde{b}\tilde{Q} = \tau_i(\mathfrak{A}, \tilde{b}'\tilde{Q}'/\tilde{P})$$

implies

$$\mathfrak{B} \models \varphi(\tilde{b}, \tilde{Q}) \quad \text{iff} \quad \mathfrak{A} \models \varphi_i^r(\tilde{b}', \tilde{Q}'; \tilde{P}).$$

Proof. It is clearly sufficient to prove the claim for basic transductions.

(i) For interpretations, the claim follows immediately from Proposition I.4.5.

(ii) Let Q_o, \dots, Q_{m-1} be the predicates for the parameters. We obtain the formula $\varphi_{o \dots o}^r(\tilde{x}, \tilde{Y}; \tilde{Z})$ by replacing in the given formula $\varphi(\tilde{x}, \tilde{Y})$ every occurrence of Q_i by the variable Z_i .

(iv) A filtering does not modify the structure. Hence, we can simply set $\varphi_{o \dots o}^r(\tilde{x}, \tilde{Y}) := \varphi(\tilde{x}, \tilde{Y})$.

(III) It remains to consider copy_k -operations. We define the desired formula $\varphi_i^\tau(\bar{x}, \bar{Y})$ by induction on φ .

$$\begin{aligned} (R\bar{x})_i^\tau &:= \begin{cases} R\bar{x} & \text{if } l_0 = \dots = l_{s-1}, \\ \text{false} & \text{otherwise,} \end{cases} & (Ixy)_i^\tau &:= x = y, \\ (x = y)_{ij}^\tau &:= \begin{cases} x = y & \text{if } i = j, \\ \text{false} & \text{otherwise,} \end{cases} & (Yx)_i^\tau &:= Y_l x, \\ (H_i x)_j^\tau &:= \begin{cases} \text{true} & \text{if } i = j, \\ \text{false} & \text{otherwise,} \end{cases} & (\neg\varphi)_i^\tau &:= \neg\varphi_i^\tau, \\ (\varphi \vee \psi)_i^\tau &:= \varphi_i^\tau \vee \psi_i^\tau, \\ (\exists x' \varphi(\bar{x}x', \bar{Y}))_i^\tau &:= \bigvee_{i < k} \exists x' \varphi_{li}^\tau(\bar{x}x', \bar{Y}), \\ (\exists Y' \varphi(\bar{x}, \bar{Y}Y'))_i^\tau &:= \exists Y'_0 \dots Y'_{k-1} \varphi_{li}^\tau(\bar{x}, \bar{Y}Y'_0 \dots Y'_{k-1}). \end{aligned} \quad \square$$

The above statement is rather technical. The following simpler version is usually sufficient.

Corollary 1.16. *Let τ be an MSO-transduction. For every MSO-formula φ , there exists an MSO-formula $\varphi^\tau(\bar{Z})$ such that*

$$\mathfrak{A} \models \varphi^\tau(\bar{P}) \quad \text{iff} \quad \tau(\mathfrak{A}/\bar{P}) \text{ is defined and } \tau(\mathfrak{A}/\bar{P}) \models \varphi,$$

for every structure \mathfrak{A} and all parameters \bar{P} .

Proof. Suppose that $\tau = \tau_o \circ \rho \circ \text{copy}_k \circ \text{exp}_m$ is in standard form and let χ be the formula used by ρ . We use Proposition 1.15 below to construct the formulae $\varphi_{\langle \rangle}^\tau(\bar{Z})$ and $\chi_{\langle \rangle}^{\text{copy}_k \circ \text{exp}_m}(\bar{Z})$ for, respectively, φ and χ . Then the formula

$$\varphi^\tau := \exists \bar{Z} [\chi_{\langle \rangle}^{\text{copy}_k \circ \text{exp}_m}(\bar{Z}) \wedge \varphi_{\langle \rangle}^\tau(\bar{Z})]$$

has the desired properties. \square

2 Tree-Decompositions

There is a close relationship between transductions and tree-decompositions which we will investigate in this section.

Proposition 2.1. *Let Σ be a signature and $k < \omega$. There exists an GSO-transduction τ_k that associates with each undirected tree \mathfrak{T} the class of all Σ -structures that have a tree-decomposition of width at most k with index tree \mathfrak{T} .*

Proof. Suppose that \mathfrak{A} is a Σ -structure which has a tree-decomposition $(U_t)_{t \in T}$ of width k . We prove that \mathfrak{A} can be defined from a suitable colouring of T where the number of colours depends only on Σ and k .

Let $\mathfrak{C}_0, \dots, \mathfrak{C}_{m-1}$ be an enumeration of all Σ -structures whose domain is a subset of $[k]$. For each $t \in T$, let \mathfrak{U}_t be the substructure of \mathfrak{A} induced by U_t . It follows that, for every $t \in T$, we can find some index $\lambda(t)$ such that $\mathfrak{U}_t \cong \mathfrak{C}_{\lambda(t)}$. Let $\pi_t : \mathfrak{U}_t \rightarrow \mathfrak{C}_{\lambda(t)}$ be the corresponding isomorphism.

Furthermore, we associate with each (directed) edge $\langle s, t \rangle$ of T the binary relation

$$R(s, t) := \{ \langle \pi_s(a), \pi_t(a) \rangle \mid a \in U_s \cap U_t \} \subseteq [k] \times [k].$$

Then we can recover \mathfrak{A} from T with the help of the vertex colouring λ and the edge colouring R . We form the disjoint union of all structures $(\mathfrak{C}_{\lambda(t)})_{\text{in}}$, for $t \in T$, and, for every edge $\langle s, t \rangle$ of T , we identify two elements $i \in C_{\lambda(s)}$ and $j \in C_{\lambda(t)}$ if $\langle i, j \rangle \in R(s, t)$. This can be performed by an n -copying MSO-transduction where n is the maximal size of the structures $(\mathfrak{C}_i)_{\text{in}}$, $i < m$. \square

Corollary 2.2. *Let $k < \omega$. The following classes have a decidable CGSO-theory.*

- (a) *The class of all successor trees.*
- (b) *The class of all countable successor trees.*
- (c) *The class of all finite successor trees.*
- (d) *The class of all Σ -structures of tree-width at most k .*
- (e) *The class of all countable Σ -structures of tree-width at most k .*

(f) The class of all finite Σ -structures of tree-width at most k .

The following classes have a decidable CMSO-theory.

(g) The class of all Σ -structures of crossing-width at most k .

(h) The class of all countable Σ -structures of crossing-width at most k .

(i) The class of all finite Σ -structures of crossing-width at most k .

Proof. (a) follows by Theorem V.4.7.

(b) There is a GSO-transduction mapping the infinite binary tree $\mathfrak{T}_{\text{bin}}$ to the infinite tree \mathfrak{T}_ω of arity ω . Furthermore, there exists a GSO-transduction mapping \mathfrak{T}_ω to the class of all countable trees. Hence, the claim follows by Corollary V.3.9.

(c) There exists a GSO-transduction mapping the infinite binary $\mathfrak{T}_{\text{bin}}$ to the class of all finite trees. Hence, the claim follows by Corollary V.3.9.

(d), (e), (f) follow by (a), (b), (c) and Proposition 2.1.

(g), (h), (i) follow by (a), (b), (c) and Corollary IX.5.6. \square

Thus, classes of bounded tree-width can be obtained via a transduction from a suitable class of trees. Conversely, every class obtained from trees in that way has bounded tree-width. The proof rests on the following theorem.

Theorem 2.3. *Let τ be a GSO-transduction. If \mathfrak{A} is a Σ -structure with a tree-decomposition $(U_t)_{t \in T}$ of width k , then $\tau(\mathfrak{A})$ has a tree-decomposition $(U'_t)_{t \in T}$ with the same index tree T and a width that depends only on k and τ .*

Proof. Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{A} of width k . By Proposition IX.4.6, the incidence structure \mathfrak{A}_{in} has a partition decomposition whose width bounded by some function of k and Σ and whose index tree T_+ is obtained from T by adding a bounded number of leaves to every vertex. Consequently, we can use Proposition IX.5.1 to obtain a partition decomposition of $\tau(\mathfrak{A})_{\text{in}}$ whose index tree S_+ is a prefix of T_+ and whose width is bounded by a constant depending only on k and τ . We can use Proposition IX.4.6 again to translate this decomposition into a tree-decomposition of $\tau(\mathfrak{A})$ whose index tree S is a prefix of T and whose width depends only on k and τ . Finally, duplicating some components we can turn this decomposition into one with index tree T and with the same width. \square

Corollary 2.4. *Let τ be a GSO-transduction. There exists a constant k such that, if \mathfrak{T} is a tree, then every structure $\mathfrak{A} \in \tau(\mathfrak{T})$ has a tree-decomposition of width k with index tree \mathfrak{T} .*

Corollary 2.5. *A class of structures has bounded tree-width if, and only if, it is GSO-interpretable in a class of coloured trees.*

We have shown that there exists a transduction mapping each tree-decomposition to the structure it denotes. The converse statement also holds, but it is much more involved.

Theorem 2.6 (Bojańczyk, Pilipczuk). *For every finite relational signature Σ and every constant $k < \omega$, there exists an MSO-transduction τ_k mapping every finite Σ -structure of tree-width at most k to the trees underlying some of its tree-decompositions of width k .*

We defer the proof to Section 5.

3 Trees of Bounded Height

Let us take a look at transductions between classes of trees of bounded height. We start by presenting several equivalences between such classes and related ones, such as certain classes of structures with nested equivalence relations.

Definition 3.1. Let $n < \omega$.

(a) An n -equivalence structure takes the form $\mathfrak{A} = \langle A, E_0, \dots, E_{n-1} \rangle$ where $E_0 \subseteq \dots \subseteq E_{n-1}$ are equivalence relations on A . We denote the class of all n -equivalence structures by \mathcal{E}_n . Given an n -equivalence structure \mathfrak{A} , we denote the E_i -class of an element $a \in A$ by $[a]_i$.

(b) A tree \mathfrak{T} has *uniform height* n if it has no infinite branches and every leaf is on level n (i.e., has distance exactly n from the root). We denote the class of all order-trees of uniform height n by \mathcal{T}_n .

(c) The n -equivalence structure \mathfrak{A} associated with an order-tree $\mathfrak{T} = \langle T, \leq \rangle$ of uniform height $n + 1$ is defined as follows. Let $p : T \rightarrow T$ be the function

mapping every vertex of \mathfrak{T} to its predecessor. (This leaves p undefined for the root of \mathfrak{T} .) The domain of \mathfrak{A} is the set of leaves of \mathfrak{T} . Two such leaves u and v are E_i -equivalent if $p^{i+1}(u) = p^{i+1}(v)$. \square

Lemma 3.2. *Let $n < \omega$.*

- (a) *Every tree \mathfrak{T} of uniform height $n + 1$ has a unique associated n -equivalence structure \mathfrak{A} .*
- (b) *Every n -equivalence structure \mathfrak{A} is associated with a unique tree \mathfrak{T} of uniform height $n + 1$.*

Proof. (a) follows immediately from the definition.

(b) Let $\mathfrak{A} = \langle A, E_0, \dots, E_{n-1} \rangle$ be an n -equivalence structure. We define the corresponding tree $\mathfrak{T} = \langle T, \leq \rangle$ as follows. The set of vertices is

$$T := A/E_{-1} \cup A/E_0 \cup \dots \cup A/E_{n-1} \cup A/E_n,$$

where $E_{-1} := \{ \langle a, a \rangle \mid a \in A \}$ is the identity relation and $E_n := A \times A$ the trivial equivalence relation. The tree order is given by

$$[a]_i \leq [b]_k \quad : \text{iff} \quad [b]_k \subseteq [a]_i. \quad \square$$

Theorem 3.3. *Let $n < \omega$. The following transductions are FO-interpretations.*

- (a) *The transduction τ mapping an order-tree to the corresponding undirected tree.*
- (b) *The transduction τ mapping a successor-tree of height at most n to the corresponding order-tree.*
- (c) *The transduction τ mapping a successor-tree of uniform height $n + 1$ to the corresponding n -equivalence structure.*

The following transductions are FO-transductions.

- (d) *The transduction τ mapping an undirected tree \mathfrak{T} with diameter at most n to all successor trees whose underlying undirected tree is \mathfrak{T} .*
- (e) *The transduction τ mapping an order-tree of height at most n to its incidence structure.*

- (f) *The transduction τ mapping an n -equivalence structure to the corresponding successor-tree.*

Furthermore, the corresponding origin maps in (a), (b), (d), and (e) are the identity, while the origin maps in (c) and (f) induce mutually inverse bijections between the leaves of the tree and the elements of the equivalence structure.

Proof. (a) We can define the undirected edge relation by

$$\varphi(x, y) := [x < y \wedge \neg \exists z(x < z < y)] \vee [y < x \wedge \neg \exists z(y < z < x)].$$

(b) We can define a formula $\psi_k(x, y)$ expressing the existence of a path of length k by

$$\psi_0(x, y) := x = y \quad \text{and} \quad \psi_{k+1}(x, y) := \exists z[Sxz \wedge \psi_k(z, y)],$$

where S is the successor relation. Consequently, we can define the order relation by

$$\varphi(x, y) := \bigvee_{k \leq n} \psi_k(x, y).$$

(c) Clearly, both the fact that a vertex v is a leaf and the fact that $p^{i+1}(u) = p^{i+1}(v)$ can be expressed in first-order logic.

(d) We can use a transduction of the form $\tau = \tau_0 \circ \rho \circ \exp_I$, where (I) we guess one parameter P ; (II) we check with ρ that P contains exactly one element r ; and then (III) we define the new edge relation in τ_0 by stating that the two given vertices x and y are connected by an edge and that the path from r to y contains the vertex x . (The latter can be done as in (b) since the length of such paths is bounded.)

(e) We define an interpretation mapping $\text{copy}_{n+2}(\mathfrak{T})$ to the incidence structure of \mathfrak{T} . Let H_0, \dots, H_{n+1} be the predicates for the various copies and let $p : T \rightarrow T$ be the (partial) function mapping a vertex to its predecessor. We encode each vertex v by its copy in H_{n+1} and each edge $\langle p^k(v), v \rangle$ by the copy of v in H_k . Hence, we can use the formulae

$$\delta_V(x) := H_{n+1}x \quad \text{and} \quad \delta_E(x) := \bigvee_{k \leq n} [H_kx \wedge "p^k(x) \text{ is defined}"]$$

for the domains, and the formula

$$\varphi_{\text{in}}(x, y) := H_{n+1}x \wedge \bigvee_{k \leq n} [H_k y \wedge x = p^k(y)]$$

for the incidence relation.

(f) We use a transduction of the form $\tau = \tau_o \circ \rho \circ \text{copy}_{n+2} \circ \text{exp}_n$, where (i) we guess sets P_o, \dots, P_n ; (ii) we check using ρ that P_n is a singleton and that P_i , for $i < n$, contains exactly one element of each E_i -class; and (iii) we then define the tree as follows. Let H_o, \dots, H_{n+1} be the predicates from the copying operation. The leaves are encoded by the elements in H_o and the internal vertices at level $i < n + 1$ by the elements in $H_{n+1-i} \cap P_{n-i}$. This leads to the formulae

$$\begin{aligned} \delta(x) &:= H_o x \vee \bigvee_{i \leq n} [H_{i+1}x \wedge P_i x], \\ \varphi_S(x, y) &:= \bigvee_{i < n} [H_{i+1}x \wedge H_i y \wedge E_i x y] \vee [H_{n+1}x \wedge H_n y]. \end{aligned} \quad \square$$

We have shown in Theorem V.1.13 that, over trees, we can reduce every MSO-formula to an FO-formula by adding a colouring. For trees of bounded height, we can do without this colouring.

Theorem 3.4. *Let $n < \omega$.*

- (a) *Over the class \mathcal{U}_n of all structures with n unary predicates, every GSO-formula $\varphi(\tilde{X}, \tilde{x})$ is equivalent to an FO-formula.*
- (b) *Over the class \mathcal{E}_n , every MSO-formula $\varphi(\tilde{X}, \tilde{x})$ is equivalent to an FO-formula.*
- (c) *Over the class \mathcal{T}_n , every GSO-formula $\varphi(\tilde{X}, \tilde{x})$ is equivalent to an FO-formula.*
- (d) *Over the class $\mathcal{C}_{n,k}$ of all Σ -structures \mathcal{A} with $\text{twd}_n \mathcal{A} < k$, every GSO-formula $\varphi(\tilde{X}, \tilde{x})$ is equivalent to an FO-formula.*

Proof. Ultimately all five of the proofs below are based on the fact that, over the empty signature, all MSO can do is to count up to some constant depending on the quantifier-rank. We start with the simplest case and work our way up to the most complicated one.

(a) has already been proved in Proposition I.4.8.

(b) Again we get rid of the parameters by including them in the structures. Hence, we work with structures of the form $\mathcal{A} = \langle A, E_o, \dots, E_{n-1}, \bar{P}, \bar{c} \rangle$. We prove by induction on n that there exists some function $f_n : \omega \rightarrow \omega$ such that, for structures \mathcal{A} and \mathcal{B} with n equivalence relations,

$$\mathcal{A} \equiv_{\text{FO}}^{f_n(m)} \mathcal{B} \quad \text{implies} \quad \mathcal{A} \equiv_{\text{MSO}}^m \mathcal{B}, \quad \text{for every } m < \omega.$$

For $n = o$, the claim follows by (a). Hence, suppose that $n > o$. Note that we can write each structure \mathcal{A} as a disjoint union of all E_{n-1} -classes

$$\mathcal{A} \cong \bigoplus_{[a]_{n-1} \in A/E_{n-1}} \mathcal{A}|_{[a]_{n-1}}.$$

We can replace this disjoint union by a generalised sum plus a quantifier-free interpretation that deletes the unneeded relations added by the generalised sum.

$$\mathcal{A} \cong \sigma \left(\sum_{[a]_{n-1} \in A/E_{n-1}} \mathcal{A}|_{[a]_{n-1}} \right).$$

By Theorem I.4.24, it follows that, for every MSO-formula φ of quantifier-rank m , there exists some MSO-formula ψ of some quantifier-rank $g(m)$ such that

$$\mathcal{A} \models \varphi \quad \text{iff} \quad \langle A/E_{n-1}, \bar{Q} \rangle \models \psi,$$

where the predicates \bar{Q} encode the MSO_m -theory of the corresponding E_{n-1} -class.

$$\mathcal{Q}_\theta := \{ [a]_{n-1} \in A/E_{n-1} \mid \text{Th}_{\text{MSO}}^m(\mathcal{A}|_{[a]_{n-1}}) = \theta \}.$$

By inductive hypothesis, there exists a first-order interpretation τ of quantifier-rank $f_{n-1}(m)$ mapping \mathcal{A} to $\langle A/E_{n-1}, \bar{Q} \rangle$. It follows that

$$\begin{aligned} & \mathcal{A} \equiv_{\text{FO}}^{f_{n-1}(m)+g(m)} \mathcal{B} \\ \Rightarrow & \langle \mathcal{A}, \bar{Q} \rangle \equiv_{\text{FO}}^{g(m)} \langle \mathcal{B}, \bar{Q}' \rangle \\ \Rightarrow & \mathcal{A} \equiv_{\text{MSO}}^m \mathcal{B}. \end{aligned}$$

(c) As order-trees of height at most n are n^2 -sparse, it follows by Theorem VII.3.6 that we can translate every GSO-formula into an MSO-formula. Hence, the claim follows by Proposition V.1.17.

(d) We work with tree-decompositions of a special form. Let us call a tree-decomposition $(U_t)_{t \in T}$ *connected* if, for every edge $\langle s, t \rangle$, the set

$$\bigcup_{u \in S} U_u \setminus U_s \quad \text{is connected in the Gaifman graph,}$$

where S is the component of $T - \langle s, t \rangle$ containing the vertex t . Note that we can transform every tree-decomposition into a connected one without increasing the width of the decomposition or the height of its index tree: if the subtree S rooted at some vertex t does not satisfy the above condition, i.e., the set $\bigcup_{u \in S} U_u \setminus U_s$ has several connected components C_0, \dots, C_{k-1} , we can map k copies of this subtree whose components are of the form $U_u \cap C_i$.

We will prove by induction on the height n that, for every MSO_m -theory θ and every $k < \omega$, there exists an FO-formula $\psi_{k\theta}^n(\bar{x})$ such that

$$\begin{aligned} \mathfrak{A} \models \psi_{k\theta}^n(\bar{a}) \quad \text{iff} \quad & \text{Th}_{\text{MSO}}^m(\mathfrak{A}) = \theta \text{ and } \mathfrak{A} \text{ has a connected} \\ & \text{tree-decomposition of height at most } n \text{ and} \\ & \text{width at most } k \text{ such that the root component} \\ & \text{contains the tuple } \bar{a}. \end{aligned}$$

Then it follows that

$$\mathfrak{A} \equiv_{\text{FO}}^r \mathfrak{B} \quad \text{implies} \quad \mathfrak{A} \equiv_{\text{MSO}}^m \mathfrak{B},$$

where r is the quantifier-rank of the formulae $\psi_{k\theta}^n$.

Hence, it remains to construct the formulae $\psi_{k\theta}^n$. For $n = 0$, we have to say that the structure has at most k elements and that it is isomorphic to one of the finitely many structure of that size whose MSO_m -theory is equal to θ . This can be done in first-order logic.

For the inductive step, suppose that $n > 0$. To express the existence of a tree-decomposition, we have to say that there exist k elements \bar{c} such that $\bar{a} \subseteq \bar{c}$ and every connected component of $A \setminus \bar{c}$ has a tree-decomposition of

height at most $n - 1$. We have shown in Proposition VIII.2.9 that no graph of n -height tree-width at most k can contain a path of length $(k + 1)^n$. Thus the length of paths in the Gaifman graph of \mathfrak{A} is bounded and there exists an FO-formula expressing that there exists a path from x to y . Modifying this formula we obtain a FO-formula $\chi(x, y, \bar{z})$ stating that x and y belong to the same connected component of $A \setminus \bar{z}$. Consequently, the following formula checks that \mathfrak{A} has a tree-decomposition of the desired form:

$$\varphi(\bar{z}) := \forall y [y \notin \bar{z} \rightarrow \bigvee_{\rho} (\psi_{k\rho}^{n-1}(\bar{z}))^{\chi(x, y, \bar{z})}],$$

where $\psi^{(\chi(x, y, \bar{z}))}$ denotes the relativisation of ψ to the set of all elements x satisfying $\chi(x, y, \bar{z})$.

It remains to compute the MSO_m -theory of \mathfrak{A} . Let $(U_t)_{t \in T}$ be a connected tree-decomposition of \mathfrak{A} with root component $U_r = \bar{c}$, where we consider T as an order-tree. Let $S \subseteq T$ be the set of successors of the root r . For $s \in S$, let \mathfrak{A}'_s be the substructure of \mathfrak{A} induced by the set

$$\bigcup_{t \geq s} U_t \setminus \bar{c},$$

and let $\mathfrak{A}_s := \langle \mathfrak{A}'_s, \bar{R} \rangle$ be its expansion by the relations

$$R_\sigma := \{ \bar{a} \in (U_s \setminus \bar{c})^k \mid \text{atp}(\bar{a}\bar{c}) = \sigma \},$$

where $\text{atp}(\bar{a})$ denotes the atomic type of the tuple \bar{a} and σ ranges over all atomic types of $2k$ -tuples. Then

$$\mathfrak{A} = \tau \left(\langle \bar{c} \rangle_{\mathfrak{A}} \oplus \bigoplus_{s \in S} \mathfrak{A}_s \right),$$

where the quantifier-free interpretation τ adds the missing relations between elements of \bar{c} and tuples in \mathfrak{A}_s based on the information stored in the relations R_σ . As usual we replace the disjoint union by a generalised sum and it follows that there is some MSO-formula ϑ such that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{A}) = \theta \quad \text{iff} \quad \langle S, \bar{Q} \rangle \models \vartheta,$$

where

$$Q_\rho := \{s \in S \mid \text{Th}_{\text{MSO}}^m(\mathcal{A}_s) = \rho\}.$$

By inductive hypothesis, the sets Q_ρ can be computed in first-order logic. Furthermore, by (a), we may assume that ϑ is a first-order formula. Replacing each predicates Q_ρ in ϑ by a first-order formula defining it, we obtain the desired formula ϑ' computing the theory of \mathcal{A} .

Finally, we can define the formula $\psi_{k\theta}^n$ by

$$\psi_{k\theta}^n(\bar{x}) := \exists \bar{z} [\bar{x} \subseteq \bar{z} \wedge \varphi(\bar{z}) \wedge \vartheta'(\bar{z})]. \quad \square$$

We can deduce the following Compactness Theorem for GSO over trees of bounded height.

Corollary 3.5. *A set Φ of GSO-formulae is satisfiable over \mathcal{T}_n if, and only if, every finite subset $\Phi_o \subseteq \Phi$ is satisfiable over \mathcal{T}_n .*

Proof. (\Rightarrow) is trivial.

(\Leftarrow) We have seen in Theorem 3.4 that, over the class \mathcal{T}_n , every GSO-formula is equivalent to an FO-formula. Hence, we may w.l.o.g. assume that Φ is a set of FO-formulae. Let ϑ be an FO-formula stating that the given model is a tree of height at most n . By assumption, every finite subset of $\{\vartheta\} \cup \Phi$ has a model. By the Compactness Theorem for first-order logic, it follows that there is some structure \mathfrak{T} satisfying all of $\{\vartheta\} \cup \Phi$. Since $\mathfrak{T} \models \vartheta$, we have $\mathfrak{T} \in \mathcal{T}_n$. Hence, \mathfrak{T} is our desired model of Φ . \square

Remark. In the above theorem, the restriction to trees with bounded height is essential. The statement fails for classes \mathcal{T} containing trees of unbounded height. For instance, the following set of FO-formulae is finitely satisfiable over such a class, but not satisfiable over it:

- ♦ one FO-formula stating that the tree contains a vertex with label a ; and
- ♦ for every $n < \omega$, an FO-formula stating that no vertex with a distance of at most n from the root has label a .

Similarly, the theorem fails for MSO over every class \mathcal{G} of graphs that contains all finite paths. We can take:

- ♦ one MSO-formula stating that the graph is a (connected) path with two ends; and
- ♦ for every $n < \omega$, an FO-formula stating that the graph has at least n vertices. $\quad \text{J}$

We can prove a converse to the above theorem if we make a few additional assumptions. If we only consider classes of graphs that are closed under non-induced subgraphs, we obtain the following strong statement.

Corollary 3.6. *Let \mathcal{C} be a class of undirected graphs that is closed under taking (non-induced) subgraphs. The following statements are equivalent.*

- (1) *Over \mathcal{C} , every MSO-sentence is equivalent to an FO-sentence.*
- (2) *Over \mathcal{C} , every GSO-sentence is equivalent to an FO-sentence.*
- (3) *\mathcal{C} has bounded height n tree-width, for some $n < \omega$.*

Proof. (3) \Rightarrow (2) follows by Theorem 3.4 (d) and (2) \Rightarrow (1) is trivial.

(1) \Rightarrow (3) Suppose that \mathcal{C} has unbounded height n tree-width, for all n . Then the Excluded Path Theorem (Theorem VIII.2.8) implies that $\text{Min}(\mathcal{C})$ contains the class of all finite paths. As \mathcal{C} is closed under subgraphs, it follows that \mathcal{C} contains all finite paths. The claim follows since the expressive power of FO and MSO is distinct over paths (e.g., one cannot express that a path has even length in FO). \square

We can replace closure under arbitrary subgraphs by closure under induced subgraphs, if we also replace MSO by GSO. To do so, we need the following Ramsey lemma for graphs.

Lemma 3.7. *For every $k < \omega$, there exists a number $N < \omega$ such that every undirected graph \mathfrak{G} containing a path (not necessarily induced) of length at least N contains one of the following induced subgraphs:*

- (I) *a complete graph of size k ;*
- (II) *a complete bipartite graph where both classes have size at least k ;*
- (III) *a path of length k .*

Proof. Let $\mathfrak{G} = \langle V, E \rangle$ be a graph containing a path v_0, \dots, v_{n-1} of length n . W.l.o.g. we may assume that \mathfrak{G} has no vertices that do not lie on this path. Let \leq be the ordering induced on V by the sequence v_0, \dots, v_{n-1} .

We call a path u_0, \dots, u_{m-1} of \mathfrak{G} *increasing* if $u_0 < \dots < u_{m-1}$. A path u_0, \dots, u_{m-1} is a *shortest increasing path* from v to w if it is increasing, it leads from $u_0 = v$ to $u_{m-1} = w$, and every other such path has length at least m . For every pair $u < w$ of vertices, we fix some shortest increasing path $P[u, w]$ from u to w . (Such a path exists since, for $u = v_i$ and $w = v_k$, the path v_i, v_{i+1}, \dots, v_k connects u and w .)

Note that every shortest increasing path is an induced path of \mathfrak{G} . In particular, if there exists a shortest increasing path of length k , we are done. Consequently, we may assume that the length of every path $P[u, w]$ is less than k . With each increasing sequence $u < u' < w < w'$ of four vertices, we associate the substructure

$$\Omega[u, u'; w, w'] \subseteq \langle \mathfrak{G}, \leq, u, u', w, w' \rangle$$

that is induced by the vertices of $P[u, u'] \cup P[w, w']$. As every shortest increasing path has at most k vertices, it follows that

$$|\Omega[u, u'; w, w']| \leq 2k.$$

Up to isomorphism, there are therefore only finitely many such structures and we obtain a finite colouring of the increasing 4-tuples in V . By the Theorem of Ramsey, there exists some number N such that, if $n \geq N$, then there exists a subset $H \subseteq V$ of size $|H| = 4k$ such that every 4-tuple from H has the same colour. We claim that this constant N satisfies the conditions of the lemma.

Hence, assume that $n \geq N$. Then there exists a set H of size $4k$ and a structure Ω such that

$$\Omega \cong \Omega[u, u'; w, w'], \quad \text{for all } u < u' < w < w' \text{ in } H.$$

By definition Ω consists of two paths u_0, \dots, u_{l-1} and $w_0, \dots, w_{l'-1}$ with possibly some edges between them. Let l and l' be the lengths of these paths.

Since

$$\Omega[x, x'; y, y'] \cong \Omega[y, y'; z, z'], \quad \text{for } x < x' < y < y' < z < z' \text{ in } H,$$

it follows that $l = l'$.

Let $z_0 < \dots < z_{4k-1}$ be an enumeration of H and set $P_i := P[z_i, z_{i+1}]$. We distinguish several cases, depending on what kind of edges exist between the two paths of Ω .

(I) First, suppose that there are no edges between the two paths. Let \mathfrak{F} be the subgraph of \mathfrak{G} induced by $P_0 \cup \dots \cup P_{4k-2}$. By assumption, there are no edges between P_i and P_j , for $j - i > 1$. Consequently, the shortest increasing path from z_0 to z_{4k-1} must contain some vertex from P_1 . By the same argument, this path must contain some vertex from P_2 , and so on. Therefore, the shortest increasing path from z_0 to z_{4k-1} must have length at least $4k - 1$. As this path is an induced subgraph of \mathfrak{G} the claim follows.

(II) Next, suppose that there exists an edge between u_m and w_m , for some $m < l$. Let x_i be the m -th vertex of the path P_i . By assumption, there are edges between x_i and x_j for $j - i > 1$. Consequently, the set $\{x_{2i} \mid i < 2k\}$ induces a complete subgraph of \mathfrak{G} of size $2k$.

(III) Finally, suppose that there exist edges between the two paths, but each such edge connects vertices u_m and $w_{m'}$ with $m \neq m'$. Let x_i and y_i be the m -th and the m' -th vertex of P_i . For $j - i > 1$, it follows that exist an edge between x_i and y_j , but not between x_i and x_j or between y_i and y_j . Consequently, the set $\{x_{2i} \mid i < k\} \cup \{y_{2i} \mid k \leq i < 2k\}$ induces a complete bipartite subgraph of \mathfrak{G} where each class has size k . \square

Theorem 3.8. *Let \mathcal{C} be a class of graphs that is closed under taking induced subgraphs. The following statements are equivalent.*

- (1) FO and GSO have the same expressive power over \mathcal{C} .
- (2) \mathcal{C} has bounded height n tree-width, for some $n < \omega$.

Proof. (2) \Rightarrow (1) follows by Theorem 3.4 (d).

(1) \Rightarrow (2) Suppose that \mathcal{C} has unbounded height n tree-width, for all n . Then the Excluded Path Theorem (Theorem VIII.2.8) implies that $\text{Min}(\mathcal{C})$

contains the class of all finite paths. By Lemma 3.7, this means that \mathcal{C} contains at least one of

- (I) the class of all finite complete graphs;
- (II) the class of all finite complete bipartite graphs; or
- (III) the class of all finite paths.

Over all three of these classes, FO and GSO have different expressive powers. For instance, in all three cases there exists an GSO-formula expressing that the cardinality of the graph is even. \square

Example. This statement does not hold if we replace GSO by MSO. For instance, the class \mathcal{K} of all finite complete graphs has unbounded tree-width, but FO and MSO have the same expressive power over \mathcal{K} .

Finally we present a variant where we can get rid of the closure requirement altogether. In exchange we have to allow formulae with free variables.

Corollary 3.9. *Let \mathcal{C} be a class of undirected graphs. The following statements are equivalent.*

- (1) *Over \mathcal{C} , every MSO-formula (possibly with free GSO-variables) is equivalent to an FO-formula.*
- (2) *Over \mathcal{C} , every GSO-formula (possibly with free GSO-variables) is equivalent to an FO-formula.*
- (3) *\mathcal{C} has bounded height n tree-width, for some $n < \omega$.*

Proof. (3) \Rightarrow (2) follows by Theorem 3.4 (d); and (2) \Rightarrow (1) is trivial.

(1) \Rightarrow (3) Suppose that \mathcal{C} has unbounded height n tree-width, for all n . Then the Excluded Path Theorem (Theorem VIII.2.8) implies that $\text{Min}(\mathcal{C})$ contains the class of all finite paths. By Lemma 3.7, this means that the graphs in \mathcal{C} contain as induced subgraphs

- (I) unboundedly large finite complete graphs;
- (II) unboundedly large finite complete bipartite graphs; or
- (III) unboundedly long finite paths.

Let $\varphi(X, Z)$ be an MSO-formula stating that the subgraph induced by the set X is of the form (I), (II), or (III) and that Z is a Hamiltonian path of this subgraph. Furthermore, let $\psi(Z)$ be an MSO-formula stating that the path Z has even length. The conjunction $\varphi(X, Z) \wedge \psi(Z)$ is not equivalent to any FO-formula. \square

We conclude this section with a few technical results that can be used to prove that a given class cannot be interpreted in a class of trees of bounded height. We start by introducing a way to control the origin map, i.e., to put restrictions on which elements of an input structure are used to encode a given element of the output structure. The idea is simple: we use a definable bijection $\varphi : T \rightarrow T$ on the input structure \mathfrak{T} and then modify τ such that its origin map o is replaced by $o \circ \varphi$ (see also Lemma VII.4.10).

Definition 3.10. Let $d < \omega$ be a constant, τ a k -copying MSO-transduction, and $(\psi_i(x, y))_{i < k}$ a sequence of formulae. The d -rearrangement of τ defined by the family $(\psi_i)_i$ is the $(k + d)$ -copying MSO-transduction σ that, given an input structure \mathfrak{T} ,

- ♦ guesses nk new parameters P_j^i , for $i < k$ and $j < d$;
- ♦ checks that
 - \mathfrak{T} is an order-tree,
 - each formula $\psi_i(x, y)$ defines a function $f_i : T \rightarrow T$ such that $|f_i^{-1}(w)| \leq d$, for all $w \in T$, and
 - every $u \in f_i^{-1}(w)$ belongs to exactly one of the sets P_0^i, \dots, P_{d-1}^i ;
- ♦ if the check fails, σ returns the same structure as τ ;
- ♦ otherwise, σ returns $\pi(\tau(\mathfrak{T}))$ where π is the isomorphism mapping u to

$$\pi(\langle i, u \rangle) := \begin{cases} \langle i, u \rangle & \text{if } f_i(u) = u, \\ \langle k + j, f_i(u) \rangle & \text{if } f_i(u) \neq u \text{ and } u \in P_j^i. \end{cases}$$

Below we will use the following normal form for transductions on trees.

Proposition 3.11. *For every MSO-transduction τ and every constant $d < \omega$, there exists an MSO-transduction σ with the following properties.*

- ♦ σ defines the same relation on structures as τ .
- ♦ For every order-tree \mathfrak{T} such that $\sigma(\mathfrak{T})$ is a graph of indegree at most d , we have

$$o(u) \leq o(v) \quad \text{or} \quad o(v) \leq o(u), \quad \text{for every edge } \langle u, v \rangle \text{ of } \sigma(\mathfrak{T}),$$

where o is the origin map of σ and \leq the tree-ordering of \mathfrak{T} .

Proof. Let \mathfrak{T} be an order-tree such that $\tau(\mathfrak{T})$ is a graph of indegree at most d . Suppose that τ is k -copying. W.l.o.g. we may assume that the formula defining the new edge-relation is of the form

$$\varphi_E(x, y) = \bigvee_{i, j < k} [H_i x \wedge H_j y \wedge \varphi_{ij}(x, y)],$$

where H_0, \dots, H_{k-1} are the predicates for the various copies and the formulae φ_{ij} do not contain any of them. Let m be the maximal quantifier-rank of the formulae φ_{ij} . For $j < k$, we set

$$f_j(u) := \inf \left(\{u\} \cup \{u \sqcap v \mid u \not\leq v, v \not\leq u, \text{ and } \mathfrak{T} \models \varphi_{jl}(u, v) \text{ for some } l < k\} \right).$$

We start by proving that, for every $w \in T$ and every MSO_m -theory θ , there are at most $d(d+1)$ vertices $u \geq w$ such that

$$f_j(u) = w \quad \text{and} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{T}_v, u) = \theta,$$

where z is the immediate successor of w with $z \leq u$. For a contradiction, suppose that there are at least $d(d+1)+1$ such vertices $u_0, \dots, u_{d(d+1)}$. We distinguish two cases.

First, suppose that there exists an immediate successor z of w such that T_z contains more than d of the u_i , say, $u_0, \dots, u_d \in T_z$. By definition of f_j , there exists a vertex $v \in T \setminus T_z$ such that $\mathfrak{T} \models \varphi_{jl}(u_0, v)$, for some l . Hence,

$$\mathfrak{T}_z, u_i \equiv_{\text{MSO}}^m \mathfrak{T}_z, u_0 \quad \text{implies} \quad \mathfrak{T} \models \varphi_{jl}(u_i, v), \quad \text{for all } 0 \leq i \leq d.$$

Consequently, the vertex $\langle l, v \rangle$ has indegree at least $d+1$ in $\tau(\mathfrak{T})$. A contradiction.

It follows that each subtree T_z contains at most d of the u_i . Then we can find immediate successors z_0, \dots, z_{d+1} of w and vertices $u_{l_0}, \dots, u_{l_{d+1}}$ with $z_i \leq u_{l_i}$, for all i . W.l.o.g. we may assume that $l_i = i$, for all i . As above, there exists a vertex $v \in T \setminus T_{z_0}$ such that $\mathfrak{T} \models \varphi_{jl}(u_0, v)$, for some l . This vertex is contained in at most one of the subtrees $T_{z_0}, \dots, T_{z_{d+1}}$. W.l.o.g. we may assume that $v \notin T_{z_0}, \dots, T_{z_d}$. It therefore follows that

$$\mathfrak{T}_{z_i}, u_i \equiv_{\text{MSO}}^m \mathfrak{T}_{z_0}, u_0 \quad \text{implies} \quad \mathfrak{T} \models \varphi_{jl}(u_i, v), \quad \text{for all } 0 \leq i \leq d.$$

Again a contradiction to the fact that v has indegree at most d .

This concludes the proof of the above claim. It now follows that

$$|f_j^{-1}(v)| \leq d(d+1)N, \quad \text{for all } v \in f_j(T) \text{ and all } j,$$

where N is the number of Th_{MSO}^m -theories. Let σ be the $d(d+1)N$ -rearrangement of τ defined by the functions $(f_j)_j$. To show that σ has the desired properties, consider an edge $\langle \langle i, u \rangle, \langle j, v \rangle \rangle$ of $\sigma(T)$. We distinguish four cases.

$(i, j < k)$ In this case, both u and v were not moved from their original positions in $\tau(\mathfrak{T})$. Consequently, u and v are fixed-points of, respectively, f_i and f_j , which implies that $u \leq v$ or $v \leq u$.

$(i < k \text{ and } j \geq k)$ In this case v was moved to a new position, whereas u remained where it was. Hence, $f_i(u) = u$, which implies that $u \leq v$ or $v \leq u$.

$(i \geq k \text{ and } j < k)$ In this case u was moved to a new position, whereas v remained where it was. Fix an element $u' \neq u$ with $f_{i-k}(u') = u$. By definition of f_{i-k} , we have $u = f_{i-k}(u') \leq v'$, for all v' , with

$$\mathfrak{T} \models \varphi_{(i-k)j}(u', v').$$

Since v satisfies this condition, it follows that $u \leq v$.

$(i, j \geq k)$ Both vertices were moved to new positions. Hence, there are vertices $u' \neq u$ and $v' \neq v$ with

$$f_{i-m}(u') = u, \quad f_{k-m}(v') = v, \quad \text{and} \quad \mathfrak{T} \models \varphi_{(i-m)(k-m)}(u', v').$$

As in the previous case, we have $u = f_{i-k}(u') \leq v'$. Since $v = f_{j-k}(v') \leq v'$ it follows that $v \leq u$ or $u \leq v$. \square

Corollary 3.12. *Let $n < \omega$ and let τ be an MSO-transduction. There exists an MSO-transduction σ such that*

$$\sigma(\mathfrak{A}) \cong \tau(\mathfrak{A}), \quad \text{for all structures } \mathfrak{A},$$

and such that, for every order-tree \mathfrak{T} such that $\tau(\mathfrak{T})$ is an order-tree of height at most n , we have

$$u \leq v \text{ in } \sigma(\mathfrak{T}) \quad \text{implies} \quad o(u) \leq o(v) \text{ in } \mathfrak{T},$$

where o is the origin map of σ and \leq the tree-ordering of \mathfrak{T} and $\sigma(\mathfrak{T})$.

Proof. Suppose that

$$\tau = \tau_o \circ \rho \circ \text{copy}_k \circ \exp_m,$$

let $\varphi_{\leq}(x, y)$ be the formula of τ_o defining the new order relation. W.l.o.g. we may assume that

$$\varphi_{\leq}(x, y) = \bigvee_{i,j < k} [H_i x \wedge H_j y \wedge \varphi_{ij}(x, y)],$$

where H_0, \dots, H_{k-1} are the predicates for the k copies and the formulae φ_{ij} do not contain any of them. Suppose that $\langle \mathfrak{T}, \mathfrak{S} \rangle \in \tau$ are order trees of height at most n and let \tilde{P} be the parameters used by τ to define \mathfrak{S} out of \mathfrak{T} .

Note that in an order-tree of height at most n , every vertex has at most $n+1$ predecessors (with respect to the relation \leq). Consequently, we can use Proposition 3.11 with $d := n+1$ to make sure that

$$u \leq v \text{ in } \tau(\mathfrak{T}) \quad \text{implies} \quad u \leq v \text{ or } v \leq u.$$

For $i < k$ and $u \in T$, we set

$$f_i(u) := \inf \{ v \leq u \mid \langle \mathfrak{T}, \tilde{P} \rangle \models \varphi_{ij}(u, v) \text{ for some } j \}.$$

Note that we have

$$|f_i^{-1}(v)| \leq k(n+1), \quad \text{for all } v \in T,$$

since, in $\mathfrak{S} = \tau(\mathfrak{T})$, there are at most $k(n+1)$ vertices that lie on a path from the root to some vertex in $o^{-1}(v)$. It follows that the $k(n+1)$ -rearrangement of τ by $(f_i)_i$ has the desired properties. \square

Let us use this normal form for transductions to prove that certain classes of trees cannot be obtained via a transduction from other ones.

Definition 3.13. Let \mathfrak{S} and \mathfrak{T} be order-trees. A k -embedding of \mathfrak{S} in \mathfrak{T} is a function $h : S \rightarrow T$ such that

- ♦ $u \leq v$ implies $h(u) \leq h(v)$,
- ♦ $|h^{-1}(u)| \leq k$,

for all $u, v \in S$. \lrcorner

Proposition 3.14. *Let \mathcal{C} and \mathcal{D} be classes of trees of bounded height. Then the following statements are equivalent.*

- (1) *There exists an MSO-transduction with $\mathcal{C} \subseteq \tau[\mathcal{D}]$.*
- (2) *There exists a GSO-transduction with $\mathcal{C} \subseteq \tau[\mathcal{D}]$.*
- (3) *There exists a constant $k < \omega$ such that every $\mathfrak{S} \in \mathcal{C}$ has a k -embedding into some $\mathfrak{T} \in \mathcal{D}$.*

Proof. (1) \Leftrightarrow (2) follows by Lemma 4.6 since trees are 1-sparse.

(1) \Rightarrow (3) Fix an MSO-transduction τ with $\mathcal{C} \subseteq \tau[\mathcal{D}]$ and let $n < \omega$ be a bound on the height of trees in $\mathcal{C} \cup \mathcal{D}$. Let σ be the MSO-transduction from Lemma 3.12 and suppose that σ is k -copying. Then $\mathcal{C} \subseteq \tau[\mathcal{D}]$ and, for every pair $\langle \mathfrak{C}, \mathfrak{D} \rangle \in \sigma$, the origin map o induces a k -embedding $\mathfrak{C} \rightarrow \mathfrak{D}$.

(3) \Rightarrow (1) Given a tree \mathfrak{T} we can encode an k -embedding $h : \mathfrak{S} \rightarrow \mathfrak{T}$ as follows. Let $p : T \rightarrow T$ be the function mapping each vertex $v \in T$ to its predecessor. (We set $p(v) := v$ if v is the root of T .) For $v \in T$, let v_0, v_1, \dots be an enumeration of $h^{-1}(v)$. We encode h by the following predicates.

$$P_n := \{ v \in T \mid |h^{-1}(v)| = n \},$$

$$Q_{ijn} := \{ v \in T \mid v_k \text{ is a successor of } p^n(v_i) \text{ in } \mathfrak{S} \}.$$

We can recover \mathfrak{S} from $\text{copy}_k(\langle \mathfrak{T}, \bar{P}, \bar{Q} \rangle)$ using the interpretation defined by

$$\begin{aligned}\delta(x) &:= \bigvee_{i \leq n} [H_i x \wedge P_n x], \\ \varphi(x, y) &:= \bigvee_{i, j, n} [H_i x \wedge H_j y \wedge Q_{i j n}(y) \wedge x = p^n(y)],\end{aligned}$$

where the predicates H_0, \dots, H_{k-1} denote the various copies of \mathfrak{T} . \square

As a simple application let us show that there are no transductions that can encode trees of a certain height in trees of smaller height.

Lemma 3.15. *Let $k, d, n < \omega$ and let \mathfrak{T}_d^n be the complete d -ary tree of height n .*

- (a) *There is no k -embedding of \mathfrak{T}_k^{n+1} into \mathfrak{T}_d^n .*
- (b) *If \mathfrak{S} is a tree of height at most $n + 1$ such that the tree \mathfrak{T}_d^{n+1} cannot be k -embedded into \mathfrak{S} , then there exists a d -embedding of \mathfrak{S} into some tree of height n .*

Proof. (a) We prove the claim by induction on n . For a contradiction, suppose that there exists a k -embedding $h : \mathfrak{T}_k^{n+1} \rightarrow \mathfrak{T}_d^n$. Let v be the root of \mathfrak{T}_k^{n+1} and let u_0, \dots, u_{k-1} be its successors. Then h induces a k -embedding of $\mathfrak{T}_k^{n+1}|_{u_i}$ into $\mathfrak{T}_d^n|_{h(u_i)}$. By inductive hypothesis, it follows that $\mathfrak{T}_d^n|_{h(u_i)}$ has height at least n . Consequently, h maps u_i to the root w of \mathfrak{T}_d^k . This implies that also $h(v) = w$. Consequently, $h^{-1}(w)$ contains at least $k + 1$ elements v, u_0, \dots, u_{k-1} . A contradiction.

(b) We define the desired embedding h as follows. Let $P \subseteq T$ be the minimal (w.r.t. \subseteq) set of vertices that contains

- ♦ every leaf of \mathfrak{S} at level $n + 1$ and
- ♦ every vertex that has at least d successors in P .

This implies that, if $v \in P$ is a vertex on level k of \mathfrak{S} , we can embed \mathfrak{T}_d^{n-k} into the subtree attached at v . As \mathfrak{T}_d^{n+1} cannot be embedded into \mathfrak{S} , it follows that P does not contain the root of \mathfrak{S} . Let F be the set of all edges of \mathfrak{S} linking some vertex in $S \setminus P$ to a vertex in P . By definition of P it follows that

- (I) every vertex of \mathfrak{S} has less than d F -successors;
- (II) every path of \mathfrak{S} from the root to some leaf on level $n + 1$ contains at least one edge from F ; and
- (III) no such path contains two consecutive edges from F .

Let \mathfrak{R} be the tree obtained from \mathfrak{S} by contracting every edge in F and let $h : \mathfrak{S} \rightarrow \mathfrak{R}$ be the function mapping each vertex of \mathfrak{S} to the corresponding vertex of \mathfrak{R} . By (II), \mathfrak{R} has height at most n . Furthermore, by (I) and (III), every vertex v of \mathfrak{R} has at most d preimages under h . Hence, h is a d -embedding. \square

Corollary 3.16. *There does not exist a GSO-transduction τ mapping the class of all trees of height n to the class of all trees of height $n + 1$.*

4 The Transduction Hierarchy

For two classes \mathcal{C} and \mathcal{D} , we can regard a transduction τ with $\mathcal{C} \subseteq \tau(\mathcal{D})$ as a way to *encode* every structure $\mathfrak{C} \in \mathcal{C}$ in some $\mathfrak{D} \in \mathcal{D}$. The transduction τ then provides the corresponding *decoding map*. In this section we study the encoding power of such classes with respect to GSO-transductions. It turns out that there are not that many possibilities and we can work out the complete list. Formally, we are interested in the equivalence classes associated with the following preorder.

Definition 4.1. Let L be either MSO or GSO. For two classes \mathcal{C}, \mathcal{D} of structures, we define

$$\mathcal{C} \sqsubseteq_L \mathcal{D} \quad : \text{iff} \quad \mathcal{C} \subseteq \tau(\mathcal{D}), \quad \text{for some } L\text{-transduction } \tau.$$

The relation \sqsubseteq_L is obviously reflexive and transitive, that is, it forms a preorder. The classes figuring in our characterisation below are the following.

Definition 4.2. We denote by

- ♦ \mathcal{P} the class of all finite paths;
- ♦ \mathcal{G} the class of all finite grids;

- ♦ \mathcal{T}_0 the class containing only the empty structure; and
- ♦ \mathcal{T}_n the class of all finite trees of height strictly less than n , for $0 < n \leq \omega$.

These classes are closely related to the various variants of tree-width we have introduced above.

Proposition 4.3. *Let \mathcal{C} be a class of finite structures.*

- (a) $\mathcal{C} \sqsubseteq_{\text{GSO}} \mathcal{T}_\omega$ iff \mathcal{C} has bounded tree-width.
- (b) $\mathcal{C} \sqsubseteq_{\text{GSO}} \mathcal{P}$ iff \mathcal{C} has bounded path-width.
- (c) $\mathcal{C} \sqsubseteq_{\text{GSO}} \mathcal{T}_{n+1}$ iff \mathcal{C} has bounded height- n tree-width.

Proof. The implication (\Leftarrow) follows by Proposition 2.1, while (\Rightarrow) follows by Corollary 2.4. \square

We obtain the following explicit description of the relation \sqsubseteq_{GSO} .

Theorem 4.4 (GSO-Transduction Hierarchy). *We have the following hierarchy.*

$$\emptyset \sqsubseteq_{\text{GSO}} \mathcal{T}_0 \sqsubseteq_{\text{GSO}} \mathcal{T}_1 \sqsubseteq_{\text{GSO}} \cdots \sqsubseteq_{\text{GSO}} \mathcal{T}_n \sqsubseteq_{\text{GSO}} \cdots \sqsubseteq_{\text{GSO}} \mathcal{P} \sqsubseteq_{\text{GSO}} \mathcal{T}_\omega \sqsubseteq_{\text{GSO}} \mathcal{G}.$$

Furthermore, every class of finite structures is \sqsubseteq_{GSO} -equivalent to exactly one of these classes.

Proof. We start with the two extremal classes. Clearly, $\emptyset \sqsubseteq_{\text{GSO}} \mathcal{C}$, for every non-empty class \mathcal{C} . Furthermore, it follows by Lemma 1.10 that $\mathcal{C} \sqsubseteq_{\text{GSO}} \mathcal{G}$, for all classes \mathcal{C} .

Next, we establish the basic ordering between the classes. We have seen in Lemma 1.8 that $\mathcal{T}_n \sqsubseteq_{\text{GSO}} \mathcal{P}$, for all n . Together with the obvious transductions mapping \mathcal{T}_ω to \mathcal{P} and \mathcal{T}_{n+1} to \mathcal{T}_n , it follows that

$$\mathcal{T}_0 \sqsubseteq_{\text{GSO}} \mathcal{T}_1 \sqsubseteq_{\text{GSO}} \cdots \sqsubseteq_{\text{GSO}} \mathcal{T}_n \sqsubseteq_{\text{GSO}} \cdots \sqsubseteq_{\text{GSO}} \mathcal{P} \sqsubseteq_{\text{GSO}} \mathcal{T}_\omega \sqsubseteq_{\text{GSO}} \mathcal{G}.$$

For strictness, note that

- ♦ we have $\mathcal{T}_{n+1} \not\sqsubseteq_{\text{GSO}} \mathcal{T}_n$ by Corollary 3.16,
- ♦ \mathcal{T}_ω has unbounded path-width, by Proposition VIII.1.7, and
- ♦ \mathcal{G} has unbounded tree width, by Proposition VIII.3.6.

By Proposition 4.3, this implies that $\mathcal{P} \sqsubseteq_{\text{GSO}} \mathcal{T}_\omega$ and $\mathcal{T}_\omega \sqsubseteq_{\text{GSO}} \mathcal{G}$.

Hence, it remains to show that every class \mathcal{C} is equivalent to one of the classes in the above hierarchy. We consider several cases. If $\mathcal{C} \not\sqsubseteq_{\text{GSO}} \mathcal{T}_\omega$, then \mathcal{C} has unbounded tree-width and the Excluded Grid Theorem implies that $\mathcal{G} \subseteq \text{Min}(\mathcal{C})$. By Lemma 1.7, it follows that $\mathcal{G} \sqsubseteq_{\text{GSO}} \mathcal{C}$.

If, on the other hand, $\mathcal{C} \sqsubseteq_{\text{GSO}} \mathcal{T}_\omega$ then \mathcal{C} has bounded tree-width. Let τ be the transduction from Theorem 2.6 mapping a structure to its tree-decomposition and let σ be the inverse transduction from Proposition 2.1. Setting $\mathcal{S} := \tau(\mathcal{C})$ it follows that

$$\mathcal{C} = \sigma[\mathcal{S}] \sqsubseteq_{\text{GSO}} \mathcal{S} = \tau[\mathcal{C}] \sqsubseteq_{\text{GSO}} \mathcal{C}.$$

Consequently, \mathcal{C} is equivalent to a class \mathcal{S} of trees. We have to show that \mathcal{S} is equivalent to one of \mathcal{P} or \mathcal{T}_n , $n \leq \omega$. This follows from the following three statements.

- (a) $\mathcal{T}_\omega \not\sqsubseteq_{\text{GSO}} \mathcal{S}$ implies $\mathcal{S} \sqsubseteq_{\text{GSO}} \mathcal{P}$.
- (b) $\mathcal{P} \not\sqsubseteq_{\text{GSO}} \mathcal{S}$ implies $\mathcal{S} \sqsubseteq_{\text{GSO}} \mathcal{T}_n$, for some $n < \omega$.
- (c) $\mathcal{T}_{n+1} \not\sqsubseteq_{\text{GSO}} \mathcal{S}$ implies $\mathcal{S} \sqsubseteq_{\text{GSO}} \mathcal{T}_n$.

(a) Suppose that $\mathcal{S} \not\sqsubseteq_{\text{GSO}} \mathcal{P}$. By Proposition 4.3, this implies that $\text{pwd } \mathcal{S}$ is unbounded. Consequently, we can use the Excluded Tree Theorem (Theorem VIII.2.14) to show that $\mathcal{T}_\omega \subseteq \text{Min}(\mathcal{S})$. Hence, $\mathcal{T}_\omega \sqsubseteq_{\text{GSO}} \mathcal{S}$.

(b) If \mathcal{S} contains a tree of height h , we can define a path of length h in it. Since $\mathcal{P} \not\sqsubseteq_{\text{GSO}} \mathcal{S}$, the height of the trees in \mathcal{S} must therefore be bounded by some constant $h < \omega$. Consequently, $\mathcal{S} \sqsubseteq_{\text{GSO}} \mathcal{T}_{h+1}$.

(c) Note that $\mathcal{T}_{n+1} \not\sqsubseteq_{\text{GSO}} \mathcal{S}$ implies $\mathcal{P} \not\sqsubseteq_{\text{GSO}} \mathcal{S}$, since $\mathcal{T}_{n+1} \sqsubseteq_{\text{GSO}} \mathcal{P}$. By (b), it follows that $\mathcal{S} \sqsubseteq_{\text{GSO}} \mathcal{T}_m$, for some m . According to Proposition 3.14, there exists a constant l such that, every tree $\mathfrak{S} \in \mathcal{S}$ has an l -embedding into some $\mathfrak{T} \in \mathcal{T}_m$. Let $\mathcal{S}' \subseteq \mathcal{T}_m$ be the image of this embedding. Then $\mathcal{S}' \sqsubseteq_{\text{GSO}} \mathcal{S}$ and it is sufficient to prove that $\mathcal{S}' \sqsubseteq_{\text{GSO}} \mathcal{T}_n$.

We proceed by induction on m . If $m \leq n$, there is nothing to do. For $m > n$, note that $\mathcal{T}_{n+1} \not\equiv_{\text{GSO}} \mathcal{S}'$ implies $\mathcal{T}_m \not\equiv_{\text{GSO}} \mathcal{S}'$. According to Proposition 3.14, this means that, for every $k < \omega$, we can find some tree $\mathfrak{T}_k \in \mathcal{T}_m$ such that there is no k -embedding of \mathfrak{T}_k to any $\mathfrak{S} \in \mathcal{S}'$. W.l.o.g. we may assume that \mathfrak{T}_k is the complete d -ary tree of height m . By Lemma 3.15 (b), this implies that every $\mathfrak{S} \in \mathcal{S}'$ can be d -embedded into some tree from \mathcal{T}_{m-1} . Hence, $\mathcal{S}' \equiv_{\text{GSO}} \mathcal{T}_{m-1}$ and the claim follows by inductive hypothesis. \square

One way to interpret this theorem is as follows. As each of the classes above correspond to a variant of tree-width, the theorem provides a complete list of all the complexity measures for graphs that are compatible with guarded second-order logic. There are no such measures besides: (I) tree-width; (II) path-width; and (III) height- n tree-width.

Corollary 4.5. *Let \mathcal{C} be a class of relational structures. Then exactly one of the following two cases holds.*

- (1) \mathcal{C} has bounded tree-width and, over \mathcal{C} , GSO coincides with MSO.
- (2) \mathcal{C} has unbounded tree-width and the GSO-theory of \mathcal{C} interprets the full second-order theory of all finite sets.

Proof. By Lemma VIII.1.8, if \mathcal{C} has bounded tree-width, there is some constant k such that every structure in \mathcal{C} is k -sparse. Hence it follows by Theorem VII.3.6 that GSO coincides with MSO.

If \mathcal{C} has unbounded tree-width, it follows by the Excluded Grid Theorem that there exists a GSO-transduction mapping \mathcal{C} to the class \mathcal{G} of all finite grids. Furthermore, the MSO-theory of \mathcal{G} interprets the second-order theory of all finite sets since

- on an $n \times n$ -grid we can encode binary relations on a set of size n by unary relations, and
- every quantifier over an m -ary relation on a set of size n can be simulated by quantifiers over m binary relations on a set of size $n^m + n$ (by encoding the corresponding incidence structure). \square

Complete descriptions of the corresponding hierarchies for other logics, like MSO or FO are still missing. At this point, let us just mention the following result that allows us to transfer results for GSO to ones for MSO.

Lemma 4.6. *If \mathcal{C} and \mathcal{D} are classes of k -sparse structures, then*

$$\mathcal{C} \equiv_{\text{MSO}} \mathcal{D} \quad \text{iff} \quad \mathcal{C} \equiv_{\text{GSO}} \mathcal{D}.$$

Proof. It is straightforward to construct a transduction showing that $\mathcal{C} \equiv_{\text{MSO}} \mathcal{C}_{\text{in}}$, for every class \mathcal{C} . Conversely, if \mathcal{C} contains only k -sparse structures, it follows by Theorem VII.3.6 that $\mathcal{C}_{\text{in}} \equiv_{\text{MSO}} \mathcal{C}$.

Consequently, if both classes consist of only k -sparse structures, then

$$\begin{aligned} \mathcal{C} \equiv_{\text{MSO}} \mathcal{D} & \text{ implies } \mathcal{C}_{\text{in}} \equiv_{\text{MSO}} \mathcal{C} \equiv_{\text{MSO}} \mathcal{D} \equiv_{\text{MSO}} \mathcal{D}_{\text{in}}, \\ \mathcal{C}_{\text{in}} \equiv_{\text{MSO}} \mathcal{D}_{\text{in}} & \text{ implies } \mathcal{C} \equiv_{\text{MSO}} \mathcal{C}_{\text{in}} \equiv_{\text{MSO}} \mathcal{D}_{\text{in}} \equiv_{\text{MSO}} \mathcal{D}. \end{aligned}$$

Since $\mathcal{C} \equiv_{\text{GSO}} \mathcal{D}$ if, and only if, $\mathcal{C}_{\text{in}} \equiv_{\text{MSO}} \mathcal{D}_{\text{in}}$, the claim follows. \square

In particular, it follows that every level of the GSO-transduction hierarchy also appears in the version for MSO (although individual classes might belong to different levels of the two hierarchies). But it might still be possible that the hierarchy for MSO has additional levels.

As an application of the transduction hierarchy, let us take a look at transductions between infinite structures. We will show below that there is no transduction mapping some linear order to the infinite binary tree. The first step is to reduce the statement to one concerning classes of finite structures. For a class \mathcal{C} of Σ -structures, let \mathcal{C}_{fin} denote the class of all finite substructures of structures in \mathcal{C} .

Theorem 4.7. *Let \mathcal{T} be an MSO-definable class of generalised trees. Then*

$$\mathcal{C} \equiv_{\text{MSO}} \mathcal{T} \text{ implies } \mathcal{C}_{\text{fin}} \equiv_{\text{MSO}} \mathcal{T}_{\text{fin}}, \text{ for every class } \mathcal{C}.$$

Proof. Let τ be an MSO-transduction of quantifier-rank m with $\mathcal{C} \subseteq \tau[\mathcal{T}]$. Note that, for each $\mathfrak{C} \in \mathcal{C}_{\text{fin}}$, there are structures $\mathfrak{D} \in \mathcal{C}$ and $\mathfrak{T} \in \mathcal{T}$ such that $\mathfrak{C} \subseteq \mathfrak{D} \in \tau(\mathfrak{T})$. The origin map o maps \mathfrak{C} to some subset $o[\mathfrak{C}] \subseteq T$

(when considering o as a map $D \rightarrow T$ and ignoring the other component). Let $\mathfrak{S} := \langle\langle o[C] \rangle\rangle_{\mathfrak{T}} \subseteq \mathfrak{T}$ be the substructure generated by $o[C]$. We will construct an MSO-transduction σ such that $\mathfrak{C} \in \sigma(\mathfrak{S})$.

Given \mathfrak{S} , this transduction has to recover \mathfrak{T} and then apply τ . To see how we can make this work, consider a generalised tree $\mathfrak{T} \in \mathcal{T}$ and a finite substructure $\mathfrak{S} \subseteq \mathfrak{T}$. Let $(T_s)_{s \in S}$ be the decomposition of \mathfrak{T} induced by \mathfrak{S} as in Definition V.1.10 and $(\theta_s)_{s \in S}$ the associated labelling by MSO_m-theories. By Proposition V.1.11, we can evaluate every MSO_m-formula $\varphi(\bar{x})$ over \mathfrak{T} if we have access to $(\theta_s)_{s \in S}$.

We can now define a transduction σ that, given \mathfrak{S} , (i) guesses unary predicates \tilde{Q} encoding the theories $(\theta_s)_s$; (ii) checks that the parameters \tilde{Q} corresponds to a structure $\mathfrak{T} \supseteq \mathfrak{S}$ in \mathcal{T} ; and (iii) applies τ to \mathfrak{T} . The formal definition is as follows. With the help of the predicates \tilde{Q} we can translate every MSO_m-formula $\varphi(\bar{x})$ into a formula $\varphi^*(\bar{x})$ such that

$$\mathfrak{T} \models \varphi(\bar{a}) \quad \text{iff} \quad (\mathfrak{S}, \tilde{Q}) \models \varphi^*(\bar{a}), \quad \text{for all } \bar{a} \subseteq S.$$

Let τ^* be the transduction obtained from τ by replacing every formula φ by its translation φ^* . The desired transduction σ has the form

$$\sigma := \tau^* \circ \rho \circ \exp_n,$$

where \exp_n guesses the additional labelling \tilde{Q} and ρ is a filtering verifying that the labelling \tilde{Q} corresponds to a structure in \mathcal{T} . We have to check to conditions.

(i) The structure \mathfrak{T} corresponding to the labelling \tilde{Q} belongs to \mathcal{T} . (If \mathcal{T} is defined by ψ , we can use the formula ψ^* for this.)

(ii) The labelling \tilde{Q} corresponds to some structure \mathfrak{T} . For this we have to express that the sets Q_θ are disjoint, their union is all of T , and that

$$t \in Q_\theta \quad \text{implies} \quad \theta = \text{MTh}_m(\mathfrak{C}),$$

for some structure of the form $\mathfrak{C} = \langle C, \leq, \sqcap, P, c \rangle$ where

- ♦ $\langle C, \leq, \sqcap \rangle$ is a generalised tree;
- ♦ $P = \downarrow c$; and

- ♦ if $s \in S$ is non-maximal, then c is a maximal element of C .

Note that we do not need to express these properties using formulae. It is sufficient to fix two sets $\Theta_o \subseteq \Theta_i$ of MSO_m-theories where Θ_o contains all theories θ satisfying all three of the above conditions, while Θ_i contains those satisfying the first two of them. Then we can use the formula

$$\forall x \bigvee_{\theta \in \Theta_i} Q_{\theta} x \wedge \forall x \left['x \text{ maximal}' \rightarrow \bigvee_{\theta \in \Theta_o} Q_{\theta} x \right]. \quad \square$$

Corollary 4.8. *There does not exist an MSO-transduction τ that maps some linear order to the infinite binary tree.*


Proof. If there were such a transduction, then $\{\mathfrak{T}\} \sqsubseteq_{\text{MSO}} \mathcal{L}$, where \mathfrak{T} is the infinite binary tree and \mathcal{L} is the class of all linear orders. By the preceding theorem, this would imply that $\mathcal{T}_\omega \sqsubseteq_{\text{MSO}} \mathcal{P}$. A contradiction. \square

Finally, let us also briefly take a look at the hierarchy for GSO over classes of infinite structures. In this case the hierarchy is no longer linear since we have to distinguish between different cardinalities in the branching of the trees.

Definition 4.9. Let Φ be the class of all finite trees of height 1 and Ω the class containing the unique infinite tree of height 1. For classes \mathcal{C} and \mathcal{D} of trees, we define the following operations.

- ♦ $\mathcal{C} \oplus \mathcal{D}$ consists of all trees obtained from a tree $\mathfrak{S} \in \mathcal{C}$ and a tree $\mathfrak{T} \in \mathcal{D}$ by identifying the roots of \mathfrak{S} and \mathfrak{T} .
- ♦ $\mathcal{C} \cdot \mathcal{D}$ consists of all trees obtained from some tree $\mathfrak{S} \in \mathcal{C}$ by replacing every leaf of \mathfrak{S} by some tree from \mathcal{D} .
- ♦ $\mathcal{C} : \mathcal{D}$ consists of all trees obtained from some tree $\mathfrak{S} \in \mathcal{C}$ by replacing every leaf of \mathfrak{S} by the same tree $\mathfrak{T} \in \mathcal{D}$. \lrcorner

For classes of countable trees of height one and two, the GSO-hierarchy looks as follows. We omit the proof, which consists of a straightforward, but lengthy argument based on Proposition 3.14.



5 Defining Tree-Decompositions

Let us finally present the proof of Theorem 2.6 promised above. To simplify terminology, we will identify Σ -structures with their associated hypergraphs and with their Gaifman graphs. In particular, we call guarded sets *hyperedges*. A *path* in \mathfrak{A} is a path in its Gaifman graph, and we will call \mathfrak{A} *connected* if its Gaifman is.

We start by figuring out how to encode a tree-decomposition of a structure by a relation so we can define it in MSO. This does not work for arbitrary tree-decompositions, but only those in a certain normal form. In the following, we will work with tree-decompositions $(U_t)_{t \in T}$ indexed by an *order-tree* instead of an *undirected* one. This is no restriction since we can always define an order by picking a root for T .

Definition 5.1. Let \mathfrak{A} be a Σ -structure and $(U_t)_{t \in T}$ a tree-decomposition.

(a) The *adhesion set* of a component U_t is

$$\text{Ad}(t) := U_t \cap U_s,$$

where s is the parent of t . For the root t , we set $\text{Ad}(t) := \emptyset$.

(b) The decomposition $(U_t)_{t \in T}$ is *reduced* if

- $U_s \neq U_t$, for all $s \neq t$, and

- for every $a \in U_t$, there is some hyperedge e such that $a \in e$ and $e \not\subseteq \text{Ad}(t)$.

(c) A tree-decomposition $(U_t)_{t \in T}$ is *internally connected* if, for every subtree $S \subseteq T$ with root t , the set $\bigcup_{s \in S} U_s \setminus \text{Ad}(t)$ is connected in \mathfrak{A} .

(d) We say that an element $a \in A$ is *introduced* in a component U_t (or simply *at the vertex t*) if

$$a \in U_t \setminus \text{Ad}(t).$$

We denote by $\mu : A \rightarrow T$ the function mapping every element to the vertex it is introduced at.


(e) The *introduction order* \sqsubseteq of $(U_t)_t$ is defined by

$$a \sqsubseteq b \quad : \text{iff} \quad \mu(a) \leq \mu(b).$$

We denote the corresponding equivalence relation by \equiv .

(f) We say that a formula $\varphi(x, y; \bar{Z})$ *defines* $(U_t)_t$ if there are GSO-parameters \bar{P} such that the relation defined by $\varphi(x, y; \bar{P})$ on \mathfrak{A} is the introduction order of $(U_t)_t$.

Example. Not all tree-decompositions are definable, even if they are reduced and internally connected. The cycle of length 6 has the following two tree-decompositions. The bottom one is MSO-definable, but the top one is not. (In order to determine whether two vertices belong to the same component, we would need to count how far away they are for the left-most vertex.)



For reduced tree-decompositions, the introduction order \sqsubseteq contains sufficient information to reconstruct the index tree and the components.

Lemma 5.2. *Let \mathfrak{A} be a Σ -structure and \sqsubseteq the introduction order of a reduced tree-decomposition $(U_t)_{t \in T}$.*

- (a) $\langle T, \leq \rangle \cong \langle A, \sqsubseteq \rangle / \equiv$.
- (b) $U_{\mu(a)} = \{ b \sqsubseteq a \mid b \equiv a \text{ or } b \sqsubseteq a \text{ and there is some } c \sqsupseteq a \text{ and a hyperedge } e \text{ such that } b, c \in e \}$.
- (c) *If \sqsubseteq is definable, there exists a GSO-formula $\psi(X; \bar{Q})$ (with parameters \bar{Q}) such that*

$$\mathfrak{A} \models \psi(C; \bar{Q}) \quad \text{iff} \quad C = U_t, \quad \text{for some } t \in T.$$

Proof. (a) follows from the definition of \sqsubseteq and (b) from the fact that $(U_t)_t$ is reduced. For (c), note that, given a formula $\varphi(x, y; \bar{P})$ defining \sqsubseteq , we can express the condition from (b) in GSO. \square

When constructing formulae to define tree-decompositions below, we will frequently have to distinguish several cases, each of which yields a different formula. The next remark explains how we can combine such formulae into a single one.

Lemma 5.3. *Let $\varphi_o(x, y; \bar{Z}_o)$ and $\varphi_i(x, y; \bar{Z}_i)$ be two GSO-formulae and let \mathcal{C}_i be the class of all structures \mathfrak{A} such that, for some choice of parameters \bar{P} , the formula $\varphi_i(x, y; \bar{P})$ defines a tree-decomposition of \mathfrak{A} of width at most k . Then there exists a GSO-formula $\psi(x, y; \bar{Z})$ that defines a tree-decomposition of width at most k on every structure in $\mathcal{C}_o \cup \mathcal{C}_i$.*

Proof. We can use the formula φ_o to construct a formula $\vartheta(\bar{Z}_o)$ such that

$$\mathfrak{A} \models \vartheta(\bar{P}) \quad \text{iff} \quad \varphi_o(x, y; \bar{P}) \text{ defines a tree-decomposition of } \mathfrak{A} \text{ of width at most } k.$$

Then the desired formula is

$$\psi(x, y; \bar{Z}_o \bar{Z}_i) := [\vartheta(\bar{Z}_o) \wedge \varphi_o(x, y; \bar{Z}_o)] \vee [\neg \vartheta(\bar{Z}_o) \wedge \varphi_i(x, y; \bar{Z}_i)]. \quad \square$$

Also note that we can always assume that our tree-decomposition is reduced, and usually also that it is internally connected.

Lemma 5.4. (a) *For every tree-decomposition $(U_t)_{t \in T}$, there exists a reduced tree-decomposition $(W_s)_{s \in S}$ of the same width where the index tree S is a minor of T . Furthermore, $(W_s)_s$ has the same introduction order as $(U_t)_t$ and, if $(U_t)_t$ is internally connected, so is $(W_s)_s$.*

(b) *Every connected structure \mathfrak{A} has a tree-decomposition of width $\text{twd } \mathfrak{A}$ that is internally connected and reduced.*

Proof. (a) We obtain $(W_s)_s$ from $(U_t)_t$ in two steps. First we remove all elements $a \in U_t$ such that there is no hyperedge e as above. Then we contract every edge $\langle s, t \rangle$ of the resulting tree-decomposition $(U'_t)_t$ such that $U'_s = U'_t$. This operation does not change the connectedness of the decomposition nor its introduction order.

(b) Let S be a subtree such that $\bigcup_{s \in S} U_s$ has several connected components C_0, \dots, C_{n-1} . By assumption, this implies that $S \neq T$. Let S_0, \dots, S_{n-1} be disjoint copies of S and set

$$U_s^i := U_s \cap S_i, \quad \text{for } s \in S_i.$$

In the decomposition $(U_t)_{t \in T}$, we replace the subtree $(U_s)_{s \in S}$ by the forest consisting of the trees $(U_s^o)_{s \in S_0}, \dots, (U_s^{n-1})_{s \in S_{n-1}}$. We repeat this construction until the resulting tree-decomposition is internally connected. Finally, we can use (a) to make the tree-decomposition we have obtained in this way reduced. \square

The formal statement we will prove in the rest of this section is the following theorem.

Theorem 5.5. *Let Σ be a finite relational signature and $k < \omega$. There exists a GSO-formula $\varphi_k(x, y; \bar{Z})$ and a constant $l < \omega$ with the following property: for every finite Σ -structure \mathfrak{A} with $\text{twd } \mathfrak{A} \leq k$, there are GSO-parameters \bar{P} such that $\varphi_k(x, y; \bar{P})$ defines the introduction order \sqsubseteq of some reduced tree-decomposition $(U_t)_{t \in T}$ of \mathfrak{A} of width at most l .*

Before given the proof, let us explain how to use this result to prove Theorem 2.6.

Proof of Theorem 2.6. By the preceding theorem, there exists a formula defining the introduction ordering Ξ . According to Lemma 5.2, we can construct the associated tree by

- ♦ picking one representative from every Ξ -equivalence class and
- ♦ adding a new element as the root if there are several minimal Ξ -classes.

All of this can be done using a GSO-transduction. Since structures of tree-width at most k are r -sparse, for some r , this GSO-transduction can be transformed into an MSO-transduction using Theorem VII.3.6. \square

For the proof, we need a way to orient hyperedges e . (If e corresponds to a guarded tuple $\bar{c} \in R$, we of course already have an ordering on the elements c_0, \dots, c_{n-1} . But below we will also use hyperedges not associated with guarded tuples.) We can do so by choosing a suitable colouring of \mathfrak{A} and order the elements of e in the order of their colours.

Definition 5.6. Let $(U_t)_{t \in T}$ be a tree-decomposition of a Σ -structure \mathfrak{A} of width k . A *local colouring* of \mathfrak{A} is a function $\chi : A \rightarrow [k]$ that is injective on every set U_t . \downarrow

Lemma 5.7. Every tree-decomposition has a local colouring.

Proof. Let $(U_t)_{t \in T}$ be a tree-decomposition of width k . We define $\chi(a)$ inductively as follows. We start at an arbitrary vertex $t \in T$ by choosing some injective function $\chi : U_t \rightarrow [k]$. For the inductive step, suppose that χ is already defined for all elements in the component U_t . Let s be a neighbour of t and let $D \subseteq U_s$ be the set of elements a such that $\chi(a)$ is already defined. Then $|U_s| \leq k$ implies that we can define χ on the remaining elements in $U_s \setminus D$ such that $\chi \upharpoonright U_s$ is injective. \square

Merging Decompositions

We start by showing how to assemble partial tree-decompositions into one of the whole structure.

Definition 5.8. Let $(U_t)_{t \in T}$ be a tree-decomposition of a structure \mathfrak{A} and let $\mu : A \rightarrow T$ be the corresponding introduction map.

(a) A *factor* of T is a connected subset $F \subseteq T$. The minimal element of F is called its *root*. The *border* of F consists of all vertices $t \in T \setminus F$ whose predecessor belongs to F .

(b) The *fragment* of \mathfrak{A} corresponding to a factor F with root t is the hypergraph $\mathfrak{A}[F]$ obtained from the restriction of \mathfrak{A} to the set $\bigcup_{s \in F} U_s$ by adding new hyperedges containing the adhesion sets of t and of every component in the border. We call such a new hyperedge e corresponding to the component U_s an *adhesion edge representing s* . For $s = t$, we call e the *root adhesion edge* and for s in the border of F a *border adhesion edge*.

We can encode these edges by adding a new relation R to \mathfrak{A} whose arity is equal to the width of $(U_t)_t$ such that R contains one tuple for every adhesion edge (we can order each edge increasingly with respect to some fixed local colouring of \mathfrak{A}).

Finally, we denote by $\mathfrak{A}^*[F]$ the structure obtained from $\mathfrak{A}[F]$ by removing the root adhesion edge.

(c) We say that a formula $\psi(x, y; \bar{Z})$ defines a *refinement* of $(U_t)_t$ if, for every $t \in T$, $\psi(x, y; \bar{Z})$ defines (in the structure $\mathfrak{A}[\{t\}]$) a tree-decomposition of the fragment $\mathfrak{A}[\{t\}]$. We call the tree-decompositions defined by ψ the *local decompositions*. \downarrow

Proposition 5.9. Let $\varphi(x, y; \bar{Z})$ and $\psi(x, y; \bar{Z}')$ be two GSO-formulae and $k < \omega$ a constant. There exists a GSO-formula $\vartheta(x, y; \bar{Z}'')$ with the following property. If $\varphi(x, y; \bar{Z})$ defines a tree-decomposition $(U_t)_{t \in T}$ of some Σ -structure \mathfrak{A} and $\psi(x, y; \bar{Z}')$ defines a refinement of $(U_t)_t$ where every local decomposition has width at most k , then $\vartheta(x, y; \bar{Z}'')$ defines a tree-decomposition of \mathfrak{A} of width at most k .

Proof. Let $(U_t)_{t \in T}$ be the tree-decomposition defined by φ , Ξ its introduction order, and μ the corresponding function. Let $(W_s^t)_{s \in S^t}$ be the local tree-decomposition of the fragment at $\{t\}$ defined by ψ . We can define a second tree-order on S^t such that the minimal component (with respect to this new ordering) contains the root adhesion edge representing t . Since this new ordering is definable in GSO (using the old ordering and the in-




Figure 3: Merging tree-decompositions by their adhesion edges

introduction ordering \sqsubseteq of $(U_t)_t$, we can modify the formula ψ such that the decompositions $(W_s^t)_{s \in S^t}$ it defines are already ordered in this way. We define the desired tree-decomposition $(W_s^*)_{s \in S^*}$ of \mathcal{A} as follows. Let S^* be the tree obtained from the disjoint union of the trees S^t , $t \in T$, by adding, for every edge $\langle t_o, t_i \rangle$ of T , an edge from $s \in S^{t_o}$ to the root of S^{t_i} , where s is the vertex such that $W_s^{t_o}$ contains the adhesion edge representing t_i . We define the component at a vertex $s \in S^t \subseteq S^*$ by $W_s^* := W_s^t$.

Clearly, the resulting family $(W_s^*)_{s \in S^*}$ is a tree-decomposition of width at most k . We claim that it is definable. First, note that we can define its introduction ordering \sqsubseteq^* by

$$a \sqsubseteq^* b \quad : \text{iff} \quad \exists z I(a, b; z) \\ \vee \exists x \exists y [x \sqsubseteq y \wedge I(a, a; x) \wedge I(b, b; y)],$$

where \sqsubseteq denotes the introduction order of $(U_t)_{t \in T}$ and the relation I is

defined by

$$I(a, b; c) \quad : \text{iff} \quad a \sqsubseteq^{\mu(c)} b,$$

where \sqsubseteq^t is the introduction order of the local decomposition $(W_s^t)_{s \in S^t}$.

By assumption we have a formula φ defining \sqsubseteq . Hence it remains to show that I is definable. We can check whether a and b are introduced in the same component as c using the order \sqsubseteq . What is left is to verify that $a \sqsubseteq^t b$ holds. Let \bar{P}^t be the parameters used by the formula $\psi(x, y; \bar{Z}')$ to define the tree-decomposition $(W_s^t)_{s \in S^t}$. When trying to use $\psi(x, y; \bar{Z}')$ to define \sqsubseteq^t , we face the problem that we have to reconstruct the correct values of \bar{P}^t . Thus, it remains to show that we can encode all the families \bar{P}^t , $t \in T$, by finitely many predicates \bar{R} in such a way, that \bar{P}^t can be recovered from \bar{R} given some element $c \in \mu^{-1}(t)$ to specify the vertex t .

Note that each fragment and, thus, the sets P_i^t , can contain three different kinds of elements and hyperedges: (I) elements in $\mu^{-1}(t)$; (II) elements in $\text{Ad}(t)$; and (III) adhesion edges. As the adhesion edges are not present in \mathcal{A} we have to encode them by suitable elements of A . If e is the adhesion edge representing a vertex s in the border, we pick some element $c_e \in \mu^{-1}(s)$ and say that c_e represents e . We will consider each of the three types of elements above separately by proving that the following relations are definable (with suitable GSO-parameters).

$$\begin{aligned} K_i(a; c) & \quad : \text{iff} \quad a \in P_i^{\mu(c)} \cap \mu^{-1}(\mu(c)), \\ L_i(a; c) & \quad : \text{iff} \quad a \in P_i^{\mu(c)} \cap \text{Ad}(\mu(c)), \\ M_i(a; c) & \quad : \text{iff} \quad a \text{ represents an adhesion edge } e \in P_i^{\mu(c)}. \end{aligned}$$

We start with the relation M_i . Set

$$\begin{aligned} Q_i & := \{ c_e \mid e \text{ a border adhesion edge in } P_i^t \text{ for some } t \in T \}, \\ Q'_i & := \{ c_e \mid e \text{ the root adhesion edge in } P_i^t \text{ for some } t \in T \}. \end{aligned}$$

Since every adhesion edge belongs to a exactly two fragments, once as a

border edge and once as the root edge, it follows that

$$M_i = \{ \langle a, c \rangle \mid a \in Q_i, c \sqsubset a, \text{ and there is no } b \text{ with } c \sqsubset b \sqsubset a \} \\ \cup \{ \langle a, c \rangle \mid a \in Q'_i, c \equiv a \}.$$

For the relation K_i we can similarly use the parameter

$$R_i := \bigcup_{t \in T} [P_i^t \setminus \text{Ad}(t)].$$

Then

$$K_i = \{ \langle a, c \rangle \mid a \in R_i, a \equiv c \}.$$

Hence, it remains to consider the relations L_i . Fix a local colouring $\chi : A \rightarrow [k]$ of the combined decomposition $(W_s^*)_{s \in S^*}$. By our assumption on the order Ξ^t , the root component W_\emptyset^t contains $\text{Ad}(t)$. Since W_\emptyset^t is also a component of $(W_s^*)_s$ it follows that χ is injective on $\text{Ad}(t) \subseteq W_\emptyset^t$. For a colour $j \in [k]$, let $g_j(c)$ be the unique vertex in $\chi^{-1}(j) \cap \text{Ad}(\mu(c))$ (if it exists). Using the parameters

$$S_{ij} := \{ c \mid g_j(c) \in P_i^{\mu(c)} \} \quad \text{and} \quad S'_j := \chi^{-1}(j),$$

we can define L_i by

$$L_i = \{ \langle a, c \rangle \mid a \in S'_j, a \in \text{Ad}(\mu(c)), \text{ and } c \in S_{ij}, \text{ for some } j \}.$$

(Note that, since $(U_t)_t$ is reduced, we can define $\text{Ad}(\mu(c))$ from c using the introduction order Ξ and the relations of \mathfrak{A} .) \square

Bounded Path-Width

To show that tree-decompositions are definable, we first consider the simple case where the structure in question has bounded path-width. The reason is that we can equip path-decompositions with the structure of a semigroup, which allows us to use tools from semigroup theory.

Definition 5.10. (a) For a signature Σ and a constant $k < \omega$, we denote by \mathcal{P}_k the set of all path-decompositions $(U_i)_{i < m}$ of some finite Σ -structure where the index set m is linearly ordered and the decomposition is equipped with a local colouring.

(b) The *composition* of two path-decompositions $(U_i)_{i < m}$ and $(W_j)_{j < n}$ is the path-decomposition obtained by merging the last component U_{m-1} of $(U_i)_i$ with the first component W_0 of $(W_j)_j$ where the elements in these two components are identified with each other according to their colours. \lrcorner

The set \mathcal{P}_k together with composition forms an infinite semigroup. For what follows, we have to reduce it to a finite one by forming a suitable quotient.

Definition 5.11. We say that two path-decompositions $(U_i)_{i < m}, (W_j)_{j < n}$ in \mathcal{P}_k have *the same profile* if they satisfy the following conditions.

- The same colours appear in U_0 and in W_0 .
- The same colours appear in U_{m-1} and in W_{n-1} .
- For every pair $\langle c, d \rangle$ of colours, there is a path between the c -coloured element in U_0 and the d -coloured one in U_{m-1} if, and only if, there is such a path between W_0 and W_{n-1} .
- There is a c -coloured element in $U_0 \cap U_{m-1}$ if, and only if, there is such an element in $W_0 \cap W_{n-1}$. \lrcorner

Lemma 5.12. *The relation of having the same profile is a congruence of the semigroup \mathcal{P}_k with finitely many classes.*

We can use the resulting quotient semigroup of profiles to construct definable tree-decompositions as follows. Before doing so, let us introduced the following terminology used in the proof.

Definition 5.13. Let $(U_i)_{i \in I}$ be a path-decomposition of \mathfrak{A} .

- (a) An element $a \in A$ is *ubiquitous* if it belongs to every component U_i .
- (b) An *inner path* is a path in \mathfrak{A} that does not contain any ubiquitous elements. An *inner component* of \mathfrak{A} is a maximal set of elements that is connect with respect to inner paths. \lrcorner

Proposition 5.14. *For every $k < \omega$, there exist a constant N and a GSO-formula $\varphi(x, y; \bar{Z})$ that defines (with suitable GSO-parameters) a tree-decomposition of width at most N on every finite structure \mathfrak{A} with $\text{pwd } \mathfrak{A} \leq k$.*

Proof. Let $(U_i)_{i \in I}$ be a path-decomposition of \mathfrak{A} of width at most k . We can regard every component U_i as a 1-component path-decomposition of some substructure of \mathfrak{A} . Then $(U_i)_i$ is just the composition of these 1-component decompositions. Let a_i be the profile of U_i . By Theorem II.3.7, the sequence a_0, \dots, a_{m-1} has a Simon tree T of bounded height h . By induction on h , we will construct a formula that defines a tree-decomposition for all structures whose associated Simon tree has height at most h .

We claim that every inner component of \mathfrak{A} has a definable tree-decomposition of width at most $N := k(h+1)$ (by a single formula only depending on k, h , and the signature). Furthermore, \mathfrak{A} itself has a tree-decomposition consisting of a root component C containing the ubiquitous elements (of which there are at most k) together with one leaf $D \cup C$, for every inner connected component D of \mathfrak{A} . This decomposition is clearly definable (with parameters C). As each leaf $D \cup C$ has a definable tree-decomposition of width at most $N+k$ that can be obtained from that of D by adding the set C to every component, we can use Proposition 5.9 to construct a definable tree-decomposition of \mathfrak{A} of width at most $N+k$.

Hence, it remains to prove the claim. Fix an inner component D of \mathfrak{A} and let \mathfrak{D} be the substructure induced by it. If $h = 0$, the Simon tree consists of a single vertex. Consequently, the structure has at most k elements and we can use a tree-decomposition consisting of a single component covering all of them. For the inductive step, suppose that the Simon tree has height $h > 0$. Then there exists a factorisation F_0, \dots, F_{n-1} of the index set I such that each subdecomposition $(U_s)_{s \in F_i}$ has an associated Simon tree of height at most $h-1$. Furthermore, if there are more than two factors, we may assume that all of them have the same profile and that this profile is idempotent. Set $C_i := D \cap \bigcup_{s \in F_i} U_s$, for $i < n$. (In the following the terms ‘ubiquitous’ and ‘inner path’ will always refer to the original decomposition $(U_i)_{i \in I}$, never to the factors $(U_s)_{s \in F_i}$ or to $(C_i)_{i < n}$.) We say that a path p of \mathfrak{A} visits an index $i < n$ if p contains some element from C_i . We will establish the following

four claims. Then the desired tree-decomposition of \mathfrak{D} can be obtained from (I) and (II) via Proposition 5.9. (The other two claims are needed to prove (II).)

- (I) Every fragment $\mathfrak{D}[F_i]$ associated with one of the above factors F_i has a definable tree-decomposition of width at most $k(h+1)$.
- (II) \mathfrak{D} has a definable path-decomposition of the form $(C_i)_{i < n}$.
- (III) If $n > 2$, then every element of D belongs to at most two sets C_i .
- (IV) Suppose that $n > 2$ and that $a, b \in D$ are introduced in C_i and C_j , respectively, where $i \leq j$. Let P be the set of all inner paths from a to b .
 - ♦ Every $p \in P$ visits each index between i and j .
 - ♦ Some $p \in P$ does not visit any index less than $i-1$ or greater than $j+1$.

(I) The fragment $\mathfrak{D}[F_i]$ consists of the substructure \mathfrak{C}_i of \mathfrak{D} induced by the set C_i plus one adhesion edge e . By inductive hypothesis, the substructure \mathfrak{C}_i has a definable tree-decomposition $(W_s^i)_{s \in S_i}$ of width at most kh . We obtain a tree-decomposition $(W_s^i \cup e)_{s \in S_i}$ of $\mathfrak{D}[F_i]$ by adding all vertices from the edge e to every component. The resulting decomposition is definable and its width is at most $kh + |e| \leq kh + k = k(h+1)$.

(III) For a contradiction, suppose that $a \in D$ belongs to at least 3 components C_i . By the connectedness condition of a tree-decomposition, it follows that we can find an index i such that $a \in C_{i-1} \cap C_i \cap C_{i+1}$. Consequently, we have $a \in U_s \cap U_t$ where s is the first vertex in F_i and t is the last one. Let c be the colour of a in the local colouring. Then the profile of F_i contains the information that the c -coloured vertices in the first and last components are the same. Since $n > 2$, the profiles of all factors $F_j \cup \dots \cup F_m$ are equal and idempotent. This implies that the c -coloured vertices of the first and last components of every F_j are equal to a . Hence, $a \in U_s$, for every $s \in I$. But D contains no ubiquitous elements. A contradiction.

(IV) We start with the first claim. For a contradiction, suppose that there is some index $i \leq l \leq j$ not visited by p . Let $\langle c, d \rangle$ be the last edge of the path p with $c \in C_0 \cup \dots \cup C_{l-1}$, and let s be the smallest index with $d \in C_s$. By choice of c we have $s \geq l+1$. As the hyperedge containing c and d is

covered by some component U_w , it follows that $c \in C_t$, for some $t \geq s$. By the connectedness condition of a tree-decomposition, this implies that $c \in C_r$, for all $l-1 \leq r \leq t$. But these are $t - (l-1) + 1 \geq l+1 - (l-1) + 1 = 3$ indices. A contradiction to (III).

For the second claim, fix some path $p \in P$. We can factorise p into parts of three types: (1) paths contained in $C_0 \cup \dots \cup C_{i-1}$ and starting and ending in some vertex of C_{i-1} , (2) paths contained in $C_i \cup \dots \cup C_j$, and (3) paths contained in $C_{j+1} \cup \dots \cup C_{n-1}$ and starting and ending in some vertex of C_{j+1} . Since the profiles of F_{i-1} and of $F_0 \cup \dots \cup F_{i-1}$ are the same, we can replace every segment of p of type (1) by a path entirely contained in C_{i-1} (without changing the end-points). Similarly, we can replace every segment of type (3) by a path entirely contained in C_{j+1} . The resulting path has the desired property.

(II) If there are only two factors, we can define the path-decomposition using two unary predicates, one for each component. Hence, we may assume that $n > 2$. As explained above, this means that the profiles of the decompositions $(U_s)_{s \in F_i}$ are the same and idempotent. We use 5 unary predicates P_0, \dots, P_4 where P_j contains all elements introduced in C_i with $i \equiv j \pmod{5}$. We call j the *sequential colour* of these vertices (to distinguish these colours from those given by the local colouring). Below we will prove the following two claims.

- (II.1) Two elements a, b are introduced in the same component C_i if, and only if, they have the same sequential colour and there exists a path from a to b that does not contain all colours.
- (II.2) To elements a and b are introduced in consecutive components if, and only if, a has sequential colour j , b sequential colour $j+1 \pmod{5}$, and there exists a path from a to b that does not contain all colours.

Since the introduction ordering is the transitive closure of these two relations, it can therefore be defined with the help of the parameters P_0, \dots, P_4 .

(II.1) (\Leftarrow) follows from the first claim in (IV): if a and b have the same sequential colour but are introduced in C_i and C_j with $i < j$, then $j \geq i+5$ which implies that every path between them visits every sequential colour.

(\Rightarrow) follows from the second claim in (IV): if $a, b \in C_i$ there is some path

between them containing only the colours $i-1, i, i+1 \pmod{5}$.

(II.2) Is proved in the same way except that, for (\Rightarrow), the path can now contain colours $i-1, i, i+1, i+2$. \square

Unbounded Path-Width

For the general case, we will decompose every tree-decomposition into a family of path-decompositions. Then we can use the result from the previous section to define tree-decompositions for each part, and we can fuse them into a single one with the help of Proposition 5.9. We start with a few technical lemmas about paths.

Lemma 5.15. *Let \mathcal{P} be a set of paths in a structure \mathcal{A} such that no element $a \in A$ belongs to more than k paths from \mathcal{P} . There exists a GSO-formula $\varphi(x, y; \bar{Q})$ (with parameters \bar{Q}) such that*

$$\mathcal{A} \models \varphi(a, b; \bar{Q}) \quad \text{iff} \quad a \text{ and } b \text{ are the end-points of some path in } \mathcal{P}.$$

This formula only depends on the signature Σ and the constant k .

Proof. By assumption we can choose colourings $\mu_p : \text{dom}(p) \rightarrow [k]$, for $p \in \mathcal{P}$, such that

$$a \in \text{dom}(p) \cap \text{dom}(q) \quad \text{implies} \quad \mu_p(a) \neq \mu_q(a).$$

(Just fix a well-ordered enumeration $(p_i)_{i < \alpha}$ of \mathcal{P} and choose μ_{p_i} by induction on i .) We consider the relations

$$\begin{aligned} E_{ij} &:= \{ \langle a, b \rangle \mid a \text{ and } b \text{ are consecutive vertices on some } p \in \mathcal{P} \\ &\quad \text{with } \mu_p(a) = i \text{ and } \mu_p(b) = j \}, \\ E &:= \{ \langle \langle a, i \rangle, \langle b, j \rangle \rangle \mid \langle a, b \rangle \in E_{ij} \}. \end{aligned}$$

Note that the relations E_{ij} are already guarded, so we can use them as parameters. Furthermore, the transitive reflexive closure E^* of E can be defined by

$$\psi_{ij}(x, y) := \forall \bar{Z} \left[Z_i x \wedge \forall uv \bigwedge_{m,n} [Z_m u \wedge E_{mn} uv \rightarrow Z_n v] \rightarrow Z_j y \right].$$

It follows that we can construct the desired formula φ by expressing that

- ♦ there are $i, j < k$ such that $\langle \langle x, i \rangle, \langle y, j \rangle \rangle \in E^*$,
- ♦ there are no z and l such that $\langle z, x \rangle \in E_{li}$,
- ♦ there are no z and l such that $\langle y, z \rangle \in E_{jl}$.

□

Lemma 5.16. *Let $(U_t)_{t \in T}$ be a reduced internally connected tree-decomposition of \mathfrak{A} of width k and let $a, b \in U_t$ be two elements in the same component. There exists a factor F with root t such that*

- ♦ $\text{pwd } \mathfrak{A}[F] \leq 2k$
- ♦ a and b are connect in $\mathfrak{A}^*[F]$ by two paths p and q (not necessarily distinct) such that each border adhesion edge is used by at most one them.

Proof. Given a factor F of T and two sets $X, Y \subseteq \bigcup_{s \in F} U_s$, we call a hyperedge e X - Y -separating if every path between X and Y in $\mathfrak{A}^*[F]_{\text{in}}$ goes through the vertex e . We say that a path-decomposition of $\mathfrak{A}^*[F]$ is X - Y -separating if


$X \subseteq W_0$, $Y \subseteq W_{n-1}$, and, for every X - Y -separating adhesion edge e , there is some index i with $W_i = e$ and $W_{i-1} \cap W_{i+1} = \emptyset$.

We split the proof into the following four steps.

- (I) For every factor F , the fragment $\mathfrak{A}^*[F]$ is connected.
- (II) For every $s \in T$ and all $X, Y \subseteq U_s$, there exists an X - Y -separating path-decomposition V_0, \dots, V_n of $\mathfrak{A}^*[s]$ of width at most k .
- (III) There exists a factor F with root t such that $\mathfrak{A}[F]$ has an a - b -separating path-decomposition $(W_i)_{i < n}$ of width at most $2k$.
- (IV) There exists a factor F as in the statement of the lemma.

(I) Let r be the root of F and let S the subtree of T rooted at r . Since $(U_t)_{t \in T}$ is reduced and internally connected, the substructure $\mathfrak{A}^*[S]$ contains some path p_0 between X and Y that does not contain any elements from $\text{Ad}(r)$, except possibly for the end-points. Replacing in p_0 every sub-path outside of $\mathfrak{A}^*[F]$ by the corresponding border adhesion edge, we obtain a path p in $\mathfrak{A}^*[F]$ between a and b .

(II) If $\mathfrak{A}^*[t]$ does not contain any X - Y -separating adhesion edges, we can take the decomposition consisting of a single component $V_0 := U_s$. Otherwise, we use (I) to choose a path p in $\mathfrak{A}^*[t]$ between X and Y . Let e_0, \dots, e_{m-1} be an enumeration of all X - Y -separating adhesion edges, enumerated in the order they appear on p . (Note that this enumeration does not depend on the choice of p .) For odd indices, we set $V_{2j+1} := e_j$ while, for the even components V_0, V_2, \dots, V_{2m} , we choose the connected components of the graph obtained from U_s by removing the edges e_0, \dots, e_{m-1} . (For V_0 and V_{2m} , we take the union of all connected components containing elements from, respectively, X and Y .)



Then $X \subseteq V_0$ and $Y \subseteq V_{2m}$. Furthermore, $|V_i| \leq |U_s| \leq k$. Hence, V_0, \dots, V_{2m} is the desired path decomposition.

(III) Let $(V_i)_{i < n}$ be the path-decomposition from (II) for $s := t$, $X := \{a\}$, and $Y := \{b\}$. Setting $W_i := V_i \cup \text{Ad}(t)$, we obtain a path-decomposition $(W_i)_{i < n}$ of $\mathfrak{A}[t]$ with the desired properties.

(IV) Since the set of factors with the property from (III) is closed under unions of chains, there exists a maximal such factor F . We claim that $\mathfrak{A}^*[F]$ contains two paths between a and b such that no adhesion edge is used by both of them. (Note that, except for the adhesion edges, both paths might intersect or even be identical.) For a contradiction, suppose otherwise. Then, for every pair of paths from a to b in $\mathfrak{A}^*[F]$, there is some adhesion edge that is used by both of them. By the Theorem of Menger (applied to the graph $\mathfrak{A}^*[F]_{\text{in}}$), it then follows that there exists an a - b -separating hyperedge e that is an adhesion edge. Let $s \in T$ be the vertex represented by e . It is sufficient to prove that the factor $F' := F \cup \{s\}$ also satisfies the above conditions (a contradiction to the maximality of F). By assumption, there is some index i such that $W_i = e$ and $W_{i-1} \cap W_{i+1} = \emptyset$. Let V_0, \dots, V_m be the

path-decomposition of $\mathfrak{A}^*[s]$ from (II) for $X := e \cap W_{i-1}$ and $Y := e \cap W_{i+1}$. Then

$$W_0, \dots, W_{i-1}, V_0, \dots, V_m, W_{i+1}, \dots, W_{n-1}$$

is a path-decomposition of $\mathfrak{A}[F']$. To see that it is a - b -separating, let f be an a - b -separating adhesion edge. If f is an adhesion edge in $\mathfrak{A}[F]$, we can find the desired component W_j with $W_j = f$ by choice of F . (Note that $f \neq e$ implies $j \neq i$.) If, on the other hand, f is a new adhesion edge in $\mathfrak{A}^*[s]$, the claim follows by construction of V_0, \dots, V_m . \square

Lemma 5.17. *Let $(U_t)_{t \in T}$ be a reduced internally connected tree-decomposition of \mathfrak{A} of width k . There exists a partition \mathcal{F} of T into factors with the following properties.*

- (a) $\text{pwd } \mathfrak{A}[F] \leq 2k$, for all $F \in \mathcal{F}$.
- (b) *There exists a GSO-formula $\varphi(x, y; \bar{P})$ (with GSO-parameters \bar{P}) that only depends on k and the signature Σ such that*

$$\mathfrak{A} = \varphi(a, b; \bar{P}) \quad \text{iff} \quad \begin{array}{l} \text{the factor } F \in \mathcal{F} \text{ introducing } a \text{ is less} \\ \text{than or equal (in the tree order) to the} \\ \text{factor introducing } b. \end{array}$$

Proof. Let \mathcal{F} be a maximal family of disjoint factors of T with the following properties.

- ♦ The union $\bigcup \mathcal{F}$ is a prefix of T .
- ♦ For every $F \in \mathcal{F}$, we have $\text{pwd } \mathfrak{A}[F] \leq 2k$.
- ♦ The fragment $\mathfrak{A}[\bigcup \mathcal{F}]$ contains a family \mathcal{P} of paths such that
 - (I) for every $F \in \mathcal{F}$ with root t there is some element a introduced in F that is connected to every $b \in U_t$ by some path in \mathcal{P} ;
 - (II) every element of $\mathfrak{A}[\bigcup \mathcal{F}]$ lies on at most $2k^3 + k$ paths from \mathcal{P} ;
 - (III) every border adhesion edge of $\mathfrak{A}[\bigcup \mathcal{F}]$ is used by at most $2k^3$ paths from \mathcal{P} .

To see that such a family exists, take $\mathcal{F} = \{F\}$ where $F = \{s\}$ consists of just the root of T . Then $\text{pwd } \mathfrak{A}[s] \leq |U_s| \leq k$. Furthermore, since $(U_t)_t$ is reduced and internally connected, $\mathfrak{A}[s]$ is connected. Hence, we can pick some element $a \in U_s$ and, for every $b \in U_s$, some path p_b from b to a . Then $\mathcal{P} := \{p_b \mid b \in U_s\}$ is a family of paths satisfying (I)–(III).

Thus, there exists a maximal such family \mathcal{F} . We claim that it has the desired properties. By choice of \mathcal{F} , we know that every factor has path-width at most $2k$.

To see that the ordering between the factors is definable, we use the family \mathcal{P} from above. For every $F \in \mathcal{F}$, fix an element a_F introduced in F that is connected to every element of the root component of F by some path in \mathcal{P} . We know by Lemma 5.15 that we can define the relation

$$\{ \langle a, b \rangle \mid \mathcal{P} \text{ contains a path between } a \text{ and } b \}.$$

(We will show below that $\bigcup \mathcal{F} = T$. Hence, the paths cannot contain any adhesion edges.) Consequently, we can also define

$$R := \{ \langle a, b \rangle \mid a = a_F \text{ and } b \in \text{Ad}(t) \text{ for some } F \in \mathcal{F} \text{ with root } t \}.$$

Let $F \in \mathcal{F}$ with root t . As $(U_s)_s$ is internally connected, the fragment associated with the subtree S attached to t is connected. Since a_F does not belong to $\text{Ad}(t)$, it follows that an element b is introduced in S if, and only if, there is some path between a_F and b that contains no vertex in $\text{Ad}(t)$. This property can be expressed in GSO using the definable relation R above. Consequently, we can also define the relation

$$S := \{ \langle a, b \rangle \mid \text{there is some } F \in \mathcal{F} \text{ with root } t \text{ such that } a = a_F \text{ and } b \text{ is introduced in the subtree attached to } t \}.$$

It follows that we can order the factors by the formula

$$\varphi(x, y) := (\exists v. Svy)(\forall u. Sux) \forall z [Svz \rightarrow Suz].$$

To finish the proof it remains to show that \mathcal{F} is indeed a partition of T , i.e., that $\bigcup \mathcal{F} = T$. For a contradiction, suppose otherwise and fix a minimal

vertex $t \in T \setminus \bigcup \mathcal{F}$. Let e be the adhesion edge of $\mathfrak{A}[\bigcup \mathcal{F}]$ representing t . We define the e -trace of a path $p \in \mathcal{P}$ to be the set of all edges $\langle a, b \rangle$ of p such that $a, b \in e$. Let us call a path p *minimal* if it cannot be shortened by removing some of its vertices without changing its end-points or it becoming disconnected. Note that, by removing some vertices if necessary, we may assume that every path in \mathcal{P} is minimal without violating (i)–(iii). This implies that the e -trace of every path in \mathcal{P} contains at most one edge.

We pick two vertices $a, b \in e$ such that the number m of paths $p \in \mathcal{P}$ whose trace contains $\langle a, b \rangle$ is maximal. By Lemma 5.16, there exists a factor F' with root t such that $\text{pwd } \mathfrak{A}[F'] \leq 2k$ and the fragment $\mathfrak{A}^*[F']$ contains two paths between a and b such that every border adhesion edge of $\mathfrak{A}[F']$ lies on at most one of them. We claim that $\mathcal{F} \cup \{F'\}$ satisfies the above conditions, which contradicts the maximality of \mathcal{F} . Clearly, $\bigcup \mathcal{F} \cup F'$ is a prefix of T and $\text{pwd } \mathfrak{A}[F'] \leq 2k$. Hence, we only have to check the third condition.

We fix an element c introduced in U_t . For every $d \in U_t$, we choose a path between c and d in $\mathfrak{A}^*[\{t\}]$. Let \mathcal{Q}_0 be the set of these paths. We cannot simply use the union $\mathcal{P} \cup \mathcal{Q}_0$, since the paths in \mathcal{P} might use the adhesion edge e , which does not exist in $\mathfrak{A}[\bigcup \mathcal{F} \cup F']$. Therefore, we have to modify the paths in \mathcal{P} .

Let \mathcal{Q}_1 be the set of paths $p \in \mathcal{P}$ that do not use the edge e . The other paths are modified as follows. For paths $p \in \mathcal{P}$ whose e -trace $\langle c, d \rangle$ is different from $\langle a, b \rangle$, we replace the edge $\langle c, d \rangle$ by some path in $\mathfrak{A}^*[F']$ that connects c and d . (Such a path exists, since $(U_s)_s$ is reduced and internally connected.) Let \mathcal{Q}_2 be the set of paths obtained this way. For paths $p \in \mathcal{P}$ with e -trace $\langle a, b \rangle$ we proceed as follows. We use Lemma 5.16 to find two paths q_0 and q_1 between a and b in $\mathfrak{A}^*[F']$ such that every border adhesion edge is used by at most one of them. In half of the paths, we replace the edge $\langle a, b \rangle$ by q_0 and in the other half, we replace it by q_1 . Let \mathcal{Q}_3 be the set of paths obtained this way. It follows that, by choice of a and b , every adhesion edge in $\mathfrak{A}^*[F']$ is used by at most $\lceil m/2 \rceil$ of the paths in \mathcal{Q}_3 .

We claim that $\mathcal{Q} := \mathcal{Q}_0 \cup \mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \mathcal{Q}_3$ is the desired set of paths for $\mathcal{F} \cup \{F'\}$. Condition (i) is satisfied by choice of \mathcal{Q}_0 .

For (ii), let $c \in \mathfrak{A}[\bigcup \mathcal{F} \cup \{F'\}]$. If $c \in \mathfrak{A}[\bigcup \mathcal{F}]$, it belongs to at most $2k^3 + k$ paths from \mathcal{P} by (ii). Since every path $q \in \mathcal{Q}$ belongs to $\mathfrak{A}[F']$ or is

obtained from some $p \in \mathcal{P}$ by inserting subpaths that entirely lie in $\mathfrak{A}[F']$, it follows that such an element also belongs to at most $2k^3 + k$ paths from \mathcal{Q} . If c is introduced in F' it belongs to at most $|\mathcal{Q}_0| \leq k$ paths in \mathcal{Q}_0 , and to at most $2k^3$ paths from \mathcal{P} by (iii). This implies that it belongs to at most $2k^3$ paths from $\mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \mathcal{Q}_3$.

Finally, for (iii), let f be an adhesion edge in $\mathfrak{A}[\bigcup \mathcal{F} \cup \{F'\}]$. If f is an adhesion edge in $\mathfrak{A}[\bigcup \mathcal{F}]$, it is contained in at most $2k^3$ paths from \mathcal{P} and, hence, also in at most that many paths from \mathcal{Q} . If f belongs to $\mathfrak{A}[F']$, it is used by at most $|\mathcal{Q}_0| \leq k$ paths from \mathcal{Q}_0 , it is not used by any path in \mathcal{Q}_1 , and, as we have seen above, by at most $\lceil m/2 \rceil$ paths from \mathcal{Q}_3 . Let n be the number of paths in \mathcal{Q}_2 using f . We distinguish two cases.

If $m \geq 2k$, we have $\lceil m/2 \rceil + k \geq 2k \leq m$. Hence, $k + n + \lceil m/2 \rceil \leq n + m$ which is at most the number of paths in \mathcal{P} using the adhesion edge e . By (iii), this number is bounded by $2k^3$.

If $m < 2k$, there are at most $m \leq 2k - 1$ paths in \mathcal{P} with any given e -trace $\langle c, d \rangle$, and there are at most $|\text{Ad}(t)|^2 \leq k^2$ possible non-empty e -traces. Since every path in $\mathcal{Q}_2 \cup \mathcal{Q}_3$ using f is obtained from some path in \mathcal{P} with non-empty e -trace, follows that there are at most $k^2(2k - 1)$ such paths. Since $k^2(2k - 1) + k = 2k^3 - k^2 + k \leq 2k^3$, the claim follows. \square

Proof of Theorem 5.5. Let \mathfrak{A} be a structure of tree-width k . By Lemma 5.4 every connected component C of \mathfrak{A} has a reduced internally connected tree-decomposition of width k . Let \mathcal{F} be the factorisation of this decomposition from Lemma 5.17. This factorisation induces a tree-decomposition of C that, according to Lemma 5.17 (b), is definable. Furthermore, by (a), each fragment of this decomposition has path-width at most $2k$. Hence, we can use Proposition 5.14 to find definable tree-decompositions for every fragment $\mathfrak{A}[F]$, and Proposition 5.9 then produces a definable tree-decomposition of C . Finally, note that the decomposition of \mathfrak{A} into its connected components can be considered as the introduction ordering of a tree-decomposition (with an empty root and one leaf for every connected component). Since this decomposition is definable, we can use Proposition 5.9 once more to obtain the desired definable tree-decomposition of \mathfrak{A} . \square

Notes

Monadic second-order transductions here were introduced by Courcelle and Engelfriet in [59]. The Transduction Hierarchy is from [23].

Theorem 2.6 was originally announced by Lapoir [128], but he was never able to provide a satisfactory proof. The first complete proof is due to Bojańczyk and Pilipczuk [29]. The original proof shows how to define tree-decompositions whose width might be larger than the tree-width of the given structure. In a subsequent article [30] the same authors show how to transform such a decomposition into one of minimal width.

The fact that GSO collapses to FO over classes of bounded height- n tree-width was first noted in [77].

Corollary 4.8 was proved by Rabinovitch [157] using different methods.

Part C

Applications

XI The Caucal Hierarchy

1 Finitely Presentable Structures

WHEN DEVELOPING ALGORITHMS THAT work with infinite structures, the main problem to be solved is how to encode such a structure in a finite way, so we can input it to the algorithm. As every kind of encoding can represent only countably many objects, there is no single encoding that works for all countable structures. That means that, for every application, one has to find a representation that works well for the given task. In particular, (I) every structure of interest must have an encoding and (II) the problem we want to solve for these structures needs to be decidable when using this representation. The problem we are mostly interested in here is the model-checking problem: given a structure \mathcal{A} and a formula φ , decide whether or not $\mathcal{A} \models \varphi$. Many algorithmic problems can be phrased in this way, depending on which logic we use.

One quite general way to obtain such representations is to start with a set of basic structures with certain desirable properties, like a decidable model-checking problem, and to close it under operations preserving these properties. Then every structure in the resulting class has the desired properties by design. Furthermore, each of them also has a finite representation: a finite term over the chosen operations. Our main examples will consist of operations compatible with some logic like FO or MSO that are applied to some canonical base structure like the infinite binary tree.

When constructing representations using this recipe, the algorithmic properties of the resulting structures usually follow by the choice of the representation. But it can be quite hard to characterise *which* structures have a representation of this kind. This question is not purely theoretical

since, when choosing a representation for a given application, one needs to make sure that all the structures occurring as possible inputs can actually be encoded.

Tree-Interpretable Structures

Our first example is a class of simple structures with a decidable MSO-theory.

Definition 1.1. We call a structure \mathfrak{A} *tree-interpretable* if there exists an MSO-interpretation τ mapping the infinite binary tree $\mathfrak{T}_{\text{bin}}$ to \mathfrak{A} . J

We can represent each such structure by the corresponding interpretation τ . Furthermore, given τ and an MSO-formula φ , we can decide whether $\mathfrak{A} \models \varphi$ by checking whether $\mathfrak{T}_{\text{bin}} \models \varphi^\tau$.

Examples. (a) For every $0 < k < \omega$, the k -ary infinite tree

$$\mathfrak{T} = \langle [k]^*, \leq_{\text{pf}}, \leq_{\text{lex}}, (\text{suc}_i)_{i < k} \rangle$$

is tree-interpretable.

(b) We have shown in Theorem VI.4.10 that a linear order is tree-interpretable if, and only if, it is regular. Similarly, Theorem ?? tells us that a tree is tree-interpretable if, and only if, it is regular. J

We can use Theorem V.1.13 to replace the MSO-interpretation representing a tree-interpretable structure by an FO-interpretation.

Theorem 1.2. A structure \mathfrak{A} is tree-interpretable if, and only if, it is FO-interpretable in $\mathfrak{T}_{\text{bin}} = \langle \{0, 1\}^*, \leq_{\text{pf}}, \text{suc}_0, \text{suc}_1 \rangle$.

Proof. By Corollary V.1.14, we can write

$$\mathfrak{A} \cong \sigma(\rho(\mathfrak{T}_{\text{bin}})),$$

for an FO-interpretation σ and an MSO-colouring ρ . Let C be the set of colours used by ρ and let $\varphi_c(x)$ be the MSO-formula defining the colour $c \in C$. By Lemma 2.3, the set $K_c := \varphi_c^{\mathfrak{T}_{\text{bin}}} \subseteq \{0, 1\}^*$ of all vertices with colour c

forms a regular language. Fix a homomorphism $\eta : \{0, 1\}^* \rightarrow M$ to a finite monoid recognising all languages K_c , for $c \in C$. Let $P_c \subseteq M$ be the subset such that $K_c = \eta^{-1}[P_c]$ and let $S \subseteq M^*$ be the set of all sequences $a_0 \cdots a_{n-1}$ such that

$$a_0 = \eta(c_0) \quad \text{and} \quad a_{i+1} := a_i \cdot \eta(c_{i+1}), \quad \text{for some } c_i \in \{0, 1\}.$$

Since this condition is local, there exists an FO-formula defining the set S in the tree

$$\mathfrak{U} := \langle M^*, \leq_{\text{pf}}, (\text{suc}_a)_{a \in M} \rangle.$$

Let L_c be the set of all sequences $a_0 \cdots a_{n-1} \in S$ such that $a_{n-1} \in P_c$. It follows that the subtree

$$\mathfrak{S} := \langle S, \leq_{\text{pf}}, (\text{suc}_a)_{a \in M}, (L_c)_c \rangle$$

of \mathfrak{U} is isomorphic to $\langle \mathfrak{T}_{\text{bin}}, (K_c)_c \rangle$. Furthermore, each predicate L_c is FO-definable in \mathfrak{U} . Consequently, there exists an FO-interpretation τ of \mathfrak{S} in \mathfrak{U} . Combined with the interpretation σ from above, we obtain an FO-interpretation of \mathfrak{A} in \mathfrak{U} .

Hence, it remains to see that \mathfrak{U} is FO-interpretable in $\mathfrak{T}_{\text{bin}}$. Fix an enumeration a_0, \dots, a_{m-1} of M and let $g : M \rightarrow \{0, 1\}^{m+2}$ be the function with $g(a_i) = 0^{i+1}1^{m-i+1}$. We extend g to a homomorphism $M^* \rightarrow \{0, 1\}^*$. The existence of the desired interpretation now follows from the fact that the image of M^* under this homomorphism is FO-definable in $\mathfrak{T}_{\text{bin}}$. \square

In the sections below we will prove several characterisations of tree-interpretable structures. For convenience, let us summarise them here. All undefined terminology will be introduced later on.

Theorem 1.3. Let \mathfrak{A} be a Σ -structure. The following statements are equivalent.

- (1) \mathfrak{A} is tree-interpretable.
- (2) \mathfrak{A} is FO-interpretable in $\mathfrak{T}_{\text{bin}}$ (with tree-order).
- (3) \mathfrak{A} is prefix-recognisable.

(4) \mathfrak{A} is VR-equational.

The proofs can be found in Theorems 1.2, 2.2 and 3.1 below. For graphs, we have two additional characterisations.

Theorem 1.4. Let $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ be a directed edge-labelled graph. The following statements are equivalent.

- (1) \mathfrak{G} is tree-interpretable.
- (2) \mathfrak{G} is FO-interpretable in $\mathfrak{T}_{\text{bin}}$ (with tree-order).
- (3) \mathfrak{G} is prefix-recognisable.
- (4) \mathfrak{G} is VR-equational.
- (5) $\mathfrak{G} \cong \alpha(\mathfrak{T}_{\text{bin}})$, for regular path-interpretation α .
- (6) \mathfrak{G} is an ε -pushdown graph.

The respective proofs are given by Theorems 4.2 and 5.3.

The Caucal Hierarchy

We can obtain a larger class with decidable MSO-model checking by adding a second MSO-compatible operation: the Muchnik iteration.

Definition 1.5. The *Caucal hierarchy* is the sequence

$$\text{CH}_0 \subseteq \text{CH}_1 \subseteq \text{CH}_2 \subseteq \dots,$$

where CH_0 is the class of all finite structures and CH_{n+1} consists of all structures of the form $\tau(\mathfrak{A}^*)$, where $\mathfrak{A} \in \text{CH}_n$ and τ is an MSO-interpretation. We denote the union of this hierarchy by CH .

Examples. (a) Let \mathfrak{C} be a structure with a single element and no relations. Then

$$\mathfrak{C}^* \cong \langle \omega, \text{suc}, P \rangle, \quad \text{where} \quad P := \{ n < \omega \mid n > 1 \}.$$

Applying an interpretation that forgets the predicate P , it follows that

$$\langle \omega, \text{suc} \rangle \in \text{CH}_1.$$

(b) We have $\langle \omega, \text{suc} \rangle^* \cong \langle \omega^*, \text{suc}, \text{suc}', P \rangle$, where

$$\text{suc} := \{ \langle wk, w(k+1) \rangle \mid w \in \omega^*, k < \omega \},$$

$$\text{suc}' := \{ \langle w, wk \rangle \mid w \in \omega^*, k < \omega \},$$

$$P := \{ wkk \mid w \in \omega^*, k < \omega \}.$$

In this structure we can define the order $\leq := \text{suc}^*$, the transitive reflexive closure of suc , the set


$$T := \{ \langle k, n \rangle \in \omega^2 \mid n \leq k \},$$

and the edge relation

$$E := \{ \langle \langle k, 0 \rangle, \langle k+1, 0 \rangle \rangle \in T \times T \mid k < \omega \}$$

$$\cup \{ \langle \langle k, n+1 \rangle, \langle k, n \rangle \rangle \in T \times T \mid n < k \}.$$

Consequently, the tree $\mathfrak{T} := \langle T, E \rangle$ can be interpreted in $\langle \omega, \text{suc} \rangle^*$, which implies that $\mathfrak{T} \in \text{CH}_2$. Note that $\mathfrak{T} \notin \text{CH}_1$.



(c) For $n < \omega$, we consider the trees

$$\mathfrak{T}_n := \langle T_n, \text{suc}_a, \text{suc}_b \rangle \quad \text{and} \quad \mathfrak{S}_n := \langle S_n, \text{suc}_a, \text{suc}_o, \text{suc}_1 \rangle$$

with domains

$$T_n := \{ a^i b^j \mid i < \omega, j < \exp_n(i) \},$$

$$S_n := \{ a^i w \mid i < \omega, w \in \{0, 1\}^{<i} \},$$

where \exp_n denotes the n -fold iteration of the exponentiation function

$$\exp_0(k) := k \quad \text{and} \quad \exp_{i+1}(k) := 2^{\exp_i(k)}.$$

We claim that

$$\mathfrak{T}_n \in \text{CH}_{n+2} \quad \text{and} \quad \mathfrak{S}_n \in \text{CH}_{n+3}.$$

The proof proceeds by induction on n . By (b), we have $\mathfrak{T}_0 \in \text{CH}_2$. For the inductive step, note that we can obtain \mathfrak{S}_n from \mathfrak{T}_n by applying an interpretation that replaces every b -edge by two parallel edges with labels 0 and 1, respectively, and then unravelling the resulting graph. As the unravelling is interpretable in the Muchnik iteration, it follows by inductive hypothesis that

$$\mathfrak{S}_n \in \text{CH}_{n+3}.$$

Finally, we can obtain \mathfrak{T}_{n+1} from \mathfrak{S}_n by an MSO-interpretation that orders each subtree of the form $\{ a^i w \mid w \in \{0, 1\}^{<i} \}$ lexicographically and then replaces the ordering by the corresponding successor relation. It follows that $\mathfrak{T}_{n+1} \in \text{CH}_{n+3}$. \square

As both operations used to generate the Caucal hierarchy are CMSO-compatible, we immediately obtain the following decidability result.

Theorem 1.6. *Every structure in the Caucal hierarchy has a decidable CMSO-theory.*

Remark. In the example on page 804 below, we will present a tree with a decidable CMSO-theory that does not belong to the Caucal hierarchy. \square

The following normal form is frequently helpful. We denote the n -fold iteration of the Muchnik iteration by \mathfrak{Q}^{*n} , that is $\mathfrak{Q}^{*0} := \mathfrak{Q}$ and $\mathfrak{Q}^{*(n+1)} := (\mathfrak{Q}^{*n})^*$.

Lemma 1.7. *A structure belongs to CH_n if, and only if, it is of the form $\tau(\mathfrak{C}^{*n})$, for some MSO-interpretation τ and a finite structure \mathfrak{C} .*

Proof. We have shown in Lemma V.3.12 that MSO-interpretations and Muchnik iterations commute (in one direction). Consequently, the claim follows by induction on the level n . \square

We can use this normal form to compare the structures in CH with tree-interpretable structures. As it turns out, the only difference is that we are allowed to use a certain non-regular colouring of the binary tree.

Lemma 1.8. *Let \mathfrak{C} be a finite structure and $0 < n < \omega$. There exists an MSO-interpretation τ and (non-regular) sets $P_1, \dots, P_{n-1} \subseteq \{0, 1\}^*$ such that*

$$\mathfrak{C}^{*n} \cong \tau(\mathfrak{T}_{\text{bin}}, \vec{P}).$$

Proof. Note that

$$\mathfrak{C}^{*n} = \langle C^{*n}, \text{succ}_0, \dots, \text{succ}_{n-1}, \text{cl}_0, \dots, \text{cl}_{n-1} \rangle.$$

We encode this structure in $\mathfrak{T}_{\text{bin}}$ as follows. We define an injective function $\mu_n : C^{*n} \rightarrow \{0, 1\}^*$ by induction on n . For $n = 0$, let c_0, \dots, c_{m-1} be an enumeration of C . We set

$$\mu_0(c_i) := 0^{i+1}.$$

For $n > 0$, we set

$$\mu_{n+1}(\langle w_0, \dots, w_{k-1} \rangle) := \mu_n(w_0) 1^{n+1} \mu_n(w_1) 1^{n+1} \dots 1^{n+1} \mu_n(w_{k-1}).$$

Let

$$A := 0 + 00 + \dots + 0^m \quad \text{and} \quad B_n := 1 + 11 + \dots + 1^n.$$

The range of the function μ_n is

$$\text{rng } \mu_n = (AB_n)^* A,$$

which is MSO-definable. Furthermore, under this encoding the successor functions are given by

$$\text{suc}_i = \{ \langle w, w1^{i+1}v \rangle \mid v \in (AB_{i+1})^* A \}.$$

These relations are also MSO-definable. Finally, the first clone relation

$$\text{cl}_0 = \{ w1u1u \mid u \in A, w \text{ is empty or it ends in } 1 \}$$

is also MSO-definable since the set A is finite. Consequently, the structure

$$\langle C^{*n}, \text{suc}_0, \dots, \text{suc}_{n-1}, \text{cl}_0 \rangle$$

is MSO-interpretable in $\mathfrak{T}_{\text{bin}}$. Adding the relations $\text{cl}_1, \dots, \text{cl}_{n-1}$ as additional monadic predicates \bar{P} , we obtain an interpretation of C^{*n} in $\langle \mathfrak{T}_{\text{bin}}, \bar{P} \rangle$. \square

Corollary 1.9. *Every $\mathfrak{A} \in \text{CH}$ is of the form $\mathfrak{A} \cong \tau(\mathfrak{T}_{\text{bin}}, \bar{P})$ for some MSO-interpretation τ and some monadic parameters \bar{P} .*

Corollary 1.10. *CH_1 is the class of all tree-interpretable structures.*

Most algorithmic properties of structures in the Caucal hierarchy follow directly from the definition. We have already seen that they have decidable CMSO-theories. It is also straightforward to determine the closure properties of CH.

Lemma 1.11. *Each level CH_n of the Caucal hierarchy is closed under*

- (1) CMSO-interpretations,
- (2) isomorphisms,
- (3) finite unions,
- (4) the copy operation,

- (5) definable expansions,
- (6) expansion by finitely many constants, and
- (7) substructures with definable universe.

Let us also establish closure of tree-interpretable structures under quotients by MSO-definable equivalence relations. The following lemma is a weak version of Lemma ?? below.

Lemma 1.12. *Let $\mathfrak{T}_{\text{bin}}$ be the binary tree, $K \subseteq \{0, 1\}^*$ a regular language, and $\approx \subseteq K \times K$ an MSO-definable equivalence relation. There exists a regular language $K_0 \subseteq K$ such that K_0 contains exactly one element of each \approx -class.*

Proof. Let p be the function mapping $x \in K$ to the infimum (with respect to \leq_{pf}) of its \approx -class $[x]_{\approx}$. Note that the function p is MSO-definable. Let m be the quantifier-rank of the formula defining \approx and k be the number of MSO $_m$ -theories. We claim that each class $[x]_{\approx}$ contains an element of length at most $|p(x)| + k$. This implies that the length-lexicographically least element of an \approx -class $[x]_{\approx}$ is definable. Hence, we can set

$$K_0 := \{ x \in K \mid x \text{ the } \leq_{\text{lex}}\text{-least element of } [x]_{\approx} \}.$$

To prove the above claim, consider elements $x_0, x_1 \in [x]_{\approx}$ such that $x_0 \sqcap x_1 = p(x)$. We choose x_0 such that $|x_0|$ is minimal. Let s be the suffix of x_0 such that $x_0 = p(x_0)s$. We claim that $|s| \leq k$. For a contradiction, suppose otherwise. Then there are two non-empty prefixes $z <_{\text{pf}} z' \leq_{\text{pf}} s$ such that

$$\langle \mathfrak{T}_{\text{bin}}, s \rangle|_{\upharpoonright z} \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, s \rangle|_{\upharpoonright z'}.$$

Replacing the subtree $\mathfrak{T}_{\text{bin}}|_z$ by $\mathfrak{T}_{\text{bin}}|_{z'}$, we obtain an element x'_0 such that $|x'_0| < |x_0|$ and

$$\langle \mathfrak{T}_{\text{bin}}, x_0 x_1 \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, x'_0 x_1 \rangle.$$

Consequently, $x'_0 \approx x_1$. A contradiction to our choice of x_0 . \square

Corollary 1.13. *Let \mathcal{A} be a tree-interpretable structure and \approx an MSO-definable equivalence relation on A . Then \mathcal{A}/\approx is tree-interpretable.*

While the algorithmic properties of structures in the Caucal hierarchy are relatively easy to determine, a harder problem is to decide whether a given structure belongs to the hierarchy. As already mentioned above such characterisations are also relevant for applications since, if we want to use tree-interpretations to encode infinite structures, we need to know that all the structures arising in practice possess such an encoding. For instance, it turns out that the state spaces of most programs are not in CH, which means that the results of this section are not that useful for software verification.

Not much is known about which structures belong to the Caucal hierarchy, in particular for the higher levels. But we can make two observations. First, all structures in CH have finite crossing-width, and second, the hierarchy does not contain all structures with a decidable MSO-theory.

Lemma 1.14. *Every structure in the Caucal hierarchy has finite crossing-width.*

Proof. We prove the claim by induction on the level. Clearly, every finite structure has finite crossing-width. Furthermore, it follows by Propositions IX.5.1 and IX.5.10 that, if a structure \mathcal{A} has finite crossing-width, then so does every structure MSO-interpretable in the Muchnik iteration of \mathcal{A} . \square

While structures in the Caucal hierarchy have a decidable MSO-theory, their GSO-theory is usually undecidable. The remaining ones can be characterised as follows.

Proposition 1.15. *Let $\mathcal{A} \in \text{CH}_n$. The following statements are equivalent.*

- (1) \mathcal{A} has a decidable CGSO-theory.
- (2) \mathcal{A} has a decidable GSO-theory.
- (3) $\mathcal{A}_{\text{in}} \in \text{CH}_n$.
- (4) \mathcal{A} has finite tree-width.
- (5) \mathcal{A} is k -sparse, for some $k < \omega$.
- (6) The Gaifman graph of \mathcal{A} does not contain the subgraph $\mathfrak{R}_{n,n}$, for some $n < \omega$.

Proof. (1) \Rightarrow (2) is trivial.

(2) \Rightarrow (4) Suppose that \mathcal{A} has infinite tree-width. Then Theorem VIII.4.1 implies that \mathcal{A} contains arbitrarily large grids as minors. Consequently, the GSO-theory of \mathcal{A} is undecidable.

(4) \Rightarrow (1) According to Lemma VIII.1.8 and Theorem VII.3.6, if \mathcal{A} has finite tree-width, it is sparse and its CGSO-theory collapses to the CMSO-theory. Hence, the claim follows by Theorem 1.6.

(4) \Leftrightarrow (5) \Leftrightarrow (6) follow by Theorem IX.4.5 as we have seen in Lemma 1.14 that \mathcal{A} has finite crossing-width.

(3) \Rightarrow (4) follows by Proposition IX.4.6 since \mathcal{A} has finite crossing-width by Lemma 1.14.

(5) \Rightarrow (3) By assumption, there exists an MSO-interpretation of \mathcal{A} in \mathfrak{C}^{*n} , for some finite structure \mathfrak{C} . Furthermore, we can obtain an MSO-interpretation of \mathcal{A}_{in} in \mathcal{A} by Theorem VII.3.6. Composing these we obtain our desired interpretation of \mathcal{A}_{in} in \mathfrak{C}^{*n} . \square

2 Prefix-Recognisable Structures

In this and the following sections we will derive several alternative characterisations of when a structure is tree-interpretable or when it belongs to the Caucal hierarchy. This will give us more precise tools to answer the question of whether a given structure is tree-interpretable or whether it belongs to CH. We start with a characterisation in terms of regular languages that is much more concrete than the interpretation-based one.

Definition 2.1. (a) Let Σ be an alphabet. The class of *prefix-recognisable* relations over Σ^* is inductively defined by the following closure properties.

- ♦ Every regular language $L \subseteq \Sigma^*$ is prefix-recognisable.
- ♦ If $L \subseteq \Sigma^*$ is regular and $R \subseteq \Sigma^* \times \cdots \times \Sigma^*$ is prefix-recognisable, then

$$LR := \{ \langle wu_0, \dots, wu_{n-1} \rangle \mid w \in L, \bar{u} \in R \}$$

is prefix-recognisable.

- ♦ If R and S are prefix-recognisable, then so are $R \times S$ and $R \cup S$.

- ♦ If R is a prefix-recognisable relation of arity n and $\sigma : [n] \rightarrow [n]$ is a permutation, then ${}^\sigma R := \{ \langle a_{\sigma(o)}, \dots, a_{\sigma(n-1)} \rangle \mid \bar{a} \in R \}$ is also prefix-recognisable.

(b) A structure \mathfrak{A} is *prefix-recognisable* if it is isomorphic to a structure of the form $\langle L, \bar{R} \rangle$ where L is a regular language and each relation R_i is prefix-recognisable. J

Example. For each $0 < k < \omega$, the infinite k -ary tree

$$\mathfrak{T} = \langle [k]^*, (\text{suc}_i)_{i < k}, \leq_{\text{pf}}, \leq_{\text{lex}} \rangle$$

is prefix-recognisable since

$$\begin{aligned} \text{suc}_i &= K^*(\langle \rangle \times i), \\ \leq_{\text{pf}} &= K^*(\langle \rangle \times K^*), \\ \leq_{\text{lex}} &= \leq_{\text{pf}} \cup \bigcup_{i < j} K^*(iK^* \times jK^*), \end{aligned}$$

where we have set $K := [k]$ and we write $\langle \rangle, i, j$ for the languages $\{\langle \rangle\}, \{i\}$, and $\{j\}$. J

Exercise 2.1. Prove that every binary prefix-recognisable relation R is of the form

$$R = \bigcup_{i < n} W_i(U_i \times V_i),$$

where $n < \omega$, and $U_i, V_i, W_i \subseteq \Sigma^*$ are regular languages. J

Remark. Given a binary prefix-recognisable relation

$$R = \bigcup_{i < n} W_i(U_i \times V_i)$$

and a pair $\langle u, v \rangle \in R$, we see that the word v is obtained from u by replacing some suffix. Hence, R can be interpreted as a certain kind of *suffix-rewriting system* where each rewriting rule is given by regular languages. In the article

where the term ‘prefix-recognisable’ was coined, the definition used the opposite ordering

$$R = \bigcup_{i < n} (U_i \times V_i) W_i.$$

Such a relation corresponds to a prefix-rewriting systems with recognisable rules. (‘recognisable’ is an alternative name for ‘regular.’) The modern ordering $\bigcup_i W_i(U_i \times V_i)$ has been chosen to better fit with our conventions regarding trees. Some authors have therefore changed the terminology to ‘suffix-recognisable’. J

The aim of this section is to prove the following characterisation of the class of tree-interpretable structures.

Theorem 2.2. *A structure \mathfrak{A} be a structure is tree-interpretable if, and only if, it is prefix-recognisable.*

The proof consists of the following technical lemma and the proposition below.

Lemma 2.3. *Let $\mathfrak{T}_{\text{bin}}$ be the infinite binary tree and θ an MSO_m -theory. The language*

$$K_\theta := \{ v \in \{0, 1\}^* \mid \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\{0, 1\}^* \setminus \uparrow\{v0, v1\}}, \downarrow v) = \theta \}$$

is regular.

Proof. Suppose that $v = d_0 \dots d_{n-1}$, set $c_i := 1 - d_i$, let

$$S := \{0, 1\}^* \setminus \uparrow\{v0, v1\},$$

and define

$$U_n := \{v\} \quad \text{and} \quad U_i := d_0 \dots d_{i-1}(\langle \rangle + c_i(0 + 1)^*), \quad \text{for } i < n.$$

Let \mathfrak{U}_u be the substructure of $\mathfrak{T}_{\text{bin}}$ induced by U_u . Then we can write $\mathfrak{T}_{\text{bin}}|_S$ as a generalised sum followed by a quantifier-free interpretation:

$$\mathfrak{T}_{\text{bin}}|_S \cong \sigma \left(\sum_{i \leq n} \mathfrak{U}_i \right).$$

Consequently, there exist MSO-formulae $\psi, \chi_0, \dots, \chi_{n-1}$ such that

$$\begin{aligned} \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_S, \Downarrow v) = \theta \quad \text{iff} \\ \langle [n+1], \leq \rangle \models \psi(\llbracket \chi_0 \rrbracket, \dots, \llbracket \chi_{n-1} \rrbracket). \end{aligned}$$

Note that $d_i = d_j$ implies $\mathbb{U}_j \cong \mathbb{U}_i$. Setting

$$P_0 := \{i < n \mid d_i = 0\} \quad \text{and} \quad P_1 := \{i < n \mid d_i = 1\}$$

it follows that each set $\llbracket \chi_j \rrbracket$ is of one of the following eight forms.

$$\begin{aligned} \emptyset, \quad \{n\}, \quad P_0, \quad \{n\} \cup P_0, \\ [n], \quad [n+1], \quad P_1, \quad \{n\} \cup P_1. \end{aligned}$$

Since each of these sets is definable in $\mathfrak{T} := \langle [n+1], \leq, P_0, P_1 \rangle$, we obtain a formula ψ' such that

$$\langle [n+1], \leq \rangle \models \psi(\llbracket \chi_0 \rrbracket, \dots, \llbracket \chi_{n-1} \rrbracket), \quad \text{iff} \quad \mathfrak{T} \models \psi'.$$

Note that \mathfrak{T} is obtained from the word structure for $v \in \{0, 1\}^*$ by adding one position at the end. Thus,

$$\mathfrak{T} = \tau(v \oplus 1),$$

for some interpretation τ , and we obtain a formula ψ'' such that

$$\mathfrak{T} \models \psi' \quad \text{iff} \quad v \models \psi''.$$

It follows that $K_\theta = \{v \mid v \models \psi''\}$. This language is MSO-definable and, therefore, regular. \square

The following proposition concludes the proof of Theorem 2.2.

Proposition 2.4. *Let $\mathfrak{T}_{\text{bin}} := \langle \{0, 1\}^*, \leq_{\text{pf}}, \text{suc}_0, \text{suc}_1 \rangle$ be the infinite binary tree and R a relation over $\{0, 1\}^*$. The following statements are equivalent.*

(1) R is prefix-recognisable.

(2) R is MSO-definable in $\mathfrak{T}_{\text{bin}}$.

(3) R is WMSO-definable in $\mathfrak{T}_{\text{bin}}$.

Proof. (3) \Rightarrow (2) is trivial.

(1) \Rightarrow (3) We prove the claim by induction on the term constructing R . First, suppose that $R \subseteq \{0, 1\}^*$ is a regular language. As R is regular, there exists an MSO-formula φ such that

$$w \models \varphi \quad \text{iff} \quad w \in R.$$

Furthermore, there exists an MSO-interpretation τ that, given a vertex $w \in \{0, 1\}^*$, maps $\langle \mathfrak{T}_{\text{bin}}, w \rangle$ to the word structure associated with w . This interpretation maps the formula φ to some MSO-formula $\psi(w) := \varphi^\tau$ which states that φ holds on the path from the root to w . It follows that ψ defines R .

For the inductive step note that, if R and S are defined by the formulae $\varphi(\bar{x})$ and $\psi(\bar{y})$, respectively, then $R \cup S$ is defined by

$$\vartheta(\bar{x}) := \varphi(\bar{x}) \vee \psi(\bar{x}),$$

$R \times S$ is defined by

$$\vartheta(\bar{x}, \bar{y}) := \varphi(\bar{x}) \wedge \psi(\bar{y}),$$

and ${}^\sigma R$ by

$$\vartheta(\bar{x}) := \varphi(\bar{x}'),$$

where $\bar{x}' = \langle x_{\sigma^{-1}(0)}, \dots, x_{\sigma^{-1}(n-1)} \rangle$ is the tuple obtained from \bar{x} by applying the permutation σ^{-1} . Finally, if R is 1-ary, we can define RS by


$$\vartheta(\bar{x}) := \exists z \left[\bigwedge_i z \leq_{\text{pf}} x_i \wedge \varphi(z) \wedge \psi'(\bar{x}; z) \right],$$

where the formula $\psi'(\bar{x}; z) := \psi^{(\uparrow z)}(\bar{x})$ is the relativisation of ψ to the subtree attached at z .

(2) \Rightarrow (1) To show that R is prefix-recognisable, we have to express it as a term over the operations $\cup, \times, {}^\sigma$, and concatenation. The idea of the

construction is quite simple, but making it precise turns out to be a bit technical. Hence, let us start with a bit of intuition. Given a tuple \vec{a} , we consider the closure S of \vec{a} under meets. This set forms a tree, for each edge of which we associate a regular language. We obtain the term for R by assembling these languages according to the structure of the tree. In general, R might contain tuples whose associated trees are different. In this case, we have to take a union of the expressions obtained for each tree.

For instance, for the tuple $\vec{a} = \langle \text{IIII}, \text{IOII}, \text{IOIOII}, \text{I} \rangle$, we obtain the following tree.



The components of \bar{a} are indicated by the brackets. Note that, strictly speaking, the root $\langle \rangle$ does not belong to S , but we still consider the edge leading to it (dotted in the figure). If U, V, W, X, Y are the languages associated with those edges, we obtain for R the expression

$$R = X(Y \times W(U \times V) \times \{\langle \rangle\}).$$

The formal construction is as follows. Let $\varphi(\vec{x})$ be an MSO-formula defining R , let m be its quantifier-rank, and let n be the arity of R . Given a tuple \vec{a} , let S be its closure under meets and let $\mathfrak{S} := \langle S, \leq_{\text{pf}}, \vec{a} \rangle$ be the tree induced by it. We have seen in Proposition V.1.11 that the theory $\text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, \vec{a})$ can be computed from (the isomorphism type) of \mathfrak{S} and the theories

$$\theta_s := \text{Th}_{\text{MSO}}^m(\langle T_s, \leq_{\text{pf}}, \Downarrow s \rangle), \quad \text{for } s \in S.$$

For a tree $\mathfrak{S} = \langle S, \leq_{\text{pf}}, \bar{s} \rangle$ and a family of theories $\bar{\theta}$, we denote by $R_{\mathfrak{S}, \bar{\theta}} \subseteq R$ the set of all tuples $\bar{a} \in R$ associated with \mathfrak{S} and $\bar{\theta}$. It follows that

$$R = \bigcup_{\mathfrak{S}, \bar{\theta}} R_{\mathfrak{S}, \bar{\theta}} ,$$

where the union ranges over all finite binary trees \mathfrak{S} equipped with a tuple \vec{s} of n constants such that every element is a meet of some components of \vec{s} , and over all families of theories $\hat{\theta} = (\theta_i)_{i \in \mathfrak{S}}$ such that the associated theory $\text{Th}_{\text{MSO}}^m(\mathfrak{S}_{\text{bin}}, \vec{a})$ contains the formula φ . Note that this union is finite. Since prefix-recognisable relations are closed under finite unions, it is therefore sufficient to prove that each relation $R_{\mathfrak{S}, \vec{s}, \hat{\theta}}$ is prefix-recognisable.

Hence, fix \mathfrak{S} and $\bar{\theta}$. We prove the claim by induction on the size $|S|$. Let r be the root of S . For $i \in \{0, 1\}$, we set

$$\mathfrak{S}_i := \langle S_i, \leq_{\text{pf}}, \bar{s}_i \rangle \quad \text{and} \quad \bar{\theta}_i := (\theta_t)_{t \in S_i},$$

where

$$S_i := \{ t \in S \mid \text{suc}_i(r) \leq_{\text{pf}} t \} \quad \text{and} \quad \bar{s}_i := \bar{s} \cap S_i,$$

By inductive hypothesis (applied to the subtrees $\mathfrak{T}_{\text{bin}}|_{r_0}$ and $\mathfrak{T}_{\text{bin}}|_{r_1}$, respectively), we know that the relations $R_0 := R_{\mathbb{C}_0, \hat{\theta}_0}$ and $R_1 := R_{\mathbb{C}_1, \hat{\theta}_1}$ are prefix recognisable. (If $S_i = \emptyset$, we take for R_i the \mathbf{o} -ary relation containing the empty tuple.) Given a theory θ , we denote by K_θ the regular language obtained by Lemma 2.3. We distinguish several cases.

First, suppose that $r \notin \bar{s}$. Then $r = s_i \sqcap s_j$, for some i, j , which implies that $S_0 \neq \emptyset$ and $S_1 \neq \emptyset$. It follows that

$$R_{\mathfrak{S}, \bar{\theta}} = K_{\theta_r}{}^\sigma(\circ R_o \times \mathbf{I} R_i),$$

where σ is the permutation induced by the bijection $\bar{s}_0 \bar{s}_1 \rightarrow \bar{s}$.

Next, suppose that $r \in \bar{s}$ and let k be the number of indices i with $s_i = r$. If $S_0 \neq \emptyset$ and $S_1 \neq \emptyset$, we obtain

$$R_{\mathfrak{S}, \bar{\theta}} = K_{\theta_r}{}^\sigma(oR_o \times {}_1R_I \times \langle \rangle \times \cdots \times \langle \rangle),$$

with k factors of the language $\{\langle \rangle\}$ and the permutation σ that is induced by the bijection $\bar{s}_0 \bar{s}_1 r \dots r \rightarrow \bar{s}$.

If $S_0 \neq \emptyset$ and $S_1 = \emptyset$, we face the problem that the theory θ_r from Proposition V.1.11 does not refer to the set $\{o, i\}^* \setminus \uparrow\{ro, ri\}$, but to the set $\{o, i\}^* \setminus \uparrow ro$. By a composition argument, we can compute

$$\theta_r = \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}} | \{o, I\}^* \setminus \uparrow r o, r),$$

from the theories

$$\tau := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\{o, I\}^* \setminus \uparrow\{ro, rI\}}, r) \quad \text{and} \quad v := \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\uparrow rI}).$$

We denote the corresponding operation on theories by $\tau \cdot_I v$. Since the theory $v = \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}})$ does not depend on r , we can therefore write

$$R_{\mathfrak{S}, \bar{\theta}} = \bigcup_{\tau} K_{\tau}^{\sigma}(\circ R_o \times \langle \rangle \times \cdots \times \langle \rangle),$$

where σ is the same function as above and the union ranges over all theories τ such that $\tau \cdot_I v = \theta_r$.

If $S_o = \emptyset$ and $S_I \neq \emptyset$, we obtain a similar expression. Finally, suppose that $S_o = \emptyset$ and $S_I = \emptyset$. Again we have to compute the theory

$$\theta_r = \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\{o, I\}^* \setminus \uparrow ro}, r),$$

from the theories

$$\begin{aligned} \tau &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\{o, I\}^* \setminus \uparrow\{ro, rI\}}, r), \\ v &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\uparrow ro}), \\ v' &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}|_{\uparrow rI}). \end{aligned}$$

We denote the corresponding operation on theories by $\tau \cdot_{o, I} (v, v)$. Since $v = v'$ does not depend on r , we obtain

$$R_{\mathfrak{S}, \bar{\theta}} = \bigcup_{\tau} K_{\tau}(\langle \rangle \times \cdots \times \langle \rangle),$$

where the union ranges over all theories τ such that $\tau \cdot_{o, I} (v, v) = \theta_r$. \square

3 VR-equational and HR-equational Structures

We can also characterise the structures in the Caucal hierarchy using the \oplus -terms that we introduced in Section IX.3 to build structures with a given

crossing-width. If we require these terms to be regular (as coloured trees), the resulting structures are tree-interpretable; if we use terms in CH_n , we get all structures in CH_n .

A convenient way to specify terms is via a systems of equations of the form

$$x_i = t_i, \quad \text{for } i \in I,$$

where the x_i are variables and the t_i are finite terms over these variables (similarly to the way we have used systems of equations to define regular linear orders in Theorem VI.4.10). If no term t_i consists of just a variable, there exists a unique n -tuple of infinite terms satisfying all equations. To find this n -tuple, we can treat the system as a directed graph with the variables as vertices. The unravelling of this graph produces the desired solution.

Theorem 3.1. *A structure \mathfrak{A} belongs to CH_n if, and only if, it is the value of a \oplus -term that belongs to CH_n .*

Proof. (\Leftarrow) Let T be a \oplus -term denoting \mathfrak{A} with $T \in \text{CH}_n$. We have seen in Lemma IX.5.4 that there exists an MSO-interpretation τ mapping T to the structure it denotes. Since CH_n is closed under MSO-interpretations, it follows that $\mathfrak{A} = \tau(T) \in \text{CH}_n$.

(\Rightarrow) We consider the following three operations on trees. For a tree \mathfrak{T} with domain C^* , vertices v, \bar{a} , and directions $c, d \in C$, $c \neq d$, we set

$$\begin{aligned} \langle \mathfrak{T}, \bar{a} \rangle|_{\uparrow vc} \oplus \langle \mathfrak{T}, \bar{a} \rangle|_{\uparrow vd} &:= \langle \mathfrak{T}, \bar{a} \rangle|_{\uparrow \{vc, vd\}}, \\ \langle \mathfrak{T}, \downarrow v, \bar{a} \rangle|_{T \setminus \uparrow v} \ll_* \langle \mathfrak{T}, \bar{a} \rangle|_{\uparrow v} &:= \langle \mathfrak{T}, \bar{a} \rangle, \\ \langle \mathfrak{T}, \downarrow v, \bar{a} \rangle|_{T \setminus \uparrow vC} \ll_+ \langle \mathfrak{T}, \bar{a} \rangle|_{\uparrow vC} &:= \langle \mathfrak{T}, \bar{a} \rangle. \end{aligned}$$

Since each of these operations can be expressed as a disjoint union followed by a quantifier-free interpretation, we obtain corresponding operations on MSO_m -theories such that

$$\begin{aligned} \text{Th}_{\text{MSO}}^m(\mathfrak{S}) \oplus \text{Th}_{\text{MSO}}^m(\mathfrak{T}) &= \text{Th}_{\text{MSO}}^m(\mathfrak{S} \oplus \mathfrak{T}), \\ \text{Th}_{\text{MSO}}^m(\mathfrak{S}) \ll_* \text{Th}_{\text{MSO}}^m(\mathfrak{T}) &= \text{Th}_{\text{MSO}}^m(\mathfrak{S} \ll_* \mathfrak{T}), \\ \text{Th}_{\text{MSO}}^m(\mathfrak{S}) \ll_+ \text{Th}_{\text{MSO}}^m(\mathfrak{T}) &= \text{Th}_{\text{MSO}}^m(\mathfrak{S} \ll_+ \mathfrak{T}). \end{aligned}$$

We can use these operations to generate the infinite binary tree $\mathfrak{S}_{\text{bin}}$ since

$$\mathfrak{S}_{\text{bin}}|_v = 1 \ll_+ (\mathfrak{S}_{\text{bin}}|_{v_0} \oplus \mathfrak{S}_{\text{bin}}|_{v_1}), \quad \text{for every } v \in \{0, 1\}^*.$$

It follows that $\mathfrak{S}_{\text{bin}}$ is the unique solution of the equation

$$x = 1 \ll_+ (x \oplus x).$$

Given an interpretation χ and a finite structure \mathfrak{C} with $\mathfrak{A} = \chi(\mathfrak{C}^{*n})$, we obtain a similar term with value \mathfrak{A} as follows. Let c_0, \dots, c_{l-1} be an enumeration of C , let m be the maximal quantifier-rank of the formulae in χ , and let r be the maximal arity of a relation of \mathfrak{A} . For $v \in C^*$, let \mathfrak{B}_v be the substructure of \mathfrak{A} with universe $A \cap \uparrow v$ where we have added the auxiliary relations

$$U_\theta := \{ \bar{a} \in A^n \mid \text{Th}_{\text{MSO}}^m(\mathfrak{S}_{\text{bin}}|_v, \bar{a}) = \theta \},$$

for each MSO_m -theory θ of n -tuples with $0 < n < r$. These structures satisfy the equations

$$\mathfrak{B}_v = d_v \oplus^{\sigma_v} \bigoplus_{c \in C}^{\tau_v} \mathfrak{B}_{vc}, \quad \text{if } v \in A,$$

$$\mathfrak{B}_v = \bigoplus_{c \in C}^{\tau_v} \mathfrak{B}_{vc}, \quad \text{if } v \notin A,$$

where d_v is a suitable constant and the (non-symmetric) update specifications σ_v and τ_v are chosen as follows. τ_v adds a tuple \bar{a}

- ♦ to an auxiliary relation U_θ if

$$\theta_0 \oplus \dots \oplus \theta_{l-1} = \theta,$$

- ♦ and to the relation R defined by the formula φ_R from χ if

$$\varphi_R \in \xi_v^+ \ll_+ (\theta_0 \oplus \dots \oplus \theta_{l-1}),$$

where $\theta_0, \dots, \theta_{l-1}$ are the theories such that $\bar{a} \cap \uparrow v c_i \in U_{\theta_i}$ and

$$\xi_v^+ := \text{Th}_{\text{MSO}}^m(\mathfrak{S}_{\text{bin}}|_{\{0,1\}^* \setminus \uparrow v C}).$$

Similarly, σ_v adds a tuple \bar{a}

- ♦ to an auxiliary relation U_θ if

$$\theta_0 \ll_+ \theta_1 = \theta,$$

- ♦ and to the relation R defined by the formula φ_R from χ if

$$\varphi_R \in \xi_v^* \ll_* (\theta_0 \ll_+ \theta_1),$$

where $\theta_0, \theta_1, \xi_v^*$ are the theories such that

$$\bar{a} \cap \{v\} \in U_{\theta_0}, \quad \bar{a} \cap \uparrow v C \in U_{\theta_1},$$

and $\xi_v^* := \text{Th}_{\text{MSO}}^m(\mathfrak{S}_{\text{bin}}|_{\{0,1\}^* \setminus \uparrow v})$.

Let t be the \oplus -term defined by the equations

$$x_v = d_v \oplus^{\sigma_v} \bigoplus_{c \in C}^{\tau_v} x_{vc}, \quad \text{if } v \in A,$$

$$x_v = \bigoplus_{c \in C}^{\tau_v} x_{vc}, \quad \text{if } v \notin A.$$

Then t can be obtained from the tree \mathfrak{C}^* by an MSO-transduction that replaces every vertex $v \in C^*$ by the term

$$d_v \oplus^{\sigma_v} \bigoplus_{c \in C}^{\tau_v} x_{vc} \quad \text{or} \quad \bigoplus_{c \in C}^{\tau_v} x_{vc}.$$

Since CH_n is closed under copying operations and MSO-interpretations, it follows that $t \in \text{CH}_n$.

The term t we have constructed above uses non-symmetric update specifications, but \oplus -terms were defined using symmetric specifications only. Since every \oplus^τ -operation in t has exactly $l < \omega$ arguments, we can transform t into a proper \oplus -term by using l copies of each colour, one copy for each argument. \square

Of particular interest are structures that can be denoted by a term that is regular. By the preceding theorem, these are exactly the tree-interpretable ones.

Definition 3.2. A structure is *VR-equational* if it is denoted by a \oplus -term whose underlying tree is regular.

Originally, VR-equational structures were introduced in the study of certain graph grammars with *vertex replacement*. Hence, the name.

We have seen above that a structure in the Caucal hierarchy has a decidable GSO-theory if, and only if, it has finite tree-width. We can also characterise such structures via certain terms.

Definition 3.3. Let Σ be a finite relational signature and $k < \omega$ a constant.

(a) An HR_k -term is a term using

- ♦ binary operations $\parallel_{\xi, \zeta}$, for all $\xi, \zeta \subseteq [k]$ with $\xi \cup \zeta = [k]$,
- ♦ constant symbols $R_{\bar{\alpha}}$, for each relation symbol $R \in \Sigma$ and every tuple $\bar{\alpha} \in [k]^n$ with $\alpha_i \neq \alpha_j$, where n is the arity of R , and
- ♦ a constant symbol 1.

(b) The value of an HR_k -term is a $(\Sigma + [k])$ -structure defined as follows.

We consider the numbers $i \in [k]$ as constant symbols which we call *handles*.

The constant term 1 denotes a structure with a single element where all relations are empty and all handles are equal to the only element.

For every n -ary relation symbol $R \in \Sigma$ and every n -tuple of distinct handles $\bar{\alpha} \in [k]^n$, the constant term $R_{\bar{\alpha}}$ denotes an n -element structure \mathbb{C}_R whose elements are denoted by the handles $\alpha_0, \dots, \alpha_{n-1}$. The handles $\beta \in [k] \setminus \bar{\alpha}$ are set to $\beta := \alpha_0$. Finally, all relations of \mathbb{C}_R are empty, except for the relation R which contains the tuple $\langle \alpha_0, \dots, \alpha_{n-1} \rangle$.

Given two $(\Sigma + [k])$ -structures \mathfrak{A} and \mathfrak{B} , the operation $\mathfrak{A} \parallel_{\xi, \zeta} \mathfrak{B}$ constructs the structure obtained from the disjoint union of the $(\Sigma + \xi)$ -reduct of \mathfrak{A} and the $(\Sigma + \zeta)$ -reduct of \mathfrak{B} by identifying the elements $\alpha^{\mathfrak{A}}$ and $\alpha^{\mathfrak{B}}$, for all $\alpha \in \xi \cap \zeta$.

Finally, given a possibly infinite HR_k -term t , we define the value $\text{val}(t)$ of t as the $(\Sigma + [k])$ -structure obtained by

- ♦ forming the disjoint union of all structures associated with the constant symbols in t

- ♦ and then identifying elements according to the operations $\parallel_{\xi, \zeta}$ in t .

(c) We say t is an *HR-term* for a Σ -structure \mathfrak{A} if \mathfrak{A} is isomorphic to the Σ -reduct of $\text{val}(t)$.


(d) A structure \mathfrak{A} is *HR-equational* if it is denoted by a regular HR_k -term, for some $k < \omega$.

Similarly to VR-equational structures, the name HR-equational refers to certain graph grammars with *hyperedge replacement*.

Example. The HR_3 -term

$$(E_{12} \parallel_{1,3} E_{23}) \parallel_{\emptyset, \emptyset} (E_{12} \parallel_{1,3} E_{23})$$

produces the square



Theorem 3.4. Let \mathfrak{A} be a Σ -structure. Then $\mathfrak{A}_{\text{in}} \in \text{CH}_n$ if, and only if, there is some $k < \omega$ such that \mathfrak{A} is the value of a HR_k -term t that belongs to CH_n .

Proof. (\Leftarrow) The function val mapping an HR_k -term to its value is an k -copying GSO-transduction. (We need the copy operation to turn every leaf labelled by a constant symbol $R_{\bar{\alpha}}$ into n elements, where $n \leq k$ is the arity of R . The construction is similar to the one in the proof of Proposition X.2.1.) Furthermore, for every successor-tree $\mathfrak{T} \in \text{CH}_n$, there exists an GSO-interpretation mapping some successor-tree $\mathfrak{S} \in \text{CH}_n$ to $\text{copy}_k(\mathfrak{T})$. It follows that \mathfrak{A}_{in} can be MSO-interpreted in \mathfrak{S} . Since CH_n is closed under MSO-interpretations, the claim follows.

(\Rightarrow) Let s be an \otimes -term denoting \mathfrak{A}_{in} with $s \in \text{CH}_n$. According to Lemma IX.4.3, there exists an MSO-interpretation τ mapping s to a tree-decomposition $(W_t)_{t \in T}$ of bounded width whose index tree T is the same

as the underlying tree of s . Note that $s \in \text{CH}_n$ implies that (the tree encoding) $(W_t)_{t \in T}$ also belongs to CH_n . Let k be the width of the decomposition $(W_t)_{t \in T}$. We turn $(W_t)_{t \in T}$ into the desired HR_k -term as follows.

Let $\chi: A \rightarrow [k]$ be a k -colouring of A such that the restriction $\chi \upharpoonright W_t$ is injective, for every t . Since $(W_t)_{t \in T}$ belongs to CH_n , we can choose χ such that the coloured tree encoding $\langle (W_t)_{t \in T}, \chi \rangle$ belongs to CH_n . Below we will define HR_k -terms u_t , for $t \in T$, such that $\text{val}(u_t) = \mathfrak{A}|_{W_t}$ and each element $a \in W_t$ is denoted by the handle $\chi(a)$. (We do not care which elements the handles outside of $\chi[W_t]$ are assigned to.) Using these terms u_t , we consider the equations

$$x_t = u_t \parallel_{[k], M_0} x_{v_0} \parallel_{[k], M_1} x_{v_1} \parallel_{[k], M_2} \cdots \parallel_{[k], M_{n-1}} x_{v_{n-1}},$$

where v_0, \dots, v_{n-1} are the successors of t in T and $M_i := \chi[W_t \cap W_{v_i}]$. Let \mathfrak{B}_t denote the value of x_t in the unique solution of these equations. It follows that \mathfrak{B}_t is isomorphic to the substructure of \mathfrak{A} induced by the set $\bigcup_{v \in \uparrow t} W_v$. In particular, we have $\mathfrak{B}_r|_\Sigma \cong \mathfrak{A}$, for the root r of T .

Finally, note that we can obtain the solution of the equations by applying a parameter-less MSO-transduction to the tree encoding $\langle (W_t)_{t \in T}, \chi \rangle$. This transduction replaces every vertex $t \in T$ by the term

$$u_t \parallel_{[k], M_0} x_{v_0} \parallel_{[k], M_1} x_{v_1} \parallel_{[k], M_2} \cdots \parallel_{[k], M_{n-1}} x_{v_{n-1}}.$$

Since CH_n is closed under copying operations and interpretations, it follows that the solution also belongs to CH_n .

It remains to define the terms u_t . More generally we will prove that, for every structure \mathfrak{C} and every injective function $\chi: C \rightarrow [k]$, there exists an HR_k -term u such that $\text{val}(u) = \mathfrak{C}$ and every element $c \in C$ is denoted by the handle $\chi(c)$. Let $\bar{a}^0, \dots, \bar{a}^{n-1}$ be an enumeration of all tuples that belong to a relation in \mathfrak{C} . We construct u by induction on n . First suppose that $n = 0$ and let $\alpha_0, \dots, \alpha_{m-1}$ be an enumeration of $\chi[C]$. We set

$$u := 1 \parallel_{[k] \setminus \alpha_1, \alpha_1} \cdots \parallel_{[k] \setminus \alpha_{n-1}, \alpha_{n-1}} 1,$$

where we assume that $\parallel_{\xi, \zeta}$ associates to the left (i.e., $a \parallel b \parallel c = (a \parallel b) \parallel c$)

For the inductive step, suppose that there is some tuple \bar{a} in the relation R . Let \mathfrak{C}' be the structure obtained from \mathfrak{C} by removing \bar{a} from R . By inductive hypothesis, we have a HR_k -term u' denoting \mathfrak{C}' . Then we can set

$$u := R_{\chi(\bar{a})} \parallel_{\chi(\bar{a}), [k]} u'. \quad \square$$

Exercise 3.1. Prove that a countable Σ -structure \mathfrak{A} has tree-width at most k if, and only if, it is the value of an HR_k -term. \lrcorner

4 Regular Path-Interpretations

For particular types of structures, there are additional characterisations of tree-interpretability. We have already seen in Theorem VI.4.10 that the tree-interpretable linear orders are exactly the regular ones. In this section we consider graphs, more specifically, edge-labelled directed graphs.

Remark. The results in this section also hold for graphs with vertex labels. To simplify the exposition, we have refrained from stating them in this generality. For the general versions, you can add monadic predicates to all graphs with the requirement that each of them forms a regular language. \lrcorner

For graphs, we can replace MSO-interpretations by the following two more concrete operations.

Definition 4.1. Let A and B be two sets of edge labels and let

$$\bar{A} := \{ \bar{a} \mid a \in A \}$$

be a disjoint copy of A . A *regular path-interpretation* (from A to B) is a tuple $\alpha = \langle K, (L_b)_{b \in B} \rangle$ of regular languages $K \subseteq A^*$ and $L_b \subseteq (A + \bar{A})^*$, for $b \in B$.

Given an A -labelled directed graph $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ and a vertex $v_0 \in V$, such an interpretation produces a new graph

$$\alpha(\mathfrak{G}, v_0) := \langle W, (F_b)_{b \in B} \rangle$$

as follows. The universe W consists of all vertices $u \in V$ that can be reached (in \mathfrak{G}) from v_o by some path whose labelling belongs to K . The edge relation F_b consists of those pairs $\langle u, v \rangle \in V \times V$ such that in the expansion of \mathfrak{G} by the inverse edge relations

$$E_{\bar{a}} := (E_a)^{-1}, \quad \text{for } \bar{a} \in \bar{A},$$

there is a path from u to v labelled by some word in L_b .

Example. (a) We have $\langle \omega, \leq, \text{suc} \rangle \cong \alpha(\mathfrak{T}_{\text{bin}}, \langle \rangle)$, where $\alpha = \langle K, L_{\text{suc}}, L_{\leq} \rangle$ with

$$K := 1^*, \quad L_{\text{suc}} := 1, \quad \text{and} \quad L_{\leq} := 1^*.$$

(b) Let \mathfrak{G} be the directed $\mathbb{Z} \times \mathbb{Z}$ grid with edge relations

$$E_h := \{ \langle \langle x, y \rangle, \langle x + 1, y \rangle \rangle \mid x, y \in \mathbb{Z} \},$$

$$E_v := \{ \langle \langle x, y \rangle, \langle x, y + 1 \rangle \rangle \mid x, y \in \mathbb{Z} \},$$

and let $\alpha = \langle K, L_h, L_v, L_v \rangle$ be the path-interpretation where

$$K := (h + v)^*, \quad L_h := (h + \bar{h})^*,$$

$$L_d := (hv)^* + (\bar{h}\bar{v})^*, \quad L_v := (v + \bar{v})^*.$$

Then $\alpha(\mathfrak{G}, \langle o, o \rangle)$ is the graph with vertices $\mathbb{Z} \times \mathbb{Z}$ and edge relations

$$E_h := \{ \langle \langle x, y \rangle, \langle x', y \rangle \rangle \mid x, x', y \in \mathbb{Z} \},$$

$$E_v := \{ \langle \langle x, y \rangle, \langle x, y' \rangle \rangle \mid x, y, y' \in \mathbb{Z} \},$$

$$E_d := \{ \langle \langle x, y \rangle, \langle x + k, y + k \rangle \rangle \mid x, y, k \in \mathbb{Z} \}.$$

Exercise 4.1. Let α and β be two regular path-interpretations (from A to B and from B to C , respectively). Prove that $\beta \circ \alpha$ is also a regular path-interpretation.

Theorem 4.2. A graph $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ is tree-interpretable if, and only if,

$$\mathfrak{G} \cong \alpha(\mathfrak{T}_{\text{bin}}, \langle \rangle), \quad \text{for some regular path-interpretation } \alpha.$$

Proof. (\Leftarrow) Let $\alpha = \langle K, (L_b)_b \rangle$. For every regular language L over the alphabet $\{o, 1, \bar{o}, \bar{1}\}$, there exists a GSO-formula $\psi_L(u, v)$ stating that there exists a path from u to v labelled by some word in L . Since trees are sparse, we can translate these formulae into MSO. Thus, we obtain an MSO-interpretation $\tau = \langle \delta, (\varphi_a)_{a \in A} \rangle$ of \mathfrak{G} in $\mathfrak{T}_{\text{bin}}$ by

$$\delta(x) := \psi_K(\langle \rangle, x) \quad \text{and} \quad \varphi_a(x, y) := \psi_{L_a}(x, y).$$

(\Rightarrow) Suppose that $E_a = \bigcup_{i < n} W_i^a(U_i^a \times V_i^a)$ where all languages are over the alphabet Σ . Note that, replacing V and all W_i^a by oV and oW_i^a , respectively (which results in a graph isomorphic to \mathfrak{G}), we may assume that $\langle \rangle \notin W_i^a$. To create an edge $\langle x, y \rangle$ we have to check that $x = wu$ and $y = wv$, for some $w \in W_i^a$, $u \in U_i^a$, and $v \in V_i^a$. To do so, we use the path from x to y . Starting at x we go the path u backwards and then the path v forwards to arrive at y . But we have to check that the vertex w where we change directions belongs to W_i^a . To do so, we label the incoming edge at w with additional information indicating that the vertex w belongs to W_i^a . Then we can traverse this edge twice: first backwards and then immediately forwards again, to check that the vertex really is in W_i^a . (This is where we need our assumption that $\langle \rangle \notin W_i^a$.)

The formal definition is as follows. Let $\eta : \Sigma^* \rightarrow \mathfrak{M}$ be a homomorphism to a finite monoid that recognises the languages V and W_i^a , for all a, i , and let $P_i^a \subseteq M$ be the set such that $W_i^a = \eta^{-1}[P_i^a]$. Let \mathfrak{S} be the tree with domain $(\Sigma \times M)^*$. We start by defining a path-interpretation $\alpha = \langle K, (L_b)_b \rangle$ with

$$\mathfrak{G} \cong \alpha(\mathfrak{S}, \langle \rangle).$$

For K we take the language of all words

$$\langle d_o, s_o \rangle \cdots \langle d_{m-1}, s_{m-1} \rangle \in (\Sigma \times M)^*$$

such that

$$d_o \cdots d_{m-1} \in V, \quad s_o = \eta(d_o), \quad \text{and} \quad s_{j+1} = s_j \cdot \eta(d_{j+1}),$$

for $j < m - 1$. Note that these conditions can be expressed in MSO. Hence, K is a regular language. The languages L_b are defined by

$$L_b := \bigcup_{i < n} p^{-1}[\overline{U_i^a}] z_i^a z_i^a p^{-1}[V_i^a],$$

where

- $p : \Sigma \times M \rightarrow \Sigma$ is the projection to the first component,
- \overline{U} denotes the language obtained from U by reversing the order of every word and replacing every letter c by its inverse \bar{c} , and
- $z_i^a = \sum \{ \langle d, s \rangle \in \Sigma \times M \mid s \in P_i^a \}$.

To conclude the proof, let g be arbitrary injective function $\Sigma \times M \rightarrow \{0, 1\}^m$ with sufficiently large $m < \omega$. Extending g to a homomorphism $(\Sigma \times M)^* \rightarrow \{0, 1\}^*$, we obtain

$$\mathfrak{S} = \beta(\mathfrak{T}_{\text{bin}}, \langle \rangle), \quad \text{where} \quad \beta = \langle \text{rng } g, (g(c))_{c \in \Sigma \times M} \rangle.$$

Hence,

$$\mathfrak{S} = \alpha(\beta(\mathfrak{T}_{\text{bin}}, \langle \rangle), \langle \rangle) = \gamma(\mathfrak{T}_{\text{bin}}, \langle \rangle),$$

where $\gamma := \langle g[K], (g[L_b])_{b \in B} \rangle$. \square

If we want a similar characterisation of the higher levels of the Caucal hierarchy, we need an operation to replace the Muchnik iteration with. This turns out to be the unravelling operation on graphs.

Definition 4.3. The *unravelling* $\mathcal{U}(\mathfrak{S}, v_o)$ of a graph $\mathfrak{S} = \langle V, (E_a)_{a \in A} \rangle$ from a vertex $v_o \in V$ is the graph

$$\mathcal{U}(\mathfrak{S}, v_o) := \langle U, (F_a)_{a \in A} \rangle$$

whose universe U consists of all finite paths (including the empty one) in \mathfrak{S} that start at the vertex v_o , and where the edge relation F_a contains all pairs $\langle p, q \rangle$ of paths where q is obtained from p by appending a single E_a -edge. \dashv

Theorem 4.4. A graph \mathfrak{S} belongs to CH_{n+1} if, and only if, it is of the form

$$\mathfrak{S} \cong \alpha(\mathcal{U}(\mathfrak{S}, v_o), v_o), \quad \text{for some regular path-interpretation } \alpha \text{ and} \\ \text{some } \mathfrak{S} \in \text{CH}_n.$$

We split the proof into several lemmas. In the same way as in Lemma 1.7, we can generate each level of the hierarchy from graphs of a particularly simple kind using regular path-interpretations. These graphs are trees with the additional property that every vertex has at most one outgoing edge with a given label.

Definition 4.5. A graph $\mathfrak{S} = \langle V, (E_a)_{a \in A} \rangle$ is *deterministic* if every vertex has at most one outgoing edge with label a , for all $a \in A$. \dashv

We have seen in the proof above that, for regular trees, we can replace an MSO-interpretation by a regular path-interpretation. For non-regular trees, this is only possible if we annotate the tree first with additional information. We use an MSO-interpretation for this. It sounds like we have not made any progress by this, replacing an interpretation by an interpretation followed by a path-interpretation. But the interpretation we construct is of a simple special form that commutes with the unravelling operation. This is all we will need to make the inductive argument in the theorem below go through.

Lemma 4.6. Let $\mathfrak{T} \in \text{CH}_n$ be an A -labelled deterministic tree and τ an MSO-interpretation mapping A -labelled graphs to B -labelled ones. There exist a set of labels A^+ , an A^+ -labelled deterministic tree $\mathfrak{S} \in \text{CH}_n$, and a regular path-interpretation α such that

$$\tau(\mathcal{U}(\mathfrak{T}, \langle \rangle)) \cong \alpha(\mathcal{U}(\mathfrak{S}, \langle \rangle), \langle \rangle).$$

Proof. If \mathfrak{T} has a single vertex, the claim is trivial. (Each of the languages in α is either empty or contains the empty word.) Hence, we may assume that there are at least two vertices, which means that every vertex has at least one incoming or outgoing edge. Let m be the maximal quantifier-rank of the formulae from τ and let Θ be the set of all MSO_m -theories with one

additional first-order parameter. Let $\mathfrak{T} \in \text{CH}_n$ be a deterministic tree with edge labels from A .

The tree \mathfrak{S} is obtained by annotating every edge $\langle u, v \rangle$ of \mathfrak{T} by the theories

$$\begin{aligned}\theta &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u \setminus \uparrow v}, u), \\ \rho &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v}, v), \\ \sigma &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow v}, u), \\ v_a &:= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u \setminus \uparrow \{v, w\}}, u), \quad \text{where } w \text{ is the } a\text{-successor of } u.\end{aligned}$$

(If u has no a -successor, we choose an arbitrary theory for v_a .) Hence, we set

$$A^\dagger := A \times \Theta \times \Theta \times \Theta \times \Theta^A.$$

Let \mathfrak{S} be the graph obtained from \mathfrak{T} by replacing every edge label a , by $\langle a, \theta, \rho, \sigma, (v_a)_{a \in A} \rangle$. Clearly, \mathfrak{S} is a deterministic tree.

It therefore remains to prove that it belongs to CH_n . To do so, we will define an MSO-interpretation σ with $\mathfrak{S} \cong \sigma(\mathfrak{T})$. As CH_n is closed under MSO-interpretations, it then follows that $\mathfrak{S} \in \text{CH}_n$. To construct σ note that, given an edge $\langle u, v \rangle$, the theories

$$\begin{aligned}\theta &= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u \setminus \uparrow v}, u), \\ \rho &= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v}, v), \\ \sigma &= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow v}, u), \\ v_a &= \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u \setminus \uparrow \{v, w\}}, u)\end{aligned}$$

are MSO-definable in $\mathcal{U}(\mathfrak{T}, \langle \rangle)$. Using the Theorem of Muchnik, it follows that they are also definable in \mathfrak{T} . Consequently, we can write down an MSO-formula stating that there is an $\langle a, \theta, \rho, \sigma \rangle$ labelled edge from u to v . These formulae form the interpretation σ .

Having defined \mathfrak{S} , it remains to construct $\alpha = \langle K, (L_b)_{b \in B} \rangle$. Suppose that $\tau = \langle \delta, (\varphi_b)_{b \in B} \rangle$. We have to define the language K such that it selects all vertices that satisfy δ , that is, all vertices v such that

$$\delta \in \text{Th}_{\text{MSO}}^m(\mathfrak{T}, v).$$

This theory is determined by the labelling of the path from the root to the vertex v in \mathfrak{S} . Hence, let K be the set of all words

$$\langle a_0, \theta_0, \rho_0, \sigma_0 \rangle \cdots \langle a_n, \theta_n, \rho_n, \sigma_n \rangle$$

such that

$$\delta \in \sum_{i < n} \theta_i + \rho_n.$$

(We are quite liberal with our notation here. We use the symbol $+$ for any operation on theories corresponding to a disjoint union followed by some quantifier-free interpretation. But this interpretation does not need to be the same at each occurrence of $+$.) Furthermore, we include the empty word $\langle \rangle$ in K if $\mathfrak{T} \models \delta(\langle \rangle)$. Otherwise, we omit it.

To construct h we have to compute the theory $\text{Th}_{\text{MSO}}^m(\mathfrak{T}, uv)$ by following a path between the given vertices u and v . This path takes a different form depending on the relative position of u and v . Consequently, we have

$$h(b) := L_b^0 \cup L_b^1 \cup L_b^2 \cup L_b^3 \cup L_b^4, \quad \text{for } b \in B,$$

with one language L_b^i for each of the following five cases.

If $u = v$ and u has a predecessor w , we can compute the theory of uv from the theories

$$\begin{aligned}\text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow u}, w), \\ \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u}, u).\end{aligned}$$

Hence, we can use for L_b^0 the set of all words

$$\overline{\langle a, \theta, \rho, \sigma, \bar{v} \rangle} \langle a, \theta, \rho, \sigma, \bar{v} \rangle \quad \text{with} \quad \varphi_b(x, y) \in \sigma + \rho.$$

Similarly, if $u = v$ and u has a successor w , we can compute the theory of uv from

$$\begin{aligned}\text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow w}, u), \\ \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow w}, w).\end{aligned}$$

Hence, we can use for L_b^1 the set of all words

$$\langle a, \theta, \rho, \sigma, \bar{v} \rangle \overline{\langle a, \theta, \rho, \sigma, \bar{v} \rangle} \quad \text{with} \quad \varphi_b(x, y) \in \sigma + \rho.$$

If $u <_{\text{pf}} v$, we can compute the theory of uv from

$$\begin{aligned} & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow u, \downarrow u}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u \setminus \uparrow v, \downarrow v, u}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v, v}). \end{aligned}$$

This can be done by the language L_b^2 of all words

$$\langle a_o, \theta_o, \rho_o, \sigma_o, \bar{v}_o \rangle \cdots \langle a_n, \theta_n, \rho_n, \sigma_n, \bar{v}_n \rangle$$

such that

$$\varphi_b(x, y) \in \sigma_o + \sum_{o < i < n} \theta_i + \rho_n.$$

If $v <_{\text{pf}} u$, we use the theories

$$\begin{aligned} & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{T \setminus \uparrow v, \downarrow v}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v \setminus \uparrow u, \downarrow u, v}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u, u}). \end{aligned}$$

The language L_b^3 is therefore defined analogously to L_b^2 but we use inverse edges and reverse their ordering.

Finally, suppose that u and v are incomparable. Set $w := u \sqcap v$ and let u' and v' be the successors of, respectively, u and v with $u' \leq_{\text{pf}} v$ and $v' \leq_{\text{pf}} u$. Then we can compute the theory of uv from

$$\begin{aligned} & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow w \setminus \uparrow \{u', v'\}}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u' \setminus \uparrow u, \downarrow u}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v' \setminus \uparrow v, \downarrow v}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow u, u}), \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}|_{\uparrow v, v}). \end{aligned}$$

Let N_b be the set of all words

$$\overline{\langle a_m, \theta_m, \rho_m, \sigma_m, \bar{v}_m \rangle} \cdots \overline{\langle a_o, \theta_o, \rho_o, \sigma_o, \bar{v}_o \rangle} \\ \langle a'_o, \theta'_o, \rho'_o, \sigma'_o, \bar{v}'_o \rangle \cdots \langle a'_n, \theta'_n, \rho'_n, \sigma'_n, \bar{v}'_n \rangle$$

such that $(v_o)_{a'_o} = (v'_o)_{a_o}$ and $\varphi_b(x, y) \in (v_o)_{a'_o}(\xi, \xi')$, where

$$\xi := \sum_{o < i < m} \theta_i + \rho_m \quad \text{and} \quad \xi' := \sum_{o < i < n} \theta'_i + \rho'_m. \quad \square$$

We would like to express the Muchnik iteration by a combination of path-interpretations and an unravelling operation. But there is one minor technical issue: the latter operations are only defined for edge-labelled graphs, while the Muchnik iteration adds a unary relation cl . For this reason, we introduce a variant of the Muchnik iteration that replaces cl by an edge relation instead.

Definition 4.7. Let $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ be a graph. We denote by $\mathfrak{G}^{*'}$ the graph obtained from the Muchnik iteration \mathfrak{G}^* by omitting the relations suc and cl and adding the new relation

$$\text{cl}' := \{ \langle u, v \rangle \mid \langle u, v \rangle \in \text{suc} \text{ and } v \in \text{cl} \}. \quad \lrcorner$$

Remark. Note that, for connected graphs \mathfrak{G} , we can translate between $\mathfrak{G}^{*'}$ and \mathfrak{G}^* using an MSO-interpretation. Consequently, the hierarchy generated by interpretations and this variant of the Muchnik iteration coincides with the Caucal hierarchy level by level. \lrcorner

Lemma 4.8. Let A be a set of edge-labels, let A' be a copy of A , $\# \notin A$ a new label, and set $B := A + A' + \{\#\}$. There exist regular path-interpretations α and β such that

$$\mathfrak{T}^{*'} = \alpha(\mathcal{U}(\beta(\mathfrak{T}, v_o), v_o), v_o),$$

for every deterministic tree $\mathfrak{T} = \langle V, (E_a)_{a \in A} \rangle$ and every vertex $v_o \in V$.

Proof. Let \mathfrak{T}' be the graph obtained from \mathfrak{T} by adding an inverse to each edge and a $\#$ -labelled loop to every vertex. We use the labels $a' \in A'$ for

the inverse of an edge with label $a \in A$. Then $\mathfrak{T}' = \beta(\mathfrak{T})$ where the regular path-interpretation

$$\beta = \langle A^*, L_{\#}, (L_a)_{a \in A}, (L_{a'})_{a \in A} \rangle$$

is defined by the languages

$$L_{\#} := \langle \rangle, \quad L_a := a, \quad \text{and} \quad L_{a'} := \bar{a}, \quad \text{for } a \in A.$$

Set

$$\mathfrak{S} := \mathcal{U}(\mathfrak{T}', v_o).$$

By definition, each vertex of \mathfrak{S} corresponds to a path through \mathfrak{T}' starting at v_o . Every such path is of the form

$$p_o \# p_1 \# \dots \# p_{n-2} \# p_{n-1},$$

where each p_o is a path without the label $\#$. Let $\rho : S \rightarrow T^*$ be the function mapping each such path to the word

$$\langle u_o, u_1, \dots, u_{n-1} \rangle,$$

where u_i is the last vertex of p_i . Hence, using the function ρ , we can encode the elements of \mathfrak{T}^* by the vertices of \mathfrak{S} . The problem is that ρ is not injective. We have to find a subset $U \subseteq S$ such that the restriction $\rho \upharpoonright U$ is bijective. For an example of the definition of U see Figure 1. (For simplicity we have omitted the $\#$ -edges. The actual graph \mathfrak{S} consists of infinitely many copies of the graph in the figure.),

For each vertex u of \mathfrak{T} , we need to find a set of paths starting at v such that there is exactly one path ending at every given other vertex of \mathfrak{T} . Such a set of paths consists of all paths that do not contain two consecutive edges labelled by a and a' , or by a' and a . Hence, let $K \subseteq (A + \bar{A})^*$ be the set of all words that do not contain a factor of the form $a\bar{a}$ or $\bar{a}a$, for $A \in A$. We claim that the restriction of ρ to the vertices reachable by a path in K is bijective.




Figure 1: A tree \mathfrak{T} on the left and (part of) $\mathcal{U}(g_o^{-1}(\mathfrak{T}), v_o)$ on the right, where g_o is the restriction of g without the $\#$ -edges and v_o is the vertex denoted by the filled circle. The broken arrows represent the edges deleted by K .

To see this first note that, given two vertices v and u of \mathfrak{T} , there is a path from u to v whose labelling consists of a sequence of inverse edges followed by a sequence of ordinary ones. Since \mathfrak{T} is deterministic, the corresponding sequence of labels belongs to K . Conversely, if p is a path without $\#$ -edges whose labelling belongs to K , the, since we are dealing with a tree, the path p must consist of a sequence of inverse edges followed by one of ordinary ones.

It follows that we can use the path-interpretation $\alpha = \langle K, L_{cl}, (L_a)_{a \in A} \rangle$ where

$$L_{cl} := \# \quad \text{and} \quad L_a := a + \bar{a}', \quad \text{for } a \in A. \quad \square$$

Proof of Theorem 4.4. (\Leftarrow) Suppose that there exists a regular path-interpretation $\alpha = \langle K, (L_a)_a \rangle$ such that

$$\mathfrak{S} \cong \alpha(\mathcal{U}(\mathfrak{H}, v_o), v_o)$$

where $\mathfrak{S} = \langle V, (E_a)_{a \in A} \rangle$ and $\mathfrak{H} = \langle W, (F_b)_{b \in B} \rangle$.

For every regular language L over the alphabet $B \cup \bar{B}$, there exists a GSO-formula $\psi_L(u, v)$ stating that there exists a path from u to v labelled by some word in L . Since trees are sparse, we can translate these formulae into MSO. Thus, we obtain an MSO-interpretation $\sigma = \langle \delta, (\varphi_a)_{a \in A} \rangle$ of \mathfrak{S} in $\mathcal{U}(\mathfrak{H}^*, v_o)$ by

$$\delta(x) := \psi_K(v_o, x) \quad \text{and} \quad \varphi_a(x, y) := \psi_{L_a}(x, y).$$

Furthermore, we have seen in the exercise on page 228 that there exists an MSO-interpretation τ such that $\mathcal{U}(\mathfrak{H}) = \tau(\mathfrak{H}^*)$. Hence,

$$\mathfrak{G} = (\sigma \circ \tau)(\mathfrak{H}^*) \in \text{CH}_{n+1}.$$

(\Rightarrow) Let $\mathfrak{G} \in \text{CH}_{n+1}$. By Lemma 1.7, this means that

$$\mathfrak{G} = \tau(\mathfrak{C}^{*(n+1)}),$$

for some finite structure \mathfrak{C} and an MSO-transduction τ . If \mathfrak{P} is a path of sufficient length, we can interpret \mathfrak{C} in \mathfrak{P} . By Lemma V.3.12, it follows that

$$\mathfrak{G} = \sigma(\mathfrak{P}^{*(n+1)}),$$

for some MSO-interpretation σ . Replacing the iteration $-^*$ by the variant $-^{*'}$, we can construct an MSO-interpretation σ' such that

$$\mathfrak{G} = \sigma'(\mathfrak{P}^{*'(n+1)}).$$

Since $-^{*'}$ maps deterministic trees to deterministic trees, it follows that $\mathfrak{T} := \mathfrak{P}^{*'n}$ is a deterministic tree. By Lemma 4.8, we can write

$$\mathfrak{T}^{*'} = \alpha(\mathcal{U}(\beta(\mathfrak{T}, \langle \rangle), \langle \rangle), \langle \rangle),$$

for regular path-interpretations α and β . Since α is equal to an MSO-interpretation ρ , we obtain

$$\mathfrak{G} \cong (\sigma' \circ \rho)(\mathcal{U}(\beta(\mathfrak{T}, \langle \rangle), \langle \rangle)).$$

As CH_n is closed under regular path-interpretations, we have $\beta(\mathfrak{T}, \langle \rangle) \in \text{CH}_n$. Consequently, we can use Lemma 4.6 to find a deterministic tree $\mathfrak{S} \in \text{CH}_n$ and a regular path-interpretation γ such that

$$\mathfrak{G} \cong \gamma(\mathcal{U}(\mathfrak{S}, \langle \rangle), \langle \rangle). \quad \square$$

5 Pushdown Graphs

The second description specific to graphs is based on a suitable automaton model. We start with tree-interpretable graphs.

Definition 5.1. (a) A *pushdown automaton with ε -transitions* is a tuple

$$\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, \gamma_o, F \rangle$$

consisting of a finite set Q of *states*, an *input alphabet* A , a *stack alphabet* Γ , an *initial state* $q_o \in Q$, an *initial stack symbol* $\gamma_o \in \Gamma$, a set $F \subseteq Q$ of *final states*, and a finite *transition relation*

$$\Delta \subseteq Q \times A^{\leq 1} \times \Gamma^* \times \Gamma^* \times Q.$$

A transition $\langle p, a, u, v, q \rangle \in \Delta$ is called an *ε -transition* if $a = \langle \rangle$. We assume that, for every state $p \in Q$, if there are transitions

$$\langle p, a, u, v, q \rangle, \langle p, a', u', v', q' \rangle \in \Delta \quad \text{then} \quad a = \langle \rangle \Leftrightarrow a' = \langle \rangle.$$

(b) The *configuration graph* of \mathcal{A} is the graph with vertex set $Q \times \Gamma^+$. A *configuration* of \mathcal{A} is a pair $\langle q, z \rangle \in Q \times \Gamma^+$. For each $a \in A^{\leq 1}$, there is an *a -labelled edge* from $\langle p, z \rangle$ to $\langle q, z' \rangle$ if there exists a transition

$$\langle p, a, u, v, q \rangle \in \Delta,$$

such that $z = wu$ and $z' = wv$. The *initial configuration* is $\langle q_o, \gamma_o \rangle$.

(c) An *ε -pushdown graph* is a graph \mathfrak{G} obtained from the configuration graph of a pushdown automaton with ε -transitions by


- adding an *a -labelled edge* between two configurations if there exists a path between them that consists of an *a -transition* followed by an arbitrary number of ε -transitions,
- removing all vertices with outgoing ε -transitions, and
- restricting the resulting graph to the set of configurations reachable from the initial one.

(d) A *restricted pushdown automaton* is a pushdown automaton \mathcal{A} such that every transition $\langle p, a, u, v, q \rangle \in \Delta$ is of one of the following forms.

- ♦ a *pop*: $|u| = 1$ and $v = \langle \rangle$,
- ♦ a *nop*: $|u| = 1$ and $u = v$,
- ♦ a *push*: $u = \gamma$ and $v = \gamma\delta$, for some $\gamma, \delta \in \Gamma$.

A *restricted ε -pushdown graph* is an ε -pushdown graph associated with a restricted pushdown automaton. \downarrow

Example. A pushdown automaton for $\langle \omega, \leq, \text{suc} \rangle$ has two states p, q , one stack letter γ , and the following configuration graph.



This automaton is restricted. \downarrow

Restricted automata are the more commonly used automaton model. We have introduced unrestricted ones as they are more convenient to work with. We start by observing that both models have the same expressive power.

Lemma 5.2. *Every configuration graph of a pushdown automaton is also the configuration graph of a restricted pushdown automaton.*

Proof. Fix a pushdown automaton $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_0, \gamma_0, F \rangle$. Let $n < \omega$ be some constant such that

$$\langle p, a, u, v, q \rangle \in \Delta \quad \text{implies} \quad |u|, |v| \leq n.$$

We construct the desired restricted pushdown automaton

$$\mathcal{B} = \langle Q', A, \Gamma', \Delta', q'_0, \gamma'_0, F' \rangle$$

as follows. The automaton \mathcal{B} combines n stack symbols of \mathcal{A} into a single symbol. Furthermore, \mathcal{B} remembers the top-most stack symbols (up to $2n$ many) in its state. Consequently, we set

$$\Gamma' := \Gamma^n, \quad Q' := Q \times \Gamma^{\leq 2n}, \quad \text{and} \quad q'_0 := \langle q'_0, \langle \rangle \rangle.$$

The transition relation Δ' contains all transitions

$$\langle \langle p, x \rangle, a, u', v', \langle q, y \rangle \rangle$$

that can be obtained from a transition $\langle p, a, u, v, q \rangle \in \Delta$ in one of the following ways.

- ♦ $x = zu$, $zv = y$, $|zv| < n$, $v' = u' \in \Gamma'$.
- ♦ $x = zu$, $zv = wy$, $|w| = n$, $u' \in \Gamma$, $v' = u'w$.
- ♦ $|x| < |u|$, $u'x = zu$, $zv = y$, $v' = \varepsilon$.

To conclude the proof, note that in all three cases above we have

$$|x|, |y| \leq 2n, \quad u'x = wu \quad \text{and} \quad v'y = wv, \quad \text{for some } w.$$

Consequently, the canonical map $(\Gamma^n)^* \rightarrow \Gamma^*$ induces an isomorphism between the configuration graph of \mathcal{B} and that of \mathcal{A} . \square

Theorem 5.3. *Let $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ be a graph. The following statements are equivalent.*

- (1) \mathfrak{G} is tree-interpretable.
- (2) \mathfrak{G} is isomorphic to an ε -pushdown graph.
- (3) \mathfrak{G} is isomorphic to a restricted ε -pushdown graph.

Proof. (2) \Rightarrow (3) follows immediately by Lemma 5.2.

(3) \Rightarrow (1) Let \mathfrak{G} be a graph associated with a restricted ε -pushdown automaton $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_0, \gamma_0, F \rangle$, let $Q_0 \subseteq Q$ be the set of all states with outgoing non- ε -transitions, and let \mathfrak{T} be the tree with domain $(\Gamma + Q)^*$. We encode a configuration $\langle q, z \rangle \in Q \times \Gamma^*$ by the word $zq \in (\Gamma + Q)^*$. Then

$$\mathfrak{G} \cong \tau(\alpha(\mathfrak{T}, \langle \rangle)),$$

where $\tau = \langle \delta, (\varphi_a)_{a \in A} \rangle$ is the MSO-interpretation with

$$\begin{aligned}\delta(x) &:= x \in \Gamma^+ Q_o \wedge 'x \text{ is reachable from the configuration } \gamma q_o.' \\ \varphi_a(x, y) &:= E_a x y,\end{aligned}$$

and $\alpha = \langle K, (L_a)_{a \in A} \rangle$ is the path interpretation where $K := \Gamma^+ Q$ and L_a is the set of all paths between two configurations u and v such that, when starting in u , \mathcal{A} can go to v by making an a -transition followed by an arbitrary number of ε -transitions. To see that the languages L_a are regular, it is sufficient to prove that the language of all paths consisting entirely of ε -transitions is regular. This language can be recognised by the following (non-pushdown) automaton \mathcal{B} with ε -transitions. The set of states is the same as that of \mathcal{A} . For every ε -transition $\langle p, \langle \rangle, u, v, q \rangle$ of \mathcal{A} the automaton \mathcal{B} has a transition

$$\begin{aligned}\langle p, c, q \rangle, & \quad \text{if } v = uc, \\ \langle p, \langle \rangle, q \rangle, & \quad \text{if } u = v, \\ \langle p, \bar{c}, q \rangle, & \quad \text{if } u = c \text{ and } v = \langle \rangle.\end{aligned}$$

(1) \Rightarrow (2) Suppose that $\mathfrak{G} = \alpha(\mathfrak{T}_{\text{bin}}, \langle \rangle)$, for a regular path-interpretation $\alpha = \langle K, (L_a)_{a \in A} \rangle$. Let $\mathcal{B} = \langle P, \{o, i\}, \Delta', p_o, F' \rangle$ be an automaton recognising K and $\mathcal{C}_a = \langle Q_a, \{o, i, \bar{o}, \bar{i}\}, \Delta_a, q_o^a, F_a \rangle$ one recognising L_a . We construct a pushdown automaton \mathcal{A} with input alphabet A , stack alphabet $\Gamma := \{\$, o, i\}$, and set of states

$$Q := * + P + \sum_{a \in A} Q_a, \quad \text{for some new state } *.$$

The initial configuration is $\langle p_o, \$ \rangle$. The transitions are as follows.

- $\langle p, z \rangle \xrightarrow{\langle \rangle} \langle q, zc \rangle$, for every transition $\langle p, c, q \rangle \in \Delta'$,
- $\langle p, z \rangle \xrightarrow{\langle \rangle} \langle *, z \rangle$, if $p \in F'$,
- $\langle *, z \rangle \xrightarrow{a} \langle q_o^a, z \rangle$, for $a \in A$,

- $\langle q, z \rangle \xrightarrow{\langle \rangle} \langle q', z' \rangle$, for every transition $\langle q, c, q' \rangle \in \Delta_a$ where,
 - $z' = zc$, if $c \in \{o, i\}$ (we push the letter c),
 - $z = z'c$, if $c \in \{\bar{o}, \bar{i}\}$ (we pop the letter c),

- $\langle q, z \rangle \xrightarrow{\langle \rangle} \langle *, z \rangle$, if $q \in F_a$,
- $\langle q, z \rangle \xrightarrow{\langle \rangle} \langle q, z \rangle$, if $q \neq *$.

(The transitions of the last kind are just needed to ensure that every configuration with a state $q \neq *$ has an outgoing ε -transition.) It follows that, in the corresponding ε -pushdown graph, there is an edge $\langle *, \$u \rangle \rightarrow^a \langle *, \$v \rangle$ if, and only if, there is a path from u to v in $\mathfrak{T}_{\text{bin}}$ labelled by some word in L_a . Furthermore, the reachable configurations with outgoing non- ε -transitions are precisely those of the form $\langle *, \$z \rangle$ with $z \in K$. It follows that \mathfrak{G} is equal to the ε -pushdown graph associated with \mathcal{A} . \square

The reason why we need ε -transitions is the fact that tree-interpretable graphs can have vertices with infinite degree. For finite degree, we can use pushdown automata without ε -transitions.

Definition 5.4. A graph $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ is a *pushdown graph* if it is isomorphic to the reachable subset of the configuration graph of an (unrestricted) pushdown automaton \mathcal{A} without ε -transitions. If \mathcal{A} is restricted, we call \mathfrak{G} a *restricted pushdown graph*. \lrcorner

Proposition 5.5. Let \mathfrak{G} be a directed edge-labelled graph. The following statements are equivalent.

- (1) \mathfrak{G} is a pushdown graph.
- (2) \mathfrak{G} is a restricted pushdown graph.
- (3) \mathfrak{G} is tree-interpretable and every vertex has finite degree.
- (4) \mathfrak{G} is isomorphic to a structure of the form $\langle V, (E_a)_{a \in A} \rangle$ where V is a regular language over some alphabet Σ and each edge relation is of the form

$$E_a = \bigcup_{i < n} W_i(\{u_i\} \times \{v_i\}),$$

for a regular language $W_i \subseteq \Sigma^*$ and words $u_i, v_i \in \Sigma^*$.

Proof. (1) \Leftrightarrow (2) follows by Lemma 5.2, and (4) \Rightarrow (3) by Theorem 2.2.

(3) \Rightarrow (4) Suppose that \mathfrak{G} is prefix-recognisable and let

$$E_a = \bigcup_{i < n} W_i(U_i \times V_i).$$

If every vertex has degree at most $k < \omega$, it follows that $|U_i|, |V_i| \leq k$. Hence, let

$$U_i = \{u_o^i, \dots, u_{m-1}^i\} \quad \text{and} \quad V_i = \{v_o^i, \dots, v_{n-1}^i\}.$$

Then

$$W_i(U_i \times V_i) = \bigcup_{j,l} W_i(\{u_j^i\} \times \{v_l^i\}).$$

(1) \Rightarrow (3) Every pushdown graph is an ε -pushdown graph. Furthermore, in every configuration c of a pushdown automaton, there are only finitely many possible transitions. (There are only finitely many possible new states and finitely many stack operations.) Similarly, there are only finitely many configurations with a transition leading to c . (There are only finitely many possible old states and finitely many stack operations producing the stack of c .)

(3) \Rightarrow (1) Suppose that $\mathfrak{G} = \alpha(\mathfrak{T}_{\text{bin}})$, for some path interpretation $\alpha = \langle K, (L_a)_a \rangle$. If \mathfrak{G} has finite degree, every language L_a is finite.

Let \mathcal{A} be the pushdown automaton with input alphabet A , stack alphabet $\{o, \mathfrak{i}\}$, a single state p and transitions $\langle p, z \rangle \xrightarrow{a} \langle p, z' \rangle$, where z' is the stack obtained from z by applying some path $w \in L_a$. Note that this transition is of the form $\langle p, xu \rangle \rightarrow \langle p, xv \rangle$ where u and v depend on w . Hence, it can indeed be performed by an (unrestricted) pushdown automaton. \square

Exercise 5.1. Prove that a tree-interpretable graph \mathfrak{G} has only finitely many nonisomorphic (i) connected components and (ii) strongly connected components. \downarrow

Higher-Order Pushdown Automata

The corresponding automaton model for the higher levels of the Caucal hierarchy uses nested stacks. That is, stacks are words in Γ^{+n} and we have push and pop operations for each level. Since pushing a fixed element of $\Gamma^{+(i-1)}$ onto the stack at level i does not make much sense, the push operation at higher levels will duplicate the top element of the stack instead.

Definition 5.6. Let Γ be an alphabet and $n < \omega$.

(a) For $z \in (\Gamma^{+k})^*$ and $c \in \Gamma^{+k}$, we use the notation

$$z : c := zc.$$

We assume that the operation $:$ is right-associative, i.e., $x : y : z = x : (y : z)$. It follows that we can write every $z \in \Gamma^{+n}$ uniquely in the form

$$z = z_n : z_{n-1} : \dots : z_1 : z_o, \quad \text{with } z_o \in \Gamma \text{ and } z_{k+1} \in (\Gamma^{+k})^*.$$

(b) The *level- n stack operations* over an alphabet Γ are the following operations on Γ^{+n} .

$$\text{nop}(z_n : \dots : z_o) := z_n : \dots : z_o,$$

$$\text{pop}_k(z_n : \dots : z_o) := z_n : \dots : z_k,$$

$$\text{push}_\gamma(z_n : \dots : z_o) := z_n : \dots : (z_1 z_o) : \gamma,$$

$$\text{clone}_l(z_n : \dots : z_o) := z_n : \dots : z_{l+1} : (z_l : z_{l-1} : \dots : z_o) : z_{l-1} : \dots : z_o,$$

for $\gamma \in \Gamma$, $1 \leq k \leq n$, and $2 \leq l \leq n$. We denote the set of these operations by Op_n . We also set

$$\text{top}(z_n : \dots : z_o) := z_o.$$

(c) A *pushdown automaton of level n* is a tuple

$$\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, \gamma_o, F \rangle$$

consisting of a finite set Q of states, an input alphabet A , a stack alphabet Γ , an initial state $q_o \in Q$, an initial stack symbol $\gamma_o \in \Gamma$, a set $F \subseteq Q$ of final states, and a transition relation

$$\Delta \subseteq Q \times A^{\leq 1} \times \Gamma \times \text{Op}_n \times Q.$$

A transition $\langle p, a, \gamma, o, q \rangle \in \Delta$ is called an ε -transition if $a = \langle \rangle$. We assume that, for every state $p \in Q$, if there are transitions

$$\langle p, a, \gamma, o, q \rangle, \langle p, a', \gamma', o', q' \rangle \in \Delta \quad \text{then} \quad a = \langle \rangle \Leftrightarrow a' = \langle \rangle.$$

If we do not want to specify the level n , we speak of a *higher-order pushdown automaton*.

(d) A configuration of a pushdown automaton \mathcal{A} of level n is a pair $\langle q, z \rangle$ consisting of a state $q \in Q$ and a stack $z \in \Gamma^{+n}$. Given a transition $\langle p, a, \gamma, o, q \rangle \in \Delta$ there is an a -labelled edge from the configuration $\langle p, z \rangle$ to $\langle q, o(z) \rangle$ provided that $\text{top}(z) = \gamma$.

(e) A ε -pushdown graph of level n is a graph \mathfrak{G} obtained from the configuration graph of a pushdown automaton of level n by

- adding an a -labelled edge between two configurations if there exists a path between them that consists of an a -transition followed by an arbitrary number of ε -transitions,
- removing all vertices with outgoing ε -transitions, and
- restricting the resulting graph to the set of configurations reachable from the initial one.

Example. For every n , there exists an automaton \mathcal{A}_n of level $n+1$ recognising the language

$$L_n := \{ a^{\exp_n(k)} \mid k < \omega \},$$

where $\exp_n(k)$ is the function defined by

$$\exp_o(k) := k \quad \text{and} \quad \exp_{n+1}(k) = 2^{\exp_n(k)}.$$

Informally, the automaton \mathcal{A}_n starts by guessing the number k and writing an encoding of $\exp_n(k)$ onto its stack. Then it enters a loop where in each iteration it decrements the number stored in the stack and reads one input letter. \mathcal{A}_n stops when the number on the stack becomes 0.

To encode such huge numbers into a stack of level $n+1$ we use the tack alphabet $\Gamma = \{1, \dots, n, a\}$. The bottom of a stack of level k will be marked by the level $k-1$ word

$$[k] := \langle \rangle : \dots : \langle \rangle : 1 \, 2 \, \dots \, (k-1) \quad k \in \Gamma^{+(k-1)}.$$

By induction on n , we define a coding function $\mu_n : \omega \rightarrow \Gamma^{+n}$ based on the binary encoding of integers.

$$\begin{aligned} \mu_1(m) &:= [1] a^m, \\ \mu_{n+1}(m) &:= [n+1] \mu_n(k_o) \cdots \mu_n(k_l), \\ &\text{where } m = 2^{k_o} + \dots + 2^{k_l} \text{ and } k_o > \dots > k_l. \end{aligned}$$

Instead of writing down the actual transition table of the automaton, we use pseudo-code to specify it. We need a predicate $\text{zero}_k(z)$ that is true if the top-most level k stack in z is empty, and we need a function $\text{dec}_k(z)$ that decrements the top-most level k stack of z . The predicate zero_k can be defined with the help of the markers $[k]$.

$$\text{zero}_k(z) \quad : \text{iff} \quad \text{top}(z) = [k].$$

For level 1 the numbers are stored on the stack in unary. Hence, the decrementation procedure only needs to remove one symbol.

$$\text{dec}_1(z) := \text{pop}_1(z).$$

For $k > 1$, we use binary encoding and $\text{dec}_k(z)$ has to distinguish two cases. If the last digit is 1 then we change it to 0. Otherwise, the number ends with

a sequence of digits $10 \cdots 0$ that we have to replace by $01 \cdots 1$.

```

deck+1(z) := if not zerok(z) then           // last digit is 1
               return popk+1(z)
            else                               // last digit is 0
               z := deck(z)                   // change 10 to 01
               while not zerok(z) do
                  z := (deck ◦ clonek+1)(z) // change 10 to 11
               end
            return z
        end
    end

```

The automaton \mathcal{A}_n works as follows. First, it creates the stack content

$$[n+1] : \cdots : [1].$$

Then nondeterministically it performs m push_a-operations. The stack contents now is

$$[n+1] : \cdots : [1] a^m = \mu_{n+1}(\exp_n(m)).$$

Finally, it enters a loop where in each iteration it reads one input letter and then calls dec_n.

The analogue of Theorem 5.3 for the Caucal hierarchy take the following form.

Theorem 5.7. *A graph $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ belongs to CH_n if, and only if, it is an ε -pushdown graph of level n .*

We split the proof of the theorem into several steps. As usual, one direction is straightforward.

Lemma 5.8. *Every ε -pushdown graph of level n belongs to CH_n .*

Proof. Let $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_0, \gamma_0, F \rangle$ be a pushdown automaton of level n and let \mathfrak{G} be the corresponding ε -pushdown graph. Set $s := |Q|$ and let $\mathfrak{C} := \langle \Gamma, (P_\gamma)_{\gamma \in \Gamma} \rangle$ be the finite structure with universe Γ and unary predicates $P_\gamma := \{\gamma\}$. It is sufficient to show that the configuration graph of \mathcal{A} is MSO-interpretable in $\mathfrak{A} := \text{copy}_s(\mathfrak{C}^{*n})$. Since \mathfrak{G} can be MSO-interpreted in the configuration graph, the lemma then follows from the closure properties of CH_n .

First, note that the configurations of \mathcal{A} correspond to a definable subset of \mathfrak{A} . (The elements not corresponding to configuration are those in $\text{C}^{*n} \setminus \text{C}^{+n}$.) Hence, it remains to define the edge relations. It is sufficient to show that, for every transition $\delta := \langle p, a, \gamma, o, q \rangle \in \Delta$, there exists an MSO-formula $\psi_\delta(x, y)$ expressing that the configuration y can be reached from x using the transition δ . To do so, note that the structure \mathfrak{A} has the form

$$\mathfrak{A} = \langle \Gamma^{*n}, (\text{suc}_i)_{i < n}, (\text{cl}_i)_{i < n}, (P_\gamma)_{\gamma \in \Gamma}, (H_q)_{q \in Q}, I \rangle,$$

where the relations suc_i and cl_i stem from the i -th application of the Muchnik iteration, the predicates H_q encode the s copies, I relates all copies of the same element, and P_γ checks that the top stack symbol is γ . The formula ψ_δ for $\delta := \langle p, a, \gamma, o, q \rangle \in \Delta$ consequently takes the following form.

- ♦ If $o = \text{pop}_k$, we set

$$\psi_\delta(x, y) := H_p x \wedge H_q y \wedge \exists y' [I y y' \wedge \text{suc}_{k-1}(y', x)].$$

- ♦ If $o = \text{push}_\gamma$, we set

$$\psi_\delta(x, y) := H_p x \wedge H_q y \wedge \exists y' [I y y' \wedge \text{suc}_o(x, y') \wedge P_\gamma y'].$$

- ♦ If $o = \text{clone}_k$, we set

$$\psi_\delta(x, y) := H_p x \wedge H_q y \wedge \exists y' [I y y' \wedge \text{suc}_{k-1}(x, y') \wedge \text{cl}_{k-1}(y')].$$

□

The converse is based on the following lemma.

Lemma 5.9. Let \mathfrak{G} be an ε -pushdown graph of level n and v_o a vertex. Then the expansion of the unravelling $\mathcal{U}(\mathfrak{G}, v_o)$ by inverse edges is an ε -pushdown graph of level $n + 1$.

Before giving the proof, which is a bit involved, let us show how to use it to prove Theorem 5.7.

Proof of Theorem 5.7. (\Rightarrow) has already been proved in Lemma 5.8.

(\Leftarrow) By the characterisation of CH_n from Theorem 4.4 and induction, it is sufficient to prove that, if \mathfrak{G} is an ε -pushdown graph of level n and α is a regular path-interpretation, then $\alpha(\mathcal{U}(\mathfrak{G}, v_o), v_o)$ is an ε -pushdown graph of level $n + 1$.

Hence, suppose that $\alpha = \langle K, (L_b)_{b \in B} \rangle$ and let

$$\mathcal{B}_b = \langle P_b, A + \bar{A}, \Delta_b, p_o, \{f\} \rangle$$

be a nondeterministic automaton recognising L_b . We may assume that

- the initial state p_o has no incoming transitions, and
- the unique final state f has no outgoing transitions.

Furthermore, we may assume that all automata \mathcal{B}_b share the same initial state p_o and the same final state f but that, apart from these, their state spaces are disjoint.

By Lemma 5.9, the extension of $\mathcal{U}(\mathfrak{G}, v_o)$ by inverse edges is an ε -pushdown graph of level $n + 1$. Let

$$\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, c_o, F \rangle$$

be the corresponding automaton.

We construct a pushdown automaton for $\alpha(\mathcal{U}(\mathfrak{G}, v_o), v_o)$ in three steps. First, we build a pushdown automaton \mathcal{A}_1 with states

$$Q_1 := Q + Q \times \bigcup_{b \in B} (P_b \setminus \{f\}).$$

The transitions are those of \mathcal{A} together with ε -transitions

$$\langle q, \langle \rangle, c, \text{nop}, \langle q, p \rangle \rangle, \quad \text{for all } q \in Q, p \in \bigcup_{b \in B} (P_b \setminus \{f\}), c \in \Gamma^{\leq 1}.$$




Figure 2: The construction of \mathcal{A}_2 for $L_c = a(a+b)b$. The automaton \mathcal{B}_c is on the top and (part of) the configuration graph of \mathcal{A}_2 is on the bottom. ε -transitions are dashed.

Consequently, the configuration graph of \mathcal{A}_1 consists of several disjoint copies of the graph of \mathcal{A} together with ε -transitions $\langle q, z \rangle \rightarrow \langle \langle q, p \rangle, z \rangle$ from every of the original configurations to each of its copies.

In the second step, we add transitions simulating the automata \mathcal{B}_b . For an example of the construction, see Figure 2. Let \mathcal{A}_2 be the automaton obtained from \mathcal{A}_1 by adding all transitions

$$\langle \langle q, p \rangle, \bar{a}, c, o, \langle q', \tilde{p} \rangle \rangle, \quad \text{with } \langle q, a, c, o, q' \rangle \in \Delta, \langle p, a, p' \rangle \in \Delta_b,$$

$$\bar{a} := \begin{cases} a & \text{if } p = p_o, \\ \varepsilon & \text{otherwise,} \end{cases} \quad \text{and} \quad \tilde{p} := \begin{cases} p_o & \text{if } p = f, \\ p & \text{otherwise.} \end{cases}$$

In case L_b contains the empty word, we also add the transition

$$\langle \langle q, p_o \rangle, b, c, \text{nop}, \langle q', p_o \rangle \rangle,$$

The construction ensures that there exists a path from configuration $\langle \langle q, p_o \rangle, z \rangle$ to $\langle \langle q', p_o \rangle, z' \rangle$ in the configuration graph of \mathcal{A}_2 that is labelled by $b \in B$ if, and only if, there is and a path from $\langle q, z \rangle$ to $\langle q', z' \rangle$ in the configuration graph of \mathcal{A}_1 that is labelled by some word $w \in L_b$.

In the final step, we replace all transitions $\langle q, a, c, o, q' \rangle$ of \mathcal{A}_2 that correspond to a transition of \mathcal{A} (that is, with $q, q' \in Q$) by ε -transitions $\langle q, \langle \rangle, c, o, q' \rangle$. Let \mathcal{A}_3 be the resulting automaton. Note that \mathcal{A}_3 is really a valid pushdown automaton since all configurations whose state belongs to $Q \times \{p_o\}$ have only outgoing non- ε -transitions, while all other configurations have only outgoing ε -transitions.

Contracting all ε -transitions of the configuration graph of \mathcal{A}_3 does not quite produce the graph $\alpha(\mathcal{U}(\mathfrak{G}, v_o), v_o)$ since there are still some paths starting in a configuration with state $\langle q, p_o \rangle$ that end in a configuration $\langle q', p' \rangle$ with $p' \neq p$. To eliminate these paths during the contraction, we add loops

$$\langle \langle q, p \rangle, \langle \rangle, c, \text{nop}, \langle q, p \rangle \rangle, \quad \text{for all } p \neq p_o.$$

The pushdown graph corresponding to the resulting automaton is then isomorphic to $\alpha(\mathcal{U}(\mathfrak{G}, v_o), v_o)$. \square

EC-Pushdown Graphs

It remains to provide the missing proof of Lemma 5.9. To do so, we introduce the following variant of our automaton model.

Definition 5.10. (a) A *pushdown automaton \mathcal{A} of level n with equality checks* has the same form as a pushdown automaton of level n , but it can only make a pop_k -transition with $k > 1$ if its stack is of the form $\text{clone}_k(z)$, for some z . (There is no restriction on pop_1 -transitions.) To avoid confusion, we will write such pop operations with a superscript pop_k^{ec} , and we write Op_n^{ec} for the version of Op_n with equality checks.

(b) An *ec-pushdown graph* is the graph obtained from the configuration graph of a pushdown automaton with equality checks by contracting ε -transitions as above. \square

We start by proving that there is no difference between ec-pushdown graphs and ε -pushdown graphs. One direction is relatively straightforward.

Lemma 5.11. *Every higher-level ε -pushdown graph is an ec-pushdown graph.*

Proof. Let $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, \gamma_o, F \rangle$ be an ordinary pushdown automaton of level n . We construct a pushdown automaton \mathcal{B} with equality checks with the same associated graph as follows. The transitions of \mathcal{B} are the same as those of \mathcal{A} , except when \mathcal{A} would perform a pop_k -transition with $k > 1$. In this case, \mathcal{B} first performs a sequence of ε -transitions that arbitrarily modify the top k levels of the stack before it can do the desired pop_k^{ec} -transition. It remains to prove that there exists a sequence of stack operations that leads to a stack of the form $\text{clone}_k(z)$, so that we really can perform the pop_k^{ec} .

We will prove the stronger statement that, for all words $x, y \in \Gamma^{+(k-1)}$ such that every level 1 word inside x and y starts with the initial stack symbol γ_o , there exists a sequence of stack operations transforming

$$z_n : \cdots : z_k : x \quad \text{into} \quad z_n : \cdots : z_k : y.$$

(Note that every configuration of \mathcal{A} reachable from the initial one is of this form, since we can never remove the bottom-most stack symbol γ_o .)

We proceed by induction on k . If $k = 1$, we can use a sequence of pop_1 operations to reduce x to γ_o , followed by a sequence of push-operations that add the missing symbols to γ_o to get y .

For the inductive step, suppose that $k > 1$. Again we first reduce x to the word $\langle \rangle : \cdots : \langle \rangle : \gamma_o$ and then construct the new word y . We construct the sequence of operations for the first part by induction on $|x|$. If $|x| = 1$, we have $x = \langle \rangle : u$ and we can use the inductive hypothesis on k to find a sequence of stack operations turning u into $\langle \rangle : \cdots : \langle \rangle : \gamma_o$. If $|x| > 1$, we have $x = w : u$ with $w \neq \langle \rangle$. By inductive hypothesis on k , there exists a sequence of operations producing $\text{clone}_{k-1}(w)$. After a $\text{pop}_{k-1}^{\text{ec}}$ -operation

we obtain w , which can be turned into $\langle \rangle : \cdots : \langle \rangle : \gamma_o$ by inductive hypothesis on $|x|$.

It remains to do the second part: transforming $\langle \rangle : \cdots : \langle \rangle : \gamma_o$ to y . Again we proceed by $|y|$. If $|y| = 1$, the claim follows immediately by inductive hypothesis. Hence, suppose that $|y| > 1$, say, $y = w : u$. By inductive hypothesis, we can turn $\langle \rangle : \cdots : \langle \rangle : \gamma_o$ into w . Performing a clone_{k-1} -operation, we get $\text{clone}_{k-1}(w)$. By inductive hypothesis on k , we can turn this word into $w : u$. \square

The other direction is more involved. We start by organising the stack configurations in a tree as follows. Let \mathbb{U} be the tree with domain Γ^* and let $\mathbb{U}_n := \mathbb{U}^{*'(n-1)}$, where $-^{*'}$ is the modified Muchnik iteration we introduced in Definition 4.7 above. Note that the vertices of \mathbb{U}_n are the words in Γ^{*n} and every edge corresponds to either a push_d or a clone_k operation. Let \mathfrak{S}_n be the tree obtained from \mathbb{U}_n by

- orienting every edge away from the respective root,
- removing all vertices that do not belong to Γ^{+n} ,
- removing all vertices that are not reachable from $\$ _n := \langle \rangle : \cdots : \langle \rangle : \$$, where $\$ \in \Gamma$ is the initial stack symbol,
- labelling every of the remaining edges by the corresponding operation in Op_n^{ec} .

Note that the resulting tree \mathfrak{S}_n has no edges labelled by pop_i^{ec} or by nop .

Lemma 5.12. *Let $z \in S_n$ and let $o \in \text{Op}_n^{\text{ec}}$ be an operation different from nop . If $o(z)$ is defined, it is either the predecessor of z in \mathfrak{S}_n or one of its successors.*

We will show that, for each stack z , there is a unique shortest sequence of operations generating z . First, let us make the following observations.

- Every stack operation except for pop_1 is injective.
- If $o \neq o'$ are distinct operations such that $o(z)$ and $o(z')$ are defined, then $o(z) \neq o'(z)$ for every z .

- For every stack operation $o \in \text{Op}_n^{\text{ec}}$, there exists an inverse $o^{-1} \in \text{Op}_n^{\text{ec}}$.

$$(\text{push}_c)^{-1}(z) = \text{pop}_1(z), \quad (\text{pop}_1^{\text{ec}})^{-1}(z) = \text{push}_{\text{top}(z)}(z)$$

$$(\text{clone}_k)^{-1}(z) = \text{pop}_k^{\text{ec}}(z), \quad (\text{pop}_k^{\text{ec}})^{-1}(z) = \text{clone}_k(z).$$

Note that $(\text{pop}_i^{\text{ec}})^{-1}$ is not uniquely defined since it depends on the top-most stack symbol. In the following we will use the short-hand $o = (\text{pop}_i^{\text{ec}})^{-1}$ to mean that $o = \text{push}_\gamma$, for some γ .

Every sequence of operations mapping z to z' forms a walk in \mathfrak{S}_n between these to vertices. The shortest such sequence is the unique path from z to z' .

Lemma 5.13. *Let $z, z' \in S_n$. There exists a unique shortest sequence of stack operations mapping z to z' .*

Proof. Let $(z_i)_{i \leq m}$ be the unique path from z to z' in \mathfrak{S}_n (ignoring edge directions), and let o_i be the operation mapping z_i to z_{i+1} . By Lemma 5.12, there is no shorter sequence mapping z to z' than $(o_i)_{i < m}$. To prove uniqueness, consider a second sequence $(o'_i)_{i < m}$ of m operations mapping z to z' . Let k be the least index with $o'_k \neq o_k$. Then $o'_k(z_k) \neq o_k(z_k) = z_{k+1}$ implies that either $o'_k(z_k) = z_{k-1}$ or $o'_k(z_k)$ does not lie on the path between z and z' . In both cases, it follows that the path from $o'_k(z_k)$ to z' in \mathfrak{S}_n has the length $m - k + 1$. Consequently, we need at least $m - k + 1$ operations mapping $o'_k(z_k)$ to z' . But $(o'_i)_{k < i < m}$ is supposed to map $o'_k(z_k)$ to z' . A contradiction. \square

Example. Without equality checks there are pairs of stacks that are connected by several shortest sequences. For instance, we have

$$ab : abc : d \xrightarrow{\text{pop}_1} ab : ab : c \xrightarrow{\text{pop}_1} ab : a : b,$$

$$ab : abc : d \xrightarrow{\text{pop}_2} \langle \rangle : a : b \xrightarrow{\text{clone}_a} ab : a : b.$$

Those sequences that are not the shortest can be characterised as follows.

Definition 5.14. The set of *Dyck sequences* is the subset of $(\text{Op}_n^{\text{ec}})^*$ that is inductively defined as follows.

- The empty sequence is a Dyck sequence.

- ♦ If w and w' are Dyck sequences and $o \in \text{Op}_n^{\text{ec}}$, then owo^{-1} and ww' are also Dyck sequences. \square

Lemma 5.15. Let $(o_i)_{i < m}$ be a sequence of stack operations. The following statements are equivalent.

- (1) $(o_i)_i$ is a Dyck sequence.
- (2) $\emptyset \neq o_{m-1} \circ \dots \circ o_0 \subseteq \text{id}$.
- (3) $o_{m-1} \circ \dots \circ o_0$ has a fixed point.

Proof. (1) \Rightarrow (2) The second inclusion follows by a straightforward induction on m . For the first one, note that it follows by Lemma 5.12 that every Dyck sequence is the label of some closed walk p in the tree \mathfrak{S}_n . Let z be the vertex of p where this walk starts and ends. Then $(o_{m-1} \circ \dots \circ o_0)(z) = z$.

(2) \Rightarrow (3) is trivial.

(3) \Rightarrow (1) Let $z \in \Gamma^{+n}$ be a fixed point of $(o_i)_i$. For $i \leq m$, set

$$z_i := (o_{i-1} \circ \dots \circ o_0)(z).$$

We prove the claim by induction on m . If $m = 0$, then $(o_i)_{i < m}$ is a Dyck sequence. If there is some index $0 < k < m$ with $z_k = z$, it follows by inductive hypothesis that $(o_i)_{i < k}$ and $(o_i)_{k \leq i < m}$ are Dyck sequences. Hence, so is their concatenation $(o_i)_{i < m}$. Similarly, if there are indices $0 < k < l < m$ with $z_k = z_l$, it follows by inductive hypothesis that $o_0, \dots, o_{k-1}, o_l, \dots, o_{m-1}$ and $(o_i)_{k \leq i < l}$ are Dyck sequences. Hence, so is the sequence obtained by inserting the latter into the former. Consequently, we may assume that $m > 0$ and that $z_i \neq z_j$, for all $0 \leq i < j < m$. We claim that $o_{m-1} = o_0^{-1}$. Then $z_{m-1} = z_0$ and it follows by inductive hypothesis that $(o_i)_{0 < i < m-1}$ is a Dyck sequence. Hence, so is $(o_i)_{i < m}$.

It therefore remains to prove the claim. We denote by $d(x, y)$ the distance between x and y in the tree \mathfrak{S}_n . By Lemma 5.12, the sequence z_0, \dots, z_m forms a (closed) walk in \mathfrak{S}_n . This implies that

$$d(\$_n, z_i) - d(\$_n, z_{i+1}) = \pm 1, \quad \text{for all } i.$$

Let $k < m$ be the index such that $d(\$_n, z_k)$ is maximal. Then

$$d(\$_n, z_{k-1}) = d(\$_n, z_k) - 1 = d(\$_n, z_{k+1}).$$

This implies that there are two shortest paths from $\$n$ to z_k : the shortest paths from $\$n$ to z_{k-1} followed by o_k , and the shortest paths from $\$n$ to z_{k+1} followed by o_{k+1}^{-1} . These paths are really different since $z_{k-1} \neq z_{k+1}$. A contradiction to the uniqueness result from Lemma 5.13. \square

Lemma 5.16. Let $z, z' \in S_n$. Every sequence of stack operations mapping z to z' can be obtained from the shortest one by inserting a Dyck sequence after every operation.

Proof. Let $(o_i)_{i < m}$ be some sequence mapping z to z' and set

$$z_i := (o_{i-1} \circ \dots \circ o_0)(z), \quad \text{for } i \leq m.$$

We prove the claim by induction on m . It follows by Lemma 5.12 that the sequence $(z_i)_i$ forms a walk in the tree \mathfrak{S} . In particular, every stack $w \in S_n$ that lies on the path from z to z' must occur in $(z_i)_i$. The subsequence between the first occurrence of w and the last one forms a Dyck sequence by Lemma 5.15. Removing all these subsequences, we obtain a path that contains each of the stacks w exactly once. Hence, the remaining sequence forms the shortest path from z to z' . \square

To simulate an automaton with equality checks by one without, we annotate each stack by additional information that can be used to determine whether we are allowed to perform a pop_k^{ec} operation. Since every operation of the original automaton is simulated by possibly several operations, we introduce the following convenient notation for sequences of operations that can be performed by a pushdown automaton.

Definition 5.17. Let Γ be an alphabet and let $n < \omega$.

(a) A *stack-aware* stack operation of level n is an expression built up from the operations in Op_n using (I) composition \circ , (II) union $+$, and (III) tests $?\gamma$, for $\gamma \in \Gamma$.

(b) Let o be a stack-aware operation and $z \in \Gamma^{+n}$. We define the value $o(z)$ by induction on o as follows.

$$\begin{aligned} o(z) &:= o(z), & \text{for } o \in \text{Op}_n, \\ (p \circ q)(z) &:= p(q(z)), \end{aligned}$$

$$\begin{aligned} ?_y(z) &:= \begin{cases} z & \text{if } \text{top}(z) = y, \\ \text{undefined} & \text{otherwise,} \end{cases} \\ (p + q)(z) &:= \begin{cases} p(z) & \text{if } p(z) \text{ is defined but } q(z) \text{ is not,} \\ q(z) & \text{if } q(z) \text{ is defined but } p(z) \text{ is not,} \\ \text{undefined} & \text{otherwise.} \end{cases} \end{aligned}$$

Remark. Clearly, every stack-aware stack operation can be performed by a pushdown automaton.

We define an encoding $\llbracket z \rrbracket_n$ of stacks $z \in \Gamma^{+n}$ and an encoding $\llbracket o \rrbracket_n$ of operations $o \in \text{Op}_n^{\text{ec}}$ such that

$$\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) = \llbracket o(z) \rrbracket_n.$$

To encode a stack z , we consider the shortest sequence of operations producing z and annotate the sequence of intermediate stacks by the operation producing them. Consequently, we will use the stack alphabet $\Gamma + \Delta$, where Γ is the original stack alphabet and Δ is the set of all stack operations.

For technical reasons, we assume that the automaton never pushes a second copy of the initial stack symbol $\$$ onto the stack. (We can transform every pushdown automaton into this form by adding a second copy $\$'$ of $\$$ to Γ .)

Definition 5.18. Let Γ be an alphabet, $n < \omega$, and let $\$ \in \Gamma$ be the initial stack symbol.

- (a) We write $\$ _n := \langle \rangle : \dots : \langle \rangle : \$$.
- (b) Set $\Delta := \text{Op}_n^{\text{ec}} \setminus \{\text{pop}_n, \text{nop}\}$.
- (c) Let $z \in \Gamma^{+n}$ and let $(o_i)_{i < k}$ be the shortest sequence of stack operations mapping $\$ _n$ to z . We denote by y_i the top-most level $n - i$ stack of $(o_{i-1} \circ \dots \circ o_0)(\$ _n)$. We define the encoding $\llbracket z \rrbracket_n \in (\Gamma + \Delta)^{+n}$ by induction on n .

$$\begin{aligned} \llbracket z \rrbracket_1 &:= z, \\ \llbracket z \rrbracket_n &:= \langle \$_{n-1}, \text{push}_{o_0}(\llbracket y_0 \rrbracket_{n-1}), \dots, \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}) \rangle. \end{aligned}$$

- (d) We encode a stack operation $o \in \text{Op}_n^{\text{ec}}$ as a stack-aware operation $\llbracket o \rrbracket_n$ by induction on n as follows.

$$\begin{aligned} \llbracket o \rrbracket_1 &:= o. \\ \llbracket \text{clone}_n \rrbracket_n &:= \sum_{p \in \Delta} \text{push}_{\text{clone}_n} \circ \text{pop}_1 \circ \text{clone}_n \circ ?_p \\ &\quad + \text{push}_{\text{clone}_n} \circ \text{clone}_n \circ ?_{\$ _n}, \\ \llbracket \text{pop}_n^{\text{ec}} \rrbracket_n &:= \text{pop}_n \circ ?_{\text{clone}_n}, \\ \llbracket \text{nop} \rrbracket_n &:= \text{nop}, \end{aligned}$$

and, for $o \in \text{Op}_{n-1}^{\text{ec}}$ different from nop ,

$$\begin{aligned} \llbracket o \rrbracket_n &:= \sum_{p \in \Delta \setminus \{o^{-1}\}} \text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{pop}_1 \circ \text{clone}_n \circ ?_p \\ &\quad + \text{pop}_n \circ ?_{o^{-1}} \\ &\quad + \text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{clone}_n \circ ?_{\$ _n}. \end{aligned}$$

Lemma 5.19. Let $z \in \Gamma^{+n}$ be a stack and $o \in \text{Op}_n^{\text{ec}}$ an operation.

$$\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) = \begin{cases} \llbracket o(z) \rrbracket_n & \text{if } o(z) \text{ is defined,} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Proof. We prove the claim by induction on n . If $n = 1$, we have

$$\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) = o(z) = \llbracket o(z) \rrbracket_n.$$

Hence, suppose that $n > 1$. We distinguish two cases.

- (1) First, suppose that $z = \$ _n$ is the empty stack. Then

$$\begin{aligned} \llbracket \text{clone}_n \rrbracket_n(\llbracket z \rrbracket_n) &= (\text{push}_{\text{clone}_n} \circ \text{clone}_n)(\$ _n) \\ &= \langle \$_{n-1}, \text{push}_{\text{clone}_n}(\$_{n-1}) \rangle \\ &= \llbracket \text{clone}_n(z) \rrbracket_n, \end{aligned}$$

$$\begin{aligned}
\llbracket \text{pop}_n^{\text{ec}} \rrbracket_n(\llbracket z \rrbracket_n) &= \text{undefined} \\
&= \llbracket \text{pop}_n^{\text{ec}}(z) \rrbracket_n, \\
\llbracket \text{nop} \rrbracket_n(\llbracket z \rrbracket_n) &= \llbracket z \rrbracket_n \\
&= \llbracket \text{nop}(z) \rrbracket_n.
\end{aligned}$$

For $o \in \text{Op}_{n-1}^{\text{ec}}$, we have

$$\begin{aligned}
&\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) \text{ is defined} \\
&\text{iff } (\text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{clone}_n \circ ?_{\$_n})(\$_n) \text{ is defined} \\
&\text{iff } \llbracket o \rrbracket_{n-1}(\langle \$_{n-1}, \$_{n-1} \rangle) \text{ is defined} \\
&\text{iff } \llbracket o \rrbracket_{n-1}(\$_{n-1}) \text{ is defined} \\
&\text{iff } \llbracket o \rrbracket_{n-1}(\llbracket \$_{n-1} \rrbracket_{n-1}) \text{ is defined} \\
&\text{iff } o(\$_{n-1}) \text{ is defined} \\
&\text{iff } o(\$_n) \text{ is defined.}
\end{aligned}$$

Furthermore, if $z' := o(\$_n)$ is defined, the shortest sequence of operations mapping e to z' consists just of the operation o and we have

$$\begin{aligned}
\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) &= (\text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{clone}_n \circ ?_{\$_n})(\$_n) \\
&= \langle \$_{n-1}, \text{push}_o(\llbracket o \rrbracket_{n-1}(\$_{n-1})) \rangle \\
&= \langle \$_{n-1}, \text{push}_o(\llbracket o(\$_{n-1}) \rrbracket_{n-1}) \rangle \\
&= \llbracket z' \rrbracket_n \\
&= o(\$_n).
\end{aligned}$$

(II) Finally, suppose that $z \neq \$_n$ and let

$$\llbracket z \rrbracket_n := \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}) \rangle.$$

for some $y_o, \dots, y_{k-1} \in \Gamma^{+(n-1)}$ and $o_o, \dots, o_{k-1} \in \text{Op}_{n-1}^{\text{ec}}$. We distinguish several cases.

(II.1) First, suppose that $o = \text{pop}_n^{\text{ec}}$. Note that, if

$$\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) = (\text{pop}_n \circ ?_{\text{clone}_n})(\llbracket z \rrbracket_n)$$

is defined, we must have $o_{k-1} = \text{clone}_n$ and it follows that

$$\text{pop}_n^{\text{ec}}(z) = \text{pop}_n^{\text{ec}}(\text{clone}_n(z')), \quad \text{for some } z',$$

is also defined.

Conversely, suppose that $\text{pop}_n^{\text{ec}}(z)$ is defined. We have to show that so is $\llbracket o \rrbracket_n(\llbracket z \rrbracket_n)$ and that the corresponding values are the same. First, note that $o_{k-1} = \text{clone}_n$ since, otherwise, the shortest sequence mapping $\$ _n$ to $\text{pop}_n^{\text{ec}}(z)$ would be o_o, \dots, o_{k-1}, o . But the tree \mathfrak{S}_n has no edges labelled by pop_n^{ec} . A contradiction. Hence, $\text{pop}_n^{\text{ec}}(z)$ has the shortest sequence o_o, \dots, o_{k-2} and we have

$$\begin{aligned}
\llbracket o(z) \rrbracket_n &= \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-2}}(\llbracket y_{k-2} \rrbracket_{n-1}) \rangle \\
&= \text{pop}_n(\llbracket z \rrbracket_n) \\
&= \llbracket \text{pop}_n^{\text{ec}} \rrbracket_n(\llbracket z \rrbracket_n).
\end{aligned}$$

Hence, $\llbracket o(z) \rrbracket_n$ is defined and its value is equal to $\llbracket \text{pop}_n^{\text{ec}} \rrbracket_n(\llbracket z \rrbracket_n)$.

(II.2) Suppose that $o = \text{clone}_n$. Note that

$$\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) = (\text{push}_{\text{clone}_n} \circ \text{pop}_1 \circ \text{clone}_n)(\llbracket z \rrbracket_n)$$

and $\text{clone}_n(z)$ are always defined. Furthermore, the shortest sequence mapping $\$ _n$ to $\text{clone}_n(z)$ is o_o, \dots, o_{k-1}, o since, otherwise, we would have $o_{k-1} = o^{-1} = \text{pop}_n^{\text{ec}}$. A contradiction to the fact that the tree \mathfrak{S}_n contains no edges with label pop_n^{ec} . Hence,

$$\begin{aligned}
\llbracket o(z) \rrbracket_n &= \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}), \\
&\quad \text{push}_{\text{clone}_n}(\llbracket y_{k-1} \rrbracket_{n-1}) \rangle \\
&= (\text{push}_{\text{clone}_n} \circ \text{pop}_1 \circ \text{clone}_n)(\llbracket z \rrbracket_n) \\
&= \llbracket \text{clone}_n \rrbracket_n(\llbracket z \rrbracket_n).
\end{aligned}$$

(II.3) If $o = \text{nop}$, the claim is trivial.

(11.4) Suppose that $o = o_{k-1}^{-1} \in \text{Op}_{n-1}^{\text{ec}}$. Then $o(z)$ is defined and the shortest sequence mapping $\$ _n$ to $o(z)$ is o_o, \dots, o_{k-2} . Hence,

$$\begin{aligned} \llbracket o(z) \rrbracket_n &= \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-2}}(\llbracket y_{k-2} \rrbracket_{n-1}) \rangle \\ &= \text{pop}_n(\langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \\ &\quad \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}) \rangle) \\ &= \llbracket o \rrbracket_n(\llbracket z \rrbracket_n). \end{aligned}$$

It follows that $\llbracket o \rrbracket_n(\llbracket z \rrbracket_n)$ is also defined and equal to $\llbracket o(z) \rrbracket_{n-1}$.

(11.5) Finally, suppose that $o \in \text{Op}_{n-1}^{\text{ec}}$ and $o \neq o_{k-1}^{-1}$, and let $z = z' : x$. Then

$$\begin{aligned} &\llbracket o \rrbracket_n(\llbracket z \rrbracket_n) \text{ is defined} \\ \text{iff } &(\text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{pop}_1 \circ \text{clone}_n \circ ?_{o_{k-1}})(\llbracket z \rrbracket_n) \\ &\text{is defined} \\ \text{iff } &\llbracket o \rrbracket_{n-1}(\llbracket y_{k-1} \rrbracket_{n-1}) \text{ is defined} \\ \text{iff } &\llbracket o \rrbracket_{n-1}(\llbracket x \rrbracket_{n-1}) \text{ is defined} \\ \text{iff } &o(x) \text{ is defined} \\ \text{iff } &o(z) \text{ is defined.} \end{aligned}$$

Furthermore, if $z'' := o(z)$ is defined, then $o_{k-1} \neq o^{-1}$ implies that the shortest sequence of operations mapping $\$ _n$ to z'' is o_o, \dots, o_{k-1}, o . Hence,

$$\begin{aligned} \llbracket o \rrbracket_n(\llbracket z \rrbracket_n) &= (\text{push}_o \circ \llbracket o \rrbracket_{n-1} \circ \text{pop}_1 \circ \text{clone}_n \circ ?_{o_{k-1}})(\llbracket z \rrbracket_n) \\ &= \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}), \\ &\quad \text{push}_o(\llbracket o \rrbracket_{n-1}(\llbracket y_{k-1} \rrbracket_{n-1})) \rangle \\ &= \langle \$_{n-1}, \text{push}_{o_o}(\llbracket y_o \rrbracket_{n-1}), \dots, \text{push}_{o_{k-1}}(\llbracket y_{k-1} \rrbracket_{n-1}), \\ &\quad \text{push}_o(\llbracket o(y_{k-1}) \rrbracket_{n-1}) \rangle \\ &= \llbracket o(z) \rrbracket_n. \end{aligned} \quad \square$$

Corollary 5.20. Every ec-pushdown graph is an ε -pushdown graph.

Proof. Let $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, \$, F \rangle$ be a pushdown automaton with equality checks of level n . W.l.o.g. we may assume that \mathcal{A} never pushes a second copy of the initial stack symbol $\$$ onto the stack. We construct an ordinary pushdown automaton \mathcal{B} with the same associated graph as follows. The automaton \mathcal{B} simulates \mathcal{A} by replacing every a -transition using an operation $o \in \text{Op}^{\text{ec}}$ by a sequence of an a -transition followed by several ε -transitions that together perform the operation $\llbracket o \rrbracket_n$. It follows by Lemma 5.19 that the configuration graph of \mathcal{B} is obtained from that of \mathcal{A} by replacing every transition by a sequence of transitions (the first one with the original label and all other ones ε -transitions). Consequently, the pushdown graph encoded by \mathcal{B} is isomorphic to that of \mathcal{A} . \square

We are finally able to give the missing proof of Lemma 5.9.

Proof of Lemma 5.9. By Lemma 5.11, we may assume that \mathfrak{G} is an ec-pushdown graph. Hence, let $\mathcal{A} = \langle Q, A, \Gamma, \Delta, q_o, \$, F \rangle$ be a pushdown automaton with equality checks of level n that corresponds to \mathfrak{G} . We have to construct a new automaton \mathcal{B} representing $\mathcal{U}(\mathfrak{G}, v_o)$. By Corollary 5.20, it is sufficient to construct an automaton with equality checks. This automaton \mathcal{B} will encode a path

$$\langle p_o, z_o \rangle \xrightarrow{a_o} \dots \xrightarrow{a_{n-1}} \langle p_n, z_n \rangle$$

in the configuration graph of \mathcal{A} by the single configuration

$$\langle p_n, \hat{z}_o \dots \hat{z}_{n-1} z_n \rangle \quad \text{where} \quad \hat{z}_i := \text{push}_{p_i}(\text{push}_{a_i}(z_i)).$$

To do so, we use the stack alphabet $\Sigma := \Gamma + Q + A + \{e\}$. (The new symbol e is used for the case where $a_i = \langle \rangle$.) We have to define \mathcal{B} such that, for every a -transition of \mathcal{A} from $\langle p, z \rangle$ to $\langle p', z' \rangle$, there exists a sequence of transitions in \mathcal{B} from

$$\langle p, wz \rangle \quad \text{to} \quad \langle p', w\hat{z}z' \rangle, \quad \text{where} \quad \hat{z} := \text{push}_p(\text{push}_a(z)),$$

whose labelling is a . In addition we need a sequence of transitions in the other direction with label \bar{a} . This can be done by replacing every transition

$\langle p, a, \gamma, o, q \rangle$ of \mathcal{A} by a sequence of 6 transitions

$\langle p, a, \gamma, \text{push}_a, p_1 \rangle,$	(push a)
$\langle p_1, \langle \rangle, a, \text{push}_p, p_2 \rangle,$	(push p)
$\langle p_2, \langle \rangle, p, \text{clone}_n, p_3 \rangle,$	(extend the path)
$\langle p_3, \langle \rangle, p, \text{pop}_1, p_4 \rangle,$	(remove the p from the copy)
$\langle p_4, \langle \rangle, a, \text{pop}_1, p_5 \rangle,$	(remove the a from the copy)
$\langle p_5, \langle \rangle, \gamma, o, q \rangle,$	(perform the original stack operation)

where p_1, \dots, p_5 are new states. An ε -transition is translated analogously with the first transition now being

$$\langle p, \langle \rangle, \gamma, \text{push}_e, p_1 \rangle.$$

(We use a special symbol e since we cannot push $\langle \rangle$.) For the inverse direction, we add the transitions

$$\begin{aligned} &\langle q, \bar{a}, d, o^{-1}, p'_1 \rangle, \\ &\langle p'_1, \langle \rangle, \gamma, \text{push}_a, p'_2 \rangle, \\ &\langle p'_2, \langle \rangle, a, \text{push}_p, p'_3 \rangle, \\ &\langle p'_3, \langle \rangle, p, \text{pop}_n^{\text{ec}}, p'_4 \rangle, \\ &\langle p'_4, \langle \rangle, p, \text{pop}_1, p'_5 \rangle, \\ &\langle p'_5, \langle \rangle, a, \text{pop}_1, q \rangle, \end{aligned}$$

where p'_1, \dots, p'_5 are new states, $d \in \Gamma^{\leq}$ is arbitrary, and o^{-1} denotes the inverse to the stack operation o .

As defined above, the configuration graph of \mathcal{B} represents the unravelling of \mathcal{G} from the vertex corresponding to the initial configuration of \mathcal{A} . To obtain the unravelling from the specified vertex v_o , we modify \mathcal{B} to first construct the configuration corresponding to v_o using a sequence of ε -transitions.

Finally, note that the construction of \mathcal{B} does not introduce new pop_k^{ec} -operations with $k > 1$, except for the inverse edges where we use a pop_n^{ec} to

undo a clone_n . Here, the equality check is needed to ensure that the current stack is really of the form $\text{clone}_n(z)$. \square

6 Context-Free Groups

The first examples of algorithmic questions regarding infinite structures come from group theory. Consequently various kinds of finite presentations of groups are well-studied. Here we present several characterisations of the groups that belong to the Caucal hierarchy.

There are two different ways to represent finitely generated groups as structures. Either multiplication is given as binary function or one just includes several unary functions for multiplication by one of the generators. If the first version is chosen, it turns out that only finite groups are in the Caucal hierarchy.

Proposition 6.1. *A group $\mathcal{G} = \langle G, \cdot \rangle$ is in the Caucal hierarchy if, and only if, it is finite.*

Proof. We have shown in Corollary IX.7.3 that no infinite group has finite crossing-width, and in Lemma 1.14 that every structure in the Caucal hierarchy has finite crossing-width. \square

Example. $\langle \mathbb{Z}, + \rangle \notin \text{CH}$. \lrcorner

The second type of presentation is called the *Cayley-graph* of a group.

Definition 6.2. Let \mathcal{G} be a group and $S \subseteq G$ a set generating \mathcal{G} as a semi-group. The *Cayley-graph* of \mathcal{G} is the structure

$$\Gamma(\mathcal{G}, S) := \langle G, (f_a)_{a \in S} \rangle \quad \text{with} \quad f_a(x) := xa.$$

Example. Let \mathcal{G} be the free group with two generators a and b and set $S := \{a, b, a^{-1}, b^{-1}\}$. The Cayley-graph $\Gamma(\mathcal{G}, S)$ is tree-interpretable. The universe consists of all words over S which are reduced, that is, they do not contain one of the following factors:

$$aa^{-1}, \quad a^{-1}a, \quad bb^{-1}, \quad b^{-1}b.$$

Multiplication by a takes words w not ending in a^{-1} to wa and words of the form $w = ua^{-1}$ to u . Hence, we can write f_a in the form

$$f_a = (\langle \rangle \times a) + S^*(a \times aa) + S^*(a^{-1} \times \langle \rangle) \\ + S^*(b \times ba) + S^*(b^{-1} \times b^{-1}a),$$

which is prefix-recognisable. The other generators can be defined similarly. It follows that $\Gamma(\mathfrak{G}, S)$ is tree-interpretable. \downarrow

As $\Gamma(\mathfrak{G}, S)$ is MSO-interpretable in \mathfrak{G} , the requirement that the Cayley-graph is tree-interpretable is weaker than the one that \mathfrak{G} is. It turns out that we indeed obtain a non-trivial class of groups using this representation. We will show that the class of tree-interpretable groups coincides with the class of so-called *context-free groups*.

Definition 6.3. Let \mathfrak{G} be a group.

(a) \mathfrak{G} is *context-free* if there exists a finite set $S \subseteq G$ of semigroup generators such that the set $\{w \in S^* \mid \pi(w) = 1\}$ forms a context-free language, where $\pi : S^* \rightarrow G$ is the function multiplying a sequence of generators.

(b) \mathfrak{G} is *virtually free* if it has a free group of finite index. \downarrow

The class of context-free groups is well investigated and has many different characterisations.

Theorem 6.4. Let \mathfrak{G} be a finitely generated group and S a finite set of semigroup generators. The following statements are equivalent.

- (1) \mathfrak{G} is context-free.
- (2) \mathfrak{G} is virtually free.
- (3) $\Gamma(\mathfrak{G}, S)$ has only finitely many ends up to isomorphism.
- (4) $\Gamma(\mathfrak{G}, S)$ is a pushdown graph.
- (5) $\Gamma(\mathfrak{G}, S)$ is κ -triangulable, for some finite κ .
- (6) $\Gamma(\mathfrak{G}, S)$ is tree-interpretable.
- (7) $\Gamma(\mathfrak{G}, S)$ belongs to the Caucal hierarchy.
- (8) $\Gamma(\mathfrak{G}, S)$ has finite tree-width.

(9) $\Gamma(\mathfrak{G}, S)$ has finite crossing-width.

(10) $\Gamma(\mathfrak{G}, S)$ has a decidable MSO-theory.

(11) $\Gamma(\mathfrak{G}, S)$ has a decidable CGSO-theory.

Remark. There exist similar characterisations for finitely generated groups \mathfrak{G} whose word problem $W := \pi^{-1}(e)$ has a different complexity.

- (a) W is regular if, and only if, \mathfrak{G} is finite. [5, 144]
- (b) W is decidable if, and only if, \mathfrak{G} can be embedded in a simple subgroup of some finitely presented group. [31]
- (c) W is computably enumerable if, and only if, \mathfrak{G} can be embedded in some finitely presented group. [101] \downarrow

We split the proof of the theorem into several lemmas, which consitute the remainder of this section.

Triangulations

Our main tool in the proof of Theorem 6.4 is the fact that every cycle in the Cayley graph has a triangulation of the following form.

Definition 6.5. Let \mathfrak{G} be an undirected graph and $\kappa < \omega$.

(a) For vertices u, v , we denote by $d(u, v)$ the length of the shortest path between u and v .

(b) A *triangulation* of a sequence v_0, \dots, v_{n-1} is a tree-decomposition $(U_t)_{t \in T}$ of the graph with vertices v_0, \dots, v_{n-1} and edges $\langle v_i, v_{i+1} \rangle$, for $i < n$ (index arithmetic modulo n), such that $|U_t| = 3$, for all t . For $n < 3$, we drop the requirement that $|U_t| = 3$ and allow the decomposition with a single component $\{v_0, \dots, v_{n-1}\}$. We call each component U_t a *triangle* of the triangulation.

(c) Let v_0, \dots, v_{n-1} be vertices of \mathfrak{G} . A κ -*triangulation* is a triangulation $(U_t)_{t \in T}$ of v_0, \dots, v_{n-1} such that

$$u, v \in U_t \text{ implies } d(u, v) \leq \kappa, \text{ for all } t \in T.$$

(d) We say that \mathfrak{G} is κ -*triangulable* if every cycle of \mathfrak{G} has a κ -triangulation.

(e) Let T be an undirected tree and $\langle s, t \rangle$ an edge of T . We denote by T_{st} the subtree of T consisting of all vertices $u \in T$ such that the path from u to s does not contain the vertex t . \downarrow

The following technical remark will be used in the proof below.

Lemma 6.6. *Let $(U_t)_{t \in T}$ be a triangulation of v_0, \dots, v_{n-1} with $n \geq 3$. For all $0 < k < l < n$, there exists a triangle U_t that intersects each of the sets*

$$A := \{v_0, \dots, v_k\}, \quad B := \{v_k, \dots, v_l\}, \quad C := \{v_l, \dots, v_{n-1}, v_0\}.$$

Proof. Let $s, t \in T$ be indices such that $v_0 \in U_s$ and $v_k \in U_t$. Then U_s intersects A and C . If it also intersects B , we are done. Hence, suppose otherwise. Let $u \in T$ be the first index on the path from s to t such that $U_u \cap B \neq \emptyset$, and let u' be the vertex before u on this path. Since B is connected, it follows that all elements of B appear in the subtree $T_{uu'}$. In particular, it follows that there are $w, w' \in T_{uu'}$ with $v_k \in U_w$ and $v_l \in U_{w'}$. Hence, U_w intersects A and $U_{w'}$ intersects C . As A and C are connected and U_s intersects both A and C , it follows that every component in between also intersects, respectively, A or C . In particular, U_u does. Hence, U_u intersects all sets A, B, C . \square

Let us also prove that finite tree-width implies triangulability. We will use the following special form of a tree-decomposition.

Definition 6.7. Let $\mathfrak{G} = \langle V, E \rangle$ be an undirected graph.

(a) A *separator* of \mathfrak{G} is a subset S such that the complement $V \setminus S$ is not connected. A separator S is *minimal* if there exists no separator S' with $S' \subset S$.

(b) A tree-decomposition $(U_t)_{t \in T}$ of an undirected graph \mathfrak{G} is *tight* if its index tree T is directed and, for every directed edge $\langle s, t \rangle$ of T ,

- ♦ $U_s \subseteq U_t$ or $U_s \supseteq U_t$ implies $s = t$,
- ♦ the set $\bigcup_{u \in T_{st}} U_u \setminus U_t$ is connected, and
- ♦ the intersection $U_s \cap U_t$ is a minimal separator of \mathfrak{G} . \downarrow

Lemma 6.8. *Let \mathfrak{G} be a connected undirected graph of tree-width k and let v_0 be a vertex. Then \mathfrak{G} has a tight tree-decomposition of width k whose root component contains v_0 .*

Proof. Let $(U_t)_{t \in T}$ be a tree-decomposition of \mathfrak{G} of width k and fix an index $t_0 \in T$ with $v_0 \in U_{t_0}$. We consider T as a directed tree with root t_0 . We construct a tight tree-decomposition $(W_s)_{s \in S}$ of \mathfrak{G} such that every component W_s is contained in some component U_t . We proceed in three steps.

(i) We start with the connectivity condition. We perform the following operation for every vertex $t \in T$, starting at the root t_0 and then proceeding by induction on the distance between t and t_0 . Given $v \in T$, we consider each successor s of t in turn. Let $(C_i)_{i \in I}$ be the family of all connected components of $\bigcup_{u \in T_{st}} U_u \setminus U_t$. We replace the subtree of T rooted at s by several copies, one for every $i \in I$, and we change the labels U_u in the i -th copy by $U_u \cap (C_i \cup U_t)$.

(ii) In the second step we make every separator minimal. Let $(W_s)_{s \in S}$ be the decomposition obtained as the limit of the construction in step (i). For each edge $\langle s, t \rangle$ set

$$A_{st} := \bigcup_{u \in S_{st}} W_u \setminus W_t, \quad B_{st} := \bigcup_{u \in S_{ts}} W_u \setminus W_s, \quad C_{st} := W_s \cap W_t.$$

We modify $(W_s)_{s \in S}$ as follows. For every element $v \in C_{st}$ that has no neighbour in A_{st} , we remove v from all components W_u with $u \in S_{st}$. Similarly, for every element $v \in C_{st}$ that has no neighbour in B_{st} , we remove v from all components W_u with $u \in S_{ts}$. The resulting family $(W'_s)_{s \in S}$ is a tree-decomposition such that every separator $W'_s \cap W'_t$ is minimal. Furthermore, note that the above operation does not invalidate the connectivity condition established in (i).

(iii) It remains to make distinct components incomparable. Let $(W_s)_{s \in S}$ be the decomposition obtained in step (ii). We modify it in three steps.

First, let \approx be the equivalence relation on S defined by

$$s \approx t \quad \text{iff} \quad U_s = U_t.$$

We replace the index tree S by the quotient S/\approx . Since every \approx -class is connected, this quotient forms a tree. Hence, we obtain a tree-decomposition $(W_s)_{s \in S/\approx}$.

Next, for every maximal directed path s_0, \dots, s_n with $W_{s_0} \supset \dots \supset W_{s_n}$, we delete the vertices s_1, \dots, s_n and attach their successors directly to s_0 .

Finally, for every maximal directed path s_0, \dots, s_n with $W_{s_0} \subset \dots \subset W_{s_n}$, we delete the vertices s_0, \dots, s_{n-1} and attach their successors directly to s_n .

Note that the resulting tree-decomposition $(W'_s)_{s \in S'}$ still has the properties from steps (i) and (ii) since, for every edge $\langle s', t' \rangle$ of S' ,

- ♦ the set $\bigcup_{u \in S'_{t'}} W'_u \setminus W'_{t'}$ is of the form $\bigcup_{u \in S_{st}} W_u \setminus W_t$, for some $s, t \in S$, and
- ♦ the separator $W'_s \cap W'_{t'}$ is equal to $W_s \cap W_t$, for some $s, t \in S$. \square

Lemma 6.9. *Let \mathfrak{G} be a connected undirected graph, v_o a vertex of finite degree, and $n < \omega$. There exist only finitely many minimal separators C of \mathfrak{G} with $v_o \in C$ and $|C| = n$.*

Proof. We distinguish three cases. If $n = 1$, there is only one such set C .


Next, suppose that $n = 2$. For a contradiction, suppose that there are infinitely many minimal separators C of size 2 containing v_o . Then there exists an infinite set W such that, for every $w \in W$, $\{v_o, w\}$ is a minimal separator of \mathfrak{G} . We distinguish two cases.

(i) First, suppose that there exist an infinite subset $W_o \subseteq W$ and, for every $w \in W_o$, some connected component P_w of $V \setminus \{w, v_o\}$ such that $P_w \cap P_{w'} = \emptyset$, for all $w, w' \in W_o$. As v_o has finite degree, there exists some $w \in W_o$ such that no vertex of P_w is a neighbour of v_o . This implies that P_w is a connected component of $V \setminus \{w\}$, which means that $\{w\}$ is a separator of \mathfrak{G} . A contradiction to the minimality of $\{v_o, w\}$.

(ii) It remains to consider the case where there is no set W_o as in (i). Then there exist an infinite subse $W_o \subseteq W$ and a linear order \leq on W_o such that, for every $w \in W_o$,

- ♦ $V \setminus \{w, v_o\}$ has exactly 2 connected components P_w and Q_w and
- ♦ $w < w'$ implies $P_w \subset P_{w'}$ and $Q_w \supset Q_{w'}$.

Furthermore, $\langle W_o, \leq \rangle$ must either contain an infinite increasing sequence $w_o < w_1 < \dots$ or an infinite decreasing one. By symmetry, we may assume the former. To simplify notation, set $P_i := P_{w_i}$ and $Q_i := Q_{w_i}$.



As v_o has finite degree, there exists a largest index k such that some vertex in $P_{k+1} \setminus P_k$ is a neighbour of v_o . Since $V \setminus \{w_{k+1}\}$ is connected, there exists a path p from w_{k+2} to v_o that does not contain the vertex w_{k+1} . The last vertex of p before v_o has to belong to $\cap_i Q_i$. Consequently, p must contain all vertices w_i with $i \geq k + 2$. A contradiction, since p is finite.

Finally, let $n > 2$. For a contradiction, suppose that there exists an infinite set \mathcal{S} of minimal separators such that $|C| = n$ and $v_o \in C$, for all $C \in \mathcal{S}$. Replacing \mathcal{S} by an infinite subset, we may assume that there is some set D such that

$$C \cap C' = D, \quad \text{for all } C \neq C' \text{ in } \mathcal{S}.$$

Let U be a set containing exactly one element of $C \setminus D$, for every $C \in \mathcal{S}$. Let \mathfrak{G}' be the graph of obtained from \mathfrak{G} by deleting all vertices in

$$(D \setminus \{v_o\}) \cup \bigcup_{C \in \mathcal{S}} C \setminus (U \cup D).$$

It follows that, in the graph \mathfrak{G}' , we obtain infinitely many minimal separators $\{u, v_o\}$, for $u \in U$. A contradiction to the case $n = 2$ above. \square

Lemma 6.10. *Let $(U_t)_{t \in T}$ be a tight tree-decomposition of finite width of a Cayley graph $\Gamma(\mathfrak{G}, S)$. There exists a constant $L < \omega$ such that*

$$u, v \in U_t \quad \text{implies} \quad d(u, v) \leq L, \quad \text{for all } t \in T.$$

Proof. Let k be the width of $(U_t)_t$. By definition, the intersection $U_s \cap U_t$ is a minimal separator, for each edge $\langle s, t \rangle$ of T . Let \mathcal{S} be the set of all minimal separators C such that $|C| \leq k$ and $e \in C$ (where e denotes the neutral element of \mathfrak{G}). It follows by Lemma 6.9 that \mathcal{S} is finite. We claim that the desired constant is

$$L := k(k + l + 1), \quad \text{where } l := \max_{u, v \in W} d(u, v), \quad W := \bigcup \mathcal{S}.$$

Before giving the full proof, we start with a special case. Let $\langle s, t \rangle$ be an edge of T and let $C := U_s \cap U_t$ be the associated separator. Fixing a vertex $w \in C$, there is an automorphism σ of $\Gamma(\mathfrak{G}, S)$ with $\sigma(w) = e$. It follows that $\sigma[C]$ is a minimal separator with $e \in \sigma[C]$. Hence, $\sigma[C] \in \mathcal{S}$ and $\sigma[C] \subseteq W$. This implies that

$$d(u, v) \leq l, \quad \text{for all } u, v \in C.$$

To conclude the proof, we show that

$$d(u, v) \leq k(k + l + 1), \quad \text{for all } u, v \in U_t, \quad t \in T.$$

Fix $t \in T$ and $u, v \in U_t$. Let $s_0, \dots, s_{n-1} \in T$ be the neighbours of t and let p be a path of minimal length between u and v . This path induces a walk on T . Since T is a tree, we can therefore factorise p as

$$p_0, q_0, p_1, q_1, \dots, p_{m-1}, q_{m-1}, p_m,$$

where each p_i is contained in U_t and each q_i connects two vertices of $U_t \cap U_{s_{j_i}}$, for some $j_i < n$. Note that $|U_t| \leq k$ implies $|p_i| \leq k$. Furthermore, since the path p has minimal length, so do the paths q_i . As the end-points of q_i belong to $U_t \cap U_{s_{j_i}}$, it therefore follows by the special case above that $|q_i| \leq l$. Finally, since each vertex of U_t appears at most once on p , we have $m \leq |U_t| \leq k$. Thus,

$$d(u, v) = |p| \leq (m + 1)m + ml \leq (k + 1)k + kl. \quad \square$$

Lemma 6.II. *If a Cayley graph $\Gamma(\mathfrak{G}, S)$ has finite tree-width, it is κ -triangulable, for some $\kappa < \omega$.*

Proof. Suppose that $\Gamma(\mathfrak{G}, S)$ has finite tree-width. By Lemmas 6.8 and 6.10, $\Gamma(\mathfrak{G}, S)$ has a tight tree-decomposition $(U_t)_{t \in T}$ of width $k < \omega$ such that $u, v \in U_t$ implies $d(u, v) \leq L$. Set $\kappa := 2L + 1$. To see that $\Gamma(\mathfrak{G}, S)$ is κ -triangulable, we will prove that every sequence v_0, \dots, v_{n-1} of distinct vertices such that $d(v_i, v_{i+1}) \leq L$, for all $i < n$ (index arithmetic modulo n), has an L -triangulation.

Given v_0, \dots, v_{n-1} , we construct this triangulation as follows. For each index i , fix a path p_i of length at most L from v_i to v_{i+1} , and let S be the subtree of the decomposition which contains edges of these paths. (Since the union of the paths is connected, so is S . Hence, S indeed forms a subtree.) The triangulation is constructed by induction on the size of S . Consider a leaf U_t of S containing v_i, \dots, v_{i+m} but neither v_{i-1} nor v_{i+m+1} . Note that any two vertices in U_t are connected by a path of length at most L . If $m = 1$, the edge $\langle v_i, v_{i+1} \rangle$ is also contained in the predecessor of U_t . Consequently, can remove t from S and the existence of a L -triangulation follows by inductive hypothesis.

Suppose that $m > 1$. Since $d(v_i, v_{i+m}) \leq L$, we can construct a smaller cycle by replacing the sequence v_i, \dots, v_{i+m} by a path of length at most L . Then we can use the inductive hypothesis to find an L -triangulation of the resulting cycle $v_0, \dots, v_i, v_{i+m}, \dots, v_n$. We can complete it to a triangulation of the original cycle by adding the components $\{v_i, v_{i+j}, v_{i+j+1}\}$, for $j < m - 1$. (These components form a path which we can attach to the partial triangulation by adding an edge between U_t and the new component $\{v_i, v_{i+m-1}, v_{i+m}\}$.) \square

Open Question. *Is it true that every tree-interpretable graph is κ -triangulable, for some $\kappa < \omega$?*

Exercise 6.I. Find a graph in CH_2 that is not κ -triangulable for any $\kappa < \omega$. \downarrow

Ends

A second notion we will need is that of an *end* of a graph, which can be considered a generalisation of the notion of a branch from trees to more

general graphs.

Definition 6.12. Let $\mathfrak{G} = \langle V, (E_a)_{a \in A} \rangle$ be a directed, edge-labelled graph and fix a vertex v_o of \mathfrak{G} , which we call the *origin*.

(a) For a vertex u , we denote by $\text{End}(u)$ the connected component of the set

$$\{ w \in V \mid d(v_o, w) \geq d(v_o, u) \}$$

that contains the vertex u .

(b) The corresponding *frontier* is the set

$$\text{Fr}(u) := \{ w \in \text{End}(u) \mid d(v_o, w) = d(v_o, u) \}.$$

(c) An *end* of \mathfrak{G} is a subgraph induced by a set of the form $\text{End}(u)$.

(d) The *number of ends* is the limit

$$\sup_{n \rightarrow \infty} |\{ \text{End}(u) \mid d(v_o, u) = n \}|.$$

(e) We say that two ends $\text{End}(u)$ and $\text{End}(v)$ are *isomorphic* if there exists a bijection $\text{End}(u) \rightarrow \text{End}(v)$ that preserves each edge relation and that maps $\text{Fr}(u)$ bijectively to $\text{Fr}(v)$. J

Remark. One can show that a finitely generated group \mathfrak{G} has either 0, 1, 2, or infinitely many ends. Furthermore, this number is 0 if, and only if, \mathfrak{G} is finite; and it is 2 if, and only if, \mathfrak{G} has a subgroup of finite index that is isomorphic to \mathbb{Z} . Every free group with more than one generator has infinitely many ends, all of them isomorphic. J

We will need the following two observations about ends of κ -triangulable Cayley graphs.

Lemma 6.13. *If a Cayley graph $\Gamma(\mathfrak{G}, S)$ is κ -triangulable, then*

$$u' \in \text{Fr}(u) \quad \text{implies} \quad d(u, u') \leq 3\kappa.$$

Proof. Fix vertices u, u' with $d(e, u) = d(e, u')$. Let p and p' be the shortest paths from e to u and u' , respectively. Since $\text{End}(u) = \text{End}(u')$, there exists a path q from u to u' that is contained in $\text{End}(u)$. Without loss of generality, we may assume that p and p' are of the form $p = rp_o$ and $p' = rp'_o$, for some r, p_o, p'_o , where p_o and p'_o are disjoint (except for the first vertex). By assumption, the cycle consisting of p_o, q, p'_o (with p'_o in inverse direction) has a κ -triangulation. (Note that these three paths really form a cycle since p_o and p'_o are disjoint by definition, and q being contained in $\text{End}(u)$ does not intersect either of them.) Hence, we can use Lemma 6.6 to find a triangle $\{x, y, x'\}$ with $x \in p_o, y \in q$, and $x' \in p'_o$. Let p_1 and p'_1 be the suffixes of p_o and p'_o between, respectively, x and u and between x' and u' . Note that $d(u, x) \leq \kappa$ since, otherwise,

$$\begin{aligned} d(e, y) &\leq d(e, x) + d(x, y) \\ &\leq d(e, x) + \kappa \\ &= d(e, u) - d(u, x) + \kappa < d(e, u) - \kappa + \kappa = d(e, u). \end{aligned}$$

This contradicts the fact that $y \in \text{End}(u)$. Similarly, we obtain $d(u', x') \leq \kappa$. It follows that

$$d(u, u') \leq d(u, x) + d(x, x') + d(x', u') \leq 3\kappa,$$

as desired. □

Lemma 6.14. *No κ -triangulable Cayley graph $\Gamma(\mathfrak{G}, S)$ has exactly 1 end.*

Proof. If \mathfrak{G} is finite, its number of ends is 0. Hence, we may suppose that it is infinite. Then $\Gamma(\mathfrak{G}, S)$ contains an infinite path w_o, w_1, w_2, \dots such that $d(w_i, w_o) = i$, for all i . Since the automorphism group of \mathfrak{G} acts transitively on $\Gamma(\mathfrak{G}, S)$, there exists an automorphism σ_i with $\sigma_i(w_i) = e$. Set $u_i := \sigma_i(w_o)$ and $v_i := \sigma_i(w_{2i})$. It follows that $d(e, u_i) = d(e, v_i) = i$ and $d(u_i, v_i) = 2i$. Set $n := 2\kappa + 1$. If $\text{Fr}(u_i) = \text{Fr}(v_i)$, it follows by Lemma 6.13 that

$$2i = d(u_i, v_i) \leq 3\kappa < 4\kappa + 2 = 2n.$$

Hence, $\text{Fr}(u_i) \neq \text{Fr}(v_i)$, for all $i \geq n$, which implies that $\text{End}(u_i) \neq \text{End}(v_i)$. In particular, the number of ends of $\Gamma(\mathfrak{G}, S)$ is at least 2. □

Virtually Finite Groups

Finally, we need some results about virtually finite groups. We start by showing that every virtually finite group is tree-interpretable. The following lemmas are needed in the proof. First, let us note that the choice of generators does not matter when proving that a Cayley graph is tree-interpretable.

Lemma 6.15. *If $\Gamma(\mathfrak{G}, S)$ is tree-interpretable, so is $\Gamma(\mathfrak{G}, T)$, for every other set T of generators.*

Proof. We show that there exists an MSO-interpretation mapping $\Gamma(\mathfrak{G}, T)$ to $\Gamma(\mathfrak{G}, S)$. For each generator $t \in T$, we fix a word $w_t \in S^*$ with product t . It follows that $ut = v$ if, and only if, there exists a path from u to v in $\Gamma(\mathfrak{G}, S)$ whose label is w_t . This property can be expressed in MSO. \square

Let us also note that, in the definition of a virtually free group, we can choose the free subgroup to be normal.

Lemma 6.16. *Let \mathfrak{G} be a group. For every subgroup $\mathfrak{H} \subseteq \mathfrak{G}$ of finite index there exists a normal subgroup $\mathfrak{N} \subseteq \mathfrak{G}$ of finite index with $\mathfrak{N} \subseteq \mathfrak{H}$.*

Proof. Let $\alpha : \mathfrak{G} \rightarrow \text{Sym}(S)$ be the action of \mathfrak{G} on $S := \{gH \mid g \in H\}$ defined by

$$a \cdot gH := (ag)H.$$

We claim that the set $N := \ker \alpha$ induces the desired subgroup. First, note that N is normal in \mathfrak{G} since it is the kernel of a homomorphism. Furthermore, its index is bounded by $|\text{Sym}(S)| = |\text{Sym}(n)| = n!$, where $n := |S|$ is the index of \mathfrak{H} . To see that $N \subseteq H$, note that

$$\begin{aligned} a \in N = \alpha^{-1}(\text{id}) & \quad \text{iff} \quad a \cdot gH = gH, \quad \text{for all } g \in G, \\ & \quad \text{iff} \quad ag \in H, \quad \text{for all } g \in G. \end{aligned}$$

In particular, $a \in N$ implies $a = ae \in H$. \square

Corollary 6.17. *If \mathfrak{G} is a finitely generated virtually free group, it has a normal subgroup of finite index that is finitely generated and free.*

Proof. Let \mathfrak{H} be a free subgroup of \mathfrak{G} of finite index and let \mathfrak{N} be the normal subgroup from Lemma 6.16. Since subgroups of free groups are themselves free, it follows that \mathfrak{N} is a free group. It is also finitely generated since every subgroup of a finitely generated group that has finite index is itself finitely generated. \square

After these preparations we can show that virtually free groups have tree-interpretable Cayley graphs.

Lemma 6.18. *Let \mathfrak{G} be a finitely generated virtually free group. Then $\Gamma(\mathfrak{G}, S)$ is tree-interpretable.*

Proof. Let \mathfrak{G} be finitely generated and virtually free. By Corollary 6.17, \mathfrak{G} has a normal subgroup \mathfrak{N} of finite index that is finitely generated and free. Let C be a finite set of generators of \mathfrak{N} , fix an enumeration Nb_0, \dots, Nb_{m-1} (without repetitions) of G/N , and set $B := \{b_0, \dots, b_{m-1}\}$. Then every element of G can be written uniquely as a product hb with $h \in N$ and $b \in B$. By Lemma 6.15, it is sufficient to prove that $\Gamma(\mathfrak{G}, C \cup B)$ is prefix-recognisable.

As universe we take the (regular) set $C^*B \subseteq (C + B)^*$. By the above remark the product function $C^*B \rightarrow G$ is a bijection. Hence, it remains to show that multiplication by a generator are prefix-recognisable. For the proof, fix $s \in C \cup B$. For $b, b' \in B$ and $c \in C$, choose $\alpha(b, b'), \gamma(b, c) \in C^*$ and $\beta(b, b'), \delta(b, c) \in B$ such that

$$bb' = \alpha(b, b')\beta(b, b') \quad \text{and} \quad bc = \gamma(b, c)\delta(b, c).$$

(By the above remark, these choices are unique.) For $s \in B$, it follows that

$$ubs = u\alpha(b, s)\beta(b, s), \quad \text{for } u \in C^* \text{ and } b \in B.$$

For $s \in C$, we have

$$ubs = u\gamma(b, s)\delta(b, s), \quad \text{for } u \in C^* \text{ and } b \in B.$$

To see that these relations are prefix-recognisable note that the words $\alpha(b, s)$ and $\gamma(b, s)$ have bounded length. (There are only finitely many of them.)

Consequently, going from ub to $u\alpha(b, s)\beta(b, s)$ or $u\gamma(b, s)\delta(b, s)$ involves removing a bounded suffix and replacing it with another one. This operation can be expressed in MSO. \square

It remains to show that every context-free group is virtually free. To do so we need some machinery from group theory. We omit the proofs as they would lead us too far afield. We will use two different operations to decompose groups. The exact definitions are only given for completeness sake, we will not use them below. So the reader may skip them if desired.

Definition 6.19. (a) Let \mathcal{G} and \mathcal{H} be two groups with a common subgroup \mathcal{C} . We denote by $\mathcal{G} *_{\mathcal{C}} \mathcal{H}$ the *push-out* (in the category of groups) of the embeddings $\mathcal{C} \rightarrow \mathcal{G}$ and $\mathcal{C} \rightarrow \mathcal{H}$. Thus $\mathcal{G} *_{\mathcal{C}} \mathcal{H}$ is the ‘freest’ group admitting embeddings of \mathcal{G} and \mathcal{H} that agree on \mathcal{C} . We call \mathcal{C} the *base* of the push-out. A push-out $\mathcal{G} *_{\mathcal{C}} \mathcal{H}$ is *trivial* if $\mathcal{C} = \mathcal{G}$ or $\mathcal{C} = \mathcal{H}$.

(b) Let \mathcal{G} be a group with presentation $\langle S \mid R \rangle$ and let $\alpha : \mathcal{C} \rightarrow \mathcal{D}$ be an isomorphism between two subgroups $\mathcal{C}, \mathcal{D} \subseteq \mathcal{G}$. The *HNN-extension* of \mathcal{G} with respect to α is the group $\mathcal{G} *_{\alpha} \mathbb{Z}$ (this is non-standard notation) with presentation

$$\langle S, t \mid R, tct^{-1} = \alpha(c) \text{ (for all } c \in \mathcal{C}) \rangle,$$

where t is a new symbol. We call \mathcal{C} the *base* of the extension. (The letters HNN stand for Higman, Neumann, and Neumann.) \downarrow

Lemma 6.20 (Gregorac [91], Karrass, Pietrowski, Solitar [114]). *The class of finitely generated virtually free groups is closed under push-outs with finite base and HNN-extensions with finite base.*

Definition 6.21. A group \mathcal{G} *splits* if at least one of the following conditions holds.

- (I) \mathcal{G} can be written as a non-trivial push-out $\mathcal{H} *_{\mathcal{C}} \mathcal{K}$ with a finite base \mathcal{C} .
- (II) \mathcal{G} is an HNN-extension $\mathcal{H} *_{\alpha} \mathbb{Z}$ with a finite base.

In this case we say that \mathcal{G} *splits into* the groups \mathcal{H}, \mathcal{K} and into \mathcal{H} , respectively. \downarrow

The main structure theorem our proof below is based on is the following result of Stallings.

Theorem 6.22 (Stallings [197]). *A finitely generated group \mathcal{G} has more than one end if, and only if, it splits.*

The proof below is by induction on the number of times a group splits. This number is called its *accessibility length*.

Definition 6.23. Let \mathcal{G} be a group. By induction on $s < \omega$, we say that \mathcal{G} has *accessibility length* if \mathcal{G} only splits into groups of accessibility length less than s . If no such number exists, we say that \mathcal{G} has no accessibility length. We call \mathcal{G} *accessible* if it has an accessible length. \downarrow

The existence of an accessibility length follows from the following result of Dunwoody.

Theorem 6.24 (Dunwoody [74]). *Every finitely presentable group is accessible.*

Corollary 6.25. *If $\Gamma(\mathcal{G}, S)$ is κ -triangulable, then \mathcal{G} is accessible.*

Proof. By Theorem 6.24, it remains to prove that \mathcal{G} is finitely presentable. Let $\pi : S^* \rightarrow \mathcal{G}$ be the function multiplying a sequence of generators and set

$$R := \{ w \in S^* \mid \pi(w) = e \text{ and } |w| \leq 3\kappa \}.$$

We claim that $\langle S \mid R \rangle$ is a presentation of \mathcal{G} . Clearly, every $w \in R$ evaluates to e . Conversely, let $w \in S^*$ be a word with $\pi(w) = e$. Then w induces a cycle in $\Gamma(\mathcal{G}, S)$. By assumption, this cycle has a κ -triangulation. Every triangle of this triangulation is a cycle of length at most 3κ , which implies that its label belongs to R . Thus each triangle evaluates to e in the group presented by $\langle S \mid R \rangle$. This implies that w also evaluates to e in this group. \square

Lemma 6.26. *If $\Gamma(\mathcal{G}, S)$ is κ -triangulable, then \mathcal{G} is virtually free.*

Proof. Suppose that $\Gamma(\mathcal{G}, S)$ is κ -triangulable. By Corollary 6.25, \mathcal{G} has an accessibility length s . We proceed by induction on s . If $s = 0$, it follows by

Theorem 6.22 and Lemma 6.14 that the number of ends of $\Gamma(\mathfrak{G}, S)$ is 0. Hence, \mathfrak{G} is finite and, therefore, virtually free.

For the inductive step, suppose that $s > 0$. This implies that $\mathfrak{G} = \mathfrak{H} *_{\mathfrak{C}} \mathfrak{K}$ or $\mathfrak{G} = \mathfrak{H} *_{\alpha} \mathbb{Z}$ where \mathfrak{H} and \mathfrak{K} have accessibility length less than s , \mathfrak{C} is finite, and $\alpha : \mathfrak{C} \rightarrow \mathfrak{D}$. By inductive hypothesis, it follows that \mathfrak{H} and \mathfrak{K} are virtually free. Hence, we can use Lemma 6.20 to show that so is \mathfrak{G} . \square

The Proof of the Theorem

Proof of Theorem 6.4. Let $\pi : S^* \rightarrow G$ be the function mapping a sequence of generators to its product.

(4) \Leftrightarrow (6) follows by Proposition 5.5 as Cayley graphs have bounded degree.

(6) \Rightarrow (8), (9), (11) Suppose that $\Gamma(\mathfrak{G}, S)$ is tree-interpretable. Then it follows by Lemma 1.14 that it has finite crossing-width. Since $\Gamma(\mathfrak{G}, S)$ is $|S|$ -sparse, we can use Theorem IX.4.5 to show that it even has finite tree-width. Hence, the claim follows by Proposition 1.15.

(11) \Rightarrow (10) is trivial.

(10) \Rightarrow (8) Suppose that $\Gamma(\mathfrak{G}, S)$ has infinite tree-width. Then Theorem VIII.4.1 implies that $\Gamma(\mathfrak{G}, S)$ contains arbitrarily large grids as minors. Consequently, the GSO-theory of $\Gamma(\mathfrak{G}, S)$ is undecidable. Since $\Gamma(\mathfrak{G}, S)$ is $|S|$ -sparse, it follows that its MSO-theory is also undecidable.

(8) \Rightarrow (5) has been proved in Lemma 6.11.

(5) \Rightarrow (3) Suppose that $\Gamma(\mathfrak{G}, S)$ is κ -triangulable. For each $u \in G$, fix an automorphism α_u of $\Gamma(\mathfrak{G}, S)$ mapping u to e , and set

$$W := \bigcup_{u \in G} \alpha_u[\text{Fr}(u)].$$

We have shown in Lemma 6.13 that

$$u' \in \text{Fr}(u) \text{ implies } d(u, u') \leq 3\kappa.$$

This implies that every element of W has distance at most 3κ from e . Since $\Gamma(\mathfrak{G}, S)$ has bounded degree, it follows that W is a finite set. Furthermore, $\alpha_u[\text{Fr}(u)] = \alpha_{u'}[\text{Fr}(u')]$ implies that $\alpha_{u'} \circ \alpha_u^{-1}$ maps $\text{Fr}(u)$ to $\text{Fr}(u')$. Since

W has only finitely many subsets, it therefore follows that the equivalence relation

$$u \sim u' : \text{iff there is an automorphism mapping } \text{Fr}(u) \text{ to } \text{Fr}(u')$$

has finite index. Finally, note that $u \sim u'$ implies that $\text{End}(u)$ and $\text{End}(u')$ are isomorphic. (Any automorphism mapping $\text{Fr}(u)$ to $\text{Fr}(u')$ must map $\text{End}(u)$ to $\text{End}(u')$.)

(3) \Rightarrow (6) Let

$$T := \{ \text{Fr}(u) \mid u \in G \} \quad \text{and} \quad N := \max \{ |C| \mid C \in T \}.$$

(Note that $N < \omega$ since there are only finitely many frontiers, up to isomorphism.) We define an edge relation $E \subseteq T \times T$ by making two frontiers $C, D \in T$ adjacent if there are adjacent vertices $u \in C$ and $v \in D$. The resulting graph $\mathfrak{T} := \langle T, E \rangle$ forms a tree. Furthermore since, up to isomorphism, there are only finitely many ends, it follows that this tree is regular. Consequently, we can interpret the structure $\text{copy}_N(\mathfrak{T})$ in the infinite binary tree. As we can interpret $\Gamma(\mathfrak{G}, S)$ in $\text{copy}_N(\mathfrak{T})$, the claim follows.

(6) \Rightarrow (7) is trivial.

(7) \Rightarrow (9) follows by Lemma 1.14.

(9) \Rightarrow (8) follows from Theorem IX.4.5 since Cayley graphs are $|S|$ -sparse.

(5) \Rightarrow (2) \Rightarrow (6) follow by Lemmas 6.26 and 6.18, respectively.

(5) \Rightarrow (1) Suppose that $\Gamma(\mathfrak{G}, S)$ is κ -triangulable. We construct a context-free grammar for the set $\pi^{-1}[e]$ as follows. Let $C \subseteq G$ be the set of all elements of the form $\pi(w)$ with $|w| \leq \kappa$. The grammar has one non-terminal A_c , for every $c \in C$, and the rules

$$A_c \rightarrow aA_d, \quad A_c \rightarrow A_d b, \quad \text{and} \quad A_c \rightarrow A_d A_{d'},$$

for all $a, b \in G$ and $c, d, d' \in C$ with $ad = c$, $db = c$, and $dd' = c$. In addition, we use the rule

$$A_e \rightarrow \langle \rangle.$$

The start symbol is A_e . By induction on the length of a derivation it now follows that

$$A_e \Rightarrow^* w_0 A_{d_1} w_1 \cdots w_{n-2} A_{d_{n-1}} w_{n-1}$$

implies

$$c = \pi(w_0) \cdot d_1 \cdot \pi(w_1) \cdots \pi(w_{n-2}) \cdot d_{n-1} \cdot \pi(w_{n-1}).$$

For $n = 0$, it follows that $A_e \Rightarrow w$ implies $\pi(w) = e$.

Conversely, suppose that $\pi(w) = e$ where $w = a_0 \cdots a_{n-1}$. Let v_0, \dots, v_{n-1} be vertices corresponding to w (i.e., $v_i := \pi(a_0 \cdots a_{i-1})$), and fix a κ -triangulation $(U_t)_{t \in T}$ of w in $\Gamma(\mathfrak{G}, S)$. We prove that $A_e \Rightarrow^* w$ by induction on $|T|$. If $T = \{t\}$, the triangulation consists of a single triangle with edges a_0, a_1, a_2 . Consequently, we have the derivation

$$A_e \Rightarrow a_0 A_{a_1 a_2} \Rightarrow a_0 a_1 A_{a_2} \Rightarrow a_0 a_1 a_2 A_e \Rightarrow a_0 a_1 a_2.$$

For the inductive step, fix a leaf t of T . Then $U_t = \{v_i, v_{i+1}, v_{i+2}\}$, for some i (index arithmetic modulo n), and

$$A_{a_i a_{i+1}} \Rightarrow a_i A_{a_{i+1}} \Rightarrow a_i a_{i+1} A_e \Rightarrow a_i a_{i+1}.$$

Applying the inductive definition to the sequence

$$a_0, \dots, a_{i-1}, c, a_{i+2}, \dots, a_{n-1} \quad \text{with} \quad c := a_i a_{i+1}$$

we obtain a derivation

$$A_e \Rightarrow^* a_0 \cdots a_{i-1} c a_{i+2} \cdots a_{n-1}.$$

Replacing in this derivation the step introducing the terminal c by the above derivation $A_e \Rightarrow^* a_i a_{i+1}$, we obtain a derivation of w .

(i) \Rightarrow (5) Let S' be a set of generators such that the corresponding preimage $\pi^{-1}(e) \subseteq T^*$ is context-free. We claim that $\Gamma(\mathfrak{G}, S')$ is κ' -triangulable, for some κ' . Since we have already proved the equivalence (5) \Leftrightarrow (6), it

then follows that $\Gamma(\mathfrak{G}, S')$ is tree-interpretable. Hence, so is $\Gamma(\mathfrak{G}, S)$, by Lemma 6.15. This implies that $\Gamma(\mathfrak{G}, S)$ is κ -triangulable, for some κ .

Hence, fix a grammar in Chomsky normal form for the language $\pi^{-1}(e)$. We assume that it is reduced in the sense that we cannot remove any rule without altering the language and that, for every non-terminal A , the language

$$L(A) := \{w \in S'^* \mid A \Rightarrow^* w\}$$

is non-empty. For each non-terminal A , fix a word $z_A \in L(A)$ of minimal length and set

$$\kappa' := \max_A |z_A|.$$

We claim that $\Gamma(\mathfrak{G}, S')$ is κ' -triangulable.

Let w be (the label of) a cycle in $\Gamma(\mathfrak{G}, S')$. We construct a κ' -triangulation of w by induction on $n := |w|$. If $n \leq 3$ there is nothing to do. Hence, suppose otherwise. Let Z be the start symbol of our grammar. By assumption, there exists a derivation $Z \Rightarrow^* w$. Let $Z \Rightarrow AB$ be its first step and let u, v be words such that $w = uv$, $A \Rightarrow^* u$, and $B \Rightarrow^* v$.

First, suppose that $|u|, |v| > 1$. Let x be the first vertex of u and y the first one of v . We split the cycle w into two cycles by inserting an edge with label z_B from x to y . Below, we will show that $\pi(z_B) = \pi(v)$. Consequently,

$$\pi(uz_B) = \pi(w) = e \quad \text{and} \quad \pi(v(z_B)^{-1}) = e.$$

By inductive hypothesis, the two resulting cycles have κ' -triangulations. Each of them must have a component containing the edge $\langle x, y \rangle$. To obtain a κ' -triangulation of w we can combine these two triangulations by adding an edge between these two components.

Next, suppose that $|u| = 1$. (The case where $|v| = 1$ is handled analogously.) Let $B \Rightarrow C_0 C_1$ be the second step in the derivation $Z \Rightarrow^* w$ and let v_0, v_1 be words such that $v = v_0 v_1$ and $C_i \Rightarrow^* v_i$. Since $n \geq 4$, there is some $i < 2$ with $|v_i| > 1$. Let x be the first vertex of the path v_i and y be the last one. We split the cycle w into two cycles by inserting an edge with label z_{C_i} from x

to y . Then we can use the inductive hypothesis as above to construct the desired triangulation of w .

It remains to prove the above claim. More generally, we will show that $u, v \in L(A)$ implies $\pi(u) = \pi(v)$. Since the grammar is reduced, there exists a derivation $Z \Rightarrow^* xAy$. Since the grammar describes the language $\pi^{-1}(e)$, it follows that

$$\pi(xuy) = e \quad \text{and} \quad \pi(xvy) = e.$$

Consequently, $\pi(u) = \pi(x^{-1}y^{-1}) = \pi(v)$. □

7 Graphs and Paths

Let us collect a few technical results about the structure of graphs in the Causal hierarchy, in particular, about the existence of certain paths and the distances between vertices. Such results can be used to prove that a structure does not belong to a given level of the Causal hierarchy. We will also use them for an axiomatisation result in Section 8 below.

Graphs of Finite Outdegree

We start with graphs in the Causal hierarchy where each vertex has finite out-degree. Note that the universe of a structure $\mathfrak{A} \in \text{CH}_n$ in the n -th level of the hierarchy has the form $A \subseteq \Gamma^{*n}$, for some finite set Γ . We define a norm $\|z\|_k$ on such sets by taking the maximal length of a level k word contained in z .

Definition 7.1. Let Γ be a finite set. For $k \leq n$ and $z = x_0 \cdots x_{m-1} \in \Gamma^{*n}$ with $x_i \in \Gamma^{*(n-1)}$, we define, by induction on k ,

$$\|z\|_k := \begin{cases} 0 & \text{if } m = 0, \\ \|z\| & \text{if } k = n, \\ \max \{ \|x_i\|_k \mid i < m \} & \text{if } k < n \text{ and } m > 0. \end{cases}$$

Lemma 7.2. Let Γ be a finite set with at least two elements and let k_1, \dots, k_n be numbers. There are less than $|\Gamma|^{k_1 \cdots k_n}$ words $z \in \Gamma^{*n}$ such that $\|z\|_i < k_i$, for all $i \leq n$.

Proof. The claim follows by a straightforward induction on n . For $n = 1$, we have

$$\sum_{i < k_1} |\Gamma|^i = \frac{|\Gamma|^{k_1} - 1}{|\Gamma| - 1} < |\Gamma|^{k_1}$$

words $z \in \Gamma^{*n}$ with $|z| < k_1$. For $n > 1$, we can employ the inductive hypothesis to obtain the bound

$$\sum_{i < k_n} (|\Gamma|^{k_1 \cdots k_{n-1}})^i < |\Gamma|^{k_1 \cdots k_n}. \quad \square$$

It follows by Lemma 1.7 that every graph $\mathfrak{G} = \langle V, (E_a)_a \rangle \in \text{Ch}_n$ is of the form

$$\begin{aligned} V &= \{ z \in C^{*n} \mid \mathfrak{C}^{*n} \models \delta(z) \}, \\ E &= \{ (z, z') \in C^{*n} \times C^{*n} \mid \mathfrak{C}^{*n} \models \varphi(z, z') \}, \end{aligned}$$

for some finite structure \mathfrak{C} and suitable MSO-formulae $\delta(x)$ and $\varphi_a(x, y)$. Therefore, we will consider a structure of the form \mathfrak{C}^{*n} and an MSO-formula $\varphi(x, y)$ with two free first-order variables.

Definition 7.3. Let \mathfrak{A} be a structure and $\varphi(x, y) \in \text{MSO}$ a formula. The φ -out-degree of $a \in A$ in \mathfrak{A} is the number of elements $b \in A$ such that $\mathfrak{A} \models \varphi(a, b)$. J

We obtain the following bound on the φ -out-degree.

Theorem 7.4. Let $\varphi(x, y)$ be an MSO-formula and $n < \omega$. There are constants $s_1, \dots, s_n < \omega$ such that, for every finite structure \mathfrak{C} with at least two elements and every element $u \in C^{*n}$ of finite φ -out-degree in \mathfrak{C}^{*n} , we have

$$\mathfrak{C}^{*n} \models \varphi(u, v) \quad \text{implies} \quad \|v\|_i \leq L_i(u) \quad \text{for all } i \leq n,$$

where $L_i(u) := \|u\|_i + s_i |C|^{L_1(u) \cdots L_{i-1}(u)}$.

Proof. We prove the claim by induction on n . Fix an element $u \in C^{*n}$ of finite φ -out-degree. Let m be the quantifier-rank of φ . We denote by $\pi : C^{*n} \rightarrow C^{*(n-1)}$ the function mapping a sequence $w : c$ to its last element c . (For the root we set $\pi(\langle \rangle) := \perp$.) Finally, for $w \leq_{\text{pf}} u$, set

$$V(u, w) := \left\{ z \mid w <_{\text{pf}} z \leq_{\text{pf}} v' \text{ for some } v' \text{ such that } u \sqcap v' = w \text{ and } \mathfrak{C}^{*n} \models \varphi(u, v') \right\}.$$

We will prove below that the projection $\pi[V(u, w)]$ is definable in $\mathfrak{C}^{*(n-1)}$ given the element $\pi(w)$. It then follows by inductive hypothesis that

$$\|\pi(z)\|_i \leq L_i(\pi(w)) \leq L_i(u), \quad \text{for all } z \in V(u, w) \text{ and } i < n.$$

Consequently,

$$\begin{aligned} \|v\|_i &= \max \left\{ \|\pi(z)\|_i \mid z \leq_{\text{pf}} v \right\} \\ &= \max \left(\left\{ \|u \sqcap v\|_i \right\} \cup \left\{ \|\pi(z)\|_i \mid u \sqcap v <_{\text{pf}} z \leq_{\text{pf}} v \right\} \right) \\ &\leq L_i(u), \quad \text{for all } i < n. \end{aligned}$$

For $i = n$ note that, if there are vertices $u \sqcap v <_{\text{pf}} z <_{\text{pf}} z' \leq_{\text{pf}} v$ with

$$\pi(z) = \pi(z') \quad \text{and} \quad \mathfrak{C}^{*n} \upharpoonright_{\uparrow z}, v \equiv_{\text{MSO}}^m \mathfrak{C}^{*n} \upharpoonright_{\uparrow z'}, v,$$

we can replace the subtree $\mathfrak{C}^{*n} \upharpoonright_{\uparrow z'}$ by $\mathfrak{C}^{*n} \upharpoonright_{\uparrow z}$ and obtain a vertex v' with

$$\mathfrak{C}^{*n} \models \varphi(u, v') \quad \text{and} \quad |v'| > |v|.$$

Repeating this construction we obtain an infinite number of vertices v' with $\mathfrak{C}^{*n} \models \varphi(u, v')$. A contradiction.

As $V(u, w)$ is finite, it therefore follows by inductive hypothesis that

$$\begin{aligned} |v| - |u \sqcap v| &= |V(u, u \sqcap v)| \\ &\leq r \cdot |\pi[V(u, u \sqcap v)]| \\ &\leq r \cdot s' \cdot |C|^{L_1(u \sqcap v) \cdots L_{n-i}(u \sqcap v)} \leq r s' \cdot |C|^{L_1(u) \cdots L_{n-i}(u)}, \end{aligned}$$

where r the number of MSO_m -theories with one first-order parameter.

It remains to prove the above claim about the definability of $\pi[V(u, w)]$. Fix $w <_{\text{pf}} u$ and let $a \in C^{*(n-1)}$ be the element such that $wa \leq_{\text{pf}} u$. (The case where $w = u$ is handled analogously.) Note that, given $w \leq_{\text{pf}} z <_{\text{pf}} v$ and elements $b, c \in C^{*(n-1)}$ with $wb \leq_{\text{pf}} zc \leq_{\text{pf}} v$, we can compute the theory $\text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}, u, v)$ from the theories

$$\begin{aligned} \rho &:= \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{C^{*n} \setminus \uparrow wb}, u), \\ \sigma &:= \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow wb \setminus \uparrow zc}, wb, z), \\ \tau &:= \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow zc}, zc, v). \end{aligned}$$

We use the notation $\rho + \sigma + \tau := \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}, u, v)$ for the composed theory. Below we will construct MSO -formulae $\chi_\rho(x, y)$, $\psi_\sigma(x, y)$, $\vartheta_\tau(x)$ such that

$$\begin{aligned} \mathfrak{C}^{*(n-1)} \models \chi_\rho(a, b) &\quad \text{iff} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{C^{*n} \setminus \uparrow wb}, u) = \rho, \\ \mathfrak{C}^{*(n-1)} \models \psi_\sigma(b, c) &\quad \text{iff} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow wb \setminus \uparrow zc}, wb, z) = \sigma, \\ \mathfrak{C}^{*(n-1)} \models \vartheta_\tau(c) &\quad \text{iff} \quad \text{there is some } v \geq_{\text{pf}} zc \text{ such that} \\ &\quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow zc}, zc, v) = \tau. \end{aligned}$$

It then follows that

$$\begin{aligned} c \in \pi[V(u, w)] &\quad \text{iff} \quad \text{there exist elements } w <_{\text{pf}} z \leq_{\text{pf}} v \text{ such that } \pi(z) = c \text{ and} \\ &\quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{C^{*n} \setminus \uparrow wb}, u) + \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow wb \setminus \uparrow zc}, wb, z) \\ &\quad + \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n} \upharpoonright_{\uparrow zc}, zc, v) \models \varphi \\ &\quad \text{iff} \quad \text{there exist theories } \rho, \sigma, \tau \text{ such that} \\ &\quad \rho + \sigma + \tau \models \varphi \quad \text{and} \quad \tau \models \text{cl}(c) \\ &\quad \text{iff} \quad \text{there exist theories } \rho, \sigma, \tau \text{ and an element } b \text{ such that} \\ &\quad \rho + \sigma + \tau \models \varphi \quad \text{and} \quad \mathfrak{C}^{*(n-1)} \models \chi_\rho(a, b) \wedge \psi_\sigma(b, c) \wedge \vartheta_\tau(c) \\ &\quad \text{iff} \quad \mathfrak{C}^{*(n-1)} \models \bigvee_{\rho, \sigma, \tau} [\chi_\rho(a, b) \wedge \psi_\sigma(b, c) \wedge \vartheta_\tau(c)], \end{aligned}$$

where the disjunction above ranges over all theories ρ, σ, τ that satisfy $\rho + \sigma + \tau \models \varphi$.

It remains to construct the formulae χ_ρ, ψ_σ and ϑ_τ . Given a theory τ , it follows by Theorem V.3.7 that there exists a formula ϑ_τ such that

$$\langle \mathfrak{C}^{*n}, \langle c \rangle \rangle \models \exists z \bigwedge \tau \quad \text{iff} \quad \mathfrak{C}^{*(n-1)} \models \vartheta_\tau(c).$$

To define ψ_σ , note that set of all the theories of the form

$$\text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow y}, x, y'),$$

where $x <_{\text{pf}} y$ and y' is the predecessor of y , carry a natural semigroup structure. Consequently, we have

$$\text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow y}, x, y') = \sigma$$

if, and only if, there exist c_0, \dots, c_n and $\sigma_0, \dots, \sigma_n$ such that $\pi(y) = c_n, \sigma_0 + \dots + \sigma_n = \sigma$, and

$$\text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow x c_i}, x) = \sigma_i, \quad \text{for all } i \leq n.$$

(Strictly speaking we have to replace x by $x c_0 \dots c_{i-1}$ here, but since the corresponding structures are isomorphic, we have used just x to simplify notation.) It is therefore sufficient to construct formulae $\psi_\sigma^o(x, y)$ such that

$$\mathfrak{C}^{*(n-1)} \models \psi_\sigma^o(\pi(x), c) \quad \text{iff} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow x c}, x) = \sigma,$$

for all $x \in C^{*n}$ and $c \in C^{*(n-1)}$. Then the above property can be expressed in MSO by guessing sets $(P_{\sigma'})_{\sigma'}$; stating that these are the least sets satisfying

- ♦ $\mathfrak{C}^{*(n-1)} \models \psi_{\sigma_0}^o(\pi(x), c)$ implies $c \in P_{\sigma_0}$,
- ♦ $\mathfrak{C}^{*(n-1)} \models \psi_{\sigma_0}^o(c, d)$ and $c \in P_{\sigma_1}$ implies $d \in P_{\sigma_1 + \sigma_0}$;

and then checking that $\pi(y) \in P_\sigma$.

To construct ψ_σ^o , note that we can express the structure $\langle \mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow x c}, x \rangle$ as a generalised sum

$$\sum_{d \in \langle \mathfrak{C}^{*(n-1)}, b \rangle} \mathfrak{D}_d \quad \text{where} \quad \mathfrak{D}_d := \begin{cases} \langle \mathfrak{C}^{*n}, \langle d \rangle \rangle & \text{if } d \neq c, \\ \langle \{d\}, d \rangle & \text{otherwise.} \end{cases}$$

where $b := \pi(x)$, followed by an interpretation. Consequently, for each theory σ , there exist formulae ξ and $\eta_0, \dots, \eta_{n-1}$ such that

$$\begin{aligned} \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow x c}, x) &= \sigma \\ \text{iff} \quad \langle \mathfrak{C}^{*(n-1)}, b, c \rangle &\models \xi(\llbracket \eta_0 \rrbracket, \dots, \llbracket \eta_{n-1} \rrbracket). \end{aligned}$$

Furthermore, we can use the Theorem of Muchnik to construct formulae η'_i such that

$$\langle \mathfrak{C}^{*n}, \langle d \rangle \rangle \models \eta_i \quad \text{iff} \quad \langle \mathfrak{C}^{*(n-1)}, c, d \rangle \models \eta'_i.$$

Substituting these into ξ , we obtain a formula ψ_σ^o such that

$$\text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow x \setminus \uparrow x c}, x) = \sigma \quad \text{iff} \quad \mathfrak{C}^{*(n-1)} \models \psi_\sigma^o(b, c).$$

The construction of χ_ρ is similar to that of ψ_σ^o . We write $\langle \mathfrak{C}^{*n}|_{\uparrow w \setminus \uparrow w b}, u \rangle$ as a generalised sum

$$\sum_{d \in \langle \mathfrak{C}^{*(n-1)}, a \rangle} \mathfrak{D}_d$$

followed by an interpretation. The only difference is that the theory of \mathfrak{D}_a (which depends on u) is fixed. This way we obtain formulae $\xi_{\rho'}$ such that

$$\mathfrak{C}^{*(n-1)} \models \xi_{\rho'}(a, b) \quad \text{iff} \quad \text{Th}_{\text{MSO}}^m(\mathfrak{C}^{*n}|_{\uparrow w \setminus \uparrow w b}, u) = \rho'.$$

Since we can compute the theory of $\langle \mathfrak{C}^{*n}|_{C^{*n} \setminus \uparrow w b}, u \rangle$ from the theories of $\mathfrak{C}^{*n}|_{C^{*n} \setminus \uparrow w}$ (which is fixed) and $\langle \mathfrak{C}^{*n}|_{\uparrow w \setminus \uparrow w b}, u \rangle$, we can use these formulae $\xi_{\rho'}$ to build the desired formula χ_ρ . \square

We can derive several corollaries from this result that are easier to apply in practice.

Definition 7.5. (a) We denote by \exp_n the n -fold iteration of the exponentiation function

$$\exp_o(k) := k \quad \text{and} \quad \exp_{i+1}(k) := 2^{\exp_i(k)}.$$

(b) Let \mathfrak{A} be a structure and $\varphi(x, y)$ a formula. A φ -path is a sequence a_0, \dots, a_{n-1} of elements of A such that

$$\mathfrak{A} \models \varphi(a_i, a_{i+1}), \quad \text{for all } i < n-1.$$

Corollary 7.6. Let \mathfrak{C} be a finite structure and $\varphi(x, y) \in \text{MSO}$ a formula that defines a relation of finite out-degree in \mathfrak{C}^{*n} . Every φ -path $u_0, u_1, \dots \in C^{*n}$ satisfies

$$\|u_k\|_i \leq \|u_0\|_i + \exp_{i-1}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-1})),$$

for all $i \leq n$.

Proof. By Theorem 7.4, we have

$$\|u_k\|_1 \leq \|u_{k-1}\|_1 + c_1 \leq \|u_0\|_1 + c_1 k \leq \|u_0\|_1 + \exp_0(\mathcal{O}(k)).$$

It follows by induction that

$$\begin{aligned} \|u_k\|_i &\leq \|u_{k-1}\|_i + c_i |C|^{L_1(u_{k-1}) \dots L_{i-1}(u_{k-1})} \\ &\leq \|u_0\|_i + \sum_{l < k} c_i |C|^{L_1(u_l) \dots L_{i-1}(u_l)}. \end{aligned}$$

Since

$$\begin{aligned} &L_1(u_l) \dots L_{i-1}(u_l) \\ &\leq [\|u_0\|_1 + \exp_0(\mathcal{O}(l))] \\ &\quad \dots [\|u_0\|_{i-1} + \exp_{i-2}(\mathcal{O}(l + \|u_0\|_1 + \dots + \|u_0\|_{i-2}))] \\ &\leq [\|u_0\|_1 + \dots + \|u_0\|_{i-1} + \exp_{i-2}(\mathcal{O}(l + \|u_0\|_1 + \dots + \|u_0\|_{i-2}))]^{i-1} \\ &\leq [\|u_0\|_1 + \dots + \|u_0\|_{i-1} + \exp_{i-2}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2}))]^{i-1} \\ &\leq [k + \|u_0\|_1 + \dots + \|u_0\|_{i-1} \\ &\quad + \exp_{i-2}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1}))]^{i-1} \\ &\leq \exp_{i-2}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1}))^{i-1} \\ &\leq \exp_{i-2}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1})), \end{aligned}$$

we obtain

$$\begin{aligned} &\|u_k\|_i \\ &\leq \|u_0\|_i + \sum_{l < k} c_i 2^{\exp_{i-2}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1}))} \\ &\leq \|u_0\|_i + c_i k \exp_{i-1}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1})) \\ &\leq \|u_0\|_i + \exp_{i-1}(\mathcal{O}(k + \|u_0\|_1 + \dots + \|u_0\|_{i-2} + \|u_0\|_{i-1})). \quad \square \end{aligned}$$

Corollary 7.7. Let \mathfrak{C} be a finite structure and $\varphi(x, y) \in \text{MSO}$ a formula that defines a relation of finite out-degree in \mathfrak{C}^{*n} . The size of every k -neighbourhood

$$N_k(u) := \{v \in C^{*n} \mid d(u, v) \leq k\}$$

of an element $u \in C^{*n}$ is bounded by

$$|N_k(u)| \leq \exp_n(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_n)).$$

Proof. If the distance between u and v is at most k , we know by Corollary 7.6 that

$$\|v\|_i \leq \|u\|_i + \exp_{i-1}(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_{i-1})).$$

It therefore follows by Lemma 7.2 that there are less than

$$\begin{aligned} &|C|^{(\|u\|_1 + \mathcal{O}(k)) \dots (\|u\|_n + \exp_{n-1}(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_{n-1})))} \\ &\leq |C|^{(\|u\|_1 + \dots + \|u\|_n + \exp_{n-1}(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_{n-1})))^n} \\ &\leq |C|^{\exp_{n-1}(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_{n-1} + \|u\|_n))^n} \\ &\leq |C|^{\exp_{n-1}(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_{n-1} + \|u\|_n))} \\ &= \exp_n(\mathcal{O}(k + \|u\|_1 + \dots + \|u\|_n)) \end{aligned}$$

such words v . \square

Corollary 7.8. Let \mathfrak{C} be a finite structure, $\varphi(x, y) \in \text{MSO}$, and $u \in C^{*n}$. If the φ -out-degree of u in \mathfrak{C}^{*n} is finite, it is bounded by

$$\exp_n(\mathcal{O}(\|u\|_1 + \dots + \|u\|_n)).$$

We can use the above results to show that the Caucal hierarchy is strict and that it does not contain *all* structures with a decidable MSO-theory.

Examples. (a) For $n < \omega$, we consider the tree $U_n := \langle U_n, \text{succ} \rangle$ with domain

$$U_n := \{ o^i j \in \omega^* \mid i < \omega, j < \exp_n(i) \}.$$

We can interpret U_n in the tree \mathfrak{U}_n from the example on page 718. Consequently, $U_n \in \text{CH}_{n+2}$. We use the above results to show that $U_{2n} \notin \text{CH}_n$. This implies in particular that the hierarchy is strict. For a contradiction, suppose otherwise. Let $w_n \in C^{*k}$ be the word encoding the element $o^n \in U_{2k}$. According to Corollary 7.6, we have

$$\|w_n\|_i \leq \exp_{i-1}(\mathcal{O}(n)).$$

Hence, Corollary 7.8 implies that the out-degree of w_n is bounded by

$$\exp_{2k-1}(\mathcal{O}(n)).$$

A contradiction.

(b) Similarly, we can show that the tree $U_\omega := \langle U_\omega, \text{succ} \rangle$ with domain $U_\omega := \{ o^n i \in \omega^* \mid i < \exp_{2n}(n) \}$ is not contained in any level of the hierarchy. But we can show that U_ω has a decidable MSO-theory since Proposition 1.4.3 implies that, for every $m < \omega$, we can compute a constant $k < \omega$ such that

$$U_\omega \equiv_{\text{MSO}}^m U_k.$$

Paths in Tree-Interpretable Graphs

For tree-interpretable graphs, we can derive more precise bounds on the degree of vertices. The following notation comes in handy.

Definition 7.9. Given a word w , we denote by w/k its prefix of length $|w| - k$.

We start with two combinatorial lemmas. The first one allows us to obtain information about the words encoding an element.

Lemma 7.10. Let \mathfrak{A} be a tree-interpretable structure and $\varphi(\bar{x}, y)$ a CMSO-formula. There exists a constant $N < \omega$ with the following properties.

$$(I) \quad \mathfrak{A} \models \varphi(\bar{a}, b) \quad \text{implies} \quad \mathfrak{A} \models \varphi(\bar{a}, b'),$$

for some b' such that

$$\blacklozenge \quad b' \sqcap a_i = b \sqcap a_i, \quad \text{for all } i,$$

$$\blacklozenge \quad b'/N <_{\text{pf}} a_i, \quad \text{for some } i.$$

$$(II) \quad \text{If } \bar{a} \in A^n \text{ has finite } \varphi\text{-out-degree, then}$$

$$\mathfrak{A} \models \varphi(\bar{a}, b) \quad \text{implies} \quad b/N <_{\text{pf}} a_i, \quad \text{for some } i < n.$$

In particular, if the φ -out-degree of $\bar{a} \in A^n$ is finite, it is bounded by a linear function in $\max_i |a_i|$.

Proof. Let $\varphi \in \text{CMSO}_m$ and let N be the number of CMSO_m -theories. Fix $\bar{a} \in A^n$ and $b \in A$ with $\mathfrak{A} \models \varphi(\bar{a}, b)$. For (I), choose $b' \in A$ such that

$$\mathfrak{A} \models \varphi(\bar{a}, b), \quad b' \sqcap a_i = b \sqcap a_i, \quad \text{for all } i,$$

and $|b|$ is minimal. For (II), we set $b' := b$.

Let $i < n$ be an index such that $|a_i \sqcap b'|$ is maximal. It is sufficient to prove that

$$|b'| < |a_i \sqcap b'| + N.$$

For a contradiction, suppose otherwise. Then we can find words

$$b'/N \leq_{\text{pf}} x <_{\text{pf}} y \leq_{\text{pf}} b' \quad \text{such that} \quad \langle \mathfrak{A}, \bar{a}, x \rangle \equiv_{\text{CMSO}}^m \langle \mathfrak{A}, \bar{a}, y \rangle.$$

Setting $u := x^{-1}y$ and $z := y^{-1}b'$, it follows that

$$\mathfrak{A} \models \varphi(\bar{a}, b') \quad \text{implies} \quad \mathfrak{A} \models \varphi(\bar{a}, xu^jz), \quad \text{for all } j < \omega.$$

For (II), this contradicts the fact that \bar{a} has finite φ -out-degree. For (I), we also obtain a contradiction since $|xz| < |b'|$ and $|b'|$ is minimal. \square

Our second lemma states that the class of tree-interpretable structures is closed under expansions by Skolem functions.

Lemma 7.11. *Let \mathcal{A} be tree-interpretable and $\varphi(x, \bar{y}) \in \text{CMSO}$. There exists a function $f : A^n \rightarrow A$ such that the expansion $\langle \mathcal{A}, f \rangle$ is tree-interpretable and*

$$\mathcal{A} \models \exists x \varphi(x, \bar{c}) \quad \text{implies} \quad \mathcal{A} \models \varphi(f(\bar{c}), \bar{c}), \quad \text{for all } \bar{c} \in A^n.$$

Furthermore, we can choose f such that, for all sequences $(\bar{c}_i)_{i < \omega}$ satisfying

$$\mathcal{A} \models \forall x [\varphi(x, \bar{c}_i) \rightarrow \varphi(x, \bar{c}_{i+1})],$$

the sequence of images $(f(\bar{c}_i))_{i < \omega}$ is eventually constant.

Proof. Let R be the relation defined by φ , fix a tuple $\bar{a} \in A^n$, and let N be the constant from Lemma 7.10. Then there is some element $b \in A$ with

$$\langle b, \bar{a} \rangle \in R \quad \text{and} \quad |b| < |b \sqcap a_i| + N, \quad \text{for some } i < n.$$

Let $\bar{a} \in A^n$ and $b \in A$. We define a linear order \sqsubseteq_b on \bar{a} by

$$a_i \sqsubseteq_b a_j \quad : \text{iff} \quad b \sqcap a_i <_{\text{pf}} b \sqcap a_j, \\ \text{or both are equal and } a_i \leq_{\text{lex}} a_j.$$

Let $h(b, \bar{a})$ be the \sqsubseteq_b -maximal element of $\{a_0, \dots, a_{n-1}\}$.

For $\bar{a} \in A^n$, let $f(\bar{a})$ be the element b such that, in the following order,

- (1) $\mathcal{A} \models \varphi(\bar{a}, b)$,
- (2) the element $h(b, \bar{a})$ is \leq_{lex} -maximal,
- (3) the prefix $b \sqcap h(b, \bar{a})$ is \leq_{pf} -minimal,
- (4) the suffix $(b \sqcap h(b, \bar{a}))^{-1}b$ is \leq_{lex} -minimal.

As this function is MSO-definable in $\mathfrak{T}_{\text{bin}}$, the expansion $\langle \mathcal{A}, f \rangle$ is tree-interpretable.

To conclude the proof, suppose that there exists a sequence $(\bar{a}^i)_{i < \omega}$ such that

$$\langle b, \bar{a}^i \rangle \in R \quad \text{implies} \quad \langle b, \bar{a}^{i+1} \rangle \in R$$

and the sequence $b^i := f(\bar{a}^i)$ is not eventually constant. Note that $b^i = b^j$ implies $b^k = b^j$ for all $i \leq k \leq j$. Hence, by considering an appropriate subsequence of $(\bar{a}^i)_{i < \omega}$, we may assume that all the b^i are different and, since for each b there are only finitely many b' with $b'/N \leq_{\text{pf}} b/N$, that

$$b^i/N \not\leq_{\text{pf}} b^j/N, \quad \text{for all } j < i < \omega.$$

Since $b^i/N \leq_{\text{pf}} h(b^i, \bar{a}^i)$, it follows that $h(b^j, \bar{a}^i) \neq h(b^i, \bar{a}^i)$ and, hence,

$$h(b^j, \bar{a}^i) \sqcap h(b^i, \bar{a}^i) \leq_{\text{pf}} b^j \quad \text{for } j < i < \omega.$$

By induction on k , we construct infinite sets $J_k \subseteq \omega$, indices $j_k \in J_k$, words $w_k \in \Sigma^{<\omega}$, and letters $c_k \neq d_k$ satisfying

$$J_k \supseteq J_{k+1}, \quad j_k < j_{k+1}, \quad w_k <_{\text{pf}} w_{k+1}, \\ w_k c_k \leq_{\text{pf}} h(b^{j_k}, \bar{a}^{j_k}) \quad \text{for all } i \in J_k \text{ with } i > j_k, \\ w_k d_k \leq_{\text{pf}} h(b^{j_k}, \bar{a}^{j_k}) \quad \text{for all } i \in J_k \text{ with } i > j_k.$$

These conditions imply that $h(b^{j_k}, \bar{a}^i) \neq h(b^{j_{k'}}, \bar{a}^i)$, for all $k, k' \leq n$ with $k \neq k'$ and for every $i \in J_n$ such that $i > j_n$. As h can take only n different values this yields the desired contradiction. We start the construction by setting $J_0 := \omega$, $j_0 := 0$, and c_0 is the first letter of b^0 while $d_0 \neq c_0$ is arbitrary.

Given J_k, j_k, c_k , and d_k , we construct the next stage as follows. Since

$$h(b^{j_k}, \bar{a}^i) \sqcap h(b^i, \bar{a}^i) \leq_{\text{pf}} b^{j_k}$$

for all $i \in J_k, i > j_k$, there is some word w_{k+1} with

$$h(b^{j_k}, \bar{a}^i) \sqcap h(b^i, \bar{a}^i) = w_{k+1} \quad \text{for infinitely many } i \in J_k, i > j_k.$$

Thus, we can find symbols c_{k+1}, d_{k+1} such that

$$w_{k+1} c_{k+1} \leq_{\text{pf}} h(b^i, \bar{a}^i) \quad \text{and} \quad w_{k+1} d_{k+1} \leq_{\text{pf}} h(b^{j_k}, \bar{a}^i)$$

for infinitely many $i \in J_k, i > j_k$. Finally, let J_{k+1} be the set of these indices and j_{k+1} the first element of J_{k+1} . \square

We also need a bit more terminology concerning paths.

Definition 7.12. Let \mathfrak{G} be a prefix-recognisable graph where each edge relation is of the form $W(U \times V)$, for regular languages W, U, V .

(a) The *base point* of an edge $\langle a, b \rangle$ is the word $a \sqcap b$. The *spine* of a path is the sequence of the base points of its edges.

(b) A *path above* c is a path a_0, \dots, a_n such that $c \leq_{\text{pf}} a_i$ for all i .

(c) A path a_0, \dots, a_n is *bounded by* l if $|a_i| \leq l$ for all i .

(d) A sequence a_0, \dots, a_n is *k-increasing* if $|a_j| \geq |a_i| - k$ for all $i < j$.

(e) A path a_0, \dots, a_n with spine w_0, \dots, w_{n-1} is called *k-normal* if both the path and its spine are *k-increasing* and we have $a_i/k \leq_{\text{pf}} a_j$, for all $i \leq j$. \lrcorner

Our aim is to show that every vertex can be reached by a *k-normal* path. The importance of such paths stems from the fact that, by following a *k-normal* path to a vertex x , we can compute certain information about x like, e.g., its theory. This fact will be needed in Section 8 below. We start with some immediate observations.

Lemma 7.13. Let a_0, \dots, a_n be a path with spine w_0, \dots, w_{n-1} .

(a) For all $i < n-1$, either $w_i \leq_{\text{pf}} w_{i+1}$ or $w_i \geq_{\text{pf}} w_{i+1}$.

(b) If w_0, \dots, w_{n-1} is *k-increasing* then $w_i/k \leq_{\text{pf}} w_j$ for all $i < j$.

The next two lemmas can be used to find a *k-normal* path once we have shown how to obtain a path with a *k-increasing* spine.

Lemma 7.14. Let $\mathfrak{G} = \langle V, E \rangle$ be a tree-interpretable graph. There exists a constant N with the following properties.

(a) Given a path a_0, a_1, a_2 with spine w_0, w_1 , there exists a vertex a'_1 of length

$$|a'_1| < \max\{|w_0|, |w_1|\} + N$$

such that a_0, a'_1, a_2 is a path with spine w_0, w_1 .

(b) Given a path from x to y with a *k-increasing* spine w_0, \dots, w_{n-1} , there exists a path a_0, \dots, a_n from x to y with the same spine such that

$$a_i/(k+N-1) \leq_{\text{pf}} w_j \quad \text{for all } 0 < i \leq j < n.$$

Proof. Let N be the constant from Lemma 7.10.

(a) Applying Lemma 7.10 to the formula

$$\varphi(xx', y) := Exy \wedge Eyx',$$

we obtain an element a' such that a_0, a', a_2 forms a path,

$$a' \sqcap a_0 = a_1 \sqcap a_0 = w_0 \quad \text{and} \quad a' \sqcap a_2 = a_1 \sqcap a_2 = w_1,$$

and $a'/N \leq_{\text{pf}} a_i$, for $i = 0$ or $i = 2$.

(b) By (a), we can replace each a_i for $0 < i < n$ by some a'_i with

$$|a'_i| < \max\{|w_{i-1}|, |w_i|\} + N.$$

Since $w_{i-1}/k \leq_{\text{pf}} w_i$, it follows that $a'_i/(k+N-1) \leq_{\text{pf}} w_i \leq_{\text{pf}} w_j$, for all $j \geq i$. \square

In the proofs below, we frequently need to remove parts of a path and glue the remaining pieces together. The following construction is the main tool in this process.

Definition 7.15. Let a_0, \dots, a_n be a path with spine w_0, \dots, w_{n-1} and let x be a word such that $x \leq_{\text{pf}} w_i$, for all $i < n$. *Shifting* the path from x to some other word y yields the sequences a'_0, \dots, a'_n and w'_0, \dots, w'_{n-1} where a'_i and w'_i are obtained from, respectively, a_i and w_i by replacing the prefix x by y . \lrcorner

Lemma 7.16. Let \mathfrak{G} be a tree-interpretable graph and let m be the quantifier-rank of the formula defining the edge relation of \mathfrak{G} . Let a'_0, \dots, a'_n and w'_0, \dots, w'_{n-1} be the sequences obtained by shifting a path a_0, \dots, a_n with spine w_0, \dots, w_{n-1} from x to y . Then

$$\langle \mathfrak{T}_{\text{bin}}, x \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, y \rangle$$

implies that a'_0, \dots, a'_n is a path with spine w'_0, \dots, w'_{n-1} .

Proof. $\langle \mathfrak{T}_{\text{bin}}, x \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, y \rangle$ implies that

$$\langle \mathfrak{T}_{\text{bin}}, a_i a_{i+1} \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, a'_i a'_{i+1} \rangle, \quad \text{for all } i < n.$$

Consequently, $\langle a_i, a_{i+1} \rangle \in E$ implies $\langle a'_i, a'_{i+1} \rangle \in E$. \square

Now we are ready to prove the main result needed to obtain k -normal paths.

Proposition 7.17. *Let \mathfrak{G} be a tree-interpretable graph. There is a constant K_o such that, for all paths a_o, \dots, a_n with spine w_o, \dots, w_{n-1} , there exists a path of length $m \leq n$ from a_o to a_n with spine w'_o, \dots, w'_{m-1} satisfying*

$$|w'_i| < \max\{|w_o|, |w_{n-1}|\} + K_o, \quad \text{for all } i < m.$$

Proof. Let N be the constant from Lemma 7.10. Recall that N is the number of MSO_m -theories, where m is the quantifier-rank of the formula defining the edge relation. We set

$$K_o := N^4.$$

Let a'_o, \dots, a'_m be a path from a_o to a_n such that $\sum_{i < m} |w'_i|$ is minimal. We will prove the following claims.

(I) For all k, l with $k + 1 < l$, there is some $k < i < l$ such that

$$|w'_i| < \max\{|w'_k|, |w'_l|\} + N.$$

(II) $|w'_{i+1}| < |w'_i| + N^2$, for all $i < m - 1$.

(III) $|w'_i| < \max\{|w_o|, |w_{n-1}|\} + K_o$, for all $i < n$.

(I) For a contradiction, suppose that there is no such index i . Since $w_{i-1}, w_i \leq_{\text{pf}} a_i$, we have

$$w'_k, w'_l \leq_{\text{pf}} w'_i, \quad \text{for all } k < i < l.$$

Hence, either $w'_k \leq_{\text{pf}} w'_i$ or $w'_l \leq_{\text{pf}} w'_i$. W.l.o.g. assume the latter. Then there exists a word x of length N such that $w'_k x \leq_{\text{pf}} w'_i$, for all $0 < i < n - 1$. Since $|x| = N$, there are prefixes $y <_{\text{pf}} z \leq_{\text{pf}} x$ with

$$\langle \mathfrak{T}_{\text{bin}}, w'_o y \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, w'_o z \rangle.$$

Let $a''_{k+1}, \dots, a''_{l-1}$ be the path obtained by shifting the path $a'_{k+1}, \dots, a'_{l-1}$ from $w'_o z$ to $w'_o y$. Then

$$a'_o, \dots, a'_k, a''_{k+1}, \dots, a''_{l-1}, a'_l, \dots, a'_m$$

is a path from a_o to a_n with spine w''_o, \dots, w''_{m-1} such that $\sum_i |w''_i| < \sum_i |w'_i|$. A contradiction.

(II) For a contradiction, suppose that there is some index k with $|w'_{k+1}| \geq |w'_k| + N^2$. Let i_o, \dots, i_t be the sequence of all indices $i > k$ such that

$$|w'_i| \geq |w'_k| \quad \text{and} \quad |w'_i| < |w'_j|, \quad \text{for all } k < j < i.$$



By (I) and choice of i_{s+1} , we have

$$|w'_{i_{s+1}}| < |w'_{i_s}| < |w'_{i_{s+1}}| + N, \quad \text{for all } s < t - 1.$$

This implies that $t \geq N$. Consequently, there are $i < j$ in $\{i_o, \dots, i_t\}$ such that


$$\langle \mathfrak{T}_{\text{bin}}, w'_i \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, w'_j \rangle.$$

Let x be the word with $w'_i = w'_j x$. Then $w'_j x \leq_{\text{pf}} w'_i$, for all $k < l < i$. Let $(a''_i)_i$ be the path obtained from $(a'_i)_i$ by (1) shifting the subpath a_{k+1}, \dots, a_i from w'_i to w'_j and (2) removing the subpath a_{i+1}, \dots, a_j . A contradiction to the minimality of $\sum_i |w'_i|$.

(III) Fix some base point w'_k such that $|w'_k|$ is maximal, and consider the sequences $i_o < \dots < i_s < k < j_o < \dots < j_t$ of indices such that

$$|w'_{i'}| > |w'_{i_l}| \quad \text{for all } i_l < i' \leq k,$$

$$|w'_{j'}| > |w'_{j_l}| \quad \text{for all } k \leq j' < j_l.$$



For a contradiction, suppose that $|w'_k| \geq \max \{|w'_0|, |w'_{m-1}|\} + K_o$. By (II), it follows that

$$w_{i_{l+1}} \leq w_{i_{l+1}} < w_{i_l} + N^2 \quad \text{and} \quad w_{j_{l-1}} \leq w_{j_{l-1}} < w_{i_l} + N^2,$$

for all l . Consequently,

$$s \geq K_o/N^2 = N^2 \quad \text{and} \quad k \geq K_o/N^2 = N^2.$$

For $l \leq s$, let $f(l) \in \{j_0, \dots, j_t\}$ be the index such that $|w'_{f(l)}|$ is minimal with $|w'_{f(l)}| \geq |w'_{i_l}|$. Since $s \geq N^2$, we can find indices $l < l'$ such that

$$\begin{aligned} \langle \mathfrak{S}_{\text{bin}}, w_{i_l} \rangle &\equiv_{\text{MSO}}^m \langle \mathfrak{S}_{\text{bin}}, w_{i_{l'}} \rangle, \\ \langle \mathfrak{S}_{\text{bin}}, w_{f(l)} \rangle &\equiv_{\text{MSO}}^m \langle \mathfrak{S}_{\text{bin}}, w_{f(l')} \rangle. \end{aligned}$$

Let x be the word such that $w'_{i_{l'}} = w'_{i_l} x$. Then

$$w'_{i_l} x \leq_{\text{pf}} w'_{i_{l'}}, \quad \text{for } i_{l'} \leq i < f(l').$$

Hence, we obtain a new path from a'_0, \dots, a'_m by

- removing the subpaths $a'_{i_{l+1}}, \dots, a'_{i_{l'}}$ and $a'_{f(l')+1}, \dots, a'_{f(l)}$, and
- by shifting the subpath $a_{i_{l'}+1}, \dots, a_{f(l')}$ from $w'_{i_{l'}}$ to w'_{i_l} .

The resulting path contradicts the minimality of $\sum_i |w'_i|$. \square

Corollary 7.18. *Let \mathfrak{G} be a tree-interpretable graph. There exists a constant K such that all vertices a, b in the same connected component of \mathfrak{G} are connected by a path that is bounded by $\max\{|a|, |b|\} + K$.*

Proof. Let K_o be the constant from Proposition 7.17 and N the constant from Lemma 7.14 (a). Given vertices a and b in the same connected component, it follows that there exists a path a_0, \dots, a_n from a to b whose spine w_0, \dots, w_{n-1} satisfies

$$|w_i| < \max\{|w_0|, |w_{n-1}|\} + K_o, \quad \text{for all } i < n.$$

Applying Lemma 7.14 (a), we obtain a path a'_0, \dots, a'_n from a to b with the same spine that satisfies

$$|a'_i| \leq \max\{|w_{i-1}|, |w_i|\} + N < \max\{|w_0|, |w_{n-1}|\} + K_o + N. \quad \square$$

With these preparations we are able to prove the existence of k -normal paths.

Proposition 7.19. *Let \mathfrak{G} be a tree-interpretable graph. There is a constant K such that each connected component C of \mathfrak{G} contains a vertex v_o such that there are K -normal paths from v_o to all other vertices of C .*

Proof. Let K_o be the constant from Proposition 7.17 and N the constant from Lemma 7.14. We set $K := K_o + N - 1$.

Given a connected component C we choose $v_o \in C$ such that $|v_o|$ is minimal. Let a_0, \dots, a_n be a path starting at v_o and let w_0, \dots, w_{n-1} be its spine. We transform it into a path with a K_o -increasing spine as follows. Suppose there are indices $i < j$ such that $|w_j| < |w_i| - K_o$. Let $k < i$ be the greatest index such that $|w_k| < |w_j|$. By Proposition 7.17, there is a path b_0, \dots, b_m from a_k to a_j whose spine is bounded by $|w_j| + K_o$. We replace the path a_k, \dots, a_j by b_0, \dots, b_m . Iterating this operation, we obtain a path with a K_o -increasing spine. Then we can use Lemma 7.14 (b) to construct a path $a'_0, \dots, a'_{n'}$ from v_o to a_n with

$$a'_i/K \leq_{\text{pf}} w_j \leq_{\text{pf}} a'_j, a'_{j+1}, \quad \text{for all } 0 < i \leq j < n'.$$

It remains to prove that $a'_o/K = v_o/K \leq_{\text{pf}} a'_1/K$. Since $|a'_1| \geq |v_o|$, it is sufficient to show $v_o/K \leq_{\text{pf}} a'_1$. For a contradiction, suppose that $|v_o \sqcap a'_1| < |v_o| - K$. Then $|a'_1| \geq |v_o| \geq |v_o \sqcap a'_1| + K$. Recall that the constant N from

Lemma 7.14 is the number of MSO_m -theories, where m is the quantifier-rank of the formula defining the edge relation E . Therefore, there exists a vertex b with

- ♦ $v_o \sqcap a'_1 \leq_{\text{pf}} b \leq_{\text{pf}} a'_1$,
- ♦ $|b| < |v_o \sqcap a'_1| + N < |v_o|$,
- ♦ $\langle \mathfrak{T}_{\text{bin}}, x \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, y \rangle$,

where x and y are the words such that

$$b = (v_o \sqcap a'_1)x \quad \text{and} \quad a'_1 = (v_o \sqcap a'_1)y.$$

Consequently, $\langle v_o, a'_1 \rangle \in E$ implies $\langle v_o, b \rangle \in E$. In particular $b \in C$, which contradicts the fact that $|b| < |v_o|$. \square

We conclude this section with some applications bounding the length of paths and the degree of vertices.

Lemma 7.20. *Let $k > 1$. If a_o, \dots, a_n is a path with k -increasing spine w_o, \dots, w_{n-1} then its length is bounded by*

$$n < (|w_{n-1}| - |w_o| + k + 1)2^{k+1}.$$

Proof. By assumption we have $w_i/k \leq_{\text{pf}} w_j$ for all $i < j$, i.e., there are words $x_{ij} \in \{0, 1\}^*$ such that $w_j = (w_i/k)x_{ij}$. Since there are

$$\Delta := 2^{k+1} - 1$$

words of length at most k (we are working in the binary tree, so the alphabet is binary), all sets of the form $\{x_{ii}, \dots, x_{im}\}$ with $m \geq i + \Delta$ contain some word x_{ij} of a greater length. It follows that $|w_j| > |w_i|$ and, by induction, there is some $\delta \leq \Delta$ such that

$$|w_{n-1-\delta}| \geq |w_o| + \lfloor (n-1)/\Delta \rfloor.$$

Thus, $|w_{n-1}| \geq |w_o| + \lfloor (n-1)/\Delta \rfloor - k$

and $n \leq (|w_{n-1}| - |w_o| + k + 1)\Delta < (|w_{n-1}| - |w_o| + k + 1)2^{k+1}$. \square

Proposition 7.21. *Let \mathfrak{G} be a tree-interpretable graph. There exists a constant K such that every pair of vertices a, b in the same connected component of \mathfrak{G} is connected by a path of length less than*

$$(|a| + |b| + 2K + 2)2^{K+1}.$$

Proof. By Proposition 7.19, there are K -normal paths from some vertex v to a and to b . Their concatenation yields a path from a to b whose length is bounded by

$$(|a| + K + 1)2^{K+1} + (|b| + K + 1)2^{K+1}$$

according to Lemma 7.20. \square

Lemma 7.22. *Let $\mathfrak{G} = \langle V, E, P \rangle$ be a tree-interpretable graph with an additional unary predicate P . There exists a constant N such that, if there exists a k -normal path from a vertex x to some element in P , we can find such a path of length less than $(N + 1)2^{k+1}$.*

Proof. Fix some k -normal path a_o, \dots, a_n with spine w_o, \dots, w_{n-1} starting in $a_o = x$, ending in $a_n \in P$, and such that n is minimal. Let i_o, \dots, i_s be the subsequence of all indices i such that $|w_j| > |w_i|$, for all $j > i$. By (the proof of) Lemma 7.20, it follows that $i_{l+1} - i_l < 2^{k+1}$.

Let N be the number of MSO_m -theories, where m is the quantifier-rank of the formulae defining the edge relation E and the predicate P . We claim that $s < N$. Then that $n < (s + 2)2^{k+1} \leq (N + 1)2^{k+1}$, as desired.

For a contradiction, suppose otherwise. Then we can find indices $j < j'$ in $\{i_o, \dots, i_s\}$ such that

$$\langle \mathfrak{T}_{\text{bin}}, w_j \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, w_{j'} \rangle.$$

Hence, we can obtain a shorter path by deleting the subpath $a_{j+1}, \dots, a_{j'}$ and shifting the path $a_{j'+1}, \dots, a_n$ from $w_{j'}$ to w_j . This new path ends in some vertex a' satisfying

$$\langle \mathfrak{T}_{\text{bin}}, a' \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, a_n \rangle.$$

In particular, $a' \in P$. \square

8 Axiomatisations

Each tree-interpretable structure can be encoded by a finite amount of information, e.g., by an MSO-interpretation in the binary tree. It is therefore not very surprising that every such structure can be axiomatised in a suitable logic. The goal of this section is to prove the following theorem. Note that we have proved similar statements for linear orders in Theorem VI.4.10 and for trees in Theorem ??.

Theorem 8.1. *Every tree-interpretable structure is the unique countable model of a suitable GSO-formula.*

Before giving the proof, let us mention a few consequences. We start with a weak form of the converse statement.

Proposition 8.2. *Let \mathfrak{A} be a structure of finite crossing-width. If \mathfrak{A} is the unique countable model of a CMSO-formula, it is tree-interpretable.*

Proof. Let φ be the CMSO-formula whose only countable model is \mathfrak{A} . Since \mathfrak{A} has finite crossing-width, there exists an MSO-interpretation τ and a tree \mathfrak{T}_0 such that $\mathfrak{A} \cong \tau(\mathfrak{T}_0)$, for some tree \mathfrak{T} . The class \mathcal{T} of trees \mathfrak{T} with $\tau(\mathfrak{T}) \cong \mathfrak{A}$ is finitely CMSO-axiomatisable since

$$\tau(\mathfrak{T}) \cong \mathfrak{A} \quad \text{iff} \quad \tau(\mathfrak{T}) \models \varphi \quad \text{iff} \quad \mathfrak{T} \models \varphi^\tau.$$

By Theorem V.2.17, it follows that \mathcal{T} contains a regular tree \mathfrak{C} . Since every regular tree can be interpreted in $\mathfrak{T}_{\text{bin}}$, we can find an MSO-interpretation σ mapping $\mathfrak{T}_{\text{bin}}$ to \mathfrak{C} . It follows that

$$\mathfrak{A} \cong \tau(\mathfrak{C}) = (\tau \circ \sigma)(\mathfrak{T}_{\text{bin}})$$

is tree-interpretable. \square

For HR-equational structures, where GSO collapses to MSO, we obtain a precise characterisation.

Corollary 8.3 (Courcelle). *A structure \mathfrak{A} is HR-equational if, and only if, it has finite tree-width and it is the unique countable model of a suitable MSO-formula.*

It is currently an open question whether the corresponding statement for tree-interpretable structures is also true.

Conjecture. *A structure \mathfrak{A} is tree-interpretable if, and only if, it has finite crossing-width and it is the unique countable model of a suitable MSO-formula.*

Let us also mention that the strengthening of Proposition 8.2 to GSO does not hold.

Lemma 8.4. *There exists a countable structure \mathfrak{A} with finite crossing-width that is finitely GSO-axiomatisable but not tree-interpretable.*

Proof. Let $\mathfrak{A} := \langle \mathbb{N}, \leq, P \rangle$ where $P \subseteq \mathbb{N}$ is any arithmetical but non-recursive set. Then \mathfrak{A} has finite crossing width since $\langle \mathbb{N}, \leq \rangle$ is tree-interpretable and the crossing-width is not affected by expansions by unary predicates. On the other hand, \mathfrak{A} is not tree-interpretable since its MSO-theory is undecidable.

Finally note that, because of \leq , the expressive power of GSO equals full second-order logic. In particular, we can define addition and multiplication. Consequently, there exists a GSO-formula defining P based on the FO-definition of P in $\langle \mathbb{N}, \cdot, + \rangle$. \square

We can also prove an analogue to Theorem VI.4.7.

Theorem 8.5. *Let $m < \omega$. For every countable structure \mathfrak{A} with finite crossing-width, there exists a tree-interpretable structure \mathfrak{B} with $\mathfrak{B} \equiv_{\text{MSO}}^m \mathfrak{A}$.*

Proof. As \mathfrak{A} has finite crossing-width, there exist some MSO-interpretation τ and some tree \mathfrak{T} such that $\mathfrak{A} = \tau(\mathfrak{T})$. Let k be the quantifier-rank of τ . By Corollary V.2.18, we can find a regular tree $\mathfrak{C} \equiv_{\text{MSO}}^{m+k} \mathfrak{T}$. It follows that $\mathfrak{A} = \tau(\mathfrak{T}) \equiv_{\text{MSO}}^m \tau(\mathfrak{C})$. \square

Corollary 8.6. *Let $m < \omega$. For every structure \mathfrak{A} with finite tree-width, there exists an HR-equational structure \mathfrak{B} with*

$$\mathfrak{A} \equiv_{\text{GSO}}^m \mathfrak{B}.$$

Finally, let us take a look at the isomorphism problem.

Theorem 8.7 (Courcelle [56]). *It is decidable whether two given HR-equational structures are isomorphic.*

Proof. Although not stated explicitly, the construction of the formula in Theorem 8.1 is effective. Thus, in order to determine whether $\mathfrak{A} \cong \mathfrak{B}$, one can construct the GSO-formula $\varphi_{\mathfrak{A}}$ for \mathfrak{A} and check whether $\mathfrak{B} \models \varphi_{\mathfrak{A}}$. \square

Again, the case of tree-interpretable structures is still open.

Open Question. *Is the isomorphism problem for tree-interpretable structures decidable?*

It remains to prove Theorem 8.1. The idea is as follows. Given a tree-interpretable structure \mathfrak{A} , we guess guarded relations that encode a regular tree \mathfrak{T} which \mathfrak{A} can be interpreted in. Since regular trees are axiomatisable the result then follows. Hence, the proof is very similar to the statement of Theorem X.2.6. The difference is that we require the resulting tree \mathfrak{T} to be regular and that we do not need the mapping $\mathfrak{A} \mapsto \mathfrak{T}$ to be an interpretation. As a consequence, we cannot use Theorem X.2.6 directly. But the proof is fortunately a lot simpler than the one of that theorem.

Sorefts

We start slowly by showing that forests are finitely axiomatisable. We regard forests as partial orders such that the elements below any given one form a finite linear order.

As a technical tool in the proof we use colourings of single elements and of pairs of elements the following form.

Definition 8.8. (a) Let u be a word. Recall that u/k denotes the prefix of u of length $|u| - k$. We denote by $\text{suf}_k(u)$ the suffix of u of length k .

(b) For each $m < \omega$, we define the equivalence relation

$$x \equiv_m y \quad : \text{iff} \quad \langle \mathfrak{T}_{\text{bin}}, x \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, y \rangle.$$

We denote the \equiv_m -class of a word x by $[x]_m$.

(b) Let $m, k < \omega$. The m, k -colouring χ_k^m maps each word $x \in \{0, 1\}^*$ to the pair

$$\chi_k^m(x) := \langle [x/k]_m, \text{suf}_k(x) \rangle,$$

and each pair of words $\langle x, y \rangle \in \{0, 1\}^* \times \{0, 1\}^*$ to the pair

$$\chi_k^m(x, y) := \langle \chi_k^m(w^{-1}x), \chi_k^m(w^{-1}y) \rangle, \quad \text{where } w := x \sqcap y.$$

(c) The m, k -expansion $\langle \mathfrak{A}, \chi_k^m \rangle$ of a structure \mathfrak{A} expands \mathfrak{A} by unary and binary relations for each colour class where the binary colour classes consists only of pairs $\langle x, y \rangle$ which are guarded. \lrcorner

The restriction to guarded pairs is essential since GSO allows only quantification over relations of this form. But below we frequently will need to obtain the value $\chi(x, y)$ for pairs $\langle x, y \rangle$ which are not guarded. These values must be computed explicitly from available data. This is where the k -normal paths of Section 7 come into play.

Remark. Note that the m, k -expansion of a structure \mathfrak{A} is MSO-interpretable in \mathfrak{A} . \lrcorner

Definition 8.9. We say that a set P of vertices *codes* a path between x and y if P is connected and every element of P except for x and y has exactly two neighbours in P , whereas x and y have only one such neighbour. \lrcorner

Clearly, not every path can be coded in this way. Fortunately, for our purposes it is sufficient that, if there exists a k -normal path between two vertices, then we can obtain a codable k -normal path between them by removing some vertices.

Lemma 8.10. *Let $m, k < \omega$, and let c be a colour. There exists an MSO-formula $\varphi_c(P, x, y)$ such that, for all graphs \mathfrak{G} ,*

$$\langle \mathfrak{G}, \chi_k^m \rangle \models \varphi_c(P, x, y) \quad \text{iff} \quad \begin{aligned} &P \text{ codes a } k\text{-normal path from } x \text{ to } y \\ &\text{and } \chi_k^m((x \sqcap y)^{-1}y) = c. \end{aligned}$$

Proof. We label each elements $z \in P$ by $\chi_k^m((x/k')^{-1}z)$, for some $k \leq k' \leq 2k$. Since $x/k \leq_{\text{pf}} y$, we can compute the colour $\chi_k^m((x \sqcap y)^{-1}y)$ from $\chi_k^m(x)$ and the label of y . To decide whether a given labelling is correct note that, if $\langle z, z' \rangle$ is an edge of the path and z is labelled $\langle [u]_m, w \rangle$, the label of z' consists of the suffix w' of z' of length $\min\{2k, |w| + |z'| - |z|\}$ and the \equiv_m -class of $(x/k)^{-1} \cdot z' \cdot (w')^{-1}$, both of which can be calculated from the colour of z . (Note that $|w| + |z'| - |z| \geq 2k - k = k$ since the path is k -normal.) \square

As a first application of the colourings, let us note that the colour of a vertex determines the attached subtree up to isomorphism.

Lemma 8.11. *Let $\mathfrak{T} := \langle T, \leq \rangle$ be a tree-interpretable forest. There exists constants $m, k < \omega$ such that*

$$\chi_k^m(x) = \chi_k^m(y) \quad \text{implies} \quad \uparrow x \cong \uparrow y, \quad \text{for } x, y \in T.$$

Proof. Let m be the quantifier-rank of the formula defining the ordering \leq and let N be the constant from Lemma 7.10. For each $b \in T$, there are only finitely many $a \leq b$. By choice of N , it follows that $a \leq b$ implies $a/N \leq_{\text{pf}} b$. Hence, we have $\uparrow a \subseteq (a/N)\{0, 1\}^*$, and the function

$$f : (x/N)\{0, 1\}^* \rightarrow (y/N)\{0, 1\}^* : (x/N)z \mapsto (y/N)z$$

is the desired isomorphism. \square

This property can be used to prove the special case of Theorem 8.1 for forests.

Proposition 8.12. *Let $\mathfrak{T} := \langle T, \leq \rangle$ be a tree-interpretable forest. There exist constants $m_0, k_0 < \omega$ such that, for all $m \geq m_0$ and $k \geq k_0$, the expansion $\langle \mathfrak{T}, \chi_k^m \rangle$ is the unique countable model of some MSO-formula.*

Proof. Let m_0, k_0 be the constants from Lemma 8.11 and let $T_0 \subseteq T$ be the set of minimal elements of \mathfrak{T} . We denote by $\text{Suc}(x)$ the set of immediate successors of $x \in T$. For $X \subseteq T$, let $\mu(X)$ be the function which maps each colour c to the number of elements $x \in X$ coloured c .

We claim that a countable structure $\mathfrak{X} := \langle X, \leq, \chi' \rangle$ is isomorphic to $\langle \mathfrak{T}, \chi \rangle$ if, and only if,

- (I) \leq is a partial order such that, for all $x \in X$, the set $\downarrow x$ forms a finite linear order,
- (II) $\mu(X_0) = \mu(T_0)$, where $X_0 \subseteq X$ is the set of minimal elements, and
- (III) $\chi'(x) = \chi_k^m(u)$ implies $\mu(\text{Suc}(x)) = \mu(\text{Suc}(u))$, for all $x \in X$ and $u \in T$.

Clearly, all these conditions can be expressed in MSO.

To prove the non-trivial direction we construct an isomorphism $h : X \rightarrow T$ given some order \mathfrak{X} that satisfies the above conditions. Note that (I) implies that \mathfrak{X} is a forest. Let $\text{ht}(x) := |\downarrow x| - 1$. We construct h as the limit of partial isomorphisms

$$h_i : \{x \in X \mid \text{ht}(x) \leq i\} \rightarrow \{u \in T \mid \text{ht}(u) \leq i\}, \quad i < \omega,$$

as follows.

($i = 0$) Since $\mu(X_0) = \mu(T_0)$, there is a bijection $h_0 : X_0 \rightarrow T_0$ that preserves the colouring.

($i > 0$) For every $x \in X$ with $\text{ht}(x) = i - 1$, we can use (III) to find a colour preserving bijection $g_x : \text{Suc}(x) \rightarrow \text{Suc}(h_{i-1}x)$. Then h_i is the extension of h_{i-1} by all these functions g_x .

Using Lemma 8.11, it is now straightforward to show that the limit h is well-defined and indeed an isomorphism. \square

Partial Orders

The next step consists in extending the result to tree-interpretable certain partial orders $\mathfrak{A} = \langle A, \leq \rangle$. To do so we define a forest in \mathfrak{A} and reduce the problem to the preceding proposition. When speaking of paths we always consider undirected paths in this section, that is, we ignore the direction of the edges.

Definition 8.13. Let $\mathfrak{A} = \langle A, \leq \rangle$ be a partial order and $K < \omega$. We write $x \sqsubseteq y$ if there exists an undirected \leq -path z_0, \dots, z_n from x to y with

$$x/K \leq_{\text{pf}} z_i/K, \quad \text{for all } i \leq n.$$

We set $x \equiv y$ if $x \sqsubseteq y$ and $y \sqsubseteq x$.

Lemma 8.14. Let $\mathfrak{A} = \langle A, \leq \rangle$ be a partial order and $K < \omega$.

- (a) $\langle A, \sqsubseteq \rangle$ is a forest.
- (b) $|[a]_{\equiv}| \leq 2^K$, for all $a \in A$.
- (c) For every $m < \omega$, the set $A_o \subseteq A$ of lexicographically minimal elements of each \equiv -class is MSO-definable in $\langle A, \sqsubseteq, \chi_K^m \rangle$.

Proof. (a) It is sufficient to show that $\downarrow[x]_{\equiv}$ is a linear order, for all $[x]_{\equiv} \in A/\equiv$. Suppose that $[y]_{\equiv}, [z]_{\equiv} \sqsubseteq [x]_{\equiv}$. Then $y/K, z/K \leq_{\text{pf}} x/K$ and, by symmetry, we may assume that $y/K \leq_{\text{pf}} z/K$. We claim that $[y]_{\equiv} \sqsubseteq [z]_{\equiv}$.

By definition, there are undirected \leq -paths y_0, \dots, y_l from y to x and z_0, \dots, z_m from z to x such that

$$y/K \leq_{\text{pf}} y_i/K \quad \text{and} \quad z/K \leq_{\text{pf}} z_i/K, \quad \text{for all } i.$$

Furthermore, $y/K \leq_{\text{pf}} z/K$ implies $y/K \leq_{\text{pf}} z_i/K$. The concatenated path $y_0, \dots, y_l, z_{m-1}, \dots, z_0$ leading from y to z therefore witnesses that $y \sqsubseteq z$.

(b) By definition, $x \equiv y$ implies $x/K = y/K$. Hence, $[x]_{\equiv} \subseteq w\{o, i\}^K$, for $w := x/K$.

(c) Since $x \equiv y$ implies $x/K = y/K$, one can determine whether $x \leq_{\text{lex}} y$ by looking at the suffices $\text{suf}_K(x)$ and $\text{suf}_K(y)$. This information is contained in the colouring χ_K^m . \square

Using the result of the previous section we first prove that $\langle A, \sqsubseteq, \chi \rangle$ is axiomatisable by defining a suitable copy of $\langle A, \sqsubseteq \rangle/\equiv$ inside of it. Then we transfer the result to $\langle A, \leq, \chi \rangle$.

Proposition 8.15. Let $\mathfrak{A} = \langle A, \leq \rangle$ be a partial order and $K < \omega$. There exist constants $m_o, k_o < \omega$ such that, for all $m \geq m_o$ and $k \geq k_o$, the order $\langle A, \sqsubseteq, \chi_k^m \rangle$ is the unique countable model of some MSO-formula.

Proof. Set $m := m'$ and $k := \max\{K, k'\}$, where m', k' are the constants from Proposition 8.12 (applied to the structure \mathfrak{B}/\equiv). By Lemma 8.14 (c), there exists an MSO-formula $\delta(x)$ defining the set A_o in $\langle \mathfrak{B}, \chi_k^m \rangle$, where $\mathfrak{B} := \langle A, \sqsubseteq \rangle$.

A structure $\langle \mathfrak{X}, \chi' \rangle := \langle X, \sqsubseteq', \chi' \rangle$ is isomorphic to $\langle \mathfrak{B}, \chi_k^m \rangle$ if, and only if, there is an isomorphism $f : \langle \mathfrak{X}/\equiv', \chi' \rangle \rightarrow \langle \mathfrak{B}/\equiv, \chi_k^m \rangle$ such that $[x]_{\equiv'} \cong f[x]_{\equiv}$, for all $x \in X$, where $\equiv' := \sqsubseteq' \cap \sqsupset'$. This condition is equivalent to the following ones.

- (I) $\delta^{\mathfrak{X}}$ contains exactly one element of each \equiv' -class of X .
- (II) $\langle \delta^{\mathfrak{X}}, \sqsubseteq', \chi' \rangle \cong \langle \delta^{\mathfrak{B}}, \sqsubseteq, \chi_k^m \rangle$.
- (III) $\chi'(x) = \chi_k^m(a)$ implies $[x]_{\equiv'} \cong [a]_{\equiv}$, for $x \in X$ and $a \in A$.

These statements can be expressed in MSO. For (II), note that the forest $\langle \delta^{\mathfrak{B}}, \sqsubseteq, \chi_k^m \rangle \cong \langle \mathfrak{B}/\equiv, \chi_k^m \rangle$ can be axiomatised using Proposition 8.12. For (III), note that Lemma 8.14 (b) provides a bound on the size of each \equiv -class. \square

In order to transfer this result from $\langle A, \sqsubseteq \rangle$ to \mathfrak{A} , we have to show that each of the structures is definable in the other one.

Lemma 8.16. Let $\mathfrak{A} = \langle A, \leq \rangle$ be a partial order and $K < \omega$. There exist constants $m_o, k_o < \omega$ with the following properties.

- (a) If $x \sqsubseteq y$, there is a k_o -normal path z_0, \dots, z_n from x to y with $x/K \leq_{\text{pf}} z_i/K$ for all i .
- (b) If $x \equiv y$, there exists an undirected \leq -path z_0, \dots, z_n from x to y with

$$x/K \leq_{\text{pf}} z_i/K \quad \text{and} \quad |z_i| \leq |x| + k_o, \quad \text{for all } i \leq n.$$

- (c) The relation \sqsubseteq is MSO-definable in $\langle \mathfrak{A}, \chi_{k_o}^{m_o} \rangle$.
- (d) $\langle A, \sqsubseteq, \chi_k^m \rangle$ is MSO-definable in $\langle \mathfrak{A}, \chi_k^m \rangle$, for all $m \geq m_o$ and $k \geq k_o$.

Proof. Let m_o be arbitrary and set $k_o := \max\{K, K', K''\}$ where K' is the constant from Proposition 7.19 and K'' the one from Corollary 7.18.

(a) is implied by Proposition 7.19, while (b) follows from Corollary 7.18 since x and y are connected by a path above $x/K = y/K$.

(c) By (a), we have $x \sqsubseteq y$ if, and only if, there is a k_o -normal undirected path z_0, \dots, z_n from x to y with $x/K \leq_{\text{pf}} z_i/K$, for all i . Thus, $x \sqsubseteq y$ if, and only if, there is a k_o -normal path p such that each initial segment p' of p leads to some vertex z with $|z| \geq |x|$. It follows from Lemma 8.10 that the condition $|z| > |x|$ can be expressed by an MSO-formula.

(d) By (c), it remains to define the colours $\chi_k^m(x, y)$, for $x \sqsubseteq y$. This can be done using Lemma 8.10 since we have shown in (a) that there exists a k -normal path from x to y . \square

Lemma 8.17. *Let $K < \omega$ and let \mathfrak{A} be a tree-interpretable partial order such that $x \leq y$ implies $x/K \leq_{\text{pf}} y/K$. There exist constants $m_o, k_o < \omega$ such that $\langle \mathfrak{A}, \chi_k^m \rangle$ is MSO-definable in $\langle A, \sqsubseteq, \chi_k^m \rangle$, for all $m \geq m_o$ and $k \geq k_o$.*

Proof. Set $k_o := K$ and let m_o be the quantifier-rank of the formula defining \leq . Fix $m \geq m_o$ and $k \geq k_o$. Since \leq is MSO_m -definable and since $x \leq y$ implies $x/K \leq_{\text{pf}} y/K$, it follows that there exists a set Θ of triples $\langle \rho, \sigma, \tau \rangle$ of MSO_m -theories such that

$$\begin{aligned} x \leq y \quad \text{iff} \quad & x/K \leq_{\text{pf}} y/K \text{ and there are } \langle \rho, \sigma, \tau \rangle \in \Theta \text{ such that} \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, w) = \rho, \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, w^{-1}x) = \sigma, \\ & \text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, w^{-1}y) = \tau, \\ & \text{where } w := x \sqcap y. \end{aligned}$$

As $x \leq y$ implies $x \sqsubseteq y$, all of the above conditions can be expressed in MSO using $\chi_k^m(x)$, $\chi_k^m(y)$, and $\chi_k^m(x, y)$. The colouring of $\langle \mathfrak{A}, \chi_k^m \rangle$ is definable for the same reason. \square

Proposition 8.18. *Let $K < \omega$ and let \mathfrak{A} be a tree-interpretable partial order such that $x \leq y$ implies $x/K \leq_{\text{pf}} y/K$. There exist constants $m_o, k_o < \omega$ such that, for all $m \geq m_o$ and $k \geq k_o$, the order $\langle \mathfrak{A}, \chi_k^m \rangle$ is the unique countable model of some MSO-formula.*

Proof. Let m_1, k_1 be the constants from Proposition 8.15. m_2, k_2 be the constants from Lemma 8.17, and m_3, k_3 those from Lemma 8.16. We set

$$m_o := \max \{m_1, m_2, m_3\} \quad \text{and} \quad k_o := \max \{K, k_1, k_2, k_3\}.$$

Let τ be the MSO-definition of $\langle A, \sqsubseteq, \chi_k^m \rangle$ in $\langle \mathfrak{A}, \chi_k^m \rangle$. It follows by Lemmas 8.17 and 8.16, a structure $\langle \mathfrak{X}, \chi' \rangle$ is isomorphic to $\langle \mathfrak{A}, \chi_k^m \rangle$ if, and only if, $\tau(\mathfrak{X}, \chi') \cong \tau(\mathfrak{A}, \chi_k^m)$. The claim therefore follows by applying Proposition 8.15 to $\tau(\mathfrak{A}, \chi_k^m)$. \square

The General Case

Finally, we consider an arbitrary tree-interpretable structure \mathfrak{A} . For the reduction to the previous case we define, as above, a partial order \leq and show that the structures $\langle A, \leq \rangle$ and \mathfrak{A} are definable within each other.

Definition 8.19. Let \mathfrak{A} be a structure and $K < \omega$. We denote by \leq_K the reflexive and transitive closure of the relation

$$x \vdash_K y \quad \text{: iff} \quad x/K \leq_{\text{pf}} y/K \text{ and the pair } \langle x, y \rangle \text{ is guarded.}$$

Lemma 8.20. *Let \mathfrak{A} be a structure and $K < \omega$. The structure $\langle A, \leq_K, \chi_K^m \rangle$ is MSO-definable in $\langle \mathfrak{A}, \chi_K^m \rangle$, for all $m < \omega$.*

Proof. The relation \vdash is MSO-definable since, for a guarded tuple, one can tell whether $x/K \leq_{\text{pf}} y/K$ holds by looking at $\chi_K^m(x, y)$. Thus, the reflexive and transitive closure \leq_K is also MSO-definable.

To show that the colouring is definable we prove that, for each colour c of χ_K^m , there is a formula $\varphi_c(x, y)$ such that

$$\langle \mathfrak{A}, \chi_K^m \rangle \models \varphi_c(x, y) \quad \text{iff} \quad x \leq_K y \text{ and } \chi_K^m(x, y) = c.$$

If $x \vdash y$, there is a relation R of \mathfrak{A} and a tuple $\bar{a} \in R$ with $x, y \in \bar{a}$. Hence, the colour $\chi_K^m(x, y)$ is available in $\langle \mathfrak{A}, \chi_K^m \rangle$ and we can construct a formula $\varphi_c^+(x, y)$ expressing that $x \vdash y$ and $\chi_K^m(x, y) = c$. We have $x \leq_K y$ if, and only if, there is a path $x = z_0 \vdash \dots \vdash z_n = y$. Note that $z_i \vdash z_{i+1}$ implies $z_i/K \leq_{\text{pf}} z_{i+1}/K$. Therefore, we can compute $\chi_K^m(x, z_{i+1})$ from $\chi_K^m(x, z_i)$ and $\chi_K^m(z_i, z_{i+1})$. \square

The proof of the converse is more involved. It follows by Proposition V.1.11 that we can compute the MSO_m -theory of an n -tuple \bar{a} in the binary tree from the MSO_m -theories of the words $(a_i \sqcap a_l)^{-1}(a_j \sqcap a_l)$, for all $i, j, l < n$. To do so, we need access to the words $a_i \sqcap a_j$. The problem is that these words do not have to be elements of the given tree-interpretable structure \mathfrak{A} . But using a pumping argument we can find elements b_{ij} of \mathfrak{A} which are ‘sufficiently close’ to $a_i \sqcap a_j$ that we can use them instead. The formal definition is as follows.

Definition 8.21. Let $\bar{a} \in A^n$ and $K < \omega$.

(a) A K -code of \bar{a} is a family $(b_{ij})_{i,j < n}$ of elements of A such that, for all $i, j, l < n$,

- ♦ $b_{ii} = a_i$,
- ♦ $b_{ij}/K <_{\text{pf}} a_i \sqcap a_j \leq_{\text{pf}} b_{ij}$,
- ♦ $a_i \sqcap a_l <_{\text{pf}} a_j \sqcap a_l$ implies $b_{il} \vdash_K b_{jl}$ or $b_{il} = b_{jl}$.

(b) The *branching type* of \bar{a} is the set

$$\{ \langle i, j, l \rangle \in [n]^3 \mid a_i \sqcap a_l \leq_{\text{pf}} a_j \sqcap a_l \}.$$

In particular, given $b_{il} \vdash_k b_{jl}$, we can compute the theory of the word $(a_i \sqcap a_l)^{-1}(a_j \sqcap a_l)$ from the colour $\chi(b_{il}, b_{jl})$.

Lemma 8.22. Let \mathfrak{A} be a structure, $n, m, k < \omega$ constants and $H \subseteq [n]^3$. There exists an MSO -formula $\beta(\bar{x}, \bar{y})$ such that, for every $\bar{a} \in A^n$,

$$\langle A, \leq_k, \chi_k^m \rangle \models \beta(\bar{a}, \bar{b})$$

if, and only if, \bar{b} is a k -code of \bar{a} and H is its branching type.

Proof. It is straightforward to express the facts that $b_{ii} = a_i$ and that $\langle i, j, l \rangle \in H$ implies $b_{il} \vdash_k b_{jl}$.

We also have to say that $b_{ij}/s_{ij} = a_i \sqcap a_j$, for some $s_{ij} < k$. Since $b_{ij} \vdash_k b_{ii} = a_i$ and $b_{ij} \vdash_k b_{jj} = a_j$, this can be determined by looking at $\chi_k^m(b_{ij}, a_i)$ and $\chi_k^m(b_{ij}, a_j)$.

Finally, we have to check that H is the branching type of \bar{a} . If $b_{il} \not\vdash_k b_{jl}$, we must have $\langle i, j, l \rangle \notin H$. For $b_{il} \vdash_k b_{jl}$, we can check whether or not

$$a_i \sqcap a_l = b_{il}/s_{il} \leq_{\text{pf}} b_{jl}/s_{jl} = a_j \sqcap a_l$$

by looking at the colour $\chi_k^m(b_{il}, b_{jl})$. □

Lemma 8.23. Let \mathfrak{A} be tree-interpretable. There exists a constant $K < \omega$ such that every guarded n -tuple \bar{a} has an nK -code $(b_{ij})_{i,j}$.

Proof. Let K be the number of MSO_m -theories where m is the minimal number such that every relation of \mathfrak{A} is MSO_m -definable.

We call two n -tuples \bar{a}, \bar{b} in $\{0, 1\}^*$ *similar* if they have the same branching type and

$$\begin{aligned} \langle \mathfrak{S}_{\text{bin}}, a_i \sqcap a_j \rangle &\equiv_{\text{MSO}}^m \langle \mathfrak{S}_{\text{bin}}, b_i \sqcap b_j \rangle \\ \langle \mathfrak{S}_{\text{bin}}, (a_i \sqcap a_l)^{-1}(a_j \sqcap a_l) \rangle &\equiv_{\text{MSO}}^m \langle \mathfrak{S}_{\text{bin}}, (b_i \sqcap b_l)^{-1}(b_j \sqcap b_l) \rangle, \end{aligned}$$

for all i, j, l . An n -tuple \bar{b} is a *reduct* of \bar{a} if it is similar to \bar{a} and

$$|b_i| \leq nK, \quad \text{for all } i < n.$$

More generally, we say that \bar{b} is a reduct of \bar{a} over some word $w \in \{0, 1\}^*$ if there are words \bar{c}, \bar{d} such that

$$a_i = wc_i, \quad b_i = wd_i, \quad \text{and} \quad \bar{d} \text{ is a reduct of } \bar{c}.$$

We start by proving the following claims.

(I) If \bar{a} and \bar{b} are similar, then

$$\langle \mathfrak{S}_{\text{bin}}, \bar{a} \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{S}_{\text{bin}}, \bar{b} \rangle.$$

(II) Every tuple \bar{a} has a reduct \bar{b} .

(I) follows by Proposition V.1.11.

(11) Let \bar{a} be an n -tuple. We prove the claim by induction on n . If $n = 0$, the claim is trivial. Hence, suppose that $n > 1$. Set $w := a_0 \sqcap \dots \sqcap a_{n-1}$. By choice of K , there exists a word w' of length less than K such that

$$\langle \mathfrak{T}_{\text{bin}}, w' \rangle \equiv_{\text{MSO}}^m \langle \mathfrak{T}_{\text{bin}}, w \rangle.$$

Let \bar{a}'_i be the word obtained from a_i by replacing the prefix w by w' . Then \bar{a}' is similar to \bar{a} . For $d \in \{0, 1\}$, let

$$I_d := \{ i < n \mid wd \leq_{\text{pf}} a_i \}.$$

By inductive hypothesis, the tuple $\bar{a}'|_{I_d}$ has a reduct \bar{c}_d over w' . Let \bar{b} be the tuple obtained from \bar{a}' by replacing the subtuples $\bar{a}'|_{I_0}$ and $\bar{a}'|_{I_1}$ by, respectively, \bar{c}_0 and \bar{c}_1 . Then \bar{b} is a reduct of \bar{a} .

To conclude the proof, fix a guarded n -tuple. We define $(b_{ij})_{i,j}$ as follows. For $i, j < n$, set

$$J_{ij} := \{ l < n \mid a_j \sqcap a_i <_{\text{pf}} a_l \sqcap a_i \}.$$

By (11), we can find a reduct \bar{c}_{ij} of $\bar{a}|_{J_{ij}}$ over the word $a_i \sqcap a_j$. If $J_{ij} = \emptyset$, we set $b_{ij} := a_j$. Otherwise, $b_{ij} := (c_{ij})_l$, for some $l \in J_{ij}$. We choose both \bar{c}_{ij} and the index l such that $|b_{ij}|$ is minimal. It follows that $J_{ij} \subseteq J_{i'j'}$ implies $|b_{ij}| \leq |b_{i'j'}|$.

We claim that the family $(b_{ij})_{i,j}$ is the desired nK -code of \bar{a} . By construction, $J_{ii} = \emptyset$ implies that $b_{ii} = a_i$. Furthermore,

$$b_{ij}/nK = (c_{ij})_l/nK <_{\text{pf}} a_i \sqcap a_j \leq_{\text{pf}} (c_{ij})_l = b_{ij}.$$

For the last condition, fix indices i, j, l such that $a_i \sqcap a_l <_{\text{pf}} a_j \sqcap a_l$. Note that the tuple \bar{d} obtained from \bar{a} by replacing the subtuple $\bar{a}|_{J_{il}}$ by \bar{c}_{il} is similar to \bar{a} . In the same way, the tuple \bar{d}' obtained from \bar{d} by replacing the subtuple $\bar{d}|_{J_{jl}}$ by \bar{c}_{jl} is similar to \bar{d} and, hence, to \bar{a} . Hence, it follows by (1) that $\bar{d}' \in R$. Since $b_{il}, b_{jl} \in \bar{d}'$, the pair $\langle b_{il}, b_{jl} \rangle$ is guarded. Finally,

$$b_{il}/K <_{\text{pf}} a_i \sqcap a_l <_{\text{pf}} a_j \sqcap a_l \leq_{\text{pf}} b_{jl},$$

which implies that $b_{il}/K \leq_{\text{pf}} b_{jl}/K$ since, as mentioned above, we have $|b_{il}| \leq |b_{jl}|$. Consequently, $b_{il} \vdash_{nK} b_{jl}$. \square

At last, we are able to prove the other direction.

Lemma 8.24. *Let \mathfrak{Q} be tree-interpretable. There exist constants $m_0, k_0 < \omega$ such that the structure $\langle \mathfrak{Q}, \chi_k^m \rangle$ is MSO-definable in $\langle A, \leq_k, \chi_k^m \rangle$, for all $m \geq m_0$ and $k \geq k_0$.*

Proof. Let r be the maximal arity of a relation of \mathfrak{Q} , fix m_0 such that every relation of \mathfrak{Q} is MSO_{m_0} -definable, and set $k_0 := rK$ where K is the constant from Lemma 8.23.

Let $m \geq m_0, k \geq k_0$, and let R be an n -ary relation of $\langle \mathfrak{Q}, \chi_k^m \rangle$. We prove that R is definable in $\langle A, \leq_k, \chi_k^m \rangle$ by induction on n .

If $n = 1$, we can decide whether an element a belongs to R by looking at $\text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, a)$. This theory can be computed from the colour $\chi_k^m(a)$.

Suppose that $n > 1$ and consider a tuple $\bar{a} \in A^n$. By Proposition V.1.11 it follows that, in order to determine whether or not $\bar{a} \in R$ holds, it is sufficient to know

- the branching type H of \bar{a} and
- the theories $\text{Th}_{\text{MSO}}^m(\mathfrak{T}_{\text{bin}}, (a_i \sqcap a_l)^{-1}(a_j \sqcap a_l))$, for $\langle i, j, l \rangle \in H$.

By Lemma 8.23, the tuple \bar{a} has some nK -code $(b_{ij})_{i,j}$. Furthermore, this nK -code and the branching type are definable from \bar{a} using Lemma 8.22. Finally note that, since we have $b_{il} \vdash_{nK} b_{jl}$, for $\langle i, j, l \rangle \in H$, the above theories can be computed from the colours $\chi_k^m(b_{il}, b_{jl})$, for $\langle i, j, l \rangle \in H$.

It remains to prove that the colouring of \mathfrak{Q} is definable from that of $\langle A, \leq_k \rangle$. The colours of single elements is the same in both structures. Hence, consider a guarded pair $\langle a, a' \rangle$ in \mathfrak{Q} . By Lemma 8.23, there exists an element $b \in A$ such that

$$b \vdash_{2K} a, \quad b \vdash_{2K} a', \quad \text{and} \quad b/2K \leq_{\text{pf}} a \sqcap a' \leq_{\text{pf}} b.$$

Consequently, the colour $\chi_k^m(a, a')$ can be computed from $\chi_k^m(b, a)$ and $\chi_k^m(b, a')$. \square

The following results concludes the proof of Theorem 8.1.

Theorem 8.25. *Let \mathfrak{A} be a tree-interpretable structure. There exist constants $m, k < \omega$ and an MSO-formula φ such that $\langle \mathfrak{A}, \chi_k^m \rangle$ is the unique countable model of φ .*

Proof. The proof is analogous to the one of Proposition 8.18. Let m_o, k_o be the constants from Lemma 8.24 and m_i, k_i those from Proposition 8.18 (for the structure $\langle A, \leq_{k_o} \rangle$). We set

$$m := \max \{m_o, m_i\} \quad \text{and} \quad k := \max \{k_o, k_i\}.$$

We have shown in Lemmas 8.20 and 8.24 that the structures $\langle \mathfrak{A}, \chi_k^m \rangle$ and $\langle A, \leq_k, \chi_k^m \rangle$ are MSO-definable within each other. Let τ be an interpretation mapping the former to the latter.

Given a structure \mathfrak{X} , it follows that

$$\mathfrak{X} \cong \langle \mathfrak{A}, \chi_k^m \rangle \quad \text{iff} \quad \tau(\mathfrak{X}) \cong \tau(\langle \mathfrak{A}, \chi_k^m \rangle).$$

The later condition can be expressed in MSO with the help of Proposition 8.18. \square

Proof of Theorem 8.1. Let φ be the formula from the preceding theorem and let \bar{P} be the predicates encoding the colouring χ_k^m . Then the GSO-formula $\exists \bar{P} \varphi$ has a unique countable model: the structure \mathfrak{A} . \square

Notes

A good survey on finitely presentable structures can be found in [9]. The class of prefix-recognisable relations were introduced in [131] under the name of *special relations*. The main result of that paper is the characterisation in Proposition 2.4. Theorem 1.2 is from [49]. Prefix-recognisable graphs were independently introduced in [44, 45] using the point of view of formal language theory. Their characterisation in terms of pushdown automata is due to [198], while the connection to VR-grammars is from [11]. Regular path-interpretations were introduced in [44] under the name ‘inverse rational substitution’.

Originally, VR-equational and HR-equational structures and classes were introduced in [55, 54, 61] in the study of graph grammars. For more information see [60].

Higher-order pushdown automata were introduced in [138, 139]. The characterisations of the Caucal Hierarchy in terms of path-interpretations and higher-order pushdown automata is due to [42, 209, 38]. There is also a considerable amount of work relating higher-order pushdown automata and certain recursion schemes for the λ -calculus. Damm and Goerdt [65, 66] have shown that higher-order pushdown automata correspond to a subclass of such schemes called ‘safe’. To capture the full class of all recursion schemes, so-called ‘collapsible pushdown automata’ are needed; see, e.g., [41] for an overview. Collapsible pushdown automata have been omitted from this chapter, since their configuration graphs have an undecidable MSO-theory [99].

Context-free groups were introduced in [144, 145], where the equivalences (1)–(5) of Theorem 6.4 were proved (under the additional assumption of accessibility which can be eliminated by a result of [74]). The equivalences between (5), (6), (8), and (10) are from [124]. Theorem 6.22 is a consequence of Theorems 4.A.6.5 and 5.A.9 in [197], see also [189, 73, 147].

The pumping lemma for the Muchnik iteration is from [18]. (Note that the second part of that article concerning higher-order pushdown automata is wrong. The correct result can be found in [152].)

The axiomatisation of tree-interpretable structures is from [16] generalising the corresponding result for HR-equational graphs in [56].

XII Automatic Structures

1 Automatic Presentations

WE HAVE ALREADY STUDIED certain finitely-presentable infinite structures in Chapter XI. In the following we take a look at other such classes where this time the structures are encoded by automata of some kind. To do so we will need automata that recognise relations, i.e., sets of tuples, instead of sets of single words. We could use so-called multi-head automata for this task, but it is much simpler to encode tuples of words as a single word and use ordinary automata instead. The idea is to simply lay all input words (or trees) on top of each other and simultaneously read a letter from each of them. Thus, a k -tuple of words over an alphabet Σ is turned into a single word over the alphabet Σ^k . One technicality we have to deal with is supporting words of different length. To do so we pad the shorter words by a special blank symbol \square to make all words of the tuple the same length. The same idea also works for trees. The formal definition is as follows.

Definition 1.1. (a) For an alphabet Σ , we denote by $\Sigma_{\square} := \Sigma + \square$ the extension by a new *blank symbol* $\square \notin \Sigma$.

(b) Let s_0, \dots, s_{k-1} be either words or trees (finite or infinite) over the alphabet Σ , which we regard as functions $s_i : \text{dom}(s_i) \rightarrow \Sigma$. The *convolution*

$$\begin{bmatrix} a \\ b \\ a \end{bmatrix} \begin{bmatrix} b \\ a \\ a \end{bmatrix} \begin{bmatrix} b \\ \square \\ a \end{bmatrix} \begin{bmatrix} a \\ \square \\ c \end{bmatrix} \begin{bmatrix} c \\ \square \\ c \end{bmatrix} \begin{bmatrix} \square \\ \square \\ b \end{bmatrix} \begin{bmatrix} \square \\ \square \\ a \end{bmatrix} \begin{bmatrix} \square \\ \square \\ b \end{bmatrix} \begin{bmatrix} \square \\ \square \\ a \end{bmatrix}.$$

Figure 1: $abbac \otimes ba \otimes aaaccbaba$

$s_o \otimes \cdots \otimes s_{k-1}$ is the word/tree

$$s_o \otimes \cdots \otimes s_{k-1} : \bigcup_{i < k} \text{dom}(s_i) \rightarrow \Sigma_{\square}^k$$

over the alphabet Σ_{\square}^k where the labelling is given by

$$(s_o \otimes \cdots \otimes s_{k-1})(x) := \begin{bmatrix} s_o(x) \\ s_1(x) \\ \vdots \\ s_{k-1}(x) \end{bmatrix}, \quad \text{for } x \in \bigcup_{i < k} \text{dom}(s_i),$$

with the convention that $s_i(x) = \square$, if $x \notin \text{dom}(s_i)$. Usually, we write the elements of Σ_{\square}^k as columns as above, but sometimes we use rows $[c_o, \dots, c_{k-1}]$ instead to save vertical space. We also use the shorthand

$$\bar{s}^{\otimes} := s_o \otimes \cdots \otimes s_{k-1}.$$

(c) We identify a relation $R \subseteq (\Sigma^*)^k$ with the language

$$R^{\otimes} := \{ \bar{s}^{\otimes} \mid \bar{s} \in R \} \subseteq (\Sigma_{\square}^k)^*,$$

and similarly for relations on ω -words or on trees. In particular, we call R *regular* if R^{\otimes} is a regular language. \downarrow

Remark. Note that, strictly speaking, the operation \otimes is not associative. But we treat it as such by identifying $\bar{s}^{\otimes} \otimes \bar{t}^{\otimes} \in ((\Sigma_{\square}^m)_{\square} \times (\Sigma_{\square}^n)_{\square})^*$ with $\bar{s}\bar{t}^{\otimes} \in (\Sigma_{\square}^{m+n})^*$. \downarrow

Examples. The following relations are regular.

(a) The *equality* relation $=$ is given by

$$\left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ a \end{bmatrix} \right)^*.$$

(b) The *total relation* $\Sigma^* \times \Sigma^*$ is given by

$$\left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ b \end{bmatrix} \right)^* \cdot \left[\left(\sum_{a \in \Sigma} \begin{bmatrix} \square \\ a \end{bmatrix} \right)^* + \left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ \square \end{bmatrix} \right)^* \right],$$

and similarly for $\Sigma^* \times \cdots \times \Sigma^*$.

(c) The *prefix order* \leq_{pf} has the expression

$$\left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ a \end{bmatrix} \right)^* \cdot \left(\sum_{a \in \Sigma} \begin{bmatrix} \square \\ a \end{bmatrix} \right)^*.$$

(d) The *equal length* relation $=_{\text{len}} := \{ \langle u, v \rangle \mid |u| = |v| \}$ has the regular expression

$$\left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ b \end{bmatrix} \right)^*.$$

(e) The *lexicographic ordering* \leq_{lex} is given by

$$\begin{aligned} & \left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ a \end{bmatrix} \right)^* \cdot \left(\sum_{a \in \Sigma} \begin{bmatrix} \square \\ a \end{bmatrix} \right)^* \\ & + \left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ a \end{bmatrix} \right)^* \cdot \sum_{\substack{a, b \in \Sigma \\ a < b}} \begin{bmatrix} a \\ b \end{bmatrix} \cdot \left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ b \end{bmatrix} \right)^* \cdot \left[\left(\sum_{a \in \Sigma} \begin{bmatrix} \square \\ a \end{bmatrix} \right)^* + \left(\sum_{a \in \Sigma} \begin{bmatrix} a \\ \square \end{bmatrix} \right)^* \right]. \end{aligned}$$


(f) The *length-lexicographic ordering* \leq_{lex} has the regular expression

$$\left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ b \end{bmatrix} \right)^* \cdot \left(\sum_{a \in \Sigma} \begin{bmatrix} \square \\ a \end{bmatrix} \right)^+ + \left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ a \end{bmatrix} \right)^* \cdot \sum_{\substack{a, b \in \Sigma \\ a < b}} \begin{bmatrix} a \\ b \end{bmatrix} \cdot \left(\sum_{a, b \in \Sigma} \begin{bmatrix} a \\ b \end{bmatrix} \right)^*.$$

(g) For the alphabet $\Sigma = \{0, \dots, p-1\}$, several forms of addition are regular. We can encode natural numbers in base p with the least significant digit first, i.e.,

$$\sum_{i < n} a_i p^i \in \mathbb{N} \quad \text{is encoded by} \quad a_0 \cdots a_{n-1} \in [p]^*.$$

Then the relation $\{ \langle a, b, c \rangle \in \mathbb{N}^3 \mid a + b = c \}$ is regular. It is recognised by the following automaton (which keeps track of the carry bit).



(We have only given the part of the automaton that does not mention blanks \square . The full automaton is obtained from the above one by adding several more copies where some labels 0 are replaced by \square .)

Similarly, for the p -adic rationals $\mathbb{Z}[\frac{1}{p}] := \{k/p^n \mid k \in \mathbb{Z}, n \in \mathbb{N}\}$ we can encode a number

$$\sum_{i < n} a_i p^{-i} \in \mathbb{Z}[\frac{1}{p}] \cap [0, 1) \quad \text{by} \quad a_0 \cdots a_{n-1}.$$

Then the two relations

$$\begin{aligned} \{ \langle a, b, c \rangle \in (\mathbb{Z}[\frac{1}{p}] \cap [0, 1))^3 \mid a + b = c \} \\ \{ \langle a, b, c \rangle \in (\mathbb{Z}[\frac{1}{p}] \cap [0, 1))^3 \mid a + b = 1 + c \} \end{aligned}$$

are regular and can be recognised by automata similar to the one above. \lrcorner

Let us take a look at a few operations that allow us to construct regular relations.

Definition 1.2. A homomorphism $\varphi : \Sigma^* \rightarrow \Gamma^*$ is *uniform* if

$$|\varphi(a)| = |\varphi(b)|, \quad \text{for all } a, b \in \Sigma.$$

Similarly, a homomorphism for ω -words/trees/ ω -trees is called uniform if

$$\text{dom}(\varphi(a)) = \text{dom}(\varphi(b)), \quad \text{for all } a, b \in \Sigma. \quad \lrcorner$$

Lemma 1.3. The class of regular relations over Σ^* is closed under (I) boolean operations, (II) direct products, (III) projections, (IV) inverse uniform homomorphisms, and (V) uniform homomorphisms. The same holds for relations over ω -words, trees, and ω -trees.

The proofs are based on the standard constructions for regular languages and are left to the reader.

Exercise 1.1. Show that regular relations are not closed under homomorphisms that are not uniform. \lrcorner

We are interested in structures where every relation is regular. More precisely, structures where every element can be encoded by a finite word (several words are allowed to encode the same element) and where every relation is regular when using these encodings.

Definition 1.4. Let \mathfrak{A} be a Γ -structure, for some relational signature Γ .

(a) An *automatic presentation* of \mathfrak{A} (over the alphabet Σ) is a surjective partial function $\pi : \Sigma^* \rightarrow A$ such that the languages

$$\begin{aligned} L_\delta &:= \text{dom}(\pi), \\ L_&:= \{ u \otimes v \mid \pi(u) = \pi(v) \}, \\ L_R &:= \pi^{-1}[R]^\otimes, \quad \text{for } R \in \Gamma, \end{aligned}$$

are all regular.

(b) The structure \mathfrak{A} is *automatic* if it has an automatic presentation. A structure with functions is automatic, if it is automatic when we replace every function $f : A^n \rightarrow A$ by its graph $G_f \subseteq A^{n+1}$.

(c) If, instead of languages of finite words, we use languages of infinite words, finite trees, or infinite trees, we speak of, respectively, ω -automatic, tree automatic, and ω -tree automatic presentations or structures. \lrcorner

Remark. (a) Sometimes we will be sloppy and also use the term automatic presentation for the tuple $\langle L_\delta, L_&, (L_R)_R \rangle$, leaving the map π implicit. Note that π can be recovered from $\langle L_\delta, L_&, (L_R)_R \rangle$ (up to isomorphism) since

$$\mathfrak{A} \cong \langle L_\delta, (L_R)_R \rangle / L_&.$$

(b) If not explicitly stated otherwise, all trees in this chapter will be assumed to be successor-ordered binary trees, i.e., trees with two successor

relations suc_o and suc_i where every non-leaf has exactly two successors. Infinite trees are assumed not to have leaves.

(c) When specifying presentations, having to deal with blanks \square is often a nuisance. When defining the language L_R for a relation R we will therefore often be sloppy and only present some language L' such that

$$L_R = L' \cap L_\delta \otimes \cdots \otimes L_\delta.$$

Regularity of L_R then follows by Lemma 1.3.

Frequently, it is also the case that the blank symbol \square is 'equivalent' to some letter c of the alphabet. For instance, when encoding numbers \square is usually identified with the digit o . In such cases we will use the letter c throughout with the implicit understanding that, at the end of every word, an arbitrary number of occurrences of c can be changed to \square . We can do so since regular relations are closed under inverse uniform homomorphisms.

Let us give some examples to get a feel for automatic structures. Note that, while some of the following claims are quite easy to prove, others require a major effort to establish. The missing proofs will be given later in this chapter.

Examples. (a) $\langle \mathbb{N}, +, \leq \rangle$, $\langle \mathbb{Z}, +, \leq \rangle$, and $\langle \mathbb{N} \cup \{\infty\}, \min, + \rangle$ are automatic. To see this, note that we can encode every natural number in binary with the least significant bit first. Using this encoding, we have

$$L_\delta := (o + i)^+, \\ L = ([\begin{smallmatrix} o \\ o \end{smallmatrix}] + [\begin{smallmatrix} i \\ i \end{smallmatrix}])^* ([\begin{smallmatrix} o \\ o \end{smallmatrix}]^* + [\begin{smallmatrix} i \\ i \end{smallmatrix}]^*).$$

The language L_+ for addition was already given in the example on page 834. It is straightforward to check that the relations \leq and \min are also regular with this encoding. For \mathbb{Z} , we have to add a sign bit at the beginning of every word (an alternative encoding will be given in Lemma 2.6 below).

(b) Let $1 < p < \omega$. The structure $\langle \mathbb{Z}[1/p], +, |_p, \leq \rangle$ is automatic, where

$$x |_p y \quad : \text{iff} \quad x \text{ is a power of } p \text{ dividing } y.$$

Every element of $\mathbb{Z}[1/p]$ has a finite expansion $\pm a_k \dots a_o \cdot b_o \dots b_k$ in base p . We can encode such a number by the word

$$s \begin{bmatrix} a_o \\ b_o \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix} \cdots \begin{bmatrix} a_{k-1} \\ b_{k-1} \end{bmatrix} \begin{bmatrix} a_k \\ b_k \end{bmatrix}$$

over the alphabet $[p]^2$, where $s \in [p]^2$ encodes the sign. Using similar expressions as those in the example on page 834, it follows that the relations $+$, $|_p$, and \leq are regular when using this encoding.

(c) In the same way, one can show that the quotient $\langle \mathbb{Z}[1/p]/\mathbb{Z}, + \rangle$ is automatic, and that $\langle \mathbb{R}, + \rangle$ and the additive group of the p -adic numbers $\langle \mathbb{Z}_p, + \rangle$ are ω -automatic.

(d) We will show in Proposition 4.1 below that the configuration graph of every Turing machine is automatic.

(e) We will prove in Section 9 that $\langle \mathbb{Q}, + \rangle$ is not automatic.

(f) We will show in Section 8 that $\langle \mathbb{N}, \cdot \rangle$ is tree-automatic but not automatic.

(g) We will prove in Section 8 that the random graph is not automatic.

The following relationships follow immediately from the definitions.

Proposition 1.5.

- (a) Every finite structure is automatic.
- (b) Every automatic structure is countable, ω -automatic, and tree automatic.
- (c) Every tree-automatic structure is countable and ω -tree automatic.
- (d) Every ω -automatic structure is ω -tree automatic.

Proposition 1.6. Every tree-interpretable structure is automatic. There are automatic structures that are not tree-interpretable.

Proof. Every prefix-recognisable relation is obtained from regular languages using concatenation, union, and direct products. By Lemma 1.3, the class of regular relations is closed under these operations. Hence, every prefix-recognisable relation is regular.

The second statement follows from the fact that there are automatic structures like $\langle \mathbb{Z}, + \rangle$ with an undecidable MSO-theory. \square

The main reason why one is interested in automatic structures is that their first-order theory is decidable. In fact, as we will show in Section 7, this remains true if we add counting quantifiers.

Theorem 1.7. *Given an automatic presentation of a structure \mathfrak{A} and an FO-formula $\varphi(\bar{x})$, one can effectively construct an automaton recognising the relation $\varphi^{\mathfrak{A}}$.*

Proof. We construct the desired automaton \mathcal{A} by induction on φ . Suppose that $\bar{x} = \langle x_0, \dots, x_{m-1} \rangle$. If φ is an atomic formula $Rx_{i_0} \dots x_{i_{n-1}}$ or $x_{i_0} = x_{i_1}$, we have

$$\varphi^{\mathfrak{A}} = h^{-1}[L_R] \quad \text{or} \quad \varphi^{\mathfrak{A}} = h^{-1}[L_=],$$

where L_R and $L_ =$ are the languages from the presentation π and $h : \Sigma_{\square}^m \rightarrow \Sigma_{\square}^n$ is the homomorphism defined by

$$h([c_0, \dots, c_{m-1}]) := [c_{i_0}, \dots, c_{i_{n-1}}], \quad \text{for } c_0, \dots, c_{m-1} \in \Sigma_{\square}.$$

By Lemma 1.3, the regular relations are closed under all operations of FO. For quantifiers, there is one technical issue we have to pay attention to. Given a formula of the form $\exists y \varphi(\bar{x}, y)$, we can take the language for $\varphi(\bar{x}, y)$ and project it down to the components corresponding to the variables \bar{x} . For languages of ω -words and ω -trees this is sufficient, but for finite words or finite trees, we have to be more careful since the resulting language L might contain all-blank symbols $[\square, \dots, \square]$ (if the word/tree y can be larger than all the \bar{x}). For words, we can remove these symbols by first taking the quotient LK^{-1} where K is the set of all words containing only all blank letters. Then the intersection with $\Sigma^* \times \dots \times \Sigma^*$ (which is regular as we have seen in the example on page 834) produces the desired language for $\exists y \varphi(\bar{x}, y)$. For languages of finite trees, we can use a similar construction. \square

Corollary 1.8. *Every ω -tree automatic structure has a decidable first-order theory.*

Proof. Given a first-order formula φ , we can compute the corresponding automaton \mathcal{A} and check that it recognises a non-empty language. \square

Finally, let us make a small observation that sometimes can be used to simplify proofs involving automatic structures.

Lemma 1.9. *Every automatic structure has a presentation over a binary alphabet. The same holds for ω -automatic, tree automatic, and ω -tree automatic structures.*

Proof. We present the proof for a tree automatic structure \mathfrak{A} . The other cases are analogous. Let π be a presentation over an alphabet Σ of size k . Fix some finite tree t with at least $\log_2 k$ vertices and let T be the set of all labellings of t with labels 0 and 1. Then there exists an injective function $\Sigma \rightarrow T$ which induces a uniform homomorphism between the set of Σ -labelled trees and the set of $[2]$ -labelled ones. By Lemma 1.3, the image of π under this homomorphism is again regular. \square

Exercise 1.2. Prove that every tree automatic structure has a tree automatic presentation over a unary alphabet. \dashv

2 Interpretations

The representation of automatic structures by lists of automata is mainly useful for algorithmic purposes. But when proving statements about them other encodings are often more convenient. Here we derive a purely logical representation based on interpretations of the following form.

Definition 2.1. (a) A k -dimensional first-order interpretation (from the signature Σ to Γ) is given by a list of FO $[\Sigma]$ -formulae

$$\tau = \langle \delta(\bar{x}), \varepsilon(\bar{x}, \bar{y}), (\varphi_R(\bar{x}_0, \dots, \bar{x}_{n_R-1}))_{R \in \Gamma} \rangle$$

where $\bar{x}, \bar{y}, \bar{x}_i$ are k -tuples of variables and n_R is the arity of the relation R . Given a Σ -structure \mathfrak{A} , it produces the Γ -structure

$$\tau(\mathfrak{A}) := \langle \delta^{\mathfrak{A}}, (\varphi_R^{\mathfrak{A}})_R \rangle / \approx,$$

where \approx is the equivalence relation on A^k generated by the relation $\varepsilon^{\mathfrak{A}}$, i.e.,

$$\approx := (\varepsilon^{\mathfrak{A}} \cup (\varepsilon^{\mathfrak{A}})^{-1})^*.$$

In case we do not want to specify the number k , we speak of a *multi-dimensional* interpretation.

(b) We write $\mathfrak{A} \leq_{\text{FO}} \mathfrak{B}$ if, for some $k < \omega$, there exists a k -dimensional FO-interpretation mapping \mathfrak{B} to \mathfrak{A} .

(c) We denote the infinite binary tree by $\mathfrak{T}_{\text{bin}} := \langle [2]^*, \leq_{\text{pf}}, \text{suc}_o, \text{suc}_i \rangle$.

We start with the observation that automatic structures are closed under first-order interpretations.

Theorem 2.2. *The classes of (I) automatic structures, (II) ω -automatic structures, (III) tree automatic structures, and (IV) ω -tree automatic structures are closed under multi-dimensional FO-interpretations.*

Proof. Let \mathfrak{A} be an automatic structure (of any of the four kinds) and let $\tau = \langle \delta, \varepsilon, (\varphi_R)_R \rangle$ be a k -dimensional FO-interpretation. By Theorem 1.7, the languages $\delta^{\mathfrak{A}}$, $\varepsilon^{\mathfrak{A}}$, and $\varphi_R^{\mathfrak{A}}$ are regular. Hence, $\langle \delta^{\mathfrak{A}}, \varepsilon^{\mathfrak{A}}, (\varphi_R^{\mathfrak{A}})_R \rangle$ is a presentation of $\tau(\mathfrak{A})$ \square

The equivalence of regularity and MSO-definability immediately leads to the following characterisation of automatic structures in terms of interpretations.

Theorem 2.3. *Let \mathfrak{A} be a structure.*

- (a) \mathfrak{A} is automatic if, and only if, $\mathfrak{A} \leq_{\text{FO}} \mathcal{P}_{\text{fin}}(\omega, \leq)$.
- (b) \mathfrak{A} is ω -automatic if, and only if, $\mathfrak{A} \leq_{\text{FO}} \mathcal{P}(\omega, \leq)$.
- (c) \mathfrak{A} is tree automatic if, and only if, $\mathfrak{A} \leq_{\text{FO}} \mathcal{P}_{\text{fin}}(\mathfrak{T}_{\text{bin}})$.
- (d) \mathfrak{A} is ω -tree automatic if, and only if, $\mathfrak{A} \leq_{\text{FO}} \mathcal{P}(\mathfrak{T}_{\text{bin}})$.

Proof. (\Leftarrow) By Theorem 2.2, it is sufficient to find automatic presentations of the above structures.

(b) We encode each set $X \subseteq \omega$ by its characteristic function. This leads to the languages

$$L_\delta := (o + i)^\omega,$$

$$L_=: \left(\begin{bmatrix} o \\ o \end{bmatrix} + \begin{bmatrix} i \\ i \end{bmatrix} \right)^\omega$$

$$L_\subseteq := \left(\begin{bmatrix} o \\ o \end{bmatrix} + \begin{bmatrix} o \\ i \end{bmatrix} + \begin{bmatrix} i \\ i \end{bmatrix} \right)^\omega$$

$$L_\leq := \begin{bmatrix} o \\ o \end{bmatrix}^* \left(\begin{bmatrix} i \\ i \end{bmatrix} + \begin{bmatrix} i \\ o \end{bmatrix} \begin{bmatrix} o \\ o \end{bmatrix}^* \begin{bmatrix} o \\ i \end{bmatrix} \right) \begin{bmatrix} o \\ o \end{bmatrix}^\omega.$$

(a) is similar to (b), except that finite sets can be encoded by finite words. When ignoring blanks \square , the corresponding expressions for the languages look like in (b), except that the ω -power is replaced by a star $*$. Adding support for blanks can then be done as in the remark above.

(c) and (d) are analogous, except that we use trees. L_δ contains all trees; $L_=:$ contains all trees with letters $\begin{bmatrix} o \\ o \end{bmatrix}$ and $\begin{bmatrix} i \\ i \end{bmatrix}$; L_\subseteq contains all trees with letters $\begin{bmatrix} o \\ o \end{bmatrix}$, $\begin{bmatrix} o \\ i \end{bmatrix}$, and $\begin{bmatrix} i \\ i \end{bmatrix}$; and L_{suc_o} contains all trees where every position is labelled $\begin{bmatrix} o \\ o \end{bmatrix}$, except for two which are labelled by, respectively, $\begin{bmatrix} i \\ o \end{bmatrix}$ and $\begin{bmatrix} o \\ i \end{bmatrix}$, and where the second one is the o -successor of the first one. The language L_{suc_i} is defined analogously. All of the above properties are clearly regular.

(\Rightarrow) Fix a presentation $\langle L_\delta, L_=: (L_R)_R \rangle$ of \mathfrak{A} over the alphabet $\{o, i\}$, and let $\varphi_\delta, \varphi_=: \varphi_R$ be MSO-formulae defining these languages.

(b) Let us start with ω -automatic structures, since this is the easiest case. We encode a word $w \in \{o, i\}^\omega$ by the set

$$U_w := \{ i < \omega \mid a_i = i \}.$$

We can then turn the formulae $\varphi_\delta, \varphi_=: \varphi_R$ from the presentation into the desired interpretation as follows. Let φ be one of these formulae and let n be the arity of the corresponding relation. We will construct an MSO-formula $\varphi^*(\tilde{X})$ such that

$$\langle \omega, \leq, (P_c)_{c \in \Sigma^n} \rangle \models \varphi \quad \text{iff} \quad \langle \omega, \leq \rangle \models \varphi^*(U_{w_o}, \dots, U_{w_{n-1}}),$$

where $\langle \omega, \leq, (P_c)_{c \in \Sigma^n} \rangle$ is the structure encoding the word $w_o \otimes \dots \otimes w_{n-1}$. This formula $\varphi^*(\tilde{X})$ can be obtained from φ by replacing every atomic formula of the form $P_{[c_o, \dots, c_{n-1}]}z$ by

$$\bigwedge_{i < n} \vartheta_{c_i}(X_i, z),$$

where

$$\vartheta_1(X, z) := z \in X \quad \text{and} \quad \vartheta_0(X, z) := z \notin X.$$

This results in MSO-formulae $\varphi_\delta^*(X)$, $\varphi_=(X, Y)$, and $\varphi_R^*(\tilde{X})$ over the structure $\langle \omega, \leq \rangle$ with the desired property. According to Proposition I.2.2, we can translate them into FO-formulae $\hat{\varphi}_\delta(x)$, $\hat{\varphi}_=(x, y)$, and $\hat{\varphi}_R(\tilde{x})$ over the structure $\mathcal{P}\langle \omega, \leq \rangle$. The desired interpretation is $\tau = \langle \hat{\varphi}_\delta, \hat{\varphi}_=, (\hat{\varphi}_R)_R \rangle$.

(d) is entirely analogous to (b). We encode a tree t by the set of all vertices labelled 1. Then we can use the same construction for the formulae $\varphi^*(\tilde{X})$ and $\hat{\varphi}(\tilde{x})$.

(a) We proceed similarly to (b), but since we are now dealing with finite words, we have to take the length of a word into account. In particular, the alphabet might now contain \square . We therefore encode a word $w = a_0 \dots a_{n-1}$ by the set

$$\{n\} \cup \{i < n \mid a_i = 1\},$$

where the additional element n marks the end of the word. For the construction of $\varphi^*(\tilde{X})$ we replace atoms $P_{[c_0, \dots, c_{n-1}]}z$ by

$$\bigwedge_{i < n} \vartheta_{c_i}(X_i, z),$$

where

$$\vartheta_1(X, z) := z \in X \wedge \exists y[z < y \wedge y \in X],$$

$$\vartheta_0(X, z) := z \notin X \wedge \exists y[z < y \wedge y \in X],$$

$$\vartheta_\square(X, z) := \forall y[y \in X \rightarrow y \leq z].$$

In addition we have to relativise every quantifier to the elements that are less than the maximal element of $\bigcup_i X_i$. The rest of the construction then proceeds as in (b).

(c) We adapt the construction from (a). We encode a finite tree t by the set

$$\begin{aligned} & \{v \in \text{dom}(t) \mid t(v) = 1\} \\ & \cup \{v \in [2]^{<\omega} \mid v \notin \text{dom}(t), \text{ but } u \in \text{dom}(t) \text{ for all } u <_{\text{pf}} v\}, \end{aligned}$$

where the second set marks the bottom of the tree. Then the formulae for the letters are

$$\vartheta_1(X, z) := z \in X \wedge \neg \vartheta_\square(X, z),$$

$$\vartheta_0(X, z) := z \notin X \wedge \neg \vartheta_\square(X, z),$$

$$\vartheta_\square(X, z) := \forall y[z <_{\text{pf}} y \rightarrow y \in X].$$

Besides replacing the predicates $P_{\vec{c}}$ by these formulae we also have to

- ♦ relativise every quantifier to the set of all elements z with $y \in \bigcup_i X_i$, for some $y > z$, and
- ♦ add the condition that every free variable X_i is a valid encoding of a finite tree, i.e., that every branch β contains some vertex z with $z \in X_i$ and $y \notin X_i$, for all $y > z$.

The rest of the construction then proceeds as in the other cases. \square

We can use this theorem to derive other interpretation results. Let us start by introducing the corresponding structures.

Definition 2.4. (a) For $k, m, p \in \mathbb{N}$ with $p > 1$, we write

$$k \mid_p m \quad : \text{iff} \quad k = p^n \mid m, \quad \text{for some } n \in \mathbb{N}.$$

(b) The *equal-length relation* on finite words is defined by

$$u =_{\text{len}} v \quad : \text{iff} \quad |u| = |v|.$$

Theorem 2.5. Let \mathfrak{A} be a structure. The following statements are equivalent.

- (1) \mathfrak{A} is automatic.
- (2) $\mathfrak{A} \leq_{\text{FO}} \mathcal{P}_{\text{fin}}\langle \omega, \leq \rangle$
- (3) $\mathfrak{A} \leq_{\text{FO}} \langle [p]^*, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle$, for some number $p \geq 2$.
- (4) $\mathfrak{A} \leq_{\text{FO}} \langle \mathbb{N}, +, \mid_p \rangle$, for some number $p \geq 2$.

Proof. (1) \Leftrightarrow (2) was already proved in Theorem 2.3

(3) \Rightarrow (1) By Theorem 2.2, it is sufficient to prove that the structure $\langle [p]^*, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle$ is automatic. We obtain a presentation over the alphabet $\Sigma := [p]$ with languages

$$\begin{aligned} L_{\leq_{\text{pf}}} &= \left(\begin{bmatrix} \circ \\ \circ \end{bmatrix} + \cdots + \begin{bmatrix} p-1 \\ p-1 \end{bmatrix} \right)^* \cdot \left(\begin{bmatrix} \square \\ \square \end{bmatrix} + \cdots + \begin{bmatrix} \square \\ p-1 \end{bmatrix} \right)^*, \\ L_{\text{suc}_k} &= \left(\begin{bmatrix} \circ \\ \circ \end{bmatrix} + \cdots + \begin{bmatrix} p-1 \\ p-1 \end{bmatrix} \right)^* \cdot \begin{bmatrix} \square \\ k \end{bmatrix}, \\ L_{=_{\text{len}}} &= \left(\sum_{i,j < p} \begin{bmatrix} i \\ j \end{bmatrix} \right)^*. \end{aligned}$$

(4) \Rightarrow (3) It is sufficient to construct an interpretation

$$\langle \mathbb{N}, +, |_p \rangle \leq_{\text{FO}} \langle [p]^*, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle.$$

We encode a number $n \in \mathbb{N}$ in base p with the least significant digit first. To construct the formulae of the interpretation we start with a formula to access the digits in a word. The formula

$$\psi_k(x, y) := (\exists z. z =_{\text{len}} x) [\text{suc}_k(z) \leq_{\text{pf}} y \vee (k = \circ \wedge y \leq_{\text{pf}} z)]$$

states that the $|x|$ -th letter of $y \circ^w$ is equal to k . It follows that we can define p -divisibility by

$$\begin{aligned} \varphi_{|_p}(x, y) &:= \forall z [z <_{\text{pf}} x \rightarrow [\text{suc}_1(z) = x \\ &\quad \vee (\text{suc}_0(z) <_{\text{pf}} x \wedge \psi_0(z, y))]] \end{aligned}$$

and addition by

$$\begin{aligned} \varphi_+(x, y, z) &:= \exists u \forall v \left[\text{suc}_0(\langle \rangle) \leq_{\text{pf}} u \right. \\ &\quad \wedge \bigvee_{\langle k, l, c, d, m \rangle \in H} [\psi_k(v, x) \wedge \psi_l(v, y) \wedge \psi_c(v, u) \\ &\quad \left. \wedge \psi_d(\text{suc}_0(v), u) \wedge \psi_m(v, z)] \right], \end{aligned}$$

where H is the set of all tuples $\langle k, l, c, d, m \rangle \in [p]^2 \times [2]^2 \times [p]$ such that $k + l + c = dp + m$. (c is the old carry and d is the new one.)

(2) \Rightarrow (4) It is sufficient to find an interpretation

$$\mathcal{P}_{\text{fin}} \langle \omega, \leq \rangle \leq_{\text{FO}} \langle \mathbb{N}, +, |_p \rangle.$$

To do so, we encode a finite set $w \subseteq \omega$ by the number

$$n_w := \sum_{i \in w} p^i.$$

To construct the formulae of the interpretation we start with a formula for set membership. Setting

$$\begin{aligned} \psi(x, y) &:= \exists y_0 \exists y_1 [y = y_0 + y_1 \wedge y_0 < x \wedge x |_p y_1 \\ &\quad \wedge \neg(p x |_p y_1)] \end{aligned}$$

we have

$$\mathbb{N} \models \psi(k, n_w) \quad \text{iff} \quad k = p^i \text{ and } i \in w.$$

Consequently, we can define set inclusion and the order by

$$\begin{aligned} \varphi_{\subseteq}(x, y) &:= \forall z [\psi(z, x) \rightarrow \psi(z, y)], \\ \varphi_{\leq}(x, y) &:= \exists u \exists v [u \leq_{\text{pf}} v \wedge \forall u' [\psi(u', x) \leftrightarrow u' = x] \\ &\quad \wedge \forall v' [\psi(v', y) \leftrightarrow v' = v]]. \quad \square \end{aligned}$$

Lemma 2.6. $\langle \mathbb{Z}, + \rangle$ is automatic.

Proof. By the preceding theorem, it is sufficient to present an interpretation of $\langle \mathbb{Z}, + \rangle$ in $\langle \mathbb{N}, + \rangle$. We encode every integer $k \in \mathbb{Z}$ by the pairs $\langle m, n \rangle \in \mathbb{N}^2$ with $k = m - n$. Then addition and equality are given by

$$\begin{aligned} \langle m, n \rangle + \langle m', n' \rangle &:= \langle m + m', n + n' \rangle, \\ \langle m, n \rangle = \langle m', n' \rangle &\quad \text{iff} \quad m + n' = m' + n. \end{aligned}$$

This is expressible in first-order logic. \square

For ω -automatic structures we obtain a similar characterisation.

Theorem 2.7. *Let \mathcal{A} be a structure. The following statements are equivalent.*

- (1) \mathcal{A} is ω -automatic.
- (2) $\mathcal{A} \leq_{\text{FO}} \mathcal{P}(\omega, \leq)$.
- (3) $\mathcal{A} \leq_{\text{FO}} \langle [p]^\infty, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle$, for some number $p \geq 2$.

Proof. (1) \Leftrightarrow (2) was already proved in Theorem 2.3

(3) \Rightarrow (1) It is sufficient to prove that $\langle [p]^\infty, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle$ is ω -automatic. We obtain a presentation over the alphabet $\Sigma_\square := [p] + \square$ by encoding (I) infinite words $w \in [p]^\omega$ by themselves and (II) finite words $w \in [p]^*$ by the ω -word $w\square^\omega$. Then the set of these encodings and the relations \leq_{pf} , suc_k , and $=_{\text{len}}$ are all regular.

(2) \Rightarrow (3) It is sufficient to construct an interpretation

$$\mathcal{P}(\omega, \leq) \leq_{\text{FO}} \langle [2]^\infty, \leq_{\text{pf}}, (\text{suc}_k)_{k < p}, =_{\text{len}} \rangle.$$

We encode a set $X \subseteq \omega$ by its characteristic function. The relation $|x| \in y$ can then be defined by the formula

$$\psi(x, y) := \exists z [z =_{\text{len}} x \wedge \text{suc}_1(z) <_{\text{pf}} y].$$

Hence, we obtain formulae for inclusion and order by

$$\begin{aligned} \psi_\subseteq(x, y) &:= \forall z [\psi(z, x) \rightarrow \psi(z, y)], \\ \psi_\leq(x, y) &:= \exists u \exists v [u \leq_{\text{pf}} v \wedge \forall u' [\psi(u', x) \leftrightarrow u' = x] \\ &\quad \wedge \forall v' [\psi(v', y) \leftrightarrow v' = v]]. \end{aligned} \quad \square$$

For automatic and tree automatic structures, the closure result for interpretations extends to interpretations with parameters.

Lemma 2.8. *If \mathcal{A} is automatic or tree automatic, so is $\langle \mathcal{A}, \bar{a} \rangle$, for all $\bar{a} \in A^n$.*

Proof. Let $\langle L_\delta, L_-, (L_R)_R \rangle$ be a presentation of \mathcal{A} . For each $i < n$, we fix an encoding $u_i \in \pi^{-1}(a_i)$. Then the relations

$$\pi^{-1}(a_i) = \{ w \mid w \otimes u_i \in L_- \}$$

are regular. \square

Remark. The characterisation in Theorem 2.3 can be used to generalise the notion of an automatic structure as follows. A *higher-order automatic structure* is a structure \mathcal{A} such that

$$\mathcal{A} \leq_{\text{FO}} \mathcal{P}(\mathcal{C}^{* \cdots *}), \quad \text{for some finite structure } \mathcal{C},$$

where $-^*$ denotes the Muchnik iteration. Clearly, every ω -tree automatic structure is higher-order automatic. Furthermore, it follows immediately by Theorem V.3.5, that every higher-order automatic structure has a decidable first-order theory. Apart from that, not much is known about such structures. \downarrow

3 Closure Properties

Besides interpretations there are other natural operations the class of automatic structures is closed under. We start with disjoint unions and direct products.

Proposition 3.1. *If \mathcal{A} and \mathcal{B} are both automatic, ω -automatic, tree automatic, or ω -tree automatic, so are $\mathcal{A} \oplus \mathcal{B}$ and $\mathcal{A} \times \mathcal{B}$.*

Proof. Fix presentations

$$\langle L_\delta, L_-, (L_R)_R \rangle \quad \text{and} \quad \langle K_\delta, K_-, (K_R)_R \rangle$$

of, respectively, \mathcal{A} and \mathcal{B} over the alphabet $\Sigma := \{0, 1\}$. We present constructions for languages of finite words. The other cases are exactly the same.

We construct a presentation of the disjoint union $\mathcal{A} \oplus \mathcal{B}$ as follows. We encode elements of $A + B$ by prefixing them with 0 or 1 depending on which component they belong to. The corresponding presentation is

$$\begin{aligned} M_\delta &:= 0L_\delta + 1K_\delta, \\ M_- &:= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \cdot L_- + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \cdot K_-, \\ M_R &:= \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \cdot L_R + \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \cdot K_R. \end{aligned}$$

For the direct product, we encode pairs by their convolution. This yields the presentation

$$\begin{aligned} M_\delta &:= L_\delta \otimes K_\delta, \\ M_+ &:= h(L_+ \otimes K_+), \\ M_R &:= h(L_R \otimes K_R). \end{aligned}$$

where h is the relabelling by the canonical bijection $(\Sigma_\square^n)_\square^2 \rightarrow (\Sigma_\square^2)_\square^n$. \square

Definition 3.2. Let \mathfrak{A} be a structure and $c \in A$ some fixed element.

- (a) The ω -power \mathfrak{A}^ω of \mathfrak{A} is the infinite direct product $\prod_{n < \omega} \mathfrak{A}$.
- (b) The *infinite direct sum* $\mathfrak{A}^{(\omega)}$ is the substructure of the ω -power \mathfrak{A}^ω consisting of all sequence $(a_n)_{n < \omega}$ such that $a_n = c$, for all but finitely many n . \downarrow

Remark. We have omitted the element c from the notation for the direct sum. In most cases, the choice of c is obvious, e.g., if \mathfrak{A} contains a neutral element. \downarrow

Proposition 3.3. If \mathfrak{A} is finite, then $\mathfrak{A}^{(\omega)}$ is automatic and \mathfrak{A}^ω is ω -automatic.

Proof. For \mathfrak{A}^ω , there exists an ω -automatic presentation over the alphabet A consisting of the languages

$$L_\delta := A^\omega, \quad L_+ := \left(\sum_{a \in A} \begin{bmatrix} a \\ a \end{bmatrix} \right)^\omega, \quad L_R := R^\omega.$$

For $\mathfrak{A}^{(\omega)}$, we can use the same construction with slight modifications. Let c be the neutral element. Ignoring some special cases concerning blanks \square , we have

$$L_\delta := A^* \quad \text{and} \quad L_+ := \left(\sum_{a \in A} \begin{bmatrix} a \\ a \end{bmatrix} + \begin{bmatrix} c \\ \square \end{bmatrix} + \begin{bmatrix} \square \\ c \end{bmatrix} \right)^*.$$

For each relation R , we can set

$$L_R := \begin{cases} R^* & \text{if } \langle c, \dots, c \rangle \in R, \\ \emptyset & \text{otherwise.} \end{cases} \quad \square$$


Corollary 3.4. Every countably-dimensional vector space over a finite field is automatic.

Proof. Let \mathbb{K} be a finite field. Every finitely-dimensional vector space over \mathbb{K} is finite and, thus, automatic. If \mathfrak{V} is a countably infinitely dimensional vector space over \mathbb{K} , then $\mathfrak{V} \cong \mathbb{K}^{(\omega)}$ which, according to the preceding proposition, is also automatic. \square

Proposition 3.5.

- (a) If \mathfrak{A} is tree automatic or ω -tree automatic, then so is $\mathfrak{A}^{(\omega)}$.
- (b) If \mathfrak{A} is ω -tree automatic, so is \mathfrak{A}^ω .

Proof. (b) We encode a sequence $(a_n)_{n < \omega}$ where a_n is represented by the tree t_n , by the tree t consisting of an infinite branch where we attach to the n -th vertex the tree t_n as a subtree.



Given MSO-formulae defining the relations of \mathfrak{A} it is now straightforward to construct corresponding formulae for \mathfrak{A}^ω .

(a) is similar to (b), except that we use a finite path with the understanding that all elements not on the path are equal to the neutral element c . \square

Examples. Using the above closure properties, we can find automatic presentations for several groups and semigroups.

(a) $\langle \mathbb{N}, \cdot \rangle$ is tree automatic. We can interpret $\langle \mathbb{N}, \cdot \rangle$ in the structure $\langle \mathbb{N}, + \rangle^{(\omega)} \oplus \{0\}$ by encoding a product of primes $p_0^{k_0} \dots p_n^{k_n}$ by the tuple $\langle k_0, \dots, k_n \rangle$.

(b) The following structures are ω -automatic.

$$\begin{array}{lll} \text{(I)} \langle \mathbb{R}, + \rangle & \text{(III)} \langle \mathbb{R}, \cdot \rangle & \text{(IV)} \langle \mathbb{C}, + \rangle \\ \text{(II)} \langle \mathbb{R}/\mathbb{Z}, + \rangle & \text{(V)} \langle \mathbb{C}, \cdot \rangle & \end{array}$$

For the proof, note that we have seen in Lemma 2.6 that $\mathfrak{Z} := \langle \mathbb{Z}, +, \leq \rangle$ is automatic. It is also straightforward to construct an ω -automatic presentation of $\mathfrak{Z} := \langle \mathbb{R}/\mathbb{Z}, +, \leq \rangle$. This establishes (II). (I) now follows since we can interpret $\langle \mathbb{R}, + \rangle$ in $\mathfrak{Z} \oplus \mathfrak{Z}$. For (III), note that the exponential map provides an isomorphism $\exp : \langle \mathbb{R}, + \rangle \rightarrow \langle \mathbb{R}_+, \cdot \rangle$ (where \mathbb{R}_+ denotes the set of positive real numbers), and that we can interpret $\langle \mathbb{R}, \cdot \rangle$ in the structure $\langle \mathbb{R}_+, \cdot \rangle \oplus \{0\}$. For (IV), note that

$$\langle \mathbb{C}, + \rangle \cong \langle \mathbb{R}, + \rangle \times \langle \mathbb{R}, + \rangle.$$

Finally, for (V), we can use polar coordinates to interpret $\langle \mathbb{C}, \cdot \rangle$ in

$$\{0\} \oplus \langle \mathbb{R}/\mathbb{Z}, + \rangle \times \langle \mathbb{R}_+, \cdot \rangle.$$

Open Question. Are there other ω -automatic Lie groups? In particular, for which $n > 1$ are $\text{SO}(n)$ and $\text{SU}(n)$ are ω -automatic?

Finally, let us make a few remarks about closure under elementary substructures.

Exercise 3.1. Let \mathfrak{A} be an ω -automatic structure and let $\mathfrak{A}_0 \subseteq \mathfrak{A}$ be the substructure consisting of all ultimately periodic words. Prove that \mathfrak{A}_0 is an elementary substructure of \mathfrak{A} (but not necessarily automatic).

Exercise 3.2. Let \mathfrak{A} be an ω -tree automatic structure and let $\mathfrak{A}_0 \subseteq \mathfrak{A}$ be the substructure consisting of all regular trees. Prove that \mathfrak{A}_0 is an elementary substructure of \mathfrak{A} (not necessarily automatic).

Example. Let \approx_* be the equivalence relation on $\wp(\omega)$ saying that the symmetric difference of the given sets is finite. The atomless boolean algebra $\langle \wp(\omega), \cap, \cup \rangle / \approx_*$ is ω -automatic, but it follows by Theorem 10.2 below that it has no countable elementary substructure that is automatic.

Exercise 3.3. The Stupp-iteration $\mathfrak{A}^\#$ of a structure \mathfrak{A} is obtained from its Muchnik-iteration by removing the clone predicate. Prove that the Stupp-iteration of an automatic structure is automatic.

4 Undecidability

Unfortunately, most logics stronger than FO, like monadic second-order logic or various fixed-point logics, are undecidable. The reason is that reachability is undecidable for automatic structures.

Proposition 4.1. The configuration graph of every Turing machine is automatic.

Proof. A configuration consists of a triple $\langle q, h, w \rangle$ where q is the current state, h the head position, and w the tape contents. We can encode such a configuration by the word uqv where $uv = w$ and $|u| = h$. The set of all words of the form uqv is clearly regular. As each transition modifies a configuration only locally, the transition relation is also regular. For the latter, suppose that $u = u'a$ and $v = bv'$. Then we have transitions

$$\begin{array}{ll} u'apbv' \vdash u'acqv' & \text{if } \langle p, b, c, 1, q \rangle \in \Delta, \\ u'apbv' \vdash u'acqv' & \text{if } \langle p, b, c, 0, q \rangle \in \Delta, \\ u'apbv' \vdash u'qacv' & \text{if } \langle p, b, c, -1, q \rangle \in \Delta. \end{array}$$

These conditions can easily be checked by an automaton. \square

Proposition 4.2. There exists an automatic directed graph such that it is undecidable whether, for two given vertices x and y , there exists a path from x to y .

Proof. Let \mathcal{M} be a universal Turing machine. We can modify \mathcal{M} such that, before accepting, it erases its tape and moves its head to the beginning. Consequently, we may assume that \mathcal{M} as a unique accepting configuration c_{acc} . Let c_w be the initial configuration of \mathcal{M} for the input w . Then it is undecidable whether the configuration graph \mathfrak{G} of \mathcal{M} contains a path from c_w to c_{acc} . Furthermore, we have seen in Proposition 4.1 that \mathfrak{G} is automatic. \square

From an algebraic point of view, the class of automatic structures is unfortunately not that well-behaved. In the remainder of this section, we will present a first piece of evidence for this fact: the isomorphism problem for automatic structures is extremely complicated. Other, more concrete examples will be given in Sections 8 and 9 below.

Theorem 4.3. *The isomorphism problem for automatic structures is Σ_1^1 -complete.*

The proof consists of a reduction of the isomorphism problem for a certain class of trees to that for automatic structures. Note that every countably-branching successor tree can be encoded as a prefix-closed subset $T \subseteq \omega^{<\omega}$. (We do *not* require that $wk \in T$ and $i < k$ implies $wi \in T$, for $w \in \omega^{<\omega}$ and $i, k < \omega$.) We call such a tree *recursively enumerable* if T is a recursively enumerable subset of $\omega^{<\omega}$. We will make use of the following fact from computable model theory.

Theorem 4.4 (Goncharov, Knight [88]). *The isomorphism problem for recursively enumerable trees is Σ_1^1 -complete.*

Proposition 4.5. *Let \mathfrak{S} be the tree obtained from $\omega^{<\omega}$ by attaching to every vertex infinitely many paths of length n , for every $n < \omega$. Then \mathfrak{S} is automatic.*

Proof. We describe a 4-dimensional FO-interpretation of \mathfrak{S} in the tree $\langle \{0, 1\}^*, \text{suc}_0, \text{suc}_1, \leq_{\text{pf}}, =_{\text{len}} \rangle$. We can encode a vertex $\langle k_0, \dots, k_{n-1} \rangle \in \omega^{<\omega}$ as tuple

$$\langle 10^{k_0} 10^{k_1} \dots 10^{k_{n-2}} 10^{k_{n-1}}, \langle \rangle, \langle \rangle, \langle \rangle \rangle.$$

The edges between such tuples are

$$\langle w, \langle \rangle, \langle \rangle, \langle \rangle \rangle \rightarrow \langle w10^n, \langle \rangle, \langle \rangle, \langle \rangle \rangle,$$

which is clearly FO-definable. The attached paths consists of the vertices

$$\langle w, u, 0^n, 0^k \rangle, \quad \text{for } w, u \in \{0, 1\}^* \text{ and } k < n < \omega,$$

(n is the length of the path, w is the vertex it is attached to, and u specifies which of the infinitely many copies of the path the vertex belongs to) with edges

$$\begin{aligned} \langle w, \langle \rangle, \langle \rangle, \langle \rangle \rangle &\rightarrow \langle w, u, 0^n, 0^n \rangle \\ \text{and } \langle w, u, 0^n, 0^{k+1} \rangle &\rightarrow \langle w, u, 0^n, 0^k \rangle. \end{aligned}$$

Again these transitions are FO-definable. \square

For the reduction, we consider configuration graphs of Turing machines. These can be considerably simplified if we assume that the machine is reversible.

Definition 4.6. A Turing machine \mathcal{M} is *reversible* if its configuration graph consists of a disjoint union of finite paths and (one-way) infinite paths, and if every finite path ends in an accepting configuration. \urcorner

Lemma 4.7. *Every deterministic Turing machine \mathcal{M} can be transformed into an equivalent reversible one.*

Proof. This is a standard construction from computability theory. We only give a sketch. The main idea is that the machine maintains a log of all irreversible operations it performs. For instance, if the machine overwrites a cell on the tape it adds a note to the log containing the old contents of the cell. This ensures that every configuration has at most one predecessor.

To ensure that all finite computations are accepting, it is sufficient for the new machine to enter an infinite loop whenever the old one would reject the input. (To make this reversible it can, for instance, move the head to the right in each iteration of the loop, so no configuration repeats.) \square

Proof of Theorem 4.3. To see that the problem is in Σ_1^1 note that in existential second-order logic we can express that there exists a binary relation $I \subseteq \mathbb{N} \times \mathbb{N}$ that is the graph of an isomorphism between the two structures encoded by the given presentations.

Hence, it remains to prove hardness. For a reversible Turing machine \mathcal{M} , let $\text{Conf}(\mathcal{M})$ be its configuration graph and let $\mathcal{T}(\mathcal{M})$ be the graph obtained from the disjoint union of $\text{Conf}(\mathcal{M})$ and the tree \mathfrak{S} in Proposition 4.5 by adding edges from $w \in \omega^{<\omega} \subseteq S$ to the initial configuration $\langle q_0, 0, w \rangle$ of \mathcal{M} associated with the input word w .

The following two claims conclude the proof by providing a reduction of the isomorphism problem for recursively enumerable trees (which is Σ_1^1 -complete by Theorem 4.4) to the isomorphism problem for automatic structures.

Claim. *Let \mathcal{M} be a reversible Turing machine. $\mathcal{T}(\mathcal{M})$ is automatic and a presentation can be computed from \mathcal{M} .*

For the proof note that \mathfrak{S} and $\text{Conf}(\mathcal{M})$ are both automatic and automatic structures are closed under disjoint unions. Furthermore, with the encoding of configuration in the proof of Proposition 4.1, the additional edges belong to the relation

$$\{ \langle w, q_0 w \rangle \mid w \in \omega^{<\omega} \},$$

which is regular.

Claim. *Let S and T be recursive enumerable trees and let \mathcal{M} and \mathcal{N} be reversible Turing machines recognising S and T . Then*

$$S \cong T \quad \text{iff} \quad \mathcal{T}(\mathcal{M}) \cong \mathcal{T}(\mathcal{N}).$$

By construction of $\mathcal{T}(\mathcal{M})$, we have

$$w \in S \quad \text{iff} \quad \text{there is no infinite path attached to } w \text{ in } \mathcal{T}(\mathcal{M}).$$

It follows that every isomorphism $S \rightarrow T$ can be extended to one $\mathcal{T}(\mathcal{M}) \rightarrow \mathcal{T}(\mathcal{N})$. Conversely, every isomorphism $\mathcal{T}(\mathcal{M}) \rightarrow \mathcal{T}(\mathcal{N})$ restricts to one $S \rightarrow T$. \square

For ω -tree automatic structures, the isomorphism problem is even harder: it turns out that it is independent of set theory!

Definition 4.8. For $U, V \subseteq [2]^*$, we define

$$\begin{aligned} U \approx_* V &: \text{iff} \quad \text{the symmetric difference of } U \text{ and } V \text{ is finite,} \\ U \approx_{\text{ac}} V &: \text{iff} \quad \text{the symmetric difference of } U \text{ and } V \text{ does not} \\ &\quad \text{contain an infinite anti-chain with respect to the} \\ &\quad \text{prefix ordering.} \end{aligned}$$

Theorem 4.9 (Finkel, Todorćević [81]). *The question of whether*

$$\langle \wp([2]^*), \subseteq \rangle / \approx_* \cong \langle \wp([2]^*), \subseteq \rangle / \approx_{\text{ac}}$$

is independent of ZFC.

Theorem 4.10. *The isomorphism problem for ω -tree automatic structures is independent of ZFC.*

Proof. The two boolean algebras in the preceding theorem are ω -tree automatic. \square

5 Injective Presentations

Sometimes it can be convenient to choose a presentation with particularly useful properties. But one has to be careful when doing so since not all presentations are equal: some relations are only regular with respect to certain presentations, but not with respect to others.

Definition 5.1. Two presentations π_0 and π_1 are *equivalent* if, for every relation R ,

$$\pi_0^{-1}[R] \text{ is regular} \quad \text{iff} \quad \pi_1^{-1}[R] \text{ is regular.}$$

We say that π_0 and π_1 are *effectively equivalent* if, when given an automaton for $\pi_i^{-1}[R]$, one can effectively compute an automaton for $\pi_{1-i}^{-1}[R]$. \square

Example. The structure $\langle \omega, \text{succ} \rangle$ has several non-equivalent presentations.

(I) We can encode $n < \omega$ in unary 1^n . Then the successor function is

$$L_{\text{suc}} = \{ \langle 1^n, 1^{n+1} \rangle \mid n < \omega \},$$

which is regular.

(II) We can use binary encoding (with the least significant bit first). Then the successor function is

$$L_{\text{suc}} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}^* \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right)^*.$$

(III) We can represent a number $n < \omega$ by the pair $\langle i, j \rangle \in \omega \times \omega$ with $n = \frac{1}{2}(i+j)(i+j+1) + j$. Then the successor function becomes

$$\langle i+1, j \rangle \mapsto \langle i, j+1 \rangle \quad \text{and} \quad \langle 0, j \rangle \mapsto \langle j+1, 0 \rangle.$$

We obtain an automatic presentation by encoding each pair $\langle i, j \rangle$ by all words of the form $u\#v$ where u and v are binary representations of, respectively, i and j (as usual with the least significant bit first) and $\#$ is some separator.

To see that these presentations are not equivalent, note that the function $f : n \mapsto 2n$ is regular in the second presentation, but not in the first one. Furthermore, the order relation \leq is regular in the first two presentations, while in the third one

$$L_{\leq} = \{ \langle i, j \rangle \otimes \langle k, l \rangle \mid i+j < k+l \text{ or } [i+j = k+l \text{ and } j \leq l] \}$$

is not. J

One way to show that two presentations are equivalent is the following lemma. Note that this statement is far from the best possible, but it suffices for our needs.

Lemma 5.2. *Let π_0 and π_1 be presentations of the same structure \mathcal{A} . If the relation*

$$\pi_1^{-1} \circ \pi = \{ \langle w_0, w_1 \rangle \mid \pi_0(w_0) = \pi_1(w_1) \}$$

is regular, then π_0 and π_1 are effectively equivalent.

Exercise 5.1. Find effectively equivalent presentations π_0 and π_1 such that $\pi_1^{-1} \circ \pi_0$ is not regular. J

Our next goal is to extend Theorem 1.7 to logics with counting quantifiers. As a preparation, we take a look at presentations where each element has only a single encoding, which makes counting much easier. For automatic and tree automatic structures, we can always find such a presentation, but unfortunately not for ω -automatic or ω -tree automatic ones.

Definition 5.3. An (ω -/tree/ ω -tree) automatic presentation π is called *injective* if the encoding function $\pi : L_{\delta} \rightarrow A$ is injective. J

Remark. We can define injective presentations also in terms of interpretations. The presentation associated with an interpretation $\langle \delta, \varepsilon, (\varphi_R)_R \rangle$ in one of the structures from Theorems 2.3, 2.5, or 2.7 is injective if, and only if, the formula $\varepsilon(x, y)$ is equivalent to $x = y$. J

In the proof we make use of the orders \leq_{lex} and \leq_{llex} from Definitions VI.1.18 and VI.6.15, respectively (which we have already shown to be regular in the example on page 834).

Lemma 5.4. *Let Σ be a finite alphabet. The lexicographic order \leq_{lex} and the length-lexicographic order \leq_{llex} on Σ^* are regular.*

Corollary 5.5. *For every automatic structure \mathcal{A} , there exists a well-order \leq such that $\langle \mathcal{A}, \leq \rangle$ is also automatic and with an effectively equivalent presentation.*

Proof. Fix a presentation π of \mathcal{A} . Note that the length-lexicographic order \leq_{llex} is a regular well-ordering. Hence, so is the relation

$$a \leq b \quad : \text{iff} \quad \begin{array}{l} \text{the } \leq_{\text{llex}}\text{-least element of } \pi^{-1}(a) \text{ is } \leq_{\text{llex}}\text{-smaller} \\ \text{than the } \leq_{\text{llex}}\text{-least element of } \pi^{-1}(b). \end{array}$$

Furthermore, \leq is first-order definable in $\langle \mathcal{A}, \leq_{\text{llex}} \rangle$ and, therefore, regular. □

Proposition 5.6. *For every automatic presentation π , there exists an equivalent injective one. The same holds for tree automatic presentations.*

Proof. First, suppose that \mathfrak{A} has an automatic presentation $\langle L_\delta, L_-, (L_R)_R \rangle$. By Corollary 5.5, we can extend this presentation (or an equivalent one) by a well order \leq . Let U be the set of \leq -minimal representatives of every L_- -class. Then U is FO-definable and, hence, regular. We obtain the injective presentation with universe U and relations $L_R \cap (U \otimes \cdots \otimes U)$. Furthermore, the map taking an element of L_δ to its representative in U is FO-definable and, hence, regular. Consequently, the new presentation is effectively equivalent to the old one.

For the second claim, consider a tree automatic structure \mathfrak{A} with presentation π . We cannot use the same trick as for automatic structures since there is no regular well-ordering on the set of all trees. What we do instead is to represent an element of \mathfrak{A} by the ‘intersection’ of all trees encoding it. For $a \in A$, we call the intersection

$$C(a) := \bigcap_{t \in \pi^{-1}(a)} \text{dom}(t)$$

the *core* of a and the set $F(a)$ of minimal elements of $[2]^* \setminus C(a)$ its *frontier*. For trees t , we also write $C(t)$ and $F(t)$ instead of $C(\pi(t))$ and $F(\pi(t))$.

Let \mathcal{A} be a tree automaton recognising L_- and Q its set of states. The *profile* of a finite Σ -labelled tree t is the triple $\chi(t) := \langle F(t), t \upharpoonright C(t), \lambda \rangle$ where $\lambda : F(a) \rightarrow \wp(Q)$ maps each vertex $v \in F(a)$ to the set of states from which \mathcal{A} accepts the tree $\square \otimes t|_v$ (where \square denotes the tree entirely labelled by blanks). We claim that

$$\chi(t) = \chi(t') \quad \text{implies} \quad \pi(t) = \pi(t').$$

Hence, suppose that $\chi(t) = \langle F, s, \lambda \rangle = \chi(t')$.

By transitivity of equality and the fact that F is finite, it is sufficient to consider the case where there is a single vertex $v \in F$ with $t|_v \neq t'|_v$. Since $v \in F(t)$, there exists some tree r with $\pi(r) = \pi(t)$ such that $v \notin \text{dom}(r)$. Fix an accepting run ρ of \mathcal{A} on $r \otimes t$. To prove the claim, it is sufficient to find an accepting run ρ' of \mathcal{A} on $r \otimes t'$. Then the claim follows by transitivity and symmetry. Since $\rho(v) \in \lambda(v)$, there exists a run σ of \mathcal{A} on the tree $\square \otimes t'|_v$ with $\sigma(v) = \rho(v)$. We obtain the desired run ρ' by replacing in ρ

the subtree $\rho|_v$ by σ . (Note that $r \otimes t$ and $r \otimes t'$ agree on all vertices not in the subtree attached to v .)

Having established the above claim, we can conclude the proof as follows. Note that we can encode every profile $\langle F, s, \lambda \rangle$ by a finite tree over a suitable alphabet. Furthermore, for every element $a \in A$, the set

$$\chi(a) := \{ \chi(t) \mid t \in \pi^{-1}(a) \}$$

of possible profiles is finite (since the first component $F(t)$ is the same for all of them). Finally, we can define a linear ordering on $\chi(a)$ by

$$\langle F, s, \lambda \rangle < \langle F, s', \lambda' \rangle$$

iff $s(v) < s'(v)$ or $(s = s' \text{ and } \lambda(v) < \lambda'(v))$,

where v is the \leq_{lex} -least vertex with $s(v) \neq s'(v)$ or $\lambda(v) \neq \lambda'(v)$.

As $\chi(a)$ is finite and non-empty, it has a minimal element. We use the tree encoding this element as the new representative of a . Again, the resulting presentation is equivalent to the original one since the function mapping a tree to the encoding of its equivalence class is regular. \square

For ω -automatic structures we obtain a negative result. The proof is based on a translation of the problem into the realm of topology. Recall the definitions of a Borel set and of the Cantor topology from Section V.5.

Definition 5.7. Let Σ be an alphabet. A *Borel structure* is a structure \mathfrak{D} where the universe $D \subseteq \Sigma^\omega$ is Borel and every n -ary relations R is a Borel subsets of $(\Sigma^\omega)^n$. A *Borel presentation* of a Γ -structure \mathfrak{A} is a Borel structure $\langle D, \sim, (R^\mathfrak{D})_{R \in \Gamma} \rangle$ such that $\langle D, (R^\mathfrak{D})_R \rangle / \sim \cong \mathfrak{A}$. \lrcorner

Below we will construct an ω -automatic structure that has no injective Borel presentation. Before doing so, let us show that every injective ω -automatic presentation is Borel (and injective).

Lemma 5.8. Every ω -automatic presentation is Borel.

Proof. Let $\mathfrak{D} = \langle L_\delta, L_-, (L_R)_R \rangle$ be an ω -automatic presentation of some structure \mathfrak{A} over the alphabet Σ . According to Lemma V.5.23, the languages L_δ, L_-, L_R are all Borel. Hence, \mathfrak{D} is a Borel presentation. \square

The example below makes use of the observation that there is basically only one Borel presentation of the boolean algebra $\langle \wp(\omega), \subseteq \rangle$.

Lemma 5.9. *Let $\mathfrak{D} = \langle D, \sim, \leq \rangle$ and $\mathfrak{D}' = \langle D', \sim', \leq' \rangle$ be two Borel presentations of $\mathfrak{B} := \langle \wp(\omega), \subseteq \rangle$. Every isomorphism $\varphi : \mathfrak{D}/\sim \rightarrow \mathfrak{D}'/\sim'$ is Borel.*

Proof. By Lemma V.5.20, it is sufficient to prove that the graph of φ is Borel. Let $(a_n)_{n < \omega}$ be an enumeration of all the atoms of \mathfrak{D}/\sim and set $b_n := \varphi(a_n)$. Then

$$\varphi(x) = y \quad \text{iff} \quad a_n \leq x \Leftrightarrow b_n \leq' y, \quad \text{for all } n < \omega.$$

Consequently, the graph of φ is equal to $\bigcap_n R_n$, where

$$R_n := \{ \langle x, y \rangle \mid a_n \leq x \Leftrightarrow b_n \leq y \}.$$

As Borel sets are closed under countable intersections, it therefore remains to prove that each relation R_n is Borel. By Lemma V.5.19 (c), the relations

$$P_n := \{ x \mid a_n \leq x \} \quad \text{and} \quad Q_n := \{ y \mid b_n \leq y \}$$

are both Borel. Hence, it follows by Lemma V.5.19 (b) that so is

$$R_n = (P_n \times Q_n) \cup (B \setminus P_n) \times (B \setminus Q_n). \quad \square$$

Theorem 5.10. *There exists an ω -automatic structure that does not have an injective presentation.*

Proof. Set $\mathfrak{B} := \langle \wp(\omega), \subseteq \rangle$ and $\mathfrak{A} := \text{copy}_2(\mathfrak{B})$, and let \approx be the equivalence relation that is equal to the identity on the first copy of \mathfrak{B} and equal to \approx_* on the second one. Then the structure $\langle \mathfrak{A}, \approx \rangle$ is ω -automatic. Hence, so is the quotient \mathfrak{A}/\approx .

We claim that \mathfrak{A}/\approx has no injective ω -automatic presentation. According to Lemma 5.8, it is sufficient to prove that \mathfrak{A}/\approx has no injective Borel presentation. For a contradiction, suppose otherwise. Let $\langle D, =, \leq, L, R, I \rangle$ be such a presentation where L encodes the left copy of \mathfrak{B} , R the right one, and $I \subseteq L \times R$ relates the two copies of each element of \mathfrak{B} . Let $i : L \rightarrow R$

be the function whose graph is I and let $\varphi_0 : L \rightarrow \mathfrak{B}$ and $\varphi_1 : R \rightarrow \mathfrak{B}/\approx_*$ be the restrictions of the isomorphism $\varphi : \langle D, \leq, L, R, I \rangle \rightarrow \mathfrak{A}/\approx$. Since I is Borel, so is i , by Lemma V.5.20. Let $\langle [2]^\omega, =, \leq \rangle$ be the canonical Borel presentation of \mathfrak{B} and let $\psi : [2]^\omega \rightarrow \wp(\omega)$ be the associated isomorphism. By Lemma 5.9, the map $\varphi^{-1} \circ \psi : [2]^\omega \rightarrow D$ is Borel. Consequently, $f := i \circ \varphi_0^{-1} \circ \psi : 2^\omega \rightarrow D$ is also Borel. Let $p : \mathfrak{B} \rightarrow \mathfrak{B}/\approx_*$ be the quotient map. Then

$$f = i \circ \varphi_0^{-1} \circ \psi = \varphi_1^{-1} \circ p \circ \psi$$

implies that

$$\begin{aligned} f(x) = f(y) & \quad \text{iff} \quad (\varphi_1^{-1} \circ p \circ \psi)(x) = (\varphi_1^{-1} \circ p \circ \psi)(y) \\ & \quad \text{iff} \quad (p \circ \psi)(x) = (p \circ \psi)(y) \\ & \quad \text{iff} \quad x \approx_* y. \end{aligned}$$

A contradiction to Theorem V.5.22. \square

Using a bit more machinery from Descriptive Set Theory, the preceding theorem can be strengthened as follows. We omit the proof, a sketch of which can be found in [103].

Theorem 5.11. *There exists an ω -automatic structure that does not have an injective ω -tree automatic presentation.*

6 Partition Theorems

A partition theorem is a statement of the form: every sufficiently large structure of a given kind contains a large configuration of a certain type. A typical example is the Theorem of Ramsey saying that every sufficiently large coloured set has a large homogeneous subset. In this section we collect several such results specific to ω -automatic structures. We start with the following auxiliary relation, which turns out to be central in our proofs.

Definition 6.1. For two infinite words or trees s and t , we define

$$s \approx_* t \quad : \text{iff} \quad \text{there are only finitely many positions } x \text{ with } s(x) \neq t(x).$$

We denote the \approx_* -class of s by $[s]_*$.

We have shown in the previous section that, for every automatic equivalence relation, we can find a regular set of representatives, while this is no longer the case for ω -automatic relations. In particular, there does not exist a regular well-ordering on Σ^ω . The following lemma provides a weak substitute for such an ordering, which sometimes can be used instead.

Lemma 6.2. *There exists a regular ternary relation $x \sqsubseteq^z y$ on Σ^ω such that, for every $u \in \Sigma^\omega$, the induced binary relation \sqsubseteq^u well-orders the \approx_* -class $[u]_*$.*

Proof. Fix an order \leq on Σ . For $x, y \in \Sigma^\omega$, we set

$$\delta(x, y) := \sup \{ i < \omega \mid x(i) \neq y(i) \},$$

and

$$\begin{aligned} x \sqsubseteq^u y \quad : \text{iff} \quad & x = y, \quad \text{or} \quad \delta(u, x) < \delta(u, y), \\ & \text{or} \quad \delta(u, x) = \delta(u, y) \text{ and} \\ & x(\delta(u, x)) \leq y(\delta(u, x)). \end{aligned}$$

Clearly, \sqsubseteq^u is reflexive. For transitivity, suppose that $x \sqsubseteq^u y \sqsubseteq^u z$. If $\delta(u, x) < \delta(u, y)$ or $\delta(u, y) < \delta(u, z)$, we have $\delta(u, x) < \delta(u, z)$ and $x \sqsubseteq^u z$, as desired. In the other case, we have $\delta(u, x) = \delta(u, y) = \delta(u, z)$. Setting $k := \delta(u, x)$ and $l := \delta(u, z)$, we distinguish three cases.

- ♦ If $k = l$, then $x(k) < y(k) < z(k)$ implies $x \sqsubseteq^u z$.
- ♦ If $k < l$, then $x(l) = y(l) < z(l)$ implies $x \sqsubseteq^u z$.
- ♦ If $k > l$, then $x(k) < y(k) = z(k)$ implies $x \sqsubseteq^u z$.

For linearity, suppose that $x \not\sqsubseteq^u y$, for $x \approx_* u \approx_* y$. Then $\delta(u, x) > \delta(u, y)$, or $\delta(u, x) = \delta(u, y)$ and $x(\delta(u, x)) > y(\delta(u, x))$. In both cases, it follows that $y \sqsubseteq^u x$.

To prove well-foundedness, suppose that $[u]_*$ contains an infinite decreasing sequence $x_0 \supseteq^u x_1 \supseteq^u \dots$. Then

$$\delta(u, x_0) \geq \delta(u, x_1) \geq \dots$$

Consequently, there is some index n such that

$$\delta(u, x_i) = \delta(u, x_n), \quad \text{for all } i \geq n.$$

Then the words x_n, x_{n+1}, \dots differ only in the first $k := \delta(u, x_n)$ letters. Consequently, there can be at most $|\Sigma|^k$ indices $i \geq n$ with $x_i \sqsubseteq^u x_{i+1}$ and the sequence x_0, x_1, \dots is ultimately constant. \square

The following remark is a first indication of why the relation \approx_* is useful when dealing with cardinality questions.

Lemma 6.3. *A regular set $L \subseteq \Sigma^\omega$ is infinite if, and only if, $L \cap [w]_*$ is infinite, for some \approx_* -class $[w]_*$.*

Proof. (\Leftarrow) is trivial. For (\Rightarrow), let $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ be a homomorphism into a finite ω -semigroup that recognises L . Then there exists elements $a_0, \dots, a_{n-1}, e_0, \dots, e_{n-1} \in S$ such that $\eta[L] = \{a_0 e_0^\omega, \dots, a_{n-1} e_{n-1}^\omega\}$. This implies that

$$L = \bigcup_{i < n} U_i V_i^\omega, \quad \text{where } U_i := \eta^{-1}(a_i) \text{ and } V_i := \eta^{-1}(e_i).$$

Furthermore, as L is infinite, one of the languages $U_i V_i^\omega$ is also infinite. It is consequently sufficient to prove the claim for languages of the form $L = UV^\omega$, where $U, V \subseteq \Sigma^*$ are non-empty and regular.

First, consider the case where $V^\omega = \{v^\omega\}$, for some word v . Since

$$uv^\omega \approx_* u'v^\omega, \quad \text{for all } u, u' \in U \text{ with } |u| \equiv |u'| \pmod{|v|},$$

it follows that UV^ω is contained in the union of at most $|v|$ \approx_* -classes. One of these must have an infinite intersection with UV^ω .

It remains to consider the case where there are words $v, w \in V$ with $v^\omega \neq w^\omega$. Setting $U' := Uw^*$ we have

$$U'v^\omega \subseteq UV^*V^\omega = UV^\omega.$$

If we can show that $U'v^\omega$ is infinite, the claim follows therefore by the first case. For a contradiction, suppose that $U'v^\omega$ is finite. Then w^*v^ω is also finite and there are $i < j$ with $w^i v^\omega = w^j v^\omega$. Consequently, $w^k = v^n$, for some $k, n > 0$, which implies that $w^\omega = v^\omega$. A contradiction. \square

Equivalence Relations

The fact that ω -automatic structures do not necessarily have injective presentations makes counting elements of such structures more complicated as we have to count equivalence classes instead of single elements. We therefore start with a partition theorem for structures with equivalence relations that allows us to count equivalence classes.

Definition 6.4. Let $E, \sim \subseteq A \times A$ be equivalence relations.

- (a) We denote the E -class of a by $[a]_E$.
- (b) We say that a set $S \subseteq A$ is E -covered by a set $W \subseteq A/\sim$ of \sim -classes if, for every $s \in S$, there is some $[w]_\sim \in W$ such that

$$[s]_E \cap [w]_\sim \neq \emptyset.$$

- (c) If the relation E is clear from the context, we simply speak of S being covered by W . In particular, if \mathfrak{A} is an ω -automatic structure, we say that a set $S \subseteq A$ is covered by W if it is $L_=_$ -covered by it. \downarrow

Definition 6.5. Let $\eta : \Sigma^\omega \rightarrow \mathfrak{S}$ be a homomorphism into a finite ω -semigroup and let $w \in \Sigma^\omega$ be a word. We say that an infinite subset $H \subseteq \omega$ induces an η -homogeneous factorisation of w of type $\langle b, e \rangle$ if

$$\begin{aligned} \eta(w[o, i]) &= b, \quad \text{where } i := \min H, \\ \text{and } \eta(w[i, j]) &= e, \quad \text{for all } i < j \text{ in } H. \end{aligned} \quad \downarrow$$

A homogeneous factorisation is thus similar to a Ramsey factorisation (see Section III.3), except that we do not require that $\eta(w[o, i]) = b$, for all $i \in H$. In particular, Lemma III.3.4 (a) implies that every word $w \in \Sigma^*$ has an η -factorisation of some type $\langle b, e \rangle$. Furthermore, it follows immediately from the definition that e is idempotent.

The combinatorial core of our argument is contained in the following rather technical lemma which generalises the construction in the proof of Lemma V.6.21. Intuitively, it states that sets not E -covered by a few \approx_* -classes contain two words that share a common, highly homogeneous factorisation.

Lemma 6.6. Let $K \subseteq (\Sigma \times \Gamma)^\omega$ and $E \subseteq (\Sigma \times \Sigma)^\omega$ be regular languages recognised by the homomorphisms $\eta_K : (\Sigma \times \Gamma)^\omega \rightarrow \mathfrak{T}$ and $\eta_E : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{S}$, respectively, and assume that E is an equivalence relation on Σ^ω . Set

$$K(\gamma) := \{ \alpha \in \Sigma^\omega \mid \alpha \otimes \gamma \in K \}, \quad \text{for } \gamma \in \Gamma^\omega.$$

There exists a constant k with the following property.

For every word $\gamma \in \Gamma^\omega$ such that $K(\gamma)$ is not E -covered by k \approx_* -classes, there exist words $\alpha_o, \alpha_1 \in K(\gamma)$, an infinite set $H \subseteq \omega$, and elements $c, f \in T$ and $b, e_<, e_=:, e_> \in S$ with the following properties.

- (I) $\alpha_o, \alpha_1 \in K(\gamma)$
- (II) $\alpha_o \otimes \alpha_1 \notin E$
- (III) $\alpha_o[o, i] = \alpha_1[o, i]$, where $i := \min H$.
- (IV) $\alpha_o[i, j] \neq \alpha_1[i, j]$, for all $i < j$ in H .
- (V) H induces η_K -homogeneous factorisations of $\alpha_o \otimes \gamma$ and $\alpha_1 \otimes \gamma$ of the same type $\langle c, f \rangle$.
- (VI) H induces η_E -homogeneous factorisations of $\alpha_i \otimes \alpha_j$, for all $i, j < 2$, whose types are

$$\begin{aligned} \langle b, e_=\rangle, \quad \text{for } \alpha_o \otimes \alpha_o, \quad & \langle b, e_<\rangle, \quad \text{for } \alpha_o \otimes \alpha_1, \\ \langle b, e_=\rangle, \quad \text{for } \alpha_1 \otimes \alpha_1, \quad & \langle b, e_>\rangle, \quad \text{for } \alpha_1 \otimes \alpha_o. \end{aligned}$$

- (VII) $e_<, e_=:, e_>$ are idempotent and satisfy the equations

$$e_<e_=: = e_<, \quad e_>e_=: = e_>, \quad be_=: = b.$$

Proof. Set $k := |S| \cdot |T|$ and suppose that there are words $\beta_0, \dots, \beta_k \in K(\gamma)$ such that

$$[\beta_i]_E \cap [\beta_j]_* = \emptyset, \quad \text{for all } i \neq j.$$

We label each pair $m < n$ of natural numbers by the semigroup elements

$$\langle \eta_K((\beta_i \otimes \gamma)[m, n]) \rangle_{i \leq k} \quad \text{and} \quad \langle \eta_E((\beta_i \otimes \beta_j)[m, n]) \rangle_{i \leq j \leq k}$$

By the Theorem of Ramsey, there exists an infinite set $I \subseteq \omega$ and elements $\tilde{f} \in T^{k+1}$ and $\tilde{t} \in S^{k(k+1)/2}$ such that

$$\eta_K((\beta_i \otimes \gamma)[m, n]) = f_i \quad \text{and} \quad \eta_E((\beta_i \otimes \beta_j)[m, n]) = t_{ij},$$

for all $m < n$ in I . Since there are only k different possible pairs $\langle f_i, t_{ii} \rangle$, we can find $i \neq j$ with $f_i = f_j$ and $t_{ii} = t_{jj}$. W.l.o.g. we may assume that $i = 0$ and $j = 1$. Then I induces


- an η_K -homogeneous factorisation of $\beta_i \otimes \gamma$ of type $\langle c_i, f_o \rangle$, for each $i < 2$, and
- an η_E -homogeneous factorisation of $\beta_i \otimes \beta_j$ of type $\langle s_{ij}, t_{ij} \rangle$, for all $i, j < 2$.

Since $\beta_0 \not\approx_* \beta_1$, there exists an infinite subset $I_o \subseteq I$ such that

$$\beta_0[i, j] \neq \beta_1[i, j], \quad \text{for all } i < j \text{ in } I_o.$$

Let $g_0 < g_1 < \dots$ be an enumeration of I_o . We define α_o and α_1 by

$$\begin{aligned} \alpha_o &:= \beta_1[0, g_1) \beta_o[g_1, \infty), \\ \alpha_1[0, g_1) &:= \beta_1[0, g_1), \\ \alpha_1[g_{2i+1}, g_{2i+3}) &:= \beta_1[g_{2i+1}, g_{2i+2}) \beta_o[g_{2i+2}, g_{2i+3}). \end{aligned}$$



Fix a number m such that $a^m = a^\pi$, for all elements a of \mathfrak{S} or \mathfrak{T} , and set

$$h_i := g_{2mi+1} \quad \text{and} \quad H := \{h_i \mid i < \omega\} \subseteq I_o.$$

Then (III) and (IV) follow by choice of I_o . To prove (i) and (v), it is sufficient to note that

$$\begin{aligned} \eta_K(\alpha_o \otimes \gamma) &= c_1 f_1 \cdot f_o^\omega = c_1 f_o^\omega = \eta_K(\beta_1 \otimes \gamma), \\ \eta_K(\alpha_1 \otimes \gamma) &= c_1 f_1 \cdot (f_1 f_o)^\omega = c_1 f_o^\omega = \eta_K(\beta_1 \otimes \gamma). \end{aligned}$$

In particular, the factorisation induced by H is η_K -homogeneous of type $\langle c_1, f_o \rangle$. Concerning (II), we suppose for a contradiction that $\alpha_o \otimes \alpha_1 \in E$. Then

$$\begin{aligned} \eta_E(\alpha_o \otimes \alpha_1) &= s_{11} t_{11} \cdot (t_{01} t_{00})^\omega \\ &= s_{11} t_{11} \cdot (t_{01} t_{11})^\omega \\ &= s_{11} \cdot (t_{11} t_{01})^\omega = \eta_E(\alpha_1 \otimes \beta_1) \end{aligned}$$

and $\alpha_1 \otimes \beta_1 \in E$ implies by transitivity that $\alpha_o \otimes \beta_1 \in E$. Thus, $\beta_1 \in \alpha_o \approx_* \beta_o$. A contradiction.

It remains to prove (VI) and (VII). Set

$$\begin{aligned} r_+ &:= t_{11}, & r_- &:= t_{01} t_{00}, & r_< &:= t_{10} t_{00}, \\ e_+ &:= r_+, & e_- &:= r_+^m, & e_< &:= r_<^m, \quad \text{and} \quad b := s_{11} t_{11}. \end{aligned}$$

Then we have

$$\begin{aligned} r_+ r_- &= t_{11} t_{11} = t_{11} = r_+, \\ r_- r_+ &= t_{01} t_{00} t_{11} = t_{01} t_{00} t_{00} = t_{01} t_{00} = r_-, \\ r_< r_+ &= t_{10} t_{00} t_{11} = t_{10} t_{00} t_{00} = t_{10} t_{00} = r_<, \end{aligned}$$

which, by choice of m , implies that

$$\begin{aligned} e_=_ \cdot e_=_ &= r_=_ \cdot r_=_ = r_=_ = e_=_ , \\ e_> \cdot e_> &= r_>^m \cdot r_>^m = r_>^m = e_> , \\ e_< \cdot e_< &= r_<^m \cdot r_<^m = r_<^m = e_< , \\ e_> \cdot e_=_ &= r_>^m \cdot r_=_^m = r_>^m = e_> , \\ e_< \cdot e_=_ &= r_<^m \cdot r_=_^m = r_<^m = e_< , \\ b \cdot e_=_ &= s_{\text{II}} t_{\text{II}} \cdot t_{\text{II}} = s_{\text{II}} t_{\text{II}} = b . \end{aligned}$$

This establishes (vii). For (vi), note that

$$\begin{aligned} \eta_E((\alpha_i \otimes \alpha_j)[o, h_o]) &= s_{\text{II}} t_{\text{II}} = b_=_ , \\ \eta_E((\alpha_o \otimes \alpha_o)[h_i, h_j]) &= r_=_^{2m(j-i)} = r_=_ = e_=_ , \\ \eta_E((\alpha_o \otimes \alpha_i)[h_i, h_j]) &= r_<^{m(j-i)} = r_<^m = e_< , \\ \eta_E((\alpha_i \otimes \alpha_o)[h_i, h_j]) &= r_>^{m(j-i)} = r_>^m = e_> , \\ \eta_E((\alpha_i \otimes \alpha_i)[h_i, h_j]) &= r_=_^{2m(j-i)} = r_=_ = e_=_ . \end{aligned} \quad \square$$

The main result of this section is the following characterisation of regular languages that intersect at most countably many \approx_* -classes.

Theorem 6.7. *Let $K \subseteq (\Sigma \times \Gamma)^\omega$ and $E \subseteq (\Sigma \times \Sigma)^\omega$ be regular languages such that E is an equivalence relation on Σ^ω , and set*

$$K(\gamma) := \{ \alpha \in \Sigma^\omega \mid \alpha \otimes \gamma \in K \}, \quad \text{for } \gamma \in \Gamma^\omega .$$

We can compute a constant k such that, for every $\gamma \in \Gamma^\omega$, the following statements are equivalent.

- (1) $|K(\gamma)/E| \leq \aleph_o$
- (2) $|K(\gamma)/E| < 2^{\aleph_o}$
- (3) $K(\gamma)$ is E -covered by at most k \approx_* -classes.

Proof. Fix homomorphisms $\eta_E : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{S}$ and $\eta_K : (\Sigma \times \Gamma)^\omega \rightarrow \mathfrak{T}$ recognising E and K , and let k be the constant from Lemma 6.6.

(1) \Rightarrow (2) is trivial.

(3) \Rightarrow (1) Note that every \approx_* -class is countable. Hence $K(\gamma)$ can contain only countably many E -classes.

(2) \Rightarrow (3) If $K(\gamma)$ is not covered by k \approx_* -classes, we can find words α_o, α_i , and a set $H \subseteq \omega$ as in Lemma 6.6.

Let $h_o < h_1 < \dots$ be an enumeration of H . For every word $\sigma \in [2]^\omega$, we denote by $\hat{\alpha}_\sigma$ the word with

$$\hat{\alpha}_\sigma[o, h_o) := \alpha_i[o, h_o) \quad \text{and} \quad \hat{\alpha}_\sigma[h_i, h_{i+1}) := \alpha_{\sigma(i)}[h_i, h_{i+1}) .$$

Then

$$\eta_K(\hat{\alpha}_\sigma \otimes w) = c f^\omega = \eta_K(\alpha_i \otimes w) \quad \text{implies} \quad \hat{\alpha}_\sigma \otimes w \in K ,$$

for all σ . Furthermore, $\alpha_o[h_i, h_{i+1}) \neq \alpha_i[h_i, h_{i+1})$ means that

$$\sigma \not\approx_* \tau \quad \text{implies} \quad \hat{\alpha}_\sigma \not\approx_* \hat{\alpha}_\tau .$$

Hence, the set $\{ \hat{\alpha}_\sigma \mid \sigma \in [2]^\omega \}$ has size 2^{\aleph_o} . To conclude the proof it is therefore sufficient to show that

$$\sigma \not\approx_* \tau \quad \text{implies} \quad \hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau \notin E .$$

Let $\beta_o := \hat{\alpha}_{(o1)^\omega}$ and $\beta_i := \hat{\alpha}_{(io)^\omega}$. It is sufficient to establish the following two claims.

- (a) $\sigma \not\approx_* \tau$ implies $\eta_E(\hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau) = \eta_E(\beta_i \otimes \beta_j)$ for some $i \neq j$.
- (b) $\beta_o \otimes \beta_i \notin E$
- (a) Set

$$m := \min \{ i \mid \sigma(i) \neq \tau(i) \} , \quad c_i := \begin{cases} e_< & \text{if } \tau(i) > \sigma(i) , \\ e_> & \text{if } \tau(i) < \sigma(i) , \\ e_=_ & \text{if } \tau(i) = \sigma(i) . \end{cases}$$

By symmetry, we may assume that $\sigma(m) = 1$. Since $e_<$ and $e_>$ are idempotent and the elements $b, e_<, e_>$ absorb $e_=_$ from the right, it follows that

$$\eta_E(\hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau) = b_=_ \cdot \prod_{i < \omega} c_i = b_=(e_< e_>)^\omega = \eta_E(\beta_o \otimes \beta_1) .$$

(b) Set $\gamma_{ijl} := \hat{\alpha}_{(ijl)}^\omega$. By idempotency of $e_<$ and $e_>$ it follows that

$$\begin{aligned}\eta_E(\gamma_{oio} \otimes \gamma_{ioi}) &= b(e_<e_>e_<)^\omega = b(e_<e_>)^\omega = \eta_E(\beta_o \otimes \beta_i), \\ \eta_E(\gamma_{ioi} \otimes \gamma_{iio}) &= b(e_=e_<e_>)^\omega = b(e_<e_>)^\omega = \eta_E(\beta_o \otimes \beta_i), \\ \eta_E(\gamma_{oio} \otimes \gamma_{iio}) &= b(e_<e_=e_>)^\omega = be_<^\omega = \eta_E(\alpha_o \otimes \alpha_i).\end{aligned}$$

For a contradiction, suppose that $\beta_o \otimes \beta_i \in E$. Then the first two of the above equations imply that

$$\gamma_{oio} \otimes \gamma_{ioi} \in E \quad \text{and} \quad \gamma_{ioi} \otimes \gamma_{iio} \in E.$$

By transitivity of E and the third equation, we therefore have

$$\gamma_{oio} \otimes \gamma_{iio} \in E \quad \text{and} \quad \alpha_o \otimes \alpha_i \in E.$$

A contradiction. \square

Corollary 6.8. *Let \mathfrak{A} be an ω -automatic structure and $\varphi(x; \bar{z})$ a first-order formula. Given a presentation of \mathfrak{A} we can compute a constant k such that, for all tuples \bar{c} in A the following two statements are equivalent.*

- (1) *The set defined by $\varphi(x; \bar{c})$ is non-empty and countable.*
- (2) *There exist k elements $a_o, \dots, a_{k-1} \in \varphi(x; \bar{c})^{\mathfrak{A}}$ such that $\varphi(x; \bar{c})^{\mathfrak{A}}$ is covered by $[a_o]_*, \dots, [a_{k-1}]_*$.*

Corollary 6.9. *A regular language $K \subseteq \Sigma^\omega$ is countable if, and only if, it is a finite union of languages of the form Uw^ω with regular $U \subseteq \Sigma^*$ and $w \in \Sigma^+$.*

Proof. (\Leftarrow) is trivial as every set of the form Uw^ω is countable.

(\Rightarrow) We consider the equivalence relation

$$u \sim v \quad : \text{iff} \quad u \in K \Leftrightarrow v \in K.$$

Clearly, \sim is regular. Fix $u \in K$. Then $[u]_\sim = K$ is countable and it follows by Theorem 6.7 that $[u]_\sim / \approx_*$ is finite. Consequently, there are finitely many words w_o, \dots, w_{n-1} such that every $u \in K$ is \approx_* -equivalent to some w_i . Note that every regular language of ω -words can be written in the form $K = \bigcup_{i < m} U_i V_i^\omega$, for regular languages $U_i, V_i \subseteq \Sigma^+$. It follows that $V_i^\omega = \{v_i^\omega\}$, for some $v_i \in \Sigma^+$ such that v_i^ω is \approx_* -equivalent to one of the w_j . \square

Corollary 6.10. *Let \sim be a regular equivalence relation on Σ^ω with countably many classes. Then there exists a regular language $U \subseteq \Sigma^\omega$ containing exactly one representative of every \sim -class.*

Proof. Applying Theorem 6.7 to the language $K = \Sigma^\omega$ we obtain words $u_o, \dots, u_{k-1} \in \Sigma^\omega$ such that every \sim -class intersects some \approx_* -class $[u_i]_*$. These words satisfy the formula

$$\varphi(\bar{z}) := \forall x \exists x' \bigvee_{i < k} [x \sim x' \wedge x' \approx_* z_i].$$

Hence, φ defines a non-empty regular language over Σ^k . This language must contain some ultimately periodic word \bar{u} . Then the set $P := \bigcup_{i < k} [u_i]_*$ is of the form

$$P := \bigcup_{i < k} \Sigma^* p_i^\omega, \quad \text{for some } p_o, \dots, p_{k-1} \in \Sigma^+.$$

Using the well-order \sqsubseteq^u from Lemma 6.2, we can now define our set of representatives by

$$\begin{aligned}\varphi(x) &:= \bigvee_{i < k} [x \in P_i \setminus P_{<i} \\ &\quad \wedge \forall x' [x \sim x' \wedge x' \in P \rightarrow x' \notin P_{<i} \wedge x \sqsubseteq^{p_i^\omega} x']],\end{aligned}$$

where

$$P_i := \Sigma^* p_i^\omega \quad \text{and} \quad P_{<i} := P_o \cup \dots \cup P_{i-1}.$$

(Note that P_i and p_i^ω are definable.) \square

Linear Orders

In the presense of a linear order, we can strengthen Theorem 6.7 as follows.

Theorem 6.11. *Let $K \subseteq \Sigma^\omega$, $E \subseteq (\Sigma \times \Sigma)^\omega$, and $R \subseteq (\Sigma \times \Sigma)^\omega$ be regular languages such that K is uncountable, E is an equivalence relation on Σ^ω , and R is a strict linear order on K/E . There exists a regular set $M \subseteq K$ such that*

$$\langle M, R \rangle \cong \langle [2]^\omega, <_{\text{lex}} \rangle \quad \text{and} \quad \alpha \otimes \beta \notin E, \quad \text{for all } \alpha \neq \beta \text{ in } M.$$

Proof. We fix homomorphisms

$$\kappa : \Sigma^\omega \rightarrow \mathfrak{S}, \quad \varepsilon : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{T}, \quad \text{and} \quad \rho : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{U}$$

into finite ω -semigroups recognising, respectively, K , E , and R . Then

$$\eta := \langle \varepsilon, \rho \rangle : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{T} \times \mathfrak{U}$$

also recognises E . Since K is uncountable, it is not covered by countably many \approx_* -classes. Therefore, we can use Lemma 6.6 to find words α_o and α_i and an infinite set $H \subseteq \omega$ with the properties from the lemma.

Let $h_o < h_i < \dots$ be an enumeration of H and set

$$u := \alpha_o[o, h_o], \quad v_o := \alpha_o[h_o, h_i], \quad \text{and} \quad v_i := \alpha_i[h_o, h_i].$$

We claim that $M := u\{v_o, v_i\}^\omega$ is the desired set. For a word $\sigma \in [2]^\omega$, set

$$\hat{\alpha}_\sigma := uv_{\sigma(0)}v_{\sigma(1)}v_{\sigma(2)}\cdots \in \Sigma^\omega.$$

Then

$$\kappa(\hat{\alpha}_\sigma) = cf^\omega = \kappa(\alpha_o) \quad \text{implies} \quad \hat{\alpha}_\sigma \in K, \quad \text{for all } \sigma.$$

We have therefore found a regular set $M \subseteq K$. It remains to check that it has the desired properties.

First, set $\beta_{ij} := u(v_i v_j)^\omega$ and $\gamma_{ijl} := u(v_i v_j v_l)^\omega$. By idempotency of $e_<$ and $e_>$ it follows that

$$\begin{aligned} \eta(\gamma_{oio} \otimes \gamma_{ioi}) &= b(e_< e_> e_<)^\omega = b(e_< e_>)^\omega = \eta(\beta_{oi} \otimes \beta_{io}), \\ \eta(\gamma_{ioi} \otimes \gamma_{iio}) &= b(e_< e_< e_>)^\omega = b(e_< e_>)^\omega = \eta(\beta_{oi} \otimes \beta_{io}), \\ \eta(\gamma_{oio} \otimes \gamma_{iio}) &= b(e_< e_< e_<)^\omega = b e_<^\omega = \eta(\alpha_o \otimes \alpha_i). \end{aligned}$$

We claim that $\beta_{oi} \otimes \beta_{io} \in E$. For a contradiction, suppose otherwise. Then the first two of the above equations imply that

$$\gamma_{oio} \otimes \gamma_{ioi} \in E \quad \text{and} \quad \gamma_{ioi} \otimes \gamma_{iio} \in E.$$

By transitivity of E and the third equation, we therefore have

$$\gamma_{oio} \otimes \gamma_{iio} \in E \quad \text{and} \quad \alpha_o \otimes \alpha_i \in E.$$

A contradiction.

Having shown that $\beta_{oi} \otimes \beta_{io} \notin E$, it follows that

$$\beta_{oi} \otimes \beta_{io} \in R \quad \text{or} \quad \beta_{io} \otimes \beta_{oi} \in R.$$

By symmetry, we may assume the former. To conclude the proof, it is sufficient to show that

$$\hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau \in R \quad \text{iff} \quad \sigma <_{\text{lex}} \tau.$$

(If $\beta_{io} \otimes \beta_{oi} \in R$, the ordering would be reversed.) Fix $\sigma \neq \tau$. Then we can write $\eta(\hat{\alpha}_\sigma, \hat{\alpha}_\tau)$ as the product of some sequence in $b\{e_<, e_< e_<, e_>\}^\omega$. By the idempotency and absorption relations between the elements $b, e_<, e_< e_<, e_>$, we can transform every such product into one of the following forms:

$$\begin{aligned} a_0 &:= b(e_< e_>)^\omega & a_5 &:= b(e_> e_<)^\omega \\ a_1 &:= b(e_< e_>)^\omega e_<^\omega & a_6 &:= b(e_> e_<)^\omega e_>^\omega \\ a_2 &:= b(e_< e_>)^\omega e_<^\omega & a_7 &:= b(e_> e_<)^\omega e_<^\omega \\ a_3 &:= b(e_< e_>)^\omega e_< e_>^\omega & a_8 &:= b(e_> e_<)^\omega e_> e_<^\omega \\ a_4 &:= b(e_< e_>)^\omega e_< e_<^\omega & a_9 &:= b(e_> e_<)^\omega e_> e_<^\omega. \end{aligned}$$

Furthermore, it is one of the values of the left side if, and only if, we have $\sigma(i) < \tau(i)$, where i is the least index with $\sigma(i) \neq \tau(i)$. Consequently, we have

$$\sigma <_{\text{lex}} \tau \quad \text{iff} \quad \eta(\hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau) = a_i, \quad \text{for some } i < 5.$$

As R is anti-symmetric, it is therefore sufficient to find, for every $i < 5$, words $x_i \otimes y_i \in R$ with $\eta(x_i \otimes y_i) = a_i$. We already know that

$$\beta_{oi} \otimes \beta_{io} \in R \quad \text{and} \quad \eta(\beta_{oi} \otimes \beta_{io}) = b(e_< e_>)^\omega = a_0.$$

Since η recognises R , this implies that

$$\eta(x \otimes y) = a_o \quad \text{implies} \quad x \otimes y \in R.$$

For the remaining cases, it is therefore sufficient to find words x_i, y_i, z_i such that

$$\eta(x_i \otimes y_i) = a_o = \eta(y_i \otimes z_i) \quad \text{and} \quad \eta(x_i \otimes z_i) = a_i.$$

By transitivity of R , it then follows that $x_i \otimes z_i \in R$.

($i = 1$) Setting

$$\begin{aligned} x &:= u(v_o v_i)^n (v_o v_o v_i)^\omega, \\ y &:= u(v_i v_o)^n (v_o v_i v_o)^\omega, \\ z &:= u(v_i v_o)^n (v_i v_o v_i)^\omega, \end{aligned}$$

we have

$$\begin{aligned} \eta(x \otimes y) &= b(e_{<} e_{>})^n (e_{=} e_{<} e_{>})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(y \otimes z) &= b(e_{=} e_{=})^n (e_{<} e_{>} e_{<})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(x \otimes z) &= b(e_{<} e_{>})^n (e_{<} e_{=} e_{=})^\omega = b(e_{<} e_{>})^n e_{<}^\omega = a_1. \end{aligned}$$

($i = 2$) Setting

$$\begin{aligned} x &:= u(v_o v_i)^{n+1} (v_o v_i)^\omega, \\ y &:= u(v_o v_i)^{n+1} (v_i v_o)^\omega, \\ z &:= u(v_i v_o)^{n+1} (v_o v_i)^\omega, \end{aligned}$$

we have

$$\begin{aligned} \eta(x \otimes y) &= b(e_{=} e_{=})^{n+1} (e_{<} e_{>})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(y \otimes z) &= b(e_{<} e_{>})^{n+1} (e_{>} e_{<})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(x \otimes z) &= b(e_{<} e_{>})^{n+1} (e_{=} e_{=})^\omega = b(e_{<} e_{>})^{n+1} e_{=}^\omega = a_2. \end{aligned}$$

($i = 3$) Setting

$$\begin{aligned} x &:= u(v_o v_i)^n v_o (v_i v_o v_i)^\omega, \\ y &:= u(v_o v_i)^n v_o (v_i v_i v_o)^\omega, \\ z &:= u(v_i v_o)^n v_i (v_o v_o v_i)^\omega, \end{aligned}$$

we have

$$\begin{aligned} \eta(x \otimes y) &= b(e_{=} e_{=})^n e_{=} (e_{=} e_{<} e_{>})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(y \otimes z) &= b(e_{<} e_{>})^n e_{<} (e_{>} e_{>} e_{<})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(x \otimes z) &= b(e_{<} e_{>})^n e_{<} (e_{>} e_{=} e_{=})^\omega = b(e_{<} e_{>})^n e_{<} e_{>}^\omega = a_3. \end{aligned}$$

($i = 4$) Setting

$$\begin{aligned} x &:= u(v_o v_i)^n v_o (v_o v_i)^\omega, \\ y &:= u(v_o v_i)^n v_o (v_i v_o)^\omega, \\ z &:= u(v_i v_o)^n v_i (v_o v_i)^\omega, \end{aligned}$$

we have

$$\begin{aligned} \eta(x \otimes y) &= b(e_{=} e_{=})^n e_{=} (e_{<} e_{>})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(y \otimes z) &= b(e_{<} e_{>})^n e_{<} (e_{>} e_{<})^\omega = b(e_{<} e_{>})^\omega = a_o, \\ \eta(x \otimes z) &= b(e_{<} e_{>})^n e_{<} (e_{=} e_{=})^\omega = b(e_{<} e_{>})^n e_{<} e_{=}^\omega = a_4. \quad \square \end{aligned}$$

Functions

Finally, we derive a partition theorem based on functions. We start with a simple lemma.

Lemma 6.12. *Let $R \subseteq (\Sigma \times \Gamma)^\omega$ be regular and set*

$$R(\gamma) := \{ \alpha \in \Gamma^\omega \mid \gamma \otimes \alpha \in R \}, \quad \text{for } \gamma \in \Sigma^\omega.$$

There exist constants c, d such that, for every ultimately periodic word $\gamma = wv^\omega \in \Sigma^\omega$ with period $p = |v|$ and $R(\gamma) \neq \emptyset$, there is some ultimately periodic word $\alpha = zu^\omega \in R(\gamma)$ with $|z| \leq |w| + cp$ and $|u| = ip$, for some $i \leq d$.

Proof. Fix a homomorphism $\eta : (\Sigma \times \Gamma)^\omega \rightarrow \mathfrak{S}$ into a finite ω -semigroup recognising R , let $\gamma = wv^\omega \in \Sigma^\omega$ be an ultimately periodic word with period $p := |v|$, and suppose that $\beta \in R(\gamma)$. Set $n := |w|$. By the Theorem of Ramsey, there exists an infinite set $H \subseteq \omega$ and an element $a \in S$ such that

$$\eta((\gamma \otimes \beta)[n + ip, n + jp]) = a, \quad \text{for all } i < j \text{ in } H.$$

Let u be a word of minimal length kp such that

$$\eta(v^k \otimes u) = a.$$

If $k > |S|$, there would be two prefixes of u with the same image under η and we could shorten the word u . Hence, minimality implies that $k \leq |S|$. Set

$$z' := \beta[o, h) \quad \text{where } h := \min H,$$

and let z be the shortest word of length $|w| + ip$, for some i , such that

$$\eta(\gamma[o, |w| + ip) \otimes z) = \eta(\gamma[o, h) \otimes z').$$

By the same pumping argument as above, it follows that $i \leq |S|$. It follows that $\alpha := zu^\omega$ has period kp and

$$\eta(\gamma \otimes \alpha) = \eta(\gamma \otimes z'u^\omega) = \eta((\gamma \otimes \beta)[o, h)) \cdot a^\omega = \eta(\gamma \otimes \beta),$$

implies that $\alpha \in R(\gamma)$. \square

Theorem 6.13. *Let \mathfrak{A} be an uncountable ω -automatic structure. There exists a number $k > 0$ such that, for every FO-definable function $f : A^{k+1} \rightarrow A$, there exists uncountable regular sets $U \subseteq A^k$ and $V \subseteq A$ such that*

$$f(u, v) = f(u', v), \quad \text{for all } u, u' \in U, v \in V.$$

Proof. We fix a presentation $\langle L_\delta, L_-, (L_R)_R \rangle$ of \mathfrak{A} and let $\kappa : \Sigma^\omega \rightarrow \mathfrak{S}_\delta$ and $\eta : (\Sigma \times \Sigma)^\omega \rightarrow \mathfrak{S}_-$ be homomorphisms into finite ω -semigroups recognising, respectively, L_δ and L_- . Since A is uncountable, it is not covered by countably many \approx_* -classes. Therefore, we can use Lemma 6.6 to find

words α_o and α_i and an infinite set $H \subseteq \omega$ with the properties from the lemma. We proceed in several steps.

(I) Let $h_o < h_i < \dots$ be an enumeration of H and set

$$u := \alpha_o[o, h_o), \quad v_o := \alpha_o[h_o, h_i), \quad \text{and} \quad v_i := \alpha_i[h_o, h_i).$$

Note that $|v_o| = |v_i|$. Let $M := u\{v_o, v_i\}^\omega$. As in the proof of Theorem 6.7 it follows that

$$\beta \approx_* \beta' \quad \text{implies} \quad \beta \otimes \beta' \notin L_-, \quad \text{for } \beta, \beta' \in M.$$

Set $k := |v_o| + 1$ and suppose that $f : A^{k+1} \rightarrow A$ is FO-definable. Fix a homomorphism $\varphi : (\Sigma^{k+2})^\omega \rightarrow \mathfrak{S}_f$ recognising f and set $m := |S_f|$. We consider the sets

$$\begin{aligned} U_o &:= \{uv^\omega \mid v \in (v_o + v_i)^m x_o\}, & \text{where } x_o &:= v_o^{m(2k+1)}, \\ V_o &:= \{uv^\omega \mid v \in x_i(v_o + v_i)^{m(2k+1)}\}, & \text{where } x_i &:= v_o^m. \end{aligned}$$

Then $U_o, V_o \subseteq M$ and $\beta \approx_* \beta'$, for all $\beta \neq \beta'$ in $U_o \cup V_o$.

Given $\bar{\alpha} \in U_o^k$ and $\beta \in V_o$, let $\gamma := f(\bar{\alpha}, \beta)$. Since $\bar{\alpha} \otimes \beta$ is ultimately periodic with a period dividing $2m(k+1)$, it follows by Lemma 6.12 that γ is ultimately periodic with period at most $r := d \cdot 2m(k+1)$ and a prefix of length at most $s := |u| + c \cdot |v_o|$, for some constants c, d (independent of $\bar{\alpha}$ and β). Set $h_i := s + ri$.

For an infinite word γ and $i < \omega$, we set

$$\begin{aligned} p_i(\gamma) &:= \gamma[h_i, h_i + m|v_o|), \\ q_i(\gamma) &:= \gamma[h_i + m|v_o|, h_i + 2m(k+1)|v_o|). \end{aligned}$$

Thus,

$$\begin{aligned} \bar{\alpha} &= u \cdot p_o(\bar{\alpha}) \cdot x_o \cdot p_1(\bar{\alpha}) \cdot x_o \cdot p_2(\bar{\alpha}) \dots = u \cdot (p_o(\bar{\alpha}) \cdot x_o)^\omega \\ \beta &= u \cdot x_i \cdot q_o(\beta) \cdot x_i \cdot q_1(\beta) \cdot x_i \cdot q_2(\beta) \dots = u \cdot (x_i \cdot q_o(\beta))^\omega. \end{aligned}$$

(II) We fix words $\bar{\alpha}, \bar{\alpha}' \in U_o^k$ and $\beta, \beta' \in V_o$ as follows. By definition,

$$|U_o^k| = |p_o[U_o^k]| = |\Sigma|^{mk} \quad \text{and} \quad |V_o| = |q_o[V_o]| = |\Sigma|^{m(2k+1)}.$$

By choice of k , it follows that

$$|U_o^k| = |\Sigma|^{mk} \geq |\Sigma|^m \cdot |\Sigma|^{m \cdot |v_o|} > m \cdot |\Sigma|^{m \cdot |v_o|} = |S_f| \cdot |\Sigma|^{m \cdot |v_o|}.$$

Consequently, for every $\beta \in V_o$, there exist two tuples $\tilde{\alpha}_\beta, \tilde{\alpha}'_\beta \in U_o^k$ such that, for some/all $i < \omega$,

$$p_i(f(\tilde{\alpha}_\beta, \beta)) = p_i(f(\tilde{\alpha}'_\beta, \beta))$$

and $\varphi(p_i(\tilde{\alpha}_\beta) \otimes x_i \otimes p_i(f(\tilde{\alpha}_\beta, \beta))) = \varphi(p_i(\tilde{\alpha}'_\beta) \otimes x_i \otimes p_i(f(\tilde{\alpha}'_\beta, \beta)))$.

Then there is some pair $\tilde{\alpha}, \tilde{\alpha}'$ that is assigned to at least

$$\frac{|V_o|}{|U_o^k \times U_o^k|} = \frac{|\Sigma|^{m(2k+1)}}{|\Sigma|^{mk} \cdot |\Sigma|^{mk}} = |\Sigma|^m > m$$

words $\beta \in V_o$. Consequently, we can find two words $\beta, \beta' \in V_o$ both associated with $\tilde{\alpha}, \tilde{\alpha}'$ such that

$$\varphi(x_o \otimes q_i(\beta) \otimes q_i(f(\tilde{\alpha}, \beta))) = \varphi(x_o \otimes q_i(\beta') \otimes q_i(f(\tilde{\alpha}, \beta')))$$

(III) For every word $\sigma \in [2]^\omega$, we define words $\hat{\alpha}_\sigma \in M^k$ and $\hat{\beta}_\sigma \in M$ by

$$\begin{aligned} \hat{\alpha}_\sigma[o, h_o] &:= \tilde{\alpha}[o, h_o], \\ \hat{\alpha}_\sigma[h_i, h_{i+1}] &:= \begin{cases} \tilde{\alpha}[h_i, h_{i+1}] & \text{if } \sigma(i) = o, \\ \tilde{\alpha}'[h_i, h_{i+1}] & \text{if } \sigma(i) = 1, \end{cases} \\ \hat{\beta}_\sigma[o, h_o] &:= \beta[o, h_o], \\ \hat{\beta}_\sigma[h_i, h_{i+1}] &:= \begin{cases} \beta[h_i, h_{i+1}] & \text{if } \sigma(i) = o, \\ \beta'[h_i, h_{i+1}] & \text{if } \sigma(i) = 1. \end{cases} \end{aligned}$$

Then $\tilde{\alpha}[h_i, h_{i+1}] \neq \tilde{\alpha}'[h_i, h_{i+1}]$ implies that

$$\sigma \neq \tau \Rightarrow \hat{\alpha}_\sigma \not\approx_* \hat{\alpha}_\tau \Rightarrow \hat{\alpha}_\sigma \otimes \hat{\alpha}_\tau \notin L_-^k.$$

Similarly, $\beta[h_i, h_{i+1}] \neq \beta'[h_i, h_{i+1}]$ implies that

$$\sigma \neq \tau \Rightarrow \hat{\beta}_\sigma \not\approx_* \hat{\beta}_\tau \Rightarrow \hat{\beta}_\sigma \otimes \hat{\beta}_\tau \notin L_-.$$

For $\rho \in [2]^\omega$, set $\gamma_\rho := f(\tilde{\alpha}, \beta_\rho)$. Then

$$\begin{aligned} & \varphi(\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho) \\ &= \varphi((\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho)[o, h_o]) \cdot \prod_{i < \omega} \varphi((\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho)[h_i, h_{i+1}]) \\ &= \varphi((\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho)[o, h_o]) \\ & \quad \cdot \left[\prod_{i < \omega} [\varphi(p_i(\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho)) \cdot \varphi(q_i(\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho))] \right]^d \\ &= \varphi((\tilde{\alpha}_\sigma \otimes \beta_\rho \otimes \gamma_\rho)[o, h_o]) \\ & \quad \cdot \prod_{i < \omega} [\varphi(p_i(\tilde{\alpha}_\sigma) \otimes x_i \otimes p_i(f(\tilde{\alpha}, \beta_\rho))) \\ & \quad \cdot \varphi(x_o \otimes q_i(\beta_\rho) \otimes q_i(f(\tilde{\alpha}, \beta_\rho)))]^d \\ &= \varphi((\tilde{\alpha} \otimes \beta_\rho \otimes f(\tilde{\alpha}, \beta_\rho))[o, h_o]) \\ & \quad \cdot \prod_{i < \omega} [\varphi(p_i(\tilde{\alpha}) \otimes x_i \otimes p_i(f(\tilde{\alpha}, \beta_\rho))) \\ & \quad \cdot \varphi(x_o \otimes q_i(\beta_\rho) \otimes q_i(f(\tilde{\alpha}, \beta_\rho)))]^d \\ &= \varphi(\tilde{\alpha} \otimes \beta_\rho \otimes f(\tilde{\alpha}, \beta_\rho)) \end{aligned}$$

implies that

$$f(\tilde{\alpha}_\sigma, \beta_\rho) = f(\tilde{\alpha}, \beta_\rho), \quad \text{for all } \sigma \in [2]^\omega.$$

Consequently,

$$U := \{ \hat{\alpha}_\sigma \mid \sigma \in [2]^\omega \} \quad \text{and} \quad V := \{ \hat{\beta}_\rho \mid \rho \in [2]^\omega \}$$

are the desired uncountable regular sets. \square

Remark. This theorem also holds if the function f is FOC-definable. The proof is exactly the same using Theorem 7.4 to show that the graph of f is regular. \dashv

Countable Structures

Let us conclude this section with a remark about countable ω -automatic structures.

Theorem 6.14. *Let \mathcal{A} be a countable structure. The following conditions are equivalent.*

- (1) \mathcal{A} is ω -automatic
- (2) \mathcal{A} is automatic.
- (3) \mathcal{A} has an injective ω -automatic presentation.

Proof. (3) \Rightarrow (1) is trivial; (1) \Rightarrow (2) follows by Corollary 6.10; and (2) \Rightarrow (3) holds since we can turn every injective automatic presentation of \mathcal{A} into an injective ω -automatic one. (We can turn every finite word w encoding an element of \mathcal{A} into an ω -word by appending, say, \square^ω .) \square

For ω -tree automatic structures, we would like to have results analogous to those above. But due to the much more involved combinatorics we are only able to handle the case of injective presentations, i.e., we cannot count equivalence classes but only elements. This results in the following weaker result.

Theorem 6.15. *A countable structure \mathcal{A} has an injective ω -tree automatic presentation if, and only if, \mathcal{A} is tree automatic.*

Proof. (\Rightarrow) Fix an injective tree automatic presentation of \mathcal{A} . We can turn every finite tree into an infinite one by appending some padding symbol \square . This operation preserves regularity.

(\Leftarrow) Let $\langle L_\delta, (L_R)_R \rangle$ be an injective ω -automatic presentation of \mathcal{A} . We can use the function from Theorem V.6.27 to encode every tree in L_δ by a finite one. Clearly, the images of L_δ and L_R under this map are regular and form a tree automatic presentation of \mathcal{A} . \square

There are still several open questions regarding the existence of injective presentations. For instance, every injective ω -automatic presentation can be equipped with a linear order (say, the lexicographic one). Does the converse hold?

Open Questions.

- (a) Is it true that every linearly ordered ω -automatic structure has an injective presentation?
- (b) Does an analogue of Theorem 6.14 hold for countable ω -tree automatic structures?
- (c) Is it true that an ω -automatic structure has an injective ω -automatic presentation if, and only if, it has an injective ω -tree automatic one?
- (d) Is there an ω -tree automatic structure that is not ω -automatic?

7 Counting Quantifiers

We have shown in Theorem 1.7 that every automatic structure has a decidable first-order theory. This result can be extended to the following version of first-order logic with counting quantifiers.

Definition 7.1. *First-order logic with counting quantifiers* FOC is the extension of FO by quantifiers of the form

- ♦ $\exists^\kappa x$ ‘There exist at least κ many x ’, for every infinite cardinal κ , and
- ♦ $\exists^{k,m} x$ ‘The number of values x is finite and congruent k modulo m ’, for $0 \leq k < m < \omega$.

Using injective presentations it is easy to eliminate such quantifiers.

Lemma 7.2. *Let L be a regular language of finite words or finite trees over the alphabet $\Sigma_\square \times \Gamma_\square$. The following languages are also regular, and one can effectively compute automata recognising them.*

- (a) $K_\infty := \{ w \mid \text{there are only finitely many } u \text{ with } w \otimes u \in L \}$.
- (b) $K_{k,m} := \{ w \mid \text{the number of } u \text{ with } w \otimes u \in L \text{ is finite and congruent } k \text{ modulo } m \}$.

Proof. (a) We start with the case of words. Note that $w \in K_\infty$ if, and only if,

$$w \otimes u \in L \quad \text{implies} \quad |u| < |w| + n,$$

where n is the number of states of some automaton recognising L . (If u is longer, we can pump some factor to obtain infinitely many words u' with $w \otimes u' \in L$. This argument is spelt out in detail in Lemma 8.2 below.) Let $h : (\Sigma_{\square} \times \Gamma_{\square})^* \rightarrow \Sigma_{\square}^*$ be the projection to the first component. It follows that

$$w \in K_{\infty} \quad \text{iff} \quad \text{there is no } k \geq n \text{ with } w \square^k \in h[L].$$

Let \mathcal{A} be an automaton for $h[L]$. We can construct an automaton for K_{∞} by

- ♦ removing from \mathcal{A} all transitions with the letter \square , and
- ♦ marking a state p as final if in \mathcal{A} there is no path labelled by \square^n starting at p .

The proof for trees is similar. There are infinitely many u with $w \otimes u \in L$ if, and only if, some tree u with $w \otimes u \in L$ contains a path that is at least n edges longer than the corresponding path of w . An automaton can recognise the complement of K_{∞} by guessing such a tree.

(b) We start with the case of word languages. Suppose that $L = \varphi^{-1}[P]$ for some homomorphism $\varphi : (\Sigma_{\square} \times \Gamma_{\square})^* \rightarrow \mathfrak{M}$ into a finite monoid \mathfrak{M} . We will construct a homomorphism $\psi : \Sigma^* \rightarrow \mathfrak{N}$ where $N := [m]^M$ such that

$$\psi(w) = \sigma \quad \text{iff} \quad \text{for all } a \in M, \text{ the number of } u \text{ with } \varphi(w \otimes u) = a \text{ is congruent } \sigma(a) \text{ modulo } m.$$

To do so, we define ψ by

$$\psi(c)(a) := |\{d \in \Gamma \mid \varphi(c \otimes d) = a\}| \bmod m, \quad \text{for } c \in \Sigma \text{ and } a \in M.$$

The monoid multiplication is given by

$$(\sigma \cdot \tau)(a) := \sum \{ \sigma(b) \cdot \tau(c) \mid b, c \in M, b \cdot c = a \} \bmod m.$$

It is straightforward to check that $[m]^M$ together with this product forms a monoid. For $a \in M$, let n_a be the number of words $u \in \Gamma^*$ such that

$\varphi(\langle \rangle \otimes u) = a$. Then we have $K_{km} = \psi^{-1}[Q] \cap K_{\infty}$ where $Q \subseteq [m]^M$ is the set of all functions σ such that

$$\sum \{ \sigma(a) \cdot n_b \mid a \cdot b \in P, n_b \text{ finite} \} \equiv k \pmod{m}.$$

The proof for trees is similar. Given a homomorphism into a tree algebra recognising L , we can use the same construct as above to find a homomorphism recognising K_{km} . Alternative, we can use an automata construction. Given an automaton \mathcal{A} for L , we can build an automaton \mathcal{B} that computes, for each vertex v of the input tree w and each state q of \mathcal{A} , the number of trees u (modulo m) such that from state q \mathcal{A} accepts the tree $w|_v \otimes u$. These numbers can be computed by \mathcal{B} while reading the input tree w from the leaves to the root. \square

Lemma 7.3. *Let $R \subseteq \Sigma^{\omega} \times \Gamma^{\omega}$ and $E \subseteq \Sigma^{\omega} \times \Sigma^{\omega}$ be regular and suppose that E is an equivalence relation. The following languages are also regular, and one can effectively compute automata recognising them.*

$$(a) \quad K_{\infty} := \{ w \mid \text{there are infinitely many } E\text{-classes } [u]_E \text{ with}$$

$$u \otimes w \in R \}.$$

$$(b) \quad K_{lm} := \{ w \mid \text{the number of } E\text{-classes } [u]_E \text{ with } u \otimes w \in R \text{ is finite and congruent } l \text{ modulo } m \}.$$

Proof. (a) By Theorem 6.7, $w \in K_{\infty}$ if, and only if, there is some \approx_* -class $[u]_*$ such that the set

$$\{ [v]_E \mid v \otimes w \in R, [v]_E \cap [u]_* \neq \emptyset \}$$

is infinite. Let \sqsubseteq'' be the well-ordering from Lemma 6.2. It follows that $w \in K_{\infty}$ if, and only if, the set

$$\{ v \mid v \otimes w \in R, v \text{ is the } \sqsubseteq''\text{-minimal element of } [v]_E \cap [u]_* \}$$

is infinite. This is equivalent to the condition that, for every $n < \omega$, there exists some $v \in \Gamma^{\omega}$ such that

- $v \otimes w \in R$,
- v is the \sqsubseteq^u -minimal element of $[v]_E \cap [u]_*$, and
- $v(i) \neq u(i)$, for some $i \geq n$.

These conditions are all regular.

(b) Let k be the constant from Theorem 6.7 and let \sqsubseteq^u be the well-ordering from Lemma 6.2. For $u_0, \dots, u_{k-1} \in \Sigma^\omega$, let $S_i(\vec{u})$ be the set of all words $v \in \Sigma^\omega$ such that

- $v \otimes w \in R$,
- v is the \sqsubseteq^{u_i} -minimal element of $[v]_E \cap [u_i]_*$, and
- $[v]_E \cap ([u_0]_* \cup \dots \cup [u_{i-1}]_*) = \emptyset$.

Then $w \in K_{lm}$ if, and only if, $w \notin K_\infty$ and there are $u_0, \dots, u_{k-1} \in \Sigma^\omega$ such that

- $v \otimes w \in R$ implies $[v]_E \cap [u_i]_* \neq \emptyset$, for some i , and
- $|S_0(\vec{u})| + \dots + |S_{k-1}(\vec{u})| \equiv l \pmod{m}$.

It remains show that the latter condition is regular. It is sufficient to show that a statement of the form $|S_i(\vec{u})| \equiv l \pmod{m}$ is regular. Note that all ω -words $v \in [u_i]_*$ can be represented by the finite prefix where they differ from u_i . Consequently, $S_i(\vec{u})$ corresponds to a finite set of finite words and we can use the same construction as in Lemma 7.2 to count them modulo m . \square

Theorem 7.4. *Given an automatic presentation π for some structure \mathfrak{A} and an FOC-formula $\varphi(\vec{x})$, one can effectively compute an automaton recognising the relation $\varphi^{\mathfrak{A}}$ defined by φ . The same holds for ω -automatic and tree automatic presentations, and for injective ω -tree automatic ones.*

Proof. We start with the automatic and tree automatic case. Then we can use Proposition 5.6 to replace the given presentation π by an injective one π' . Since the function mapping between these presentations is regular, we can translate every automaton for π' back into one for π .

Hence, let us assume that π is injective. We construct the desired automaton by induction on φ . For boolean operations and the usual first-order

quantifiers, we proceed as in Theorem 1.7. The quantifiers \exists^{\aleph_0} and $\exists^{k,m}$ can be eliminated using Lemma 7.2. Finally, note that, for uncountable cardinals κ , a formula of the form $\exists^\kappa y \varphi$ is always false. Hence, the relation defined by it is empty and, therefore, regular.

For ω -automatic structures, we cannot assume that the presentation is injective. Instead, we use Corollary 6.8 to eliminate the quantifier \exists^κ , for uncountable κ , and Lemma 7.3 for the quantifiers \exists^{\aleph_0} and $\exists^{k,m}$.

Finally, let \mathfrak{A} be an ω -tree automatic structures with an injective presentation. Fixing an interpretation $\mathfrak{A} \leq_{\text{FO}} \wp(\mathfrak{T}_{\text{bin}})$, we can translate every FOC-formula $\varphi(\vec{x})$ for \mathfrak{A} into an C_2MSO -formula $\hat{\varphi}(\vec{X})$ for $\mathfrak{T}_{\text{bin}}$. By Theorem V.6.4 (c), this formula is equivalent to an MSO-formula $\hat{\psi}(\vec{X})$. Finally, translating this formula back to the structure $\wp(\mathfrak{T}_{\text{bin}})$, we obtain an FO-formula $\psi(\vec{x})$ that is equivalent to $\varphi(\vec{x})$. We can translate this formula to the desired automaton using Theorem 1.7. \square

Theorem 7.5. *Let \mathfrak{A} be a structure with an ω -automatic presentation or an injective ω -tree automatic one. Then \mathfrak{A} has a decidable FOC-theory.*

Proof. We can use Theorem 7.4 to translate a given FOC-sentence φ into an automaton \mathcal{A} . As a sentence φ has no free variables, this automaton is over the unary alphabet $\Sigma^\circ = \{\langle \rangle\}$. Furthermore, \mathcal{A} accepts some input if, and only if, φ holds. As emptiness for tree automata is decidable, the claim follows. \square

Besides counting quantifiers, regular relations are also closed under several other generalised quantifiers. Let us take a look at two of them.

Definition 7.6. (a) For a formula $\varphi(\vec{x}_0 \dots \vec{x}_{n-1}, \vec{y})$ where $|\vec{x}_0| = \dots = |\vec{x}_{n-1}| = k$, a structure \mathfrak{A} , and a tuple \vec{c} in A , we define the *Ramsey quantifier* H by setting

$$\mathfrak{A} \models H[\vec{x}_0, \dots, \vec{x}_{n-1}] \varphi(\vec{x}_0, \dots, \vec{x}_{n-1}, \vec{c})$$

if, and only if, there exists an infinite relation $H \subseteq A^k$ such that

$$\mathfrak{A} \models \varphi(\vec{a}_0, \dots, \vec{a}_{n-1}, \vec{c}), \quad \text{for all pairwise distinct } \vec{a}_0, \dots, \vec{a}_{n-1} \in H.$$

(b) For a formula $\varphi(X)$ in which the second-order variable X (not necessarily monadic) occurs only negatively, we denote by $\text{UX}\varphi(X)$ the statement that there exists an infinite relation X satisfying φ .

(c) We denote by $\text{FOC}(\text{H})$ and $\text{FOC}(\text{U})$ the corresponding extensions of FOC by these quantifiers. J

Example. (a) We can express that an order tree $\mathfrak{T} = \langle T, \leq_{\text{pf}} \rangle$ contains an infinite branch by the formula

$$\varphi := \text{H}[x, x'] [x <_{\text{pf}} x' \vee x' <_{\text{pf}} x].$$

(b) We can express that an undirected graph $\mathfrak{G} = \langle V, E \rangle$ contains an infinite clique by the formulae

$$\varphi := \text{H}[x, x'] [Exx']$$

$$\text{or } \psi := \text{UZ} \forall xy [x \in Z \wedge y \in Z \wedge x \neq y \rightarrow Exy]. \quad \text{J}$$


Remark. Note that we can express the quantifier H with the help of U since

$$\text{H}[\bar{x}_0, \dots, \bar{x}_{n-1}] \varphi \equiv \text{UZ} \forall \bar{x}_0 \dots \forall \bar{x}_{n-1} \left[\bigwedge_{i < n} Z\bar{x}_i \wedge \bigwedge_{i < j} \bar{x}_i \neq \bar{x}_j \rightarrow \varphi \right]. \quad \text{J}$$

To prove that these quantifiers preserve regularity, we start with the following observation.

Definition 7.7. Let Σ be an alphabet.

(a) A *comb* of words is a sequence $\gamma = \langle s_i, t_i \rangle_{i < \omega}$ where $s_i, t_i \in \Sigma^+$ are words with $|t_i| \leq |s_i|$.



The set *encoded* by a comb γ is

$$\{s_0 \dots s_{i-1} t_i \mid i < \omega\}.$$

We usually do not distinguish between a comb and the set it represents. In particular, we write $w \in \gamma$ for $w = s_0 \dots s_{i-1} t_i$.

(b) A *comb* of trees is a pair $\gamma = \langle s, \langle t_i \rangle_{i < \omega} \rangle$ where s is an infinite tree over Σ and $\langle t_i \rangle_i$ is a sequence of finite trees over Σ such that

$$\text{dom}(t_i) \subseteq s \sqcap t_{i+1}, \quad \text{for all } i < \omega,$$

$$\text{where } s \sqcap t := \{v \in \text{dom}(t) \mid s(u) = t(u) \text{ for all } u \leq_{\text{pf}} v\}.$$

Lemma 7.8. A set P of finite words or finite trees is infinite if, and only if, there exists a comb γ with $\gamma \subseteq P$.

Proof. (\Leftarrow) is trivial since every comb encodes an infinite set. For (\Rightarrow) , let P be an infinite set of finite words or finite trees. The closure of P under prefixes together with the prefix-ordering \leq_{pf} forms an infinite finitely-branching tree which contains an infinite branch, by the Lemma of König. Consequently, we can find an infinite word/tree ζ such that every finite prefix $z <_{\text{pf}} \zeta$ is a prefix of some $p \in P$. Using ζ , we can pick an infinite sequence $(p_i)_{i < \omega}$ of elements of P such that

$$\text{dom } p_i \subseteq \zeta \sqcap p_{i+1}, \quad \text{for all } i < \omega.$$

In the case of trees, the pair $\langle \zeta, \langle p_i \rangle_i \rangle$ forms the desired comb. For words, we inductively construct two sequences $(s_i)_i$ and $(t_i)_i$ of words by

$$p_i \sqcap \zeta = s_0 \dots s_i \quad \text{and} \quad p_i = s_0 \dots s_i t_i.$$

Then $\langle s_{i+1}, t_i \rangle_{i < \omega}$ is the desired comb. \square

The advantage of combs over arbitrary infinite sets is that they can easily be encoded as ω -words/ ω -trees. Therefore they can be used as a weak kind of power-set operation on automatic structures.

Definition 7.9. Let \mathfrak{A} be an automatic or tree-automatic Γ -structure with an injective presentation π .

(a) For $k < \omega$, we denote by $k \times \mathfrak{A}$ the structure with universe $A + A^k$, all relations from \mathfrak{A} (on the universe A), and the additional relations

$$\text{pr}_i := \{ \langle a, \bar{b} \rangle \in A \times A^k \mid a = b_i \}, \quad \text{for } i < k.$$

(b) We denote by $\text{cb}(\mathfrak{A})$ the $(\Gamma + \{C, \text{in}\})$ -structure with universe $A + C$, where C is the set of all combs γ such that $\gamma \subseteq A$. The relations are those of \mathfrak{A} , the predicate C , and the element relation

$$\text{in} := \{ \langle a, \gamma \rangle \in A \times C \mid a \in \gamma \}.$$

Lemma 7.10. *Let \mathfrak{A} be an automatic structure or a tree-automatic one with an injective presentation $\langle L_\delta, (L_R)_R \rangle$ over the alphabet Σ .*

- (a) *The structure $k \times \mathfrak{A}$ has an automatic/tree-automatic presentation such that the set A is represented by L_δ .*
- (b) *$\text{cb}(\mathfrak{A})$ has an injective ω -automatic/ ω -tree automatic presentation such that the set A is represented by $L_\delta \square^\omega$, for some new symbol $\square \notin \Sigma$.*

Proof. For the first claim, we can encode elements $a \in A$ as in the given presentation while tuples $\bar{a} \in A^k$ are encoded by words/trees of the form $\$(w_0 \otimes \dots \otimes w_{k-1})$ over the alphabet $\Sigma^k + \{\$\}$ where w_i encodes a_i and $\$$ is a new marker symbol.

For the second one, first note that \mathfrak{A} has an ω -automatic presentation where each $a \in A$ is encoded by the ω -word $a \square^\omega$. Furthermore, we can encode a comb $\gamma = \langle s_i, t_i \rangle_i$ by the ω -word

$$(s_0 \otimes t_0 \otimes \$)(s_1 \otimes t_1 \otimes \$) \dots \in (\Sigma_\square^2 \times \{\$, \square\})^\omega,$$

s_0	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8
t_0	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8
$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$	$\$ \square \dots \square$

where the third component is used as a marker for the beginning of a new segment. Clearly, this language is regular and so is the element relation in . A similar encoding works for trees. \square

Proposition 7.11. *Let \mathfrak{A} be an automatic or tree-automatic structure. For every FOC(U)-formula $\varphi(\bar{x})$ (without free second-order variables), we can effectively construct an FOC-formula $\psi(\bar{x})$ such that*

$$\mathfrak{A} \models \varphi(\bar{a}) \quad \text{iff} \quad \text{cb}(k \times \mathfrak{A}) \models \psi(\bar{a}), \quad \text{for all } \bar{a} \text{ in } A,$$

where k is the maximal arity of a second-order variable in φ .

Proof. Let ψ be the formula obtained from φ by replacing

- ♦ every subformula of the form $\text{UZ}\vartheta$ by $\exists z\vartheta$ and
- ♦ every atomic subformula of the form $Z\bar{y}$ (with $n := |\bar{y}|$) by

$$\exists u \left[\text{in}(u, z) \wedge \bigwedge_{i < n} \text{pr}_i(y_i, u) \wedge \bigwedge_{n \leq i < k} \text{pr}_i(y_{n-1}, u) \right].$$

To show that this translation works, first note that every infinite relation $R \subseteq A^n$, with $n \leq k$, can be encoded as a subset $R' \subseteq A^k$: we can turn every n -tuple into a k -tuple by repeating the last component $(k - n)$ times. Consequently, every such relation can be represented by a comb over A^k .

Therefore we only have to show that we can restrict second-order quantifiers to subsets of A^k that can be encoded by a comb without changing the meaning of the formula. A straightforward induction on φ using Lemma 7.8 establishes that, for every FOC-formula $\varphi(\bar{x}, \bar{Y}, Z)$ where the variables \bar{Y}, Z occur only negatively, for all parameters \bar{a} in A , and all combs \bar{y} ,

$$\mathfrak{A} \models \text{UZ}\varphi(\bar{a}, \bar{y}, Z) \quad \text{iff} \quad \mathfrak{A} \models \varphi(\bar{a}, \bar{y}, \delta), \quad \text{for some comb } \delta. \quad \square$$

Theorem 7.12. *Given an injective automatic or tree-automatic presentation π for some structure \mathfrak{A} and an FOC(U)-formula $\varphi(\bar{x})$ (without free second-order variables), one can effectively compute an automaton recognising the relation $\varphi^{\mathfrak{A}}$ defined by φ .*

Proof. By Proposition 7.11, there exists an FOC-formula $\psi(\bar{x})$ and a constant k such that

$$\mathfrak{A} \models \varphi(\bar{a}) \quad \text{iff} \quad \text{cb}(k \times \mathfrak{A}) \models \psi(\bar{a}), \quad \text{for all } \bar{a} \text{ in } A.$$

We have shown in Lemma 7.10 that $\text{cb}(k \times \mathfrak{A})$ is ω -automatic/ ω -tree automatic, and that the corresponding presentation restricts on the set A to π . Therefore, we can use Theorem 7.12 to construct an automaton \mathcal{A} recognising $\psi^{\text{cb}(k \times \mathfrak{A})}$. Since the presentations of the two structures are compatible, this automaton also recognises $\psi^{\mathfrak{A}}$. \square

Corollary 7.13. *The classes of automatic structures and tree-automatic structures are closed under many-dimensional $\text{FOC}(\text{U})$ -interpretations.*

8 Proving Non-Automaticity

Proving that a certain structure is automatic is rather easy: one just has to find a presentation. But proving that it is not automatic is much harder: one has to show that no presentation can exist. The goal of this section is to develop tools to make this task easier. The central idea of all what follows is contained in the following observation (which is a variant of Lemma XI.7.10).

Definition 8.1. Let \mathfrak{A} be a relational structure.

(a) Given a presentation π and an element $a \in A$, we write

$$\|a\| := \min \{ |w| \mid \pi(w) = a \}.$$

For tuples \bar{a} and sets S , we set

$$\|\bar{a}\| := \max_i \|a_i\| \quad \text{and} \quad \|S\| := \max_{a \in S} \|a\|.$$

(b) We say that a relation $R \subseteq A^{k+l}$ has *finite out-degree* if, for every $\bar{a} \in A^k$, there are only finitely many $\bar{b} \in A^l$ with $\bar{a}\bar{b} \in R$. A formula $\varphi(\bar{x}, \bar{y})$ has finite out-degree if the relation $\varphi^{\mathfrak{A}}$ defined by it has. \urcorner

Lemma 8.2. *Let \mathfrak{A} be an automatic structure with an injective presentation. For every $\text{FOC}(\text{U})$ -formula $\varphi(\bar{x}; \bar{z})$, there exists a constant k with the following two properties.*

(I) *For all tuples $\bar{a}, \bar{c} \in A$, there exist a tuple \bar{a}' such that*

$$\mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \leftrightarrow \varphi(\bar{a}'; \bar{c}) \quad \text{and} \quad \|\bar{a}'\| \leq \|\bar{c}\| + k.$$

(II) $\|\bar{a}\| > \|\bar{c}\| + k$ *implies that there are infinitely many tuples \bar{a}' with*

$$\mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \leftrightarrow \varphi(\bar{a}'; \bar{c}).$$

Proof. Fix a morphism $\eta : (\Sigma_{\square}^m)^* \rightarrow \mathfrak{M}$ into a finite monoid recognising the relation $\varphi^{\mathfrak{A}}$. We claim that $k := |M|$ is the desired constant. For the proof, fix tuples \bar{a} and \bar{c} . If $\|\bar{a}\| \leq \|\bar{c}\| + k$, there is nothing to do. Hence, we may assume that $\|\bar{a}\| > \|\bar{c}\| + k$. We can factorise the word $\bar{a}_{\circ}^{\otimes} \otimes \bar{c}^{\otimes}$ as wu where $|w| = \|\bar{c}\|$. Then $|u| > k$ implies that there exists a factorisation $u = u_{\circ}u_1u_2$ such that $\eta(u_{\circ}) = \eta(u_{\circ}u_1)$.

a_{\circ}				
a_1				
a_2				
c_{\circ}				
c_1				
	w	u_{\circ}	u_1	u_2

Consequently,

$$\eta(wu_{\circ}u_1u_2) = \eta(wu_{\circ}(u_1)^nu_2), \quad \text{for all } n < \omega.$$

Let $\bar{a}_n\bar{c}_n$ be the tuple encoded by $wu_{\circ}(u_1)^nu_2$. By choice of w it follows that all components of u corresponding to the elements c_i contain only blanks. Hence, $\bar{c}_n = \bar{c}$ and we have

$$\mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \leftrightarrow \varphi(\bar{a}_n; \bar{c}), \quad \text{for all } n < \omega.$$

Furthermore, $\|\bar{a}_{\circ}\| < \|\bar{a}\|$. Replacing \bar{a} by \bar{a}_{\circ} and repeating this construction we will therefore eventually obtain a tuple \bar{a}_{\circ} with $\|\bar{a}_{\circ}\| \leq \|\bar{c}\| + k$. \square

The fact that we will mainly use below is the following consequence of this lemma.

Corollary 8.3. *Let \mathfrak{A} be an automatic structure with an injective presentation. For every $\text{FOC}(\text{U})$ -formula $\varphi(\bar{x}; \bar{z})$ of finite out-degree, there exists a constant k such that*

$$\mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \quad \text{implies} \quad \|\bar{a}\| \leq \|\bar{c}\| + k, \quad \text{for all } \bar{a}, \bar{c}.$$

A similar lemma holds for tree automatic structures. Instead of parametrising the encodings by their length, we can now bound their number of vertices or their height.

Definition 8.4. Let \mathfrak{A} be a Γ -structure and $U \subseteq A$.

(a) Suppose that \mathfrak{A} has an injective tree automatic presentation. We set

$$D(U) := \left| \bigcup_{c \in U} \text{dom}(c) \right|,$$

$$h(U) := \max \{ n \mid \text{there is some tree } c \in U \text{ with height } n \}.$$

For tuples \bar{a} or single elements a , we also use the notation $D(\bar{a})$, $D(a)$, etc.

(b) For a set Δ of formulae of the form $\varphi(\bar{x}; \bar{z})$, a set $U \subseteq A$, and two tuples $\bar{a}, \bar{b} \in A^m$, we define

$$\bar{a} \simeq_U^{\Delta} \bar{b} \quad : \text{iff} \quad \mathfrak{A} \models \varphi(\bar{a}; \bar{c}) \leftrightarrow \varphi(\bar{b}; \bar{c}),$$

$$\text{for all } \bar{c} \in U^n \text{ and all } \varphi(\bar{x}; \bar{z}) \in \Delta.$$

Lemma 8.5. Let \mathfrak{A} be a tree automatic structure and $\varphi(\bar{x}; \bar{z})$ an FOC(U)-formula. There exist constants k, m with the following two properties.

(I) For all tuples $\bar{a} \subseteq A$ and all $U \subseteq A$, there exist a tuple $\bar{a}' \simeq_U^{\varphi} \bar{a}$ such that

$$D(\bar{a}') \leq k \cdot D(U) \quad \text{and} \quad h(\bar{a}') \leq h(U) + m.$$

(II) If $D(\bar{a}) > k \cdot D(U)$ or $h(\bar{a}) > h(U) + m$, then there are infinitely many tuples \bar{a}' with $\bar{a}' \simeq_U^{\varphi} \bar{a}$.

Proof. We fix an injective presentation of the structure \mathfrak{A} and two automata $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, F \rangle$ and $\mathcal{A}' = \langle Q', \Sigma, \Delta', q'_0, F' \rangle$ recognising the relations $\varphi^{\mathfrak{A}}$ and $\neg \varphi^{\mathfrak{A}}$, respectively. Set $s := |Q| + |Q'|$. We claim that the desired constants are $m := s^s$ and $k := 2^{m+2}$. For the proof, fix a tuple \bar{a} and a set U . If $D(\bar{a}) \leq k \cdot |D(U)|$ and $h(\bar{a}) \leq h(U) + m$, there is nothing to do. Hence, we may assume that

$$D(\bar{a}) > k \cdot |D(U)| \quad \text{or} \quad h(\bar{a}) > h(U) + m.$$

We claim that in both cases the set $\text{dom}(\bar{a}^{\otimes}) \setminus D(U)$ contains a path of length at least m .

If $h(\bar{a}) > h(U) + m$ this is clear. Hence, we may assume that $D(\bar{a}) > k \cdot |D(U)|$. Note that every connected component of the subforest of \bar{a}^{\otimes} induced by the set $\text{dom}(\bar{a}^{\otimes}) \setminus D(U)$ is attached to some leaf of $D(U)$. Furthermore, there are at most two components attached to each leaf. Since the tree $D(U)$ has $(|D(U)| + 1)/2$ leaves, it follows that $\text{dom}(\bar{a}^{\otimes}) \setminus D(U)$ has at most $|D(U)| + 1$ components. Consequently, there exists a component C of size

$$\frac{k \cdot |D(U)|}{|D(U)| + 1} \geq \frac{k}{2} = 2^{m+1}.$$

As C is a binary tree of size $|C| \geq 2^{m+1}$, it contains the desired path of length at least m .

Having found our path p , we proceed as follows. For every tuple $\bar{c} \subseteq U$, we fix an accepting run $\rho_{\bar{c}}$ of \mathcal{A} or \mathcal{A}' on $\bar{a} \otimes \bar{c}$. We may assume that, if $\rho_{\bar{c}}$ and $\rho_{\bar{d}}$ agree on the first vertex of p , they agree on all vertices of p . As the length of p is at least $m = s^s$, there exist two vertices $u <_{\text{pf}} v$ of p such that $\rho_{\bar{c}}(u) = \rho_{\bar{d}}(v)$, for all $\bar{c} \subseteq U$. (We can apply the Pigeon Hole Principle to the colouring assigning to each vertex v of p the function mapping $\rho_{\bar{c}}(w)$ to $\rho_{\bar{c}}(v)$, where w is the first vertex of p .) Let \bar{a}_i be the tuple obtained from \bar{a} by repeating the part between u and v i -times. Then $\bar{a}_0, \bar{a}_1, \dots$ are infinitely many tuples with $\bar{a}_i \simeq_U^{\varphi} \bar{a}$. Furthermore, $D(\bar{a}_0) < D(\bar{a})$ and $h(\bar{a}_0) \leq h(\bar{a})$, and at least one branch of \bar{a}_0 is strictly shorter than the corresponding one of \bar{a} . Hence, repeating this construction (if necessary) we obtain a tuple \bar{a}_0 with $D(\bar{a}_0) \leq k \cdot D(U)$ and $h(\bar{a}_0) \leq h(U) + m$. \square

The following dual point of view is sometimes useful.

Exercise 8.1. Let \mathfrak{A} be a structure, $R \subseteq A^n \times A$ an FOC-definable relation, and $U \subseteq A$ a set. We call a subset of the form

$$\{ a \in U \mid \langle \bar{c}, a \rangle \in R \}, \quad \text{for } \bar{c} \in A^n,$$

an R -trace of U .

(a) Prove that, if \mathfrak{A} is automatic, there exists a constant k such that, for all $n < \omega$, the number of different R -traces of $U := A \cap \Sigma^{<n}$ is at most $k \cdot |U|$. *Hint.* If $\bar{c} \simeq_U \bar{c}'$, the R -traces associated with \bar{c} and \bar{c}' are the same.

(b) Prove that, if \mathfrak{A} is tree automatic, there exists a constant k such that, for all $n < \omega$, the number of different R -traces of $U := A \cap \Sigma^{<n}$ is at most $|U|^k$. \square

The basic observation in Lemmas 8.2 and 8.5 can be used in several different ways to restrict the form of automatic structures. Let us present them one by one.

Length Arguments

The simplest way to use Lemma 8.2 is to (i) compute explicit length bounds for the encodings of elements and then (ii) show that the number of elements of a given length exceeds the number of possible codes of this length. To do so, the following bound is frequently useful.

Lemma 8.6. *Let $L \subseteq \Sigma^*$ be regular. There exist constants d and m such that*

$$|L \cap \Sigma^{\leq k+1}| \leq d \cdot |L \cap \Sigma^{\leq k}|, \quad \text{for all } k \geq m.$$

Proof. Let $\eta : \Sigma^* \rightarrow \mathfrak{M}$ be a homomorphism into a finite monoid recognising L , set $m := |M|$, and let

$$P_k := \{w \in \Sigma^k \mid wu \in L \text{ for some } u\}, \quad \text{for } k < \omega.$$

Note that for every $v \in \Sigma^*$, there is some $u \in \Sigma^{<m}$ with $\eta(u) = \eta(v)$. This implies that, for every $w \in P_k$, there is some word $u \in \Sigma^{<m}$ with $wu \in L$. Hence,

$$|P_k| \leq |L \cap \Sigma^{<k+m}|.$$

Since, for every $w \in P_k$, there are at most $|\Sigma^m|$ words u with $wu \in P_{k+m}$, we have

$$|P_{k+m}| \leq |\Sigma^m| \cdot |P_k| \leq |\Sigma^m| \cdot |L \cap \Sigma^{<k+m}|, \quad \text{for all } k < \omega.$$

For $k \geq m$, we obtain

$$|P_k| = |P_{(k-m)+m}| \leq |\Sigma^m| \cdot |L \cap \Sigma^{<(k-m)+m}| = |\Sigma^m| \cdot |L \cap \Sigma^{<k}|.$$

Similarly there are, for every $w \in P_k$, at most $|\Sigma|$ letters c with $wc \in L \cap \Sigma^{k+1}$. For $k \geq m$, it therefore follows that

$$\begin{aligned} |L \cap \Sigma^{\leq k+1}| &= |L \cap \Sigma^{\leq k}| + |L \cap \Sigma^{k+1}| \\ &\leq |L \cap \Sigma^{\leq k}| + |\Sigma| \cdot |P_k| \\ &\leq |L \cap \Sigma^{\leq k}| + |\Sigma| \cdot |\Sigma|^m \cdot |L \cap \Sigma^{\leq k}| \\ &= (1 + |\Sigma|^{m+1}) \cdot |L \cap \Sigma^{\leq k}|. \quad \square \end{aligned}$$

As an example let us use arguments based on length bounds in the context of semigroups. In particular, we are interested in the question of which semigroups can and cannot be embedded into a given structure. To be applicable not only to semigroups, we use the following slightly more general notion.

Definition 8.7. Let $R \subseteq A^3$ be a ternary relation on a set A . We say that a semigroup \mathfrak{S} can be *embedded* into R if there exists a function $f : S \rightarrow A$ such that

$$\langle f(s), f(t), a \rangle \in R \quad \text{iff} \quad a = f(st), \quad \text{for all } s, t \in S \text{ and } a \in A. \quad \square$$

Lemma 8.8. *Let \mathfrak{A} be an automatic structure and $R \subseteq A^3$ an $\text{FOC}(\mathfrak{U})$ -definable relation. There exists a constant k such that, for every semigroup $\mathfrak{S} = \langle S, \cdot \rangle$ that can be embedded into R , we have*

$$\|a_0 \cdots a_{n-1}\| \leq \|\bar{a}\| + k \lceil \log_2 n \rceil, \quad \text{for all } \bar{a} \in S^n.$$

Proof. By Corollary 8.3, there exists a constant k such that

$$\|a \cdot b\| \leq \max\{\|a\|, \|b\|\} + k.$$

By induction on n and setting $m := \lceil n/2 \rceil$, it now follows that

$$\begin{aligned} \|a_0 \cdots a_{n-1}\| &\leq k + \max\{\|a_0 \cdots a_{m-1}\|, \|a_m \cdots a_{n-1}\|\} \\ &\leq k + \max\left\{\max_{i < m} \|a_i\| + k \lceil \log_2 m \rceil, \right. \\ &\quad \left. \max_{m \leq i < n} \|a_i\| + k \lceil \log_2 (n - m) \rceil\right\} \end{aligned}$$

$$\begin{aligned}
&\leq \max_{i < n} \|a_i\| + k(1 + \max\{\lceil \log_2 m \rceil, \lceil \log_2(n - m) \rceil\}) \\
&\leq \|\tilde{a}\| + k(1 + \lceil \log_2 n \rceil - 1) \\
&= \|\tilde{a}\| + k\lceil \log_2 n \rceil. \quad \square
\end{aligned}$$

Proposition 8.9. *Let \mathcal{A} be an automatic structure and $R \subseteq A^3$ an $\text{FOC}(\mathcal{U})$ -definable relation. There exists a number $n < \omega$ such that the semigroup $\langle \mathbb{N}, + \rangle^n$ cannot be embedded into R .*

Proof. Fix an injective presentation of \mathcal{A} over the alphabet Σ . For a contradiction, suppose that there exist functions $f_n : \mathbb{N}^n \rightarrow A$ as above, for every $n < \omega$. Let k be the constant from Lemma 8.8. Choose n such that

$$n > 2(k+1) \log_2 |\Sigma|.$$

and let $P = \{p_0, \dots, p_{n-1}\}$ be the image of the generators of \mathbb{N}^n under f_n . We set

$$N_m := \left\{ \sum_{i < n} c_i p_i \mid 0 \leq c_i < 2^m \right\}.$$

By choice of k , it follows for $m \geq \|P\|$ that

$$\|c_i p_i\| \leq \|P\| + k\lceil \log_2 2^m \rceil \leq m(k+1), \quad \text{for all } i < n, \ c_i < 2^m.$$

By the same argument, we have

$$\|a\| \leq m(k+1) + k\lceil \log_2 n \rceil, \quad \text{for all } a \in N_m \text{ and } m < \omega.$$

Since all the sums in N_m are different, it therefore follows that

$$(2^m)^n \leq |N_n| \leq |\Sigma|^{m(k+1) + k\lceil \log_2 n \rceil + 1}.$$

Consequently,

$$\begin{aligned}
mn &\leq \log_2 |\Sigma| \cdot [m(k+1) + k\lceil \log_2 n \rceil + 1] \\
&\leq (k+1)(m + \log_2 n) \log_2 |\Sigma|,
\end{aligned}$$

which, for $m \geq \log_2 n$, implies that

$$n \leq (k+1)\left(1 + \frac{1}{m} \log_2 n\right) \log_2 |\Sigma| \leq 2(k+1) \log_2 |\Sigma|.$$

A contradiction to our choice of n . \square

Corollary 8.10. *If the semigroup $\langle \mathbb{N}, + \rangle^{(\omega)}$ can be embedded into a semigroup $\mathfrak{S} = \langle S, + \rangle$ (not necessarily commutative), then \mathfrak{S} is not automatic.*

Corollary 8.11.

- (a) *Skolem arithmetic $\langle \mathbb{N}, \cdot \rangle$ is not automatic.*
- (b) *No free group with at least two generators is automatic.*

Exercise 8.2. Let p be a prime number and let \mathfrak{S} be a semigroup into which we can embed $\mathbb{N}[1/p]^{(\omega)}$ or $(\mathbb{Z}[1/p]/\mathbb{Z})^{(\omega)}$. Prove that \mathfrak{S} is not automatic. \downarrow

Let us also give an example of a structure that is not tree-automatic. The *Rado graph* \mathfrak{R} is a countably infinite, undirected graph such that, for every pair of disjoint finite sets A and B , there is some vertex v that has edges to all vertices in A but to none in B . A simple back-and-forth argument shows that this condition uniquely determines \mathfrak{R} up to isomorphism.

Exercise 8.3. Prove that there exists exactly one countably infinite graph satisfying the above condition. \downarrow

Theorem 8.12. *The Rado graph is not tree automatic and not ω -automatic.*

Proof. It is sufficient to prove that the Rado graph $\mathfrak{R} = \langle V, E \rangle$ is not tree automatic. Since \mathfrak{R} is countable, it then follows by Theorem 6.14 that it is also not ω -automatic.

For a contradiction, suppose that \mathfrak{R} is tree automatic and fix an injective presentation of \mathfrak{R} over a binary alphabet. Let $W_n \subseteq V$ be the set of all vertices encoded by a tree of height at most n . Let m be the constant from Lemma 8.5

for the formula $\varphi(x, z) := Exz$. It follows that, for every $X \subseteq W_n$, there exists some vertex v with

$$\langle v, x \rangle \in E \quad \text{iff} \quad x \in X, \quad \text{for all } x \in W_n,$$

and $h(v) \leq h(W_n) + m = n + m$.

This implies that $v \in W_{n+m}$. Since there are $2^{|W_n|}$ subsets $X \subseteq W_n$, it follows that $|W_{n+m}| \geq 2^{|W_n|}$. By induction, we obtain

$$|W_{n+im}| \geq \exp(|W_n|, i),$$

where the iterated exponential function \exp is defined by

$$\exp(k, 0) := k \quad \text{and} \quad \exp(k, i+1) := 2^{\exp(k, i)}.$$

For large enough i , we obtain a contradiction since the number of binary $\{0, 1\}$ -labelled trees of height at most $n + im$ is at most $3^{2^{n+im+1}}$. (We can encode a $\{0, 1\}$ -labelled tree of height at most n by a $\{0, 1, \square\}$ -labelled full binary tree of height n , and such a tree has less than 2^{n+1} vertices.) \square

Counting Arguments

Length bounds can frequently be simplified by replacing them by certain counting arguments. We start with a simple version, a more elaborate one will be presented below.

Definition 8.13. Let \mathcal{A} be a structure and $R \subseteq A^{k+l}$ a relation.

(a) Suppose that R has finite out-degree. The *expansion* of R is the function $\text{ex}_R : \omega \rightarrow \omega$ defined by

$$\text{ex}_R(n) := \min \{ |R[S]| \mid S \subseteq A, |S| = n \},$$

where

$$R[S] := \{ b_i \in A \mid i < l \text{ and } \langle \bar{a}, \bar{b} \rangle \in R \text{ for some } \bar{a} \in S^k \}.$$

(b) We say that R has *bounded expansion* if there exists a constant c such that $\text{ex}_R(n) \leq cn$, for all n . \lrcorner

Proposition 8.14. Let \mathcal{A} be an automatic structure. Every $\text{FOC}(\mathcal{U})$ -definable relation of finite out-degree has bounded expansion.

Proof. Fix an injective presentation $\langle L_\delta, L_=\rangle, (L_R)_R$ of \mathcal{A} over the alphabet Σ and let $R \subseteq A^{k+l}$ be $\text{FOC}(\mathcal{U})$ -definable. Given $n < \omega$, let $m < \omega$ be the minimal number such that $|L_\delta \cap \Sigma^{\leq m}| \geq n$ and choose a set $S \subseteq L_\delta \cap \Sigma^{\leq m}$ of size $|S| = n$. Let c be the constant from Corollary 8.3. Then

$$R[S] \subseteq L_\delta \cap \Sigma^{\leq m+c},$$

which implies that

$$\text{ex}_R(n) \leq |R[S]| \leq |L_\delta \cap \Sigma^{\leq m+c}| \leq d^{c+1} \cdot |L_\delta \cap \Sigma^{\leq m-1}| \leq d^{c+1}n,$$

where d the constant from Lemma 8.6. \square

Example. Let \mathcal{A} be a structure with an $\text{FOC}(\mathcal{U})$ -definable pairing function $f : A \times A \rightarrow A$. Then $\text{ex}_f(n) = n^2$ and \mathcal{A} is not automatic. \lrcorner

Growth Arguments

In cases where the preceding proposition does not apply, the following, more involved condition is often useful.

Definition 8.15. Let \mathcal{A} be a structure and $\varphi(\bar{x}, y)$ a formula. For a set $U \subseteq A$ and a number $n < \omega$, we define the set $N_\varphi(U, n)$ of *reachable elements at distance n* by

$$N_\varphi(U, 0) := U,$$

and $N_\varphi(U, n+1) := N_\varphi(U, n) \cup \{ b \in A \mid \mathcal{A} \models \varphi(\bar{a}, b) \text{ for some } \bar{a} \subseteq U \}.$

Proposition 8.16. Let \mathcal{A} be an automatic structure, $U \subseteq A$ finite, and let $\varphi(\bar{x}, y) \in \text{FOC}(\mathcal{U})$ be a formula of finite out-degree. There exist constants d, k such that

$$|N_\varphi(U, n)| \leq 2^{dn+k}, \quad \text{for all } n < \omega.$$

Proof. Fix an injective presentation π of \mathfrak{A} over a binary alphabet and set $k := \|U\|$. Since φ has finite out-degree, we can use Corollary 8.3 to find a constant d such that

$$\|a\| \leq \|N(U, n)\| + d, \quad \text{for all } a \in N(U, n+1).$$

By induction, we obtain

$$\|a\| \leq dn + k, \quad \text{for all } a \in N(U, n).$$

Consequently,

$$|N(U, n)| \leq |\{0, 1\}^{<(dn+k)}| = 2^{dn+k+1} - 1. \quad \square$$

Let us give a few applications. First, we consider pairing functions.

Proposition 8.17. *Let \mathfrak{A} be a structure with an infinite subset $U \subseteq A$ such that there exists an FOC(U)-definable injective function $f : U \times U \rightarrow U$. Then \mathfrak{A} is not automatic.*

Proof. For a contradiction, suppose that \mathfrak{A} is automatic. Note that, for $S \subseteq U$, the set

$$N_f(S, 1) = S \cup \{f(a, b) \mid a, b \in S\}$$

has size $|N_f(S, 1)| = |S| + |S|^2 \geq |S|^2$. By induction it follows that

$$\begin{aligned} |N_f(S, n)| &= |N_f(N_f(S, n-1), 1)| \geq |N_f(S, n-1)|^2 \\ &\geq (|S|^{2^{(n-1)}})^2 = |S|^{2^n}. \end{aligned}$$

For $|S| > 1$, this contradicts Proposition 8.16. \square

Corollary 8.18.

- (a) *The structure $\langle \mathbb{N}, f \rangle$ is not automatic where $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ is a pairing function.*
- (b) *No free semigroup with at least two generators is automatic.*

- (c) *If Γ is a signature with at least one constant symbol and at least one symbol of arity at least 2, the term algebra over Γ is not automatic.*

Exercise 8.4. Let \mathfrak{A} be an automatic structure with a FOC(U)-definable equivalence relation \sim and let $[a_0]_{\sim}, [a_1]_{\sim}, \dots$ be an enumeration of all finite \sim -classes such that $|[a_0]_{\sim}| \leq |[a_1]_{\sim}| \leq \dots$. Prove that there exist constants c, d such that

$$|[a_n]_{\sim}| \leq 2^{cn+d}.$$

We can strengthen Corollary 8.11 (a) as follows.

Proposition 8.19. *The divisibility order $\langle \mathbb{N}, | \rangle$ is not automatic.*

Proof. Let $\varphi(x; zz')$ be the FO-formula stating that

- x is the \leq_{lex} -minimal prime number with $x >_{\text{lex}} z$, or
- x is the least common multiple of z and z' , or
- $x = pz$, for some prime factor p of z .

Note that the latter condition can be expressed in FO by expressing that $z \mid x$, $z \neq x$, every prime factor of x is also a prime factor of z , and $z \mid y \mid x$ implies $y = x$ or $y = z$.

Let $p_0 <_{\text{lex}} p_1 <_{\text{lex}} \dots$ be an enumeration of all prime numbers. Then

$$p_0, \dots, p_{n-1} \in N_{\varphi}(\{p_0\}, n).$$

Hence,

$$p_n^k \in N_{\varphi}(\{p_0\}, n+k).$$

Set $\gamma(n) := |N_{\varphi}(\{p_0\}, n)|$. It follows that $N_{\varphi}(\{p_0\}, n)$ contains at least

- n numbers p_0, p_0^2, \dots, p_0^n ,
- $\gamma(n-1)$ numbers that are coprime to p_0 ,
- for every $0 < k < n$, $\gamma(n-2) - 1$ numbers of the form $p_0^k m$, where m is coprime to p_0 .

It follows that

$$\begin{aligned}\gamma(n) &\geq n + \gamma(n-1) + (n-1)(\gamma(n-2) - 1) \\ &= \gamma(n-1) + (n-1)\gamma(n-2) \\ &\geq \gamma(n-2) + (n-1)\gamma(n-2) = n\gamma(n-2).\end{aligned}$$

For even $n > 0$, we obtain by induction

$$\begin{aligned}\gamma(n) &\geq n(n-2)(n-4)\cdots 4 \cdot 2 \cdot \gamma(0) \\ &= n(n-2)(n-4)\cdots 4 \cdot 2 \cdot 1 \\ &\geq (n/2)! \\ &\geq 2^{cn \log_2 n},\end{aligned}$$

for some constant c . A contradiction to Proposition 8.16. \square

There also exists a version of Propositions 8.16 for tree automatic structures.

Proposition 8.20. *Let \mathfrak{A} be a tree automatic structure, $U \subseteq A$ finite, and let $\varphi(\bar{x}, y) \in \text{FOC}(U)$ be a formula of finite out-degree. There exist constants d, k such that*

$$|N_\varphi(U, n)| \leq 2^{2^{dn+k}}, \quad \text{for all } n < \omega.$$

Proof. Fix an injective presentation π of \mathfrak{A} over a binary alphabet and set

$$k := \max \{ |\text{dom}(c)| \mid c \in U \}.$$

Since φ has finite out-degree, we can use Lemma 8.5 (ii), to find a constant d such that

$$|\text{dom}(a)| \leq d \cdot \max \{ |\text{dom}(c)| \mid c \in N(U, n) \},$$

for all $a \in N(U, n+1)$. By induction on n , it follows that

$$|\text{dom}(a)| \leq d^n k, \quad \text{for all } a \in N(U, n).$$

Since there are at most $\frac{1}{m+1} \binom{2m}{m} \leq 4^m$ unlabelled binary trees with a domain of size at most m (these are the Catalan numbers) and, therefore, at most $2^m \cdot 4^m = 6^m$ $\{0, 1\}$ -labelled binary trees, it follows that

$$|N(U, n)| \leq 6^{d^n k} = 2^{cd^n} = 2^{2^{n \log_2 d + \log_2 c}}, \quad \text{for } c := k \cdot \log_2 6. \quad \square$$

Note that the bound in the preceding proposition is rather weak. For instance, the sets N_f grow doubly exponential for every function f . Hence, N_R is only really useful for relations R .

Decomposition Arguments

A powerful argument of a different kind is based on decompositions of a given structure.

Definition 8.21. Let \mathfrak{A} be a Γ -structure and \mathcal{S} a class of Γ -structures. A *sum-decomposition* of \mathfrak{A} over \mathcal{S} consists in a finite partition $(C_i)_{i < n}$ of A such that, for every class C_i , the substructure $\mathfrak{A}|_{C_i}$ is isomorphic to some structure in \mathcal{S} . \lrcorner

Proposition 8.22. *Let \mathfrak{A} be automatic. For every formula $\varphi(x; \bar{y})$ there exists a finite class \mathcal{S} of structures such that, for every tuple $\bar{c} \in A^n$ of parameters, the substructure of \mathfrak{A} induced by $\varphi(x; \bar{c})^{\mathfrak{A}}$ has a sum-decomposition over \mathcal{S} .*

Proof. Fix an injective presentation of \mathfrak{A} over the alphabet Σ and let Γ be the signature of \mathfrak{A} . Fix homomorphisms κ and η , and ψ_R , for $R \in \Gamma$, recognising respectively, L_* , $\varphi^{\mathfrak{A}}$, and L_R . For $\bar{c} \in A^n$, and $u \in \Sigma^*$ with $|u| \geq \|\bar{c}\|$, define

$$f(u; \bar{c}) := \langle \kappa(u), \eta(u \otimes \bar{c}^{\otimes}), (\psi_R(u \otimes \cdots \otimes u))_R \rangle.$$

We claim that the isomorphism type of

$$\varphi(x; \bar{c})^{\mathfrak{A}} \cap u\Sigma^*$$

only depends on $f(u; \bar{c})$. Then we can find a finite sum-decomposition of $\varphi(x; \bar{c})^{\mathfrak{A}}$ into the sets

$$\{a\} \quad \text{and} \quad \varphi(x; \bar{c})^{\mathfrak{A}} \cap u\Sigma^*, \quad \text{for } \|a\| < \|\bar{c}\| \text{ and } u \in \Sigma^{\|\bar{c}\|}.$$

It remains to prove the claim. Suppose that $f(u, \bar{c}) = f(v, \bar{d})$, and let $\tau : u\Sigma^* \rightarrow v\Sigma^*$ be the bijection defined by $\tau(uw) := vw$. Then $\kappa(u) = \kappa(v)$ implies that τ restricts to a bijection $L_\delta \cap u\Sigma^* \rightarrow L_\delta \cap v\Sigma^*$. Similarly, $\psi_R(u \otimes \cdots \otimes u) = \psi_R(v \otimes \cdots \otimes v)$ implies that

$$\bar{b} \in R \quad \text{iff} \quad \tau(\bar{b}) \in R, \quad \text{for all } \bar{b} \text{ in } u\Sigma^*.$$

By the same argument, $\eta(u \otimes \bar{c}^\otimes) = \eta(v \otimes \bar{c}^\otimes)$ implies that τ maps the set $\varphi(x; \bar{c})^\mathfrak{A} \cap u\Sigma^*$ to $\varphi(x; \bar{c})^\mathfrak{A} \cap v\Sigma^*$. Thus, τ is an isomorphism between the substructures of \mathfrak{A} induced by the sets $\varphi(x; \bar{c})^\mathfrak{A} \cap u\Sigma^*$ and $\varphi(x; \bar{c})^\mathfrak{A} \cap v\Sigma^*$. \square

The main drawback of this result is that we do not have any information about how the parts of the decomposed structure look like. Consequently, its applicability is restricted to classes of structures where all induced substructures are sufficiently simple. A typical example is the class of linear orders.

Exercise 8.5. Let α be an ordinal and let $(C_i)_{i < m}$ be a sum-decomposition of $\langle \omega^\alpha, \leq \rangle$. Prove that there is one component $\langle C_i, \leq \rangle$ that is isomorphic to $\langle \omega^\alpha, \leq \rangle$. \lrcorner

Proposition 8.23. *The ordinal $\langle \omega^\omega, \leq \rangle$ is not automatic.*

Proof. For a contradiction, suppose that ω^ω is automatic. Applying Proposition 8.22 to the formula $\varphi(x; z) := x < z$, we then obtain a finite set \mathcal{S} of linear orders such that every subset of the form

$$\downarrow\beta := \{i < \alpha \mid i < \beta\}, \quad \text{for } \beta < \alpha,$$

can be partitioned into finitely many orders from \mathcal{S} . Since \mathcal{S} is finite, there is some number $n < \omega$, such that $\omega^n \notin \mathcal{S}$. Hence, there must exist a finite partition of $\downarrow\omega^n$ into classes whose order type is different from ω^n . This contradicts the above exercise. \square

For tree automatic structures, the above results take the following form.

Definition 8.24. Let \mathcal{S} be a class of Γ -structures.

(a) We denote by $\text{prod}(\mathcal{S})$ the class of all finite direct products of structures in \mathcal{S} .

(b) A Γ -structure \mathfrak{A} is a *superposition* of Γ -structures $\mathfrak{C}_0, \dots, \mathfrak{C}_{n-1}$ if there exist bijections $\sigma_i : C_i \rightarrow A$, for $i < n$, such that

$$R^\mathfrak{A} = \bigcup_{i < n} \sigma_i[R^{\mathfrak{C}_i}], \quad \text{for every relation } R \in \Gamma.$$

In this case we write $\mathfrak{A} = \mathfrak{C}_0 \cup \cdots \cup \mathfrak{C}_{n-1}$.

For a class \mathcal{S} of Γ -structures, we denote by $\text{supp}(\mathcal{S})$ the class of all superpositions of finitely many structures from \mathcal{S} . \lrcorner

Theorem 8.25. *Let \mathfrak{A} be a tree automatic Γ -structure. For every $\text{FOC}(\mathbf{U})$ -formula $\varphi(x; \bar{y})$ there exists a finite class \mathcal{S} of structures such that, for every tuple $\bar{c} \in A^n$ of parameters, the substructure of \mathfrak{A} induced by $\varphi(x; \bar{c})^\mathfrak{A}$ has a sum-decomposition over $\text{supp}(\text{prod}(\mathcal{S}))$.*

Proof. Fix an injective presentation of \mathfrak{A} , let \mathcal{A}_δ and \mathcal{A}_R , $R \in \Gamma$, be the corresponding automata, and let \mathcal{A}_φ be the automaton recognising $\varphi^\mathfrak{A}$. We assume that these automata are deterministic bottom-up automata. We denote by $\delta_x(t)$ the (unique) state from which the automaton \mathcal{A}_x accepts the tree t . Similarly, given a tree t with several holes and a tuple of states \bar{q} , we write $\delta_x(t, \bar{q})$ be the state from which \mathcal{A}_x accepts t when the holes are labelled by the states \bar{q} .

Fix a tuple \bar{c} of parameters. We denote by $D := \bigcup_i \text{dom}(c_i)$ the union of the domains and by F the set of minimal elements of $[2]^* \setminus D$ (with respect to the prefix order). We define the type $\tau(t)$ of a finite tree t as the triple $\tau(t) := \langle s, \lambda \rangle$ where

- ♦ s is the restriction of t to the domain $\text{dom}(t) \cap D$ and
- ♦ λ is the function mapping every vertex $v \in \text{dom}(t) \cap F$ to the state

$$\lambda(v) := \delta_\varphi(t|_v \otimes \square \otimes \cdots \otimes \square),$$

where \square denotes the empty tree.

We can partition $\varphi(x; \bar{c})^{\mathfrak{A}}$ into the sets

$$C_{s\lambda} := \{ t \in \varphi(x; \bar{c})^{\mathfrak{A}} \mid \tau(t) = (s, \lambda) \},$$

for all possible values of s and λ . Note that these sets are disjoint and there are only finitely many of them. To conclude the proof it is therefore sufficient to find a finite set \mathcal{S} such that

$$\mathfrak{C}_{s\lambda} \in \text{supp}(\text{prod}(\mathcal{S})), \quad \text{for all } s, \lambda,$$

where $\mathfrak{C}_{s\lambda}$ is the substructure of \mathfrak{A} induced by $C_{s\lambda}$.

For a state p of \mathcal{A}_φ and states $\bar{q} = (q_R)_{R \in \Gamma}$ of \mathcal{A}_R , let

$$S_{p, \bar{q}} := \{ t \mid \delta_\varphi(t \otimes \square^n) = p \}.$$

We turn $S_{p, \bar{q}}$ into a Γ -structure $\mathfrak{S}_{p, \bar{q}}$ with relations

$$R_{p, \bar{q}} := \{ \bar{t} \in (S_{p, \bar{q}})^k \mid \delta_R(\bar{t}^\otimes) = q_R \},$$

where k is the arity of $R \in \Gamma$. We claim that the class \mathcal{S} of all structures $\mathfrak{S}_{p, \bar{q}}$ has the desired properties. For a subset $H \subseteq F$, let $M(H)$ be the set of all functions μ mapping each $v \in H$ to some tuple $\mu(v) = (q_R)_{R \in \Gamma}$ where q_R is a state of \mathcal{A}_R . For $\mu \in M(H)$ and $R \in \Gamma$, we denote by μ_R the function mapping $v \in H$ to the R -component of $\mu(v)$. We claim that

$$\mathfrak{C}_{s\lambda} \cong \bigcup_{\mu} \prod_{v \in \text{dom}(\lambda)} \mathfrak{S}_{\lambda(v), \mu(v)},$$

where the union ranges over all functions $\mu \in M(\text{dom}(\lambda))$ such that

$$\delta_R(s \otimes \cdots \otimes s, \mu_R) \quad \text{is an accepting state.}$$

Let σ be the function mapping a tree $t \in C_{s\lambda}$ to the tuple $(t|_v)_{v \in \text{dom}(\lambda)}$. To see that σ is the desired isomorphism, first note that

$$\delta_\varphi(t|_v \otimes \square^n) = \lambda(v) \quad \text{implies} \quad t|_v \in S_{\lambda(v), \mu(v)}, \quad \text{for all } \mu.$$

Hence, $\sigma(t) \in \prod_v S_{\lambda(v), \mu(v)}$. Furthermore, σ has an inverse that maps a tuple $(t_v)_v$ to the tree obtain from s by attaching t_v as a subtree to the vertex v .

Finally, note that $\bar{t} \in R^{\mathfrak{C}_{s\lambda}}$ implies $\sigma(\bar{t}) \in R^{\prod_v \mathfrak{S}_{\lambda(v), \mu(v)}}$ where

$$\mu_R(v) := \delta_R(t_v|_v \otimes \cdots \otimes t_{k-1}|_v).$$

Furthermore, this function μ is part of the above union since $\bar{t} \in R^{\mathfrak{C}_{s\lambda}}$ implies that the state $\delta_R(s \otimes \cdots \otimes s, \mu_R)$ is accepting.

Conversely, $\sigma(\bar{t}) \in R^{\prod_v \mathfrak{S}_{\lambda(v), \mu(v)}}$ implies that $\bar{t} \in R^{\mathfrak{C}_{s\lambda}}$. Since

$$\delta_R(\bar{t}) = \delta_R(s \otimes \cdots \otimes s, \mu_R),$$

which is accepting. □

An application of this theorem can be found in Section II.

ω -Automatic Structures

Most of the tools we have developed for (tree) automatic structures do not work for ω -automatic (or ω -tree automatic) ones for the simple reason that we have no analogue to the length of an element. The exception is the bound on the expansion of a relation, which is based on a simple counting argument. Unsurprisingly, the proofs for the ω -automatic case turn out to be much more involved. We start with the following refinement of the relation \approx_* .

Definition 8.26. For $m < \omega$ and $u, v \in \Sigma^\omega$, we define

$$u \approx_*^m v \quad : \text{iff} \quad u[m, \infty) = v[m, \infty).$$

We denote the \approx_*^m -class of u by $[u]_*^m$. J

Recall the notion of a set being E -covered by a set of \approx_* -classes we introduced in Definition 6.4.

Lemma 8.27. Let \mathfrak{A} be an ω -automatic structure and $f : A^{k+1} \rightarrow A$ an FOC-definable function. There exists a constant d with the following property. For every $m < \omega$, every $\bar{c} \in A^l$, and every set $S \subseteq A$ that is covered by $n \approx_*^m$ -classes, the image $f(S, \bar{c})$ is covered by a set of at most $dn \approx_*^m$ -classes.

Proof. Let $\eta : (\Sigma^{k+l+1})^\infty \rightarrow \mathfrak{T}$ be a morphism into a finite ω -semigroup recognising the graph of f and set $d := |T|$. Every set S that is covered by n \approx_*^m -classes can be written as a union $S = \bigcup_{i < n} S_i$, where each S_i is covered by a single class. Since

$$f(S, \bar{c}) = \bigcup_{i < n} f(S_i, \bar{c}),$$

it is therefore sufficient to prove that each image $f(S_i, \bar{c})$ is covered by at most $d \approx_*^m$ -classes. Therefore, we may assume without loss of generality that $n = 1$.

Fix an ω -word v representing \bar{c} , let D be the \approx_*^m -class covering S , and, for each $s \in S$, choose an ω -word $u(s) \in D$ representing s and set

$$U := \{u(s) \mid s \in S\}$$

and $u(\bar{s}) := \langle u(s_0), \dots, u(s_{k-1}) \rangle$, for $\bar{s} \in S^k$.

Finally, let $w(\bar{s})$ be an ω -word representing $f(\bar{s}, \bar{c})$ and set

$$\alpha(\bar{s}) := \eta((u(\bar{s}) \otimes v \otimes w(\bar{s}))[0, m]).$$

We claim that

$$\alpha(\bar{s}) = \alpha(\bar{s}') \quad \text{implies} \quad [f(\bar{s}', \bar{c})]_{=} \cap [w(\bar{s})]_*^m \neq \emptyset,$$

where $[a]_{=}$ denotes the set of ω -words encoding a and $[w]_*^m$ is the \approx_*^m -class of w . Then it follows that, picking one tuple $\bar{s}_t \in \alpha^{-1}(t)$, for every $t \in T$, we obtain a set

$$W := \{[w(\bar{s}_t)]_*^m \mid t \in T\}$$

of at most $|T| = d \approx_*^m$ -classes covering $f(S^k, \bar{c})$. Hence, it remains to prove the claim. Note that

$$u(\bar{s}), u(\bar{s}') \in D \quad \text{implies} \quad u(\bar{s}) \otimes v \approx_*^m u(\bar{s}') \otimes v,$$

for all $\bar{s}, \bar{s}' \in S^k$. Consequently, $\alpha(\bar{s}) = \alpha(\bar{s}')$ implies that

$$\begin{aligned} & \eta(u(\bar{s}) \otimes v \otimes w(\bar{s})) \\ &= \eta((u(\bar{s}) \otimes v \otimes w(\bar{s}))[0, m]) \cdot \eta((u(\bar{s}) \otimes v \otimes w(\bar{s}))[m, \infty)) \\ &= \eta((u(\bar{s}') \otimes v \otimes w(\bar{s}'))[0, m]) \cdot \eta((u(\bar{s}') \otimes v \otimes w(\bar{s}'))[m, \infty)) \\ &= \eta(u(\bar{s}') \otimes v \otimes (w(\bar{s}') [0, m] \cdot w(\bar{s}) [m, \infty))) \end{aligned}$$

Hence,

$$w(\bar{s}') [0, m] \cdot w(\bar{s}) [m, \infty) \in [f(\bar{s}', \bar{c})]_{=} \cap [w(\bar{s})]_*^m \neq \emptyset. \quad \square$$

Lemma 8.28. *Let \mathfrak{A} be an ω -automatic structure with presentation π and let $U \subseteq A$ be an infinite regular subset such that $\pi \upharpoonright U$ is injective. There exists an infinite set $S \subseteq U$ that is covered by a single \approx_* -class.*

Proof. Since U is infinite, we can use Lemma 6.3 to find a \approx_* -class $[w]_*$ such that $U \cap [w]_*$ is infinite. Let $S \subseteq U$ be the set of all elements represented by a word in $U \cap [w]_*$. As the presentation is injective on U , it follows that S is also infinite. \square

We start by generalising Proposition 8.14 to ω -automatic structures with injective presentations.

Lemma 8.29. *Let \mathfrak{A} be an ω -automatic structure. If there exists an infinite subset $S \subseteq A$ covered by a single \approx_* -class, then every FOC-definable function has bounded expansion.*

Proof. Suppose that $S \subseteq A$ is infinite and covered by the \approx_* -class $[w]_*$. For a contradiction, suppose that there exists an FOC-definable function $f : A^k \rightarrow A$ such that ex_f grows super-linearly. Let d be the constant from Lemma 8.27 and choose n such that

$$\text{ex}_f(n) > 2dn.$$

Since $[w]_* = \bigcup_{m < \omega} [w]_m^*$ and each \approx_*^m -class has size at most 2^m , it follows that there exists some number $m < \omega$ such that some \approx_*^m -class contains

representatives of at least n different elements. We choose m minimal. Let $D \subseteq A$ be a set of maximal size that is covered by some \approx_*^m -class. By choice of m , we have $|D| \geq n$. Furthermore, since D is covered by at most $2 \approx_*^{m-1}$ -classes, it follows by minimality of m that

$$|D| \leq 2n.$$

By Lemma 8.27, the image $f(D^k)$ is covered by at most $d \approx_*^m$ -classes. One of these classes contains representatives of at least

$$\frac{f(D^k)}{d} > \frac{2dn}{d} = 2n \geq |D|$$

elements of $f(D^k)$. A contradiction to the choice of D . \square

Corollary 8.30. *Let \mathfrak{A} be an infinite structure with an injective ω -automatic presentation. Every FOC-definable function has bounded expansion.*

Proof. As the presentation is injective, we can use Lemma 8.28 to find an infinite set $S \subseteq A$ that is covered by a single \approx_* -class. Hence, the claim follows by Lemma 8.29 \square

We were only able to prove the preceding corollary for injective presentations. Next, we will strengthen this result to ω -automatic structures (not necessarily injective) that admit an FOC-definable linear order.

Theorem 8.31. *Let \mathfrak{A} be an uncountable ω -automatic structure with an FOC-definable linear order \leq . For every presentation π of \mathfrak{A} , there exists a regular subset $U \subseteq A$ such that $\langle U, \leq \rangle \cong \langle [2]^\omega, \leq_{\text{lex}} \rangle$ and the restriction of π to U is injective.*

Proof. The claim follows immediately from Theorem 6.11, except for a minor technical issue: the languages $L_=\$ and L_\leq from the presentation are only defined on L_δ , and not on all of Σ^ω . But this can easily be corrected by adding $\Sigma^\omega \setminus L_\delta$ as a new equivalence class to $L_=\$ and by making this class the least element of L_\leq . Since these operations preserve regularity, we can now apply Theorem 6.11. \square

As a consequence of this theorem, we can generalise Corollary 8.30 to ordered ω -automatic structures.

Proposition 8.32. *Let \mathfrak{A} be an ω -automatic structure with an FOC-definable linear order \leq . Every FOC-definable function has bounded expansion.*

Proof. By Lemma 8.29 it is sufficient to find an infinite subset $S \subseteq A$ that is covered by a single \approx_* -class.

If \mathfrak{A} is countable, we can use Theorem 6.14 to construct an injective presentation of \mathfrak{A} . Consequently, the existence of S follows by Lemma 8.28.

Hence, we may assume that \mathfrak{A} is uncountable. By Theorem 8.31, there exists a regular subset $U \subseteq A$ such that $\langle U, \leq \rangle \cong \langle [2]^\omega, \leq_{\text{lex}} \rangle$ and the restriction of the presentation to U is injective. Hence, we can use Lemma 8.28 to find an infinite subset $S \subseteq U$ that is covered by a single \approx_* -class. \square

A second method to prove the non-existence of an ω -automatic presentation is based on Theorem 6.13.

Lemma 8.33. *Let \mathfrak{A} be an uncountable ω -automatic structure. There exists an integer $n > 0$ such that no FOC-definable function $f : A^n \rightarrow A$ is injective.*

Proof. Let k be the constant from Theorem 6.13 and set $n := k + 1$. By choice of k , it follows that no FOC-definable function $A^{k+1} \rightarrow A$ is injective. \square

Proposition 8.34. *No ω -automatic structure has an FOC-definable pairing function.*

Proof. For a contradiction, suppose that there exists an ω -automatic structure \mathfrak{A} with an FOC-definable pairing function $f : A \times A \rightarrow A$. If \mathfrak{A} is countable, it is automatic and the claim follows by Proposition 8.14 since $\text{ex}_f(n) = n^2$. Consequently, we may assume that \mathfrak{A} is uncountable. Consider the functions f_i defined by

$$f_0(x, y) := f(x, y) \quad \text{and} \quad f_{i+1}(\tilde{x}, \tilde{y}) := f(f_i(\tilde{x}), f_i(\tilde{y})).$$

By induction on i , it follows that every f_i is injective and FOC-definable. This contradicts Lemma 8.33. \square

9 Automatic Groups

We have already seen several examples of which groups are and are not automatic. Let us take a look at this question more systematically. Before doing so, a remark is in order about terminology.

Remark. (a) Before automatic structures were introduced, group theorists had already started studying groups represented by automata, which they called ‘automatic groups.’ Unfortunately, their definition is slightly different from ours. A group \mathfrak{G} is automatic in the group-theoretic sense if it has a finite set S of generators and there exists a regular language $L \subseteq S^*$ such that the relations

$$M_c := \{ u \oplus w \mid \pi(u) = \pi(w) \cdot c \}, \quad \text{for } c \in S \cup \{1\},$$

are regular, where $\pi : S^* \rightarrow G$ is group multiplication. It follows that every automatic group has an automatic presentation in our sense, but one where the encoding map $\pi : L_\delta \rightarrow G$ has a prescribed form. Furthermore, this presentation is not one of the structure $\langle G, \cdot \rangle$, but one of the *Cayley graph* of \mathfrak{G} , i.e., of the structure $\langle G, (\mu_c)_{c \in S} \rangle$ where $\mu_c(g) = gc$ is multiplication by the generator c .

Finally, let us mention that there are examples of groups whose Cayley graph is automatic in our sense, but not in the group-theoretic ones. One such example is the *discrete Heisenberg group* $\mathfrak{H} = \langle H, \cdot \rangle$, which is the multiplicative group of all upper triangular matrices

$$\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} \in \mathbb{Z}^{3 \times 3}.$$

The group is generated by the two matrices

$$\alpha := \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \beta := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

One can prove that \mathfrak{H} is not automatic in the group-theoretic sense, but it is straightforward to give a 3-dimensional interpretation of the Cayley graph $\langle H, \mu_\alpha, \mu_\beta, \mu_{\alpha^{-1}}, \mu_{\beta^{-1}} \rangle$ in $\langle \mathbb{Z}, + \rangle$.

(b) There is another related concept used in group-theory, which is that of an *automaton group*. One way to define such groups is as follows. Let $\mathfrak{A} = \langle A, f_0, \dots, f_{n-1} \rangle$ be an automatic structure with finitely many unary functions f_0, \dots, f_{n-1} that are assumed to be bijective. Then the subgroup of $\text{Sym}(A)$ generated by f_0, \dots, f_{n-1} is an automaton group. Clearly, if the Cayley graph of a group \mathfrak{G} is automatic in our sense, then \mathfrak{G} is an automaton group. Conversely, if \mathfrak{G} is the automaton group presented by the automatic structure $\mathfrak{A} = \langle A, f_0, \dots, f_{n-1} \rangle$ and if the action of the functions f_0, \dots, f_{n-1} on the set A is transitive, then \mathfrak{A} is isomorphic to the Cayley graph of \mathfrak{G} . In particular, \mathfrak{G} has a Cayley graph that is automatic in our sense. J

In the remainder of this chapter we will use the term ‘automatic group’ for a group $\mathfrak{G} = \langle G, \cdot \rangle$ that has an automatic presentation in our sense. We have already seen a few examples of such groups. In particular, the following groups are automatic:

- ♦ every finite group;
- ♦ the additive groups \mathbb{Z} , $\mathbb{Z}[1/k]$, and $\mathbb{Z}[1/k]/\mathbb{Z}$, for $k \geq 2$;
- ♦ groups of the form $\mathfrak{G}^{(\omega)}$, for finite \mathfrak{G} ;
- ♦ finite sums of such groups.

Finately Generated Groups

It is possible to characterise all finitely generated automatic groups via the following algebraic property.

Definition 9.1. Let \mathfrak{G} be a group and P a property of groups.

- (a) The *index* $[\mathfrak{G} : \mathfrak{H}]$ of a subgroup $\mathfrak{H} \subseteq \mathfrak{G}$ is the number $|\mathfrak{G}/\mathfrak{H}|$ of cosets.
- (b) We say that \mathfrak{G} is *virtually abelian* if it has an abelian subgroup of finite index. J

Theorem 9.2. *Let \mathfrak{G} be an automatic group. Every finitely generated subgroup of \mathfrak{G} is virtually abelian.*

Corollary 9.3. *Let \mathfrak{G} be a finitely generated group. The following statements are equivalent.*

- (1) \mathfrak{G} is automatic.
- (2) \mathfrak{G} is ω -automatic.
- (3) \mathfrak{G} is virtually abelian.
- (4) There exists an FO-interpretation of \mathfrak{G} in $\langle \mathbb{N}, + \rangle$.

The rest of this section is devoted to a proof of these two results. Let us start with the observation that, in a finitely generated virtually abelian group, we can always choose the abelian subgroup to be normal.

Lemma 9.4. *Every finitely generated virtually abelian group \mathfrak{G} has a normal abelian subgroup of finite index.*

Proof. Fix an abelian subgroup $\mathfrak{A} \subseteq \mathfrak{G}$ of finite index. Since \mathfrak{G} is finitely generated, so is \mathfrak{A} . Hence, $\mathfrak{A} \cong \mathbb{Z}^m \times \mathfrak{C}$, for some finite group \mathfrak{C} and $m < \omega$. Let $\mathfrak{A}_0 \subseteq \mathfrak{A}$ be the subgroup isomorphic to \mathbb{Z}^m . Then \mathfrak{A}_0 has finite index in \mathfrak{G} . By Lemma XI.6.16, there exists a normal subgroup $\mathfrak{N} \subseteq \mathfrak{G}$ of finite index with $N \subseteq \mathfrak{A}_0$. Since \mathfrak{A}_0 is free, so is \mathfrak{N} . In particular, \mathfrak{N} is abelian. \square

This allows us already to prove one half of Corollary 9.3.

Lemma 9.5. *Every finitely generated virtually abelian group can be FO-interpreted in $\langle \mathbb{N}, + \rangle$.*

Proof. Let \mathfrak{G} be a finitely generated group and $\mathfrak{A} \subseteq \mathfrak{G}$ an abelian subgroup such that $\mathfrak{G}/\mathfrak{A}$ is finite. By Lemma 9.4, we may assume that \mathfrak{A} is normal. Then \mathfrak{A} is finitely generated and therefore of the form $\mathbb{Z}^m \times \mathfrak{C}$, for some finite group \mathfrak{C} . Since \mathbb{Z}^m also has finite index in \mathfrak{G} , we can replace \mathfrak{A} by \mathbb{Z}^m and assume that $\mathfrak{A} \cong \mathbb{Z}^m$. Let $g_0A, \dots, g_{n-1}A$ be the elements of $\mathfrak{G}/\mathfrak{A}$. Let $\mu : [n]^2 \rightarrow [n]$ be the function such that

$$g_iA \cdot g_jA = g_{\mu(i,j)}A,$$

and choose elements $a_{ij} \in A$, for $i, j < n$, such that

$$g_i \cdot g_j = g_{\mu(i,j)} a_{ij}.$$

For $b, c \in A$ it then follows that

$$g_i b \cdot g_j c = g_i g_j \cdot g_j^{-1} b g_j \cdot c = g_{\mu(i,j)} a_{ij} \cdot g_j^{-1} b g_j \cdot c.$$

Note that, for every j , the function $f_j : \mathbb{Z}^m \rightarrow \mathbb{Z}^m$ mapping b to $f_j(b) := g_j^{-1} b g_j$ forms an automorphism of \mathbb{Z}^m . Consequently, it takes the form $f_j(x) = M_j x$, for some matrix $M_j \in \text{GL}_m(\mathbb{Z})$. This implies that f_j is definable when viewed as an $2m$ -ary relation on $\langle \mathbb{Z}, + \rangle$. Thus, we obtain the following formula for multiplication in \mathfrak{G} .

$$g_i b \cdot g_j c = g_{\mu(i,j)} (a_{ij} \cdot f_j(b) \cdot c).$$

Representing an element $g_i b$ by the pair $\langle i, b \rangle \in [n] \times \mathbb{Z}^m \subseteq \mathbb{Z}^{m+1}$ we can therefore define multiplication by

$$\langle i, b \rangle \cdot \langle j, c \rangle = \langle \mu(i, j), a_{ij} \cdot f_j(b) \cdot c \rangle.$$

Since this is first-order definable in $\langle \mathbb{Z}, + \rangle$, we obtain an $m+1$ -dimensional FO-interpretation of \mathfrak{G} in $\langle \mathbb{Z}, + \rangle$. As the latter can be interpreted in $\langle \mathbb{N}, + \rangle$, so can the former. \square

To prove the converse, we need some group theory. We start by taking a look at the *discrete Heisenberg group*, which we have already encountered above. Recall that it consists of all invertible upper triangular matrices in $\mathbb{Z}^{3 \times 3}$.

Definition 9.6. Let \mathfrak{G} be a group. The *commutator* of two elements $a, b \in \mathfrak{G}$ is

$$[a, b] := a^{-1} b^{-1} a b.$$

We can extend this notation to more elements by

$$[a_0, \dots, a_n] := [[a_0, \dots, a_{n-1}], a_n].$$

For $H, K \subseteq G$, we set

$$[H, K] := \{ [h, k] \mid h \in H, k \in K \}.$$

The following equation follows immediately from the definition by a direct calculation.

Lemma 9.7. $[ab, c] = b^{-1}[a, c]b \cdot [b, c]$ and $[a^{-1}, b] = a[b, a]a^{-1}$.

We can use the Heisenberg group as a criterion for non-automaticity.

Proposition 9.8. *The Heisenberg group is generated by three elements a, b, g satisfying*

$$g = [a, b], \quad [a, g] = e, \quad \text{and} \quad [b, g] = e.$$

Furthermore, the element g has infinite order.

Proof. Let $I \in \mathbb{Z}^{3 \times 3}$ be the identity matrix and

$$A := \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad B := \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad C := \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Note that

$$A^2 = B^2 = C^2 = AC = BC = CA = CB = BA = 0 \quad \text{and} \quad AB = C.$$

We claim that the elements

$$a := I + A, \quad b := I + B, \quad \text{and} \quad g := I + C$$

have the desired properties. We start by computing their powers. For $k > 1$, we have

$$a^k = (I + A)^k = I + kA + A^2(\dots) = I + kA.$$

Hence,

$$(I - kA)(I + kA) = I - k^2A^2 = I \quad \text{implies} \quad a^{-k} = I - kA.$$

In the same way, we obtain analogous expressions for b^k and g^k . Thus,

$$a^k = I + kA, \quad b^k = I + kB, \quad g^k = I + kC, \quad \text{for all } k \in \mathbb{Z}.$$

In particular, g has infinite order. Furthermore,

$$\begin{aligned} a^i b^j c^k &= (I + iA)(I + jB)(I + kC) \\ &= I + iA + jB + kC + ijAB + ikAC + jkBC + ijkABC \\ &= I + iA + jB + kC \end{aligned}$$

implies that every group element can be uniquely written in the form $a^i b^j c^k$, for some $i, j, k \in \mathbb{Z}$. It remains to check the commutator relations.

$$\begin{aligned} [a, b] &= (I - A)(I - B)(I + A)(I + B) \\ &= (I - A - B + C)(I + A + B + C) \\ &= (I + C)^2 - (A + B)(I + C) + (I + C)(A + B) - (A + B)^2 \\ &= (I + 2C + C^2) - (AC + BC) + (CA + CB) - (A^2 + C + BA + B^2) \\ &= I + C \\ &= g, \\ [a, g] &= (I - A)(I - C)(I + A)(I + C) \\ &= (I - A - C + AC)(I + A + C + AC) \\ &= I + (A + C) - (A + C) - (A + C)(A + C) \\ &= I - A^2 - AC - CA - C^2 \\ &= I, \end{aligned}$$

and similarly for $[b, g]$. □

Corollary 9.9. *There exists an embedding of the Heisenberg group into \mathfrak{S} if and only if, \mathfrak{S} contains elements $a, b, g \in G$ such that*

$$g = [a, b], \quad [a, g] = e, \quad [b, g] = e, \quad \text{and} \quad g \text{ has infinite order.}$$

Proposition 9.10. *Let \mathfrak{G} be a group. If there exists an embedding of the Heisenberg group into \mathfrak{G} , then \mathfrak{G} is not automatic.*

Proof. Suppose that there exists an embedding of the Heisenberg group into \mathfrak{G} . Then we can find three elements $a, b, g \in G$ such that

$$g = [a, b], \quad [a, g] = e, \quad \text{and} \quad [b, g] = e,$$

and g has infinite order. Below we will prove the following statements where $0 < i, k < \omega$ and $c \in G$.

- (a) $[a, c] = g^k$ implies $[a^i, c] = g^{ki}$
- (b) $[a^i, b] = g^i, \quad [a^{-i}, b] = g^{-i}, \quad [a, b^k] = g^k, \quad [a^i, b^k] = g^{ik}.$
- (c) There exists an FO-formula $\psi(x, y, z; u, v)$ such that

$$\mathfrak{G} \models \psi(g^i, g^k, c; a, b) \quad \text{iff} \quad c = g^{ik}, \quad \text{for all } c \in G.$$

Then the fact that \mathfrak{G} is not automatic can be established as follows. Let $R := \psi(x, y, z; a, b)^{\mathfrak{G}}$ be the relation defined by the formula from (c). By Proposition 8.9, it is sufficient to find an embedding of \mathbb{N}^n into $\langle G, R \rangle$, for every $n < \omega$. Hence, fix n . Let p_0, \dots, p_{n-1} be n distinct prime numbers and set

$$P := \{ p_0^{k_0} \dots p_{n-1}^{k_{n-1}} \mid k_0, \dots, k_{n-1} < \omega \}.$$

Then $\langle P, \cdot \rangle \cong \mathbb{N}^n$ and the function $f : p \mapsto g^p$ satisfies

$$\langle f(p), f(q), h \rangle \in R \quad \text{iff} \quad h = f(pg), \quad \text{for all } p, q \in P, \quad h \in G,$$

as desired. It therefore remains to prove the above claims.

(a) We prove the claim by induction on i . For $i = 1$ the claim is trivial. For the inductive step, suppose that $[a^i, b] = g^{ki}$. Since $[a, g] = e$, we have

$$[a^{i+1}, b] = a^{-1}[a^i, b]a \cdot [a, b] = a^{-1}g^{ki}a \cdot g^k = g^{k(i+1)}.$$

(b) The first equation follows by (a) for $h = b$ and $k = 1$. For the second one, note that $[a, g] = e$ implies

$$\begin{aligned} e &= [a^i, b] \cdot [a^{-i}, b] = a^i[a^i, b]a^{-i} \cdot [a^{-i}, b] \\ &= a^i g^i a^{-i} \cdot [a^{-i}, b] = g^i \cdot [a^{-i}, b]. \end{aligned}$$

For the third one, note that $[b, a] = [a, b]^{-1} = g^{-1}$. Hence, applying (a) to $h = a$ and $k = 1$ we obtain $[b^k, a] = g^{-k}$, which implies that $[a, b^k] = [b^k, a]^{-1} = g^k$.

Finally, the last equation follows by the third equation using (a) for $h = b^k$.

(c) We define

$$\begin{aligned} \psi(x, y, z; a, b) &:= [a, x] = e \wedge [b, y] = e \\ &\quad \wedge (\exists u, v \in C(g)) [C(b) \subseteq C(v) \wedge z = [u, v] \\ &\quad \quad \quad \wedge x = [u, b] \wedge y = [a, v]], \end{aligned}$$

where $C(x) = \{ y \mid [x, y] = e \}$ and $g := [a, b]$. We claim that ψ has the desired property.

(\Leftarrow) Given $x = g^i, y = g^k$, and $z = g^{ik}$, we set $u := a^i$ and $v := b^k$. Then $a^i, b^k \in C(g)$ and $C(b) \subseteq C(b^k)$. Furthermore, we have shown in (b) that

$$[a^i, b^k] = g^{ik}, \quad [a^i, b] = g^i, \quad \text{and} \quad [a, b^k] = g^k.$$

(\Rightarrow) Suppose that $\eta(g^i, g^k, z; a, b)$ holds and let u and v the the corresponding witnesses.

Note that, by (b),

$$[u, b] \cdot [a^{-i}, b] = [u, b] \cdot g^{-i} = g^i \cdot g^{-1} = e.$$

Hence, $[u, b] \in C(a)$ implies that

$$a^i[u, b]a^{-i} \cdot [a^{-i}, b] = [ua^{-i}, b] = [u, b] \cdot [a^{-i}, b] = e.$$

Setting $c := ua^{-i}$, it follows that $s \in C(b) \subseteq C(v)$ and

$$[u, v] = [sa^i, v] = a^{-i}[s, v]a^i \cdot [a^i, v] = a^{-i}a^i \cdot [a^i, v] = [a^i, v].$$

Finally, by (a), we know that $[a, v] = g^k$ implies $[a^i, v] = g^{ik}$. Consequently,

$$z = [u, v] = [a^i, v] = g^{ik}. \quad \square$$

Finally, we need some results about nilpotent groups.

Definition 9.11. Let \mathfrak{G} be a group.

(a) Let $C \subseteq G$ be a finite set of generators such that $c \in C$ implies $c^{-1} \in C$. We say that \mathfrak{G} has *polynomial growth* if there exists a polynomial $p(x)$ such that

$$|\{c_0 \cdots c_{n-1} \mid c_0, \dots, c_{n-1} \in C\}| \leq p(n), \quad \text{for all } n < \omega.$$

(b) \mathfrak{G} is *nilpotent* if there is some $n < \omega$ such that

$$[g_1, \dots, g_n] = e, \quad \text{for all } g_1, \dots, g_n \in G.$$

It is *virtually nilpotent* if it has a nilpotent subgroup of finite index. \dashv

Exercise 9.1. Prove that, if \mathfrak{G} has polynomial growth with respect to some set C of generators, it has polynomial growth with respect to every finite set of generators. \dashv

The following result is more substantial. For the proof we refer to the literature.

Theorem 9.12 (Gromov [92]). *Every finitely generated group of polynomial growth is virtually nilpotent.*

Proposition 9.13. *Every finitely generated nilpotent group \mathfrak{G} has a normal subgroup \mathfrak{H} of finite index that is torsion-free.*

Proof. Let C be a finite set of generators of \mathfrak{G} and set $G_0 := G$ and $G_{i+1} := [G_i, G]$. Note that, since \mathfrak{G} is nilpotent, this sequence stabilises at $G_i = \{e\}$, for some i . We start by proving that G_{i-1} is generated by the set $G_i \cup D_i$, where

$$D_i := \{[c_1, \dots, c_i] \mid c_1, \dots, c_i \in C\} \subseteq G_{i-1}.$$

We proceed by induction on i . The statement is trivial for $i = 1$ since $G_0 = G$ is generated by $C = D_1$. For the inductive step, note that G_{i-1} is generated by elements of the form $[x, y]$ with $x \in G_i$ and $y \in G$. As G_{i-1} is a normal subgroup of \mathfrak{G} , we can use the inductive hypothesis to write

$$x = x_0^{e_0} \cdots x_{n-1}^{e_{n-1}} z,$$

where $e_0, \dots, e_{n-1} \in \{-1, 1\}$, $z \in G_{i-1}$ and $x_0, \dots, x_{n-1} \in D_{i-1}$. Writing z as a product of elements of C and using the equations from Lemma 9.7, it follows that we can write x as a product of elements (and their inverses) of the form

$$g^{-1}[x_j, c]g = [x_j, c][x_j, c, g] \quad \text{and} \quad g^{-1}[z, c]g, \quad \text{for } c \in C, g \in G.$$

Since $[x_j, c, g], [z, c] \in G_i$ and $[x_j, c] \in D_i$ the claim follows.

Having proved the claim, it follows that every subgroup G_i is finitely generated and each quotient $G_i/G_{i+1} = G_i/[G_i, G]$ is abelian. Hence, there exists a finite sequence

$$G_i = H_0^i \supseteq \cdots \supseteq H_{m_i}^i = G_{i+1}$$

of subgroups such that each factor H_j^i/H_{j+1}^i is cyclic and $[H_j^i, G_i]$ is a normal subgroup of H_{j+1}^i . Concatenating these sequences, we obtain a sequence

$$G = H_0 \supseteq \cdots \supseteq H_n = \{e\}$$

such that each factor H_j/H_{j+1} is cyclic and $[H_j, G]$ is a normal subgroup of H_{j+1} . We prove the claim by induction on the length n of this sequence. If $n = 0$, the claim is trivial. Hence, suppose that $n > 0$. Let j be the minimal index with $H_j \neq G$ and set $\mathfrak{H} := H_j$. By assumption, there exists some element $a \in G$ generating $\mathfrak{G}/\mathfrak{H}$. By inductive hypothesis, \mathfrak{H} has a normal subgroup \mathfrak{K} of finite index that is torsion-free. If \mathfrak{H} has finite index in \mathfrak{G} , the group \mathfrak{K} is the desired torsion-free subgroup of \mathfrak{G} . Hence, suppose that the index of \mathfrak{H} in \mathfrak{G} is infinite. Let U be the subgroup of \mathfrak{G} induced by the set

$$U := \langle\langle a \rangle\rangle \cdot K.$$

This group is torsion-free and its index is

$$|G : U| = |G : \langle\langle a \rangle\rangle \cdot K| = |H : H \cap \langle\langle a \rangle\rangle \cdot K| = |H : K|,$$

which is finite by choice of \mathfrak{R} . \square

Lemma 9.14. *Let \mathfrak{G} be a finitely generated nilpotent group. If \mathfrak{G} is not virtually abelian, it has a subgroup isomorphic to the Heisenberg group.*

Proof. By Proposition 9.13, \mathfrak{G} has a torsion-free subgroup \mathfrak{H} of finite index. As \mathfrak{G} is not virtually abelian, it follows that \mathfrak{H} is not abelian. Consider the sequence defined by $H_0 := H$ and $H_{i+1} := [H_i, H]$, for $i < \omega$. Since \mathfrak{H} is nilpotent, there exists some number n such that $H_{n+1} = \{e\}$. Let n be the least such number. As \mathfrak{H} is not abelian, we have $n \geq 1$. Furthermore, H_n induces a non-trivial abelian subgroup of \mathfrak{H} . Fix some element $g \in H_n$ with $g \neq e$. By definition $g = [a, b]$ for some $a \in H_{n-1}$ and $b \in H$. Then $[H_n, H] = H_{n+1} = \{e\}$ implies that $[g, a] = e$ and $[g, b] = e$. Furthermore, since $\mathfrak{H} \subseteq \mathfrak{G}$ is torsion-free, the order of g is infinite. Consequently, a, b , and g determine an embedding of the Heisenberg group into \mathfrak{G} . \square

The connection between automatic groups and nilpotent ones is given by the following lemma.

Lemma 9.15. *Let \mathfrak{H} be a finitely generated subgroup of an automatic group \mathfrak{G} . Then \mathfrak{H} has polynomial growth.*

Proof. Fix a finite set $C \subseteq H$ of generators such that $c \in C$ implies $c^{-1} \in C$, and set $l := \|C\|$. By Lemma 8.8, there exists a constant k such that

$$\|c_0 \cdots c_{n-1}\| \leq l + k \lceil \log_2 n \rceil, \quad \text{for all } c_0, \dots, c_{n-1} \in C \text{ and } n < \omega.$$

Thus the number of elements of H that can be written as a product of at most n generators is bounded by

$$2^{l+k \lceil \log_2 n \rceil} \leq 2^{l+k(\log_2 n+1)} = 2^{l+k} (2^{\log_2 n})^k = 2^{l+k} n^k. \quad \square$$

Using these tools, we can finally proof the theorem and its corollary.

Proof of Theorem 9.2. Let \mathfrak{H} be a finitely generated subgroup of an automatic group \mathfrak{G} . If \mathfrak{H} is finite, there is nothing to prove. Hence, we may assume that it is infinite. By Lemma 9.15, \mathfrak{H} has polynomial growth. Consequently, we can use Theorem 9.12 to find a nilpotent subgroup \mathfrak{H}_0 of \mathfrak{H} that has finite index. If \mathfrak{H}_0 is virtually abelian, so is \mathfrak{H} .

Hence, suppose otherwise. Then we can use Lemma 9.14 to find an embedding of the Heisenberg group in \mathfrak{H} and, therefore, also in \mathfrak{G} . Consequently, it follows by Proposition 9.10 that \mathfrak{G} is not automatic. \square

Proof of Corollary 9.3. (1) \Rightarrow (3) follows by Theorem 9.2; (3) \Rightarrow (4) was proved in Lemma 9.5; (4) \Rightarrow (2) holds since $\langle \mathbb{N}, + \rangle$ is ω -automatic and ω -automatic structures are closed under FO-interpretations; and (2) \Rightarrow (1) follows by Theorem 6.14 and the fact that every finitely generated group is countable. \square

Open Questions.

- (a) Does every finitely generated (subgroup of an) ω -automatic group have polynomial growth?
- (b) Does every finitely generated (subgroup of a) tree automatic group have polynomial growth?

The Rationals

What about groups that are not finitely generated? Here, we have only partial results, even for abelian groups. We have already seen that \mathbb{Z} , $\mathbb{Z}[1/k]$, and $\mathbb{Z}[1/k]/\mathbb{Z}$ are automatic, and so are all finite products of such groups. This is in contrast to the rationals, which are *not* automatic.

Definition 9.16. Let $p > 1$ be an integer. An abelian group \mathfrak{A} is *p-divisible* if, for every $a \in A$, there is some $b \in A$ with $pb = a$. \lrcorner

Theorem 9.17 (Tsankov). *Let P be an infinite set of prime numbers.*

- (a) If \mathfrak{A} is a torsion-free abelian group that is *p-divisible*, for all $p \in P$, then \mathfrak{A} is not automatic.
- (b) $\bigoplus_{p \in P} \mathbb{Z}[1/p]/\mathbb{Z}$ is not automatic.

Corollary 9.18. $\langle \mathbb{Q}, + \rangle$ and $\langle \mathbb{Q}/\mathbb{Z}, + \rangle$ are not automatic.

For the proof of Theorem 9.17, we need some results from additive combinatorics concerning progressions and lattices. We will omit some of the rather lengthy proofs. The missing material can be found in, e.g., Chapters 3 and 5 of [201].


Definition 9.19. (a) For $\bar{n} \in \mathbb{N}^d$, we set

$$\begin{aligned} [\bar{n}] &:= \{ \langle s_0, \dots, s_{d-1} \rangle \in \mathbb{N}^d \mid 0 \leq s_i < n_i \}, \\ [\bar{n}]^\pm &:= \{ \langle s_0, \dots, s_{d-1} \rangle \in \mathbb{N}^d \mid -n_i < s_i < n_i \}. \end{aligned}$$

(b) Let \mathcal{A} be an abelian group. A *progression* is a set P of the form

$$P = a + [\bar{n}] \cdot \bar{b} := \{ a + \sum_{i < d} s_i b_i \mid s_i < n_i \},$$

where $a, b_0, \dots, b_{d-1} \in A$, $d, n_0, \dots, n_{d-1} < \omega$, and \cdot denotes the scalar product.



Such a progression is *proper* if $|P| = |[\bar{n}]|$, that is, if every element of P can be uniquely expressed as a linear combination of a, b_0, \dots, b_{d-1} . We call the number d the *rank* of the progression.

(c) A *coset progression* is a subset of A of the form $P + H$, where $P \subseteq A$ is a progression and $H \subseteq A$ a finite subgroup. The *rank* of $P + H$ is the rank of P . A coset progression $P + H$ is *proper* if P is proper and the sum $P + H$ is direct, i.e., every element of $P + H$ can be uniquely written as $p + h$ with $p \in P$ and $h \in H$.

A characteristic property of progressions P is that the sets of the form $P + P := \{ p + q \mid p, q \in P \}$ are small, that is, linear in the size of P . The following theorem of Freiman states the converse: every set S such that $S + S$ is small is close to a progression.

Theorem 9.20 (Freiman). *Let \mathcal{A} be an abelian group and let $\alpha > 0$ be a constant. There exist numbers β and d with the following property. For every finite $S \subseteq A$ satisfying $|S + S| \leq \alpha \cdot |S|$, there exists a proper coset progression $P + H$ of rank at most d with*

$$S \subseteq P + H \quad \text{and} \quad |P| \cdot |H| \leq \beta \cdot |S|.$$

The next result can be used to make a progression proper without increasing its size too much.

Lemma 9.21. *Let \mathcal{A} be an abelian group. There exists a constant β such that every progression P of rank d is contained in a proper progression Q of rank at most d and size $|Q| \leq d^{\beta d^3} \cdot |P|$.*

As a technical tool we introduce norms on abelian groups.

Definition 9.22. Let \mathcal{A} be an abelian group.

(a) A *norm* on \mathcal{A} is a function $\| \cdot \| : A \rightarrow \omega$ such that

$$\|a\| = \|-a\| \quad \text{and} \quad \|a + b\| \leq \max \{ \|a\|, \|b\| \}, \quad \text{for all } a, b \in A.$$

For a set $S \subseteq A$, we set

$$\|S\| := \max \{ \|s\| \mid s \in S \}.$$

(b) Let p be a prime. A *p-norm* on \mathcal{A} is a norm $\| \cdot \|$ satisfying the following additional conditions for all $a \in A$ and $k \in \mathbb{Z}$.

- ♦ $\|ka\| < \|a\|$ implies $p \mid k$
- ♦ $\lim_{n \rightarrow \infty} \|p^n a\| = 0$
- ♦ 0 is the only accumulation point of A .

Example. The prototypical p -norm is the p -adic norm on \mathbb{Q} , which is defined by $\|o\|_p := o$ and

$$\|a\|_p := p^m \quad \text{for } a = \frac{k}{qp^m} \text{ with } k \text{ and } q \text{ coprime to } p.$$

Below we will need the following two technical results on norms and progressions.

Lemma 9.23. *Let \mathcal{A} be an abelian group with a norm $\|\cdot\|$, and let $P = a + [\bar{n}] \cdot \bar{b}$ be a progression containing an element $c \in P$ with $\|c\| < \|P\|$. Then*

$$\|b_i\| = \|P\|, \quad \text{for some } i.$$

Proof. Let d be the rank of P and let $H := \langle b_o, \dots, b_{d-1} \rangle_{\mathcal{A}}$ be the subgroup generated by b_o, \dots, b_{d-1} . By the properties of a norm, it follows that

$$\|h\| \leq \max_i \|b_i\|, \quad \text{for all } h \in H.$$

Hence, $\|H\| = \max_i \|b_i\|$ and it remains to prove that $\|H\| \geq \|P\|$.

For a contradiction, suppose that $\|H\| < \|P\|$. Then $\|P\| = \|a\|$. By assumption, there is $c \in P$ with $\|c\| < \|P\|$. Let $c = a + h$, for some $h \in H$. Then

$$\|P\| = \|a\| \leq \max \{ \|a + h\|, \|-h\| \} = \|a + h\| < \|P\|,$$

where the third step follows from the fact that $\|-h\| = \|h\| < \|a\|$. We have obtained a contradiction. \square

Lemma 9.24. *Let \mathcal{A} be an abelian group with a p -norm $\|\cdot\|$, and let*

$$T = \{ t < n \mid \|a + tb\| < \|b\| \}, \quad \text{for some } a, b \in A \text{ and } n < \omega.$$

Then

$$|T| \leq \lceil n/p \rceil \quad \text{and} \quad \frac{1}{n}|T| \leq \max \left\{ \frac{1}{n}, \frac{2}{p} \right\}.$$

Proof. For $t, t' \in T$, we have

$$\|b\| > \max \{ \|a + tb\|, \|a + t'b\| \} \geq \|(t - t')b\|.$$

Since $\|\cdot\|$ is a p -norm, it follows that

$$p \mid t - t', \quad \text{for all } t, t' \in T.$$

This implies that

$$|T| \leq \lceil n/p \rceil.$$

For the second claim, we distinguish two cases. If $n \leq p$, we have

$$\frac{\lceil n/p \rceil}{n} \leq \frac{1}{n}.$$

Otherwise,

$$\frac{\lceil n/p \rceil}{n} \leq \frac{n/p + 1}{n} \leq \frac{1}{p} + \frac{1}{n} \leq \frac{1}{p} + \frac{1}{p} = \frac{2}{p}.$$

Hence, $\frac{1}{n}|T| \leq \frac{\lceil n/p \rceil}{n} \leq \max \left\{ \frac{1}{n}, \frac{2}{p} \right\}$. \square

Lemma 9.25. *Let \mathcal{A} be an abelian group, $H \subseteq A$ a subgroup, p a prime number, and $n < \omega$, and $a, b \in A$ elements with $b \not\in_p H$. The size of the set*

$$T = \{ t < n \mid a + tb \in H \}$$

is bounded by

$$|T| \leq \lceil n/p \rceil \quad \text{and} \quad \frac{1}{n}|T| \leq \max \left\{ \frac{1}{n}, \frac{2}{p} \right\}.$$

Proof. For $t, t' \in T$,

$$a + tb_o, a + t'b_o \in H \quad \text{implies} \quad (t - t')b_o = (a + tb_o) - (a + t'b_o) \in H.$$

Proof. Note that L , considered as a subset of \mathbb{R}^d , forms a discrete subgroup of \mathbb{R}^d . Furthermore, L has rank d since

$$\lim_n \|p^n b_i\| = 0, \quad \text{for all } i,$$

there is some $n < \omega$ such that $p^n \mathbb{Z}^d \subseteq L$. By Lemma 9.28, it is therefore sufficient to show that

$$\text{vol}(\mathbb{R}^d/L) \geq p.$$

Recall that $[\mathbb{Z}^d : L]$ denotes the index of L in \mathbb{Z}^d . Since the fundamental domain of \mathbb{Z}^d is the unit cube (which has volume 1), we have

$$\text{vol}(\mathbb{R}^d/L) = \text{vol}(\mathbb{R}^d/\mathbb{Z}^d) \cdot [\mathbb{Z}^d : L] = 1 \cdot [\mathbb{Z}^d : L].$$

To compute $[\mathbb{Z}^d : L]$, we consider the sets

$$H := \text{rng } \varphi \subseteq A \quad \text{and} \quad K := \Lambda \cap H.$$

Note that H and K are closed under addition and negation. Hence, they form subgroups of \mathfrak{A} . Furthermore, we have

$$\varphi(\bar{s}) \in \Lambda \cap \text{rng } \varphi = \Lambda \cap H = K, \quad \text{for every } \bar{s} \in L = \varphi^{-1}[\Lambda].$$

Consequently,

$$[\mathbb{Z}^d : L] = [\varphi[\mathbb{Z}^d] : \varphi[L]] = [H : \varphi[L]] \geq [H : K].$$

Fix $i < d$ with $\|b_i\| > \alpha$ and set $a := b_i$. Since $\|\cdot\|$ is a p -norm, we have

$$\|a\| = \|2a\| = \cdots = \|(p-1)a\| > \alpha.$$

For $0 \leq i < j < p$, it follows that

$$\|ja - ia\| = \|(j-i)a\| > \alpha,$$

As $\|K\| \leq \alpha$, we therefore have $ja - ia \notin K$. Hence, the cosets

$$K, a + K, 2a + K, \dots, (p-1)a + K$$

are all different and we have $[H : K] \geq p$. Thus,

$$\text{vol}(\mathbb{R}^d/L) = [\mathbb{Z}^d : L] \geq [H : K] \geq p. \quad \square$$

The next result can be used to reduce the rank of a progression.

Lemma 9.30. *Let \mathfrak{A} be an abelian group with a p -norm $\|\cdot\|$, let $\alpha < \omega$, and let $P_0 = a + [\bar{n}] \cdot \bar{b}$ and $P_1 = a' + [\bar{n}'] \cdot \bar{b}'$ be proper progressions of rank, respectively, $d \geq 1$ and $d' \geq 0$ such that*

$$d + d' \geq 2, \quad |P_0| < p/d!, \quad \|\bar{b}\| > \alpha \geq \|\bar{b}'\|,$$

and $P_0 + P_1$ contains some element of norm at most α . There exists a progression Q of rank strictly less than $d + d'$ such that

$$|Q| < 4^d d^{2d^2} \cdot |P_0| \cdot |P_1| \quad \text{and} \quad Q \supseteq \{c \in P_0 + P_1 \mid \|c\| \leq \alpha\}.$$

Proof. Set $m := |P_0|$, let $\Lambda := \{c \in A \mid \|c\| \leq \alpha\}$ be the set of elements of norm at most α , define

$$\varphi : \mathbb{Z}^d \rightarrow A : \bar{s} \mapsto \bar{s} \cdot \bar{b},$$

and set

$$L := \varphi^{-1}[\Lambda] \subseteq \mathbb{Z}^d \quad \text{and} \quad B := \{\bar{x} \in \mathbb{R}^d \mid -n_i < x_i < n_i\}.$$

Note that B is open, symmetric, and convex, and its volume is

$$\text{vol}(B) = (2n_0) \cdots (2n_{d-1}) = 2^d m < \frac{2^d p}{d!}.$$

Since $\|\bar{b}\| > \alpha$, it therefore follows by Lemma 9.29 that $B \cap L$ is contained in a sublattice of L of rank $r < d$.

By assumption, we have $(P_o + P_1) \cap \Lambda \neq \emptyset$ and $\|\tilde{b}'\| \leq \alpha$. Hence, there exists a tuple $\tilde{t} \in [\tilde{n}]$ with $a + a' + \tilde{t} \cdot \tilde{b} \in \Lambda$. Setting $\hat{a} := a + \tilde{t} \cdot \tilde{b}$, we have

$$P_o + P_1 = a + P_1 + \varphi([\tilde{n}]) = \hat{a} + P_1 + \varphi([\tilde{n}] - \tilde{t}).$$

Since $\hat{a} + P_1 \subseteq \Lambda$, it follows that

$$\begin{aligned} (P_o + P_1) \cap \Lambda &= \hat{a} + P_1 + \varphi([\tilde{n}] - \tilde{t}) \cap \Lambda \\ &= \hat{a} + P_1 + \varphi([\tilde{n}] - \tilde{t} \cap L) \subseteq \hat{a} + P_1 + \varphi(B \cap L). \end{aligned}$$

In the case where $d = 1$, $\dim \text{span}(B \cap L) = r < d = 1$ implies that $B \cap L = \{0\}$. Setting $Q := \hat{a} + P_1$, we have

$$(P_o + P_1) \cap \Lambda \subseteq \hat{a} + P_1 + \varphi(B \cap L) = \hat{a} + P_1 = Q.$$

Furthermore, the rank of Q is $d' < d + d'$ and its size is

$$|Q| = |P_1| < 4 \cdot |P_o| \cdot |P_1| = 4^d d^{2d^2} \cdot |P_o| \cdot |P_1|,$$

as desired.

Hence, it remains to consider the case where $d > 1$. We use Lemma 9.27 to find a tuple $\tilde{s} \in L^r$ of linearly independent vectors and a tuple $\tilde{k} \in \mathbb{N}^r$ of positive integers such that

$$[\tilde{k}]^\pm \cdot \tilde{s} \subseteq B \cap L \subseteq [r^{2r} \tilde{k}]^\pm \cdot \tilde{s}.$$

Setting $Q := \hat{a} + P_1 + [r^{2r} \tilde{k}]^\pm \cdot \varphi(\tilde{s})$, it then follows that

$$\begin{aligned} (P_o + P_1) \cap \Lambda &\subseteq \hat{a} + P_1 + \varphi(B \cap L) \\ &\subseteq \hat{a} + P_1 + \varphi([r^{2r} \tilde{k}]^\pm \cdot \tilde{s}) \\ &\subseteq \hat{a} + P_1 + [r^{2r} \tilde{k}]^\pm \cdot \varphi(\tilde{s}) = Q. \end{aligned}$$

Furthermore, the rank of Q is $r + d' < d + d'$, and the independence of \tilde{s} implies that

$$|[\tilde{k}]^\pm| = |[\tilde{k}]^\pm \cdot \tilde{s}| \leq |B \cap L| \leq |B \cap \mathbb{Z}^d| < 2^d m.$$

Hence, we can bound the size of Q by

$$\begin{aligned} |Q| &\leq |P_1| \cdot |[r^{2r} \tilde{k}]^\pm| \\ &< |P_1| \cdot 2^r r^{2r^2} \cdot |[\tilde{k}]^\pm| < |P_1| \cdot 2^r r^{2r^2} \cdot 2^d m < 4^d d^{2d^2} \cdot |P_o| \cdot |P_1|. \quad \square \end{aligned}$$

The key argument of the proof of Theorem 9.17 is contained in the following technical lemma. We introduce a measure $\theta(S, d)$ that indicates how well a set S is approximated by a progression P of rank d . Then the lemma provides a lower bound for how θ changes when we add a new element to the set S .

Definition 9.31. Let \mathcal{A} be an abelian group. For $S \subseteq A$ and $d < \omega$, we define

$$\theta(S, d) := \min \left\{ |P|/|S| \mid P \supseteq S \text{ a proper progression of rank at most } d \right\}.$$

If this minimum does not exist, we set $\theta(S, d) := \infty$.

Lemma 9.32. For every abelian group \mathcal{A} there exists a constant $\gamma > 0$ with the following property. Given an integer $d \geq 1$, a prime $p > d!$, a p -norm $\|\cdot\|$, a set $S \subseteq A$ of size $|S| > 1$, and an element $c \in A$ with $\|c\| > \|S\|$, we have

$$\theta(S \cup \{c\}, d) \geq \min \left\{ \frac{p^{1/d}}{4d}, \frac{\theta(S, d-1)}{d^{\gamma d^3}} \right\}.$$

Proof. If $\theta(S \cup \{c\}, d) = \infty$, the claim is trivial. Hence, suppose otherwise. We choose a proper progression $P = a + [\tilde{n}] \cdot \tilde{b}$ of rank at most d such that $S \cup \{c\} \subseteq P$ and

$$\theta(S \cup \{c\}, d) = \frac{|P|}{|S \cup \{c\}|}.$$

The set

$$\Lambda := \{a \in A \mid \|a\| < \|c\|\}$$

induces a subgroup of \mathfrak{A} with $S \subseteq P \cap \Lambda$.

Since $\|S\| < \|c\| = \|S \cup \{c\}\|$, we can use Lemma 9.23 to find an index i with $\|b_i\| \geq \|c\|$. Reordering b_0, \dots, b_{d-1} if necessary, we may therefore assume that there is some index $0 < k \leq d$ such that

$$\|b_0\|, \dots, \|b_{k-1}\| \geq \|c\| > \|b_k\|, \dots, \|b_{d-1}\|.$$

Set $m := n_0 \cdots n_{k-1}$. We distinguish two cases.

(i) First suppose that $m \geq p/k!$. Then

$$n_i \geq (p/k!)^{1/k} \geq (p/k!)^{1/d}, \quad \text{for some } i < k.$$

Without loss of generality, we may assume that $i = 0$. For $s_1, \dots, s_{d-1} \in \mathbb{Z}$, set

$$T(\vec{s}) := \{t < n_0 \mid \|a + tb_0 + s_1 b_1 + \dots + s_{d-1} b_{d-1}\| < \|c\|\}.$$

By Lemma 9.24, we have

$$\frac{|T(\vec{s})|}{n_0} \leq \max\left\{\frac{1}{n_0}, \frac{2}{p}\right\} \leq \max\left\{\left(\frac{d!}{p}\right)^{1/d}, \frac{2}{p}\right\}.$$

Choose \vec{s}' such that $T(\vec{s}')$ has maximal size. As P is proper, it follows that

$$\begin{aligned} \frac{|P \cap \Lambda|}{|P|} &\leq \frac{\sum_{\vec{s}} |T(\vec{s})|}{n_0 \cdots n_{d-1}} \\ &\leq \frac{n_1 \cdots n_{d-1} \cdot |T(\vec{s}')|}{n_0 \cdots n_{d-1}} \leq \frac{|T(\vec{s}')|}{n_0} \leq \max\left\{\left(\frac{d!}{p}\right)^{1/d}, \frac{2}{p}\right\}. \end{aligned}$$

Furthermore, the bounds $n! \leq \frac{n^{n+1}}{e^{n-1}}$ and $\ln x \leq x + 1$ imply that

$$\begin{aligned} d!^{1/d} &\leq \left(\frac{d^{d+1}}{e^{d-1}}\right)^{1/d} = d \cdot \frac{(de)^{1/d}}{e} \\ &= d \cdot \frac{e^{\frac{1}{d} \ln(de)}}{e} \\ &= d \cdot \frac{e^{\frac{1}{d} (\ln d + 1)}}{e} \leq d \cdot \frac{e^{\frac{1}{d} (d-1+1)}}{e} = d. \end{aligned}$$

Therefore,

$$\frac{2}{p} \leq \frac{2}{p^{1/d}} \leq \frac{2d}{p^{1/d}} \quad \text{and} \quad \left(\frac{d!}{p}\right)^{1/d} \leq \frac{2d}{p^{1/d}} \quad \text{implies} \quad \frac{|P \cap \Lambda|}{|P|} \leq \frac{2d}{p^{1/d}}.$$

Since $S \subseteq P \cap \Lambda$ it follows that

$$\theta(S \cup \{c\}, d) = \frac{|P|}{|S \cup \{c\}|} \geq \frac{|P|}{2 \cdot |S|} \geq \frac{|P|}{2 \cdot |P \cap \Lambda|} \geq \frac{1}{2 \cdot \frac{2d}{p^{1/d}}} \geq \frac{p^{1/d}}{4d},$$

as desired.

(ii) It remains to consider the case where $m < p/k!$. First, let us check that this implies that $d > 1$. For a contradiction, suppose that $d = 1$. Applying Lemma 9.24 to the set

$$T := \{t < n_0 \mid \|a + tb_0\| < \|c\|\}$$

we obtain $|n_0/p| \geq |T| = |P \cap \Lambda| \geq |S| > 1$. Hence, $m = n_0 > p = p/k!$. A contradiction.

Consider the progressions

$$\begin{aligned} P_0 &:= [n_0, \dots, n_{k-1}] \cdot \langle b_0, \dots, b_{k-1} \rangle, \\ P_1 &:= a + [n_k, \dots, n_{d-1}] \cdot \langle b_k, \dots, b_{d-1} \rangle. \end{aligned}$$

Then $|P_0| = m < p/k!$ and $P = P_0 + P_1$ contains the elements of S , which have norm less than $\|c\|$. Consequently, we can apply Lemma 9.30 and we obtain a progression $Q \supseteq (P_0 + P_1) \cap \Lambda \supseteq S$ of rank less than d such that

$$|Q| < 4^k k^{2k^2} \cdot |P_0| \cdot |P_1| \leq 4^d d^{2d^2} \cdot |P|.$$

Using Lemma 9.21, there exists a constant β (independent of S and d) and a proper progression $Q' \supseteq Q$ whose rank is at most that of Q such that

$$|Q'| \leq d^{\beta d^3} \cdot |Q|.$$

Since

$$\begin{aligned} 2 \cdot 4^d \cdot d^{2d^2} \cdot d^{\beta d^3} &\leq d \cdot (d^2)^d \cdot d^{2d^2} \cdot d^{\beta d^3} \\ &= d^{1+2d+2d^2+\beta d^3} \\ &\leq d^{d^3+d^3+d^3+\beta d^3} = d^{(3+\beta)d^3}, \end{aligned}$$

it follows that

$$\begin{aligned} \theta(S \cup \{c\}, d) &= \frac{|P|}{|S \cup \{c\}|} \\ &\geq \frac{|Q|}{2 \cdot |S| \cdot 4^d d^{2d^2}} \\ &\geq \frac{|Q'|}{2 \cdot |S| \cdot 4^d d^{2d^2} \cdot d^{\beta d^3}} \\ &\geq \frac{\theta(S, d-1)}{2 \cdot 4^d d^{2d^2} \cdot d^{\beta d^3}} \geq d^{-(3+\beta)d^3} \cdot \theta(S, d-1). \quad \square \end{aligned}$$

We are finally able to put all the parts together.

Proposition 9.33. *Let \mathcal{A} be an abelian group and P an infinite set of primes with the following properties.*

- ♦ *Every finite subgroup of \mathcal{A} is trivial or cyclic.*
- ♦ *For every $p \in P$, there exists an unbounded p -norm $\|\cdot\|_p$ on \mathcal{A} .*
- ♦ *For every $p \in P$, the map $a \mapsto pa$ has a finite kernel.*

Then \mathcal{A} is not automatic.

Proof. For a contradiction, suppose that \mathcal{A} is automatic. We fix an injective automatic presentation of \mathcal{A} . By Lemma 8.8, there exists a constant k such that

$$\|a_0 + \cdots + a_{n-1}\| \leq \|\bar{a}\| + k \lceil \log_2 n \rceil, \quad \text{for all } \bar{a} \in A^n.$$

(Here, $\|a\|$ denotes the length of the word encoding a , not the p -norm $\|\cdot\|_p$.) Choose a number $l \geq \|\circ\|$ such that there are at least two elements $a \in A$

with $\|a\| \leq l$ and set

$$S_n := \{a \in A \mid \|a\| \leq l + nk\}, \quad \text{for } n < \omega.$$

Then

$$\circ \in S_0, \quad |S_0| \geq 2, \quad \text{and} \quad S_n + S_n \subseteq S_{n+1}.$$

Furthermore, it follows by Lemma 8.6 that there exists some constant α such that

$$|S_{n+1}| \leq \alpha |S_n|, \quad \text{for all } n.$$

Then $|S_n + S_n| \leq |S_{n+1}| \leq \alpha |S_n|$, and we can use the Theorem 9.20 to find constants β and d_0 such that every set S_n has a proper coset progression $P + H$ of rank d_0 and size $|P + H| \leq \beta \cdot |S_n|$. Note that every finite subgroup of \mathcal{A} is trivial or cyclic and that cyclic groups form progressions of rank 1. Consequently, $P + H$ can be written as a proper progression of rank at most $d_0 + 1$. This implies that

$$\theta(S_n, d_0 + 1) \leq \beta, \quad \text{for all } n < \omega.$$

Set $d := d_0 + 1$.

Since the definable functions $a \mapsto pa$, for $p \in P$, have finite kernels, it follows by Corollary 8.3 that there exists a function $m : P \rightarrow \mathbb{N}$ such that

$$p^{-1}S_n \subseteq S_{n+m(p)}, \quad \text{for every } p \in P.$$

We inductively choose a sequence $p_d < \cdots < p_0$ of prime numbers in P (starting with p_d) such that

$$p_d > (4d\alpha\beta)^d \quad \text{and} \quad p_{i-1} > p_i \alpha^{dm(p_i)} d^{\gamma d^4},$$

where γ is the constant from Lemma 9.32. Since $\bigcup_n S_n = A$ and the norms $\|\cdot\|_{p_i}$ are unbounded on A , we can now choose a sequence $n_0 < \cdots < n_d$ of natural numbers by

$$n_0 := 0 \quad \text{and} \quad n_{i+1} := \min \{n \mid \|S_n\|_{p_{i+1}} > \|S_{n_i}\|_{p_{i+1}}\}.$$

We claim that

$$\theta(S_{n_i}, i) \geq \frac{p_i^{1/d}}{4d\alpha}.$$

For $i = d$, it then follows that

$$\theta(S_{n_d}, d) \geq \frac{p_d^{1/d}}{4d\alpha} > \frac{4d\alpha\beta}{4d\alpha} = \beta.$$

A contradiction.

Hence, it remains to prove the claim. We proceed by induction on i . For $i = 0$, we have

$$\theta(S_{n_0}, 0) = \theta(S_0, 0) = \infty,$$

since $|S_0| \geq 2$ and every progression of rank 0 is a singleton. For the inductive step, suppose that we have already established the claim for $i - 1$.

Fix $b \in S_{n_{i-1}}$ with $\|S_{n_{i-1}}\|_{p_i} = \|b\|_{p_i}$. Then $p_i^{-1}b \in S_{n_{i-1}+m(p_i)}$ implies that

$$\|S_{n_{i-1}+m(p_i)}\|_{p_i} \geq \|p_i^{-1}b\|_{p_i} > \|b\|_{p_i} = \|S_{n_{i-1}}\|_{p_i}.$$

By choice of n_i , it follows that $n_i \leq n_{i-1} + m(p_i)$. Consequently,

$$|S_{n_{i-1}}| \leq \alpha^{m(p_i)-1} |S_{n_i}|.$$

Note that, by definition of θ ,

$$S \subseteq T \quad \text{implies} \quad \theta(T, i) \geq \frac{|S|}{|T|} \theta(S, i).$$

Therefore,

$$\begin{aligned} \theta(S_{n_{i-1}}, i - 1) &\geq \frac{|S_{n_{i-1}}|}{|S_{n_{i-1}}|} \theta(S_{n_{i-1}}, i - 1) \\ &\geq \frac{|S_{n_{i-1}}|}{\alpha^{m(p_i)-1} |S_{n_{i-1}}|} \theta(S_{n_{i-1}}, i - 1) \\ &= \alpha^{1-m(p_i)} \theta(S_{n_{i-1}}, i - 1) \\ &\geq \alpha^{1-m(p_i)} \frac{p_{i-1}^{1/d}}{4d\alpha} = \alpha^{-m(p_i)} \frac{p_{i-1}^{1/d}}{4d}. \end{aligned}$$

Similarly, for every $a \in S_{n_i}$, $|S_{n_i}| \leq \alpha |S_{n_{i-1}}| \leq \alpha |S_{n_{i-1}} \cup \{a\}|$ implies that

$$\begin{aligned} \theta(S_{n_i}, i) &\geq \frac{|S_{n_{i-1}} \cup \{a\}|}{|S_{n_i}|} \theta(S_{n_{i-1}} \cup \{a\}, i) \\ &\geq \frac{|S_{n_{i-1}} \cup \{a\}|}{\alpha |S_{n_{i-1}} \cup \{a\}|} \theta(S_{n_{i-1}} \cup \{a\}, i) = \alpha^{-1} \theta(S_{n_{i-1}} \cup \{a\}, i). \end{aligned}$$

By choice of n_i , there is some element $a \in S_{n_i}$ with $\|a\|_{p_i} > \|S_{n_{i-1}}\|_{p_i}$. Since $\|\cdot\|_{p_i}$ is a p_i -norm and $p_i > d^d \geq i!$, we can therefore use Lemma 9.32 to show that

$$\begin{aligned} \theta(S_{n_i}, i) &\geq \alpha^{-1} \theta(S_{n_{i-1}} \cup \{a\}, i) \\ &\geq \alpha^{-1} \min \left\{ \frac{p_i^{1/d}}{4d}, d^{-\gamma d^3} \theta(S_{n_{i-1}}, i - 1) \right\} \\ &\geq \min \left\{ \frac{p_i^{1/d}}{4d\alpha}, \alpha^{-1} d^{-\gamma d^3} \alpha^{-m(p_i)} \frac{p_{i-1}^{1/d}}{4d} \right\} = \frac{p_i^{1/d}}{4d\alpha}, \end{aligned}$$

where the last step follows from the fact that, by choice of p_{i-1} ,

$$d^{\gamma d^3} \alpha^{m(p_i)} p_i^{1/d} < p_{i-1}^{1/d}. \quad \square$$

We can now reduce both claims of Theorem 9.17 to Proposition 9.33.

Proof of Theorem 9.17. (a) For a contradiction, suppose that there exists an automatic torsion-free abelian group \mathfrak{A} and an infinite set P of primes such that \mathfrak{A} is p -divisible for all $p \in P$. If \mathfrak{A} contains an infinite linearly independent subset, it is not automatic by Proposition 8.10. Consequently, $\mathfrak{A} \subseteq \mathbb{Q}^n$, for some $n < \omega$.

We obtain the desired contradiction via Proposition 9.33. To do so we need three ingredients. First note that, \mathfrak{A} being torsion-free, the only finite subgroup is trivial. Furthermore, for each $p \in P$, the map $a \mapsto pa$ is injective and has trivial kernel. Therefore, it remains to construct unbounded p -norms $\|\cdot\|_p$, for $p \in P$. Let $\|\cdot\|_p$ be the p -adic norm on \mathbb{Q} defined by $\|o\|_p = o$ and

$$\|a\|_p := p^m \quad \text{for } a = \frac{k}{qp^m} \text{ with } k \text{ and } q \text{ coprime to } p.$$

(It is straightforward to check that $\|\cdot\|_p$ is an unbounded p -norm on A .) For $a \in A \subseteq \mathbb{Q}^n$, we then set

$$\|a\|_p := \max \{ \|a_o\|_p, \dots, \|a_{n-1}\|_p \},$$

where a_o, \dots, a_{n-1} are the components of the tuple a .

(b) For a contradiction, suppose that there exists an infinite set P of primes such that $\mathfrak{A} := \bigoplus_{p \in P} \mathbb{Z}[1/p]/\mathbb{Z}$ is automatic. Again, we obtain the desired contradiction via Proposition 9.33.

First, note that every finite subgroup of \mathfrak{A} is cyclic. Furthermore, for each $p \in P$, the map $a \mapsto pa$ has finite kernel. ($pa \in \mathbb{Z}$ implies that $a = k/p$ for some $k < p$.) Finally, for each $p \in P$, we define an unbounded p -norm on \mathfrak{A} as follows. Let $\pi_p : \bigcup_{p \in P} \mathbb{Z}[1/p]/\mathbb{Z} \rightarrow \mathbb{Z}[1/p]/\mathbb{Z}$ be the projection to the p -th component. We set

$$\|a\|_p := \text{ord}(\pi_p(a)),$$

where $\text{ord}(b)$ denotes the order of b . (Again it is straightforward to check that $\|\cdot\|_p$ is an unbounded p -norm on A .) \square

Open Question. Does there exist an infinite set P of numbers such that the group $\bigoplus_{p \in P} \mathbb{Z}/p\mathbb{Z}$ embeds into some automatic group?

Conjecture. Every torsion-free abelian automatic group is an extension of \mathbb{Z}^n by a finite direct sum of $\mathbb{Z}(p^\infty)$, for prime numbers p .

10 Automatic Semirings

After groups, we take a look at rings and semirings. Unsurprisingly it turns out that few semirings are automatic. A notable exception is the tropical semiring $(\mathbb{N} \cup \{\infty\}, \min, +, \infty, o)$, which is automatic.

Lattices and Boolean Algebras

We start with automatic boolean algebras.

Definition 10.1. Let $\mathfrak{B} = \langle B, \sqcap, \sqcup, *, \sqsubseteq, \top, \perp \rangle$ be a boolean algebra.

- (a) An *atom* is an element $a \in B$ such that $b \sqsubset a$ implies $b = \perp$.
- (b) Two $a, b \in B$ have a *finite difference* if they satisfy the relation

$$a \approx_* b \quad : \text{iff} \quad (a \sqcap b^*) \sqcup (a^* \sqcap b) \text{ is a finite supremum of atoms.}$$

Theorem 10.2. Let \mathfrak{F} be the boolean algebra of all finite and all cofinite subsets of ω . For a boolean algebra \mathfrak{B} , the following statements are equivalent.

- (1) \mathfrak{B} is automatic.
- (2) \mathfrak{B} is finite or isomorphic to \mathfrak{F}^n , for some $n < \omega$.
- (3) \mathfrak{B} is countable and \mathfrak{B}/\approx_* is finite.

Proof. (2) \Rightarrow (1) Since all finite structures are automatic and automatic structures are closed under direct products, it is sufficient to prove that \mathfrak{F} is automatic. We have trivial interpretations

$$\mathfrak{F} \leq_{\text{FO}} \langle \wp_{\text{fin}}(\omega), \subseteq \rangle \cong \wp_{\text{fin}}(\omega) \leq_{\text{FO}} \wp_{\text{fin}}(\omega, \leq).$$

Hence, the claim follows by Theorem 2.3.

(3) \Rightarrow (2) Let \mathfrak{B} be a countably infinite boolean algebra such that \mathfrak{B}/\approx_* is finite. Let $[a'_o]_*, \dots, [a'_{n-1}]_*$ be the atoms of \mathfrak{B}/\approx_* . Set

$$a_o := a'_o \sqcap (a'_1 \sqcup \dots \sqcup a'_{n-1})^* \quad \text{and} \quad a_{i+1} := a'_i \sqcap (a'_o \sqcup \dots \sqcup a'_{i-1})^*.$$

Then

$$a_i \approx_* a'_i, \quad a_0 \sqcup \dots \sqcup a_{n-1} = \top, \quad \text{and} \quad a_i \sqcap a_j = \perp, \quad \text{for } i \neq j.$$

For $b \in B$, define

$$A(b) := \{ a \mid a \text{ atom}, a \subseteq b \}.$$

For $i < n$, let $\varphi_i : \mathfrak{B} \rightarrow \mathfrak{B}$ be the function sending $b \in B$ to $b \sqcap a_i$. Clearly, φ_i is a lattice homomorphism. Since $a_i \not\approx_* \perp$, each set $A(a_i)$ is countably infinite and, for every $b \in B$, $A(b \sqcap a_i)$ is a finite or cofinite subset of $A(a_i)$. Consequently, the range of φ_i induces a sublattice of \mathfrak{B} that is isomorphic to \mathfrak{F} . We therefore obtain a lattice homomorphism

$$\varphi := \langle \varphi_i \rangle_{i < n} : \mathfrak{B} \rightarrow \mathfrak{F}^n.$$

First of all, note that φ is even a homomorphism of boolean algebras since

$$\begin{aligned} \varphi(b^*) &= \langle b^* \sqcap a_0, \dots, b^* \sqcap a_{n-1} \rangle \\ &= \langle b \sqcap a_0, \dots, b \sqcap a_{n-1} \rangle^* = \varphi(b)^*. \end{aligned}$$

For injectivity, note that

$$\begin{aligned} \varphi_0(b) \sqcup \dots \sqcup \varphi_{n-1}(b) &= (b \sqcap a_0) \sqcup \dots \sqcup (b \sqcap a_{n-1}) \\ &= b \sqcap (a_0 \sqcup \dots \sqcup a_{n-1}) \\ &= b. \end{aligned}$$

Finally, surjectivity follows from the fact that \mathfrak{F} is the range of φ_i .

(1) \Rightarrow (3) We start by introducing some terminology. Let $a \in B$.

- ♦ a is *large* if its \approx_* -class $[a]_*$ is not a finite supremum of atoms, i.e., if $[a]_* \not\approx_* \perp$ in \mathfrak{B}/\approx_* .
- ♦ a is *infinite* if there are infinitely many elements $b \subseteq a$, i.e., if $a \not\approx_* \perp$.
- ♦ a *splits* an element $b \in B$ if $a \sqcap b \neq \perp$ and $a^* \sqcap b \neq \perp$.

Note that

- ♦ every large element is infinite;
- ♦ for every large element a , there exists some $c \in B$ such that $c \sqcap a$ is large and $c^* \sqcap a$ is infinite; and
- ♦ for every infinite element a , there exists some $c \in B$ such that at least one of $c \sqcap a$ is large and $c^* \sqcap a$ is infinite.

For a contradiction, suppose that \mathfrak{B} is automatic but \mathfrak{B}/\approx_* is infinite. Then \top is large. Since every large element can be split into a large element and an infinite one, we can construct an infinite strictly descending sequence

$$\top = a_0 \supset a_1 \supset \dots$$

where, for each $n < \omega$, a_{n+1} is the \leq_{lex} -least large element such that the difference $b_n := a_n \sqcap a_{n+1}^*$ is infinite. Similarly, since every infinite element can be split into two elements at least one of which is also infinite, we can construct infinite strictly descending sequences

$$b_n = c_0^n \supset c_1^n \supset \dots$$

where, for each $i < \omega$, c_{i+1}^n is the \leq_{lex} -least infinite element such that the difference $d_i^n := c_i^n \sqcap c_{i+1}^n \neq \perp$.

Note that the properties of being large and being infinite are both FOC-definable. Hence, there exists an FOC-formula $\varphi(x, x', y)$ stating that there exists some element z such that

- ♦ $y = x \sqcap x'$; or
- ♦ $y = x \sqcup x'$; or
- ♦ x is large, z is the \leq_{lex} -least large element such that $z \sqsubset x$ and $x \sqcap z^*$ is infinite, and $[y = z \text{ or } y = x \sqcap z^*]$; or
- ♦ x is infinite, z is the \leq_{lex} -least infinite element such that $z \sqsubset x$ and $x \sqcap z^* \neq \perp$, and $[y = z \text{ or } y = x \sqcap z^*]$.

Then

$$\begin{aligned} a_n &\in N_\varphi(a_0, n), \\ c_i^n &\in N_\varphi(a_0, n + i + 1), \\ d_i^n &\in N_\varphi(a_0, n + i + 2). \end{aligned}$$

Consequently, $N_\varphi(a_o, s + 2 + s)$ contains all elements of the sublattice generated by $D_s := \{d_i^n \mid n + i < s\}$. Since $d_i^m \sqcap d_j^n = \perp$, for $\langle m, i \rangle \neq \langle n, j \rangle$, this sublattice has size $2^{|D_s|}$ and it follows that

$$|N_\varphi(a_o, s + 2 + s)| \geq 2^{|D_s|} = 2^{s(s-1)/2}.$$

A contradiction to Proposition 8.16. \square

The following result shows that tree automatic boolean algebras are more general than automatic ones.

Definition 10.3. A boolean algebra is *atomless* if it has no atoms. ,

One can use a standard back-and-forth construction like in the Theorem of Cantor to show that all countable atomless boolean algebras are isomorphic.

Theorem 10.4. *Up to isomorphism there exists a unique countable atomless boolean algebra.*

Proposition 10.5. *The countable atomless boolean algebra is tree automatic.*

Proof. The countable atomless boolean algebra \mathfrak{B} is isomorphic to the set of clopen subsets of Cantor space $[2]^\omega$. Such sets are of the form $W[2]^\omega$, for finite $W \subseteq [2]^*$. We can encode each finite set W as a $[2]$ -labelled finite tree. Then intersection, union, and complement for such sets are regular. \square

For ω -automatic lattices, we have the following partial result.

Proposition 10.6. *Let $\mathcal{L} = \langle L, \sqcup, \sqcap \rangle$ be an uncountable ω -automatic lattice. Then there are elements $a, b \in L$ such that $\Downarrow a$ and $\Uparrow b$ are both uncountable.*

Proof. Let k be the constant from Theorem 6.13 and define

$$f(x_o, \dots, x_k) := \bigwedge_{i < k+1} x_i \quad \text{and} \quad g(x_o, \dots, x_k) := \bigvee_{i < k+1} x_i.$$

By choice of k , there exist uncountable regular sets $U, U' \subseteq A^k$ and $V, V' \subseteq A$ such that the restrictions $f \upharpoonright U \times V$ and $g \upharpoonright U' \times V'$ are constant. Setting

$a := f(\bar{c}, d)$ and $b := g(\bar{c}', d')$, for $\bar{c}, \bar{c}' \in U$ and $d, d' \in V$, it follows that $f^{-1}(a)$ and $g^{-1}(b)$ are uncountable. Hence, so are the sets

$$\bigcup_{i < k+1} p_i[f^{-1}(a)] \subseteq \Downarrow a \quad \text{and} \quad \bigcup_{i < k+1} p_i[g^{-1}(b)] \subseteq \Uparrow b,$$

where $p_i : A^{k+1} \rightarrow A$ denotes the projection to the i -th component. \square

Exercise 10.1. Show that every automatic boolean algebra is finitely FOC-axiomatisable. ,

Exercise 10.2. Show that the isomorphism problem for automatic boolean algebras is decidable. ,

Integral Domains

Next, let us take a look at automatic integral domains, i.e., commutative rings without zero-divisors. (For us, rings will always be rings with identity.)

Proposition 10.7. *An integral domain is automatic if, and only if, it is finite.*

Proof. For a contradiction, suppose that there exists an infinite integral domain $\mathfrak{R} = \langle R, +, -, \cdot, 0, 1 \rangle$ that is automatic. Fix an injective presentation of \mathfrak{R} . We say that an element $c \in R$ *separates* a set $W \subseteq R$ if

$$ac + b = a'c + b' \quad \text{implies} \quad a = a' \text{ and } b = b',$$

for all $a, a', b, b' \in W$. We start by proving that every finite set $W \subseteq R$ is separated by some $c \in R$. For a contradiction, suppose there is some finite $W \subseteq R$ which is not. For every $c \in R$, it then follows that there are $a, a', b, b' \in W$ with $ac + b = a'c + b'$ and $\langle a, b \rangle \neq \langle a', b' \rangle$. Since W is finite, we can find elements $a, a', b, b' \in W$ such that $\langle a, b \rangle \neq \langle a', b' \rangle$ and

$$ac + b = a'c + b', \quad \text{for infinitely many } c \in R.$$

This implies that $(a - a')c = b - b'$. If $a = a'$, it follows that $b = b'$, contrary to our assumption. Hence, $a \neq a'$. Furthermore, by choice of a, a', b, b' , there are elements c, c' with

$$(a - a')c = b - b' = (a - a')c'.$$

Since \mathfrak{R} is an integral domain, this implies $c = c'$. A contradiction.

Fix an automatic well-ordering \leq on \mathfrak{R} and let $\varphi(x, y)$ be the formula stating that

$$'y = az + b \text{ where } a, b \leq x \text{ and } z \text{ is the } \leq\text{-least element separating } \Downarrow x.'$$

Then φ has finite out-degree and

$$|N_\varphi(U, 1)| \geq |U|^2.$$

By induction on n it follows that

$$\begin{aligned} |N_\varphi(U, n)| &= |N_\varphi(N_\varphi(U, n-1), 1)| \\ &\geq |N_\varphi(U, n-1)|^2 \geq (|U|^{2^{n-1}})^2 = |U|^{2^n}. \end{aligned}$$

A contradiction to Proposition 8.16. \square

We can extend this result to ω -automatic structures.

Theorem 10.8. *An integral domain is ω -automatic if, and only if, it is finite.*

Proof. (\Leftarrow) is trivial since all finite structures are ω -automatic.

(\Rightarrow) For a contradiction, suppose that there exists an infinite ω -automatic integral domain \mathfrak{A} . By Proposition 10.7, \mathfrak{A} is not automatic and, hence, uncountable. Let k be the constant from Theorem 6.13 and define a function $f : A^{k+1} \rightarrow A$ by

$$f(\bar{a}, x) := \sum_{i < k} a_i x^i.$$

This function is FOC-definable and, by choice of k , there exists uncountable sets $U \subseteq A^k$ and $V \subseteq A$ such that $f \upharpoonright U \times V$ is constant. In an integral domain, every polynomial of degree n has at most n roots. This implies that, for all $\bar{a} \neq \bar{b}$ in U , there are at most $k-1$ elements $x \in A$ with

$$f(\bar{a}, x) - f(\bar{b}, x) = 0.$$

Thus, $|V| < k$. A contradiction. \square

Corollary 10.9. $\langle \mathbb{R}, +, \cdot \rangle$ and $\langle \mathbb{C}, +, \cdot \rangle$ are not ω -automatic.

General Rings

Finally, we consider possibly non-commutative rings. The aim of this section is to prove the following result.

Definition 10.10. An algebra \mathfrak{A} is *locally finite* if every finitely generated subalgebra is finite. \lrcorner

Theorem 10.11. *Every automatic ring is locally finite.*

Before giving the proof let us note that this result allows us to remove the assumption of commutativity from Proposition 10.7.

Corollary 10.12. *Every automatic ring without zero-divisors is finite.*

Proof. Let \mathfrak{R} be an automatic ring without zero-divisors. By Theorem 10.11, \mathfrak{R} is locally finite. We claim that \mathfrak{R} is commutative. Then the claim follows by Proposition 10.7.

Consider two elements $r, s \in R$ and let \mathfrak{S} be the subring generated by them. By assumption, \mathfrak{S} is finite. Hence, the non-zero elements of \mathfrak{S} form a finite cancellative monoid \mathfrak{S}^\times under multiplication. This implies that \mathfrak{S}^\times is a group and \mathfrak{S} is a finite skew-field. By the Theorem of Wedderburn, all finite skew-fields are commutative. Hence, $r, s \in S$ implies that $rs = sr$. \square

For the proof of Theorem 10.11, we start by looking at finiteness conditions for rings.

Lemma 10.13. *Let \mathfrak{R} be a ring of positive characteristic $q > 0$ and let $C \subseteq R$ be finite.*

- (a) *The closure of C under addition is finite.*
- (b) *If the closure of C under multiplication is finite, so is the subring $\langle\langle C \rangle\rangle_{\mathfrak{R}}$ of \mathfrak{R} generated by C .*

Proof. (a) Let S be the closure of C under addition. Every element of S can be written as

$$\sum_{c \in S} \lambda_c c, \quad \text{for } \lambda_c \in \mathbb{Z}.$$

Since \mathfrak{R} has characteristic q it follows that it is sufficient to consider coefficients $\lambda_c \in \{0, \dots, q-1\}$. Consequently, S contains at most $q^{|S|}$ elements.

(b) Suppose that the closure M of C under multiplication is finite. The subring $\langle\langle C \rangle\rangle_{\mathfrak{R}}$ generated by C is the closure of M under addition. Hence, the claim follows by (a). \square

Lemma 10.14. *Let \mathfrak{R} be a finitely generated ring and $H \subseteq R$ an ideal such that \mathfrak{R}/H is finite and $H^2 = 0$. Then \mathfrak{R} is finite.*

Proof. We start by proving that \mathfrak{R} has positive characteristic $p > 0$. By assumption, the quotient \mathfrak{R}/H is finite. Hence, its characteristic p_0 is positive. This implies that $p_0 \in H$ (where p_0 is considered as the sum $1 + \dots + 1$ in \mathfrak{R}). Since $H^2 = 0$, it follows that $p_0 p_0 = 0$. Consequently, the characteristic p of \mathfrak{R} divides p_0^2 . In particular, $p > 0$.

Fix a finite set of generators $C \subseteq R$ of \mathfrak{R} , let $\pi : C^* \rightarrow R$ be the function mapping a sequence in C to its product, and let $q : R \rightarrow R/H$ be the quotient map. Since $q[\pi[C^*]] \subseteq R/H$ is finite, we can find some finite set $X \subseteq C^*$ such that $q[\pi[X]] = R/H$. Let $n < \omega$ be some number such that $X \subseteq C^{<n}$. For every $u \in C^*$, it follows that the quotient $q(\pi(u)) \in \mathfrak{R}/H$ can be written as a product of less than n elements of $q[C]$. Hence, there exist $h_u \in H$ and $v_u \in C^{<n}$ such that

$$\pi(u) = \pi(v_u) + h_u.$$

Let K and U be the closures of the sets

$$\begin{aligned} K_0 &:= \{0, 1\} \cup \{h_u \mid u \in C^n\}, \\ U_0 &:= \{k\pi(u) \mid k \in K, u \in C^{<n}\} \end{aligned}$$

under addition. Since C^n and $C^{<n}$ are finite, it follows by Lemma 10.13 that so are K and U . Furthermore, $H^2 = 0$ implies that K is closed under multiplication. Hence, K induces a subring of \mathfrak{R} . We claim that $\pi[C^*] \subseteq U$. Then it follows that $R = \langle\langle C \rangle\rangle_{\mathfrak{R}} = U$. In particular, \mathfrak{R} is finite.

To prove the claim, we show that, for every $u \in C^+$ of length at least n , there exists a function $\alpha : C^{<n} \rightarrow K$ such that

$$\pi(u) = \sum_{w \in C^{<n}} \alpha(w) \pi(w) \in U.$$

We proceed by induction on $|u|$. If $|u| = n$, we have

$$\pi(u) = \pi(v_u) + h_u = 1 \cdot \pi(v_u) + h_u \cdot \pi(\langle \rangle).$$

For the inductive step, suppose that $u = u_0 c$, for $u_0 \in C^*$ and $c \in C$. By inductive hypothesis, there exists a function α_0 such that

$$\pi(u) = \pi(u_0) \cdot c = \left(\sum_{w \in C^{<n}} \alpha_0(w) \pi(w) \right) \cdot c = \sum_{w \in C^{<n}} \alpha_0(w) \pi(wc).$$

Since $\pi(wc) = \pi(v_{wc}) + h_{wc}$, it follows that

$$\begin{aligned} \pi(u) &= \sum_{w \in C^{<n}} \alpha_0(w) (\pi(v_{wc}) + h_{wc}) \\ &= \sum_{w \in C^{<n}} \alpha_0(w) \pi(v_{wc}) + \sum_{w \in C^{<n}} \alpha_0(w) h_{wc} \pi(\langle \rangle) \\ &= \sum_{w \in C^{<n}} \alpha(w) \pi(w), \end{aligned}$$

where the function α takes the form

$$\alpha(w) = \sum_i \alpha_0(w'_i) + \sum_i \alpha_0(w''_i) h_{w''_i c_i}, \quad \text{for suitable } w'_i, w''_i \text{ and } c_i.$$

Note that the sum on the right-hand side belongs to $\langle\langle K \rangle\rangle_{\mathfrak{R}} = K$. \square

To find an ideal as in the previous lemma, we will use certain matrix groups.

Definition 10.15. Let \mathfrak{R} be a ring and let $I \in \mathfrak{R}^{n \times n}$ be the unit matrix. A *transvection* is a matrix of the form $E_{ij}(r) := I + e_{ij}^r$, where $i \neq j$ and e_{ij}^r is the matrix whose entries are all 0, except for the entry in row i and column j , which has value r . We denote by $\mathfrak{E}_n(\mathfrak{R})$ the subgroup of the matrix group $\text{GL}_n(\mathfrak{R})$ generated by all transvections. \lrcorner

Let us collect a few arithmetic laws for such matrices.

Lemma 10.16. *Let \mathfrak{R} be a ring, $r, s \in R$, and let $i, j, k < n$ be pairwise distinct indices.*

$$(a) \ E_{ij}(r)^{-1} = E_{ij}(-r)$$

$$(b) \ E_{ij}(r+s) = E_{ij}(r) \cdot E_{ij}(s)$$

$$(c) \ E_{ik}(rs) = [E_{ij}(r), E_{jk}(s)]$$

(Here $[x, y] = x^{-1}y^{-1}xy$ denotes the commutator for groups, not for rings.)

Proof. (a) follows by (b) for $s = -r$.

(b) Note that $e_{ij}^r \cdot e_{ij}^s = 0$ and $e_{ij}^r + e_{ij}^s = e_{ij}^{r+s}$. Hence,

$$\begin{aligned} E_{ij}(r) \cdot E_{ij}(s) &= (I + e_{ij}^r)(I + e_{ij}^s) \\ &= I + e_{ij}^r + e_{ij}^s + e_{ij}^r e_{ij}^s = I + e_{ij}^{r+s} = E_{ij}(r+s). \end{aligned}$$

(c) Since $e_{ij}^r \cdot e_{jk}^s = e_{ik}^{rs}$ and $e_{jk}^s \cdot e_{ij}^r = 0$ (as $k \neq i$), we have

$$\begin{aligned} &[E_{ij}(r), E_{jk}(s)] \\ &= [I + e_{ij}^r, I + e_{jk}^s] \\ &= (I + e_{ij}^r)(I + e_{jk}^s)(I + e_{ij}^r)(I + e_{jk}^s) \\ &= (I + e_{ij}^r + e_{jk}^s + e_{ij}^r e_{jk}^s)(I + e_{ij}^r + e_{jk}^s + e_{ij}^r e_{jk}^s) \\ &= (I - e_{ij}^r - e_{jk}^s + e_{ik}^{rs})(I + e_{ij}^r + e_{jk}^s + e_{ik}^{rs}) \\ &= (I + e_{ik}^{rs})^2 - (e_{ij}^r + e_{jk}^s)(I + e_{ik}^{rs}) + (I + e_{ik}^{rs})(e_{ij}^r + e_{jk}^s) \\ &\quad - (e_{ij}^r + e_{jk}^s)^2 \\ &= I + 2e_{ik}^{rs} + (e_{ik}^{rs})^2 - ((e_{ij}^r)^2 + e_{ij}^r e_{jk}^s + e_{jk}^s e_{ij}^r + (e_{jk}^s)^2) \\ &= I + e_{ik}^{rs}. \end{aligned} \quad \square$$

Lemma 10.17. *For every finitely generated ring \mathfrak{R} , the group $\mathfrak{E}_n(\mathfrak{R})$ is finitely generated.*

Proof. Let $S \subseteq R$ be a set of generators of \mathfrak{R} . We claim that $\mathfrak{E}_n(\mathfrak{R})$ is generated by the matrices $E_{ij}(s)$ and $E_{ij}(-s)$, for $s \in S$ and $i \neq j$. Let us call these the *elementary matrices*.

By definition, every element of $\mathfrak{E}_n(\mathfrak{R})$ can be written as a product of transvections and their inverses. By Lemma 10.16 (a), each inverse of a transvection is also a transvection. Consequently it is sufficient to show that every transvection $E_{ij}(r)$ with $r \in R$ can be written as a product of elementary ones. This follows by Lemma 10.16 (b) and (c) since r can be written as a sum of products of elements of S . \square

Proof of Theorem 10.11. Let \mathfrak{R} be an automatic ring and \mathfrak{S} a finitely generated subring. We have to show that \mathfrak{S} is finite. Representing 3×3 matrices over \mathfrak{R} by 9-tuples, we can interpret $\text{GL}_3(\mathfrak{R})$ in \mathfrak{R} . Hence, $\text{GL}_3(\mathfrak{R})$ is an automatic group. According to Lemma 10.17, the subgroup $\mathfrak{E}_3(\mathfrak{S}) \subseteq \text{GL}_3(\mathfrak{R})$ is finitely generated. Consequently it follows by Theorem 9.2 that $\mathfrak{E}_3(\mathfrak{S})$ is virtually abelian. By Lemma 9.4, this means it has a normal abelian subgroup $\mathfrak{A} \subseteq \mathfrak{E}_3(\mathfrak{S})$ of finite index.

Let $\mathfrak{S}_+ := \langle S, + \rangle$ be the additive group of \mathfrak{S} . For $i \neq j$, we consider the functions $g_{ij} : \mathfrak{S}_+ \rightarrow \mathfrak{E}_3(\mathfrak{S})$ with

$$g_{ij}(s) := E_{ij}(s),$$

and the set

$$H := \bigcap_{i \neq j} S_{ij}, \quad \text{where} \quad S_{ij} := g_{ij}^{-1}[A] = \{s \in S \mid E_{ij}(s) \in A\}.$$

Below, we will prove the following claims.

- (a) H induces a subgroup of \mathfrak{S}_+ of finite index.
- (b) H is an ideal of \mathfrak{S} and \mathfrak{S}/H is finite.
- (c) $H^2 = 0$.

Then it follows by Lemma 10.14 that \mathfrak{S} is finite.

(a) By Lemma 10.16 (b), the function $g_{ij} : \mathfrak{S}_+ \rightarrow \mathfrak{E}_n(\mathfrak{S})$ is a group homomorphism. Since the preimage of a subgroup under a homomorphism

is again a subgroup, it follows that $\mathfrak{S}_{ij} := \langle S_{ij}, + \rangle$ forms a subgroup of \mathfrak{S}_+ . Hence, so does the intersection H . Furthermore, since $\mathfrak{G}_3(\mathfrak{S})/\mathfrak{A}$ is finite, so is

$$g_{ij}^{-1}[\mathfrak{G}_3(\mathfrak{S})]/g_{ij}^{-1}[\mathfrak{A}] \cong \mathfrak{S}_+/ \mathfrak{S}_{ij}.$$

Since every subgroup \mathfrak{S}_{ij} has finite index in \mathfrak{S}_+ , so does the subgroup induced by H .

(b) The fact that the quotient \mathfrak{S}/H (if it exists) is finite, follows immediately from (a). Furthermore, we have already shown in (a) that H is closed under addition. To prove that H is an ideal, it therefore remains to prove closure under multiplication with elements of S . Hence, fix $h \in H$ and $s \in S$. We have to show that $hs, sh \in H$. Since \mathfrak{A} is a normal subgroup of $\mathfrak{G}_3(\mathfrak{S})$ (and therefore closed under conjugation) and since $E_{ij}(h) \in A$, we have

$$M^{-1} \cdot E_{ij}(h) \cdot M \in A, \quad \text{for all } M \in E_n(\mathfrak{S}).$$

This implies that

$$[E_{ij}(h), E_{ij}(s)] = E_{ij}(h)^{-1} \cdot (E_{ij}(s)^{-1} \cdot E_{ij}(h) \cdot E_{ij}(s)) \in A.$$

Hence, it follows by Lemma 10.16 (c) that, for i, j, k pairwise distinct,

$$E_{ik}(hs) = [E_{ij}(h), E_{jk}(s)] \in A,$$

$$E_{ik}(sh) = [E_{ij}(s), E_{jk}(h)] \in A.$$

Consequently, $hs, sh \in S_{ik}$. As this holds for all $i \neq k$, we have $hs, sh \in H$, as desired.

(c) Let $h, k \in H$. Then $E_{12}(h), E_{23}(k) \in A$. Since \mathfrak{A} is abelian, it follows by Lemma 10.16 (c) that

$$E_{13}(hk) = [E_{12}(h), E_{23}(k)] = I.$$

This implies that $hk = \mathbf{o}$. □

Unfortunately, this theorem does not provide a complete characterisation of automatic rings. Let us collect a few more restrictions on such rings.

Lemma 10.18. *Every automatic ring has positive characteristic.*

Proof. If \mathfrak{R} has characteristic 0, the element 1 generates an infinite subring. Hence, \mathfrak{R} cannot be automatic. □

Proposition 10.19. *For every automatic ring \mathfrak{R} , there exists a finite set Q of prime powers such that we can write*

$$\mathfrak{R} = \bigoplus_{q \in Q} \mathfrak{R}_q,$$

where \mathfrak{R}_q is an automatic ring of characteristic q .

Proof. By Lemma 10.18, \mathfrak{R} has a positive characteristic m . We prove the claim by induction on m . If m is a prime power, there is nothing to do. Otherwise, we can write $m = kl$ where k and l are coprime. Hence, there are $i, j \in \mathbb{Z}$ with $ik + jl = 1$. For every $r \in R$, it follows that

$$r = (ik + jl)r = is + jt, \quad \text{where } s := kr \text{ and } t := lr.$$

This implies that $ls = mr = \mathbf{o}$ and $kt = mr = \mathbf{o}$. Consequently, setting

$$S := \{s \in R \mid ls = \mathbf{o}\} \quad \text{and} \quad T := \{t \in R \mid kt = \mathbf{o}\}$$

we obtain $R = S + T$. Furthermore, if $u \in S \cap T$, then

$$lu = \mathbf{o} = ku \quad \text{implies} \quad u = (ik + jl)u = \mathbf{o} + \mathbf{o}.$$

Thus, $S \cap T = \{\mathbf{o}\}$. Since S and T are closed under addition and multiplication, it follows that $\mathfrak{R} = \mathfrak{S} \oplus \mathfrak{T}$.

Furthermore, the fact that the sets S and T are FO-definable in \mathfrak{R} implies that \mathfrak{S} and \mathfrak{T} are automatic. By inductive hypothesis, each of \mathfrak{S} and \mathfrak{T} has a decomposition of the desired kind. Hence, so does \mathfrak{R} . □

By this proposition, it follows that in order to characterise all automatic rings, it is sufficient to characterise those that cannot be decomposed as a direct sum. The next result takes a look at such rings.

Proposition 10.20. *Let \mathfrak{R} be an automatic ring that cannot be written as the direct sum of two non-trivial automatic rings. Then*

$$R = N \cup U,$$

where

$$N := \{ r \in R \mid r^n = 0 \text{ for some } n > 0 \},$$

$$U := \{ r \in R \mid r^n = 1 \text{ for some } n > 0 \}.$$

Proof. Let \mathfrak{R} be such a ring. We start by proving that 0 and 1 are the only idempotent elements of \mathfrak{R} .

For a contradiction, suppose that there is some idempotent e different from 0 and 1. For every $r \in R$, it follows that $r = (1 - e)r + er$. Hence,

$$R = (1 - e)R \cup eR.$$

To show that $(1 - e)R \cap eR = \{0\}$, consider an element $r \in (1 - e)R \cap eR$. Then there are $s, t \in R$ with $r = (1 - e)s = et$. Hence, $(1 - e)^2 = 1 - 2e + e^2 = 1 - e$ implies that

$$r = (1 - e)s = (1 - e)^2 s = (1 - e)r = (1 - e)et = (e - e^2)t = 0.$$

Since e and $1 - e$ are idempotent, the sets $(1 - e)R$ and eR are closed under addition and multiplication. Consequently, $\mathfrak{R} = (1 - e)\mathfrak{R} \oplus e\mathfrak{R}$.

Note that, for $c \in \{e, 1 - e\}$, the formula

$$\varphi_c(x) = \exists y[x = c \cdot y]$$

defines the set cR . Consequently, $(1 - e)\mathfrak{R}$ and $e\mathfrak{R}$ are both first-order interpretable in \mathfrak{R} and, therefore, automatic. This contradicts our assumptions on \mathfrak{R} .

We conclude the proof as follows. To show that $R = N \cup U$, consider an element $r \in R$. Since \mathfrak{R} is locally finite, we can find two numbers $k, m > 0$ with $r^k = r^{k+m}$. Fixing i with $im \geq k$, it follows that $(r^{im})^2 = r^{2im} =$

r^{im} . Hence, r^{im} is idempotent, which implies that $r^{im} = 0$ or $r^{im} = 1$. Consequently, $r \in N \cup U$, as desired.

Finally, we have to show that $N \cap U = \emptyset$. For a contradiction, suppose that there is some $r \in N \cap U$. Then there are numbers $m, n > 0$, such that $r^m = 0$ and $r^n = 1$. Consequently,

$$0 = 0^n = (r^m)^n = (r^n)^m = 1^m = 1.$$

A contradiction. \square

Exercise 10.3. Show that the set U in the preceding proposition is closed under multiplication and that it is the set of units of \mathfrak{R} . \lrcorner

For commutative rings, we obtain the following corollary.

Corollary 10.21. *Let \mathfrak{R} be a commutative automatic ring that cannot be written as the direct sum of two non-trivial automatic rings. Then \mathfrak{R} is a local ring and $\mathfrak{R}/\mathfrak{m}$ is finite, where \mathfrak{m} is the maximal ideal of \mathfrak{R} .*

Proof. Let N and U be the sets from the preceding proposition. If \mathfrak{R} is commutative, N is closed by multiplication with elements of R . Hence, N forms an ideal. To show that N is the unique maximal ideal of \mathfrak{R} , consider some maximal ideal I . If $I \neq N$, then there exists some element $u \in I \setminus N \subseteq U$. Thus, I contains a unit, which implies that $I = R$. A contradiction.

It remains to show that \mathfrak{R}/N is finite. Being the set of non-units, N is FO-definable. Consequently, we can interpret the quotient \mathfrak{R}/N in \mathfrak{R} , which means that \mathfrak{R}/N is also an automatic ring. It is even a field since the ideal N is maximal. By Corollary 10.12, every automatic field is finite. \square

11 Automatic Orders

Linear Orders

Another class of automatic structures that has extensively been studied are partial orders of various kinds. We start with linear orders.

Examples. The following partial orders are automatic: (i) every finite partial order, (ii) the orders

$$\langle \omega, \leq \rangle, \quad \langle \mathbb{Z}, \leq \rangle, \quad \langle [2]^*, \leq_{\text{lex}} \rangle, \\ \langle \omega, \geq \rangle, \quad \langle \mathbb{Q}, \leq \rangle, \quad \langle [2]^*, \leq_{\text{pf}} \rangle,$$

and (iii) the order $\langle F, \subseteq \rangle$ where F is the set of all finite and cofinite subsets of ω (see Theorem 10.2).

Given these examples we can construct new automatic linear orders using the following operations.

Proposition 11.1. (a) *The classes of automatic, ω -automatic, tree automatic, and ω -tree automatic linear orders are closed under (i) finite ordered sums; (ii) finite ordered products; and (iii) dense shuffles.*

(b) *The classes of tree automatic and ω -tree automatic linear orders are also closed under (iv) ω -ordinal powers.*

Proof. We have shown in Theorem 2.2 and Proposition 3.1 that all four classes are closed under interpretations, disjoint union, and direct products. This implies (i), (ii), and (iv). Hence, it remains to prove (iii). Let $\mathfrak{A}_0 = \langle A_0, \leq_0 \rangle, \dots, \mathfrak{A}_{n-1} = \langle A_{n-1}, \leq_{n-1} \rangle$ be automatic linear orders with injective presentations $\langle L_\delta^i, L_\leq^i \rangle$ over the same alphabet Σ . Without loss of generality, we may assume that Σ and $\{0, 1\}$ are disjoint. Note that the sets

$$K_i := \{ u \circ 1^i \mid u \in [2]^* \}, \quad \text{for } i < n,$$

are dense in $\langle \bigcup_{i < n} K_i, \leq_{\text{lex}} \rangle$. We obtain a presentation of the shuffle with universe

$$K_0 L_\delta^0 \cup \dots \cup K_{n-1} L_\delta^{n-1}$$

and ordering

$$uw \leq u'w' : \text{iff } u <_{\text{lex}} u' \text{ or } u = u' \text{ and } w \leq_i w',$$

for $u, u' \in \{0, 1\}^*$ and $w, w' \in \Sigma^*$. Both of these relations are regular. \square

Corollary 11.2. *Every regular linear order is automatic.*

To prove that certain linear orders are not automatic we use some techniques from the structure theory of linear orders. Recall the notion of the generalised FC-rank from Definition VI.1.9.

Proposition 11.3. *Every automatic linear order has finite gFC-rank.*

Proof. Let \mathfrak{A} be an automatic linear order. By Theorem VI.1.11, we can write \mathfrak{A} as an ordered sum $\sum_{i \in I} \mathfrak{B}_i$, where every \mathfrak{B}_i is scattered and I is dense. Applying Proposition 8.22 to the formula $\varphi(x; z, z') := z \leq x \leq z'$, we obtain a finite set \mathcal{S} of linear orders such that every interval $[a, b]$ can be partitioned into finitely many orders from \mathcal{S} . Let α be the maximal gFC $_*$ -rank of an order in \mathcal{S} . Then Lemma VI.1.13 implies that

$$\text{FC}_*([a, b]) \leq \alpha, \quad \text{for all } a, b \in B_i \text{ with } i \in I.$$

By definition of the gFC-rank, this means that

$$\text{cn}^{\alpha+2}(a) = \text{cn}^{\alpha+2}(b), \quad \text{for all } a, b \in B_i \text{ with } i \in I.$$

Consequently, we have $\text{gFC}(\mathfrak{B}_i) \leq \text{gFC}(\mathfrak{B}_i) \leq \alpha + 1$ for all i , and it follows that

$$\text{cn}^{\alpha+2}(\mathfrak{A}) = \langle I, \leq \rangle.$$

Note that $\text{cn}(I) = I$ since I is dense. Hence, $\text{gFC}(\mathfrak{A}) \leq \text{gFC}_*(\mathfrak{A}) + 1 \leq \alpha + 2$. \square

Open Question. *Is there an automatic linear order that is not regular?*

Exercise 11.1. Let $n < \omega$ be finite. Show that the class of all linear orders of gFC-rank n is FOC-axiomatisable. \square

For ω -automatic linear orders, we only have the following remark, which follows immediately by Theorem 8.31.

Proposition 11.4. *Every uncountable ω -automatic linear order contains a regular subset U of order type $\langle [2]^\omega, \leq_{\text{lex}} \rangle$.*

Exercise 11.2. Find a 2-dimensional interpretation

$$\langle [2]^*, \leq_{\text{pf}}, \text{suc}_0, \text{suc}_1 \rangle \leq \langle [2]^\omega, \leq_{\text{lex}} \rangle. \quad \text{J}$$

Exercise 11.3. Prove that every infinite automatic linear order \mathfrak{A} contains a regular subset U of order type ω or ω^{op} . J

Ordinals

Next, let us take a look at automatic ordinals where a complete classification is possible. We start with the remark that the class of automatic ordinals is downwards closed.

Lemma 11.5. *Let α be an ordinal. If $\langle \alpha, \leq \rangle$ is automatic or tree-automatic then so is $\langle \beta, \leq \rangle$, for every $\beta \leq \alpha$.*

Proof. For $\beta = \alpha$, the claim is trivial. Otherwise, we can interpret $\langle \beta, \leq \rangle$ in $\langle \alpha, \leq, k \rangle$, where $k := \min(\alpha \setminus \beta)$. □

Corollary 11.6. *Every ordinal $\alpha < \omega^\omega$ is automatic. Every ordinal $\alpha < \omega^{\omega^\omega}$ is tree automatic.*

Proof. By Proposition 11.1, every ordinal of the form ω^n with $n < \omega$ is automatic, and every ordinal of the form ω^{ω^n} is tree automatic. By the preceding lemma, so is every ordinal bounded by an ordinal of this form. □

For the converse, we require some terminology.

Definition 11.7. (a) The *Cantor normal form* of an ordinal ξ is the expression

$$\xi = \omega^{\alpha_0} + \cdots + \omega^{\alpha_{m-1}} \quad \text{where} \quad \alpha_0 \geq \cdots \geq \alpha_{m-1}.$$

(b) For two ordinals $\xi = \omega^{\alpha_0} + \cdots + \omega^{\alpha_{m-1}}$ and $\zeta = \omega^{\beta_0} + \cdots + \omega^{\beta_{n-1}}$ in Cantor normal form, we define the *natural sum* \boxplus and the *natural product* \boxtimes by

$$\xi \boxplus \zeta := \omega^{\gamma_0} + \cdots + \omega^{\gamma_{m+n-1}} \quad \text{and} \quad \xi \boxtimes \zeta := \bigoplus_{\substack{i < m \\ j < n}} \omega^{\alpha_i \boxplus \beta_j},$$

where $\gamma_0 \geq \cdots \geq \gamma_{m+n-1}$ is a decreasing enumeration of the sequence $\alpha_0, \dots, \alpha_{m-1}, \beta_0, \dots, \beta_{n-1}$.

(c) The *height* $h(\mathfrak{A})$ of a partial order \mathfrak{A} is the minimal ordinal α such that there exists a function $\rho : \mathfrak{A} \rightarrow \alpha$ with $a < b \Rightarrow \rho(a) < \rho(b)$. J

The operations \boxplus and \boxtimes are closely related to partitions and superpositions of orders.

Exercise 11.4. Prove the following statements.

(a) $\alpha, \beta < \omega^\gamma$ implies $\alpha \boxplus \beta < \omega^\gamma$.

(b) $\alpha, \beta < \omega^{\omega^\gamma}$ implies $\alpha \boxtimes \beta < \omega^{\omega^\gamma}$. J

Exercise 11.5. Suppose that $\alpha \boxplus \beta = \gamma$. Prove that there exists a partition $P \cup Q = \gamma$ such that the order type of P is equal to α and that of Q equal to β . J

Lemma 11.8. *Let $P_0 \cup \cdots \cup P_{n-1}$ be a finite partition of a partial order \mathfrak{A} . Then*

$$h(\mathfrak{A}) \leq h(P_0) \boxplus \cdots \boxplus h(P_{n-1}).$$

Proof. It is sufficient to prove the claim for $n = 2$. Set $\xi_i := h(P_i)$ and let $\rho_i : P_i \rightarrow \xi_i$ be the corresponding functions. We define a function $\sigma : A \rightarrow \xi_0 \boxplus \xi_1$ by

$$\sigma(a) := \rho_0(\downarrow a \cap P_0) \boxplus \rho_1(\downarrow a \cap P_1).$$

By definition of $h(\mathfrak{A})$ it is sufficient to prove that

$$a < b \quad \text{implies} \quad \sigma(a) < \sigma(b).$$

Hence, suppose that $a < b$. By symmetry, we may assume that $a \in P_0$. Then $\rho_0(\downarrow a \cap P_0) < \rho_0(\downarrow b \cap P_0)$ implies that

$$\begin{aligned} \sigma(a) &= \rho_0(\downarrow a \cap P_0) \boxplus \rho_1(\downarrow a \cap P_1) \\ &< \rho_0(\downarrow b \cap P_0) \boxplus \rho_1(\downarrow a \cap P_1) \\ &\leq \rho_0(\downarrow b \cap P_0) \boxplus \rho_1(\downarrow b \cap P_1) = \sigma(b). \end{aligned} \quad \square$$

Proposition 11.9. *Let $\mathfrak{A} \cong \mathfrak{B}_0 \cup \dots \cup \mathfrak{B}_{m-1}$ be a superposition of partial orders. Then*

$$h(\mathfrak{A}) \leq h(\mathfrak{B}_0) \boxtimes \dots \boxtimes h(\mathfrak{B}_{m-1}).$$

Proof. Since \boxtimes is monotone, it is sufficient to prove that

$$h(\mathfrak{B}_i) \leq \xi_i, \text{ for all } i, \quad \text{implies} \quad h(\mathfrak{A}) \leq \xi_0 \boxtimes \dots \boxtimes \xi_{m-1}.$$

We proceed by induction on ξ_0, \dots, ξ_{m-1} (w.r.t. the componentwise ordering). There are three cases to distinguish.

If $\xi_k = 0$ for some k , then \mathfrak{B}_k is empty and so is \mathfrak{A} .

Next, suppose that some of the ξ_k is of the form $\xi_k = \zeta \boxplus \zeta'$ for some $\zeta, \zeta' < \xi_k$. By Exercise 11.5, there exists a partition $\xi_k = P \cup P'$ such that P has order type ζ and P' has order type ζ' . Let $\rho : B_k \rightarrow \xi_k$ be the function witnessing that $h(\mathfrak{B}_k) = \xi_k$. Using the bijections between B_k, A , and B_i , obtain functions $\rho_i : B_i \rightarrow \xi_k$. Let \mathfrak{C}_i and \mathfrak{C}'_i be the substructures of \mathfrak{B}_i induced by, respectively, $\rho_i^{-1}[P]$ and $\rho_i^{-1}[P']$, and let \mathfrak{D} and \mathfrak{D}' be the corresponding substructures of \mathfrak{A} . Then

$$\mathfrak{D} = \mathfrak{C}_0 \cup \dots \cup \mathfrak{C}_{m-1} \quad \text{and} \quad \mathfrak{D}' = \mathfrak{C}'_0 \cup \dots \cup \mathfrak{C}'_{m-1}.$$

Furthermore, we have $h(\mathfrak{C}_k) = \zeta$ and $h(\mathfrak{C}'_k) = \zeta'$ and $h(\mathfrak{C}_i), h(\mathfrak{C}'_i) \leq h(\mathfrak{B}_i)$, for all i . By inductive hypothesis, it follows that

$$\begin{aligned} h(\mathfrak{D}) &\leq \xi_0 \boxtimes \dots \boxtimes \zeta \boxtimes \dots \boxtimes \xi_{m-1} \\ h(\mathfrak{D}') &\leq \xi_0 \boxtimes \dots \boxtimes \zeta' \boxtimes \dots \boxtimes \xi_{m-1}. \end{aligned}$$

Consequently,

$$\begin{aligned} h(\mathfrak{A}) &\leq h(\mathfrak{D}) \boxplus h(\mathfrak{D}') \\ &\leq (\xi_0 \boxtimes \dots \boxtimes \zeta \boxtimes \dots \boxtimes \xi_{m-1}) \boxplus (\xi_0 \boxtimes \dots \boxtimes \zeta' \boxtimes \dots \boxtimes \xi_{m-1}) \\ &\leq \xi_0 \boxtimes \dots \boxtimes (\zeta \boxplus \zeta') \boxtimes \dots \boxtimes \xi_{m-1} \\ &\leq \xi_0 \boxtimes \dots \boxtimes \xi_k \boxtimes \dots \boxtimes \xi_{m-1}. \end{aligned}$$

It remains to consider the case, where no ordinal ξ_k can be written as a natural sum of two other ordinals. For every $a \in A$, we will prove that the substructure induced by $\downarrow a$ has height

$$h(\mathfrak{A}|_{\downarrow a}) < \xi_0 \boxtimes \dots \boxtimes \xi_{m-1}.$$

Then

$$h(\mathfrak{A}) = \sup_{a \in A} h(\downarrow a) \leq \xi_0 \boxtimes \dots \boxtimes \xi_{m-1},$$

as desired.

Let \mathfrak{C}_{ik} be the substructure of \mathfrak{B}_i induced by

$$\downarrow_k a := \{ b \in B_i \mid b <_k a \},$$

where $<_k$ is the order of \mathfrak{B}_k . Set $\zeta_i := \rho_i(a)$, where $\rho_i : B_i \rightarrow \xi_i$ is the function witnessing that $h(\mathfrak{B}_i) = \xi_i$. Then

$$\begin{aligned} \zeta_i &< \xi_i \\ h(\mathfrak{C}_{kk}) &\leq \rho_k(a) = \zeta_k, \\ h(\mathfrak{C}_{ik}) &\leq h(\mathfrak{A}) = \xi_i, \quad \text{for } i \neq k. \end{aligned}$$

Note that $B_i = B_j$ implies $\mathfrak{C}_{ik} = \mathfrak{C}_{jk}$. Hence, the restriction of \mathfrak{A} to $\downarrow_k a$ is the superposition of $\mathfrak{C}_{0k}, \dots, \mathfrak{C}_{(m-1)k}$ and the inductive hypothesis implies that

$$h(\mathfrak{A}|_{\downarrow_k a}) \leq \xi_0 \boxtimes \dots \boxtimes \zeta_k \boxtimes \dots \boxtimes \xi_{m-1}.$$

Since $\downarrow a = \downarrow_0 a \cup \dots \cup \downarrow_{m-1} a$ it therefore follows by Lemma 11.8 that

$$\begin{aligned} h(\mathfrak{A}|_{\downarrow a}) &\leq \bigoplus_{k < m} (\xi_0 \boxtimes \dots \boxtimes \zeta_k \boxtimes \dots \boxtimes \xi_{m-1}) \\ &\leq \xi_0 \boxtimes \dots \boxtimes \left(\bigoplus_{k < m} \zeta_k \right) \boxtimes \dots \boxtimes \xi_{m-1} \\ &< \xi_0 \boxtimes \dots \boxtimes \xi_k \boxtimes \dots \boxtimes \xi_{m-1}. \end{aligned}$$

□

Proposition 11.10. *The linear order $\langle \omega^{\omega^{\omega}}, \leq \rangle$ is not tree-automatic.*

Proof. We apply Theorem 8.25 to the formula $\varphi(x; z) := x < z$. Let \mathcal{S} be the corresponding class of structures. It follows that, for every $n < \omega$, the initial segment ω^{ω^n} has a sum-decomposition over $\text{supp}(\text{prod}(\mathcal{S}))$. By Exercise 11.4 and Lemma 11.8, it follows that

$$\omega^{\omega^n} \in \text{supp}(\text{prod}(\mathcal{S})), \quad \text{for all } n < \omega.$$

We start by showing that we can assume without loss of generality that the structures in \mathcal{S} are partial orders. First, note that every finite partition of a linear order consists of linear orders. Hence, we are only interested in the linear orders contained in $\text{supp}(\text{prod}(\mathcal{S}))$. Furthermore, if a linear order \mathcal{A} is a superposition of structures $\mathcal{C}_0, \dots, \mathcal{C}_{n-1}$, we can replace the relation \leq_i of \mathcal{C}_i by its transitive and reflexive closure without changing the superposition. Thus, if we replace every structure in \mathcal{S} by its transitive and reflexive closure, then class of linear orders contained in $\text{supp}(\text{prod}(\mathcal{S}))$ is not changed.

Consequently, modifying \mathcal{S} in this way we may assume without loss of generality that every structure in \mathcal{S} is a partial order. We will prove below that, for every class \mathcal{C} of partial orders,

$$\omega^{\omega^n} \in \text{supp}(\mathcal{C}) \quad \text{implies} \quad \omega^{\omega^n} \in \mathcal{C}.$$

Since the only linear orders in $\text{prod}(\mathcal{S})$ are the structures in \mathcal{S} , it then follows that $\omega^{\omega^n} \in \mathcal{S}$, for all $n < \omega$. A contradiction to the finiteness of \mathcal{S} .

Hence, it remains to prove the claim. Suppose that ω^{ω^n} is the suposition of partial orders $\mathcal{A}_0, \dots, \mathcal{A}_{m-1}$. For a contradiction, suppose that no \mathcal{A}_i is isomorphic to ω^{ω^n} . This implies that $h(\mathcal{A}_i) < \omega^{\omega^n}$, for all i . As $\alpha, \beta < \omega^{\omega^n}$ implies $\alpha \boxtimes \beta < \omega^{\omega^n}$, it follows by Proposition 11.9 that

$$\omega^{\omega^n} = h(\omega^{\omega^n}) \leq h(\mathcal{A}_0) \boxtimes \dots \boxtimes h(\mathcal{A}_{m-1}) < \omega^{\omega^n}.$$

A contradiction. \square

We obtain the following classification of all automatic ordinals.

Theorem 11.11. *Let α be an ordinal.*

- (a) $\langle \alpha, \leq \rangle$ is automatic if, and only if, $\alpha < \omega^{\omega}$.
- (b) $\langle \alpha, \leq \rangle$ is ω -automatic if, and only if, $\alpha < \omega^{\omega}$.
- (c) $\langle \alpha, \leq \rangle$ is tree automatic if, and only if, $\alpha < \omega^{\omega^{\omega}}$.
- (d) $\langle \alpha, \leq \rangle$ has an injective ω -tree automatic presentation if, and only if, $\alpha < \omega^{\omega^{\omega}}$.

Proof. (a), (c) One direction has already been proved in Corollary 11.6. For the other one, note that we have shown in Proposition 8.23 that ω^{ω} is not automatic, and in Proposition 11.10 that $\omega^{\omega^{\omega}}$ is not tree-automatic.

(b) If there were an uncountable ω -automatic ordinal $\langle \alpha, \leq \rangle$ we could use Proposition 11.4 to find some subset $U \subseteq \alpha$ of order type $\langle [2]^*, \leq_{\text{lex}} \rangle$. But the latter is not a well-order. A contradiction. Thus, all ω -automatic ordinals are countable and, therefore, automatic by Theorem 6.14. Hence, the claim follows by (a).

(d) One direction follows by (c). For the other one, let α be an ordinal with an injective ω -tree automatic presentation. If α is countable, it follows by Theorem 6.15 that it is tree automatic and the claim follows by (c). Hence, it is sufficient to prove that no model V of ZFC contains uncountable ordinals with an injective ω -tree automatic presentation. For a contradiction, suppose that there is a model V in which some uncountable ordinal α has such a presentation. As the element ω_1 is FOC-definable in every uncountable ordinal $\langle \alpha, \leq \rangle$, we can interpret $\langle \omega_1, \leq \rangle$ in $\langle \alpha, \leq \rangle$. This implies that ω_1 also has an injective ω -tree automatic presentation. Let \mathcal{A}_{δ} and \mathcal{A}_{\leq} be the two automata for this presentation. Since the universe of $\langle \omega_1, \leq \rangle$ is an uncountable language of infinite trees, it follows by Theorem V.6.5 that its cardinality is 2^{\aleph_0} . This implies the Continuum Hypothesis (CH): $2^{\aleph_0} = |\omega_1| = \aleph_1$. Hence, $V \models \text{CH}$.

By a known result of set theory, every model V of ZFC + CH has an extensions $V^+ \supset V$ that is a model of ZFC + $\neg \text{CH}$. We fix such an extension V^+ . In V^+ , the automata \mathcal{A}_{δ} and \mathcal{A}_{\leq} form an injective presentation of some structure $\langle A, \leq \rangle$. Note that the decision procedure for the FOC-theory of an ω -tree automatic structure in Theorem 7.5 works entirely in ZFC. Con-

sequently, $\langle A, \leq \rangle$ has the same FOC-theory as $\langle \omega_1, \leq \rangle$. In particular, A forms an uncountable linear order where each proper initial segment is countable. By Theorem V.6.5, this implies that A has cardinality $2^{\aleph_0} > \aleph_1$. Choose a strictly increasing sequence $(a_i)_{i < \gamma}$ of maximal length γ . If γ were countable, then $A = \bigcup_{i < \gamma} \downarrow a_i$ would be a countable union of countable sets and, hence, countable. A contradiction. Hence, $\gamma \geq \omega_1$ and the union

$$I := \bigcup_{i < \omega_1} \downarrow a_i$$

forms an initial segment of A of size $|I| = \aleph_1 < 2^{\aleph_0} = |A|$. Consequently, we can find an element $b \in A$ with $I \subseteq \downarrow b$. A contradiction to the fact that every initial segment of A is countable. \square

Exercise 11.6. Show that every automatic ordinal is finitely FOC-axiomatisable. \lrcorner

Exercise 11.7. Show that the isomorphism problem for automatic ordinals is decidable. \lrcorner

Trees

Next, let us take a look at order-trees. Recall the definition of the generalised Lifsches-Shelah rank of a tree from Definition VI.1.20.


Lemma 11.12. *Let \mathfrak{T} be a finitely-branching automatic order-tree. Then we have $\text{gLS}(\mathfrak{T}) < \omega$.*

Proof. Fix an injective presentation of \mathfrak{T} over some alphabet Σ . First, note that we can use the length-lexicographic ordering on Σ^* to turn \mathfrak{T} into a successor-ordered tree. With the help of this successor-ordering, we can define the lexicographic ordering \leq_{lex} on T (which is different from the lexicographic ordering on Σ^*). Since these relations are FO-definable in $\langle T, \leq_{\text{pf}}, \leq_{\text{lex}} \rangle$, the linear order $\langle T, \leq_{\text{lex}} \rangle$ is automatic. By Proposition 11.3, it has finite gFC-rank. Consequently, the claim follows by Proposition VI.1.21. \square

Proposition 11.13. *Every automatic order-tree has a finite generalised Lifsches-Shelah rank.*

Proof. Fix an injective presentation of \mathfrak{T} over some alphabet Σ . First, note that we can use the length-lexicographic ordering on Σ^* to define well-orderings on the successors of each vertex. We denote by $p(v)$ the predecessor of v and by $s(v)$ its first successor (if they exist). Furthermore, we define the ‘first-child-next-sibling’ relation

$$u \leq' v \quad : \text{iff} \quad u \leq v \quad \text{or} \quad p(u) < v \text{ and } u <_{\text{lex}} v.$$



Note that \leq' is a partial order that contains \leq and that is contained in \leq_{lex} . Furthermore, $\mathfrak{T}' := \langle T, \leq' \rangle$ is a tree where every vertex v has at most two successors: the first successor of v in \mathfrak{T} and the next sibling of v in \mathfrak{T} . Since \leq' is FOC-definable, \mathfrak{T}' is automatic and it follows by Lemma 11.12 that its gLS-rank is finite. To conclude the proof, we show that $\text{gLS}(\mathfrak{T}) \leq \text{gLS}(\mathfrak{T}')$.

Let $[\mathfrak{T}]$ and $[\mathfrak{T}']$ be the sets of infinite branches of the respective trees. Since every infinite branch of \mathfrak{T} induces one of \mathfrak{T}' , there exists an injective function $e : [\mathfrak{T}] \rightarrow [\mathfrak{T}']$. It follows that

$$v \in \zeta \quad \text{implies} \quad s(v) \in e(\zeta), \quad \text{for all } \zeta \in [\mathfrak{T}].$$

By induction on n , we prove that

$$v \in \partial^n \mathfrak{T} \quad \text{implies} \quad s(v) \in \partial^n \mathfrak{T}', \quad \text{for all } v \text{ such that } s(v) \text{ is defined.}$$

For $n = 0$, the claim is trivial. For the inductive step, suppose that $v \in \partial^{n+1} \mathfrak{T}$. By definition of ∂ , there exists infinitely many infinite branches ζ_i , $i < \omega$, of $\partial^n \mathfrak{T}$ that contain v . Consequently, $s(v)$ belongs to $e(\zeta_i)$. Furthermore, $e(\zeta_i)$ belong to $\partial^n \mathfrak{T}'$ by inductive hypothesis. Since e is injective, it follows

that $s(v)$ belongs to infinitely many infinite branches of $\partial^n \mathfrak{T}'$. This means that $s(v) \in \partial^{n+1} \mathfrak{T}'$.

For $n := \text{gLS}(\mathfrak{T}')$ it follows that

$$s(v) \in \partial^{n+1} \mathfrak{T}' \quad \text{implies} \quad |\{ \zeta' \in [\mathfrak{T}'] \mid s(v) \in \zeta' \}| > \aleph_0.$$

Every branch in $[\mathfrak{T}'] \setminus \text{rng } e$ consists of some finite prefix leading to some vertex v with infinitely many successors (in \mathfrak{T}) followed by an enumeration of the successors of v . In particular, the number of these additional branches is countable. For v as above, it therefore follows that

$$|\{ \zeta \in [\mathfrak{T}] \mid v \in \zeta \}| > \aleph_0.$$

Hence, $\mathfrak{T}|_v$ has uncountably many infinite branches. Conversely, every $v \in T$ such that $\mathfrak{T}|_v$ has uncountably many infinite branches belongs to $\partial^{n+1} \mathfrak{T}$. It follows that

$$\partial^{n+1} \mathfrak{T} = \{ v \in T \mid \mathfrak{T}|_v \text{ has uncountably many infinite branches} \}.$$

This implies that $\partial^{n+2} \mathfrak{T} = \partial^{n+1} \mathfrak{T}$. Thus, $\text{gLS}(\mathfrak{T}) \leq n < \omega$. \square

Next, let us turn to definability questions concerning branches in automatic trees.

Lemma 11.14. *Let \mathfrak{T} be an automatic order-tree. The sets*

$$U := \{ v \in T \mid v \in \zeta \text{ for some } \zeta \in [T] \} \quad \text{and} \quad \partial^n T, \quad \text{for } n < \omega,$$

are FOC[H]-definable.

Proof. The formula

$$\varphi(x) := \text{H}[y, y'] [x \leq y \wedge x \leq y' \wedge (y \leq y' \vee y' \leq y)]$$

states that x lies on an infinite branch. Hence, $U = \varphi^{\mathfrak{T}}$ is definable.

To prove that $\partial^n T$ is definable, it is sufficient to show that ∂T is definable. Then the claim follows by induction. Let $\varphi(x)$ be the FOC[H]-formula

from Lemma 11.14. Set $U := \varphi^{\mathfrak{T}}$ and let $U_o \subseteq U$ be the subset of all vertices v such that

$$U_o := \{ u \in U \mid \text{there are incomparable vertices } v, v' \text{ with } u \leq v, v' \}.$$

Then

$$\partial T = \{ u \in U_o \mid u \in \zeta \text{ for some } \zeta \in [U_o] \}.$$

This set is definable using the formula φ . \square

Proposition 11.15. *Let \mathfrak{T} be an automatic order-tree.*

- (a) *If \mathfrak{T} has infinite branches, at least one of them is regular.*
- (b) *If \mathfrak{T} has only countably many infinite branches, every infinite branch is regular.*
- (c) *If \mathfrak{T} has only countably many infinite branches, there exists an FOC[H]-formula $\varphi(x; z)$ such that*

$$\{ \varphi(x, v)^{\mathfrak{T}} \mid v \in T \} = \{ \emptyset \} \cup \{ P \subseteq T \mid P \text{ an infinite branch of } \mathfrak{T} \}.$$

Proof. (a) Let U be the set from Lemma 11.14 and let \leq_{lex} be the lexicographic order on Σ^* . Then $\langle \mathfrak{T}, \leq_{\text{lex}}, U \rangle$ is automatic and the left-most (with respect to \leq_{lex}) infinite branch is FOC-definable in $\langle \mathfrak{T}, \leq_{\text{lex}}, U \rangle$.

(b) follows immediately from (c).

(c) We have shown in Lemma 11.14 that, for every $n < \omega$, the set

$$P_n := \{ v \in T \mid \text{gLS}(T|_v) \geq n \}$$

is FOC[H]-definable. Let $n := \text{gLS}(\mathfrak{T}) < \omega$ and set

$$\begin{aligned} \varphi(x; z) := \bigvee_{k < n} [& z \in P_k \setminus P_{k+1} \wedge x \in P_k \wedge (x \leq z \vee z \leq x) \\ & \wedge \forall y \forall y' [P_k y \wedge P_k y' \wedge z \leq y \wedge z \leq y' \\ & \rightarrow (y \leq y' \vee y' \leq y)]]. \end{aligned}$$

For every $v \in T$, it follows that $\varphi(x; v)^{\mathbb{Z}}$ is either empty or an infinite branch. Conversely, let β be an infinite branch. Then there is some $k \leq n$ such that all but finitely many elements of β belong to $P_k \setminus P_{k+1}$. Furthermore, there exists some vertex $v \in \beta$ such that all elements of $P_k \cap \beta$ that are greater than v form a branch. Then $\varphi(x; v)^{\mathbb{Z}} = \beta$. \square

12 Subclasses of Automatic Structures

As we have seen, characterising which structures have an automatic presentation is a very hard problem. To simplify the task, we introduce several subclasses of automatic structures where it is easier to prove characterisations. We start with the following one.

Definition 12.1. A language $L \subseteq \Sigma^*$ has *polynomial growth* if there exists a polynomial $p(x)$ such that

$$|\{w \in L \mid |w| \leq n\}| \leq p(n), \quad \text{for all } n < \omega.$$

Similarly, we say that an automatic structure \mathcal{A} has *polynomial growth* if there exists a polynomial $p(x)$ such that

$$|\{a \in A \mid \|a\| \leq n\}| \leq p(n), \quad \text{for all } n < \omega.$$

In this case we also say that \mathcal{A} is *poly-growth automatic*. \lrcorner

Example. The infinite grid $\langle \mathbb{Z} \times \mathbb{Z}, E_0, E_1 \rangle$ is poly-growth automatic. (We can represent a point $\langle i, k \rangle \in \mathbb{Z} \times \mathbb{Z}$ by the word $a^i b^k$.) \lrcorner

We start with a characterisation of when a regular language has polynomial growth.

Proposition 12.2. Let $L \subseteq \Sigma^*$ be regular. The following statements are equivalent.

- (1) L has polynomial growth.
- (2) L has less than exponential growth.

- (3) L is a finite union of languages of the form

$$u_0 v_0^* u_1 v_1^* \cdots u_{k-1} v_{k-1}^* u_k \quad \text{with } u_0, \dots, u_k, v_0, \dots, v_{k-1} \in \Sigma^*.$$

- (4) L is recognised by an automaton \mathcal{A} such that every strongly connected component of the transition graph of \mathcal{A} is either an induced cycle or a singleton with at most one loop.

Proof. (3) \Rightarrow (1) \Rightarrow (2) is trivial.

(2) \Rightarrow (4) Let \mathcal{A} be the minimal deterministic automaton recognising L . For a contradiction, suppose that \mathcal{A} contains a strongly connected component C that is neither a singleton nor an induced cycle. Let u be a word labelling some path from the initial state to some state $q \in C$. By choice of C there are two distinct paths in C that both start and end in q . Let v_0 and v_1 be their labels. Finally, there is some path from q to some final state. Let w be its label. Fixing some number n with

$$n \geq |u|, |v_0|, |v_1|, |w|,$$

it follows that

$$u(v_0 + v_1)^* w \subseteq L,$$

which implies that, for all $2 \leq k < \omega$, L contains at least 2^{k-2} words of length at most kn . Consequently,

$$|\{w \in L \mid |w| \leq k\}| \geq 2^{k/n-2}, \quad \text{for all } k < \omega.$$

In particular, L has exponential growth. A contradiction.

(4) \Rightarrow (3) Fix an automaton \mathcal{A} of the above form. We prove the claim by induction on the number of strongly connected components of \mathcal{A} . If \mathcal{A} has a single state, L is of one of the following forms

$$\emptyset, \quad \{()\}, \quad \text{or} \quad a^*.$$

If \mathcal{A} consists of a single cycle, we have

$$L = \bigcup_{i < n} w^* u_i,$$

where w is the word labelling the whole cycle and each u_i is the word labelling some part of the cycle starting at the initial state and ending at some final state.

For the inductive step, suppose that \mathcal{A} has more than one strongly connected component. Let C be the component containing the initial state. Set $n := |C|$ and let u_i , for $i \leq n$, be the path from the initial state to the i -th state of C . For $i < n$, let $a_{i,0}, \dots, a_{i,m_i-1} \in \Sigma$ be an enumeration of all letters such that the transition with label $a_{i,j}$ from the i -th state of C leads to some state $q_{i,j}$ that does not belong to C and let $K_{i,j}$ be the language recognised by \mathcal{A} when starting in state $q_{i,j}$. Finally, let $F \subseteq [n]$ be the set of numbers i such that the i -th state of C is accepting. Then

$$L = \bigcup_{i \in F} u_n^* u_i \cup \bigcup_{i < n} \bigcup_{j < m_i} u_n^* u_i a_{i,j} K_{i,j}.$$

Furthermore, each language $K_{i,j}$ is of the form required by (3). Hence, so is L . \square

We can use this lemma to derive a corresponding characterisation for poly-growth automatic structures. This characterisation also contains a description in logical terms whose proof rests on the following two technical lemmas.

Lemma 12.3. *For all $m, r, k < \omega$, there exists a constant $N < \omega$ with the following property. Given tuples $\tilde{n}, \tilde{n}' \in \omega^k$ satisfying*

- ♦ $n_0 \leq \dots \leq n_{k-1}$ and $n'_0 \leq \dots \leq n'_{k-1}$,
- ♦ $n_{i+1} - n_i = n'_{i+1} - n'_i$ or $n_{i+1} - n_i, n'_{i+1} - n'_i \geq N$, for all $i < k-1$,

we have

$$\langle \omega, \leq, m \mid \cdot, \tilde{n} \rangle \equiv_{\text{FO}}^r \langle \omega, \leq, m \mid \cdot, \tilde{n}' \rangle.$$

Proof. We prove the claim by induction on the quantifier-rank r . For $r = 0$, the claim is trivial. Hence, suppose that $r > 0$. Let N_0 be the constant from the inductive hypothesis and set $N := 2N_0$. Consider tuples $\tilde{n}, \tilde{n}' \in \omega^k$ as above. To show that they are r -equivalent, we check the back-and-forth

condition. Hence, fix $l < \omega$. Let i be the index such that $n_i \leq l \leq n_{i+1}$. (The cases where $l < n_0$ or $l > n_{k-1}$ are analogous.) If $n_{i+1} - n_i = n'_{i+1} - n'_i$, we can choose $l' := n'_i + (l - n_i)$ and it follows by inductive hypothesis that tuples $\tilde{n}l$ and $\tilde{n}'l$ are $(r-1)$ -equivalent. Hence, suppose that $n_{i+1} - n_i, n'_{i+1} - n'_i \geq N$. By symmetry, we may assume that $l - n_i \leq n_{i+1} - l$. If $l - n_i < N_0$, we set $l' := n'_i + (l - n_i)$. Otherwise, we use $l' := n'_i + N_0$. In both cases, it follows by inductive hypothesis that $\tilde{n}l$ and $\tilde{n}'l'$ are $(r-1)$ -equivalent. \square

Lemma 12.4. *Let $L \subseteq \Sigma^*$ be regular and $u_0, \dots, u_n, v_0, \dots, v_{n-1} \in \Sigma^*$. There exist a constant m and a first-order formula $\varphi(\bar{x})$ such that*

$$u_0 v_0^{k_0} u_1 v_1^{k_1} \dots u_{n-1} v_{n-1}^{k_{n-1}} u_n \in L \quad \text{iff} \quad \langle \omega, \leq, m \mid \cdot \rangle \models \varphi(k_0, k_0 + k_1, \dots, k_0 + \dots + k_{n-1}).$$

Proof. Let $\varphi : \Sigma^* \rightarrow \mathfrak{M}$ be a homomorphism to a finite monoid \mathfrak{M} recognising L . Setting $c_i := \varphi(u_i)$, $d_i := \varphi(v_i)$, and $P := \varphi[L]$ it follows that

$$u_0 v_0^{k_0} u_1 v_1^{k_1} \dots u_{n-1} v_{n-1}^{k_{n-1}} u_n \in L \quad \text{iff} \quad c_0 d_0^{k_0} c_1 d_1^{k_1} \dots c_{n-1} d_{n-1}^{k_{n-1}} c_n \in P.$$

It is therefore sufficient to show that, for every set $P \subseteq M$, all elements $c_0, \dots, c_n, d_0, \dots, d_{n-1} \in M$, and every number $m < \omega$ such that d_i^m is idempotent for all $i < n$, there exists an FO-formula $\varphi(\bar{x})$ such that

$$c_0 d_0^{k_0} c_1 d_1^{k_1} \dots c_{n-1} d_{n-1}^{k_{n-1}} c_n \in P \quad \text{iff} \quad \langle \omega, \leq, m \mid \cdot \rangle \models \varphi(k_0, k_0 + k_1, \dots, k_0 + \dots + k_{n-1}).$$

We construct the desired formula φ by induction on n . For $n = 1$, set

$$K := \{ k < 2m \mid c_0 d_0^k c_1 \in P \}.$$

Then we have $c_0 d_0^{k_0} c_1 \in P$ if, and only if, either

- ♦ $k_0 < m$ and $k_0 \in K$, or
- ♦ $k_0 \geq m$ and $k_0 \equiv k' \pmod{m}$ for some $k' \in K$ with $k' \geq m$.

This property can be expressed in FO.

For the inductive step, suppose that $n > 1$. Set

$$Q_i := \{ b \in M \mid b d_{n-1}^i c_n \in P \}, \quad \text{for } i < 2m.$$

By inductive hypothesis, there exist formulae ψ_i , $i < 2m$, such that

$$c_o d_o^{k_o} c_1 d_1^{k_1} \dots c_{n-2} d_{n-2}^{k_{n-2}} c_{n-1} \in Q_i \quad \text{iff} \\ \langle \omega, \leq, m \mid \cdot \rangle \models \psi_i(k_o, k_o + k_1, \dots, k_o + \dots + k_{n-2}),$$

It follows that

$$c_o d_o^{k_o} c_1 d_1^{k_1} \dots c_{n-1} d_{n-1}^{k_{n-1}} c_n \in P$$

if, and only if, there is some $i < 2m$ such that

$$c_o d_o^{k_o} c_1 d_1^{k_1} \dots c_{n-2} d_{n-2}^{k_{n-2}} c_{n-1} \in Q_i$$

and either

- $i < m$ and $k_o + \dots + k_{n-1} = k_o + \dots + k_{n-2} + i$, or
- $i \geq m$ and there is some $j < m$ such that $k_o + \dots + k_{n-2} \equiv j \pmod{m}$ and $k_o + \dots + k_{n-1} \equiv j + i \pmod{m}$.

Using the formulae ψ_i , this property can be expressed in FO. \square

Theorem 12.5. *Let \mathfrak{A} be a structure. The following statements are equivalent.*

- (1) \mathfrak{A} is an automatic structure of polynomial growth.
- (2) \mathfrak{A} has an automatic presentation whose universe is a finite union of languages of the form

$$u_o v_o^* u_1 v_1^* \dots u_{k-1} v_{k-1}^* u_k, \quad \text{with } u_o, \dots, u_k, v_o, \dots, v_{k-1} \in \Sigma^*.$$

- (3) \mathfrak{A} has an automatic presentation whose universe is a finite union of languages of the form

$$a_o^{m_o} (b_o^{k_o})^* a_1^{m_1} (b_1^{k_1})^* \dots a_{n-1}^{m_{n-1}} (b_{n-1}^{k_{n-1}})^* a_n^{m_n}$$

for distinct letters $a_o, \dots, a_n, b_o, \dots, b_{n-1}$ (with distinct members of the union using disjoint alphabets).

(4) \mathfrak{A} is (k -dimensional) FO-interpretable in $\langle \omega, \leq, m \mid \cdot \rangle$, for some m .

(5) \mathfrak{A} is (k -dimensional) FO-interpretable in $\text{copy}_m(\langle \omega, \leq \rangle)$, for some m .

Proof. (3) \Rightarrow (1) is trivial and (1) \Leftrightarrow (2) follows by Proposition 12.2.

(2) \Rightarrow (3) Let E be the relation of all pairs $\langle w, w' \rangle$ where

$$w = u_o v_o^{i_o} u_1 v_1^{i_1} \dots u_{n-1} v_{n-1}^{i_{n-1}} u_n \\ \text{and } w' = a_o^{m_o} b_o^{k_o} a_1^{m_1} b_1^{k_1} \dots a_{n-1}^{m_{n-1}} b_{n-1}^{k_{n-1}} a_n^{m_n},$$

for $i_o, \dots, i_{n-1} < \omega$ and $m_j := |u_j|$ and $k_j := |v_j|$. Note that E is regular since we can write it as

$$[u_o \otimes a_o^{m_o}] [v_o \otimes b_o^{k_o}]^* [u_1 \otimes a_1^{m_1}] [v_1 \otimes b_1^{k_1}]^* \dots \\ [u_{n-1} \otimes a_{n-1}^{m_{n-1}}] [v_{n-1} \otimes b_{n-1}^{k_{n-1}}]^* [u_n \otimes a_n^{m_n}].$$

We claim that the image of the presentation of \mathfrak{A} under E is again an automatic presentation of \mathfrak{A} . Let $R \subseteq A^n$ be a relation of \mathfrak{A} and let R' be its image under E . We have to show that R' is regular. Let Σ be the alphabet used by the given presentation of \mathfrak{A} and let Γ be the alphabet such that $E \subseteq \Sigma^* \times \Gamma^*$. By (a variant of) Theorem 2.5, there exists an FO-formula $\varphi(\bar{x})$ defining R in the tree $\langle \Sigma^*, \leq_{\text{pf}}, (\text{suc}_a)_{a \in \Sigma}, =_{\text{len}} \rangle$. Modifying φ slightly, we obtain an FO-formula $\varphi'(\bar{x})$ defining R in

$$\mathfrak{T} := \langle (\Sigma + \Gamma)^*, \leq_{\text{pf}}, (\text{suc}_a)_{a \in \Sigma + \Gamma}, =_{\text{len}} \rangle.$$

We can define the image R' inside \mathfrak{T} by the formula

$$\psi(\bar{x}) := \exists \bar{y} \left[\varphi'(\bar{y}) \wedge \bigwedge_{i < n} E y_i x_i \right].$$

This implies that R' is regular.

(3) \Rightarrow (4) We encode a word

$$w = a_o^{k_o} b_o^{l_o} a_1^{k_1} b_1^{l_1} \dots a_{n-1}^{k_{n-1}} b_{n-1}^{l_{n-1}} a_n^{k_n}$$

by a tuple $\langle x', x_0, \dots, x_{2n} \rangle \in \omega^{2n+1}$ where x' is a number encoding the tuple $\langle a_0, b_0, \dots, a_{n-1}, b_{n-1}, a_n \rangle$ (there are only finitely many possibilities), and

$$\begin{aligned} x_0 &= k_0, \\ x_1 &= k_0 + l_0, \\ &\vdots \\ x_{2n} &= k_0 + l_0 + \dots + k_{n-1} + l_{n-1} + k_n. \end{aligned}$$

Let R be an r -ary regular relation. To show that the encoding of R is definable in $\langle \omega, \leq, m \mid \cdot \rangle$, let $x'_0 \tilde{x}_0, \dots, x'_{r-1} \tilde{x}_{r-1}$ be encodings of words w_0, \dots, w_{r-1} . Then the word $w_0 \otimes \dots \otimes w_{r-1}$ (which is also of the above form, see Figure 2) has an encoding $z' \tilde{z}$ that can be obtained from $x'_0 \tilde{x}_0, \dots, x'_{r-1} \tilde{x}_{r-1}$. It follows that we can define R using Lemma 12.4.

a_0	b_0	b_0	a_1	b_1	a_2	a_2	\square	\square
c_0	c_0	d_0	d_0	c_1	c_1	d_1	d_1	c_2
x_0	y_0	x_1	x_2, y_1	x_3	y_2	x_4	y_3	y_4

Figure 2: Merging two encodings $x' \tilde{x}$ and $y' \tilde{y}$.

(4) \Rightarrow (3) We encode a tuple $\tilde{m} \in \omega^n$ by a word of the form

$$\sigma c_0^{k_0} \dots c_{n-1}^{k_{n-1}}$$

where $\sigma : [n] \rightarrow [n]$ is some permutation such that $m_{\sigma(0)} \leq \dots \leq m_{\sigma(n-1)}$ and

$$k_0 := m_{\sigma(0)} \quad \text{and} \quad k_{i+1} := m_{\sigma(i+1)} - m_{\sigma(i)}.$$

It follows by Lemma 12.3 that, for every tuple $\tilde{m} \in \omega^n$, the set of all words encoding a tuple r -equivalent to \tilde{m} is a regular language of the form required by (3). The claim follows since every FO-definable relation R is a finite union of r -equivalence classes.

(5) \Rightarrow (4) We can construct an FO-interpretation of $\text{copy}_m(\langle \omega, \leq \rangle)$ in $\langle \omega, \leq, m \mid \cdot \rangle$ by encoding the pair $\langle i, k \rangle$ as the number $mk + i$. This leads to the formulae

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= x \leq y \wedge x \equiv y \pmod{m}, \\ \varphi_I(x, y) &:= \exists z [m \mid z \wedge x, y \in [z, z + m)], \\ \varphi_{H_i}(x) &:= x \equiv i \pmod{m}. \end{aligned}$$

(4) \Rightarrow (5) We can construct an FO-interpretation of $\langle \omega, \leq, m \mid \cdot \rangle$ in $\text{copy}_m(\langle \omega, \leq \rangle)$ by encoding the number $mk + i$ by the pair $\langle i, k \rangle$. This leads to the formulae

$$\begin{aligned} \delta(x) &:= \text{true}, \\ \varphi_{\leq}(x, y) &:= \left[Ixy \wedge \bigvee_{i \leq j} (H_i x \wedge H_j y) \right] \vee \exists y' [Iyy' \wedge x < y'], \\ \varphi_{m \mid \cdot}(x) &:= H_0 x, \end{aligned}$$

where H_0, \dots, H_{m-1}, I are the relations added by the copy operation. \square

Exercise 12.1. Prove that the class of poly-growth automatic structures is closed under finite disjoint unions and finite direct products. \lrcorner

Presburger Structures

Our second class is slightly larger than that of the poly-growth automatic structures. The definition is as follows.

Definition 12.6. (a) A *Presburger structure* is a structure \mathfrak{A} for which there exists a (k -dimensional) FO-interpretation of \mathfrak{A} in $\langle \mathbb{N}, + \rangle$.

(b) We say that a subset $S \subseteq \mathbb{N}^n$ is *Presburger-definable* if it is FO-definable in $\langle \mathbb{N}, + \rangle$. \lrcorner

Proposition 12.7. Every poly-growth automatic structure is a Presburger structure.

Proof. The claim follows from Theorem 12.5 and the fact that there exists an FO-interpretation of $\langle \mathbb{N}, \leq, m \mid \cdot \rangle$ in $\langle \mathbb{N}, + \rangle$. \square

To better understand such structures, we need some results about which kinds of relations are definable in $\langle \mathbb{N}, + \rangle$.

Definition 12.8. (a) A set $S \subseteq \mathbb{N}^n$ is *semilinear* if it can be written as a finite union of sets of the form

$$u + \langle\langle v_0, \dots, v_{m-1} \rangle\rangle := \{ u + \sum_{i < m} \mu_i v_i \mid \mu_0, \dots, \mu_{m-1} \in \mathbb{N} \},$$

for $u, v_0, \dots, v_{m-1} \in \mathbb{N}^n$

(b) We call S *simple* if it is of the form $S = u + \langle\langle V \rangle\rangle$ where the set $V \subseteq \mathbb{N}^n$ is linearly independent in \mathbb{Q}^n . \lrcorner

Since it would lead us too far away, we will omit some of the more substantial proofs. The following classical result can be found, e.g., in Corollary 7.1a of [188].

Theorem 12.9 (Farkas, Minkowski, Weyl). *A set $S \subseteq \mathbb{R}^n$ is the set of solutions to a set of inequalities $Ax \geq c$ with $A \in \mathbb{R}^{m \times n}$ and $c \in \mathbb{R}^m$ if, and only if,*

$$S = \text{cone}(P) + \text{conv}(Q), \quad \text{for finite } P, Q \subseteq \mathbb{R}^n,$$

where

$$\begin{aligned} \text{cone}(P) &:= \{ \lambda_0 p_0 + \dots + \lambda_{k-1} p_{k-1} \mid p_i \in P, \lambda_i \geq 0 \}, \\ \text{conv}(Q) &:= \{ \lambda_0 q_0 + \dots + \lambda_{k-1} q_{k-1} \mid q_i \in Q, \lambda_i \geq 0, \\ &\quad \lambda_0 + \dots + \lambda_{k-1} = 1 \}. \end{aligned}$$

Another classical result is the following theorem by Cobham and Semenov (see [75] for an introduction; the relation $|_k$ has been introduced in Definition 2.4 above).

Definition 12.10. Two natural numbers $k, l \in \mathbb{N}$ are *multiplicatively independent* if the only integer solution to the equation $k^n = l^m$ is $n = 0 = m$. \lrcorner

Theorem 12.11 (Cobham, Semenov). *Let $k, l \geq 2$ be multiplicatively independent integers. A set $S \subseteq \mathbb{N}^n$ is FO-definable in both $\langle \mathbb{N}, +, |_k \rangle$ and $\langle \mathbb{N}, +, |_l \rangle$ if, and only if, it is FO-definable in $\langle \mathbb{N}, + \rangle$.*

Proposition 12.12 (Ito [110], Eilenberg, Schützenberger [76]). *Every semilinear set $S \subseteq \mathbb{N}^n$ can be written as a finite disjoint union of simple semilinear sets.*

The main characterisation of Presburger-definable sets is the following one.

Theorem 12.13. *Let $S \subseteq \mathbb{N}^n$. The following statements are equivalent.*

- (1) S is semilinear.
- (2) S is FO-definable in $\langle \mathbb{N}, + \rangle$.
- (3) S is FOC(U)-definable in $\langle \mathbb{N}, + \rangle$.
- (4) There are multiplicatively independent numbers $k, l \geq 2$ such that S is FO-definable in both $\langle \mathbb{N}, +, |_k \rangle$ and $\langle \mathbb{N}, +, |_l \rangle$.
- (5) There is some $m < \omega$ such that S is quantifier-free definable in the structure $\langle \mathbb{N}, +, \leq, m \mid \cdot, 0, 1 \rangle$.

Proof. (1) \Rightarrow (3) We can define the set $\tilde{u} + \langle\langle \tilde{v}_0, \dots, \tilde{v}_{k-1} \rangle\rangle$ by the formula

$$\varphi(\tilde{x}) := \exists \tilde{z} \left[\bigwedge_{i < n} x_i = u_i + v_{0,i} z_0 + \dots + v_{k-1,i} z_{k-1} \right],$$

where vz is an abbreviation for the term $z + \dots + z$ (v times). Since every semilinear set is a finite union of such sets, the claim follows by the closure of FO under disjunction.

(3) \Rightarrow (4) Fix an FOC(U)-definable set $S \subseteq \mathbb{N}^n$. The structures $\mathfrak{N}_k := \langle \mathbb{N}, +, |_k \rangle$ and $\mathfrak{N}_l := \langle \mathbb{N}, +, |_l \rangle$ have automatic presentations based on, respectively, the k -ary encoding and the l -ary encoding. By Theorem 7.12, these presentations can be expanded to ones of, respectively, $\langle \mathfrak{N}_k, S \rangle$ and $\langle \mathfrak{N}_l, S \rangle$. Finally, it follows by (the proof of) Theorem 2.5 that S is FO-definable in both \mathfrak{N}_k and \mathfrak{N}_l .

(4) \Rightarrow (2) follows by Theorem 12.11.

(2) \Rightarrow (5) Set

$$\mathfrak{N}_\omega := \langle \mathbb{N}, +, -, \leq, (m \mid \cdot)_{1 < m < \omega} \rangle.$$

We use the abbreviation

$$x \equiv k \pmod{m} := m \mid x + (m - k).$$

We start by proving a normal form for quantifier-free formulae $\varphi(\bar{x})$ in the structure \mathfrak{N}_ω . Transforming φ into disjunctive normal form and eliminating all negations using the equivalences

$$\neg(m \mid t) \equiv \bigvee_{i=1}^{m-1} m \mid (t + i),$$

$$\neg(s = t) \equiv s < t \vee t < s,$$

$$\neg(s < t) \equiv t < s + 1,$$

we may assume that φ is in disjunctive normal form without any negations. Furthermore, using the equivalences

$$m \mid t \equiv \bigvee \{ \bigwedge_i x_i \equiv k_i \pmod{m} \mid m \mid t(\bar{k}) \},$$

$$\begin{aligned} [x \equiv k \pmod{m} \wedge x \equiv k' \pmod{m'}] \\ \equiv \bigvee_{\substack{i < m \\ i' < m'}} [x \equiv i'm + k \pmod{mm'} \wedge x \equiv im' + k' \pmod{mm'}] \end{aligned}$$

we may assume that all atomic subformulae involving the divisibility predicate are of the form $m \mid x_i + k$, for some constants m, k , and that the constant m is the same for all such formulae. Adding disjunctions

$$\bigvee_{k < m} x_i \equiv k \pmod{m},$$

for each variable x_i , and transforming the resulting formula into disjunctive normal form, it follows that $\varphi(\bar{x})$ is equivalent to a disjunction of formulae of the form

$$\bigwedge_i x_i \equiv k_i \pmod{m} \wedge \bigwedge_j s_j = t_j \wedge \bigwedge_j u_j < v_j,$$

for terms s_j, t_j, u_j, v_j and constants $k_i, m \in \mathbb{N}$. We call this form the *normal form* of φ .

Let us next show that every formula of the form

$$\begin{aligned} \varphi(\bar{x}) = \exists y \Big[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge y \equiv l \pmod{m} \\ \wedge \bigwedge_j s_j = t_j \wedge \bigwedge_j u_j < v_j \Big] \end{aligned}$$

is equivalent to some quantifier-free formula. By multiplying each equation $s_j = t_j$ and each inequality $u_j < v_j$ by a suitable constant, we may assume that the coefficient of y is the same in all terms s_j, t_j, u_j, v_j which y appears in. Let μ be this common coefficient. It follows that

$$\begin{aligned} \varphi(\bar{x}) \equiv \exists y \Big[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge y \equiv l \pmod{m} \\ \wedge \bigwedge_j \mu y = t_j \wedge \bigwedge_j u_j < \mu y \wedge \bigwedge_j \mu y < v_j \Big], \end{aligned}$$

where t_j, u_j, v_j are terms in the variables \bar{x} only. Replacing μy by a new variable y' , we obtain

$$\begin{aligned} \varphi(\bar{x}) \equiv \exists y' \Big[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge y' \equiv \mu l \pmod{m} \\ \wedge \bigwedge_j y' = t_j \wedge \bigwedge_j u_j < y' \wedge \bigwedge_j y' < v_j \Big]. \end{aligned}$$

If φ contains an equation of the form $y' = s$, the value of y' is uniquely determined by \bar{x} and we can replace every occurrence of y' in φ by s . It follows that

$$\begin{aligned} \varphi(\bar{x}) \equiv \bigwedge_i x_i \equiv k_i \pmod{m} \wedge s \equiv \mu l \pmod{m} \\ \wedge \bigwedge_j s = t_j \wedge \bigwedge_j u_j < s \wedge \bigwedge_j s < v_j. \end{aligned}$$

Hence, we may assume that

$$\varphi(\bar{x}) \equiv \exists y' \left[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge y' \equiv \mu l \pmod{m} \right. \\ \left. \wedge \bigwedge_j u_j < y' \wedge \bigwedge_j y' < v_j \right].$$

Adding a disjunction checking which of the u_j is the largest one and which of the v_j the smallest one, we can transform φ into

$$\bigvee_{j_0, j_1} \exists y' \left[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge y' \equiv \mu l \pmod{m} \right. \\ \left. \wedge \bigwedge_j u_j \leq u_{j_0} \wedge \bigwedge_j v_{j_1} < v_j \wedge u_{j_0} < y' < v_{j_1} \right].$$

Note that the values $k_i = x_i \bmod m$ determine $p := u_{j_0} \bmod m$. Note that there is an element y' with $y' \equiv \mu l \pmod{m}$ between u_{j_0} and v_{j_1} if, and only if,

$$u_{j_0} + q < v_{j_1} \quad \text{where} \quad 1 \leq q \leq m \quad \text{and} \quad p + q \equiv \mu l \pmod{m}.$$

Hence, we can rewrite the formula into

$$\bigvee_{j_0, j_1} \left[\bigwedge_i x_i \equiv k_i \pmod{m} \wedge \bigwedge_j u_j \leq u_{j_0} \wedge \bigwedge_j v_{j_1} < v_j \right. \\ \left. \wedge u_{j_0} + q < v_{j_1} \right].$$

Having proved the special case above, we can now translate every FO-formula $\varphi(\bar{x})$ into an equivalent quantifier-free one (still in the structure \mathfrak{N}_ω). We proceed by induction on the quantifier-rank r of $\varphi(\bar{x})$. By inductive hypothesis, we can replace every subformula of φ of quantifier-rank less than r by a quantifier-free formula. Hence, we may assume that φ has quantifier-rank 1. As φ is a boolean combination of formulae of the form $\exists y \psi(\bar{x}, y)$ with quantifier-free ψ , we may further assume that $\varphi = \exists y \psi(\bar{x}, y)$ and that ψ is in normal form. Finally, using the fact that disjunctions commute

with existential quantifiers, it follows that φ is equivalent to a disjunction of formulae of the form considered in the previous step.

To conclude the proof, it remains to prove that every relation that is quantifier-free definable in \mathfrak{N}_ω is also quantifier-free definable in the structure $\langle \mathbb{N}, +, \leq, m \mid \cdot, 0, 1 \rangle$, for some m . We have seen that, in the structure \mathfrak{N}_ω , every quantifier-free formula $\varphi(\bar{x})$ can be written as a disjunction of formulae of the form

$$\bigwedge_i x_i \equiv k_i \pmod{m} \wedge \bigwedge_j s_j - s'_j = t_j - t'_j \wedge \bigwedge_j u_j - u'_j < v_j - v'_j,$$

for constants $k_i, m \in \mathbb{N}$ and terms $s_j, s'_j, t_j, t'_j, u_j, u'_j, v_j, v'_j$ without subtraction. Furthermore, we can assume that the constant m is the same for all terms of the disjunction. Such a formula can be written equivalently as

$$\bigwedge_i x_i \equiv k_i \pmod{m} \wedge \bigwedge_j s_j + t'_j = t_j + s'_j \wedge \bigwedge_j u_j + v'_j < v_j + u'_j,$$

which can be evaluated in $\langle \mathbb{N}, +, \leq, m \mid \cdot, 0, 1 \rangle$.

(5) \Rightarrow (1) We start by proving that every atomic formula defines a semi-linear set. There are two types of such formulae. First, consider an equation

$$a_0 x_0 + \cdots + a_{n-1} x_{n-1} = b_0 x_0 + \cdots + b_{n-1} x_{n-1} + c.$$

Let C be the set of minimal solutions of this equation (with respect to the component-wise ordering) and let P be the set of minimal non-zero solutions of the homogeneous equation

$$a_0 x_0 + \cdots + a_{n-1} x_{n-1} = b_0 x_0 + \cdots + b_{n-1} x_{n-1}.$$

Then C and P are finite and we can write the set of all solutions of the above equation as

$$\bigcup_{c \in C} c + \langle P \rangle.$$

The other type of atomic formula involves the modulo predicate.

$$m \mid (a_0 x_0 + \cdots + a_{n-1} x_{n-1} + b).$$

This statement is equivalent to

$$\exists z[a_0x_0 + \dots + a_{n-1}x_{n-1} + c = mz].$$

As above, we can show that the solutions S of the equation

$$a_0x_0 + \dots + a_{n-1}x_{n-1} + c = mz.$$

form a semilinear set, say,

$$S = \bigcup_{i < k} u_i + \langle\langle V_i \rangle\rangle.$$

By removing the z -component from u_i and the vectors in the set V_i , we obtain a similar expression for the solutions of the original statement.

To conclude the proof, we have to show that the class of semilinear sets is closed under boolean operations. Closure under finite unions is trivial. For closure under finite intersections, it is sufficient to prove that each set of the form $S = u + \langle\langle V \rangle\rangle \cap u' + \langle\langle V' \rangle\rangle$ is semilinear. Let A be the matrix with columns V and let A' be the matrix with columns V' . Note that

$$\begin{aligned} u + Ay = u' + A'y & \quad \text{iff} \quad Ay - A'y' = u' - u \\ & \quad \text{iff} \quad [A | -A'] \begin{bmatrix} y \\ y' \end{bmatrix} = u' - u \\ & \quad \text{iff} \quad \begin{bmatrix} y \\ y' \end{bmatrix} \in C + \langle\langle W \rangle\rangle, \end{aligned}$$

for suitable finite sets C, W . It follows that

$$\begin{aligned} x \in S & \quad \text{iff} \quad \text{there are } y, y' \text{ such that } x = u + Ay \text{ and } x = u' + A'y \\ & \quad \text{iff} \quad \text{there is some } y' \in C + \langle\langle W \rangle\rangle \text{ such that } x = u + Ay \\ & \quad \text{iff} \quad x \in u + A[C + \langle\langle W \rangle\rangle] = [u + AC] + \langle\langle AW \rangle\rangle, \end{aligned}$$

which is semilinear.

It remains to prove closure under complements. By Proposition 12.12 and the fact that we have already proved closure under intersections, it is sufficient

to show that the complement of a simple semilinear set $S = u + \langle\langle V \rangle\rangle \subseteq \mathbb{N}^n$ is semilinear. Let $\hat{S} \subseteq \mathbb{R}^n$ be the convex hull of S . We will prove that $(\hat{S} \setminus S) \cap \mathbb{N}^n$ and $\mathbb{N}^n \setminus \hat{S}$ are both semilinear.

For the latter, note that we can use Theorem 12.9 to define \hat{S} by a system of linear inequalities $Ax \geq c$ with $A \in \mathbb{R}^{m \times n}$ and $c \in \mathbb{R}^m$. Let

$$a_0x \geq c_0, \quad \dots, \quad a_{m-1}x \geq c_{m-1}$$

be the rows of this system. Then the complement $\mathbb{R}^n \setminus \hat{S}$ is given by the union of the solutions to the inequalities

$$a_0x < c_0, \quad \dots, \quad a_{m-1}x < c_{m-1}.$$

The solution to each such inequality (in \mathbb{N}^n) is semilinear. Hence, so is $(\mathbb{R}^n \setminus \hat{S}) \cap \mathbb{N}^n$.

It remains to show that $(\hat{S} \setminus S) \cap \mathbb{N}^n$ is also semilinear. For every $x \in S$, there is a unique vector $y \in \mathbb{N}^m$ with $x = u + By$, where B is the matrix with columns V . Hence,

$$\hat{S} \setminus S = \{u + By \mid y \in \mathbb{R}^m \setminus \mathbb{N}^m\}.$$

Setting

$$C := u + \bigcup \{By \cap \mathbb{N}^n \setminus \{0\} \mid y \in [0, 1)^n\}$$

(which is finite), it follows that

$$(\hat{S} \setminus S) \cap \mathbb{N}^n = C + \langle\langle V \rangle\rangle. \quad \square$$

The following relationship between regular languages and semilinear sets holds more generally for all context-free languages, but we only need the regular case where the proof is much simpler.

Theorem 12.14 (Parikh). *Let $K \subseteq \Sigma^*$ be regular and let $\Phi : \Sigma^* \rightarrow \mathbb{N}^\Sigma$ be the function mapping each word w to the tuple $\Phi(w) := (k_c)_{c \in \Sigma}$ where k_c is the number of letters c in w . Then $\Phi[K] \subseteq \mathbb{N}^\Sigma$ is semilinear.*

Proof. Let α be a regular expression denoting K . We prove the claim by induction on α . To simplify notation, we set $\Phi[\alpha] := \Phi(K)$, where K is the language denoted by α .

If $\alpha = c$ is a single letter, $\Phi[\alpha]$ contains a single tuple and is semilinear.

If $\alpha = \beta \cdot \gamma$, we have $\Phi[\alpha] = \Phi[\beta] + \Phi[\gamma]$. Since semilinear sets are closed under finite sums, the claim follows.

If $\alpha = \beta + \gamma$, we have $\Phi[\alpha] = \Phi[\beta] \cup \Phi[\gamma]$. Since semilinear sets are closed under finite unions, the claim follows.

Finally, if $\alpha = \beta^*$, we have

$$\Phi[\alpha] = \sum_{n < \omega} [\Phi[\beta] + \dots + \Phi[\beta]]. \quad (n \text{ times})$$

To show that this set is semilinear, suppose that

$$\Phi[\beta] = \{u_o + \langle V_o \rangle, \dots, u_{k-1} + \langle V_{k-1} \rangle\}.$$

It follows that the above sum is equal to

$$\begin{aligned} & \bigcup_{I \subseteq [k]} \left\{ \sum_{i \in I} \lambda_i u_i + \sum_{i \in I} \sum_{v \in V_i} \mu_{i,v} v \mid \lambda_i \in \mathbb{N} \setminus \{0\}, \mu_{i,v} \in \mathbb{N} \right\} \\ &= \bigcup_{I \subseteq [k]} \sum_{i \in I} u_i + \langle \{u_i \mid i \in I\} \cup \bigcup_{i \in I} V_i \rangle \end{aligned}$$

(with the convention that the above sums equals 0 if $I = \emptyset$). \square

Our next aim is to derive a bound on the out-degree of a semilinear relation.

Definition 12.15. (a) The *vector partition function* $\psi_A : \mathbb{N}^n \rightarrow \mathbb{N}$ associated with a matrix $A \in \mathbb{N}^{n \times m}$ maps a tuple $\tilde{x} \in \mathbb{N}^n$ to the number of tuples $\tilde{y} \in \mathbb{N}^m$ such that $\tilde{x} = A\tilde{y}$.

(b) A *generalised vector partition function* is a function of the form

$$g(\tilde{x}) := \sum_{i < s} \psi_{A_i}(\tilde{x} + \tilde{c}_i), \quad \text{for } A_i \in \mathbb{N}^{n \times m_i} \text{ and } \tilde{c}_i \in \mathbb{N}^n.$$

(c) A function $g : \mathbb{N}^n \rightarrow \mathbb{N}$ is a *piecewise polynomial* if there exists a partition \mathcal{S} of \mathbb{N}^n into semilinear sets such that, for every $S \in \mathcal{S}$, the restriction $g \upharpoonright S$ is a polynomial in $\mathbb{Q}[x_0, \dots, x_{n-1}]$. \dashv

Proposition 12.16 (Sturm [199]). *Every vector partition function is piecewise polynomial.*

Corollary 12.17. *Every generalised vector partition function is piecewise polynomial.*

Proposition 12.18. *Let $R \subseteq \mathbb{N}^k \times \mathbb{N}^l$ be a semilinear relation of finite out-degree. The function $d : \mathbb{N}^k \rightarrow \mathbb{N}$ mapping each tuple $\tilde{u} \in \mathbb{N}^k$ to its R -out-degree is a generalised vector partition function.*

Proof. With each relation $S \subseteq \mathbb{N}^{k+l}$ we associate the formal power-series

$$f_S(\tilde{x}, \tilde{y}) := \sum_{\langle \tilde{c}, \tilde{d} \rangle \in S} \tilde{x}^{\tilde{c}} \tilde{y}^{\tilde{d}}.$$

We can use Proposition 12.12 to write $R = S_o \cup \dots \cup S_{n-1}$ as a finite disjoint union of simple semilinear sets

$$S_i = \tilde{u}_i \tilde{u}'_i + \langle \tilde{v}_{i,o} \tilde{v}'_{i,o}, \dots, \tilde{v}_{i,s_i-1} \tilde{v}'_{i,s_i-1} \rangle.$$

A direct calculation shows that

$$f_{S_i}(\tilde{x}, \tilde{y}) = \frac{\tilde{x}^{\tilde{u}_i} \tilde{y}^{\tilde{u}'_i}}{(1 - \tilde{x}^{\tilde{v}_{i,o}} \tilde{y}^{\tilde{v}'_{i,o}}) \dots (1 - \tilde{x}^{\tilde{v}_{i,s_i-1}} \tilde{y}^{\tilde{v}'_{i,s_i-1}})}.$$

Hence, we obtain

$$f_R(\tilde{x}, \tilde{y}) = \sum_{i < n} \frac{\tilde{x}^{\tilde{u}_i} \tilde{y}^{\tilde{u}'_i}}{(1 - \tilde{x}^{\tilde{v}_{i,o}} \tilde{y}^{\tilde{v}'_{i,o}}) \dots (1 - \tilde{x}^{\tilde{v}_{i,s_i-1}} \tilde{y}^{\tilde{v}'_{i,s_i-1}})},$$

which implies that

$$\begin{aligned} \sum_{\tilde{c} \in \mathbb{N}^k} d(\tilde{c}) \tilde{x}^{\tilde{c}} &= \sum_{\tilde{c} \in \mathbb{N}^k} |\{ \tilde{d} \mid \langle \tilde{c}, \tilde{d} \rangle \in R \}| \cdot \tilde{x}^{\tilde{c}} \\ &= \sum_{\langle \tilde{c}, \tilde{d} \rangle \in R} \tilde{x}^{\tilde{c}} \end{aligned}$$

$$\begin{aligned}
&= \sum_{\langle \bar{c}, \bar{d} \rangle \in R} \bar{x}^{\bar{c}} \bar{1}^{\bar{d}} \\
&= f_R(\bar{x}, 1 \dots 1) \\
&= \sum_{i < n} \frac{\bar{x}^{\bar{u}_i}}{(1 - \bar{x}^{\bar{v}_{i,0}}) \dots (1 - \bar{x}^{\bar{v}_{i,s_i-1}})} .
\end{aligned}$$

Since generalised vector partition functions are closed under addition, it is therefore sufficient to prove that, given a power-series of the form,

$$\sum_{\bar{b}} g(\bar{b}) \bar{x}^{\bar{b}} = \frac{\bar{x}^{\bar{c}}}{(1 - \bar{x}^{\bar{a}_0}) \dots (1 - \bar{x}^{\bar{a}_{m-1}})} ,$$

the coefficient function g is a generalised vector partition function. In this case, we obtain

$$\begin{aligned}
\sum_{\bar{b}} g(\bar{b} + \bar{c}) \bar{x}^{\bar{b}} &= \frac{1}{(1 - \bar{x}^{\bar{a}_0}) \dots (1 - \bar{x}^{\bar{a}_{m-1}})} \\
&= \left[\sum_{\mu_0 < \omega} \bar{x}^{\mu_0 \bar{a}_0} \right] \dots \left[\sum_{\mu_{m-1} < \omega} \bar{x}^{\mu_{m-1} \bar{a}_{m-1}} \right] \\
&= \sum_{\mu_0, \dots, \mu_{m-1} < \omega} \bar{x}^{\mu_0 \bar{a}_0 + \dots + \mu_{m-1} \bar{a}_{m-1}} ,
\end{aligned}$$

which implies that $g(\bar{z} + \bar{c})$ is equal to the number of tuples $\bar{\mu} \in \mathbb{N}^m$ satisfying

$$\bar{z} = \mu_0 \bar{a}_0 + \dots + \mu_{m-1} \bar{a}_{m-1} .$$

Hence, $g(\bar{z}) = \psi_A(\bar{z} + \bar{c})$, for some A and c . \square

Example. Let $R \subseteq \mathbb{N}^2 \times \mathbb{N}$ be the set of all triples $\langle a, b, c \rangle$ such that c is an even number with $a \leq c \leq b$. This relation is definable in $\langle \mathbb{N}, + \rangle$ and, hence, semilinear. Its out-degree is

$$d(a, b) := \begin{cases} 0 & \text{if } a > b , \\ \frac{1}{2}(b - a) + 1 & \text{for } a \leq b \text{ and } a, b \text{ even,} \\ \frac{1}{2}(b - a) & \text{for } a \leq b \text{ and } a, b \text{ odd,} \\ \frac{1}{2}(b - a + 1) & \text{otherwise.} \end{cases}$$

Growth Arguments

To prove that certain structures are not poly-growth automatic, we can use the following variant of Proposition 8.16.

Lemma 12.19. *Let \mathfrak{A} be a poly-growth automatic structure, $U \subseteq A$ finite, and φ a formula of finite out-degree. Then there exist constants $d, k > 0$ such that*

$$|N_\varphi(U, n)| \leq n^d + k , \quad \text{for all } n < \omega .$$

Proof. Let $l := \max \{ \|c\| \mid c \in U \}$. By Corollary 8.3, there exists a constant c such that

$$\|a\| \leq l + cn , \quad \text{for all } a \in N_\varphi(U, n) .$$

By assumption, there exists a polynomial $p(x)$ such that the universe of \mathfrak{A} contains at most $p(n)$ words of length at most n . Consequently,

$$|N_\varphi(U, n)| \leq p(l + mn) . \quad \square$$

Our first applications of this lemma concern linear orders.

Theorem 12.20. *An ordinal $\langle \alpha, \leq \rangle$ is poly-growth automatic if, and only if, $\alpha < \omega^\omega$.*

Proof. (\Rightarrow) We have shown in Theorem 11.11 that all automatic ordinals are smaller than ω^ω .

(\Leftarrow) As in Proposition 11.1, one can show that the class of poly-growth automatic ordinals is closed under addition and multiplication. Furthermore, $\langle \omega, \leq \rangle$ is poly-growth automatic. \square

Proposition 12.21. *Let \mathfrak{A} be a coloured linear order. If \mathfrak{A} is regular and scattered, it is poly-growth automatic.*

Proof. If \mathfrak{A} is regular and scattered, it is denoted by a finite Leonard-L\"auchli term without the shuffle operation. Hence, it is sufficient to prove that each Leonard-L\"auchli operation preserves poly-growth automaticity. Clearly,

every finite linear order is poly-growth automatic and so is the ordered sum of two poly-growth automatic linear orders. Finally, note that ω and ω^{op} are poly-growth automatic and that poly-growth automatic linear orders are closed under binary products. \square

Example. The converse is not true. Let $\langle \omega, \leq, P \rangle$ be the order with

$$P := \{ n(n+1)/2 \mid n < \omega \}.$$

This order has an automatic presentation $\langle a^*b^*, \leq_{\text{lex}}, a^* \rangle$ with polynomial growth, but it is not regular. \dashv

Instead of the converse, we can only prove the following weaker statement.

Proposition 12.22. *For every poly-growth automatic linear order \mathfrak{A} , we have $\text{FC}(\mathfrak{A}) < \omega$.*

Proof. For a contradiction, suppose that $\text{FC}(\mathfrak{A}) = \infty$. We have shown in Proposition 11.3 that $n := \text{gFC}(\mathfrak{A})$ is finite. It follows that $\text{cn}^n(\mathfrak{A}) \cong \mathfrak{Q} := \langle \mathbb{Q}, \leq \rangle$. Since we can FOC-interpret $\text{cn}^n(\mathfrak{A})$ in \mathfrak{A} , the order \mathfrak{Q} is poly-growth automatic. Let $\varphi(x, y, z)$ be the formula stating that z is the \leq_{lex} -least element with $x < z < y$. Then φ has finite out-degree and

$$|N_\varphi(\{a, b\}, n)| = 2 + 2^{n-1}, \quad \text{for } a < b$$

(in the structure \mathfrak{Q}). A contradiction to Lemma 12.19. \square

Open Question. *Find a characterisation of poly-growth automatic linear orders.*

Next, let us take a look at semigroups. In Corollary 9.3 above, we have already derived a characterisation of all finitely generated Presburger groups.

Proposition 12.23. *A finitely generated group \mathfrak{G} is a Presburger structure if, and only if, \mathfrak{G} is automatic.*

Next, let us take a look at poly-growth automatic groups.

Lemma 12.24. *Let \mathfrak{S} be a semigroup such that there exists an embedding of $\langle \mathbb{N} \setminus \{0\}, + \rangle$ into \mathfrak{S} . Then \mathfrak{S} is not poly-growth automatic.*

Proof. Suppose that \mathfrak{S} is an automatic structure into which $\langle \mathbb{N} \setminus \{0\}, + \rangle$ can be embedded, and let c be the image of 1 under this embedding. By Lemma 8.8, there exists a constant k such that

$$\|nc\| \leq \|c\| + k \log_2 n, \quad \text{for all } n.$$

It follows that

$$n \leq 2^{(m-\|c\|)/k} \quad \text{implies} \quad \|nc\| \leq m.$$

Hence, the set $\{a \in S \mid \|a\| \leq m\}$ contains at least $2^{(m-\|c\|)/k}$ elements and \mathfrak{S} is not of polynomial growth. \square

It turns out that the only poly-growth automatic groups are the finite ones.

Theorem 12.25. *A group is poly-growth automatic if, and only if, it is finite.*

Proof. Let \mathfrak{G} be a poly-growth automatic group. By Corollary 8.3 there exists a constant k such that

- ♦ for every $a \in G$, there is some $b \in G$ with $\|a\| < \|b\| \leq \|a\| + k$,
- ♦ $\|ab\| \leq \max\{\|a\|, \|b\|\} + k$,
- ♦ $\|a^{-1}\| \leq \|a\| + k$.

Setting $m := \|e\|$ it follows that, for each $n < \omega$, there exists some element $a_n \in A$ of length

$$m + 4kn \leq \|a_n\| < m + 4kn + k.$$

Set $D_0 := \{e\}$,

$$\begin{aligned} C_n &:= \{a_0^{e_0} \cdots a_{n-1}^{e_{n-1}} \mid e_0, \dots, e_{n-1} \in \{0, 1\}\}, \\ D_n &:= \{a^{-1}b \mid a, b \in C_n\}. \end{aligned}$$

We claim that

$$(1) \quad \|c\| < m + 4k(n-1) + 2k, \quad \text{for all } n > 0 \text{ and } c \in C_n,$$

$$(II) \|c\| < m + 4kn, \quad \text{for all } c \in D_n,$$

$$(III) |C_n| = 2^n.$$

It follows that G contains at least 2^n elements of length at most $4kn$. A contradiction to the fact that \mathcal{G} has polynomial growth. Hence, it remains to prove the above claims.

(I) We proceed by induction on n . For $n = 1$, we have $\|e\| = m$ and $\|a_0\| < m + k \leq m + 2k$. For the inductive step, let $c \in C_n$. If $c \in C_{n-1}$, the claim follows by inductive hypothesis. Otherwise, we can write $c = da_{n-1}$ with $d \in D_{n-1}$. Then $\|d\|, \|a_{n-1}\| < m + 4k(n-1) + k$ implies, by choice of k , that

$$\|c\| < m + 4k(n-1) + 2k.$$

(II) Let $a, b \in C_n$. By (I), we have $\|a\|, \|b\| < m + 4k(n-1) + 2k$. By choice of k , this implies that

$$\|a^{-1}b\| < m + 4k(n-1) + 2k + 2k = m + 4kn.$$

(III) Suppose that

$$a_0^{d_0} \cdots a_n^{d_n} = a_0^{e_0} \cdots a_n^{e_n}, \quad \text{for } d_0, \dots, d_n, e_0, \dots, e_n \in \{0, 1\}.$$

We prove that $d_i = e_i$ by induction on n . Set

$$b := a_0^{d_0} \cdots a_{n-1}^{d_{n-1}} \quad \text{and} \quad c := a_0^{e_0} \cdots a_{n-1}^{e_{n-1}}.$$

If $d_n = e_n$, we obtain $b = c$ and the claim follows by inductive hypothesis. Otherwise, we may assume without loss of generality that $d_n = 0$ and $e_n = 1$. Hence,

$$b = ca_n \quad \text{implies} \quad a_n = c^{-1}b \in D_n.$$

By (III), it follows that $\|a_n\| < m + 4kn$. A contradiction to our choice of a_n . \square

For Cayley graphs, we obtain the following result.

Proposition 12.26. *Every finitely generated automatic group has a Cayley graph that is poly-growth automatic.*

Proof. This follows immediately by Lemma 9.15. \square

Exercise 12.2. The converse is false: the Heisenberg group has a poly-growth automatic Cayley graph but it is not automatic. \dashv

Equivalence Structures

Our aim is to prove characterisations both of Presburger equivalence relations and of poly-growth automatic equivalence relations.

Definition 12.27. (a) An *equivalence structure* is a structure of the form $\langle A, \sim \rangle$ where \sim is an equivalence relation on A .

(b) Given a function $g : \mathbb{N}^n \rightarrow \mathbb{N}$, we denote by $\mathfrak{E}(g)$ the equivalence structure with exactly $|g^{-1}(k)|$ classes of size k , for each $k \in \mathbb{N} \setminus \{0\}$.

(c) A polynomial $p \in \mathbb{Q}[x_0, \dots, x_{n-1}]$ is *positive* if the associated polynomial function $\mathbb{Q}^n \rightarrow \mathbb{Q}$ restricts to a function $\mathbb{N}^n \rightarrow \mathbb{N} \setminus \{0\}$.

(d) For $\bar{x}, \bar{a} \in \mathbb{N}^n$, we write

$$\bar{x}^{\bar{a}} := x_0^{a_0} \cdots x_{n-1}^{a_{n-1}}.$$

We start with a simple lemma that helps us to define interpretations of structures of the form $\mathfrak{E}(g)$.

Lemma 12.28. *Let \mathfrak{A} be a well-ordered structure and $g : \mathbb{N}^n \rightarrow \mathbb{N}$ a function. There exists a (k -dimensional) FO-interpretation of $\mathfrak{E}(g)$ in \mathfrak{A} if, and only if, there exist $k, m < \omega$, an injective function $\sigma : \mathbb{N}^n \rightarrow A^k$, and an FO-definable relation $R \subseteq A^{k \times m}$ such that*

$$d_R(\bar{a}) = \begin{cases} g(\bar{c}) & \text{if } \bar{a} = \sigma(\bar{c}), \\ 0 & \text{otherwise,} \end{cases}$$

where d_R is the function mapping a tuple \bar{a} to its R -out-degree.

Proof. (\Leftarrow) Given R , we set

$$\langle \bar{x}, \bar{y} \rangle \sim \langle \bar{x}', \bar{y}' \rangle \quad : \text{iff} \quad \bar{x} = \bar{x}'.$$

Then $\langle R, \sim \rangle \cong \mathfrak{E}(g)$.

(\Rightarrow) Let $\tau = \langle \delta(\bar{x}), \varphi(\bar{x}, \bar{y}) \rangle$ be a k -dimensional FO-interpretation of $\mathfrak{E}(g) = \langle E, \sim \rangle$ in \mathfrak{A} and let $v : \delta^{\mathfrak{A}} \rightarrow E$ be the corresponding isomorphism. By definition of $\mathfrak{E}(g)$, there exists a bijection $\rho : \mathbb{N}^n \rightarrow E/\sim$ such that

$$|\rho(\bar{k})| = g(\bar{k}), \quad \text{for all } \bar{k} \in \mathbb{N}^n.$$

Set $\approx := \varphi^{\mathfrak{A}}$ and let $P \subseteq \delta^{\mathfrak{A}}$ be the set containing the \leq_{lex} -minimal element of each \approx -class. Then $R := \approx \cap (P \times \mathbb{N}^m)$ is FO-definable and the R -out-degree of an element $\bar{a} \in P$ is

$$|\bar{a}|_{\approx} = |[v(\bar{a})]_{\sim}| = g(\rho^{-1}([v(\bar{a})])) = g((\rho^{-1} \circ q \circ v)(\bar{a})),$$

where $q : E \rightarrow E/\sim$ is the projection. Since the restriction of $\rho^{-1} \circ q \circ v$ to P is bijective, we obtain the desired function σ by setting

$$\sigma := (\rho^{-1} \circ q \circ v \upharpoonright P)^{-1} : \mathbb{N}^n \rightarrow A^k. \quad \square$$

Presburger equivalence structures can be characterised as follows.

Theorem 12.29. *An equivalence structure \mathfrak{A} is a Presburger structure if, and only if,*

$$\mathfrak{A} \cong \mathfrak{E}(g) + \mathfrak{C}$$

where g is a generalised vector partition function and \mathfrak{C} is a countable equivalence structure with only infinite classes.

Proof. (\Leftarrow) Note that the equivalence structures $\mathfrak{C}_1 := \langle \mathbb{N}, E_1 \rangle$ and $\mathfrak{C}_\infty := \langle \mathbb{N}^2, E_\infty \rangle$ with

$$E_1 := \mathbb{N} \times \mathbb{N} \quad \text{and} \quad E_\infty := \{ \langle \langle n, i \rangle, \langle n, j \rangle \rangle \mid n, i, j \in \mathbb{N} \}$$

are Presburger structures. (\mathfrak{C}_1 has a single infinite class and \mathfrak{C}_∞ has countably infinitely many.) Since Presburger structures are closed under finite disjoint unions, it therefore remains to show that $\mathfrak{E}(g)$ is Presburger, for every generalised vector partition function

$$g(\bar{x}) := \sum_{k < s} \psi_{A_k}(\bar{x} + \bar{c}_k), \quad \text{for } A_k \in \mathbb{N}^{n \times m_k} \text{ and } \bar{c}_k \in \mathbb{N}^n.$$

Given such a function g , set $m := \max_k m_k$. The relation

$$R := \left\{ \langle \bar{x}, \bar{y}, k \rangle \in \mathbb{N}^n \times \mathbb{N}^m \times \mathbb{N} \mid k < s, A_k \bar{y} = \bar{x} + \bar{c}_k, \right. \\ \left. y_i = 0 \text{ for } i \geq m_k \right\}$$

is Presburger definable and the out-degree of $\bar{x} \in \mathbb{N}^n$ is equal to

$$\sum_{k < s} \psi_{A_k}(\bar{x} + \bar{c}_k) = g(\bar{x}).$$

Consequently, we can use Lemma 12.28 to find an FO-interpretation of $\mathfrak{E}(g)$ in $\langle \mathbb{N}, +, \leq \rangle$.

(\Rightarrow) Suppose that there exists an k -dimensional FO-interpretation of \mathfrak{A} in $\langle \mathbb{N}, + \rangle$. Note that the substructure \mathfrak{A}_o of \mathfrak{A} consisting of all finite equivalence classes can be defined by the FOC-formula

$$\varphi(x) := \neg \exists^{\aleph_0} y [y \sim x].$$

By Theorem 12.13, it therefore follows that \mathfrak{A}_o is also a Presburger structure. Hence, it is sufficient to prove that $\mathfrak{A}_o \cong \mathfrak{E}(g)$, for some generalised vector partition function. Let $P \subseteq A \subseteq \mathbb{N}^n$ be the set containing the \leq_{lex} -minimal element of every \sim -class. Since P is definable, so is the relation $R := \sim \cap (P \times A)$. It therefore follows by Proposition 12.18 that the function $d : \mathbb{N}^n \rightarrow \mathbb{N}$ mapping a tuple \bar{k} to its R -out-degree is of the form

$$p(\bar{x}) := \sum_{i < s} \psi_{A_i}(\bar{x} + \bar{c}_i).$$

Since

$$d(\bar{k}) = \begin{cases} |\lfloor \bar{k} \rfloor| & \text{if } \bar{k} \in P, \\ 0 & \text{otherwise,} \end{cases}$$

we further have $\mathfrak{A} \cong \mathfrak{E}(d)$. \square

We can improve one direction of the preceding theorem as follows.

Lemma 12.30. *Let g be a generalised vector partition function. Then $\mathfrak{E}(g)$ can be written as a finite union of structures of the form $\mathfrak{E}(p)$, for positive polynomials p with integer coefficients.*

Proof. By Proposition 12.16, $g \in \mathbb{Q}[\bar{x}]$ is piecewise polynomial. Hence, there exists a finite partition \mathcal{S} of \mathbb{N}^n into semilinear sets and a family of polynomials $(q_S)_{S \in \mathcal{S}}$ such that

$$g \upharpoonright S = q_S, \quad \text{for all } S \in \mathcal{S}.$$

By Proposition 12.12, we may assume that every $S \in \mathcal{S}$ is simple. For $S \in \mathcal{S}$ with $S = u + \langle\langle V \rangle\rangle$, let $h_S : \mathbb{N}^V \rightarrow \mathbb{N}^n$ be the function defined by

$$h_S(\bar{\mu}) := u + \sum_{v \in V} \mu_v v.$$

Then $q_S \circ h_S$ is a polynomial in $\mathbb{Q}[\bar{x}]$. Furthermore, we have

$$\mathfrak{E}(g) \cong \sum_{S \in \mathcal{S}} \mathfrak{E}(q_S \circ h_S)$$

because simplicity of S implies that the functions h_S are injective.

To conclude the proof, it is therefore sufficient to show that every structure of the form $\mathfrak{E}(q)$ with $q \in \mathbb{Q}[\bar{x}]$ can be written as a finite disjoint union of structures $\mathfrak{E}(h)$ with positive $h \in \mathbb{Z}[\bar{x}]$. We can write $q = \frac{1}{\mu} q_0$ with $q_0 \in \mathbb{Z}[\bar{x}]$ and $0 < \mu < \omega$. For each tuple $\bar{c} \in [\mu]^n$, we obtain a polynomial

$$p_{\bar{c}}(\bar{x}) := q(\mu \bar{x} + \bar{c}) \in \mathbb{Z}[\bar{x}].$$

(Note that the constant term of $p_{\bar{c}}$ belongs to \mathbb{Z} since q induces a function $\mathbb{N}^n \rightarrow \mathbb{N}$.) Furthermore, we have

$$\mathfrak{E}(q) \cong \sum_{\bar{c} \in [\mu]^n} \mathfrak{E}(p_{\bar{c}}).$$

\square

Open Question. *Is $\mathfrak{E}(p)$ a Presburger structure, for every positive $p \in \mathbb{Z}[\bar{x}]$?*

For poly-growth automatic equivalence structures, we obtain only a partial result. Again, we start with a lemma.

Lemma 12.31. *For every $p \in \mathbb{Q}_{\geq 0}[x_0, \dots, x_{n-1}]$, the structure $\mathfrak{E}(p)$ is poly-growth automatic.*

Proof. Suppose that

$$p = \frac{1}{\mu} \sum_{j < m} \lambda_j \bar{x}^{\bar{a}_j}, \quad \text{for } \mu, \lambda_0, \dots, \lambda_{m-1} \in \mathbb{N}.$$

Let $k_i := \max_j a_{j,i}$ be the maximal exponent of x_i in p . Let R be the relation of all tuples

$$\langle \bar{x}, \bar{y}_0 \dots \bar{y}_{n-1}, z, w \rangle \in \mathbb{N}^n \times \mathbb{N}^{k_0} \times \dots \times \mathbb{N}^{k_{n-1}} \times \mathbb{N} \times \mathbb{N}$$

such that

$$\begin{aligned} z &< m, \\ w &< \lambda_z, \\ y_{i,j} &< x_i, \quad \text{for } i < n \text{ and } j < a_{j,i}, \\ y_{i,j} &= 0, \quad \text{for } i < n \text{ and } a_{j,i} \leq j < k_i. \end{aligned}$$

Then R is definable in $\langle \mathbb{N}, \leq \rangle$ and the R -out-degree of $\bar{x} \in \mathbb{N}^n$ is equal to

$$\sum_{j < m} \lambda_j \bar{x}^{\bar{a}_j}.$$

Consequently, we can use Lemma 12.28 to construct an FO-interpretation of $\mathfrak{E}(p_0)$ in $\langle \mathbb{N}, \leq \rangle$, where $p_0 := \mu p$.

Finally, note that $\mathfrak{E}(p)$ can be obtained from $\mathfrak{E}(p_\circ)$ by taking every μ -th element of each equivalence class. Hence, $\mathfrak{E}(p)$ is isomorphic to the substructure of $\mathfrak{E}(p_\circ)$ defined by the formula

$$\varphi(x) := \exists^{\circ, \mu} y [y \sim x \wedge y \leq_{\text{lex}} x].$$

We have obtained FOC-interpretations $\mathfrak{E}(p) \leq \mathfrak{E}(r) \leq \langle \mathbb{N}, \leq \rangle$. By Theorem 12.5, it follows that $\mathfrak{E}(p)$ is poly-growth automatic. \square

Exercise 12.3. Find a poly-growth automatic presentation of $\mathfrak{E}(x^2)$. \lrcorner

Theorem 12.32. *There exists a class \mathcal{P} of positive polynomials with integer coefficients such that an equivalence structure \mathfrak{A} is poly-growth automatic if, and only if, it can be written as a finite disjoint union of*

- structures of the form $\mathfrak{E}(p)$, for $p \in \mathcal{P}$, and
- countable equivalence structures where every class is infinite.

Furthermore, every polynomial with natural coefficients belongs to \mathcal{P} .

Proof of Theorem 12.32. (\Leftarrow) We have seen in Lemma 12.31 that every structure of the form $\mathfrak{E}(p)$ with $p \in \mathbb{N}[\bar{x}]$ is poly-growth automatic. Furthermore, the equivalence structures $\mathfrak{A}_1 := \langle \circ^*, E_1 \rangle$ and $\mathfrak{A}_\infty := \langle \circ^* \mathbf{1}^*, E_\infty \rangle$ with

$$E_1 := \circ^* \times \circ^* \quad \text{and} \quad E_\infty := \{ \langle \circ^n \mathbf{1}^k, \circ^n \mathbf{1}^l \rangle \mid n, k, l < \omega \}$$

are poly-growth automatic. (\mathfrak{A}_1 has a single infinite class and \mathfrak{A}_∞ has countably infinitely many.) The claim follows since the class of poly-growth automatic structures is closed under finite disjoint unions.

(\Rightarrow) According to Theorems 12.5 and 12.13, every poly-growth automatic equivalence structure is a Presburger structure. Consequently, it follows by Theorem 12.29 that $\mathfrak{A}_\circ \cong \mathfrak{E}(g) + \mathfrak{C}$, for some generalised vector partition function g and some equivalence structure \mathfrak{C} with only infinite classes. Furthermore, we can use Lemma 12.30 to write $\mathfrak{E}(g)$ as a finite union of structures of the form $\mathfrak{E}(q)$, for positive polynomials q with integer coefficients. \square

Open Question. Find an explicit description of the above class \mathcal{P} .

Unary Automatic Structures

There exists a subclass of the poly-growth automatic structure that is even simpler: those where the presentation is over a unary alphabet.

Definition 12.33. An *unary automatic structure* is a structure that has an automatic presentation over a unary alphabet. \lrcorner

For these structures the characterisation from Theorem 12.5 simplifies as follows.

Theorem 12.34. Let \mathfrak{A} be a structure. The following statements are equivalent.

- (1) \mathfrak{A} is unary automatic.
- (2) \mathfrak{A} is MSO-interpretable in $\langle \omega, \leq \rangle$.
- (3) \mathfrak{A} is (1-dimensional) FO-interpretable in $\langle \omega, \leq, m \mid \cdot \rangle$, for some $m < \omega$.
- (4) \mathfrak{A} is (1-dimensional) FO-interpretable in $\text{copy}_m(\langle \omega, \leq \rangle)$, for some m .

Exercise 12.4. Prove the theorem. \lrcorner

Unary automatic structures are too simple to be of much interest. Let us just mention the following relationship to tree-interpretable structures.

Proposition 12.35. Every unary automatic structure is tree-interpretable. There are tree-interpretable structures that are not unary automatic.

Proof. If \mathfrak{A} is unary automatic, we can use Theorem 12.34 to find an MSO-interpretation of \mathfrak{A} in $\langle \omega, \leq \rangle$. Since $\langle \omega, \leq \rangle$ can be MSO-interpreted in the binary tree $\mathfrak{T}_{\text{bin}}$, it follows that \mathfrak{A} is tree-interpretable.

For strictness, note that it follows by Corollary X.4.8 that $\mathfrak{T}_{\text{bin}}$ is tree-interpretable, but not unary automatic. \square

Remark. Note that there are poly-growth automatic structures that are not tree-interpretable, like the infinite grid $\langle \mathbb{Z} \times \mathbb{Z}, E \rangle$. \lrcorner

We can prove a converse statement for well-ordered tree-interpretable structures.

Proposition 12.36. *A structure \mathfrak{A} equipped with a well-ordering is unary automatic if, and only if, it is tree-interpretable and its order type is less than ω^2 .*

Proof. (\Rightarrow) The fact that \mathfrak{A} is tree-automatic follows immediately from Proposition 12.35. For a contradiction, suppose that the order type of \mathfrak{A} is at least ω^2 . Fix an increasing enumeration $(a_i)_{i < \alpha}$ of A and let $f : A \rightarrow A$ and $g : A \rightarrow A$ be the functions

$$f(a_i) := a_{i+1} \quad \text{and} \quad g(a_i) := a_{i+\omega},$$

respectively. Since f and g are definable, we can use Corollary 8.3 to find a constant $k < \omega$ such that

$$\|a_{i+1}\| \leq \|a_i\| + k \quad \text{and} \quad \|a_{\omega(i+1)}\| \leq \|a_{\omega i}\| + k.$$

It follows that

$$\|a_{\omega i+j}\| \leq \|a_o\| + (i+j)k.$$

Consequently, the number of elements a with $\|a\| \leq n$ is quadratic in n . But there are only linearly many words of length at most n . A contradiction.

(\Leftarrow) Let \mathfrak{A} be tree-interpretable and well-ordered of order-type α . Fix $n < \omega$ such that $\omega(n-1) \leq \alpha < \omega n$. For $l < \omega$ and $i < n$, set

$$C(l, i) := \{a \in A \mid \omega i \leq a < \omega(i+1) \text{ and } \|a\| = l\}.$$

It follows by Lemma XI.7.10 that the size of $C(i, l)$ is bounded by some constant k . We identify A with a subset of the binary tree $\mathfrak{T}_{\text{bin}}$. Applying the homomorphism $o \mapsto o^{nk}$ and $1 \mapsto 1^{nk}$, we can assume that A contains only elements whose length is a multiple of nk . Let $a(l, i, j)$ be the j -th element of $C(nkl, i)$. To construct a unary presentation of \mathfrak{A} , we encode $a(l, i, j)$ by the word $w(l, i, j) := 1^{nkl+kj}$. (Note that the mapping $a(l, i, j) \mapsto w(l, i, j)$ is injective.) We will show that this encoding is definable in the structure $\langle \mathfrak{T}_{\text{bin}}, =_{\text{len}} \rangle$. Then we can construct formulae φ_R , for each relation R of \mathfrak{A} , such that

$$\begin{aligned} \langle \mathfrak{T}_{\text{bin}}, =_{\text{len}} \rangle \models \varphi_R(w(l_o, i_o, j_o), \dots, w(l_{r-1}, i_{r-1}, j_{r-1})) \\ \text{iff} \quad \langle a(l_o, i_o, j_o), \dots, a(l_{r-1}, i_{r-1}, j_{r-1}) \rangle \in R. \end{aligned}$$

Since every FO-definable relation in $\langle \mathfrak{T}_{\text{bin}}, =_{\text{len}} \rangle$ is regular and φ_R defines a relation over words containing only the letter 1, it follows that we obtain a unary automatic presentation of \mathfrak{A} .

As all tree-interpretable structures are automatic, there exists a $(1$ -dimensional) FO-interpretation $\tau = \langle \delta(x), (\psi_R(\bar{x}))_R \rangle$ of \mathfrak{A} in $\langle \mathfrak{T}_{\text{bin}}, =_{\text{len}} \rangle$. To define the function mapping $w(l, i, j)$ to $a(l, i, j)$, we pick the $(ki+j)$ -th element of $C(nkl, o) \cup \dots \cup C(nkl, n-1)$. (This is possible since the length of these elements is equal to $|w(l, i, j)| - ki + j$.) We define a formula $\chi(x, y)$ stating that $x = w(l, i, j)$ and $y \in C(l, i)$, for some i, j, l , by

$$\vartheta_i(x) := \bigvee_{j < k} |x| \equiv ki + j \pmod{nk}$$

$$\eta_i(y) := (\exists^{=i+1} z \leq y)(\forall u < z) \exists v [u < v < z]$$

$$\text{and } \chi(x, y) := \delta(y) \wedge |x| - nk < |y| \leq |x| \wedge \bigwedge_{i < n} [\vartheta_i(x) \leftrightarrow \eta_i(y)].$$

For simplicity, consider the case of a unary relation R only. We can define the encoding of R by

$$\begin{aligned} \varphi_R(x) := \\ \bigvee_{\substack{i < n \\ m \leq k}} \exists x_o \dots x_{m-1} \left[\exists^{=m} y. \chi(x, y) \wedge \bigwedge_{i < m} \chi(x, x_i) \wedge \bigwedge_{i < j} x_i <_{\text{lex}} x_j \right. \\ \left. \wedge \bigvee_{\substack{i < n \\ j < m}} [|x| = |x_o| + ki + j \wedge \psi_R(x_j)] \right]. \end{aligned} \quad \square$$

Exercise 12.5. Prove that the class of unary automatic structures is closed under binary disjoint unions and binary ordered sums, but it is not closed under binary products. \lrcorner

Exercise 12.6. Let $\mathfrak{A} = \langle A, f \rangle$ be a countable structure with a bijective function $f : A \rightarrow A$. Prove that \mathfrak{A} is unary automatic if, and only if,

- ♦ the size of the finite orbits of f is bounded and
- ♦ there are only finitely many infinite orbits of f . \lrcorner

Exercise 12.7. Let $\mathcal{A} = \langle A, \sim \rangle$ be a countable structure where $\sim \subseteq A \times A$ is an equivalence relation. Then \mathcal{A} is unary automatic if, and only if,

- ♦ the cardinality of the finite \sim -classes is bounded and
- ♦ there are only finitely many infinite \sim -classes. J

Exercise 12.8. A linear order \mathcal{A} is unary automatic if, and only if, \mathcal{A} can be written as a finite sum of orders of the form $\mathbb{1}$, ω , or ω^{op} . J

Notes

Automatic structures were first defined by Hodgeson [104, 106, 105] and later rediscovered by Khoussainov and Nerode in [116]. The approach via interpretations is due to [14, 24]. Σ_1^1 -completeness of the isomorphism problem was proved in [117].

The constructions of injective presentations are from [14, 52], while the non-existence proof in the ω -automatic case is due to [103]. The material on ω -automatic structures is mostly taken from [125, 112, 122, 1]. Counting quantifiers were introduced in [14, 119], the quantifier H in [186], and the quantifier U in [126].

Most of the algebraic characterisations are taken from [14, 68, 120, 117, 149, 207, 83, 108, 1, 187]. The paper [32] claims to improve on the results of [207] by proving a characterisation of all automatic torsion-free abelian groups. But the proof sketched in [32] does not seem to work. So I consider this result a mere conjecture at the moment. A good treatise on automatic groups in the group-theoretic sense can be found in [79]. A short introduction to automaton groups can be found in [12]. There exists an extensive literature about which sequences and sets of numbers are regular when encoded in various numeration systems. See [3] and [4] for an overview.

Unary automatic structures were introduced in [14, 118]. The generalisation to structures of polynomial growth is due to [8]. Algebraic characterisations of such structures can be found in [107, 85]. Note that the proof in [85] contains a gap. The statement they actually prove is Theorem 12.32 above. Proposition 12.2 is from [200]. For surveys on Presburger arithmetic

and semilinear sets, see, e.g., [98, 46]. Theorem 12.13 is by [154, 86, 87]. The full version of Theorem 12.14 can be found in [151] and Proposition 12.18 is from [210].

For space reasons, several topics had to be omitted from this chapter. These include

- ♦ complexity questions related to automatic structures [24, 135, 121, 127, 2];
- ♦ questions regarding decidability and undecidability [185, 120, 117, 148, 123];
- ♦ equivalence and non-equivalence of presentations, intrinsically regular relations [119, 7, 8];
- ♦ automatic structures with non-regular advice [211];
- ♦ automatic structures over words of length ω^α and over other kinds of structures [82, 83, 111, 109].

Bibliography

- [1] F. ABU ZAID, E. GRÄDEL, L. KAISER, AND W. PAKUSA, *Model-Theoretic Properties of ω -Automatic Structures*, *Theory Comput. Syst.*, 55 (2014), pp. 856–880.
- [2] F. ABU ZAID, D. KUSKE, AND P. LINDNER, *Climbing up the elementary complexity classes with theories of automatic structures*, in 27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4–7, 2018, Birmingham, UK, D. R. Ghica and A. Jung, eds., vol. 119 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018, pp. 3:1–3:16.
- [3] B. ADAMCZEWSKI AND J. BELL, *Automata in number theory*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 913–945.
- [4] J.-P. ALLOUCHE AND J. O. SHALLIT, *Automatic Sequences – Theory, Applications, Generalizations*, Cambridge University Press, 2003.
- [5] A. V. ANISIMOV, *Group languages*, *Kibernetika*, 4 (1971), pp. 18–24. English translation: *Cybernetics* 4 (1973), pp. 594–601.
- [6] J. T. BALDWIN AND S. SHELAH, *Second-Order Quantifiers and the Complexity of Theories*, *Notre Dame Journal of Formal Logic*, 29 (1985), pp. 229–303.
- [7] V. BÁRÁNY, *Invariants of automatic presentations and semi-synchronous transductions*, in STACS 2006, 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23–25, 2006, Proceedings, B. Durand and W. Thomas, eds., vol. 3884 of Lecture Notes in Computer Science, Springer, 2006, pp. 289–300.
- [8] ———, *Automatic Presentations of Infinite Structures*, Ph. D. Thesis, RWTH Aachen, Aachen, 2007.
- [9] V. BÁRÁNY, E. GRÄDEL, AND S. RUBIN, *Automata-based presentations of infinite structures*, in Finite and Algorithmic Model Theory, J. Esparza, C. Michaux, and C. Steinhorn, eds., vol. 379 of London Mathematical Society Lecture Note Series, Cambridge University Press, 2011, pp. 1–76.

- [10] V. BARANY, L. KAISER, AND A. RABINOVICH, *Expressing cardinality quantifiers in monadic second-order logic over trees*, Fundamenta Informaticae, 100 (2010), pp. 1–18.
- [11] K. BARTHELMANN, *On equational simple graphs*, Tech. Rep. 9, Universität Mainz, Institut für Informatik, 1997.
- [12] L. BARTHOLDI AND P. V. SILVA, *Groups defined by automata*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 871–911.
- [13] M. BILKOWSKI AND M. SKRZYPCZAK, *Unambiguity and Uniformization Problems on Infinite Trees*, in Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2–5, 2013, Torino, Italy, vol. 23 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013, pp. 81–100.
- [14] A. BLUMENSATH, *Automatic Structures*, Diploma Thesis, RWTH Aachen, Aachen, 1999.
- [15] ———, *Structures of Bounded Partition Width*, Ph. D. Thesis, RWTH Aachen, Aachen, 2003.
- [16] ———, *Axiomatising tree-interpretable structures*, Theory of Computing Systems, 37 (2004), pp. 3–27.
- [17] ———, *A Model Theoretic Characterisation of Clique-Width*, Annals of Pure and Applied Logic, 142 (2006), pp. 321–350.
- [18] ———, *On the structure of graphs in the Caucal hierarchy*, Theoretical Computer Science, 400 (2008), pp. 19–45.
- [19] ———, *Guarded Second-Order Logic, Spanning Trees, and Network Flows*, Logical Methods in Computer Science, 6 (2010).
- [20] ———, *Simple Monadic Theories and Indiscernibles*, Mathematical Logic Quarterly, 57 (2011), pp. 65–87.
- [21] ———, *Simple Monadic Theories and Partition Width*, Mathematical Logic Quarterly, 57 (2011), pp. 409–431.
- [22] A. BLUMENSATH, T. COLCOMBET, AND C. LÖDING, *Logical Theories and Compatible Operations*, in Logic and Automata: History and Perspectives, J. Flum, E. Grädel, and T. Wilke, eds., Amsterdam University Press, 2007, pp. 73–106.

- [23] A. BLUMENSATH AND B. COURCELLE, *On the Monadic Second-Order Transduction Hierarchy*, Logical Methods in Computer Science, 6 (2010).
- [24] A. BLUMENSATH AND E. GRÄDEL, *Finite Presentations of Infinite Structures: Automata and Interpretations*, Theory of Computing Systems, 37 (2004), pp. 641–674.
- [25] A. BLUMENSATH, M. OTTO, AND M. WEYER, *Decidability Results for the Boundedness Problem*, Logical Methods in Computer Science, 10 (2014).
- [26] M. BOJAŃCZYK, *Languages Recognised by Finite Semigroups and their generalisations to objects such as Trees and Graphs with an emphasis on definability in Monadic Second-Order Logic*. lecture notes, arXiv:2008.11635, 2020.
- [27] M. BOJAŃCZYK AND W. CZERWIŃSKI, *An Automata Toolbox*. lecture notes, 2018.
- [28] M. BOJAŃCZYK, T. IDZIASZEK, AND M. SKRZYPCZAK, *Regular Languages of Thin Trees*, in Proc. 30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, 2013, pp. 562–573.
- [29] M. BOJAŃCZYK AND M. PILIPCZUK, *Definability equals recognizability for graphs of bounded treewidth*, in Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5–8, 2016, 2016, pp. 407–416.
- [30] ———, *Optimizing tree decompositions in MSO*, in 34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8–11, 2017, Hannover, Germany, vol. 66 of LIPIcs, 2017, pp. 15:1–15:13.
- [31] W. W. BOONE AND G. HIGMAN, *An algebraic characterization of groups with soluble word problem*, J. Austral. Math. Soc., 18 (1974), pp. 41–53.
- [32] G. BRAUN AND L. STRÜNGMANN, *Breaking up Finite Automata Presentable Torsion-Free Abelian Groups*, Int. J. Algebra Comput., 21 (2011), pp. 1463–1472.
- [33] S. BRAUNFELD AND M. LASKOWSKI, *Characterizations of Monadic NIP*, Transactions of the American Mathematical Society, Series B, 8 (2021), pp. 948–970.
- [34] J. R. BÜCHI, *Weak Second-Order Arithmetic and Finite Automata*, Z. Math. Logik und Grundl. Math., 6 (1960), pp. 66–92.

- [35] ———, *On a decision method in restricted second-order arithmetic*, in Proc. 1960 Int. Congr. for Logic, Methodology and Philosophy of Science, 1962, pp. 1–11.
- [36] J. R. BÜCHI AND L. H. LANDWEBER, *Solving Sequential Conditions by Finite-State Strategies*, Trans. Amer. Math. Soc., 139 (1969), pp. 367–378.
- [37] C. S. CALUDE, S. JAIN, B. KHOUSSAINOV, W. LI, AND F. STEPHAN, *Deciding parity games in quasipolynomial time*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19–23, 2017, 2017, pp. 252–263.
- [38] A. CARAYOL, *Automates infinis, logiques et langages*, PhD Thesis, Université de Rennes I, 2006.
- [39] A. CARAYOL AND C. LÖDING, *MSO on the Infinite Binary Tree: Choice and Order.*, in Proc. 16th Ann. Conference of the European Association for Computer Science Logic, LNCS 4646, 2007, pp. 161–176.
- [40] A. CARAYOL, C. LÖDING, D. NIWIŃSKI, AND I. WALUKIEWICZ, *Choice functions and well-orderings over the infinite binary tree*, Central European Journal of Mathematics, 8 (2010), pp. 662–682.
- [41] A. CARAYOL AND O. SERRE, *Higher-order recursion schemes and their automata models*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 1295–1341.
- [42] A. CARAYOL AND S. WÖHRLE, *The Caucal Hierarchy of Infinite Graphs in Terms of Logic and Higher-order Pushdown Automata*, in Proc. 23rd Conf. on Foundations of Software Technology and Theoretical Computer Science, FSTTCS, LNCS 2914, 2004.
- [43] O. CARTON, T. COLCOMBET, AND G. PUPPIS, *An Algebraic Approach to MSO-Definability on Countable Linear Orderings*, J. Symb. Log., 83 (2018), pp. 1147–1189.
- [44] D. CAUCAL, *On Infinite Transition Graphs Having a Decidable Monadic Theory*, LNCS, 1099 (1996), pp. 194–205.
- [45] ———, *On Infinite Terms Having a Decidable Monadic Theory*, in Proc. 27th Int. Symp. on Mathematical Foundations of Computer Science, MFCS, 2002, pp. 165–176.
- [46] D. CHISTIKOV, *An Introduction to the Theory of Linear Integer Arithmetic*, in 44th IARCS Annual Conference on Foundations of Software Technology

- and Theoretical Computer Science, FSTTCS 2024, December 16–18, 2024, Gandhinagar, Gujarat, India, S. Barman and S. Lasota, eds., vol. 323 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, pp. 1:1–1:36.
- [47] J. CHUZHUY AND Z. TAN, *Towards tight(er) bounds for the excluded grid theorem*. unpublished note, arXiv:1901.07944v1.
- [48] T. COLCOMBET, *A Combinatorial Theorem for Trees*, in Proc. 34th Int. Colloquium on Automata, Languages and Programming, ICALP, 2007, pp. 901–912.
- [49] ———, *On Factorization Forests*, Tech. Rep. hal-00125047, Iria Rennes, 2007.
- [50] ———, *Factorization forests for infinite words and applications to countable scattered linear orderings*, Theor. Comput. Sci., 411 (2010), pp. 751–764.
- [51] ———, *The Factorisation Forest Theorem*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 653–693.
- [52] T. COLCOMBET AND C. LÖDING, *Transforming Structures by Set Interpretations*, Logical Methods in Computer Science, 3 (2007).
- [53] T. COLCOMBET AND D. NIWIŃSKI, *On the positional determinacy of edge-labeled games*, Theoretical Computer Science, 352 (2006), pp. 190–196.
- [54] B. COURCELLE, *The monadic second-order logic of graphs II: Infinite Graphs of Bounded Width*, Math. System Theory, 21 (1989), pp. 187–221.
- [55] ———, *The monadic second-order logic of graphs I: Recognizable sets of finite graphs*, Information and Computation, 85 (1990), pp. 12–75.
- [56] ———, *The monadic second-order logic of graphs IV: Definability properties of equational graphs*, Annals of Pure and Applied Logic, 49 (1990), pp. 193–255.
- [57] ———, *Structural Properties of Context-Free Sets of Graphs Generated by Vertex Replacement*, Information and Computation, 116 (1995), pp. 275–293.
- [58] ———, *The monadic second-order logic of graphs XIV: Uniformly sparse graphs and edge set quantifications*, Theoretical Computer Science, 299 (2003), pp. 1–36.
- [59] B. COURCELLE AND J. ENGELFRIET, *A Logical Characterization of the Sets of Hypergraphs Defined by Hyperedge Replacement Grammars*, Math. System Theory, 28 (1995), pp. 515–552.

- [60] ———, *Graph Structure and Monadic Second-Order Logic*, Cambridge University Press, 2012.
- [61] B. COURCELLE, J. ENGELFRIET, AND G. ROZENBERG, *Handle-Rewriting Hypergraph Grammars*, *Journal of Computer and System Science*, 46 (1993), pp. 218–270.
- [62] B. COURCELLE AND S. OLARIU, *Upper bounds to the clique width of graphs*, *Discrete Applied Mathematics*, 101 (2000), pp. 77–114.
- [63] B. COURCELLE AND S.-I. OUM, *Vertex-Minors, Monadic Second-Order Logic, and a Conjecture by Seese*, *Journal Combinatorial Theory B*, 97 (2007), pp. 91–126.
- [64] B. COURCELLE AND I. WALUKIEWICZ, *Monadic Second-Order Logic, Graph Coverings and Unfoldings of Transition Systems*, *Ann. Pure Appl. Log.*, 92 (1998), pp. 35–62.
- [65] W. DAMM, *The IO- and OI-Hierarchies*, *Theoretical Computer Science*, 20 (1982), pp. 95–207.
- [66] W. DAMM AND A. GOERDT, *An Automata-Theoretical Characterization of the OI-Hierarchy*, *Information and Control*, 71 (1986), pp. 1–32.
- [67] S. DATTA, A. MUKHERJEE, T. SCHWENTICK, N. VORTMEIER, AND T. ZEUME, *A strategy for dynamic programs: Start over and muddle through*, *Logical Methods Computer Science*, 15 (2019).
- [68] C. DELHOMMÉ, *Automaticité des ordinaux et des graphes homogènes*, *Comptes Rendus de L'Académie des Sciences*, 339 (2004), pp. 5–10.
- [69] R. DIESTEL, *Graph Minors 1: A Short Proof of the Path-Width Theorem*, *Comb. Probab. Comput.*, 4 (1995), pp. 27–30.
- [70] ———, *Graph Theory*, Springer, 1st ed., 1997.
- [71] ———, *Graph Theory*, Springer, 5th ed., 2016.
- [72] J. DREIER, N. MÄHLMANN, S. SIEBERTZ, AND S. TORUNCZYK, *Indiscernibles and Flatness in Monadically Stable and Monadically NIP Classes*, in 50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10–14, 2023, Paderborn, Germany, vol. 261 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 125:1–125:18.
- [73] M. J. DUNWOODY, *Cutting up Graphs*, *Combinatorica*, 2 (1982), pp. 15–23.
- [74] ———, *The accessibility of finitely presented groups*, *Invent. Math.*, 81 (1985), pp. 449–457.
- [75] F. DURAND AND M. RIGO, *On Cobham's Theorem*, in *Handbook of Automata Theory*, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 947–986.
- [76] S. EILENBERG AND M.-P. SCHÜTZENBERGER, *Rational Sets in Commutative Monoids*, *J. Algebra*, 13 (1969), pp. 173–191.
- [77] M. ELBERFELD, M. GROHE, AND T. TANTAU, *Where First-Order and Monadic Second-Order Logic Coincide*, *ACM Trans. Comput. Log.*, 17 (2016), pp. 25:1–25:18.
- [78] C. C. ELGOT, *Decision Problems of Finite Automata Design and Related Arithmetics*, *Trans. Amer. Math. Soc.*, 98 (1961), pp. 21–51.
- [79] D. B. A. EPSTEIN, J. W. CANNON, D. F. HOLT, S. V. F. LEVY, M. S. PATERSON, AND W. P. THURSTON, *Word Processing in Groups*, Jones and Bartlett Publishers, Boston, 1992.
- [80] S. FEFERMAN AND R. L. VAUGHT, *The first order properties of products of algebraic systems*, *Fundamenta Mathematicae*, XLVII (1959), pp. 57–103.
- [81] O. FINKEL AND S. TODORČEVIĆ, *The Isomorphism Relation Between Tree-Automatic Structures*, *Cent. Eur. J. Math.*, 8 (2010), pp. 299–313.
- [82] ———, *A hierarchy of tree-automatic structures*, *J. Symb. Log.*, 77 (2012), pp. 350–368.
- [83] ———, *Automatic Ordinals*, *Int. J. Unconv. Comput.*, 9 (2013), pp. 61–70.
- [84] D. GALE AND F. M. STEWART, *Infinite Games with Perfect Information*, in *Contributions to the Theory of Games*, vol. 28 of *Annals of Mathematical Studies*, Princeton University Press, 1953, pp. 245–266.
- [85] M. GANARDI AND B. KHOUSSAINOV, *Automatic Equivalence Structures of Polynomial Growth*, in 28th EACSL Annual Conference on Computer Science Logic, CSL 2020, January 13–16, 2020, Barcelona, Spain, M. Fernández and A. Muscholl, eds., vol. 152 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, pp. 21:1–21:16.
- [86] S. GINSBURG AND E. SPANIER, *Bounded ALGOL-Like Languages*, *Trans. Amer. Math. Soc.*, 113 (1964), pp. 333–368.
- [87] ———, *Semigroups, Presburger Formulas, and Languages*, *Pacific Journal of Mathematics*, 16 (1966), pp. 285–296.

- [88] S. S. GONCHAROV AND J. F. KNIGHT, *Computable structure and non-structure theorems*, Algebra and Logic, 41 (2002), pp. 351–373.
- [89] E. GRÄDEL, C. HIRSCH, AND M. OTTO, *Back and Forth Between Guarded and Modal Logics*, ACM Transactions on Computational Logics, 3 (2002), pp. 418–463.
- [90] E. GRÄDEL AND I. WALUKIEWICZ, *Positional Determinacy of Games with Infinitely Many Priorities*, Logical Methods in Computer Science, 2 (2006).
- [91] R. GREGORAC, *On Generalized Free Products of Finite Extensions of Free Groups*, J. London Math. Soc., 41 (1966), pp. 662–666.
- [92] M. GROMOV, *Groups of polynomial growth and expanding maps*, Publications Mathématique d’IHÉS, 53 (1981), pp. 53–78.
- [93] Y. GUREVICH, *Modest Theory of Short Chains. I*, Journal of Symbolic Logic, 44 (1979), pp. 481–490.
- [94] Y. GUREVICH, M. MAGIDOR, AND S. SHELAH, *The Monadic Theory of ω_2* , J. Symb. Log., 48 (1983), pp. 387–398.
- [95] Y. GUREVICH AND S. SHELAH, *Modest Theory of Short Chains. II*, Journal of Symbolic Logic, 44 (1979), pp. 491–502.
- [96] ———, *Monadic Theory of Order and Topology in ZFC*, Ann. Math. Log., 23 (1982), pp. 179–198.
- [97] ———, *Rabin’s Uniformization Problem*, J. Symb. Log., 48 (1983), pp. 1105–1119.
- [98] C. HAASE, *A Survival Guide to Presburger Arithmetic*, ACM SIGLOG News, 5 (2018), pp. 67–82.
- [99] M. HAGUE, A. MURAWSKI, L. ONG, AND O. SERRE, *Collapsible Pushdown Automata and Recursion Schemes*, in Proc. 23rd IEEE Symp. on Logic in Computer Science, LICS, 2008, pp. 452–461.
- [100] R. HALIN, *S-functions for graphs*, Journal of Geometry, 8 (1976), pp. 171–186.
- [101] G. HIGMAN, *Subgroups of finitely presented groups*, Proc. Royal Soc. London, Series A, 262 (1961), pp. 455–475.
- [102] G. HJORTH, *Borel equivalence relations*, in Handbook of Set Theory, M. Foreman and A. Kanamori, eds., Springer Verlag, 2010, pp. 297–332.

- [103] G. HJORTH, B. KHOUSSAINOV, A. MONTALBÁN, AND A. NIES, *From Automatic Structures to Borel Structures*, in Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24–27 June 2008, Pittsburgh, PA, USA, IEEE Computer Society, 2008, pp. 431–441.
- [104] B. R. HODGSON, *Théories décidables par automate fini*, Ph. D. Thesis, Université de Montréal, Montréal, 1976.
- [105] ———, *On direct products of automaton decidable theories*, Theor. Comput. Sci., 19 (1982), pp. 331–335.
- [106] ———, *Théories décidables par automate fini*, Annales de sciences mathématiques du Québec, 7 (1983), pp. 39–57.
- [107] M. HUSCHENBETT, *Model-Theoretic Complexity of Automatic Linear Orders*, Ph. D. Thesis, Ilmenau University, 2016.
- [108] M. HUSCHENBETT, A. KARTZOW, J. LIU, AND M. LOHREY, *Tree-Automatic Well-Founded Trees*, Log. Methods Comput. Sci., 9 (2013), pp. 1–44.
- [109] M. HUSCHENBETT, A. KARTZOW, AND P. SCHLICHT, *Pumping for ordinal-automatic structures*, Comput., 6 (2017), pp. 125–164.
- [110] R. ITO, *Every Semilinear Set is a Finite Union of Disjoint Linear Sets*, J. Comput. Syst. Sci., 3 (1969), pp. 221–231.
- [111] S. JAIN, B. KHOUSSAINOV, P. SCHLICHT, AND F. STEPHAN, *Tree-automatic scattered linear orders*, Theor. Comput. Sci., 626 (2016), pp. 83–96.
- [112] L. KAISER, S. RUBIN, AND V. BÁRÁNY, *Cardinality and counting quantifiers on omega-automatic structures*, in STACS 2008, 25th Annual Symposium on Theoretical Aspects of Computer Science, 2008, pp. 385–396.
- [113] H. KAMP, *Tense Logic and the Theory of Linear Order*, Ph. D. Thesis, University of California, Los Angeles, 1968.
- [114] A. KARRASS, A. PIETROWSKI, AND D. SOLITAR, *Finite and Infinite Cyclic Extensions of Free Groups*, J. Austral. Math. Soc., 16 (1973), pp. 458–466.
- [115] A. S. KECHRIS, *Classical Descriptive Set Theory*, Springer-Verlag, 1995.
- [116] B. KHOUSSAINOV AND A. NERODE, *Automatic Presentations of Structures*, LNCS, 960 (1995), pp. 367–392.
- [117] B. KHOUSSAINOV, A. NIES, S. RUBIN, AND F. STEPHAN, *Automatic Structures: Richness and Limitations*, Logical Methods Computer Science, 3 (2007).

- [118] B. KHOUSSAINOV AND S. RUBIN, *Graphs with Automatic Presentations over a Unary Alphabet*, J. Autom. Lang. Comb., 6 (2001), pp. 467–480.
- [119] B. KHOUSSAINOV, S. RUBIN, AND F. STEPHAN, *Definability and Regularity in Automatic Structures*, in STACS 2004, 21st Annual Symposium on Theoretical Aspects of Computer Science, Montpellier, France, March 25–27, 2004, Proceedings, V. Diekert and M. Habib, eds., vol. 2996 of Lecture Notes in Computer Science, 2004, pp. 440–451.
- [120] ———, *Automatic Linear Orders and Trees*, ACM Trans. Comput. Log., 6 (2005), pp. 675–700.
- [121] D. KUSKE, *Theories of automatic structures and their complexity*, in Algebraic Informatics, Third International Conference, CAI 2009, Thessaloniki, Greece, May 19–22, 2009, Proceedings, S. Bozapalidis and G. Rahonis, eds., vol. 5725 of Lecture Notes in Computer Science, Springer, 2009, pp. 81–98.
- [122] ———, *Is Ramsey’s Theorem omega-automatic?*, in 27th International Symposium on Theoretical Aspects of Computer Science, STACS 2010, March 4–6, 2010, Nancy, France, J. Marion and T. Schwentick, eds., vol. 5 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2010, pp. 537–548.
- [123] D. KUSKE, J. LIU, AND M. LOHREY, *The isomorphism problem on classes of automatic structures*, in Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11–14 July 2010, Edinburgh, United Kingdom, IEEE Computer Society, 2010, pp. 160–169.
- [124] D. KUSKE AND M. LOHREY, *Decidable Theories of Graphs, Factorized Unfoldings and Cayley-graphs*, Tech. Rep. 37, University of Leicester, 2002.
- [125] ———, *First-order and counting theories of omega-automatic structures*, J. Symb. Log., 73 (2008), pp. 129–150.
- [126] ———, *Hamiltonicity of Automatic Graphs*, in Fifth IFIP International Conference On Theoretical Computer Science – TCS 2008, IFIP 20th World Computer Congress, TC 1, Foundations of Computer Science, September 7–10, 2008, Milano, Italy, G. Ausiello, J. Karhumäki, G. Mauri, and C. L. Ong, eds., vol. 273 of IFIP, 2008, pp. 445–459.
- [127] ———, *Automatic structures of bounded degree revisited*, J. Symb. Log., 76 (2011), pp. 1352–1380.

- [128] D. LAPOIRE, *Recognizability Equals Monadic Second-Order Definability for Sets of Graphs of Bounded Tree-Width*, in Proc. 15th Annual Symp. on Theoretical Aspects of Computer Science, STACS, LNCS, 1373, 1998, pp. 618–628.
- [129] H. LÄUCHLI, *A decision procedure for the weak second order theory of linear order*, in Proc. Logic Colloquium 1966, 1968, pp. 189–197.
- [130] H. LÄUCHLI AND J. LEONARD, *On the elementary theory of linear order*, in Fundamenta Mathematicae, 1966, pp. 109–116.
- [131] H. LÄUCHLI AND C. SAVIOZ, *Monadic Second Order Definable Relations on the Binary Tree*, J. Symb. Log., 52 (1987), pp. 219–226.
- [132] A. LEAF AND P. D. SEYMOUR, *Tree-Width and Planar Minors*, J. Comb. Theory, Ser. B, 111 (2015), pp. 38–53.
- [133] S. LIFSCHES AND S. SHELAH, *Uniformization, choice functions and well orders in the class of trees*, J. Symbolic Logic, 61 (1996), pp. 1206–1227.
- [134] ———, *Uniformization and skolem functions in the class of trees*, J. Symbolic Logic, 63 (1998), pp. 103–127.
- [135] M. LOHREY, *Automatic structures of bounded degree*, in Logic for Programming, Artificial Intelligence, and Reasoning, 10th International Conference, LPAR 2003, Almaty, Kazakhstan, September 22–26, 2003, Proceedings, M. Y. Vardi and A. Voronkov, eds., vol. 2850 of Lecture Notes in Computer Science, Springer, 2003, pp. 346–360.
- [136] J. A. MAKOWSKY, *Algorithmic aspects of the Feferman-Vaught Theorem*, Annals of Pure and Applied Logic, 126 (2004), pp. 159–213.
- [137] D. A. MARTIN, *Borel Determinacy*, Annals of Mathematics, 102 (1975), pp. 363–371.
- [138] A. N. MASLOV, *The hierarchy of indexed languages of an arbitrary level*, Doklady Akademii Nauk, 217 (1974), pp. 1013–1016.
- [139] ———, *Multilevel stack automata*, Problemy Peredachi Informatsii, 12 (1976), pp. 55–62.
- [140] O. MATZ, N. SCHWEIKARDT, AND W. THOMAS, *The monadic quantifier alternation hierarchy over grids and graphs*, Information and Computation, 179 (2002), pp. 356–383.
- [141] R. McNAUGHTON, *Testing and generating infinite sequences by a finite automaton*, Information and Control, 9 (1966), pp. 521–530.

- [142] R. McNAUGHTON AND S. A. PAPERT, *Counter-Free Automata*, MIT Press, 1971.
- [143] B. MOHAR AND C. THOMASSEN, *Graphs on Surfaces*, Johns Hopkins University Press, 2001.
- [144] D. E. MULLER AND P. E. SCHUPP, *Groups, the Theory of Ends, and Context-Free Languages*, J. of Computer and System Science, 26 (1983), pp. 295–310.
- [145] ———, *The Theory of Ends, Pushdown Automata, and Second-Order Logic*, Theoretical Computer Science, 37 (1985), pp. 51–75.
- [146] J. NEŠETŘIL, E. SOPENA, AND L. VIGNAL, *T-preserving homomorphisms of oriented graphs*, Comment. Math. Univ. Carolinae, 38 (1997), pp. 125–136.
- [147] G. A. NIBLO, *A Geometric Proof of Stallings' Theorem on Groups with More than One End*, Geometriae Dedicata, 105 (2004), pp. 61–76.
- [148] A. NIES, *Describing groups*, Bull. Symb. Log., 13 (2007), pp. 305–339.
- [149] A. NIES AND R. THOMAS, *FA-presentable groups and rings*, J. Algebra, 320 (2008), pp. 569–585.
- [150] D. NIWIŃSKI, *On the cardinality of sets of infinite trees recognizable by finite automata*, in Mathematical Foundations of Computer Science 1991, 16th International Symposium, MFCS'91, Kazimierz Dolny, Poland, September 9–13, 1991, Proceedings, A. Tarlecki, ed., vol. 520 of Lecture Notes in Computer Science, Springer, 1991, pp. 367–376.
- [151] R. PARIKH, *On Context-Free Languages*, J. ACM, 13 (1966), pp. 570–581.
- [152] P. PARYS, *A Pumping Lemma for Pushdown Graphs of Any Level*, in 29th Int. Symp. on Theoretical Aspects of Computer Science (STACS 2012), 2012, pp. 54–65.
- [153] D. PERRIN AND J.-É. PIN, *Infinite Words – Automata, Semigroups, Logic and Games*, Elsevier, 2004.
- [154] M. PRESBURGER, *Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt*, in Comptes Rendus du I congrès de Mathématiciens des Pays Slaves, 1929, pp. 92–101.
- [155] G. PUPPIS, *Automata for Branching and Layered Temporal Structures: An Investigation into Regularities of Infinite Transition Systems*, vol. 5955 of Lecture Notes in Computer Science, Springer, 2010.

- [156] M. O. RABIN, *Decidability of second-order theories and automata on infinite trees*, Trans. Amer. Math. Soc., 141 (1969), pp. 1–35.
- [157] A. RABINOVICH, *The full binary tree cannot be interpreted in a chain*, Journal of Symbolic Logic, 75 (2010), pp. 1489–1498.
- [158] ———, *A Proof of Kamp's Theorem*, Logical Methods Computer Science, 10 (2014).
- [159] ———, *A Proof of Stavi's Theorem*, Logical Methods Computer Science, 14 (2018).
- [160] N. ROBERTSON AND P. D. SEYMOUR, *Graph Minors I. Excluding a Forest*, Journal of Combinatorial Theory B, 35 (1983), pp. 39–61.
- [161] ———, *Graph Minors III. Planar Tree-Width*, Journal of Combinatorial Theory B, 36 (1984), pp. 49–64.
- [162] ———, *Graph Minors II. Algorithmic Aspects of Tree-Width*, J. Algorithms, 7 (1986), pp. 309–322.
- [163] ———, *Graph Minors V. Excluding a Planar Graph*, Journal of Combinatorial Theory B, 41 (1986), pp. 92–114.
- [164] ———, *Graph Minors VI. Disjoint paths across a disc*, J. Comb. Theory, Ser. B, 41 (1986), pp. 115–138.
- [165] ———, *Graph Minors VII. Disjoint paths on a surface*, J. Comb. Theory, Ser. B, 45 (1988), pp. 212–254.
- [166] ———, *Graph Minors IV. Tree-width and well-quasi-ordering*, J. Comb. Theory, Ser. B, 48 (1990), pp. 227–254.
- [167] ———, *Graph Minors IX. Disjoint crossed paths*, J. Comb. Theory, Ser. B, 49 (1990), pp. 40–77.
- [168] ———, *Graph Minors VIII. A Kuratowski theorem for general surfaces*, J. Comb. Theory, Ser. B, 48 (1990), pp. 255–288.
- [169] ———, *Graph Minors X. Obstructions to Tree-Decomposition*, Journal of Combinatorial Theory B, 52 (1991), pp. 153–190.
- [170] ———, *Graph Minors XI. Circuits on a Surface*, J. Comb. Theory, Ser. B, 60 (1994), pp. 72–106.
- [171] ———, *Graph Minors XII. Distance on a Surface*, J. Comb. Theory, Ser. B, 64 (1995), pp. 240–272.

- [172] ———, *Graph Minors XIII. The Disjoint Paths Problem*, J. Comb. Theory, Ser. B, 63 (1995), pp. 65–110.
- [173] ———, *Graph Minors XIV. Extending an Embedding*, J. Comb. Theory, Ser. B, 65 (1995), pp. 23–50.
- [174] ———, *Graph Minors XV. Giant Steps*, J. Comb. Theory, Ser. B, 68 (1996), pp. 112–148.
- [175] ———, *Graph Minors XVII. Taming a Vortex*, J. Comb. Theory, Ser. B, 77 (1999), pp. 162–210.
- [176] ———, *Graph Minors XVI. Excluding a non-planar graph*, J. Comb. Theory, Ser. B, 89 (2003), pp. 43–76.
- [177] ———, *Graph Minors XVIII. Tree-decompositions and well-quasi-ordering*, J. Comb. Theory, Ser. B, 89 (2003), pp. 77–108.
- [178] ———, *Graph Minors XIX. Well-quasi-ordering on a surface*, J. Comb. Theory, Ser. B, 90 (2004), pp. 325–385.
- [179] ———, *Graph Minors XX. Wagner’s conjecture*, J. Comb. Theory, Ser. B, 92 (2004), pp. 325–357.
- [180] ———, *Graph Minors XXI. Graphs with unique linkages*, J. Comb. Theory, Ser. B, 99 (2009), pp. 583–616.
- [181] ———, *Graph Minors XXIII. Nash-Williams’ immersion conjecture*, J. Comb. Theory, Ser. B, 100 (2010), pp. 181–205.
- [182] ———, *Graph Minors XXII. Irrelevant vertices in linkage problems*, J. Comb. Theory, Ser. B, 102 (2012), pp. 530–563.
- [183] N. ROBERTSON, P. D. SEYMOUR, AND R. THOMAS, *Quickly Excluding a Planar Graph*, Journal of Combinatorial Theory B, 62 (1994), pp. 323–348.
- [184] J. G. ROSENSTEIN, *Linear Orderings*, Academic Press, 1982.
- [185] S. RUBIN, *Automatic Structures*, Ph. D. Thesis, University of Auckland, Auckland, 2004.
- [186] ———, *Automata Presenting Structures: A Survey of the Finite String Case*, Bull. Symb. Log., 14 (2008), pp. 169–209.
- [187] ———, *Automatic Structures*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 1031–1070.

- [188] A. SCHRIJVER, *Theory of Linear and Integer Programming*, Wiley, 1998.
- [189] F. SCHUPP, *Reflections on HNN Extensions*, in Theory of Groups, M. F. Newman, ed., Lecture Notes in Mathematics 372, Springer Verlag, 1974, pp. 611–632.
- [190] M. P. SCHÜTZENBERGER, *On Finite Monoids Having Only Trivial Subgroups*, Information and Control, 8 (1965), pp. 190–194.
- [191] D. SEESE, *The structure of the models of decidable monadic theories of graphs*, Annals of Pure and Applied Logic, 53 (1991), pp. 169–195.
- [192] A. L. SEMENOV, *Decidability of monadic theories*, LNCS, 176 (1984), pp. 162–175.
- [193] S. SHELAK, *The Monadic Second Order Theory of Order*, Annals of Mathematics, 102 (1975), pp. 379–419.
- [194] ———, *Classification Theory and the Number of Non-Isomorphic Models*, North-Holland, 1978.
- [195] I. SIMON, *Factorization Forests of Finite Height*, Theoretical Computer Science, 72 (1990), pp. 65–94.
- [196] M. SKRZYPCZAK, *Descriptive Set Theoretic Methods in Automata Theory – Decidability and Topological Complexity*, vol. 9802 of Lecture Notes in Computer Science, Springer, 2016.
- [197] J. R. STALLINGS, *Group Theory and Three-Dimensional Manifolds*, Yale University Press, 1971.
- [198] C. STIRLING, *Decidability of Bisimulation Equivalence for Pushdown Processes*, unpublished, 2000.
- [199] B. STURMFELS, *On Vector Partition Functions*, J. Combin. Theory Ser. A, 72 (1995), pp. 302–309.
- [200] A. SZILARD, S. YU, K. ZHANG, AND J. O. SHALLIT, *Characterizing regular languages with polynomial densities*, in Mathematical Foundations of Computer Science 1992, 17th International Symposium, MFCS’92, Prague, Czechoslovakia, August 24–28, 1992, Proceedings, I. M. Havel and V. Koubek, eds., vol. 629 of Lecture Notes in Computer Science, Springer, 1992, pp. 494–503.
- [201] T. TAO AND V. VU, *Additive Combinatorics*, Cambridge University Press, 2006.

- [202] W. THOMAS, *Automata on Infinite Objects*, in Handbook of Theoretical Computer Science, J. van Leeuwen, ed., vol. B, Elsevier, Amsterdam, 1990, pp. 135–191.
- [203] ———, *Ehrenfeucht Games, the Composition Method, and the Monadic Theory of Ordinal Words*, LNCS, 1261 (1997), pp. 118–143.
- [204] ———, *Languages, Automata, and Logic*, in Handbook of Formal Languages, G. Rozenberg and A. Salomaa, eds., vol. 3, Springer, New York, 1997, pp. 389–455.
- [205] J. J. L. TOURNEAU, *Decision problems related to the concept of operation*, Ph. D. Thesis, University of California, Berkeley, 1968.
- [206] B. A. TRAKHTENBROT, *The Synthesis of Logical Nets Whose Operators are Described in Terms of One-Place Predicate Calculus*, Dokl. Akad. Nauk. SSSR, 118 (1958), pp. 646–649.
- [207] T. TSANKOV, *The additive group of the rationals does not have an automatic presentation*, J. Symb. Log., 76 (2011), pp. 1341–1351.
- [208] I. WALUKIEWICZ, *Monadic second-order logic on tree-like structures*, Theoretical Computer Science, 275 (2002), pp. 311–346.
- [209] S. WÖHRLE, *Decision Problems over Infinite Graphs – Higher-Order Pushdown Systems and Synchronized Products*, PhD Thesis, RWTH Aachen, Aachen, 2005.
- [210] K. WOODS, *Presburger Arithmetic, Rational Generating Functions, and Quasi-Polynomials*, J. Symb. Log., 80 (2015), pp. 433–449.
- [211] F. A. ZAID, E. GRÄDEL, AND F. REINHARDT, *Advice automatic structures and uniformly automatic classes*, in 26th EACSL Annual Conference on Computer Science Logic, CSL 2017, August 20–24, 2017, Stockholm, Sweden, V. Goranko and M. Dam, eds., vol. 82 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, pp. 35:1–35:20.
- [212] E. ZERMELO, *Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels*, in Proc. 5th Congress of Mathematicians, Cambridge University Press, 1913, pp. 501–504.
- [213] W. ZIELONKA, *Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees*, Theoretical Computer Science, 200 (1998), pp. 135–183.

Symbol Index

Chapter I

$[n]$	$\{0, \dots, n-1\}$, 3	$\wp_{\text{fin}}(\mathfrak{A})$	finite power-set structure, 11
\bar{a}	tuple, 3	Σ_{in}	signature of the incidence structure, 13
$\langle \rangle$	empty tuple, 3	\mathfrak{A}_{in}	incidence structure, 13
$\wp(A)$	power set, 3	in_k	incidence relation, 13
$A + B$	disjoint union, 3	$\text{qr}(\varphi)$	quantifier rank, 16
$\text{dom } f$	domain of a function, 3	$\text{Th}_l(M)$	L -theory of M , 16
$\text{rng } f$	range of a function, 3	$M \sqsubseteq_L N$	directed L -equivalence, 17
$f \upharpoonright X$	restriction to X , 3	$M \equiv_L N$	L -equivalence, 17
$\uparrow X$	elements above X , 3	$\text{Mod}(\varphi)$	models of φ , 17
$\downarrow X$	elements below X , 3	$\mathfrak{A} \oplus \mathfrak{B}$	disjoint union, 24
$M \models \varphi$	satisfaction relation, 4	$\sum_{i \in I} \mathfrak{A}_i$	ordered sum, 35
$\text{FO}[\Sigma]$	first-order logic, 5	$\mathfrak{A}_o + \mathfrak{A}_i$	finite ordered sum, 36
$\text{MSO}[\Sigma]$	monadic second-order logic, 5	\mathfrak{A}/\approx	quotient, 40
$\mathfrak{A} \models \varphi(\bar{a}, \bar{P})$	satisfaction relation, 6	$\text{fuse}_P(\mathfrak{A})$	fusion, 41
$\text{WMSO}[\Sigma]$	weak monadic-second order logic, 7	$\text{copy}_k(\mathfrak{A})$	k copies of \mathfrak{A} , 43
$\text{CMSO}[\Sigma]$	MSO with first-order counting, 7	$\sum_{i \in \mathbb{N}} \mathfrak{A}_i$	generalised sum, 45
$\text{MSO}[\text{inf}]$	MSO with finiteness predicate, 7	$\text{MSO}_n^\circ[\Sigma]$	MSO° -formulae with quantifier structure \bar{n} , 46
$\text{GSO}[\Sigma]$	guarded second-order logic, 7	$\text{Th}_{\text{MSO}^\circ}(\mathfrak{A}, \bar{P})$	MSO° -theory of quantifier structure \bar{n} , 46
$\text{MSO}^\circ[\Sigma]$	simplified syntax for MSO, 8	$\mathfrak{A}, \bar{P} \equiv_{\text{MSO}^\circ} \mathfrak{B}, \bar{Q}$	\bar{n} -equivalence, 46
$\wp(\mathfrak{A})$	power-set structure, 11	$\llbracket \chi(\bar{P}) \rrbracket$	set of indices satisfying χ , 46
		$\mathfrak{A} \times \mathfrak{B}$	direct product, 52
		\mathfrak{A}^k	power of \mathfrak{A} , 53

Chapter II

ω	first infinite ordinal, 61
$\langle \rangle$	empty sequence, 62
$ w $	length of a word, 62
Σ^*	finite words, 62
Σ^+	finite non-empty words, 62
Σ^ω	ω -words, 62
$w[i, k]$	factor of w , 63
\mathfrak{S}^1	semigroup with added neutral element, 64
$a \leq_L b$	Greene's relation, 66
$a \leq_R b$	Greene's relation, 66
$a \leq_J b$	Greene's relation, 66
$a \equiv_L b$	Greene's relation, 66
$a \equiv_R b$	Greene's relation, 66
$a \equiv_J b$	Greene's relation, 66
$a \equiv_H b$	Greene's relation, 66
$[a]_L b$	L-class of a , 67
$[a]_R b$	R-class of a , 67
$[a]_J b$	J-class of a , 67
$[a]_H b$	H-class of a , 67
Ξ_σ	relation induced by a split, 76
$p \xrightarrow{a} q$	transition, 86
$L(\mathcal{A})$	language recognised by \mathcal{A} , 86
\sim_L	syntactic congruence, 87
LTL	linear temporal logic, 99
$\varphi \cup \psi$	until operator, 99
$X\varphi$	next operator, 100
$F\varphi$	finally operator, 100
$G\varphi$	generally operator, 100
$F_*\varphi$	reflexive finally operator, 100
$G_*\varphi$	reflexive generally operator, 100

Chapter III

$\text{Chain}_J(\mathfrak{S})$	set of J-chains, 110
$\mu(\vec{a})$	last element of \vec{a} , 111
$\pi(\vec{a})$	product of \vec{a} , 111
\leq_λ	equivalent w.r.t. J-chain labelling, 111
a^ω	ω -power, 121
$a_{[i, k]}$	product $a_i \cdots a_{k-1}$, 123
$p \xrightarrow{a} q$	transition, 130
$L(\mathcal{A})$	language recognised by \mathcal{A} , 130
$\text{WMSO}[\Sigma]$	weak monadic second-order logic, 139

Chapter IV

V_\diamond	positions for \diamond , 146
V_\square	positions for \square , 146
W_σ	winning region, 146
$\diamond X$	some successor in X , 148
$\square X$	all successor in X , 148
$\bar{\sigma}$	opponent, 148
Step_σ	step function, 148
Attr_σ	attractor, 148
$\mathcal{G}(\mathfrak{S}, \varphi)$	model-checking game, 155
$\mathfrak{S} \times_\alpha M$	product game, 165
P_σ	period set, 167
L_μ	μ -calculus, 179

Chapter V

$x \leq_{\text{pf}} y$	prefix ordering, 193
$\text{Suc}(x)$	set of successors, 193

$\text{Suc}_*(x)$	pointed set of successors, 193
T_v	attached subtree, 194
$\text{dom}(t)$	domain of t , 194
\leq_{so}	successor order, 194
$u \sqcap v$	infimum, 195
Σ_Q	signature for transition logic, 207
$\mathfrak{Succ}_*(v; \rho)$	successor structure, 207
$L_{\text{nd}}(\mathcal{A})$	language recognised by a nondeterministic automaton, 208
$\mathfrak{Succ}_*(v; \rho/q)$	successor structure for alternating automata, 212
$L_{\text{alt}}(\mathcal{A})$	language recognised by an alternating automaton, 212
$\bar{\varphi}$	dual of φ , 214
$\mathfrak{S}_\Omega(Q)$	trace semigroup, 219
\mathfrak{A}^*	Muchnik iteration, 228
cl	clone relation, 228
$\mathcal{G}_o(\mathcal{A}, \mathfrak{A})$	reduced automaton game, 229
$\mathfrak{T}_{\text{bin}}$	infinite binary tree, 236
$[\mathfrak{T}]$	space of branches, 244
$\partial \mathfrak{X}$	Cantor-Bendixson derivative, 245
$\text{CB}(\mathfrak{X})$	Cantor-Bendixson rank of a space, 245
$\text{CB}(x, \mathfrak{X})$	Cantor-Bendixson rank of a point, 245
$\partial \mathfrak{T}$	Lifsches-Shelah derivative, 248
$\text{LS}(\mathfrak{T})$	Lifsches-Shelah rank, 248
$\text{LS}(x, \mathfrak{X})$	Cantor-Bendixson rank of a point, 248
C_2MSO	MSO with second-order counting, 259
\exists^k	counting quantifier, 259

$\text{Ch}(P)$	P -choice vertices, 266
$\text{CB}(P)$	P -choice branches, 266
$\tilde{U} \approx_* \tilde{V}$	finite difference, 271
t_v	subtree at v , 275
$p(\vec{s})$	substitution of trees, 275

Chapter VI

$\sum_{i \in I} \mathfrak{A}_i$	ordered sum, 283
$\mathfrak{A} \times \mathfrak{B}$	ordered product, 283
\mathfrak{A}^{op}	opposite ordering, 284
$[a, b]$	interval, 284
$[a, b)$	interval, 284
cn	condensation map, 288
$\text{FC}(\mathfrak{A})$	finite-condensation rank, 288
$\text{gFC}(\mathfrak{A})$	finite-condensation rank, 288
$\text{FC}_*(\mathfrak{A})$	variant of FC-rank, 289
$\text{gFC}_*(\mathfrak{A})$	variant of generalised FC-rank, 289
$\text{HR}(\mathfrak{A})$	Hausdorff rank, 291
$\text{HR}_*(\mathfrak{A})$	variant of Hausdorff rank, 291
$\text{len}(\mathfrak{A})$	length, 291
\leq_{lex}	lexicographic ordering, 295
$\text{cl}_2(X)$	two-sided closure, 327
$\text{cl}(X)$	topological closure, 333
$\text{len}(P)$	level in the tree, 339
$[[\varphi(\vec{P})]]$	elements where φ holds locally, 362
$\mathfrak{A} \models_{\text{loc}} \varphi(\vec{P})$	local satisfaction, 362
\leq_{lllex}	length-lexicographic order, 408
$\text{WO}(\mathfrak{A})$	well-ordering index, 425
BLTL	linear temporal logic, 441
$\varphi \cup \psi$	until operator, 442

$\varphi \mathcal{S} \psi$ since operator, 442
 LTL linear temporal logic, 442
 $X\varphi$ next operator, 442
 $Y\varphi$ yesterday operator, 442
 $F\varphi$ eventual operator, 442
 $G\varphi$ always operator, 442

Chapter VII

\mathfrak{H}_{in} incidence representation, 462
 $\mathfrak{H}_{\text{adj}}$ adjacency representation, 462
 $\mathfrak{G} \subseteq \mathfrak{H}$ subhypergraph, 462
 \mathfrak{H}_U induced subhypergraph, 463
 A_ν auxiliary nodes, 472
 $B(X/\mathfrak{F})$ vertices X is attached to, 473
 $\beta(X/\mathfrak{F})$ maximal vertex X is attached to, 473
 $\mathcal{O}_\varepsilon(\mathfrak{H})$ oriented hypergraph, 487
 $\mathfrak{B} \cup_{\mathfrak{A}} \mathfrak{C}$ amalgamation, 492
 $B_{\mathfrak{B}}(Z)$ border of Z , 497
 d_f defect of f , 498
 $\text{Ch}_k(P)$ vertices with a lot of choice, 520
 $\text{pbr}(C)$ prime branching point, 521
 $\text{mpr}_k(C)$ minimal prefix, 524
 $\text{ppr}_k(C)$ prime prefix, 525

Chapter VIII

$\text{twd } \mathfrak{A}$ tree-width, 531
 $\text{pwd } \mathfrak{A}$ path-width, 532
 $\text{twd}_n \mathfrak{A}$ height- n tree-width, 532
 Min minor transduction, 543

$\mathfrak{G} - F$ deleting edges, 547
 ∂A boundary of A , 548
 $N(P)$ set of neighbours, 554
 $\mathfrak{K}_{m,n}$ complete bipartite graph, 562
 $C(k)$ k -th column, 562
 $R(i)$ i -th row, 562
 $\hat{\sigma}$ edge-flow of a branch-decomposition, 578
 $\text{supp}(X)$ support of X , 578
 $\text{rk}(X)$ rank of a set of edges, 578
 τ° leaf labelling induced by a flow, 584
 $B(\tau)$ blocks of a k -flow, 584

Chapter IX

\simeq_U crossing equivalence, 600
 $\text{crk}(P/\mathfrak{A})$ crossing rank, 600
 $\mathfrak{A}[P]$ induced substructure with auxiliary relations, 603
 $\bigotimes_{i \in I}^{\tau} \mathfrak{A}_i$ composition with an update specification, 606
 $H^n[P]$ the classes in P , 607
 $\text{cwd } \mathfrak{A}$ crossing-width, 610
 $\text{val}(t)$ value of t , 618
 $\text{cwd}_{\text{ns}} \mathfrak{A}$ non-standard crossing width, 639

Chapter X

$\exp_m(\mathfrak{A})$ unary expansion of \mathfrak{A} , 654
 $\tau(\mathfrak{A}/\bar{P})$ transduction with parameters, 662

\mathcal{E}_n n -equivalence structures, 667
 \mathcal{T}_n trees of height n , 667
 $\mathcal{C} \sqsubseteq_L \mathcal{D}$ L -transduction from \mathcal{D} to \mathcal{C} , 685
 \mathcal{P} finite paths, 685
 \mathcal{G} finite grids, 685
 \mathcal{T}_n trees of bounded height, 686
 $\text{Ad}(t)$ adhesion set, 692
 $\mu(a)$ component introducing a , 693
 $a \sqsubseteq b$ introduction order, 693
 $a \equiv b$ equivalence associated with the introduction order, 693
 $\mathfrak{A}[F]$ fragment induced by F , 697
 $\mathfrak{A}^*[F]$ fragment without root adhesion edge, 697

Chapter XI

CH_n Caucal hierarchy, 718
 $\mathfrak{A} \parallel_{\xi, \zeta} \mathfrak{B}$ parallel sum, 736
 $\text{val}(t)$ value of a HR-term, 736
 $\mathcal{U}(\mathfrak{G}, \nu_\circ)$ unravelling, 742
 $\mathfrak{G}^{*'}$ modified Muchnik iteration, 747
 $\text{nop}(z)$ no-op operation, 757
 $\text{pop}_k(z)$ pop operation, 757
 $\text{push}_y(z)$ push operation, 757
 $\text{clone}_k(z)$ clone operation, 757
 Op_n stack operations, 757
 $\text{top}(z)$ top of stack, 757
 \mathfrak{S}_n tree of higher-order stacks, 766
 $d(u, v)$ distance, 779

T_{st} subtree attached at s , 780
 $\text{End}(u)$ end containing u , 786
 $\mathfrak{G} *_{\mathfrak{C}} \mathfrak{H}$ push-out, 790
 \exp_n iterated exponential, 801
 w/k prefix, 804
 $x \equiv_m y$ same m -theory, 818
 $[x]_m$ \equiv_m -class, 818
 $x \sqsubseteq y$ existence of path, 822

Chapter XII

Σ_\square alphabet with blank, 833
 \square blank, 833
 $s \otimes t$ convolution, 834
 \tilde{s}^{\otimes} convolution, 834
 $u \models_{\text{len}} v$ equal length, 835
 $\mathbb{Z}[\frac{1}{p}]$ p -adic rationals, 836
 $\mathfrak{A} \leq_{\text{FO}} \mathfrak{B}$ FO-interpretation, 842
 $\mathfrak{T}_{\text{bin}}$ binary tree, 842
 $k \mid_p m$ power of p divides, 845
 $u \models_{\text{len}} v$ equal length, 845
 \mathfrak{A}^ω ω -power, 850
 $\mathfrak{A}^{(\omega)}$ infinite direct sum, 850
 $s \approx_* t$ only finitely many differences, 864
 $[s]_*$ \approx_* -class, 864
 $[a]_E$ equivalence class, 866
 FOC FO with counting quantifiers, 883
 $\text{H}\bar{x}\varphi$ Ramsey quantifier, 887
 $\text{UX}\varphi$ infinite quantifier, 888
 $\|a\|$ length of the encoding, 892
 $D(U)$ size of domain, 894
 $h(U)$ maximal height, 894
 $\bar{a} \simeq_U^{\Delta} \bar{b}$ same Δ -type over U , 894

ex_R	expansion of R , 900	$[a, b]$	commutator, 917
$N_\varphi(U, n)$	reachable elements, 901	$\theta(S, d)$	coverage measure, 935
$\text{prod}(\mathcal{S})$	finite direct products, 907	$a \approx_* b$	finite difference, 943
$\mathfrak{A} \cup \mathfrak{B}$	superposition, 907	$E_{ij}(r)$	transvection, 951
$\text{supp}(\mathcal{S})$	superpositions, 907	$\mathfrak{E}_n(\mathfrak{R})$	group of transvections, 951
$u \approx_*^m v$	equal from position n , 909	Ψ_A	vector partition function, 986
$[\mathfrak{G} : \mathfrak{H}]$	index of \mathfrak{H} in \mathfrak{G} , 915	$\mathfrak{E}(g)$	equivalence structure, 993

Index

n -equivalence structure, 667	automatic presentation, 837
— associated with a tree, 667	automatic structure, 837
	automaton, 86
accepting a tree, 208	L -automaton, 207
accepting a word, 86, 130	ω -automaton, 130
accepting run, 86, 130, 208, 212	automaton game, 213
accessibility length of a group, 791	auxiliary node, 472
accessible group, 791	auxiliary relations, 603
actual relations, 603	
acyclic flow, 498	back property, 21
additive colouring, 107	base of a push-out, 790
additive labelling, 314	base of an HNN-extension, 790
additive labelling of a linear order, 75	base point of an edge, 808
adhesion edge, 697	basic temporal formula, 444
adhesion set, 692	bicofinality, 314
adjacency representation of a	bicompatible operation, 15
hypergraph, 462	bidirectional linear temporal logic, 441
admissible labelling, 270	big vertex of a tree-decomposition, 554
alphabet, 62	binary tree, 236
alternating automaton, 211	binary vertex in a Simon tree, 75
alternating game, 145	blank symbol, 833
alternating parity game, 188	block of a flow, 584
aperiodic semigroup, 94	boolean algebra, 943
atom, 943	border hyperedge, 463
attached subtree, 194	border of a factor, 697
attachment set, 473	border of a set of vertices, 497
attractor, 148	Borel determinacy, 158

Borel function, 255
 Borel hierarchy, 256
 Borel relation, 255
 Borel set, 158, 255, 861
 Borel structure, 861
 boundary of a set of vertices, 548
 bounded expansion, 900
 bounded path, 808
 bramble, 553
 branch of a generalised tree, 195
 branch of a tree, 193
 branch-decomposition, 578
 branch-width, 578
 branching degree of a generalised tree, 394
 branching type of a tuple, 826
 Büchi automaton, 131

 cancellative function, 644
 Cantor normal form, 960
 Cantor topology, 244, 861
 Cantor-Bendixson derivative, 245
 Cantor-Bendixson rank, 245
 Cantor-Bendixson rank of a tree, 250
 Caucal hierarchy, 718
 Cayley graph, 914
 Cayley-graph, 777
 choice, 393
 P -choice factor of a tree, 266
 choice function, 393
 choice function for a hypergraph, 483
 clique-width, 599
 clone relation, 228
 code, 362
 K -code of a tuple, 826
 coding a path, 819
 coherent set, 327

colour trace, 616
 colour trace of a tuple, 616
 C -coloured order, 35
 coloured tree, 194
 colouring, 107
 \mathbb{Z} -colouring, 487
 m, k -colouring, 819
 colouring property of a weak subhypergraph, 469
 colours, 603
 column of a pseudo-grid, 562
 comb, 888
 commutator of a group, 917
 compatible operation, 15
 complete bipartite graph, 562
 component of a tree-decomposition, 531
 composition method, 23
 composition of path-decompositions, 701
 composition of transductions, 653
 composition theorem, 23
 computing a split by a transducer, 116
 condensation, 288, 370
 configuration graph of a pushdown automaton, 751
 conforming to a finite-memory strategy, 164
 conforming to a strategy, 146
 congruence of a linear order, 370
 connected component, 462
 context for a tree, 275
 context-free group, 778
 contracting an edge in a graph, 543
 convex equivalence relation, 196
 convex set, 196
 convolution, 833
 copy operation, 43

copying transduction, 655
 core of a star, 465
 coset progression, 926
 counting monadic-second order logic, 7
 counting quantifier, 259
 covered by a \approx_* -class, 866
 crossing equivalence, 600
 crossing formula, 599
 crossing rank, 600
 crossing-width, 610
 cut, 372
 cyclically finite semigroup, 65

 d -rearrangement of a transduction, 679
 decomposition
 tree \rightarrow , 531
 decomposition of a generalised tree
 induced by a substructure, 198
 Dedekind-complete order, 441
 Dedekind-complete orders, 372
 defect of a flow, 498
 definability, 62
 definable distribution, 499
 definable expansion, 491
 definable flow, 499
 defining a choice function, 487
 defining a tree-decomposition, 693
 definition scheme of an interpretation, 31
 dense colouring, 285
 dense linear order, 284
 dense subset, 284
 depth-first \mathfrak{H} -forest, 474
 determined games, 146
 deterministic automaton, 131
 deterministic graph, 743
 deterministic transducer, 115

k -dimensional first-order interpretation, 841
 k -dimensional interpretation, 54
 direct product, 52
 direction, 193
 disjoint union, 24
 disjoint union of sets, 3
 disjoint weak subhypergraphs, 463
 distance in graphs, 779
 distinguishable contexts, 276
 distinguishing vertex, 276
 distribution, 498
 σ -domain, 161
 draw, 145
 dual formula, 214
 Dyck sequence, 767

 ec-pushdown graph, 765
 edge ordering, 479
 edge-bounded flow, 498
 edge-flow, 578
 effective uniformisation, 414
 effectively equivalent presentations, 857
 embedding of a semigroup into a relation, 897
 embedding of generalised trees, 247
 embedding of linear orders, 284
 k -embedding of trees, 683
 embedding of well-ordered trees, 410
 emptiness for ω -automata, 141
 empty word, 62
 encoding a colouring, 487
 end of a graph, 786
 enough successors, 410
 enriched tree, 195
 entirely \mathcal{S} -colourless set, 348
 equal length, 835

equal-length relation, 845
 equality, 834
 equality checks, 764
 L -equivalence, 17
 equivalence structure, 993
 n -equivalence structure, 667
 — associated with a tree, 667
 equivalent models, 17
 equivalent presentations, 857
 equivalent to a parity condition, 167
 even cycle, 173
 expandable tree, 276
 m -expansion, 654
 m, k -expansion, 819
 expansion of a relation, 900
 extendible sequence, 548
 externally linked sets in a graph, 572

 factor of a tree, 265, 697
 factor of a word, 63
 factorisation tree, 74
 finite colouring, 107
 finite difference, 943
 finite out-degree, 892
 finite power-set structure, 11
 finite-condensation rank, 288
 finite-memory strategy, 164
 finitely branching generalised tree, 394
 first-order counting, 7
 first-order m -equivalence, 17
 first-order interpretation, 841
 first-order logic, 5
 first-order logic over words, 93
 first-order logic with counting
 quantifiers, 883
 first-order parameter, 20
 δ -flow, 498

flow in a graph, 498, 584
 \mathfrak{F} -forest, 473
 forest of ordinal height, 463
 formula, 4
 forth property, 21
 fragment of a tree-decomposition, 697
 frame embedding, 341
 frame for a coherent set, 334
 free monoid, 64
 free semigroup, 64
 free ω -semigroup, 121
 frontier of an end, 786
 full k -ary tree, 535
 fundamental domain, 930
 fusion of a structure, 41

 Gaifman graph, 533
 Gale-Stewart game, 157
 game with memory, 165
 generalised finite-condensation rank,
 288
 generalised Lifsches-Shelah rank, 296
 generalised path-decomposition, 537
 generalised product, 55
 generalised sum, 45
 generalised tree, 195
 generalised tree-decomposition, 537
 generalised vector partition function,
 986
 globally consistent colour trace, 616
 good tree-decompositions, 554
 graph minor, 543
 Greene's relations, 66
 grid, 558
 groups of polynomial growth, 922
 GSO-transduction, 654
 GSO-transduction hierarchy, 686

guard, 329
 guarded m -equivalence, 17
 guarded parameter, 20
 guarded relation, 7
 guarded second-order logic, 7
 guarded set, 330
 guarded tuple, 7

 H-class, 67
 handle of a structure, 736
 Hausdorff rank of a linear order, 291
 height of a forest, 463
 height of a partial order, 961
 height of a split, 75
 height of an order-tree, 610
 height- n tree-width, 532
 Heisenberg group, 914, 917
 higher-order automatic structure, 849
 higher-order pushdown automaton, 757
 higher-order pushdown graph, 758
 HNN-extension, 790
 holographic relation of a tree, 640
 homogeneous factorisations, 866
 homogeneous subset, 315
 homomorphism between games, 186
 homomorphism of ω -semigroups, 121
 Horn formula, 153
 Horn satisfiability, 153
 HR-equational, 737
 HR_k -term, 736
 hyperedge, 461
 hypergraph, 461
 hypertree, 465

 idempotent element, 64
 idempotent vertex in a Simon tree, 75
 immersion of games, 186

incidence relation, 13, 461
 incidence representation of a
 hypergraph, 462
 incidence structure, 13
 increasing path, 808
 indegree of a choice function, 483
 independence property, 649
 index of a subgroup, 915
 index tree of a tree-decomposition, 531
 inducing a distribution, 514
 infinite direct sum, 850
 injective presentation, 859
 inner component, 701
 inner path, 701
 integral domain, 948
 internal hyperedge, 463
 internally connected
 tree-decomposition, 693
 interpretation, 31
 interval, 284
 introducing an element in a
 tree-decomposition, 693
 introduction order of a
 tree-decomposition, 693
 isolated branch, 248
 isolated point, 245
 isomorphism of ends, 786

 J-class, 67

 L-class, 67
 labelled tree, 194
 language, 62
 latest appearance record, 132
 lattice closed logic, 17
 lattice (in \mathbb{R}^d), 930
 L\"auchli-Leonard operation, 313

leaf of a non-standard tree, 639
 leaf of a pseudo-tree, 466
 leaf of an order-tree, 610
 leaf-based interpretation, 631
 least solution of a system of equations, 320
 left-containing a minor-embedding, 572
 left-guarded split, 76
 length of a linear order, 291
 length of a word, 62
 length-lexicographic order, 408
 length-lexicographic ordering, 835
 letter, 62
 lexicographic ordering, 295, 835
 Lifschies-Shelah derivative, 248
 Lifschies-Shelah rank, 248, 296
 linear order, 61
 linear temporal logic, 99, 441
 linkage, 562
 little choice below a vertex, 520
 local colouring induced by a tree-decomposition, 696
 local satisfaction, 362
 local tree-decomposition, 697
 locally consistent colour trace, 616
 locally finite, 949
 logic, 4
 Löwenheim-Skolem theorem, 239
 LTL-transduction, 102

 many-dimensional interpretation, 54
 meet-embedding between trees, 595
 meet-irreducible element, 639
 meet-reducible element, 639
 memory-free strategy, 146
 minimal prefix of a branch, 524
 minimal separator, 780

 minor, 8
 minor of a graph, 543
 minor-embedding, 543
 modal μ -calculus, 179
 modal logic, 155
 model, 4
 model-checking game for modal logic, 155
 model-checking game for the μ -calculus, 181
 modest linear order, 330
 monadic m -equivalence, 17
 monadic expansion of a class, 648
 monadic independence property, 649
 monadic parameter, 20
 monadic second-order logic, 5
 monadic second-order logic with first-order counting, 7
 monadic second-order logic with second-order counting, 259
 monoid, 64
 move, 145
 MSO-colouring, 203
 MSO-filtering, 654
 MSO-transduction, 654
 Muchnik iteration, 228, 635, 718, 849
 Muchnik-iteration, 853
 multiplicatively independent numbers, 978
 mutually dense family, 284

 natural product of ordinals, 960
 natural sum of ordinals, 960
 neighbours, 554
 nilpotent group, 922
 non-copying transduction, 655
 non-standard crossing-width, 639

non-standard partition decomposition, 639
 non-standard tree, 537
 nondeterministic automaton, 86, 130, 207
 norm, 927
 normal path, 808
 normal spanning tree, 544
 normalised game, 188
 number of ends, 786

 ω -automatic presentation, 837
 ω -automatic structure, 837
 ω -power, 850
 ω -tree automatic presentation, 837
 ω -tree automatic structure, 837
 opposite of an ordering, 284
 order, 61
 order of a separation, 547
 order topology, 333
 order-preserving \oplus -term, 616
 order-tree, 194
 order-trees of ordinal height, 610
 ordered product, 53, 283
 ordered sum, 35, 283
 ordered tree sum, 264
 ordering, 61
 ordinal, 61
 origin map, 655
 out-degree, 797
 outcome, 145

 p -adic norm, 928
 p -adic rationals, 836
 p -divisible group, 925
 p -norm, 927
 pairing rank, 648

 parameter, 20
 parity condition, 130, 160
 parity game, 160
 partial order, 61
 partial ordering, 61
 partial play, 145
 partition decomposition, 610
 φ -path, 802
 path above, 808
 path in a hypergraph, 462
 path-decomposition, 532
 path-interpretation, 739
 path-width, 532
 perfect subset, 245
 period set, 167
 petal of a tree-decomposition, 554
 piecewise polynomial, 986
 play, 145
 player, 145
 poly-growth automatic structure, 970
 polynomial growth, 970
 polynomial growth for a group, 922
 position, 145
 positional determinacy, 147
 positional game, 145
 positional strategy, 146
 positive polynomial, 993
 ω -power, 121, 850
 power of a structure, 53
 power-set structure, 11
 predecessor, 193
 predecessor ordering, 480
 prefix, 63
 prefix ordering, 193, 835
 prefix-closed, 193
 prefix-invariant, 166
 prefix-recognisable, 725

Presburger arithmetic, 64, 979
 Presburger structure, 977
 Presburger-definable set, 977
 prime branching point, 521
 prime prefix of a branch, 525
 principal cut, 372
 principle point of attachment, 473
 priority function, 130, 160, 207
 product game, 165
 profile of a path-decomposition, 701
 progression, 926
 proper progression, 926
 pseudo-grid, 562
 pseudo-tree, 466
 pumping argument, 91
 push-out of groups, 790
 pushdown automaton of level n , 757
 pushdown automaton with
 ε -transitions, 751
 pushdown automaton with equality
 checks, 764
 pushdown graph, 755
 ε -pushdown graph, 751
 ε -pushdown graph of level n , 758

 quantifier rank, 16
 quantifier rank of an interpretation, 31
 quantifier structure, 46
 quotient structure, 40

 R-class, 67
 Rado graph, 899
 Ramsey factorisation, 123
 Ramsey quantifier, 887
 Ramseyan split, 76
 random graph, 899
 rank for reachability games, 149

rank of a hypergraph, 462
 rank of a lattice, 930
 rank of a progression, 926
 rank of a set of edges, 578
 ray of a star, 465
 reachability game, 147
 reachable elements, 901
 realising a cut, 372
 recognising a language by an automaton,
 86, 130, 208, 212
 recognising by a semigroup, 87
 recognising by an ω -semigroup, 121
 reduced game, 229
 reduced tree-decomposition, 692
 refinement of a tree-decomposition, 697
 regular game, 159
 regular J-class, 71
 regular linear order, 314
 regular path-interpretation, 739
 regular relation, 834
 regular tree, 226
 regular tree language, 222
 regular winning condition, 167
 replete set, 327
 respecting the crossing equivalence, 603
 restricted pushdown automaton, 752
 restricted pushdown graph, 755
 reversible Turing machine, 855
 right-guarded split, 75
 root of a factor, 697
 root of a pseudo-tree, 466
 row of a pseudo-grid, 562
 run, 86, 130, 207, 211
 run of a transducer, 116

 satisfaction reaction, 4
 satisfaction relation, 6

scattered linear order, 284
 second-order counting, 259
 second-order parameter, 20
 semigroup, 64
 ω -semigroup, 120
 semilinear set, 978
 separable linear order, 327
 separation of a graph, 547
 separator, 780
 set of neighbours, 554
 shifting a path, 809
 short linear order, 327
 shuffle of linear orders, 313
 signature, 4
 Simon tree, 75
 simple semilinear set, 978
 simultaneous game, 145
 skeleton of a code, 353
 skeleton of a tree, 248
 Skolem arithmetic, 64, 851, 899
 solution of a system of equations, 320
 sort, 4
 spanning \mathfrak{H} -forest, 474
 spanning tree, 544
 sparse distribution, 498
 k -sparse hypergraph, 481
 spine of a path, 808
 split, 75
 split of a tuple, 618
 splitting a group, 790
 splitting at vertices, 618
 stack operations, 757
 stack-aware stack operation, 769
 standard temporal formula, 444
 star, 465
 star-free regular expressions, 97
 state, 86, 130, 207

step function, 148
 strategy, 146
 stratification of a tree, 339
 stratum, 353
 strong choice, 393
 strong uniformisation, 414
 structure, 4
 Stupp-iteration, 853
 subdivision, 482
 subhypergraph, 462
 submodularity of branch-width, 579
 subbranch of a tree, 248
 subtree, 194
 subtree of a generalised tree, 195
 successor, 193
 successor structure, 207
 successor-ordered tree, 194
 successor-tree, 194
 suffix, 63
 sum-decomposition, 905
 sunflower, 465
 superposition, 907
 support of a set of edges, 578
 symmetric set in \mathbb{R}^d , 930
 symmetric update specification, 606
 syntactic congruence, 87, 129
 syntactic semigroup, 96

 tame tree, 394
 tangle, 587
 \oplus -term, 615
 theory, 16
 thin tree, 247
 tight linkage, 562
 tight tree-decomposition, 780
 topological closure, 333
 touching sets of vertices, 543

trace of a run, 212
 trace semigroup, 219
 transducer, 115
 transduction, 653
 transduction hierarchy for GSO, 686
 transduction with origin information, 653
 transition function, 131, 207
 transition logic, 207
 transition ω -semigroup of an automaton, 133
 transition relation, 86, 130
 transition semigroup of an automaton, 88
 transition system, 179
 transvection, 951
 trap, 149
 tree, 194
 tree automatic presentation, 837
 tree automatic structure, 837
 tree automaton, 207, 211
 tree domain, 193
 tree-decomposition, 531
 tree-interpretable structure, 716
 tree-width, 531
 triangle, 779
 κ -triangulable graph, 779
 triangulation of a cycle, 779
 trivial push-out, 790
 tropical semiring, 943
 two-sided closure of a set, 327
 type of a Ramsey factorisation, 124

 ubiquitous elements, 701
 unary automatic structure, 999
 undirected tree, 193
 uniform condensation, 370

uniform height, 667
 uniform homomorphism, 836
 \bar{n} -uniform linear order, 369
 uniform theory, 371
 uniformisation, 414
 unravelling, 742
 update specification, 606

 value of a \oplus -term, 618
 varying at a level, 348
 vector partition function, 986
 vertex, 194, 461
 vertex with enough successors, 410
 vertex-bounded flow, 498
 very tame tree, 434
 virtually abelian, 915
 virtually free group, 778
 virtually nilpotent, 922
 VR-equational structure, 736

 weak choice, 393
 weak monadic second-order logic, 139
 weak monadic-second order logic, 7
 weak Ramseyan split, 115
 weak subhypergraph, 463
 weak uniformisation, 414
 well-formed \oplus -term, 618
 well-order, 61
 well-ordered tree, 409
 well-ordering index, 425
 well-quasi-order, 109
 well-quasi-ordering, 593
 width of a tree-decomposition, 531
 Wilke algebra, 127
 win, 145
 winning a game, 146
 winning region, 146

winning strategy, 146

 word, 62

word structure, 62

 zero-dimensional linear order, 333

The Roman and Fraktur alphabets							
<i>A</i>	<i>a</i>	\mathfrak{A}	\mathfrak{a}	<i>N</i>	<i>n</i>	\mathfrak{N}	\mathfrak{n}
<i>B</i>	<i>b</i>	\mathfrak{B}	\mathfrak{b}	<i>O</i>	<i>o</i>	\mathfrak{O}	\mathfrak{o}
<i>C</i>	<i>c</i>	\mathfrak{C}	\mathfrak{c}	<i>P</i>	<i>p</i>	\mathfrak{P}	\mathfrak{p}
<i>D</i>	<i>d</i>	\mathfrak{D}	\mathfrak{d}	<i>Q</i>	<i>q</i>	\mathfrak{Q}	\mathfrak{q}
<i>E</i>	<i>e</i>	\mathfrak{E}	\mathfrak{e}	<i>R</i>	<i>r</i>	\mathfrak{R}	\mathfrak{r}
<i>F</i>	<i>f</i>	\mathfrak{F}	\mathfrak{f}	<i>S</i>	<i>s</i>	\mathfrak{S}	\mathfrak{s}
<i>G</i>	<i>g</i>	\mathfrak{G}	\mathfrak{g}	<i>T</i>	<i>t</i>	\mathfrak{T}	\mathfrak{t}
<i>H</i>	<i>h</i>	\mathfrak{H}	\mathfrak{h}	<i>U</i>	<i>u</i>	\mathfrak{U}	\mathfrak{u}
<i>I</i>	<i>i</i>	\mathfrak{I}	\mathfrak{i}	<i>V</i>	<i>v</i>	\mathfrak{V}	\mathfrak{v}
<i>J</i>	<i>j</i>	\mathfrak{J}	\mathfrak{j}	<i>W</i>	<i>w</i>	\mathfrak{W}	\mathfrak{w}
<i>K</i>	<i>k</i>	\mathfrak{K}	\mathfrak{k}	<i>X</i>	<i>x</i>	\mathfrak{X}	\mathfrak{x}
<i>L</i>	<i>l</i>	\mathfrak{L}	\mathfrak{l}	<i>Y</i>	<i>y</i>	\mathfrak{Y}	\mathfrak{y}
<i>M</i>	<i>m</i>	\mathfrak{M}	\mathfrak{m}	<i>Z</i>	<i>z</i>	\mathfrak{Z}	\mathfrak{z}

The Greek alphabet					
<i>A</i>	α	alpha	<i>N</i>	ν	nu
<i>B</i>	β	beta	<i>E</i>	ξ	xi
<i>Γ</i>	γ	gamma	<i>O</i>	\omicron	omicron
Δ	δ	delta	<i>Π</i>	π	pi
<i>E</i>	ϵ	epsilon	<i>P</i>	ρ	rho
<i>Z</i>	ζ	zeta	Σ	σ	sigma
<i>H</i>	η	eta	<i>T</i>	τ	tau
Θ	θ	theta	<i>Υ</i>	υ	upsilon
<i>I</i>	ι	iota	Φ	ϕ	phi
<i>K</i>	κ	kappa	<i>X</i>	χ	chi
Λ	λ	lambda	Ψ	ψ	psi
<i>M</i>	μ	mu	Ω	ω	omega