

Abstract Algebraic Language Theory

Achim Blumensath

Abstract
Algebraic Language Theory

Achim Blumensath

ab

BRNO 2022

Achim Blumensath
blumens@fi.muni.cz

This document was last updated 2022-12-11.
The latest version can be found at
www.fi.muni.cz/~blumens

COPYRIGHT 2022 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

I Monads	1
1. An overview	1
2. Discrete categories	8
3. Polynomial functors	21
4. Monads	33
5. Eilenberg-Moore algebras	39
6. Lifting monads	60
7. Bialgebras	81
II Algebra	93
1. Factorisation systems	93
2. Subalgebras	104
3. Reducts	116
4. Congruences	119
5. Varieties	127
III Languages	135
1. Weights	135
2. Languages	145
3. Minimal algebras	151
4. Syntactic algebras	160
5. Varieties	182
6. The profinitary term monad	191
7. Axiomatisations	213
IV Logic	219

1. Abstract logics	219
2. Compositionality	227
3. Definable algebras	233
4. Definable languages	239
V Trees	253
1. Monads and logics for trees and forests	253
2. Finite forests	263
3. Countable chains	265
4. Counterexamples	278
5. MSO-definable algebras	284
6. First-order logic	294
VI Temporal Logics	303
1. Temporal logics	303
2. Bisimulation	308
3. The logic EF	314
4. Wreath products	334
5. Distributive algebras	348
6. Path algebras	353
Recommended Literature	375
Bibliography	377
Symbol Index	379
Index	383

I. Monads

1. An Overview

There are many different formalisms to specify formal languages based on automata, grammars, regular expressions, homomorphisms, logics, and so on. The central topic of formal language theory is the study of such formalisms. In particular, we are interested in their expressive power and their algorithmic properties, i.e., which questions are decidable for them and what the respective complexity is. Several frameworks exist for answering such questions. Here we will adopt a very general category-theoretic point of view that covers many of them. Our focus will be on algebraic and logical approaches with a special emphasis on languages of infinite trees and their monadic second-order theories.

Before starting to develop the general theory, let us shortly present some of the specific examples it is supposed to subsume. We will be rather succinct and intended mainly as a reminder to readers already familiar with the material. The reader is encouraged to ignore and/or skip any parts that look incomprehensible.

Finite Words

The prototypical example of a formal language theory is that of finite words. A *finite word* over a given *alphabet* Σ is a finite sequence (possibly empty) of elements of Σ . We denote the set of all finite words by Σ^* . (When not explicitly mentioned otherwise, we will assume alphabets to be finite.) A (*formal*) *language* is a set $L \subseteq \Sigma^*$ of such words. The main algebraic framework for such languages is based on monoids (or semigroups). A *monoid* $\mathcal{M} = \langle M, \cdot, e \rangle$ is

a structure with universe M , a binary operation $\cdot : M \times M \rightarrow M$, and a constant $e \in M$ such that

- ♦ \cdot is associative: $a \cdot (b \cdot c) = (a \cdot b) \cdot c$;
- ♦ e is a neutral element: $e \cdot a = a = a \cdot e$.

Examples include

- ♦ the set $\{0, 1\}$ with the usual multiplication and the neutral element 1;
- ♦ the natural numbers $\langle \mathbb{N}, +, 0 \rangle$ with addition;
- ♦ the natural numbers $\langle \mathbb{N}, \cdot, 1 \rangle$ with multiplication;
- ♦ the set Σ^* of all finite words with *concatenation* as the product and the empty sequence $\langle \rangle$ as the neutral element.

The monoid $\langle \Sigma^*, \cdot, \langle \rangle \rangle$ is also called the *free monoid* since it has the following *universal property*: for every monoid \mathfrak{M} and every function $f : \Sigma \rightarrow M$, there exists a unique homomorphism $\varphi : \Sigma^* \rightarrow \mathfrak{M}$ that agrees with f on the elements of Σ .

A *homomorphism* $\varphi : \mathfrak{M} \rightarrow \mathfrak{N}$ of monoids is a function $\varphi : M \rightarrow N$ between their universes that preserves products and the neutral element, that is

$$\varphi(a \cdot b) = \varphi(a) \cdot \varphi(b) \quad \text{and} \quad \varphi(e) = e.$$

We can use a homomorphism $\varphi : \Sigma^* \rightarrow \mathfrak{M}$ from the free monoid to another (usually finite) monoid \mathfrak{M} to represent languages of Σ . We say that $L \subseteq \Sigma^*$ is *recognised* by φ if $L = \varphi^{-1}[P]$, for some $P \subseteq M$.

Examples. (a) The monoid $\langle \{0, 1\}, \max, 0 \rangle$ recognises the language $\Sigma^* a \Sigma^*$ via the morphism mapping a to 1 and every other letter to 0.

(b) The monoid $\mathbb{Z}/2\mathbb{Z}$ recognises the language of all words with an even number of letters a via the morphism mapping a to 1 and every other letter to 0.

Note that we can encode every language $L = \varphi^{-1}[P]$ recognised by some finite monoid \mathfrak{M} by specifying

- ♦ the multiplication table of \mathfrak{M} ,

- ♦ the neutral element of \mathfrak{M} ,
- ♦ the set P , and
- ♦ the values $\varphi(c)$, for $c \in \Sigma$.

This is a finite amount of information. Hence, we can use this encoding for algorithms that take languages as input.

Of course, not every language can be encoded this way: there are only countably many encodings but uncountably many languages (if the alphabet contains at least two letters). So, which languages can be recognised by a homomorphism in this way? It turns out that these are exactly the well-known regular languages.

Theorem 1.1. *Let $L \subseteq \Sigma^*$. The following statements are equivalent.*

- (1) L is recognised by a finite automaton.
- (2) L is the value of a regular expression.
- (3) L is denoted by a linear grammar.
- (4) L is definable in monadic second-order logic.
- (5) L is recognised by a homomorphism to a finite monoid.
- (6) The syntactic congruence of L has only finitely many classes.

Let us briefly explain the last item in the above characterisation. The *syntactic congruence* \sim_L of a language $L \subseteq \Sigma^*$ is a binary relation on Σ^* which is defined by

$$u \sim_L v \quad \text{iff} \quad (xuy \in L \Leftrightarrow xvy \in L), \quad \text{for all } x, y \in \Sigma^*.$$

It turns out that \sim_L forms a congruence on the free monoid $\langle \Sigma^*, \cdot, \langle \rangle \rangle$ and the quotient homomorphism $\Sigma^* \rightarrow \Sigma^* / \sim_L$ recognises L . Hence, if \sim_L has only finitely many classes, the quotient Σ^* / \sim_L forms a finite monoid recognising L .

Infinite Words

An ω -word over an alphabet Σ is an infinite sequence $w = (c_n)_{n < \omega}$ of letters $c_n \in \Sigma$, i.e., a function $w : \omega \rightarrow \Sigma$. We denote the set of all ω -words by Σ^ω .

To adapt the algebraic approach to infinite words, we need a suitable kind of algebra to replace monoids. This replacement is called an ω -semigroup. An ω -semigroup $\mathfrak{S} = \langle S_1, S_\omega, \cdot, \pi \rangle$ is a two-sorted structure where S_1 contains the *finite* elements and S_ω the *infinite* ones. These two sets are equipped with three kinds of multiplication:

- ♦ a *finite product* $\cdot : S_1 \times S_1 \rightarrow S_1$,
- ♦ a *mixed product* $\cdot : S_1 \times S_\omega \rightarrow S_\omega$, and
- ♦ an *infinite product* $\pi : (S_1)^\omega \rightarrow S_\omega$.

All products are assumed to be associative. That is,

$$\begin{aligned} a \cdot (b \cdot c) &= (a \cdot b) \cdot c, \\ a \cdot (b \cdot u) &= (a \cdot b) \cdot u, \\ b \cdot \pi(a_0, a_1, \dots) &= \pi(b, a_0, a_1, \dots), \\ \pi(a_0, a_1, a_2, \dots) &= \pi((a_0 \cdots a_{k_0-1}), (a_{k_0} \cdots a_{k_1-1}), \dots), \end{aligned}$$

for all $a, b, c, a_0, a_1, \dots \in S_1$, $u \in S_\omega$, and $0 < k_0 < k_1 < \dots < \omega$. In case of a finite ω -semigroup, we can replace the infinite product by an ω -power operation

$$a^\omega := \pi(a, a, a, \dots).$$

The resulting kind of algebra is called a *Wilke algebra*. Using a straightforward Ramsey argument one can show that every finite Wilke algebra is associated with a unique ω -semigroup.

Again, a *free ω -semigroup* $\langle \Sigma^+, \Sigma^\omega, \cdot, \pi \rangle$ consists of all words over some alphabet Σ . The finite elements are the finite, non-empty words $w \in \Sigma^+$, the infinite ones are the infinite words $w \in \Sigma^\omega$.

A *homomorphism* $\varphi : \mathfrak{S} \rightarrow \mathfrak{T}$ of ω -semigroups consists of a pair of maps $\varphi_1 : S_1 \rightarrow T_1$ and $\varphi_\omega : S_\omega \rightarrow T_\omega$ that commute with all three products.

Examples. Let $\Sigma = \{a, b\}$.

(a) The language of all words $w \in \Sigma^\omega$ containing the letter a is recognised by the morphism $\varphi : \Sigma^\omega \rightarrow \mathfrak{S}$ where

$$S_1 := \{0, 1\}, \quad S_\omega := \{0, 1\}, \quad \varphi(a) := 1, \quad \varphi(b) := 0,$$

and the product is just the maximum operation.

(b) The language of all words $w \in \Sigma^\omega$ containing infinitely many occurrences of the letter a is recognised by the morphism $\varphi : \Sigma^\omega \rightarrow \mathfrak{S}$ where

$$S_1 := \{0, 1\}, \quad S_\omega := \{0, 1\}, \quad \varphi(a) := 1, \quad \varphi(b) := 0,$$

and the product is defined by

$$\begin{aligned} c \cdot d &:= \max\{c, d\}, & \pi(c_0, c_1, \dots) &:= \limsup_{n < \omega} c_n, \\ c \cdot u &:= u, \end{aligned}$$

for $c, d, c_0, c_1, \dots \in S_1$ and $u \in S_\omega$.

Again we obtain the following equivalent characterisations of the class of regular languages.

Theorem 1.2. *Let $L \subseteq \Sigma^\omega$. The following statements are equivalent.*

- (1) L is recognised by a finite automaton.
- (2) L is definable in monadic second-order logic.
- (3) L is recognised by a homomorphism to a finite ω -semigroup.
- (4) The syntactic congruence of L has only finitely many classes.

Finite and Infinite Trees

Many results about words (finite or infinite) and their languages generalise to trees. There are two main kinds of trees considered in language theory: ranked ones and unranked ones. An *unranked tree* over an alphabet Σ is a directed tree whose vertices are labelled by letters from Σ . To define *ranked trees* we have to equip each symbol $a \in \Sigma$ with an *arity* $\text{ar}(a) \in \omega$. A Σ -labelled directed tree is *ranked* if the arity of a vertex' label coincides with the number of outgoing edges. Both kinds of trees come in finite or infinite versions. For instance, every term over a signature Γ can be seen as a ranked tree.

Algebraic descriptions of tree languages turn out to be more complicated than those for words. The simplest one uses so-called *forest algebras* to recognise languages of finite trees (and forests). A forest algebra $\langle H, V, \cdot, + \rangle$ is

two-sorted: H is the domain for *forests* and V is the domain for *contexts*. It is equipped with three *horizontal products* and two *vertical ones*

$$\begin{aligned} + : H \times H &\rightarrow H, & \cdot : V \times H &\rightarrow H, \\ + : H \times V &\rightarrow V, & \cdot : V \times V &\rightarrow V, \\ + : V \times H &\rightarrow V \end{aligned}$$

satisfying various associativity laws.

A more popular formalism is based on automata. A *tree automaton* is device that labels each vertex of the given input tree with a state in a way that is consistent with the transition relation. For infinite trees, there is an additional condition for every branch concerning the states appearing on it infinitely many times.

Weighted Languages

Sometimes it is helpful to view a language not only as a set of words (or trees, or ...), but as a function associating with every word some *weight*. For instance, instead of just recording whether or not a word is accepted by a given automaton, one could count how many different ways there are for the automaton to accepting the word (if the automaton is nondeterministic). Or if it is a more complicated automaton model, one could count how many operations of a certain kind the automaton uses, or what the maximal value of some counter is, and so on. This leads to the notion of a *weighted language*, which is just a function $\kappa : \Sigma^* \rightarrow \Omega$, where Ω is some set of *weights*. Frequently, one assumes that Ω forms a semiring. In this case the set of all weighted languages $\Sigma^* \rightarrow \Omega$ also forms a semiring, which is usually denoted $\Omega\langle\Sigma^*\rangle$ and called the semiring of *formal power-series* over Ω . It is then customary to write a function $\kappa : \Sigma^* \rightarrow \Omega$ as a formal infinite sum

$$\sum_{w \in \Sigma^*} a_w w$$

with coefficients $a_w := \kappa(w) \in \Omega$. We will not use this notation, since the analogy with sums breaks down when considering weighted languages of other objects than words.

Again, we can set up some algebraic machinery to define weighted languages where homomorphisms into finite monoids are now replaced by homomorphisms into *finite-dimensional Ω -modules*.

Data Words

When considering *infinite* alphabets one can preserve many nice properties of the finite-alphabet case if one uses formalisms that cannot check for specific letters. Instead, we only allow checking whether two letters are *equal* or, more generally, whether they satisfy any of a fixed set of predefined relations. Thus, one works with alphabets with additional structure like $\langle \mathbb{N}, = \rangle$ or $\langle \mathbb{Q}, \leq \rangle$ and one requires that every language is closed under automorphisms of the alphabet. An example of such a language is the one of all words that contain some letter at least twice.

Technically this means working with sets $\langle X, \alpha \rangle$ equipped with a group action $\alpha : \mathfrak{G} \rightarrow \text{Aut}(X)$. Such sets are called \mathfrak{G} -set, where \mathfrak{G} is the group in question. Since an action is just a homomorphism $\alpha : \mathfrak{G} \rightarrow \text{Aut}(X)$, the category $\mathfrak{G}\text{-Set}$ of all \mathfrak{G} -sets is just the comma category $(\mathfrak{G} \downarrow \text{Aut})$. In this category the rôle of finite sets is played by the *orbit-finite* ones: sets with only finitely many orbits under the associated action.

One can now define automata in $\mathfrak{G}\text{-Set}$ as follows. Let Σ be an orbit-finite \mathfrak{G} -set serving as the alphabet. For the set of states, we take some \mathfrak{G} -set Q . Usually $Q = Q_0 \times \Sigma^n$ consists of a finite part Q_0 together with several registers holding letters of the alphabet. The transition relation then takes the usual form $\Delta \subseteq Q \times \Sigma \times Q$. Instead of requiring Q and Δ to be finite, we now assume that they are orbit-finite. For instance, to check that some letter appears at least twice, an automaton can non-deterministically guess some position, store the letter at this position in memory, and then compare it with each of the remaining letters until there is a match.

Prerequisites and Notation

Unfortunately, in a book like this some prior knowledge of category theory has to be assumed. While I have tried to keep the prerequisites at a minimum,

I assume that the reader has worked through some basic introductory text covering, in particular, limits and adjunctions. More advanced concepts like monads, distributive laws, etc. will be introduced below as needed.

Let us fix some basic notation. For $n < \omega$, we set $[n] := \{0, \dots, n-1\}$. In a partial order, we denote the *upwards closure* of a set X by $\uparrow X$ and its *downwards closure* by $\downarrow X$. We denote the *comma category* for two functors F and G by $(F \downarrow G)$. For a family $(A_i)_{i \in I}$ of objects in a category, we denote the *product* by $\prod_{i \in I} A_i$ and the *coproduct* by $\sum_{i \in I} A_i$. The associated canonical morphisms are

$$\begin{aligned} \prod_{i \in I} f_i : \prod_{i \in I} A_i &\rightarrow \prod_{i \in I} B_i, & \sum_{i \in I} f_i : \sum_{i \in I} A_i &\rightarrow \sum_{i \in I} B_i, \\ \langle g_i \rangle_{i \in I} : A &\rightarrow \prod_{i \in I} B_i, & \bigoplus_{i \in I} h_i : \sum_{i \in I} A_i &\rightarrow B. \end{aligned}$$

Note that the notation for the last one is non-standard. We denote the *terminal object* by 1 .

2. Discrete Categories

We start by introducing the kind of category we are working in. The most important one is of course the category \mathbf{Set} of sets. But sometimes it is useful to equip the sets under consideration with some additional structure. For instance, when characterising logics that are not closed under negation, it will be necessary to use ordered sets. Similarly, when dealing with data words, we need to equip the sets with a group action. We will therefore work more generally in some a base category \mathcal{D} that behaves sufficiently like \mathbf{Set} for the proofs below to go through, but that is general enough to cover all the cases we are interested in. Let us quickly introduce the main categories we are interested in before presenting the generalisation we will use.

Definition 2.1. We denote by \mathbf{Set} the category of all sets and functions, and by \mathbf{Pos} the category of all partial orders and monotone functions. \mathbf{Top} denotes the category of topological spaces and continuous maps, and \mathbf{Met} is

the category of bounded metric spaces and non-expansive functions, that is, metric spaces (X, d) satisfying

$$d(x, y) \leq 1, \quad \text{for all } x, y \in X,$$

and functions $f : X \rightarrow Y$ satisfying

$$d(f(x), f(y)) \leq d(x, y), \quad \text{for all } x, y \in X.$$

Finally, for a group \mathcal{G} , we denote by $\mathcal{G}\text{-Set}$ the category of all \mathcal{G} -sets, i.e., sets S equipped with an action $\mathcal{G} \times S \rightarrow S$ and functions preserving this action.

Our base category \mathcal{D} will be one of these categories or one similar to them. One of the things these categories have in common is that there exists an adjunction $\mathbb{J} \dashv \mathbb{V}$ between \mathbf{Set} and \mathcal{D} ; that is, every object $A \in \mathcal{D}$ has an underlying set $\mathbb{V}A$ and we can equip every set $X \in \mathbf{Set}$ with the ‘weakest possible’ \mathcal{D} -structure $\mathbb{J}X$. For instance, for $\mathcal{D} = \mathbf{Pos}$, $\mathbb{V}A$ is just the underlying set of the partial order A , while $\mathbb{J}B$ is the set B equipped with the trivial order $=$. Similarly for $\mathcal{D} = \mathbf{Top}$, \mathbb{V} maps a topological space to its underlying set while \mathbb{J} equips a set with the discrete topology. Furthermore, we will assume that the forgetful functor $\mathbb{V} : \mathcal{D} \rightarrow \mathbf{Set}$ is faithful and that the objects in \mathcal{D} are ‘discrete’ in the sense that \mathbb{V} commutes with coproducts. That means that the underlying set of a coproduct $\sum_i A_i$ is just the disjoint union of the sets $\mathbb{V}A_i$. Note that this rules out most algebraic categories like groups, modules, or term algebras.

Faithfulness of \mathbb{V} ensures that every morphism of \mathcal{D} is uniquely determined by the induced function on the underlying sets. But note that not every such function needs to correspond to a morphism of \mathcal{D} . (Note very function is monotone/continuous/etc..) Furthermore, \mathbb{V} induces a partial order \sqsubseteq on the objects of \mathcal{D} by

$$A \sqsubseteq B \quad \text{iff} \quad \text{there exists a morphism } f : A \rightarrow B \text{ with } \mathbb{V}f = \text{id}.$$

Intuitively, $A \sqsubseteq B$ means that A and B have the same underlying set, but the additional structure of A is more permissive/general/weaker than that of B .

For instance, if $\mathcal{D} = \mathbf{Pos}$, we have

$$\langle A, \leq \rangle \subseteq \langle A', \leq' \rangle \quad \text{iff} \quad A = A' \quad \text{and} \quad \leq \subseteq \leq'.$$

Similarly, for $\mathcal{D} = \mathbf{Top}$, we have

$$\langle X, \mathcal{T} \rangle \subseteq \langle X', \mathcal{T}' \rangle \quad \text{iff} \quad X = X' \quad \text{and} \quad \mathcal{T} \supseteq \mathcal{T}'.$$

Let us formalise these assumptions on \mathcal{D} .

Definition 2.2. (a) A category \mathcal{D} is *disjunctive* if, for every morphism $\varphi : A \rightarrow \sum_{i \in I} B_i$, there exist unique objects A_i and unique morphisms $\varphi_i : A_i \rightarrow B_i$, for $i \in I$, such that

$$A = \sum_{i \in I} A_i \quad \text{and} \quad \varphi = \sum_{i \in I} \varphi_i.$$

Uniqueness here means that, if A'_i and φ'_i are other such objects and morphisms, there exist isomorphisms $\sigma_i : A'_i \rightarrow A_i$ such that $\varphi_i \circ \sigma = \varphi'_i$, for all i .

(b) A category \mathcal{D} is *discrete* if it is disjunctive, has arbitrary coproducts, and there exists an adjunction $\mathbb{J} \dashv \mathbb{V}$ between \mathbf{Set} and \mathcal{D} such that \mathbb{V} is faithful and preserves coproducts, \lrcorner

We will develop our language theory for a base category \mathcal{D} that is (i) discrete, (ii) has arbitrary colimits, and (iii) all countable limits. In this chapter, we will still be explicit about these assumptions; in later chapters, they will be left implicit.

Examples. Let us check that the categories introduced above are discrete.

(a) For \mathbf{Set} , we can take $\mathbb{V} := \text{Id}$. To show that \mathbf{Set} is disjunctive, let $f : C \rightarrow A + B$ be a function. Then $C = f^{-1}[A] + f^{-1}[B]$ and $f = (f \upharpoonright f^{-1}[A]) + (f \upharpoonright f^{-1}[B])$ and this decomposition is clearly unique.

(b) For \mathbf{Pos} , we define the adjunction $\mathbb{J} \dashv \mathbb{V}$ by

$$\mathbb{V}\langle A, \leq \rangle := A \quad \text{and} \quad \mathbb{J}A := \langle A, = \rangle.$$

The unit $\varepsilon : A \rightarrow A$ and the counit $\iota : \langle A, = \rangle \rightarrow \langle A, \leq \rangle$ of this adjunction are given by the identity maps. The proof that \mathbf{Pos} is disjunctive is the same as for \mathbf{Set} .

A limit is computed as in \mathbf{Set} with the ordering defined component-wise, while coproducts $\sum_i \langle A_i, \leq \rangle$ are simply disjoint unions. Note that the forgetful functor $\mathbb{V} : \mathbf{Pos} \rightarrow \mathbf{Set}$ does not preserve pushouts: consider the two bijections $f, g : \langle \{a, b\}, = \rangle \rightarrow \langle \{0, 1\}, \leq \rangle$. The pushout of f and g has a single element, while the one of $\mathbb{V}f$ and $\mathbb{V}g$ has two. To compute an arbitrary colimit in \mathbf{Pos} one first forms the corresponding colimit in \mathbf{Set} . In general, this results in a preorder. To obtain the colimit in \mathbf{Pos} we have to take the quotient of this preorder by the associated equivalence relation.

(c) For \mathbf{Top} , we define the adjunction $\mathbb{J} \dashv \mathbb{V}$ by

$$\mathbb{V}\langle X, \mathcal{C} \rangle := X \quad \text{and} \quad \mathbb{J}X := \langle X, \wp(X) \rangle.$$

The unit $\varepsilon : X \rightarrow X$ and the counit $\iota : \langle X, \wp(X) \rangle \rightarrow \langle X, \mathcal{C} \rangle$ of this adjunction are given by the identity maps. The proof that \mathbf{Top} is disjunctive is the same as for \mathbf{Set} .

Note that the functor \mathbb{V} also has a right adjoint: the functor $\mathbb{R} : \mathbf{Set} \rightarrow \mathbf{Top}$ mapping each set X to the space $\mathbb{R}X := \langle X, \{\emptyset, X\} \rangle$ with the trivial topology.

A limit $\lim D \subseteq \prod_i D(i)$ is computed as in \mathbf{Set} with the topology induced by the product topology on $\prod_i D(i)$, and coproducts $\sum_i \langle X_i, \mathcal{C}_i \rangle$ are disjoint unions.

(d) For $\mathbf{G}\text{-Set}$, we define the adjunction $\mathbb{J} \dashv \mathbb{V}$ by

$$\mathbb{V}\langle X, \alpha \rangle := X \quad \text{and} \quad \mathbb{J}X := \langle G \times X, \gamma \rangle,$$

where

$$\gamma(h)(\langle g, x \rangle) := \langle hg, x \rangle, \quad \text{for } g, h \in G \text{ and } x \in X.$$

The corresponding bijection maps $f : X \rightarrow Y$ in \mathbf{Set} to

$$\langle G \times X, \gamma \rangle \rightarrow \langle Y, \alpha \rangle : \langle g, x \rangle \mapsto \alpha(g)(f(x))$$

in $\mathfrak{G}\text{-Set}$. The unit $\varepsilon : X \rightarrow G \times X$ and the counit $\iota : \langle G \times X, \gamma \rangle \rightarrow \langle X, \alpha \rangle$ of this adjunction are given by

$$\varepsilon(x) := \langle e, x \rangle \quad \text{and} \quad \iota(g, x) := \alpha(g)(x).$$

Limits are again computed as in Set with an action that is defined component-wise. Coproducts $\sum_i \langle X_i, \alpha_i \rangle$ are disjoint unions and the proof of disjointness is again similar to that of Set .

(e) For Met , we define the adjunction $\mathbb{J} \dashv \mathbb{V}$ by

$$\mathbb{V}\langle X, d \rangle := X \quad \text{and} \quad \mathbb{J}X := \langle X, d_1 \rangle,$$

where d_1 is the discrete metric

$$d_1(x, y) := \begin{cases} 0 & \text{if } x = y, \\ 1 & \text{if } x \neq y. \end{cases}$$

The unit $\varepsilon : X \rightarrow X$ and the counit $\iota : \langle X, d_1 \rangle \rightarrow \langle X, d \rangle$ are given by the identity maps.

A limit of metric spaces is computed as in Set where the metric is given by

$$d(a, b) := \sup_i d(p_i(a), p_i(b)),$$

where p_i are the projections. The coproduct $\sum_i \langle X_i, d_i \rangle$ is the disjoint union with elements from different components at distance 1 from each other. The proof of disjointness is again similar to that of Set . \lrcorner

For reference, let us collect a few useful standard facts about adjunctions.

Lemma 2.3. *Let $\mathbb{J} \dashv \mathbb{V}$ be an adjunction between Set and \mathcal{D} where \mathbb{V} is faithful, and let $\varepsilon : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \text{Id}$ be the unit and counit of the adjunction.*

(a) *The adjunction $\mathbb{J} \dashv \mathbb{V}$ maps*

$$\begin{aligned} f : A \rightarrow \mathbb{V}B & \quad \text{to} \quad \iota \circ \mathbb{J}f : \mathbb{J}A \rightarrow B, \\ g : \mathbb{J}A \rightarrow B & \quad \text{to} \quad \mathbb{V}g \circ \varepsilon : A \rightarrow \mathbb{V}B. \end{aligned}$$

(b) $\mathbb{V}\iota \circ \varepsilon = \text{id}$ and $\iota \circ \mathbb{J}\varepsilon = \text{id}$.

(c) $f : A \rightarrow B$ is a monomorphism if, and only if, $\mathbb{V}f : \mathbb{V}A \rightarrow \mathbb{V}B$ is injective.

(d) If $\mathbb{V}f : \mathbb{V}A \rightarrow \mathbb{V}B$ is surjective, then $f : A \rightarrow B$ is an epimorphism.

(e) All morphisms of the counit $\iota : \mathbb{J}\mathbb{V} \Rightarrow \text{Id}$ are surjective.

Proof. (a) follows by naturality of the bijection $\mathcal{D}(\mathbb{J}X, A) \cong \text{Set}(X, \mathbb{V}A)$.

(b) By definition, $\iota : \mathbb{J}\mathbb{V}A \rightarrow A$ is the image of $\text{id} : \mathbb{V}A \rightarrow \mathbb{V}$ under the adjunction. Hence, we must obtain id when mapping ι back. By (a), this backwards translation is equal to $\mathbb{V}\iota \circ \varepsilon$. The second equation follows in the same way.

(c) (\Leftarrow) Suppose that $\mathbb{V}f$ is injective and consider two morphisms $g, h : C \rightarrow A$ with $f \circ g = f \circ h$. Then $\mathbb{V}f \circ \mathbb{V}g = \mathbb{V}f \circ \mathbb{V}h$ and injectivity of $\mathbb{V}f$ implies that $\mathbb{V}g = \mathbb{V}h$. As \mathbb{V} is faithful, it follows that $g = h$.

(\Rightarrow) Suppose that $\mathbb{V}f$ is not injective. Then there exist two functions $s, t : 1 \rightarrow \mathbb{V}A$ with $s \neq t$ but $\mathbb{V}f \circ s = \mathbb{V}f \circ t$. Let $\varepsilon : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \text{Id}$ be the unit and counit of the adjunction, respectively, and let $\hat{s} := \iota \circ \mathbb{J}s$ and $\hat{t} := \iota \circ \mathbb{J}t$ be the morphisms $\mathbb{J}1 \rightarrow A$ corresponding to s and t under the adjunction. Then the morphism corresponding to $f \circ \hat{s}$ under the adjunction is

$$\mathbb{V}(f \circ \hat{s}) \circ \varepsilon = \mathbb{V}f \circ \mathbb{V}\iota \circ \mathbb{V}\mathbb{J}s \circ \varepsilon = \mathbb{V}f \circ \mathbb{V}\iota \circ \varepsilon \circ s = \mathbb{V}f \circ s,$$

and similarly for $f \circ \hat{t}$. Consequently,

$$\mathbb{V}f \circ s = \mathbb{V}f \circ t \quad \text{implies} \quad f \circ \hat{s} = f \circ \hat{t}.$$

But $s \neq t$ implies $\hat{s} \neq \hat{t}$. Thus, f is not a monomorphism.

(d) Suppose that the function $\mathbb{V}f$ is surjective and consider two morphisms $g, h : B \rightarrow C$ with $g \circ f = h \circ f$. Then $\mathbb{V}g \circ \mathbb{V}f = \mathbb{V}h \circ \mathbb{V}f$ and surjectivity of $\mathbb{V}f$ implies that $\mathbb{V}g = \mathbb{V}h$. As \mathbb{V} is faithful, it follows that $g = h$.

(e) By (b), $\mathbb{V}\iota$ has a right inverse. \square

Example. The converse of (e) does not hold. Let \mathcal{D} be the category of all Hausdorff spaces. Then \mathcal{D} is discrete. We claim that a continuous map

$e : \mathfrak{X} \rightarrow \mathfrak{Y}$ is an epimorphism if, and only if, $\text{rng } e$ is dense in \mathfrak{Y} . In particular, there exists non-surjective epimorphisms.

Let us prove the interesting direction. Suppose that $\text{rng } e$ is dense in \mathfrak{Y} and consider two different continuous maps $g, h : \mathfrak{Y} \rightarrow \mathfrak{Z}$. As \mathfrak{Y} is Hausdorff, the equaliser

$$E := \{ y \in Y \mid g(y) = h(y) \}$$

is closed in \mathfrak{Y} . Hence $U := Y \setminus E$ is open and non-empty. This implies that $U \cap \text{rng } e \neq \emptyset$. Consequently, $g \circ e \neq h \circ e$ and e is not an epimorphism. \dashv

Exercise 2.1. Prove that, in a discrete category \mathcal{D} , $\mathbb{J}X = \sum_{x \in X} \mathbb{J}1$ and $\mathbb{V}\mathbb{J}X = X \times \mathbb{V}\mathbb{J}1$, for every $X \in \text{Set}$. \dashv

Sorts

Discrete categories will be our generalisation of sets. But actually, we will frequently not work with ordinary sets, but with *sorted* ones. That is, we fix a set Ξ of sorts and we consider sets where each element has some sort $\xi \in \Xi$. Such a set can be formalised either (i) as a family $A = (A_\xi)_{\xi \in \Xi}$ of sets where A_ξ contains all the elements of sort ξ ; or (ii) as a set A together with a function $\sigma : A \rightarrow \Xi$ that maps each element of A to its sort. Both definitions are equivalent. We will adopt the first one as it is usually simpler to work with.

The reason for using several sorts is that, for certain kinds of languages K , we need different sorts of elements to build up the objects in K . For instance, for languages of infinite words, we need to distinguish between finite and infinite word. Similarly, for languages of trees, we need both ordinary trees and trees with one or several holes. Sorts help us to keep these types of elements apart. In our setting, we will therefore be working in the category \mathcal{D}^Ξ whose objects are families $A = (A_\xi)_{\xi \in \Xi}$ where each A_ξ is an object of \mathcal{D} . A morphism $f : A \rightarrow B$ between such families is then just a family $f = (f_\xi)_\xi$ of morphisms $f_\xi : A_\xi \rightarrow B_\xi$. Similarly, a functor $\mathbb{F} : \mathcal{C} \rightarrow \mathcal{D}^\Xi$ is given by a family $\mathbb{F} = (\mathbb{F}_\xi)_\xi$ of functors $\mathbb{F}_\xi : \mathcal{C} \rightarrow \mathcal{D}$. For a property P , we say that

$A \in \mathcal{D}^\Xi$ is *sort-wise* P if each set A_ξ has property P . In particular, *sort-wise finite* means that every A_ξ is finite.

From this point on, we will use the terms ‘set’ and ‘function’ as a shorthand for ‘object of \mathcal{D} ’ and ‘morphism of \mathcal{D} ’, or their many-sorted counterparts in \mathcal{D}^Ξ . If we mean any other kind of set or function, we will mention this explicitly. Furthermore, to keep notation simple, we will frequently not distinguish between an object $A \in \mathcal{D}$ and its underlying set $\mathbb{V}A$.

As mentioned above, we can identify a sorted object $A = (A_\xi)_{\xi \in \Xi} \in \mathcal{D}^\Xi$ with its coproduct $A = \sum_{\xi \in \Xi} A_\xi$. Using this point of view, a morphism $f : A \rightarrow B$ corresponds to a *sort-preserving* morphism between the corresponding coproducts. Let us prove that these two points of view are equivalent for the categories we are interested in. Formally, we have to show that \mathcal{D}^Ξ is equivalent to the arrow category $(\mathcal{D} \downarrow \Xi \cdot 1)$, where we have set $X \cdot A := \sum_{x \in X} A$, for $X \in \text{Set}$ and $A \in \mathcal{D}$.

Proposition 2.4. *Let \mathcal{D} be a discrete category. Then*

$$\mathcal{D}^\Xi \cong (\mathcal{D} \downarrow \Xi \cdot 1), \quad \text{for every set } \Xi.$$

Proof. We define a functor $\mathbb{F} : \mathcal{D}^\Xi \rightarrow (\mathcal{D} \downarrow \Xi \cdot 1)$ as follows. Given an object $(A_\xi)_{\xi \in \Xi} \in \mathcal{D}^\Xi$, we set

$$\mathbb{F}(A_\xi)_\xi := \sum_{\xi \in \Xi} \alpha_\xi : \sum_{\xi \in \Xi} A_\xi \rightarrow \Xi \cdot 1,$$

where $\alpha_\xi : A_\xi \rightarrow 1$ is the unique morphism into the terminal object. For a morphism $(f_\xi)_\xi : (A_\xi)_\xi \rightarrow (B_\xi)_\xi$, we set

$$\mathbb{F}(f_\xi)_\xi := \sum_{\xi} f_\xi : \sum_{\xi} A_\xi \rightarrow \sum_{\xi} B_\xi.$$

As 1 is terminal, we have $\beta_\xi \circ f_\xi = \alpha_\xi$ (where α_ξ and β_ξ are the components of, respectively, $\mathbb{F}(A_\xi)_\xi$ and $\mathbb{F}(B_\xi)_\xi$), and it follows that

$$\mathbb{F}(B_\xi)_\xi \circ \mathbb{F}(f_\xi)_\xi = \sum_{\xi} (\beta_\xi \circ f_\xi) = \sum_{\xi} \alpha_\xi = \mathbb{F}(A_\xi)_\xi.$$

Note that \mathbb{F} is faithful since $\mathbb{F}(f_\xi)_\xi = \mathbb{F}(g_\xi)_\xi$ implies that, for every $\zeta \in \Xi$,

$$f_\zeta = \left(\sum_\xi f_\xi \right) \circ i_\zeta = \mathbb{F}(f_\xi)_\xi \circ i_\zeta = \mathbb{F}(g_\xi)_\xi \circ i_\zeta = \left(\sum_\xi g_\xi \right) \circ i_\zeta = g_\zeta,$$

where $i_\zeta : A_\zeta \rightarrow \sum_\xi A_\xi$ is the corresponding inclusion morphism.

To see that \mathbb{F} is full, consider a morphism $f : \alpha \rightarrow \beta$ where $\alpha := \mathbb{F}(A_\xi)_\xi$ and $\beta := \mathbb{F}(B_\xi)_\xi$. As \mathcal{D} is disjunctive, it follows that $f = \sum_\xi f_\xi$ for (unique) $f_\xi : A_\xi \rightarrow B_\xi$. Hence, $f = \mathbb{F}(f_\xi)_\xi$.

It remains to show that every object $\alpha : A \rightarrow \Xi \cdot 1$ of $(\mathcal{D} \downarrow \Xi \cdot 1)$ is isomorphic to one of the form $\mathbb{F}(A_\xi)_\xi$. Hence, fix $\alpha : A \rightarrow \Xi \cdot 1$. As \mathcal{D} is disjunctive, we have $A = \sum_\xi A_\xi$ and $\alpha = \sum_\xi \alpha_\xi$, for suitable A_ξ and $\alpha_\xi : A_\xi \rightarrow 1$. It follows that $\alpha = \sum_\xi \alpha_\xi = \mathbb{F}(A_\xi)_\xi$. \square

Let us introduce the following operations to work with families in \mathcal{D}^Ξ . We will only be using them in a few places where the notation would otherwise become unmanageable.

Definition 2.5. Let \mathcal{C} be a category and $f : I \rightarrow K$ a function in Set . We define functors $\square_f : \mathcal{C}^K \rightarrow \mathcal{C}^I$ and $\Sigma_f, \Pi_f : \mathcal{C}^I \rightarrow \mathcal{C}^K$ by

$$\begin{aligned} \square_f(A_k)_{k \in K} &:= (A_{f(i)})_{i \in I}, \\ \Sigma_f(A_i)_{i \in I} &:= \left(\sum_{i \in f^{-1}(k)} A_i \right)_{k \in K}, \\ \Pi_f(A_i)_{i \in I} &:= \left(\prod_{i \in f^{-1}(k)} A_i \right)_{k \in K}, \end{aligned}$$

and similarly for morphisms. \lrcorner

We can also formulate these operations in the comma category $(\mathcal{C} \downarrow \Xi \cdot 1)$.

Exercise 2.2. Using the correspondence in Proposition 2.4, the functor \square_f induces a functor $(\mathcal{C} \downarrow K \cdot 1) \rightarrow (\mathcal{C} \downarrow I \cdot 1)$. Prove that this functor maps a morphism $\alpha : C \rightarrow \mathbb{J}K$ to its pullback along $\mathbb{J}f : \mathbb{J}I \rightarrow \mathbb{J}K$. \lrcorner

The relationship between the three functors is given via the following adjunctions.

Proposition 2.6. Let \mathcal{C} be a category with products and coproducts of size $|I|$ and let $f : I \rightarrow K$ be a function in Set .

$$\Sigma_f \dashv \square_f \dashv \Pi_f.$$

Proof. For the first adjunction, consider a morphism $\varphi : \Sigma_f A \rightarrow B$. We map it to the morphism $\psi : A \rightarrow \square_f B$ with components

$$\psi_i := \varphi \circ j_i, \quad \text{for } i \in I.$$

where $j_i : A_i \rightarrow \sum_{j \in f^{-1}(f(i))} A_j$ is the inclusion morphism. Conversely, we map $\psi : A \rightarrow \square_f B$ to the morphism $\varphi : \Sigma_f A \rightarrow B$ defined by

$$\varphi_k := \bigoplus_{i \in f^{-1}(k)} \psi_i, \quad \text{for } k \in K.$$

It is straightforward to check that these transformations are natural. Furthermore, they are inverses of each other since

$$\begin{aligned} \bigoplus_{i \in f^{-1}(k)} (\varphi \circ j_i) &= \varphi_k \circ \bigoplus_{i \in f^{-1}(k)} j_i = \varphi_k, \\ \left(\bigoplus_{i \in f^{-1}(k)} \psi_i \right) \circ j_i &= \psi_i. \end{aligned}$$

For the second adjunction, consider a morphism $\varphi : \square_f A \rightarrow B$. We map it to the morphism $\psi : A \rightarrow \Pi_f B$ with components

$$\psi_k := \langle \varphi_i \rangle_{i \in f^{-1}(k)}, \quad \text{for } k \in K.$$

Conversely, we map $\psi : A \rightarrow \Pi_f B$ to the morphism $\varphi : \square_f A \rightarrow B$ defined by

$$\varphi_i := p_i \circ \psi_{f(i)}, \quad \text{for } i \in I.$$

where $p_i : \prod_{j \in f^{-1}(f(i))} B_j \rightarrow B_i$ is the projection. It is again straightforward to check that these transformations are natural. Furthermore, they are inverses of each other since

$$\begin{aligned} p_i \circ \langle \varphi_j \rangle_{j \in f^{-1}(f(i))} &= \varphi_i, \\ \langle p_i \circ \psi \rangle_{i \in f^{-1}(k)} &= \langle p_i \rangle_{i \in f^{-1}(k)} \circ \psi = \psi_k. \end{aligned} \quad \square$$

Let us also mention the following commutation relation between them.

Exercise 2.3. Let $f : I \rightarrow K$ and $g : J \rightarrow K$ be functions and $u : L \rightarrow I$ and $v : L \rightarrow J$ their pullbacks. Prove that

$$\square_f \Sigma_g \cong \Sigma_u \square_v \quad \text{and} \quad \square_f \Pi_g \cong \Pi_u \square_v .$$

Finally, we can rephrase the definition of a disjunctive category in terms of the functor Σ_f as follows.

Lemma 2.7. *A category \mathcal{D} is disjunctive if, and only if, for every function $f : I \rightarrow J$ in Set , the following two conditions hold.*

- (i) *For every morphism $\varphi : A \rightarrow \Sigma_f B$ in \mathcal{D}^I , there exist $A_* \in \mathcal{D}^I$ and $\varphi_* : A_* \rightarrow B$ such that*

$$A = \Sigma_f A_* \quad \text{and} \quad \varphi = \Sigma_f \varphi_* .$$

- (ii) *$\Sigma_f A = \Sigma_f B$ and $\Sigma_f \varphi = \Sigma_f \psi$ implies $A = B$ and $\varphi = \psi$.*

The Power Operator

Having fixed a base category \mathcal{D} and a set Ξ of sorts, we also need to choose a notion of an ' A -labelled object with domain X ', for a set $X \in \text{Set}$ and an object $A \in \mathcal{D}$. This can be captured by the concept of a power operator, which we introduce next. Many categories \mathcal{D} have the property that, given an object $A \in \mathcal{D}$ and a set X , we can equip the set of all functions $X \rightarrow \mathbb{V}A$ with the structure of an \mathcal{D} -object. For instance, if A is a partial order, then so is A^X by setting

$$f \leq g \quad : \text{iff} \quad f(x) \leq g(x), \quad \text{for all } x .$$

Similarly, if A is a topological space, we can equip A^X with the product topology. Or, if A is an abelian group, we can define an addition on A^X by setting

$$(f + g)(x) := f(x) + g(x) .$$

It turns out that in the categories we are interested in, we can identify the set of functions $X \rightarrow \mathbb{V}A$ with a product $\prod_{x \in X} A$. Let us introduce the corresponding notation.

Definition 2.8. Let \mathcal{D} be a category.

- (a) For $A \in \mathcal{D}$ and $X \in \text{Set}$, we define the *power operator*

$$A^X := \prod_{x \in X} A \in \mathcal{D} .$$

Similarly, for $A \in \mathcal{D}^\Xi$ and $X \in \text{Set}^\Xi$, we set

$$A^X := \prod_{\xi \in \Xi} A_\xi^{X_\xi} \in \mathcal{D} .$$

- (b) We extend this operation to a functor as follows. For $f : A \rightarrow B$ and $g : X \rightarrow Y$, we define

$$f^X := \prod_{x \in X} f : A^X \rightarrow B^X \quad \text{and} \quad A^g := \langle p_{g(x)} \rangle_{x \in X} : A^Y \rightarrow A^X ,$$

where $p_y : \prod_{y \in Y} A \rightarrow A$ is the projection to the y -th component.

Let us collect a few properties of this operation. We start by showing that the functor $B^{(-)}$ is left adjoint to the hom-functor $\mathcal{D}(-, B)$.

Lemma 2.9. *Let \mathcal{D} be a category with products of size less than κ and let $\text{Set}_\kappa \subseteq \text{Set}$ be the full subcategory of all sets of size less than κ . Then*

$$\mathcal{D}(A, B^X) \cong \text{Set}_\kappa(X, \mathcal{D}(A, B)), \quad \text{for } A, B \in \mathcal{D} \text{ and } |X| < \kappa ,$$

and this isomorphism is natural in A , B , and X . In particular, the functor $B^{(-)} : \text{Set}_\kappa^{\text{op}} \rightarrow \mathcal{D}$ is a right adjoint to the hom-functor $\mathcal{D}(-, B) : \mathcal{D} \rightarrow \text{Set}_\kappa^{\text{op}}$.

Proof. We map a function $f : X \rightarrow \mathcal{D}(A, B)$ to the morphism

$$\varphi(f) := \langle f(x) \rangle_{x \in X} : A \rightarrow B^X ,$$

and, conversely, we map $g : A \rightarrow B^X$ to the function $\psi(g) : X \rightarrow \mathcal{D}(A, B)$ given by

$$\psi(g)(x) := p_x \circ g,$$

where $p_x : \prod_{z \in X} B \rightarrow B$ is the projection to the x -th component. Then

$$\begin{aligned} \varphi(\psi(g)) &= \langle p_x \circ g \rangle_x = g, \\ \psi(\varphi(f))(x) &= p_x \circ \langle f(z) \rangle_z = f(x). \end{aligned}$$

Hence, φ and ψ are inverses of each other. It therefore remains to prove naturality of φ . For $h : Y \rightarrow X$, we have

$$\begin{aligned} (\mathcal{D}(A, B^h) \circ \varphi)(f) &= \langle p_{h(y)} \rangle_{y \in Y} \circ \varphi(f) \\ &= \langle p_{h(y)} \rangle_{y \in Y} \circ \langle f(x) \rangle_{x \in X} \\ &= \langle f(h(y)) \rangle_{y \in Y} \\ &= \varphi(f \circ h) \\ &= (\text{Set}_\kappa^{\text{op}}(\mathcal{D}(A, B), h) \circ \varphi)(f). \end{aligned}$$

For $h : C \rightarrow A$, we have

$$\begin{aligned} (\mathcal{D}(h, B^X) \circ \varphi)(f) &= \mathcal{D}(h, B^X) \circ \varphi(f) \\ &= \mathcal{D}(h, B^X) \circ \langle f(x) \rangle_{x \in X} \\ &= \langle f(x) \circ h \rangle_{x \in X} \\ &= \varphi(\mathcal{D}(h, B)(f)) \\ &= (\text{Set}_\kappa^{\text{op}}(\mathcal{D}(h, B), X) \circ \varphi)(f). \end{aligned}$$

For $h : B \rightarrow C$, we have

$$\begin{aligned} (\mathcal{D}(A, h^X) \circ \varphi)(f) &= \mathcal{D}(A, h^X) \circ \varphi(f) \\ &= \mathcal{D}(A, h^X) \circ \langle f(x) \rangle_{x \in X} \\ &= \langle h \circ f(x) \rangle_{x \in X} \\ &= \varphi(\mathcal{D}(A, h)(f)) \\ &= (\text{Set}_\kappa^{\text{op}}(\mathcal{D}(A, h), X) \circ \varphi)(f). \end{aligned} \quad \square$$

The following consequences of the category \mathcal{D} being discrete simplify working with power operators.

Lemma 2.10. *Let $\mathbb{J} \dashv \mathbb{V}$ be an adjunction between Set and \mathcal{D} and let $*$:= $\mathbb{J}1$ where 1 is an 1-element set.*

- (a) $\mathbb{V} \cong \mathcal{D}(*, -)$.
- (b) $\mathbb{V}(A^X) \cong \text{Set}(X, \mathbb{V}A)$.

Proof. (a) For $A \in \mathcal{D}$, we have

$$\mathcal{D}(*, A) = \mathcal{D}(\mathbb{J}1, A) \cong \text{Set}(1, \mathbb{V}A) \cong \mathbb{V}A.$$

(b) Consider $X \in \text{Set}$ and $A \in \mathcal{D}$ such that A^X is defined. Setting $\kappa := |X|^+$, it follows by (a) and Lemma 2.9 that

$$\begin{aligned} \mathbb{V}(A^X) &\cong \mathcal{D}(*, A^X) \cong \text{Set}_\kappa(X, \mathcal{D}(*, A)) \\ &\cong \text{Set}_\kappa(X, \mathbb{V}A) = \text{Set}(X, \mathbb{V}A). \end{aligned} \quad \square$$

3. Polynomial Functors

In formal language theory one studies sets of labelled objects like words, trees, traces, pictures, (hyper-)graphs, and so on. To capture all these various settings we start by introducing an operation \mathbb{M} mapping a given set A of labels to the set $\mathbb{M}A$ of all A -labelled objects. A language in this context is then simply a subset $K \subseteq \mathbb{M}A$. For instance, for languages of finite words we can use $\mathbb{M}A := A^+$. Similarly, for languages of infinite words we use two sorts $\Xi = \{1, \omega\}$ where sort 1 represents the ‘finite’ elements and sort ω the ‘infinite’ ones. The operation \mathbb{M} then maps a Ξ -sorted set $A = \langle A_1, A_\omega \rangle$ to $\mathbb{M}A = \langle \mathbb{M}_1 A, \mathbb{M}_\omega A \rangle$ where

$$\mathbb{M}_1 A := A_1^+ \quad \text{and} \quad \mathbb{M}_\omega A := A_1^+ A_\omega \cup A_1^\omega.$$

(So a finite word is a finite sequence of finite elements, while an infinite word can either be a finite sequence of finite elements followed by a single infinite one, or an infinite sequence of finite elements.)

Furthermore, every function $f : A \rightarrow B$ induces an operation $\mathbb{M}f : \mathbb{M}A \rightarrow \mathbb{M}B$ which applies the function f to each label. This turns \mathbb{M} into a functor. In this section, we introduce the kind of functors $\mathbb{M} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ that can be interpreted as producing sets of labelled objects. The definition is based on the power operators $(-)^X$ we have introduced above. Note that we can interpret the elements in A^X as labellings of X positions by elements of A .

Definition 3.1. Let \mathcal{D} be a category and let Ξ be a set of sorts.

(a) A functor $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ is *polynomial* if it is of the form

$$\mathbb{F}X = (\mathbb{F}_\xi X)_{\xi \in \Xi} \quad \text{where} \quad \mathbb{F}_\xi X := \sum_{i \in I_\xi} X^{D_\xi^i},$$

for fixed sets $I_\xi \in \mathbf{Set}$ and $D_\xi^i \in \mathbf{Set}^\Xi$.

(b) The *arity* of a polynomial functor \mathbb{F} is the least infinite cardinal κ such that $|D_\xi^i| < \kappa$, for all $\xi \in \Xi$ and $i \in I_\xi$. \lrcorner

Example. (a) The word functor $\mathbb{F}X := X^+$ is a polynomial functor on \mathbf{Pos} since we can write

$$\mathbb{F}X = \sum_{n < \omega} X^{n+1}.$$

The ordering is defined componentwise:

$$u \leq v \quad \text{iff} \quad |u| = |v| \text{ and } u(i) \leq v(i) \text{ for all } i.$$

(b) Similarly, the functor for ω -words

$$\mathbb{F}\langle X, Y \rangle := \langle X^+, X^*Y + X^\omega \rangle$$

is a polynomial functor on \mathbf{Pos}^2 . \lrcorner

Our first aim is to show that we can indeed interpret $\mathbb{F}A$ as a set of A -labelled objects of some sort. First, note that we can recover the index set I from a polynomial functor $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ with

$$\mathbb{F}_\xi X = \sum_{i \in I_\xi} X^{D_\xi^i}, \quad \text{for } \xi \in \Xi,$$

as follows. Let 1 be the terminal object of \mathcal{D}^Ξ . Since right adjoints preserve limits, its image $\mathbb{V}1 = 1$ is terminal in \mathbf{Set}^Ξ . If \mathcal{D} is discrete, it follows by Lemma 2.10 (b) that

$$\begin{aligned} \mathbb{V}\mathbb{F}_\xi 1 &= \sum_{i \in I_\xi} \mathbb{V}(1^{D_\xi^i}) \\ &\cong \sum_{i \in I_\xi} \mathbf{Set}^\Xi(D_\xi^i, \mathbb{V}1) \\ &\cong \sum_{i \in I_\xi} \mathbf{Set}^\Xi(D_\xi^i, 1) = \sum_{i \in I_\xi} 1 \cong I_\xi. \end{aligned}$$

For $i \in I$, we use the notation $\text{dom}(i) := D_\xi^i$. Thus, we can write

$$\mathbb{F}_\xi X = \sum_{i \in I_\xi} X^{\text{dom}(i)}.$$

Furthermore, identifying a set S of elements of a given sort $\xi \in \Xi$ with the family $(S_\zeta)_\zeta \in \mathbf{Set}^\Xi$ where $S_\xi := S$ and $S_\zeta := \emptyset$, for $\zeta \neq \xi$, we can also write

$$\mathbb{F}X = \sum_{\xi \in \Xi} \sum_{i \in I_\xi} X^{\text{dom}(i)} = \sum_{i \in I} X^{\text{dom}(i)} \quad \text{where } I := (I_\xi)_\xi.$$

If \mathcal{D} is discrete, it follows by Lemma 2.10 (b) that

$$\mathbb{V}\mathbb{F}_\xi A = \sum_{i \in I_\xi} \mathbb{V}(A^{\text{dom}(i)}) \cong \sum_{i \in I_\xi} \mathbf{Set}^\Xi(\text{dom}(i), \mathbb{V}A).$$

Therefore, we can regard elements of $\mathbb{F}_\xi A$ as functions $s : \text{dom}(i) \rightarrow \mathbb{V}A$ (or, more precisely, as pairs $\langle i, s \rangle$, but we usually omit the index i). We write $\text{dom}(s) := \text{dom}(i)$ in this case. Thus, the elements of $\mathbb{F}A$ are functions $s : \text{dom}(s) \rightarrow A$, i.e., A -labelled objects.

Similarly, for a function $f : A \rightarrow B$ and an element $s \in \mathbb{V}\mathbb{F}A$, we obtain

$$\mathbb{V}\mathbb{F}f(s) = \mathbb{V}(f^{\text{dom}(s)})(s) = \mathbf{Set}(\text{dom}(s), \mathbb{V}f)(s) = \mathbb{V}f \circ s.$$

Hence, $\mathbb{F}f$ acts on $\mathbb{F}A$ as a relabelling $s \mapsto f \circ s$.

These remarks justify our intuition of $\mathbb{F}A$ as a set of labelled objects. Let us summarise them in the following result.

Lemma 3.2. *Let \mathcal{D} be a discrete category. For every polynomial functor $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ there exists a polynomial functor $\mathbb{F}^\circ : \mathbf{Set}^\Xi \rightarrow \mathbf{Set}^\Xi$ such that*

$$\mathbb{F}^\circ \circ \mathbb{V} = \mathbb{V} \circ \mathbb{F}.$$

Proof. Suppose that

$$\mathbb{F}X = \sum_{s \in I} X^{\text{dom}(s)}.$$

If we use the same expression

$$\mathbb{F}^\circ X = \sum_{s \in I} X^{\text{dom}(s)}$$

in the category \mathbf{Set}^Ξ , we obtain

$$\begin{aligned} \mathbb{V}\mathbb{F}X &= \mathbb{V} \sum_{s \in I} X^{\text{dom}(s)} \\ &= \sum_{s \in I} \mathbb{V}(X^{\text{dom}(s)}) = \sum_{s \in I} (\mathbb{V}X)^{\text{dom}(s)} = \mathbb{F}^\circ \mathbb{V}X, \end{aligned}$$

where the second step follows from the assumption that \mathbb{V} commutes with coproducts and the third one from the fact that right adjoints preserve limits. \square

There exists a more elegant and concise way to describe polynomial functors. We can encode $\mathbb{F} = \sum_{s \in I} (-)^{\text{dom}(s)}$ by the function $f : \sum_{s \in I} \text{dom}(s) \rightarrow I$ mapping each $v \in \text{dom}(s)$ to s . But note that this function f does not preserve sorts. Therefore, to fully specify \mathbb{F} we need two additional functions $\alpha : \sum_{s \in I} \text{dom}(s) \rightarrow \Xi$ and $\beta : I \rightarrow \Xi$ telling us which sorts the elements of the respective sets have. Conversely, given three functions

$$\Xi \xleftarrow{\alpha} D \xrightarrow{f} I \xrightarrow{\beta} \Xi$$

in \mathbf{Set} , we can define the polynomial functor $\mathbb{F} = \sum_{s \in I} (-)^{D_s}$ where $D_s := f^{-1}(s)$ and we regard I and D_s as sorted sets via α and β . A straightforward direct calculation shows that, using the functors \square_f , Σ_f , and Π_f from Proposition 2.6, we can write \mathbb{F} as the composition

$$\mathbb{F} = \Sigma_\beta \circ \Pi_f \circ \square_\alpha.$$

We will mostly not use this formalism, since the more concrete description of \mathbb{F} as coproduct of powers leads to more elementary proofs.

Lemma 3.3. *Every polynomial functor preserves monomorphisms, injective morphisms, and surjective morphisms.*

Proof. Let $\mathbb{F}X = \sum_{i \in I} X^{D_i}$ be a polynomial functor and $f : A \rightarrow B$ a morphism. First, suppose that $f : A \rightarrow B$ is injective. To show that so is $\mathbb{F}f$, consider two elements $\langle i, s \rangle, \langle j, t \rangle \in \mathbb{F}A$ with

$$\mathbb{F}f(\langle i, s \rangle) = \mathbb{F}f(\langle j, t \rangle).$$

This implies that $i = j$ and $f(s(v)) = f(t(v))$, for all $v \in \text{dom}(s)$. Since f is injective, it follows that $s(v) = t(v)$, for all v . Hence, $s = t$.

Next, suppose that $f : A \rightarrow B$ is surjective. To show that so is $\mathbb{F}f$, let $\langle i, t \rangle \in \mathbb{F}B$. As f is surjective, $\mathbb{V}f$ has a right inverse $g : \mathbb{V}B \rightarrow \mathbb{V}A$. Setting $s := g \circ t$, it follows that

$$\mathbb{F}f(\langle i, s \rangle) = \langle i, f \circ s \rangle = \langle i, f \circ g \circ t \rangle = \langle i, t \rangle.$$

Finally, preservation of monomorphisms follows from what we have already proved since a morphism is a monomorphism if, and only if, it is injective. \square

Remark. Every functor $\mathbb{F} : \mathbf{Set} \rightarrow \mathbf{Set}$ preserves epimorphisms since epimorphisms in \mathbf{Set} have right inverses. \lrcorner

Let us take a quick look at how the composition of two polynomial functors looks like with this notation.

Lemma 3.4. If $\mathbb{F} = \sum_{i \in I} (-)^{\text{dom}(i)}$ and $\mathbb{G} = \sum_{k \in K} (-)^{\text{dom}(k)}$ are polynomial functors so is $\mathbb{F} \circ \mathbb{G}$ and

$$\mathbb{F}\mathbb{G}X = \sum_{i \in I} \sum_{g: \text{dom}(i) \rightarrow K} X^{\sum_{v \in \text{dom}(i)} \text{dom}(g(v))}.$$

Proof. We have

$$\begin{aligned} \mathbb{F}\mathbb{G}X &= \sum_{i \in I} \prod_{v \in \text{dom}(i)} \mathbb{G}X \\ &= \sum_{i \in I} \prod_{v \in \text{dom}(i)} \sum_{t \in K} \prod_{w \in \text{dom}(t)} X \\ &= \sum_{i \in I} \sum_{g: \text{dom}(i) \rightarrow K} \prod_{v \in \text{dom}(i)} \prod_{w \in \text{dom}(g(v))} X \\ &= \sum_{i \in I} \sum_{g: \text{dom}(i) \rightarrow K} X^{\sum_{v \in \text{dom}(i)} \text{dom}(g(v))}. \quad \square \end{aligned}$$

Remark. We obtain the following concrete descriptions of the elements of $\mathbb{F}\mathbb{G}X$. The index set is

$$\{ \langle i, g \rangle \mid i \in I, g \in K^{\text{dom}(i)} \},$$

and the domains are

$$\begin{aligned} \text{dom}(\langle i, g \rangle) &= \sum_{v \in \text{dom}(i)} \text{dom}(g(v)) \\ &= \{ \langle v, u \rangle \mid v \in \text{dom}(i), u \in \text{dom}(g(v)) \}. \end{aligned} \quad \lrcorner$$

Next, we turn to the study of natural transformations between polynomial functors.

Definition 3.5. A *morphism* of polynomial functors is a natural transformation $\alpha: \mathbb{F} \Rightarrow \mathbb{G}$. For a category \mathcal{C} , we denote the category of all polynomial functors $\mathcal{C} \rightarrow \mathcal{C}$ and their morphisms by $\text{Poly}(\mathcal{C})$. \lrcorner

We would like to obtain a more concrete description of such morphisms. To do so, we compare $\text{Poly}(\mathcal{C})$ with the following category.

Definition 3.6. For a category \mathcal{C} , we denote by $\Pi(\mathcal{C})$ the category whose objects are families $(A_i)_{i \in I}$ of objects $A_i \in \mathcal{C}$ (for varying I). A morphism $(A_i)_{i \in I} \rightarrow (B_j)_{j \in J}$ is a pair $\langle f, (\varphi_j)_{j \in J} \rangle$ consisting of a function $f: J \rightarrow I$ and a family of morphisms $\varphi_j: A_{f(j)} \rightarrow B_j$, for $j \in J$. The composition of two such morphisms is defined by

$$\langle g, (\psi_j)_j \rangle \circ \langle f, (\varphi_i)_i \rangle = \langle g \circ f, (\varphi_i \circ \psi_{f(i)})_i \rangle. \quad \lrcorner$$

There are obvious maps between $\Pi(\text{Set})$ and $\text{Poly}(\mathcal{C})$ that associate a family $(A_i)_{i \in I}$ of sets with the functor $\sum_{i \in I} X^{A_i}$ and vice versa. Let us show that this correspondence also preserves morphisms.

Proposition 3.7. Let \mathcal{D} be a discrete category. The above correspondence induces a functor

$$\mathbb{I}: \Pi(\text{Set})^{\text{op}} \rightarrow \text{Poly}(\mathcal{D})$$

that is faithful and bijective on objects.

Proof. Let $A = (A_i)_{i \in I}$ and $B = (B_j)_{j \in J}$ be objects of $\Pi(\text{Set})$ and let $\mathbb{F}_A := \mathbb{I}A$ and $\mathbb{F}_B := \mathbb{I}B$ be the associated polynomial functors. We map a morphism $\langle f, (\varphi_i)_i \rangle: B \rightarrow A$ to the natural transformation $\tau: \mathbb{F}_A \Rightarrow \mathbb{F}_B$ defined by

$$\tau(\langle i, s \rangle) := \langle f(i), s \circ \varphi_i \rangle, \quad \text{for } i \in I \text{ and } s \in \mathbb{V}X^{A_i} \cong \text{Set}(A_i, \mathbb{V}X),$$

where we have written $s \circ \varphi_i$ instead of the formally correct $X^{\varphi_i}(s)$. To see that τ is natural, consider a function $g: X \rightarrow Y$. For $\langle i, s \rangle \in \mathbb{F}_A X$, it follows that

$$\begin{aligned} (\tau \circ \mathbb{F}_A g)(\langle i, s \rangle) &= \tau(\langle i, g \circ s \rangle) \\ &= \langle f(i), g \circ s \circ \varphi_i \rangle \\ &= \mathbb{F}_B g(\langle f(i), s \circ \varphi_i \rangle) = (\mathbb{F}_B g \circ \tau)(\langle i, s \rangle), \end{aligned}$$

as desired.

To see that \mathbb{I} is functorial, consider two morphisms $\langle f, (\varphi_i)_i \rangle : B \rightarrow A$ and $\langle g, (\psi_j)_j \rangle : C \rightarrow B$. Then

$$\langle g, (\psi_j)_j \rangle \circ \langle f, (\varphi_i)_i \rangle = \langle g \circ f, (\varphi_i \circ \psi_{f(i)})_i \rangle.$$

Let $\sigma : \mathbb{F}_A \Rightarrow \mathbb{F}_B$, $\tau : \mathbb{F}_B \Rightarrow \mathbb{F}_C$, and $\rho : \mathbb{F}_A \Rightarrow \mathbb{F}_C$ be the corresponding natural transformations. Then

$$\begin{aligned} (\tau \circ \sigma)(\langle i, s \rangle) &= \tau(\langle f(i), s \circ \varphi_i \rangle) \\ &= \langle g(f(i)), s \circ \varphi_i \circ \psi_{f(i)} \rangle = \rho(\langle i, s \rangle). \end{aligned}$$

Clearly, the function $A \mapsto \mathbb{F}_A$ induced by \mathbb{I} on objects is bijective. Hence, it remains to prove that \mathbb{I} is faithful. Suppose that there are morphisms $\langle f, (\varphi_i)_i \rangle : B \rightarrow A$ and $\langle g, (\psi_i)_i \rangle : B \rightarrow A$ with the same image. That is,

$$\langle f(i), s \circ \varphi_i \rangle = \langle g(i), s \circ \psi_i \rangle, \quad \text{for all } i \text{ and } s.$$

Equality in the first component implies that $f = g$, while equality in the second one implies that $\varphi_i = \psi_i$ (by choosing suitable values for s). \square

In general, the embedding functor $\mathbb{I} : \Pi(\text{Set})^{\text{op}} \rightarrow \text{Poly}(\mathcal{D})$ is not full. Hence, there can be natural transformations that do not correspond to a morphism of $\Pi(\text{Set})$. To get an idea of how these additional transformations look like, we derive a translation in the other direction. We start with power operators.

Lemma 3.8. *Let \mathcal{D} be a discrete category. Every natural transformation $\alpha : (-)^{\mathcal{D}} \Rightarrow (-)^{\mathcal{E}}$ is of the form*

$$\alpha(s) = \mathbb{V}(\iota \circ \mathbb{J}s) \circ \alpha_o, \quad \text{for } s : D \rightarrow \mathbb{V}X,$$

for some $\alpha_o : E \rightarrow \mathbb{V}\mathbb{J}D$, where $\iota : \mathbb{J}\mathbb{V}A \rightarrow A$ is the counit of the adjunction.

Proof. By Lemma 2.10, we have $\mathbb{V}(A^{\mathcal{D}}) \cong \text{Set}(D, \mathbb{V}A)$. Consequently, the morphism

$$\mathbb{V}\alpha_{\mathbb{J}D} : \mathbb{V}((\mathbb{J}D)^{\mathcal{D}}) \rightarrow \mathbb{V}((\mathbb{J}D)^{\mathcal{E}})$$

induces a function

$$\hat{\alpha}_{\mathbb{J}D} : \text{Set}(D, \mathbb{V}\mathbb{J}D) \rightarrow \text{Set}(E, \mathbb{V}\mathbb{J}D).$$

We set

$$\alpha_o := \hat{\alpha}_{\mathbb{J}D}(\varepsilon) \in \text{Set}(E, \mathbb{V}\mathbb{J}D),$$

where $\varepsilon : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ is the unit of the adjunction.

To show that α is of the required form, let $A \in \mathcal{D}$ and $s \in \text{Set}(D, \mathbb{V}A) \cong \mathbb{V}A^{\mathcal{D}}$. Then $\hat{s} := \iota \circ \mathbb{J}s : \mathcal{D}(\mathbb{J}D, A)$ is the morphism corresponding to s via the adjunction. Consequently, $s = \mathbb{V}\hat{s} \circ \varepsilon$ and

$$\mathbb{V}\hat{s}^{\mathcal{D}}(\varepsilon)(d) = \mathbb{V}\hat{s}(\varepsilon(d)) = s(d), \quad \text{for all } d \in D.$$

This implies that

$$\begin{aligned} \mathbb{V}\alpha_A(s) &= \mathbb{V}\alpha_A(\mathbb{V}\hat{s}^{\mathcal{D}}(\varepsilon)) \\ &= \mathbb{V}\hat{s}^{\mathcal{E}}(\mathbb{V}\alpha_{\mathbb{J}D}(\varepsilon)) \\ &= \mathbb{V}\hat{s}^{\mathcal{E}}(\alpha_o) \\ &= \mathbb{V}\hat{s} \circ \alpha_o = \mathbb{V}(\iota \circ \mathbb{J}s) \circ \alpha_o. \end{aligned} \quad \square$$

Examples. The reverse is not true in general. For categories that are sufficiently close to Set , we will provide a precise characterisation in Theorem 3.11 below. But other categories are more complicated. Here are two examples showing that, in $\mathfrak{G}\text{-Set}$, the behaviour depends on the group \mathfrak{G} in question.

(a) Let \mathfrak{A} be an abelian group and D, E two sets. The natural transformations $\tau : (-)^{\mathcal{D}} \Rightarrow (-)^{\mathcal{E}}$ on $\mathfrak{A}\text{-Set}$ are precisely the functions of the form

$$\tau(s)(v) = \beta(v) \cdot s(\alpha(v)), \quad \text{for } s \in X^{\mathcal{D}} \text{ and } v \in E,$$

where $\alpha : E \rightarrow D$ and $\beta : E \rightarrow \mathfrak{A}$ are arbitrary functions. Consequently, there exists a bijection between natural transformations $(-)^{\mathcal{D}} \Rightarrow (-)^{\mathcal{E}}$ and functions $E \rightarrow \mathfrak{A} \times D = \mathbb{V}\mathbb{J}D$.

(b) Let \mathfrak{G} be a non-abelian group, fix two elements $g, h \in G$ with $gh \neq hg$. The function $\tau : G^D \rightarrow G^D$ defined by

$$\tau(s)(v) = h \cdot s(v), \quad \text{for } s \in X^D \text{ and } v \in D,$$

is not a morphism in $\mathfrak{G}\text{-Set}$ since, for $s(v) = e$, we have

$$(g \cdot \tau(s))(v) = gh \cdot e \neq hg \cdot e = \tau(g \cdot s)(v).$$

For arbitrary polynomial functors, we obtain the following description.

Proposition 3.9. *Let \mathcal{D} be a discrete category, and let $\mathbb{F} = \sum_{i \in I} X^{D_i}$ and $\mathbb{G} = \sum_{j \in J} X^{E_j}$ be polynomial functors. For every natural transformation $\tau : \mathbb{F} \Rightarrow \mathbb{G}$, there exists a morphism $\langle f, (\varphi_i)_i \rangle : (E_j)_j \rightarrow (\mathbb{V}\mathbb{J}D_i)_i$ of $\Pi(\text{Set})$ such that*

$$\tau_A(\langle i, s \rangle) = \langle f(i), \mathbb{V}(\iota \circ \mathbb{J}s) \circ \varphi_i \rangle,$$

where $\iota : \mathbb{V}A \rightarrow A$ is the counit of the adjunction.

Proof. Let $\tau : \mathbb{F} \Rightarrow \mathbb{G}$ be a natural transformation. We start by recovering the function $f : I \rightarrow J$. Let 1 be the terminal object of \mathcal{D}^\pm . As we have shown above, we can identify the index sets I and J with the sets $\mathbb{V}\mathbb{F}1$ and $\mathbb{V}\mathbb{G}1$, respectively. In particular, $\mathbb{V}\tau_1 : \mathbb{V}\mathbb{F}1 \rightarrow \mathbb{V}\mathbb{G}1$ induces a function $f : I \rightarrow J$. Given some object A , let $u : A \rightarrow 1$ be the unique morphism. For $\langle i, s \rangle \in \mathbb{V}\mathbb{F}A$ it follows that

$$\begin{aligned} \mathbb{V}\mathbb{G}u(\mathbb{V}\tau_A(\langle i, s \rangle)) &= \mathbb{V}\tau_1(\mathbb{V}\mathbb{F}u(\langle i, s \rangle)) \\ &= \mathbb{V}\tau_1(\langle i, * \rangle) = \langle f(i), * \rangle, \end{aligned}$$

where $*$ denotes the unique elements of $\mathbb{V}(1^{D_i})$ and of $\mathbb{V}(1^{E_{f(i)}})$. This implies that

$$\tau_A(\langle i, s \rangle) = \langle f(i), t \rangle, \quad \text{for some } t : E_{f(i)} \rightarrow A.$$

It thus remains to construct the functions $f'_i : E_{f(i)} \rightarrow \mathbb{V}\mathbb{J}D_i$. We have just shown that $\tau : \mathbb{F} \Rightarrow \mathbb{G}$ induces a natural transformation

$$\sigma_i : (-)^{D_i} \Rightarrow (-)^{E_{f(i)}} : s \mapsto t.$$

By Lemma 3.8, this transformation is of the form

$$s \mapsto \mathbb{V}(\iota \circ \mathbb{J}s) \circ \varphi_i, \quad \text{for some } \varphi_i : E_{f(i)} \rightarrow \mathbb{V}\mathbb{J}D_i. \quad \square$$

Definition 3.10. Let $\tau : \mathbb{F} \Rightarrow \mathbb{G}$ be a natural transformation between polynomial functors and let $\langle f, (\varphi_i)_i \rangle : (E_j)_j \rightarrow (\mathbb{V}\mathbb{J}D_i)_i$ be the corresponding morphism of $\Pi(\text{Set})$. We call f the *shape map* of τ and φ_i its *origin maps*.

Remark. The above proposition yields a function

$$\text{Poly}(\mathcal{C})(\mathbb{I}A, \mathbb{I}B) \rightarrow \Pi(\text{Set})^{\text{op}}(\mathbb{V}\mathbb{J}A, B).$$

This function is injective, but usually not surjective. In particular, this means that the functor $\mathbb{I}A \mapsto \mathbb{V}\mathbb{J}A$ does not form a left adjoint of \mathbb{I} .

Again, depending on the category, not all transformations of the above form are natural. The following set of additional assumptions is sufficient to show that they are.

Theorem 3.11. *Suppose that \mathcal{D} is discrete, that $\mathbb{V}\mathbb{J} = \text{Id}$, and that the counit $\iota : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ is the identity. Then $\mathbb{I} : \Pi(\text{Set})^{\text{op}} \cong \text{Poly}(\mathcal{C})$ is an isomorphism.*

Proof. It remains to prove that the functor \mathbb{I} is full. Hence, let $\tau : \mathbb{I}A \rightarrow \mathbb{I}B$ be a natural transformation. By Proposition 3.9, τ is of the form

$$\tau(\langle i, s \rangle) = \langle f(i), \mathbb{V}(\iota \circ \mathbb{J}s) \circ \varphi_i \rangle,$$

for some morphism $\langle f, (\varphi_i)_i \rangle$ of $\Pi(\text{Set})$. By our assumptions, it follows that

$$\tau(\langle i, s \rangle) = \langle f(i), s \circ \varphi_i \rangle = \mathbb{I}\langle f, (\varphi_i)_i \rangle. \quad \square$$

Categories satisfying the conditions of this theorem include Set , Pos , Top , and Met , but not $\mathfrak{G}\text{-Set}$.

It is easy to find examples showing that polynomial functors do not preserve colimits or products. For instance, for the word functor X^+ , we have

$$(X + Y)^+ \neq X^+ + Y^+ \quad \text{and} \quad (X \times Y)^+ \neq X^+ \times Y^+.$$

But one can show that, for discrete categories, polynomial functors do preserve connected limits, i.e., limits of diagrams whose index category is connected (as a directed graph). We start with the operation Σ_f from Proposition 2.6.

Proposition 3.12. *Let \mathcal{D} be a disjunctive category and $f : I \rightarrow J$ a function in Set . The functor $\Sigma_f : \mathcal{D}^I \rightarrow \mathcal{D}^J$ preserves connected limits.*

Proof. Let $D : \mathcal{K} \rightarrow \mathcal{D}^I$ be a connected diagram with limit A , and let $(\lambda_k)_k$ be the corresponding limiting cone. To prove that $(\Sigma_f \lambda_k)_k$ is limiting for $\Sigma_f \circ D$, we consider an arbitrary cone $(\mu_k)_k$ from $B \in \mathcal{D}^J$ to $\Sigma_f \circ D$.

For every $n \in \mathcal{K}$, we can use Lemma 2.7 to find objects $C_k \in \mathcal{D}^I$ and morphisms $\psi_k : C_k \rightarrow D(k)$ such that

$$B = \Sigma_f C_k \quad \text{and} \quad \mu_k = \Sigma_f \psi_k.$$

For every morphism $h : k \rightarrow l$ of \mathcal{D}^I , it follows that

$$\Sigma_f(Dh \circ \psi_k) = \Sigma_f Dh \circ \mu_k = \mu_l = \Sigma_f \psi_l.$$

By Lemma 2.7 it follows that

$$C_k = C_l \quad \text{and} \quad Dh \circ \psi_k^n = \psi_l^n.$$

Since \mathcal{K} is connected, it follows that $C_k = C_l$ for all $k, l \in \mathcal{K}$. Let C be this object. Then $(\psi_k)_k$ is a cone from C to D and we can find a unique morphism $\sigma : C \rightarrow A$ such that

$$\lambda_k \circ \sigma = \psi_k, \quad \text{for all } k.$$

Thus,

$$\Sigma_f \lambda_k \circ \Sigma_f \sigma = \Sigma_f \psi_k.$$

To show that $\Sigma_f \sigma$ is the unique morphism with this property, suppose that

$$\Sigma_f \lambda_k \circ \tau = \Sigma_f \psi_k.$$

By Lemma 2.7, it follows that $\tau = \Sigma_f \tau_*$, for some τ_* . As Σ_f is faithful, this implies that $\lambda_k \circ \tau_* = \psi_k$. Hence, we have $\tau_* = \sigma$, by uniqueness of σ . \square

Corollary 3.13. *Let \mathcal{D} be a disjunctive category. Every polynomial functor $\mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ preserves connected limits.*

Proof. Every polynomial functor can be written as composition $\Sigma_\beta \circ \Pi_f \circ \square_\alpha$, for functions α, β, f . The functor $\Pi_f \circ \square_\alpha$ is nothing but a product and, therefore, commutes with all limits. Moreover, we have shown above that Σ_β preserves connected limits. \square

Remark. One can show that, for $\mathcal{D} = \text{Set}$, the converse holds as well: a functor $\mathbb{F} : \text{Set}^\Xi \rightarrow \text{Set}^\Xi$ is polynomial if, and only if, it preserves connected limits. \lrcorner

Remark. We will prove in Corollary II.2.10 below that, under certain additional assumptions on the category, every polynomial functor preserves κ -filtered colimits. \lrcorner

Exercise 3.1. We consider limits and colimits in the category Poly of polynomial functors and natural transformations.

(a) Given a family $(\mathbb{F}_s)_{s \in S}$ of polynomial functors, show that

$$\left(\sum_{s \in S} \mathbb{F}_s \right) A = \sum_{s \in S} \mathbb{F}_s A, \quad \text{for all } A \in \mathcal{D}.$$

(b) Given a diagram $\mathbb{F} : \mathcal{S} \rightarrow \text{Poly}(\mathcal{D}, \mathcal{D})$ of polynomial functors, show that

$$\left(\lim_s \mathbb{F}(s) \right) A = \lim_s \left(\mathbb{F}(s) A \right), \quad \text{for all } A \in \mathcal{D}.$$

4. Monads

To study languages of A -labelled objects we can now work with sets of the form $\mathbb{M}A$ for some polynomial functor \mathbb{M} . Of course, just having a set is not sufficient to build a meaningful theory. Usually, the objects in a formal language are subject to various composition operations, like concatenation of words, substitution for terms, etc.. To capture such operations we will


therefore introduce two more ingredients. Firstly, the concatenation operation in question is often of the form $\text{flat} : \mathbb{M}A \rightarrow \mathbb{M}A$, that is, it takes an $\mathbb{M}A$ -labelled object $s \in \mathbb{M}A$ and assembles the appearing labels into a single large object. We call $\text{flat}(s)$ the *flattening* of s . Secondly, there is usually a *singleton operation* $\text{sing} : A \rightarrow \mathbb{M}A$ that takes a label $a \in A$ and produces an object with a single position which is labelled by a . For instance, in the case of words $\text{flat} : (A^*)^* \rightarrow A^*$ is simply the concatenation operation and $\text{sing} : A \rightarrow A^*$ produces 1-letter words.

$$\begin{aligned} \text{flat}(\langle w_0, \dots, w_{n-1} \rangle) &:= w_0 \dots w_{n-1}, & \text{for } w_0, \dots, w_{n-1} \in A^*, \\ \text{sing}(a) &:= \langle a \rangle, & \text{for } a \in A. \end{aligned}$$

Usually, the flattening operation is associative, which makes the functor \mathbb{M} into a *monad*.

Definition 4.1. A *monad* on a category \mathcal{C} consists of a functor $\mathbb{M} : \mathcal{C} \rightarrow \mathcal{C}$ that is equipped with two natural transformations: a *multiplication* $\text{flat} : \mathbb{M} \circ \mathbb{M} \Rightarrow \mathbb{M}$ and a *unit morphism* $\text{sing} : \text{Id} \Rightarrow \mathbb{M}$ (where Id is the identity functor), that satisfying the following equations.

$$\text{flat} \circ \text{sing} = \text{id}, \quad \text{flat} \circ \mathbb{M}\text{sing} = \text{id}, \quad \text{flat} \circ \text{flat} = \text{flat} \circ \mathbb{M}\text{flat}.$$



Examples. (a) The word functor $\mathbb{M}X = X^*$ forms a monad where multiplication is concatenation and the unit morphism is the singleton map $a \mapsto \langle a \rangle$.

(b) The covariant power-set functor on Set is defined by $\mathbb{P}X := \wp(X)$ and $\mathbb{P}f(X) := f[X]$. It forms a monad where multiplication is the union

operation $\cup : \wp(\wp(X)) \rightarrow \wp(X)$ and the unit morphism is the singleton map $x \mapsto \{x\}$.

(c) Similarly, we can define an upper-set functor on Pos by

$$\begin{aligned} \mathbb{U}X &:= \{ I \subseteq X \mid I \text{ upwards closed} \}, \\ \mathbb{U}f(X) &:= \{ b \mid b \geq f(a), \text{ for some } a \in X \}. \end{aligned}$$

It also forms a monad with the union operation as multiplication and the principle filter map $a \mapsto \{ b \mid b \geq a \}$ as unit morphism (see Section ??).

Examples. (a) Fix an infinite cardinal κ . The *chain monad* is defined as follows. For $A \in \text{Set}$, we denote by $\mathbb{C}_\kappa A$ the set of all A -labelled linear orders of size less than κ (up to isomorphism). The product $\text{flat} : \mathbb{C}_\kappa \mathbb{C}_\kappa A \rightarrow \mathbb{C}_\kappa A$ maps a chain $(\mathcal{Q}_i)_{i \in I}$ of linear orders to their ordered sum $\sum_{i \in I} \mathcal{Q}_i$. As $\kappa^2 = \kappa$, the resulting order still has size less than κ . Let us check one of the monad laws.

$$\text{flat} \circ \mathbb{C}_\kappa \text{flat} = \text{flat} \circ \text{flat}.$$

Note that the ordered sum $\sum_{i \in I} A_i$ consists of all pairs $\langle i, a \rangle$ with $i \in I$ and $a \in A_i$, ordered in the appropriate way. Consequently, the elements of $\text{flat}(\mathbb{C}_\kappa \text{flat}((\mathcal{Q}_j^i)_{j \in J_i})_{i \in I})$ take the form

$$\langle i, \langle j, a \rangle \rangle, \quad \text{with } i \in I, j \in J_i, \text{ and } a \in A_j^i,$$

while those of $\text{flat}(\text{flat}((\mathcal{Q}_j^i)_{j \in J_i})_{i \in I})$ have the form

$$\langle \langle i, j \rangle, a \rangle, \quad \text{with } i \in I, j \in J_i, \text{ and } a \in A_j^i.$$

The natural bijection between these two sets of elements forms an isomorphism showing that

$$\text{flat}(\mathbb{C}_\kappa \text{flat}((\mathcal{Q}_j^i)_{j \in J_i})_{i \in I}) = \text{flat}(\text{flat}((\mathcal{Q}_j^i)_{j \in J_i})_{i \in I}).$$

(b) Every subclass of linear orders that is closed under flat induces a corresponding submonad of \mathbb{C}_κ . For instance, there are monads for (i) well-founded linear orders and (ii) scattered linear orders.

Example. For a semiring \mathfrak{S} , we consider the functor $\text{Lin}_{\mathfrak{S}}$ mapping a set X to the \mathfrak{S} -module $\mathfrak{S}[X]$ generated by X . That is, $\text{Lin}_{\mathfrak{S}}X$ consists of all finite linear combinations

$$s_0x_0 + \cdots + s_{n-1}x_{n-1}, \quad \text{with } s_0, \dots, s_{n-1} \in S \text{ and } x_0, \dots, x_{n-1} \in X.$$

The corresponding flattening operation maps a nested linear combination

$$s(r_0x_0 + \cdots + r_{m-1}x_{m-1}) + \cdots + t(p_0y_0 + \cdots + p_{n-1}y_{n-1})$$

to

$$sr_0x_0 + \cdots + sr_{m-1}x_{m-1} + \cdots + tp_0y_0 + \cdots + tp_{n-1}y_{n-1}.$$

Example. For $A \in \text{Set}$, set $\mathbb{M}A := A^\omega$. We can turn \mathbb{M} into a monad with unit

$$\text{sing}(a) = \langle a, a, a, \dots \rangle$$

and multiplication

$$\text{flat}((a_i^k)_{i,k}) = \langle a_0^0, a_1^1, a_2^2, \dots \rangle.$$

A subset $K \subseteq A^\omega$ is recognised by a morphism into a finite \mathbb{M} -algebra if, and only if, it is clopen.

The prototypical example of a monad is the term monad for an algebraic signature.

Definition 4.2. Let Ξ be a set of sorts.

(a) A Ξ -sorted algebraic signature Σ is a set of function symbols, each of which has an associated type of the form $\tilde{\eta} \rightarrow \xi$, where $\xi \in \Xi$ and $\tilde{\eta} \in \Xi^\kappa$, for some cardinal κ . We call κ the *arity* of f . The *arity* of a signature is the least infinite cardinal κ such that every function symbol $f \in \Sigma$ has an arity less than κ . A signature of arity \aleph_0 is also called *finitary*.

(b) Let Σ be a Ξ -sorted signature and $X \in \text{Set}^\Xi$ a set of variables. The set $\text{Term}[\Sigma, X] \in \text{Set}^\Xi$ of Σ -terms is defined inductively as follows.

- Every variable $x \in X_\xi$ is a Σ -term of sort ξ .
- If $f \in \Sigma$ is a function symbol of type $\tilde{\eta} \rightarrow \xi$ and, for each index i , $t_i \in \text{Term}_{\eta_i}[\Sigma, X]$ is a Σ -term of sort η_i , then $f(\tilde{t})$ is a Σ -term of sort ξ .

For a function $f : X \rightarrow Y$, we denote by

$$\text{Term}[\Sigma, f] : \text{Term}[\Sigma, X] \rightarrow \text{Term}[\Sigma, Y]$$

the function replacing every variable $x \in X$ in a given term $t \in \text{Term}[\Sigma, X]$ by the variable $f(x) \in Y$.

Remark. The above recursive definition of terms is equivalent to saying that a term is a well-founded tree whose vertices are labelled by elements of $\Sigma + X$ and such that the number of successors of a vertex match the arity of its label.

Examples. (a) The (one-sorted) signature of semigroups is $\Sigma := \{\cdot\}$. Some terms are

$$x, \quad x \cdot y, \quad (x \cdot y) \cdot x, \quad x \cdot (y \cdot x),$$

where x and y are variables. Note that the last two terms above are considered to be different, as there is no built in assumption of associativity.

(b) The (one-sorted) signature of rings is $\Sigma := \{+, -, \cdot, 0, 1\}$. Some terms are

$$(x + 1) + y, \quad x + (1 + y), \quad (1 + 1) \cdot x, \quad (0 \cdot x) \cdot y.$$

(c) The $\{1, \omega\}$ -sorted signature of ω -semigroups is $\Sigma := \{\cdot, \times, \pi\}$ where the finite product \cdot has the type $\Pi \rightarrow 1$, the mixed one \times has the type $1\omega \rightarrow \omega$, and the infinite product π the type $\Pi \Pi \dots \rightarrow \omega$. An example of a term is

$$x \times \pi(x, x \cdot x, (x \cdot x) \cdot x, \dots).$$

The functor $\text{Term}[\Sigma, -]$ is clearly polynomial. We can write

$$\text{Term}[\Sigma, X] \cong \sum_{t \in T} X^{D_t},$$

where $T := \text{Term}[\Sigma, \{\square\}]$ and D_t is the set of all positions of the term t that are labelled by the variable \square . We can turn $\text{Term}[\Sigma, -]$ into a monad as follows.

Definition 4.3. Let Σ be a signature.

(a) For $t \in \text{Term}[\Sigma, \text{Term}[\Sigma, X]]$, we denote by $\text{flat}(t)$ the substitution operation. Formally, we define $\text{flat}(t)$ by induction on t as follows.

$$\begin{aligned} \text{flat}(x) &:= x, & \text{for } x \in X, \\ \text{and } \text{flat}(f(\vec{s})) &:= f'((\text{flat}(s_i))_i), & \text{for } f \in \text{Term}[\Sigma, X], \end{aligned}$$

where the function f' is defined by induction on f as follows.

$$\begin{aligned} x'(\vec{s}) &:= s_x, & \text{for } x \in X, \\ (g(\vec{u}))'(\vec{s}) &:= g((u'_i(\vec{s}))_i), & \text{for } g \in \Sigma. \end{aligned}$$

(b) For $x \in X$, we set

$$\text{sing}(x) := x \in \text{Term}[\Sigma, X].$$

Proposition 4.4. Let Σ be a Ξ -sorted algebraic signature. Then

$$\langle \mathbb{T}[\Sigma, -], \text{flat}, \text{sing} \rangle$$

forms a polynomial monad on Set^Ξ . The arity of $\text{Term}[\Sigma, -]$ as a polynomial functor coincides with the arity of Σ .

Remark. Let $(\mathbb{M}, \mu, \varepsilon)$ be a monad where $\mathbb{M} = \sum_{i \in I} (-)^{D_i}$ is polynomial. We have seen in Lemma 3.4 that

$$\mathbb{M} \circ \mathbb{M} = \sum_{k \in K} (-)^{E_k},$$

where

$$K := \{ \langle i, \sigma \rangle \mid i \in I, \sigma \in I^{D_i} \} \quad \text{and} \quad E_{\langle i, \sigma \rangle} := \sum_{v \in D_i} D_{\sigma(v)}.$$

Applying Proposition 3.9 to the natural transformations $\mu : \mathbb{M}\mathbb{M} \Rightarrow \mathbb{M}$ and $\varepsilon : \text{Id} \Rightarrow \mathbb{M}$, we obtain morphisms

$$\begin{aligned} \langle f, (\varphi_i)_i \rangle &: (D_i)_{i \in I} \rightarrow (\mathbb{V}\mathbb{J}E_k)_{k \in K}, \\ \langle g, (\psi_i)_i \rangle &: (D_i)_{i \in I} \rightarrow (\mathbb{V}\mathbb{J}1)_{j < 1} \end{aligned}$$

of $\Pi(\text{Set}^\Xi)$ where

$$\begin{aligned} g : 1 &\rightarrow I & f : K &\rightarrow I \\ \psi_* : D_{g(*)} &\rightarrow \mathbb{V}\mathbb{J}1 & \varphi_k : D_{f(k)} &\rightarrow \mathbb{V}\mathbb{J}E_k, \quad \text{for } k \in K. \end{aligned}$$

With our conventions regarding polynomial functors, we can write the latter as

$$\varphi_s : \text{dom}(\mu(s)) \rightarrow \sum_{v \in \text{dom}(s)} \mathbb{V}\mathbb{J} \text{dom}(s(v)), \quad \text{for } s \in \mathbb{M}\mathbb{M}A.$$

5. Eilenberg-Moore Algebras

We can use a monad \mathbb{M} to specify what kinds of objects the languages we study contain and in which way these compose. Next we need a way to describe such a language in an algebraic way. This will be done by a morphism $\mathbb{M}\Sigma \rightarrow \mathfrak{A}$ for some suitable algebra \mathfrak{A} . The goal of this section is to define the kind of algebras we use for this. As it turns out there is a canonical notion of an algebra associated with every monad \mathbb{M} . To motivate the definition, let us first take a look at Σ -algebras. To specify an algebra for a signature Σ , we have to provide a set of elements and functions for each operation in Σ .

Definition 5.1. Let Σ be an Ξ -sorted algebraic signature. A Σ -structure

$$\mathfrak{A} = \langle A, (f^{\mathfrak{A}})_{f \in \Sigma} \rangle$$

consists of

- ♦ a universe $A \in \text{Set}^\Xi$ and

- ♦ one function $f^{\mathfrak{A}} : \prod_i A_{\eta_i} \rightarrow A_{\xi}$, for each function symbol $f \in \Sigma$ of type $\tilde{\eta} \rightarrow \xi$.

Examples. (a) For the signature $\Sigma = \{ \cdot \}$ of semigroups, a Σ -algebra $\mathfrak{S} = \langle S, \cdot^{\mathfrak{S}} \rangle$ consists of a set S together with a function $\cdot^{\mathfrak{S}} : S \times S \rightarrow S$. Such an algebra forms a semigroup, if $\cdot^{\mathfrak{S}}$ is associative.

(b) For ω -semigroups, we can use the signature $\Sigma = \{ \cdot, \cdot', \pi \}$ with sorts $\Xi = \{1, \infty\}$, where the types of \cdot , \cdot' , and π are, respectively, $11 \rightarrow 1$, $1\infty \rightarrow \infty$, and $111 \dots \rightarrow \infty$. An ω -semigroup is then a Σ -algebra where all operations are associative.

Thus, a Σ -algebra \mathfrak{A} allows one to compute all operations in the signature. Since terms are built up from these operations, it follows that we can also evaluate every Σ -term. But note that, if the term in question contains variables, we also need to know their values. In general we obtain an evaluation map

$$\text{Term}[\Sigma, X] \times A^X \rightarrow A,$$

where the second argument $\beta : X \rightarrow A$ maps each variable to its value. We can simplify this map if, instead of providing β , we replace in the given term each variable by its value. In that way, we obtain a map

$$\text{Term}[\Sigma, A] \rightarrow A.$$

Such an evaluation map contains sufficient information to completely describe the algebra \mathfrak{A} in question. For a general monad \mathbb{M} , we can use a similar description: we can use algebras consisting of a set A equipped with a *product* operation of the form $\pi : \mathbb{M}A \rightarrow A$. For instance, for words this product takes the form $\pi : A^+ \rightarrow A$, i.e., it multiplies a sequence of elements into a single element. Hence, π can be seen as a semigroup product of variable arity. But note that not every operation $\pi : A^+ \rightarrow A$ is of the form

$$\pi(\langle a_0, \dots, a_m \rangle) = a_0 \cdot a_1 \cdot \dots \cdot a_m$$

for some semigroup product $\cdot : A \times A \rightarrow A$. If we want to exactly capture the notion of a semigroup, we have to impose additional conditions on π . It

turns out, there are two such conditions: associativity requires that

$$\pi(\pi(w_0), \dots, \pi(w_m)) = \pi(w_0 \dots w_m), \quad \text{for all } w_0, \dots, w_m \in A^+,$$

and the fact that the product of a single element should return that element again requires that

$$\pi(\langle a \rangle) = a, \quad \text{for } a \in A.$$

These two conditions can be phrased more concisely as

$$\pi \circ \mathbb{M}\pi = \pi \circ \text{flat} \quad \text{and} \quad \pi \circ \text{sing} = \text{id}.$$

where *flat* and *sing* are the multiplication and the unit map of the monad. This leads us to the following definition.

Definition 5.2. Let $\mathbb{M} : \mathcal{D}^{\Xi} \rightarrow \mathcal{D}^{\Xi}$ be a monad.

(a) An *Eilenberg-Moore algebra* for \mathbb{M} , or *\mathbb{M} -algebra* for short, is a pair $\mathfrak{A} = \langle A, \pi \rangle$ consisting of an object $A \in \mathcal{D}^{\Xi}$ and a morphism $\pi : \mathbb{M}A \rightarrow A$ satisfying


$$\begin{aligned} \pi \circ \mathbb{M}\pi &= \pi \circ \text{flat}, \\ \pi \circ \text{sing} &= \text{id}. \end{aligned}$$

The first of these equations is called the *associative law* for π , the second one the *unit law*.

(b) A *morphism* $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ of \mathbb{M} -algebras, or *\mathbb{M} -morphism* for short, is a function $\varphi : A \rightarrow B$ commuting with the respective products in the sense that

$$\varphi \circ \pi = \pi \circ \mathbb{M}\varphi.$$

(c) We denote the category of all \mathbb{M} -algebras and their morphisms by $\text{Alg}(\mathbb{M})$.



(d) An algebra \mathfrak{A} is *finitary* if its universe A is sort-wise finite and \mathfrak{A} is finitely generated, i.e., there exists a finite set $C \subseteq A$ such that every element $a \in A$ can be written as $a = \pi(s)$, for some $s \in \mathbb{M}C$. \lrcorner

Examples. (a) As already explained above, \mathbb{M} -algebras for the monad $\mathbb{M}X = X^+$ are just semigroups. Similarly, algebras for the monad $\mathbb{M}X = X^*$ correspond to monoids. We obtain commutative monoids if we take the monad $\mathbb{M}X = X^*/\approx$ where

$$u \approx v \quad : \text{iff} \quad u \text{ is a permutation of } v.$$

Semilattices, i.e., idempotent commutative monoids, correspond to the finite power-set monad $\mathbb{M}X = \{w \mid w \subseteq X \text{ finite}\}$.

(b) Let us take a look at the functor

$$\mathbb{M}\langle A_1, A_\infty \rangle := \langle A_1^+, A_1^+ A_\infty \cup A_1^\omega \rangle$$

for infinite words. In this case an \mathbb{M} -algebra has two product functions

$$\pi_1 : A_1^+ \rightarrow A_1 \quad \text{and} \quad \pi_\infty : A_1^+ A_\infty \cup A_1^\omega \rightarrow A_\infty.$$

The laws of an \mathbb{M} -algebra ensure that π_1 corresponds to a semigroup product $A_1 \times A_1 \rightarrow A_1$ and π_∞ correspond to the additional products $A_1 \times A_\infty \rightarrow A_\infty$ and $A_1^\omega \rightarrow A_\infty$ of an ω -semigroup. Hence, in this case \mathbb{M} -algebras are nothing but ω -semigroups.

(c) Let $\mathfrak{M} = \langle M, \cdot \rangle$ be a monoid. The functor $\mathbb{M}X := M \times X$ forms a monad with product

$$\text{act} : M \times (M \times X) \rightarrow M \times X : \langle a, \langle b, x \rangle \rangle \mapsto \langle ab, x \rangle$$

and unit function

$$\text{unit} : X \rightarrow M \times X : x \mapsto \langle 1, x \rangle.$$

An \mathbb{M} -algebra then corresponds to a set X together with an action of M on X . In particular, if \mathfrak{M} is a group, the category $\text{Alg}(\mathbb{M})$ of \mathbb{M} -algebras is isomorphic to the category $\mathfrak{M}\text{-Set}$ of \mathfrak{M} -sets. \lrcorner

There is a natural way to turn a set of the form $\mathbb{M}A$ into an \mathbb{M} -algebra: we can chose the function $\text{flat} : \mathbb{M}\mathbb{M}A \rightarrow \mathbb{M}A$ as the product. It turns out that algebras of this form are exactly the *free algebras*.

Proposition 5.3. *For every $A \in \mathcal{D}^\Xi$, there exists a free \mathbb{M} -algebra over A . It has the form $\langle \mathbb{M}A, \text{flat} \rangle$.*

Proof. Two of the three axioms of a monad precisely express that $\langle \mathbb{M}A, \text{flat} \rangle$ is an \mathbb{M} -algebra. Hence, it remains to prove freeness. Consider a function $f : A \rightarrow B$ where B is the domain of an \mathbb{M} -algebra \mathfrak{B} . Then

$$\varphi := \pi \circ \mathbb{M}f : \mathbb{M}A \rightarrow B$$

is a morphism of \mathbb{M} -algebras since

$$\begin{aligned} \varphi \circ \text{flat} &= \pi \circ \mathbb{M}f \circ \text{flat} \\ &= \pi \circ \text{flat} \circ \mathbb{M}\mathbb{M}f \\ &= \pi \circ \mathbb{M}\pi \circ \mathbb{M}\mathbb{M}f = \pi \circ \mathbb{M}\varphi. \end{aligned}$$

For uniqueness, suppose that $\psi : \mathbb{M}A \rightarrow \mathfrak{B}$ is a morphism with $\psi \circ \text{sing} = f$. Then

$$\varphi = \pi \circ \mathbb{M}f = \pi \circ \mathbb{M}(\psi \circ \text{sing}) = \psi \circ \text{flat} \circ \mathbb{M}\text{sing} = \psi. \quad \square$$

Example. (a) The free monoid generated by X is X^* where the product is given by concatenation. This is also the free \mathbb{M} -algebra for the monad $\mathbb{M}X = X^*$.

(b) A free Σ -algebra is called a *term algebra*. of the form $\text{Term}[\Sigma, X]$, for some set X . Its elements are all terms and a function f maps a tuple of terms \vec{i} to the term $f(\vec{i})$. \lrcorner

Corollary 5.4. *For every monad $\mathbb{M} : \mathcal{C} \rightarrow \mathcal{C}$ there exists an adjunction $\mathbb{F} \dashv \mathbb{U}$ between \mathcal{C} and $\text{Alg}(\mathbb{M})$ such that $\mathbb{M} = \mathbb{U} \circ \mathbb{F}$.*

Proof. Let \mathbb{F} be the functor sending $A \in \mathcal{C}$ to the free \mathbb{M} -algebra $\langle \mathbb{M}A, \text{flat} \rangle$, and let \mathbb{U} be the functor mapping an \mathbb{M} -algebra $\langle A, \pi \rangle$ to its universe A . By

the Proposition 5.3, there exists a bijection

$$\text{Alg}(\mathbb{M})(\mathbb{F}A, \mathfrak{B}) \cong \mathcal{C}(A, \mathbb{U}\mathfrak{B}).$$

It is straightforward to check that it is natural in A and \mathfrak{B} . \square

We conclude this section with three technical results that sometimes come in handy. First, checking the associative law is often tedious. In many cases we can use the following result to avoid having to do these kind of calculations.

Lemma 5.5. *Let \mathfrak{A} be an \mathbb{M} -algebra and $f : A \rightarrow B$ and $p : \mathbb{M}B \rightarrow B$ functions such that f and $\mathbb{M}\mathbb{M}f$ are epimorphisms and*

$$f \circ \pi = p \circ \mathbb{M}f.$$

Then $\mathfrak{B} := \langle B, p \rangle$ is an \mathbb{M} -algebra and $f : \mathfrak{A} \rightarrow \mathfrak{B}$ a morphism of \mathbb{M} -algebras.

Proof. By assumption, f satisfies the equation for a morphism. Hence, we only need to check the axioms of an \mathbb{M} -algebra. We have

$$\begin{aligned} p \circ \text{sing} \circ f &= p \circ \mathbb{M}f \circ \text{sing} \\ &= f \circ \pi \circ \text{sing} \\ &= f, \\ p \circ \mathbb{M}p \circ \mathbb{M}\mathbb{M}f &= p \circ \mathbb{M}(p \circ \mathbb{M}f) \\ &= p \circ \mathbb{M}(f \circ \pi) \\ &= f \circ \pi \circ \mathbb{M}\pi \\ &= f \circ \pi \circ \text{flat} \\ &= p \circ \mathbb{M}f \circ \text{flat} = p \circ \text{flat} \circ \mathbb{M}\mathbb{M}f. \end{aligned}$$

Since f and $\mathbb{M}\mathbb{M}f$ are epimorphisms, this implies that $p \circ \text{sing} = \text{id}$ and $p \circ \mathbb{M}p = p \circ \text{flat}$. \square

Another useful result is the following observation which allows us to prove that a function is a morphism.

Lemma 5.6. *Let $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ and $\varepsilon : \mathfrak{A} \rightarrow \mathfrak{C}$ be \mathbb{M} -morphisms where $\mathbb{M}\varepsilon$ is an epimorphism. If $f : C \rightarrow B$ is a function satisfying $f \circ \varepsilon = \varphi$, then f is also an \mathbb{M} -morphism.*

Proof. Note that

$$\pi \circ \mathbb{M}f \circ \mathbb{M}\varepsilon = \pi \circ \mathbb{M}\varphi = \varphi \circ \pi = f \circ \varepsilon \circ \pi = f \circ \pi \circ \mathbb{M}\varepsilon.$$

Since $\mathbb{M}\varepsilon$ is an epimorphism, this implies that $\pi \circ \mathbb{M}f = f \circ \pi$. \square

Limits and Colimits of Algebras

Finally, let us take a look at limits and colimits of \mathbb{M} -algebras.

Proposition 5.7. *The forgetful functor $\text{Alg}(\mathbb{M}) \rightarrow \mathcal{D}^\Xi$ reflects limits. Hence, $\text{Alg}(\mathbb{M})$ has all limits that exist in \mathcal{D} .*

Proof. Suppose that \mathcal{D} has limits with index category \mathcal{I} . Then so does \mathcal{D}^Ξ . (As limits in \mathcal{D}^Ξ are computed componentwise.) Let $D : \mathcal{I} \rightarrow \text{Alg}(\mathbb{M})$ be a diagram and let $\mathbb{U} : \text{Alg}(\mathbb{M}) \rightarrow \mathcal{D}^\Xi$ be the forgetful functor. Then $\mathbb{U} \circ D$ has a limit A . Let $(\lambda_i)_i$ be the corresponding limiting cone. We claim that there exists an \mathbb{M} -algebra \mathfrak{A} with universe A such that $\mathfrak{A} = \lim D$ and each $\lambda_i : \mathfrak{A} \rightarrow D(i)$ is an \mathbb{M} -morphism. To define the product of \mathfrak{A} , we set

$$\mu_i := \pi_i \circ \mathbb{M}\lambda_i : \mathbb{M}A \rightarrow D(i), \quad \text{for } i \in \mathcal{I},$$

where $\pi_i : \mathbb{M}D(i) \rightarrow D(i)$ is the product of $D(i)$. For every morphism $f : i \rightarrow j$ of \mathcal{I} , it follows that

$$\begin{aligned} \mathbb{U}Df \circ \mu_i &= \mathbb{U}Df \circ \pi_i \circ \mathbb{M}\lambda_i \\ &= \pi_i \circ \mathbb{M}\mathbb{U}Df \circ \mathbb{M}\lambda_i = \pi_i \circ \mathbb{M}\lambda_i = \mu_i. \end{aligned}$$

Hence, $(\mu_i)_i$ is a cone from $\mathbb{M}A$ to $\mathbb{U} \circ D$. By universality, there exists a unique morphism $\rho : \mathbb{M}A \rightarrow A$ satisfying

$$\lambda_i \circ \rho = \mu_i = \pi_i \circ \mathbb{M}\lambda_i, \quad \text{for all } i.$$

This implies that λ_i is an \mathbb{M} -morphism from $\mathfrak{A} := \langle A, \rho \rangle$ to $D(i)$, where the fact that \mathfrak{A} forms an \mathbb{M} -algebra follows since

$$\begin{aligned} \lambda_i \circ \rho \circ \text{sing} &= \pi_i \circ \mathbb{M}\lambda_i \circ \text{sing} \\ &= \pi_i \circ \text{sing} \circ \lambda_i \\ &= \lambda_i, \\ \lambda_i \circ \rho \circ \text{flat} &= \pi_i \circ \mathbb{M}\lambda_i \circ \text{flat} \\ &= \pi_i \circ \text{flat} \circ \mathbb{M}\lambda_i \\ &= \pi_i \circ \mathbb{M}\pi_i \circ \mathbb{M}\lambda_i \\ &= \pi_i \circ \mathbb{M}\lambda_i \circ \mathbb{M}\rho \\ &= \lambda_i \circ \rho \circ \mathbb{M}\rho, \end{aligned}$$

for all i . Limiting cones being jointly monomorphic, we obtain $\rho \circ \text{sing} = \text{id}$ and $\rho \circ \text{flat} = \rho \circ \mathbb{M}\rho$. Consequently, $(\lambda_i)_i$ forms a cone from $\langle A, \rho \rangle$ to D in $\text{Alg}(\mathbb{M})$. This cone is limiting since, given any cone $(\nu_i)_i$ from some \mathbb{M} -algebra $\mathfrak{B} = \langle B, \pi \rangle$ to D , there exists a unique function $\sigma : B \rightarrow A$ (in \mathcal{D}^Ξ) such that

$$\lambda_i \circ \sigma = \nu_i, \quad \text{for all } i.$$

This function is in fact a morphism of \mathbb{M} -algebras since

$$\begin{aligned} \lambda_i \circ \sigma \circ \pi &= \nu_i \circ \pi \\ &= \pi_i \circ \mathbb{M}\nu_i \\ &= \pi_i \circ \mathbb{M}\lambda_i \circ \mathbb{M}\sigma = \lambda_i \circ \pi_i \circ \mathbb{M}\sigma, \end{aligned}$$

for all i . □

The existence of colimits is less straightforward and requires additional assumptions on \mathbb{M} and \mathcal{C} . We start with a technical lemma.

Lemma 5.8. *Let \mathbb{M} be a monad on \mathcal{C} and let $\mathfrak{A} = \langle A, \pi \rangle$ be an \mathbb{M} -algebra. Then $\pi : \mathbb{M}A \rightarrow A$ is the coequaliser of flat and $\mathbb{M}\pi$ in $\text{Alg}(\mathbb{M})$.*

Proof. First, note that $\pi : \mathbb{M}A \rightarrow \mathfrak{A}$ and $\text{flat}, \mathbb{M}\pi : \mathbb{M}\mathbb{M}A \rightarrow \mathbb{M}A$ are \mathbb{M} -morphisms since

$$\begin{aligned} \pi \circ \text{flat} &= \pi \circ \mathbb{M}\pi, \\ \text{flat} \circ \text{flat} &= \text{flat} \circ \mathbb{M}\text{flat}, \\ \mathbb{M}\pi \circ \text{flat} &= \text{flat} \circ \mathbb{M}\mathbb{M}\pi. \end{aligned}$$

The first of these equations also shows that π (together with $\pi \circ \mathbb{M}\pi$) forms a cocone of the diagram. Hence, we only have to prove that π is limiting. Suppose $\varphi : \mathbb{M}A \rightarrow \mathfrak{B}$ is another \mathbb{M} -morphism such that $\varphi \circ \text{flat} = \varphi \circ \mathbb{M}\pi$. The function $\sigma := \varphi \circ \text{sing}$ satisfies

$$\begin{aligned} \sigma \circ \pi &= \varphi \circ \text{sing} \circ \pi \\ &= \varphi \circ \mathbb{M}\pi \circ \text{sing} \\ &= \varphi \circ \text{flat} \circ \text{sing} = \varphi. \end{aligned}$$

Furthermore, it is unique since $\tau \circ \pi = \varphi$ implies

$$\tau = \tau \circ \pi \circ \text{sing} = \varphi \circ \text{sing} = \sigma \circ \pi \circ \text{sing} = \sigma.$$

Hence, it remains to show that σ is an \mathbb{M} -morphism.

$$\begin{aligned} \sigma \circ \pi &= \varphi \circ \text{sing} \circ \pi \\ &= \varphi \circ \mathbb{M}\pi \circ \text{sing} \\ &= \varphi \circ \text{flat} \circ \text{sing} \\ &= \varphi \\ &= \varphi \circ \text{flat} \circ \mathbb{M}\text{sing} \\ &= \text{flat} \circ \mathbb{M}\varphi \circ \mathbb{M}\text{sing} \\ &= \text{flat} \circ \mathbb{M}\sigma. \end{aligned}$$

□

The following criterion greatly simplifies existence proofs.

Proposition 5.9. *Let \mathbb{M} be a monad on a cocomplete category \mathcal{C} . Then $\text{Alg}(\mathbb{M})$ is cocomplete if, and only if, it has coequalisers.*

Proof. (\Rightarrow) is trivial. For (\Leftarrow) , suppose that $\text{Alg}(\mathbb{M})$ has coequalisers. Since a category is cocomplete if, and only if, it has coproducts and coequalisers, it is sufficient to show that $\text{Alg}(\mathbb{M})$ has coproducts. As an example of the construction let us take a look at how we can construct coproducts of groups. We cannot simply take their disjoint union since then the product of elements from different components would be undefined. What we can do instead is to take the free group generated by this disjoint union and then quotient it by all equalities that hold in each component. (This last step requires a coequaliser.)

The general case is analogous. Consider \mathbb{M} -algebras $\mathfrak{A}_i = \langle A_i, \pi_i \rangle$, for $i \in I$. To compute $\sum_i \mathfrak{A}_i$, let $s_i : A_i \rightarrow \sum_j A_j$ and $t_i : \mathbb{M}A_i \rightarrow \sum_j \mathbb{M}A_j$ be the inclusion morphisms (where the coproducts are computed in \mathcal{C}). Then $\mathbb{M}s_i$ forms a cocone from the coproduct diagram $(\mathbb{M}A_i)_i$ to $\mathbb{M}\sum_j A_j$. Consequently, there exists a unique morphism $w : \sum_j \mathbb{M}A_j \rightarrow \mathbb{M}\sum_j A_j$ such that

$$w \circ t_i = \mathbb{M}s_i, \quad \text{for all } i.$$

Let $q : \mathbb{M}\sum_j A_j \rightarrow \mathfrak{C}$ be the coequaliser (in $\text{Alg}(\mathbb{M})$) of

$$\mathbb{M}\sum_i \pi_i, \text{ flat} \circ \mathbb{M}w : \mathbb{M}\sum_j \mathbb{M}A_j \rightarrow \mathbb{M}\sum_j A_j.$$

To see that this is well-defined, note that $\mathbb{M}\sum_i \pi_i$, flat, and $\mathbb{M}w$ are indeed \mathbb{M} -morphisms since

$$\begin{aligned} \mathbb{M}\sum_i \pi_i \circ \text{flat} &= \text{flat} \circ \mathbb{M}\mathbb{M}\sum_i \pi_i, \\ \text{flat} \circ \text{flat} &= \text{flat} \circ \mathbb{M}\text{flat} \\ \mathbb{M}w \circ \text{flat} &= \text{flat} \circ \mathbb{M}\mathbb{M}w. \end{aligned}$$

For every $i \in I$, as similar computation shows that $\mathbb{M}s_i : \mathbb{M}A_i \rightarrow \mathbb{M}\sum_j A_j$ is an \mathbb{M} -morphism. Since

$$\begin{aligned} q \circ \mathbb{M}s_i \circ \text{flat} &= q \circ \text{flat} \circ \mathbb{M}\mathbb{M}s_i \\ &= q \circ \text{flat} \circ \mathbb{M}(w \circ t_i) \\ &= q \circ \mathbb{M}\sum_j \pi_j \circ \mathbb{M}t_i = q \circ \mathbb{M}(s_i \circ \pi_i), \end{aligned}$$

we can use Lemma 5.8 to find unique \mathbb{M} -morphisms $\lambda_i : \mathfrak{A}_i \rightarrow \mathfrak{C}$ satisfying

$$\lambda_i \circ \pi_i = q \circ \mathbb{M}s_i, \quad \text{for all } i.$$

We claim that $\mathfrak{C} = \sum_i \mathfrak{A}_i$ and that $(\lambda_i)_i$ is the corresponding limiting cocone.

Hence, suppose that $(\mu_i)_i$ is a cocone from $(\mathfrak{A}_i)_i$ to some algebra \mathfrak{B} . Since $(\lambda_i)_i$ form a cocone in \mathcal{C} , we obtain a unique morphism $u : \sum_j A_j \rightarrow B$ such that

$$u \circ s_i = \mu_i, \quad \text{for all } i.$$

Similarly, for the cocone $(\mathbb{M}\mu_i)_i$, we obtain a unique morphism $v : \sum_j \mathbb{M}A_j \rightarrow \mathbb{M}B$ such that

$$v \circ t_i = \mathbb{M}\mu_i, \quad \text{for all } i.$$

We will show below that the morphism $\rho := \pi \circ \mathbb{M}u : \mathbb{M}\sum_j A_j \rightarrow \mathfrak{B}$ satisfies

$$\rho \circ (\text{flat} \circ \mathbb{M}w) = \rho \circ \mathbb{M}\sum_j \pi_j.$$

Since q is the coequaliser, we then obtain a unique morphism $\varphi : \mathfrak{C} \rightarrow \mathfrak{B}$ satisfying

$$\varphi \circ q = \rho.$$

This implies that

$$\begin{aligned} \varphi \circ \lambda_i &= \varphi \circ \lambda_i \circ \pi_i \circ \text{sing} \\ &= \varphi \circ q \circ \mathbb{M}s_i \circ \text{sing} \\ &= \rho \circ \mathbb{M}s_i \circ \text{sing} \\ &= \pi \circ \mathbb{M}u \circ \mathbb{M}s_i \circ \text{sing} \\ &= \pi \circ \mathbb{M}\mu_i \circ \text{sing} \\ &= \mu_i \circ \pi_i \circ \text{sing} \\ &= \mu_i. \end{aligned}$$

Hence, the cocone $(\mu_i)_i$ factorises through $(\lambda_i)_i$ via φ . To show that this factorisation is unique, we consider some \mathbb{M} -morphisms ψ satisfying $\psi \circ \lambda_i = \mu_i$. Then

$$\begin{aligned}\psi \circ q \circ w \circ t_i &= \psi \circ q \circ \mathbb{M}s_i \\ &= \psi \circ \lambda_i \circ \pi_i \\ &= \mu_i \circ \pi_i \\ &= u \circ s_i \circ \pi_i \\ &= u \circ \sum_j \pi_j \circ t_i, \quad \text{for all } i.\end{aligned}$$

Since limiting cocones are jointly epimorphic, it follows that

$$\psi \circ q \circ w = u \circ \sum_j \pi_j.$$

Consequently,

$$\begin{aligned}\psi \circ q &= \psi \circ q \circ \mathbb{M} \sum_j (\pi_j \circ \text{sing}) \\ &= \psi \circ q \circ \text{flat} \circ \mathbb{M}w \circ \mathbb{M} \sum_j \text{sing} \\ &= \psi \circ \pi \circ \mathbb{M}q \circ \mathbb{M}w \circ \mathbb{M} \sum_j \text{sing} \\ &= \pi \circ \mathbb{M}\psi \circ \mathbb{M}q \circ \mathbb{M}w \circ \mathbb{M} \sum_j \text{sing} \\ &= \pi \circ \mathbb{M}(u \circ \sum_j \pi_j) \circ \mathbb{M} \sum_j \text{sing} \\ &= \pi \circ \mathbb{M}u \\ &= \varphi \circ q.\end{aligned}$$

Since q is an epimorphism, this implies that $\psi = \varphi$, as desired.

Hence, it remains to prove the above claim. First, ρ is indeed an \mathbb{M} -morphism since

$$\begin{aligned}\pi \circ \mathbb{M}u \circ \text{flat} &= \pi \circ \text{flat} \circ \mathbb{M}\mathbb{M}u \\ &= \pi \circ \mathbb{M}\pi \circ \mathbb{M}\mathbb{M}u = \pi \circ \mathbb{M}(\pi \circ \mathbb{M}u).\end{aligned}$$

Furthermore, for all $i \in I$, we have

$$\begin{aligned}\pi \circ \mathbb{M}u \circ w \circ t_i &= \pi \circ \mathbb{M}u \circ \mathbb{M}s_i \\ &= \pi \circ \mathbb{M}\mu_i \\ &= \pi \circ v \circ t_i \\ &= \pi \circ \mathbb{M}\mu_i \\ &= \mu_i \circ \pi_i \\ &= u \circ s_i \circ \pi_i = u \circ \sum_j \pi_j \circ t_i,\end{aligned}$$

where π is the product of \mathfrak{B} . Since the morphisms of a limiting cocone are jointly epimorphic, it follows that

$$\pi \circ \mathbb{M}u \circ w = u \circ \sum_j \pi_j.$$

Hence,

$$\begin{aligned}\rho \circ (\text{flat} \circ \mathbb{M}w) &= \pi \circ \mathbb{M}u \circ \text{flat} \circ \mathbb{M}w \\ &= \pi \circ \text{flat} \circ \mathbb{M}\mathbb{M}u \circ \mathbb{M}w \\ &= \pi \circ \mathbb{M}\pi \circ \mathbb{M}(\mathbb{M}u \circ w) \\ &= \pi \circ \mathbb{M}(u \circ \sum_j \pi_j) \\ &= \pi \circ \mathbb{M}u \circ \mathbb{M} \sum_j \pi_j = \rho \circ \mathbb{M} \sum_j \pi_j,\end{aligned}$$

as desired. \square

Exercise 5.1. (a) Let \mathbb{M} be a monad on a category \mathcal{C} with coequalisers and suppose that \mathbb{M} preserves coequalisers. Show that $\text{Alg}(\mathbb{M})$ is cocomplete.

(b) Show that every polynomial functor $\text{Set} \rightarrow \text{Set}$ preserves coequalisers.

Thus, we only have to show that the category $\text{Alg}(\mathbb{M})$ has coequalisers. We start with a lemma simplifying this task: it is sufficient to construct *weak* coequalisers.

Definition 5.10. A weak coequaliser of two morphisms $\varphi, \psi : A \rightarrow B$ is a morphism $\rho : B \rightarrow C$ such that $\rho \circ \varphi = \rho \circ \psi$ and every morphism $\sigma : B \rightarrow D$ with $\sigma \circ \varphi = \sigma \circ \psi$ factorises (not necessarily uniquely) through ρ . \square

The idea behind the following construction comes from the theory of partial orders: a supremum can be computed as the infimum of all upper bounds. Since colimits can be regarded as generalisations of suprema, we can try to compute a colimit as a limit of suitable ‘upper bounds’. In our case, we obtain a coequaliser as a limit of weak coequalisers. Before giving the formal construction, let us mention the following technical result.

Lemma 5.11. Let \mathcal{C} be a category and $A \in \mathcal{C}$. The comma category $(A \downarrow \mathcal{C})$ has all limits that exist in \mathcal{C} .

Proof. Let $U : (A \downarrow \mathcal{C}) \rightarrow \mathcal{C}$ be the functor mapping an object $f : A \rightarrow B$ to B and a morphism $\varphi : f \rightarrow g$ to the corresponding morphism between the codomains of f and g . It is sufficient to show that U creates limits. Hence, let $D : \mathcal{I} \rightarrow (A \downarrow \mathcal{C})$ be a diagram such that $U \circ D$ has a limit B and let $(\lambda_i)_i$ be the corresponding limiting cone. For $i \in \mathcal{I}$, set $\mu_i := D(i)$. For every morphism $h : i \rightarrow k$ of \mathcal{I} , the fact that Dh is a morphism $\mu_i \rightarrow \mu_k$ in $(A \downarrow \mathcal{C})$ implies that

$$UDh \circ \mu_i = \mu_k.$$

Consequently, $(\mu_i)_i$ is a cone from A to $U \circ D$. As $(\lambda_i)_i$ is limiting, there exists a unique morphism $\sigma : A \rightarrow B$ such that

$$\lambda_i \circ \sigma = \mu_i, \quad \text{for all } i.$$

We can consider λ_i as a morphism $\lambda_i : \sigma \rightarrow \mu_i$ of $(A \downarrow \mathcal{C})$. Hence, $(\lambda_i)_i$ induces a cone from σ to D . It is straightforward to check that it is limiting. \square

Lemma 5.12. Let \mathcal{C} be a complete category. Every pair of morphisms with a weak coequaliser has a coequaliser.

Proof. Fix two morphisms $\varphi, \psi : A \rightarrow B$ with a weak coequaliser $q : B \rightarrow C$. To construct their coequaliser, we consider the category $\mathcal{E} \subseteq (B \downarrow \mathcal{C})$ of morphisms σ with $\sigma \circ \varphi = \sigma \circ \psi$ and the subcategory $\mathcal{W} \subseteq \mathcal{E}$ of all weak coequalisers $\sigma : B \rightarrow C$ of φ and ψ with codomain C . We have to show that \mathcal{E} has an initial element. We do so by proving that \mathcal{W} has an initial element and that the inclusion $\mathcal{W} \rightarrow \mathcal{E}$ is final (which means that the inclusions $\mathcal{W} \rightarrow (B \downarrow \mathcal{C})$ and $\mathcal{E} \rightarrow (B \downarrow \mathcal{C})$ have the same limit).

Let $D : \mathcal{W} \rightarrow (B \downarrow \mathcal{C})$ be the inclusion functor. Since $\mathcal{W} \subseteq \mathcal{C}(B, C)$ is a set, D has a limit $\tau : B \rightarrow C$. Let $\varepsilon : \rho \rightarrow \tau$ be the equaliser of all morphisms in $\mathcal{E}(\tau, \tau)$ (which exists by Lemma 5.11). We claim that its domain ρ is the initial element of \mathcal{E} .

We start by proving that $\rho \in \mathcal{W}$. Hence, fix some $\sigma \in \mathcal{E}$. To find a morphism $\rho \rightarrow \sigma$ we choose an arbitrary $v \in \mathcal{W}$. Let $\lambda : \tau \rightarrow v$ be the corresponding component of the limiting cone from τ to D . Since $v \in \mathcal{W}$ there also exists a morphism $f : v \rightarrow \sigma$. We therefore obtain a morphism

$$f \circ \lambda \circ \varepsilon : \rho \rightarrow \tau \rightarrow v \rightarrow \sigma.$$

Hence, ρ is a weak coequaliser of φ and ψ , which implies that $\rho \in \mathcal{W}$.

It remains to prove uniqueness. Suppose that, for some $\sigma \in \mathcal{E}$, there are two morphisms $f, g : \rho \rightarrow \sigma$. Let $e : v \rightarrow \rho$ be their equaliser (which exists by Lemma 5.11). Then $v \in \mathcal{E}$ and $\rho \in \mathcal{W}$ implies that there is some morphism $h : \rho \rightarrow v$. Since $\varepsilon \circ e \circ h$ and id are both morphisms $\tau \rightarrow \tau$ and ε is the equaliser of all such morphisms, it follows that

$$\varepsilon \circ e \circ h \circ \varepsilon = \text{id} \circ \varepsilon = \varepsilon \circ \text{id}.$$

Equalisers being monomorphisms this implies that $e \circ (h \circ \varepsilon) = \text{id}$. Consequently,


$$f = f \circ e \circ (h \circ \varepsilon) = g \circ e \circ (h \circ \varepsilon) = g. \quad \square$$

Theorem 5.13. Let \mathbb{M} be a monad on a category \mathcal{C} that is complete and cocomplete. If there exists a regular cardinal κ such that \mathbb{M} preserves κ -filtered colimits, then $\text{Alg}(\mathbb{M})$ is complete and cocomplete.

Proof. Completeness follows by Proposition 5.7. For cocompleteness, we only have to check the existence of coequalisers by Proposition 5.9. Hence, fix two \mathbb{M} -morphisms $\varphi, \psi : \mathfrak{A} \rightarrow \mathfrak{B}$. By Lemma 5.12 and since $\text{Alg}(\mathbb{M})$ is complete, it is sufficient to find a weak coequaliser of φ and ψ . We construct it as the colimit of a chain $Q(i)$, $i < \kappa + 1$. In fact, we will construct

- two diagrams $P, Q : \kappa + 1 \rightarrow \mathcal{C}$,
- natural transformations $\sigma : P \Rightarrow Q$ and $\tau : P \Rightarrow \mathbb{M} \circ Q$, and
- two cones $(p_i)_i$ and $(q_i)_i$ from B to, respectively, P and Q ,

such that $Q_* := Q(\kappa)$ will be the universe of the weak coequaliser, $P_* := P(\kappa)$ will be equal to $\mathbb{M}Q_*$, the morphism $\sigma_\kappa : P_* \rightarrow Q_*$ will be the product morphism, and $q_\kappa : \mathfrak{B} \rightarrow Q_*$ will be the coequaliser morphism.



Our construction proceeds by induction on $i < \kappa + 1$. For $i = 0$, let $p_i : B \rightarrow P(0)$ be the coequaliser (in \mathcal{C}) of $\mathbb{M}\varphi$ and $\mathbb{M}\psi$ and let $q_i : B \rightarrow Q(0)$ be the one of φ and ψ . The product morphisms of \mathfrak{A} and \mathfrak{B} form a natural transformations between the coequaliser diagrams for $\mathbb{M}\varphi, \mathbb{M}\psi$ and φ, ψ . This transformation induces a unique morphism $\sigma_0 : P(0) \rightarrow Q(0)$ such that

$$\sigma_0 \circ p_0 = q_0 \circ \pi.$$

Finally, $\mathbb{M}(q_0 \circ \varphi) = \mathbb{M}(q_0 \circ \psi)$ and universality of p_0 implies that there exists a unique morphism $\tau_0 : P(0) \rightarrow \mathbb{M}Q(0)$ such that

$$\tau_0 \circ p_0 = \mathbb{M}q_0.$$

For the inductive step, suppose that $P(i), Q(i), \sigma_i, \tau_i, p_i, q_i$ are defined for all $i < \alpha$. We distinguish two cases. If α is a limit ordinal, we choose for $P(\alpha)$ the colimit of the diagram $P|_\alpha : \alpha \rightarrow \mathcal{C}$ and for $Q(\alpha)$ the colimit of $Q|_\alpha : \alpha \rightarrow \mathcal{C}$. Let $(P(i, \alpha))_{i < \alpha}$ and $(Q(i, \alpha))_{i < \alpha}$ be the corresponding limiting cocones. Set

$$p_\alpha := P(0, \alpha) \circ p_0 \quad \text{and} \quad q_\alpha := Q(0, \alpha) \circ q_0.$$

The morphisms $(\sigma_i)_{i < \alpha}$ form a natural transformation $P|_\alpha \Rightarrow Q|_\alpha$ while $(\tau_i)_{i < \alpha}$ form one $P|_\alpha \Rightarrow \mathbb{M} \circ Q|_\alpha$. Consequently, we obtain unique morphisms $\sigma_\alpha : P(\alpha) \rightarrow Q(\alpha)$ and $\tau_\alpha : P(\alpha) \rightarrow \mathbb{M}Q(\alpha)$ satisfying

$$\sigma_\alpha \circ P(i, \alpha) = Q(i, \alpha) \circ \sigma_i \quad \text{and} \quad \tau_\alpha \circ P(i, \alpha) = \mathbb{M}Q(i, \alpha) \circ \tau_i.$$

It remains to consider the case where $\alpha = i + 1$ is a successor ordinal. Set $P(i + 1) := \mathbb{M}Q(i)$ and let $\sigma_{i+1} : \mathbb{M}Q(i) \rightarrow Q(i + 1)$ be the coequaliser of

$$\mathbb{M}\sigma_i, \text{ flat} \circ \mathbb{M}\tau_i : \mathbb{M}P(i) \rightarrow \mathbb{M}Q(i).$$

We set

$$\begin{aligned} P(i, i + 1) &:= \tau_i, \\ Q(i, i + 1) &:= \sigma_{i+1} \circ \text{sing}, \\ \tau_{i+1} &:= \mathbb{M}Q(i, i + 1), \\ p_{i+1} &:= P(i, i + 1) \circ p_i, \\ q_{i+1} &:= Q(i, i + 1) \circ q_i. \end{aligned}$$

To check that these morphisms have the desired properties, note that

$$\begin{aligned} \sigma_{i+1} \circ P(i, i + 1) &= \sigma_{i+1} \circ \tau_i \\ &= \sigma_{i+1} \circ \text{flat} \circ \text{sing} \circ \tau_i \\ &= \sigma_{i+1} \circ \text{flat} \circ \mathbb{M}\tau_i \circ \text{sing} \\ &= \sigma_{i+1} \circ \mathbb{M}\sigma_i \circ \text{sing} \\ &= \sigma_{i+1} \circ \text{sing} \circ \sigma_i \\ &= Q(i, i + 1) \circ \sigma_i, \\ \tau_{i+1} \circ P(i, i + 1) &= \mathbb{M}Q(i, i + 1) \circ \tau_i. \end{aligned}$$

This concludes the construction of P, Q , and the associated morphisms. Set

$$\begin{aligned} P_* &:= P(\kappa), & \sigma_* &:= \sigma_\kappa : P_* \rightarrow Q_*, \\ Q_* &:= Q(\kappa), & q_* &:= q_\kappa : B \rightarrow Q_*. \end{aligned}$$

As $P(i+1) = \mathbb{M}Q(i)$, we have

$$\text{colim}_{i < \kappa} P(i) = \text{colim}_{i < \kappa} \mathbb{M}Q(i) \quad \text{and} \quad \tau_\kappa \text{ is an isomorphism.}$$

Furthermore, regularity of κ implies that the diagrams $P|_\kappa, Q|_\kappa : \kappa \rightarrow \mathcal{C}$ are κ -filtered. By assumption on \mathbb{M} , it therefore follows that

$$\begin{aligned} P_* &= P(\kappa) \\ &= \text{colim}_{i < \kappa} P(i) \\ &= \text{colim}_{i < \kappa} \mathbb{M}Q(i) \\ &= \mathbb{M}(\text{colim}_{i < \kappa} Q(i)) = \mathbb{M}Q(\kappa) = \mathbb{M}Q_*. \end{aligned}$$

Hence, $\sigma_* \circ \tau_*^{-1} : \mathbb{M}Q_* \rightarrow Q_*$. We claim that $q_* : \mathfrak{B} \rightarrow \langle Q_*, \sigma_* \circ \tau_*^{-1} \rangle$ is the desired weak coequaliser of φ and ψ in $\text{Alg}(\mathbb{M})$. First, note that q_* is an \mathbb{M} -morphism since

$$\begin{aligned} q_* \circ \pi &= Q(o, \kappa) \circ q_o \circ \pi \\ &= Q(o, \kappa) \circ \sigma_o \circ p_o \\ &= \sigma_* \circ P(o, \kappa) \circ p_o \\ &= \sigma_* \circ p_\kappa \\ &= (\sigma_* \circ \tau_*^{-1}) \circ \mathbb{M}q_*. \end{aligned}$$

Next, we prove that $\Omega := \langle Q_*, \sigma_* \rangle$ it is indeed an \mathbb{M} -algebra. For $i < \kappa$, let E_i be the diagram consisting of the morphisms

$$\begin{array}{ccccc} \mathbb{M}Q(i) & \xleftarrow{\mathbb{M}\tau_i} & \mathbb{M}P(i) & \xrightleftharpoons[\text{flat} \circ \mathbb{M}\tau_i]{\mathbb{M}\sigma_i} & \mathbb{M}Q(i) & \xrightarrow{\sigma_{i+1}} & Q(i+1) \end{array}$$

We take the colimit of the diagram consisting of the diagrams $E_i, i < \kappa$, and the morphisms

$$\begin{array}{ccccccc} \mathbb{M}Q(i) & \xleftarrow{\mathbb{M}\tau_i} & \mathbb{M}P(i) & \xrightleftharpoons[\text{flat} \circ \mathbb{M}\tau_i]{\mathbb{M}\sigma_i} & \mathbb{M}Q(i) & \xrightarrow{\sigma_{i+1}} & Q(i+1) \\ \downarrow \mathbb{M}Q(i, k) & & \downarrow \mathbb{M}P(i, k) & & \downarrow \mathbb{M}Q(i, k) & & \downarrow Q(i, k) \\ \mathbb{M}Q(k) & \xleftarrow{\mathbb{M}\tau_k} & \mathbb{M}P(k) & \xrightleftharpoons[\text{flat} \circ \mathbb{M}\tau_k]{\mathbb{M}\sigma_k} & \mathbb{M}Q(k) & \xrightarrow{\sigma_{k+1}} & Q(k+1) \end{array}$$

This colimit is of the form

$$\mathbb{M}Q(\kappa) \xleftarrow{\alpha} \mathbb{M}P(\kappa) \xrightleftharpoons[\gamma]{\beta} \mathbb{M}Q(\kappa) \xrightarrow{\delta} Q(\kappa)$$

By definition of a colimit of functors, α is the unique morphism $\mathbb{M}P_* \rightarrow \mathbb{M}Q_*$ satisfying

$$\alpha \circ \mathbb{M}P(i, \kappa) = \mathbb{M}Q(i, \kappa) \circ \mathbb{M}\tau_i, \quad \text{for all } i < \kappa.$$

Consequently, $\alpha = \mathbb{M}\tau_\kappa$. Similarly, it follows that

$$\beta = \mathbb{M}\sigma_\kappa, \quad \gamma = \text{flat} \circ \mathbb{M}\tau_\kappa, \quad \delta = \sigma_\kappa \circ \tau_\kappa^{-1}.$$

Hence, we have

$$\begin{aligned} (\sigma_\kappa \circ \tau_\kappa^{-1}) \circ \mathbb{M}(\sigma_\kappa \circ \tau_\kappa^{-1}) &= \delta \circ \beta \circ \mathbb{M}\tau_\kappa^{-1} \\ &= \delta \circ \gamma \circ \mathbb{M}\tau_\kappa^{-1} \\ &= (\sigma_\kappa \circ \tau_\kappa^{-1}) \circ \text{flat} \circ \mathbb{M}\tau_\kappa \circ \mathbb{M}\tau_\kappa^{-1} \\ &= (\sigma_\kappa \circ \tau_\kappa^{-1}) \circ \text{flat}, \end{aligned}$$

as desired.

As $q_* \circ \varphi = Q(o, \kappa) \circ q_o \circ \varphi = Q(o, \kappa) \circ q_o \circ \psi = q_* \circ \psi$, it remains to prove weak universality. We will show below that, for every \mathbb{M} -morphism

$\chi : \mathfrak{B} \rightarrow \mathfrak{C}$ satisfying $\chi \circ \varphi = \chi \circ \psi$, there exist a cone $(\mu_i)_i$ from Q to C such that

$$\mu_i \circ q_i = \chi \quad \text{and} \quad \mu_i \circ \sigma_i = \pi \circ \mathbb{M}\mu_i \circ \tau_i, \quad \text{for all } i \leq \kappa.$$

Note that this immediately implies weak universality: for $i = \kappa$, we obtain $\chi = \mu_\kappa \circ q_\kappa$ and μ_κ is an \mathbb{M} -morphism since

$$\mu_\kappa \circ (\sigma_\kappa \circ \tau_\kappa^{-1}) = \pi \circ \mathbb{M}\mu_\kappa \circ \tau_\kappa \circ \tau_\kappa^{-1} = \pi \circ \mathbb{M}\mu_\kappa.$$

It therefore remains to prove the above claim. We construct μ_i by induction on i . For $i = 0$, note that $\chi \circ \varphi = \chi \circ \psi$ implies that χ factorises through the coequaliser q_0 (in C). Hence, $\chi = \mu_0 \circ q_0$, for some $\mu_0 : Q(0) \rightarrow C$. It follows that

$$\begin{aligned} \mu_0 \circ \sigma_0 \circ p_0 &= \mu_0 \circ q_0 \circ \pi \\ &= \chi \circ \pi \\ &= \pi \circ \mathbb{M}\chi \\ &= \pi \circ \mathbb{M}\mu_0 \circ \mathbb{M}q_0 = \pi \circ \mathbb{M}\mu_0 \circ \tau_0 \circ p_0. \end{aligned}$$

Since p_0 is an epimorphism, this implies that $\mu_0 \circ \sigma_0 = \pi \circ \mathbb{M}\mu_0 \circ \tau_0$.

For the successor step, suppose that we have already defined μ_i . Since

$$\begin{aligned} \pi \circ \mathbb{M}(\mu_i \circ \sigma_i) &= \pi \circ \mathbb{M}(\pi \circ \mathbb{M}\mu_i \circ \tau_i) \\ &= \pi \circ \text{flat} \circ \mathbb{M}\mathbb{M}\mu_i \circ \mathbb{M}\tau_i \\ &= \pi \circ \mathbb{M}\mu_i \circ \text{flat} \circ \mathbb{M}\tau_i, \end{aligned}$$

the function $\pi \circ \mathbb{M}\mu_i : \mathbb{M}Q(i) \rightarrow C$ factorises through the coequaliser of $\mathbb{M}\sigma_i$ and $\text{flat} \circ \mathbb{M}\tau_i$, which is σ_{i+1} . Hence, there exists a function $\mu_{i+1} : Q(i+1) \rightarrow C$ such that

$$\mu_{i+1} \circ \sigma_{i+1} = \pi \circ \mathbb{M}\mu_i.$$

It follows that

$$\begin{aligned} \mu_{i+1} \circ \sigma_{i+1} &= \pi \circ \mathbb{M}\mu_i \\ &= \pi \circ \mathbb{M}(\pi \circ \text{sing}) \circ \mathbb{M}\mu_i \\ &= \pi \circ \mathbb{M}(\pi \circ \mathbb{M}\mu_i \circ \text{sing}) \\ &= \pi \circ \mathbb{M}(\mu_{i+1} \circ \sigma_{i+1} \circ \text{sing}) \\ &= \pi \circ \mathbb{M}\mu_{i+1} \circ \mathbb{M}Q(i, i+1) \\ &= \pi \circ \mathbb{M}\mu_{i+1} \circ \tau_{i+1}, \\ \mu_{i+1} \circ Q(i, i+1) &= \mu_{i+1} \circ \sigma_{i+1} \circ \text{sing} \\ &= \pi \circ \mathbb{M}\mu_{i+1} \circ \text{sing} \\ &= \pi \circ \text{sing} \circ \mu_i \\ &= \mu_i. \end{aligned}$$

Finally, suppose that we have already defined μ_i , for every $i < \delta$, where δ is some limit ordinal. Then $(\mu_i)_{i < \delta}$ forms a cocone from $Q|_\delta$ to C . Since $(Q(i, \delta))_i$ is limiting, there exists a morphism $\mu_\delta : Q(\delta) \rightarrow C$ such that

$$\mu_\delta \circ Q(i, \delta) = \mu_i, \quad \text{for all } i < \delta.$$

This implies that

$$\begin{aligned} \mu_\delta \circ q_\delta &= \mu_\delta \circ Q(i, \delta) \circ q_i \\ &= \mu_i \circ q_i \\ &= \chi, \\ \mu_\delta \circ \sigma_\delta &= \mu_\delta \circ Q(i, \delta) \circ \sigma_i \\ &= \mu_i \circ \sigma_i \\ &= \pi \circ \mathbb{M}\mu_i \circ \tau_i \\ &= \pi \circ \mathbb{M}(\mu_\delta \circ Q(i, \delta)) \circ \tau_i \\ &= \pi \circ \mathbb{M}\mu_\delta \circ \tau_\delta. \end{aligned}$$

□

6. Lifting Monads

The free monad construction allows us to turn arbitrary functors into monads. In this section we present several ways to construct monads by transferring a monad from one category to another one. We start with a transfer along an adjunction.

Proposition 6.1. *Let $\mathbb{F} \dashv \mathbb{G}$ be an adjunction between the categories \mathcal{C} and \mathcal{D} , let $e : \text{Id} \Rightarrow \mathbb{G}\mathbb{F}$ be its unit, $i : \mathbb{F}\mathbb{G} \Rightarrow \text{Id}$ the counit, and let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a monad on \mathcal{D} .*

(a) $\langle \mathbb{N}, \nu, \eta \rangle$ forms a monad on \mathcal{D} , where

$$\mathbb{N} := \mathbb{G}\mathbb{M}\mathbb{F}, \quad \nu := \mathbb{G}(\mu \circ \mathbb{M}i), \quad \text{and} \quad \eta := \mathbb{G}\varepsilon \circ e.$$

(b) The functor $\mathbb{G} : \mathcal{D} \rightarrow \mathcal{C}$ can be lifted to a functor $\widehat{\mathbb{G}} : \text{Alg}(\mathbb{M}) \rightarrow \text{Alg}(\mathbb{N})$ which maps $\mathfrak{A} = \langle A, \pi \rangle$ to $\langle \mathbb{G}A, \mathbb{G}(\pi \circ \mathbb{M}i) \rangle$.

(c) If $\text{Alg}(\mathbb{M})$ has coequalisers, then $\widehat{\mathbb{G}}$ has a left adjoint $\widehat{\mathbb{F}}$ mapping an algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the coequaliser (in $\text{Alg}(\mathbb{M})$) of

$$\mathbb{M}\mathbb{F}\pi, (\mu \circ \mathbb{M}i) : \mathbb{M}\mathbb{F}\mathbb{N}A \rightarrow \mathbb{M}\mathbb{F}A.$$

Proof. (a) Clearly, \mathbb{N} is a functor $\mathcal{C} \rightarrow \mathcal{C}$ and ν and η are natural transformations

$$\nu : \mathbb{G}\mathbb{M}\mathbb{F}\mathbb{G}\mathbb{M}\mathbb{F} \Rightarrow \mathbb{G}\mathbb{M}\mathbb{F} \quad \text{and} \quad \eta : \text{Id} \Rightarrow \mathbb{G}\mathbb{M}\mathbb{F}.$$

For the monad axioms, note that

$$\begin{aligned} \nu \circ \eta &= \mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i \circ \mathbb{G}\varepsilon \circ e \\ &= \mathbb{G}(\mu \circ \varepsilon \circ i) \circ e \\ &= \text{id} \circ \mathbb{G}i \circ e \\ &= \text{id}, \end{aligned}$$

$$\begin{aligned} \nu \circ \mathbb{N}\eta &= \mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i \circ \mathbb{G}\mathbb{M}\mathbb{F}(\mathbb{G}\varepsilon \circ e) \\ &= \mathbb{G}(\mu \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\varepsilon \circ \mathbb{F}e)) \\ &= \mathbb{G}(\mu \circ \mathbb{M}(\varepsilon \circ i \circ \mathbb{F}e)) \\ &= \mathbb{G}(\mu \circ \mathbb{M}\varepsilon) \\ &= \text{id}, \\ \nu \circ \mathbb{N}\nu &= \mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i \circ \mathbb{G}\mathbb{M}\mathbb{F}(\mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i) \\ &= \mathbb{G}(\mu \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\mu \circ \mathbb{F}\mathbb{G}\mathbb{M}i)) \\ &= \mathbb{G}(\mu \circ \mathbb{M}(\mu \circ i \circ \mathbb{F}\mathbb{G}\mathbb{M}i)) \\ &= \mathbb{G}(\mu \circ \mu \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\mathbb{M}i)) \\ &= \mathbb{G}(\mu \circ \mu \circ \mathbb{M}(\mathbb{M}i \circ i)) \\ &= \mathbb{G}(\mu \circ \mathbb{M}i \circ \mu \circ \mathbb{M}i) \\ &= \mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i \circ \mathbb{G}\mu \circ \mathbb{G}\mathbb{M}i \\ &= \nu \circ \nu. \end{aligned}$$

(b) Let $\mathfrak{A} = \langle A, \pi \rangle$ be an \mathbb{M} -algebra. To see that $\widehat{\mathbb{G}}\mathfrak{A}$ is an \mathbb{N} -algebra, note that

$$\begin{aligned} \mathbb{G}(\pi \circ \mathbb{M}i) \circ \eta &= \mathbb{G}(\pi \circ \mathbb{M}i \circ \varepsilon) \circ e \\ &= \mathbb{G}(\pi \circ \varepsilon \circ i) \circ e \\ &= \mathbb{G}i \circ e \\ &= \text{id}, \\ \mathbb{G}(\pi \circ \mathbb{M}i) \circ \nu &= \mathbb{G}(\pi \circ \mathbb{M}i \circ \mu \circ \mathbb{M}i) \\ &= \mathbb{G}(\pi \circ \mu \circ \mathbb{M}\mathbb{M}i \circ \mathbb{M}i) \\ &= \mathbb{G}(\pi \circ \mathbb{M}\pi \circ \mathbb{M}\mathbb{M}i \circ \mathbb{M}i) \\ &= \mathbb{G}(\pi \circ \mathbb{M}((\pi \circ \mathbb{M}i) \circ i)) \\ &= \mathbb{G}(\pi \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}(\pi \circ \mathbb{M}i))) \\ &= \mathbb{G}(\pi \circ \mathbb{M}i) \circ \mathbb{N}\mathbb{G}(\pi \circ \mathbb{M}i). \end{aligned}$$

If $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ is a morphism of \mathbb{M} -algebras, then $\widehat{\mathbb{G}}\varphi : \widehat{\mathfrak{A}} \rightarrow \widehat{\mathfrak{B}}$ is a morphism

of \mathbb{N} -algebras since

$$\begin{aligned} \mathbb{G}\varphi \circ \mathbb{G}(\pi \circ \mathbb{M}i) &= \mathbb{G}(\pi \circ \mathbb{M}\varphi \circ \mathbb{M}i) \\ &= \mathbb{G}(\pi \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\varphi)) = \mathbb{G}(\pi \circ \mathbb{M}i) \circ \mathbb{N}\mathbb{G}\varphi. \end{aligned}$$

(c) The idea of the proof is as follows. When constructing $\Omega := \widehat{\mathbb{F}}\mathcal{A}$ from \mathcal{A} we have to transfer the product $\pi : \mathbb{N}A \rightarrow A$ to a function $\mathbb{M}Q \rightarrow Q$. A first try would be to set $Q := \mathbb{F}A$ and to use some correspondence between elements of A and elements of Q to transfer the definition of π from \mathcal{A} to Ω . But this does not work since not every term $s \in \mathbb{M}FA$ corresponds to some $t \in \mathbb{N}A$. And for such elements we do not know how to choose the value $\pi(s)$. Our solution is to simply leave such terms unevaluated. That is, we set $\pi(s) := s$. Of course, to do so we have to also add these terms as elements to our algebra. This leads to an algebra with universe $\mathbb{M}FA$ where multiplication is just the monad multiplication μ . But doing so is not compatible with the original product π . As a final step we therefore have to take a quotient that identifies terms $s, t \in \mathbb{M}FA$ where the products of the corresponding terms in $\mathbb{N}A$ coincide. Thus, we arrive at our final definition, we take for Ω a suitable quotient of $\mathbb{M}FA$.

The formal definition is as follows. Let $\rho : \mathbb{M}FA \rightarrow \Omega$ with $\Omega = \langle Q, \sigma \rangle$ be the coequaliser (in $\text{Alg}(\mathbb{M})$) of

$$\mathbb{M}F\pi, \mu \circ \mathbb{M}i : \mathbb{M}FNA \rightarrow \mathbb{M}FA.$$

(To see that this is well-defined, note that $\mathbb{M}f \circ \mu = \mu \circ \mathbb{M}\mathbb{M}f$, for every function f . Therefore, $\mathbb{M}F\pi$ and $\mathbb{M}i$ are \mathbb{M} -morphisms.) We set $\widehat{\mathbb{F}}\mathcal{A} := \Omega$.

$$\begin{array}{ccccc} \mathbb{M}FNA & \xrightarrow[\mu \circ \mathbb{M}i]{\mathbb{M}F\pi} & \mathbb{M}FA & \xrightarrow{\rho} & Q \\ & \searrow \mathbb{M}i & \uparrow \mu & & \uparrow \sigma \\ & & \mathbb{M}\mathbb{M}FA & \xrightarrow{\mathbb{M}\rho} & \mathbb{M}Q \end{array}$$

To define the action of $\widehat{\mathbb{F}}$ on morphisms, let $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ be an \mathbb{N} -morphism. Then φ induces morphisms between the coequaliser diagrams associated with $\widehat{\mathbb{F}}\mathcal{A}$ and $\widehat{\mathbb{F}}\mathcal{B}$. (This is in fact the definition of $\widehat{\mathbb{F}}\varphi$.)

$$\begin{array}{ccccc} \mathbb{M}FNA & \xrightarrow[\mu \circ \mathbb{M}i]{\mathbb{M}F\pi} & \mathbb{M}FA & \xrightarrow{\rho_A} & Q_A \\ \mathbb{M}F\mathbb{N}\varphi \downarrow & & \downarrow \mathbb{M}F\varphi & & \downarrow \widehat{\mathbb{F}}\varphi \\ \mathbb{M}FNB & \xrightarrow[\mu \circ \mathbb{M}i]{\mathbb{M}F\pi} & \mathbb{M}FB & \xrightarrow{\rho_B} & Q_B \end{array}$$

Note that this immediately implies that $\rho : \mathbb{M}F \Rightarrow \widehat{\mathbb{F}}$ is a natural transformation since, by definition of $\widehat{\mathbb{F}}\varphi$, we have

$$\widehat{\mathbb{F}}\varphi \circ \rho = \rho \circ \mathbb{M}F\varphi.$$

Furthermore, every component ρ_A of ρ is an epimorphism (in $\text{Alg}(\mathbb{M})$): given \mathbb{M} -morphisms $f, g : \Omega \rightarrow \mathbb{C}$ with $f \circ \rho = g \circ \rho$, we can apply the universality of ρ to the morphism $\psi := f \circ \rho$, to obtain a unique morphism $h : \Omega \rightarrow \mathbb{C}$ such that $h \circ \rho = \psi$. This implies that $f = h = g$.

Having defined $\widehat{\mathbb{F}}$, it remains to show that it is the left adjoint of $\widehat{\mathbb{G}}$. We define the unit $\tau : \text{Id} \Rightarrow \widehat{\mathbb{G}}\widehat{\mathbb{F}}$ of the adjunction by $\tau := \mathbb{G}\rho \circ \eta$. Note that τ is an \mathbb{N} -morphism $\mathcal{A} \rightarrow \widehat{\mathbb{G}}\Omega$ since

$$\begin{aligned} \tau \circ \pi &= \mathbb{G}\rho \circ \eta \circ \pi \\ &= \mathbb{G}\rho \circ \mathbb{G}\varepsilon \circ e \circ \pi \\ &= \mathbb{G}(\rho \circ \varepsilon \circ \mathbb{F}\pi) \circ e \\ &= \mathbb{G}(\rho \circ \mathbb{M}F\pi \circ \varepsilon) \circ e \\ &= \mathbb{G}(\rho \circ \mu \circ \mathbb{M}i \circ \varepsilon) \circ e \\ &= \mathbb{G}\rho \circ v \circ \eta \\ &= \mathbb{G}\rho \\ &= \mathbb{G}(\rho \circ \mu \circ \mathbb{M}\varepsilon) \\ &= \mathbb{G}(\sigma \circ \mathbb{M}\rho \circ \mathbb{M}(\varepsilon \circ \text{id})) \end{aligned}$$

$$\begin{aligned}
&= \mathbb{G}(\sigma \circ \mathbb{M}((\rho \circ \varepsilon) \circ i \circ \mathbb{F}e)) \\
&= \mathbb{G}(\sigma \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}(\rho \circ \varepsilon) \circ \mathbb{F}e)) \\
&= \mathbb{G}(\sigma \circ \mathbb{M}i) \circ \mathbb{G}\mathbb{M}\mathbb{F}(\mathbb{G}\rho \circ \eta) \\
&= \mathbb{G}(\sigma \circ \mathbb{M}i) \circ \mathbb{N}\tau.
\end{aligned}$$

Furthermore, τ is natural in \mathfrak{A} since, given an \mathbb{M} -morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$,

$$\begin{aligned}
\tau \circ \varphi &= \mathbb{G}\rho \circ \eta \circ \varphi \\
&= \mathbb{G}\rho \circ \mathbb{N}\varphi \circ \eta \\
&= \mathbb{G}(\rho \circ \mathbb{M}\mathbb{F}\varphi) \circ \eta \\
&= \mathbb{G}(\mathbb{F}\varphi \circ \rho) \circ \eta \\
&= \mathbb{G}\mathbb{F}\varphi \circ \tau,
\end{aligned}$$

where the fourth step follows from the fact that $\rho : \mathbb{M}\mathbb{F} \Rightarrow \widehat{\mathbb{F}}$ is a natural transformation.

To define the counit, let $\mathfrak{A} = \langle A, \pi \rangle$ be an \mathbb{M} -algebra and let $\rho = \rho_{\mathbb{G}\mathfrak{A}} : \mathbb{M}\mathbb{F}\mathfrak{A} \rightarrow \mathfrak{A}$ be the coequaliser used in the definition of $\widehat{\mathbb{F}}(\widehat{\mathbb{G}}\mathfrak{A})$. Since

$$\begin{aligned}
(\pi \circ \mathbb{M}i) \circ \mathbb{M}\mathbb{F}(\mathbb{G}(\pi \circ \mathbb{M}i)) &= \pi \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}(\pi \circ \mathbb{M}i)) \\
&= \pi \circ \mathbb{M}(\pi \circ \mathbb{M}i \circ i) \\
&= \pi \circ \mu \circ \mathbb{M}\mathbb{M}i \circ \mathbb{M}i \\
&= \pi \circ \mathbb{M}i \circ \mu \circ \mathbb{M}i \\
&= (\pi \circ \mathbb{M}i) \circ (\mu \circ \mathbb{M}i)
\end{aligned}$$

and $\pi \circ \mathbb{M}i$ is an \mathbb{M} -morphism, we can use universality of ρ to find a unique \mathbb{M} -morphism $v : \widehat{\mathbb{F}}\widehat{\mathbb{G}}\mathfrak{A} \rightarrow \mathfrak{A}$ with

$$v \circ \rho = \pi \circ \mathbb{M}i.$$

To prove that $\widehat{\mathbb{F}}$ and $\widehat{\mathbb{G}}$ form an adjunction, it is now sufficient to show

that $v \circ \widehat{\mathbb{F}}\tau = \text{id}$ and $\widehat{\mathbb{G}}v \circ \tau = \text{id}$. For the second equation, we have

$$\begin{aligned}
\widehat{\mathbb{G}}v \circ \tau &= \mathbb{G}v \circ \mathbb{G}\rho \circ \eta \\
&= \mathbb{G}(\pi \circ \mathbb{M}i) \circ \mathbb{G}\varepsilon \circ e \\
&= \mathbb{G}(\pi \circ \varepsilon \circ i) \circ e \\
&= \mathbb{G}i \circ e \\
&= \text{id}.
\end{aligned}$$

For the first one, we have

$$\begin{aligned}
v \circ \widehat{\mathbb{F}}\tau \circ \rho &= v \circ \rho \circ \mathbb{M}\mathbb{F}\tau \\
&= \sigma \circ \mathbb{M}i \circ \mathbb{M}\mathbb{F}\tau \\
&= \sigma \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\rho \circ \mathbb{F}\eta) \\
&= \sigma \circ \mathbb{M}(\rho \circ i \circ \mathbb{F}\eta) \\
&= \rho \circ \mu \circ \mathbb{M}(i \circ \mathbb{F}\eta) \\
&= \rho \circ \mu \circ \mathbb{M}(i \circ \mathbb{F}\mathbb{G}\varepsilon \circ \mathbb{F}e) \\
&= \rho \circ \mu \circ \mathbb{M}(\varepsilon \circ i \circ \mathbb{F}e) \\
&= \rho \circ \mu \circ \mathbb{M}(\varepsilon \circ \text{id}) \\
&= \rho \circ \text{id} \\
&= \rho.
\end{aligned}$$

Since ρ is an epimorphism in $\text{Alg}(\mathbb{M})$ and $v \circ \widehat{\mathbb{F}}\tau$ is an \mathbb{M} -morphism, the claim follows. \square

Distributive Laws

Our second construction provides a way to combine two monads into a single one. This is needed when we want to expand an \mathbb{M} -algebra with operations provided by a second monad \mathbb{N} . An equivalent way of looking at such an expansion is by finding a lift of \mathbb{N} from the base category to $\text{Alg}(\mathbb{M})$. It turns out that, in order for this to work, the two monads \mathbb{M} and \mathbb{N} need to be compatible: there needs to be what is called a *distributive law* between them.

Definition 6.2. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$ be monads. A natural transformation $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ is a *distributive law* if

$$\begin{aligned} \delta \circ \mu &= \mathbb{N}\mu \circ \delta \circ \mathbb{M}\delta, & \delta \circ \varepsilon &= \mathbb{N}\varepsilon, \\ \delta \circ \mathbb{M}\nu &= \nu \circ \mathbb{N}\delta \circ \delta, & \delta \circ \mathbb{M}\eta &= \eta. \end{aligned}$$

$$\begin{array}{ccccc} \mathbb{M}\mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\mathbb{M}\delta} & \mathbb{M}\mathbb{N}\mathbb{M}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{M}\mathbb{A} \\ \mu \downarrow & & & & \downarrow \mathbb{N}\mu \\ \mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{A} & & \\ \mathbb{M}\mathbb{N}\mathbb{N}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\mathbb{N}\delta} & \mathbb{N}\mathbb{N}\mathbb{M}\mathbb{A} \\ \mathbb{M}\nu \downarrow & & & & \downarrow \nu \\ \mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{A} & & \end{array} \quad \begin{array}{ccc} \mathbb{N}\mathbb{A} & \xrightarrow{\mathbb{N}\varepsilon} & \mathbb{N}\mathbb{M}\mathbb{A} \\ \varepsilon \downarrow & & \downarrow \delta \\ \mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{A} \\ \mathbb{M}\mathbb{A} & \xrightarrow{\mathbb{M}\eta} & \mathbb{N}\mathbb{M}\mathbb{A} \\ \mathbb{M}\eta \downarrow & & \downarrow \delta \\ \mathbb{M}\mathbb{N}\mathbb{A} & \xrightarrow{\delta} & \mathbb{N}\mathbb{M}\mathbb{A} \end{array}$$

Example. Let $\mathbb{M}\mathbb{A} := A^*$ be the monad for finite words and $\mathbb{A}\mathbb{A} := A^*/\approx$ its quotient by the relation

$$u \approx v \quad : \text{iff} \quad u \text{ is a permutation of } v.$$

Then \mathbb{M} -algebras are monoids and \mathbb{A} -algebras commutative monoids. We will write elements of $\mathbb{A}\mathbb{A}$ as sums $a + \dots + b$ and those of $\mathbb{M}\mathbb{A}$ as products. A distributive law $\delta : \mathbb{M}\mathbb{A} \Rightarrow \mathbb{A}\mathbb{M}$ maps a product of sums

$$\sum_{i < n_0} a_{oi} \cdot \sum_{i < n_1} a_{1i} \cdots \sum_{i < n_{k-1}} a_{ki}$$

to the sum

$$\sum_{\sigma} a_{o, \sigma(o)} \cdots a_{k-1, \sigma(k-1)}$$

that ranges over all functions $\sigma \in [n_0] \times \dots \times [n_{k-1}]$. For instance,

$$\delta((a+b)(c+d)) = ac + ad + bc + bd.$$

The above axioms for a distributive law are not always the most convenient to work with. In the following we will present several characterisations of when a distributive law between two monads exists. One of them tells us that a distributive law is precisely what is needed to lift a monad \mathbb{N} from the base category to the category of \mathbb{M} -algebras.

Definition 6.3. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$ be monads on some category \mathcal{C} and let $\mathbb{U} : \text{Alg}(\mathbb{M}) \rightarrow \mathcal{C}$ be the forgetful functor mapping an \mathbb{M} -algebra to its universe.

(a) We say that a monad $\langle \hat{\mathbb{N}}, \hat{\nu}, \hat{\eta} \rangle$ is a *lift* of \mathbb{N} to the category of \mathbb{M} -algebras if

$$\mathbb{U} \circ \hat{\mathbb{N}} = \mathbb{N} \circ \mathbb{U}, \quad \mathbb{U}\hat{\nu} = \nu, \quad \mathbb{U}\hat{\eta} = \eta.$$

(b) The *Kleisli category* $\text{Free}(\mathbb{N})$ of \mathbb{N} is the full subcategory of $\text{Alg}(\mathbb{N})$ induced by all free \mathbb{N} -algebras. The *free functor* $\mathbb{F}_{\mathbb{N}} : \mathcal{C} \rightarrow \text{Free}(\mathbb{N})$ maps an object $C \in \mathcal{C}$ to the free \mathbb{N} -algebra generated by C , that is,

$$\begin{aligned} \mathbb{F}_{\mathbb{N}}C &:= \langle \mathbb{N}C, \nu \rangle, & \text{for objects } C \in \mathcal{C}, \\ \mathbb{F}_{\mathbb{N}}\varphi &:= \mathbb{N}\varphi, & \text{for } \mathcal{C}\text{-morphisms } \varphi : A \rightarrow B. \end{aligned}$$

(c) An *extension* of \mathbb{M} to $\text{Free}(\mathbb{N})$ is a monad $\langle \hat{\mathbb{M}}, \hat{\mu}, \hat{\varepsilon} \rangle$ on $\text{Free}(\mathbb{N})$ satisfying

$$\hat{\mathbb{M}} \circ \mathbb{F}_{\mathbb{N}} = \mathbb{F}_{\mathbb{N}} \circ \mathbb{M}, \quad \hat{\mu} = \mathbb{F}_{\mathbb{N}}\mu, \quad \hat{\varepsilon} = \mathbb{F}_{\mathbb{N}}\varepsilon.$$

Theorem 6.4 (Beck). Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$ be monads on the category \mathcal{C} . There exist bijections between the following objects:

- (1) distributive laws $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$;
- (2) liftings $\hat{\mathbb{N}}$ of \mathbb{N} to the category of \mathbb{M} -algebras;
- (3) extensions $\hat{\mathbb{M}}$ of \mathbb{M} to the Kleisli category $\text{Free}(\mathbb{N})$;
- (4) functions κ such that
 - (M1) $\langle \mathbb{N}\mathbb{M}, \kappa, \eta \circ \varepsilon \rangle$ is a monad,

(M2) the functions $\mathbb{N}\varepsilon$ and η induce morphisms of monads $\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ and $\mathbb{M} \Rightarrow \mathbb{N}\mathbb{M}$,

(M3) κ satisfies the middle unit law: $\kappa \circ \mathbb{N}(\varepsilon \circ \eta) = \text{id}$;

(5) functions κ such that

(C1) $\langle \mathbb{N}\mathbb{M}, \kappa, \eta \circ \varepsilon \rangle$ is a monad,

(C2) $\kappa \circ \mathbb{N}\mathbb{M}\eta = \mathbb{M}\mu$,

(C3) $\kappa \circ \mathbb{N}\varepsilon = \nu$,

(C4) $\nu \circ \mathbb{N}\kappa = \kappa \circ \nu$,

(C5) $\mathbb{N}\mu \circ \kappa = \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mu$.

Proof. (1) \Rightarrow (3) Let δ be a distributive law. We set

$$\widehat{\mathbb{M}}\langle \mathbb{N}C, \nu \rangle := \langle \mathbb{N}\mathbb{M}C, \nu \rangle, \quad \text{for an algebra } \langle \mathbb{N}C, \nu \rangle \in \text{Free}(\mathbb{N}),$$

$$\widehat{\mathbb{M}}\varphi := \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta), \quad \text{for a morphism } \varphi : \mathcal{A} \rightarrow \mathcal{B}.$$

We start by checking that $\widehat{\mathbb{M}}$ is a functor. Clearly, $\widehat{\mathbb{M}}\mathcal{A} \in \text{Free}(\mathbb{N})$, for every $\mathcal{A} \in \text{Free}(\mathbb{N})$. Consider a morphism $\varphi : \mathcal{A} \rightarrow \mathcal{B}$. Then $\widehat{\mathbb{M}}\varphi : \widehat{\mathbb{M}}\mathcal{A} \rightarrow \widehat{\mathbb{M}}\mathcal{B}$ is also a morphism since

$$\begin{aligned} \widehat{\mathbb{M}}\varphi \circ \nu &= \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \circ \nu \\ &= \nu \circ \nu \circ \mathbb{N}\mathbb{N}\delta \circ \mathbb{N}\mathbb{N}\mathbb{M}(\varphi \circ \eta) \\ &= \nu \circ \mathbb{N}\nu \circ \mathbb{N}\mathbb{N}\delta \circ \mathbb{N}\mathbb{N}\mathbb{M}(\varphi \circ \eta) \\ &= \nu \circ \mathbb{N}(\nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta)) = \nu \circ \mathbb{N}\widehat{\mathbb{M}}\varphi. \end{aligned}$$

Next, we show that $\widehat{\mathbb{M}}$ is an extension of \mathbb{M} , that is, that $\widehat{\mathbb{M}} \circ \mathbb{F}_{\mathbb{N}} = \mathbb{F}_{\mathbb{N}} \circ \mathbb{M}$. For an object $C \in \mathcal{C}$, we have

$$\widehat{\mathbb{M}}\mathbb{F}_{\mathbb{N}}C = \widehat{\mathbb{M}}\langle \mathbb{N}C, \nu \rangle = \langle \mathbb{N}\mathbb{M}C, \nu \rangle = \mathbb{F}_{\mathbb{N}}\mathbb{M}C.$$

Similarly, for a morphism $\varphi : \mathcal{A} \rightarrow \mathcal{B}$, we have

$$\begin{aligned} \widehat{\mathbb{M}}\mathbb{F}_{\mathbb{N}}\varphi &= \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\mathbb{N}\varphi \circ \eta) \\ &= \nu \circ \mathbb{N}(\delta \circ \mathbb{M}\mathbb{N}\varphi \circ \mathbb{M}\eta) \\ &= \nu \circ \mathbb{N}\mathbb{N}\mathbb{M}\varphi \circ \mathbb{N}(\delta \circ \mathbb{M}\eta) \\ &= \mathbb{N}\mathbb{M}\varphi \circ \nu \circ \mathbb{N}\eta = \mathbb{N}\mathbb{M}\varphi = \mathbb{F}_{\mathbb{N}}\mathbb{M}\varphi. \end{aligned}$$

Hence, it remains to show that $\langle \widehat{\mathbb{M}}, \mathbb{F}_{\mathbb{N}}\mu, \mathbb{F}_{\mathbb{N}}\varepsilon \rangle$ is a monad.

$$\begin{aligned} \mathbb{F}_{\mathbb{N}}\mu \circ \widehat{\mathbb{M}}\mathbb{F}_{\mathbb{N}}\mu &= \mathbb{N}\mu \circ \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\mathbb{N}\mu \circ \eta) \\ &= \mathbb{N}\mu \circ \nu \circ \mathbb{N}\mathbb{N}\mathbb{M}\mu \circ \mathbb{N}(\delta \circ \mathbb{M}\eta) \\ &= \mathbb{N}\mu \circ \mathbb{N}\mathbb{M}\mu \circ \nu \circ \mathbb{N}\eta \\ &= \mathbb{N}\mu \circ \mathbb{N}\mu \circ \text{id} \\ &= \mathbb{F}_{\mathbb{N}}\mu \circ \mathbb{F}_{\mathbb{N}}\mu, \end{aligned}$$

$$\mathbb{F}_{\mathbb{N}}\mu \circ \mathbb{F}_{\mathbb{N}}\varepsilon = \mathbb{N}\mu \circ \mathbb{N}\varepsilon = \mathbb{N}\text{id} = \text{id},$$

$$\begin{aligned} \mathbb{F}_{\mathbb{N}}\mu \circ \widehat{\mathbb{M}}\mathbb{F}_{\mathbb{N}}\varepsilon &= \mathbb{N}\mu \circ \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\mathbb{N}\varepsilon \circ \eta) \\ &= \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\mathbb{N}\mathbb{M}\varepsilon \circ \mathbb{N}(\delta \circ \mathbb{M}\eta) \\ &= \nu \circ \mathbb{N}\mathbb{N}\text{id} \circ \mathbb{N}\eta \\ &= \text{id}. \end{aligned}$$

(3) \Rightarrow (1) Given an extension $\widehat{\mathbb{M}}$ of \mathbb{M} to $\text{Free}(\mathbb{N})$, we set

$$\delta := \mathbb{U}\widehat{\mathbb{M}}\nu \circ \eta.$$

(Note that $\widehat{\mathbb{M}}\nu : \mathbb{N}\mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$, so this is well-typed.) To simplify notation we will identify in the following morphisms of \mathbb{N} -algebras with their images under the forgetful functor \mathbb{U} . That is, we will omit all occurrences of \mathbb{U} and we replace $\mathbb{F}_{\mathbb{N}}$ by \mathbb{N} . Then the axioms for an extension $\widehat{\mathbb{M}}$ take the form

$$\widehat{\mathbb{M}} \circ \mathbb{N} = \mathbb{N} \circ \mathbb{M}, \quad \hat{\mu} = \mathbb{N}\mu, \quad \hat{\varepsilon} = \mathbb{N}\varepsilon.$$

Since $\widehat{\mathbb{M}}\varphi$ is an \mathbb{N} -algebra morphism between free algebras whose product function is v , we also have the equation

$$v \circ \mathbb{N}\widehat{\mathbb{M}}\varphi = \widehat{\mathbb{M}}\varphi \circ v, \quad \text{for every morphism } \varphi : \mathbb{F}_{\mathbb{N}}A \rightarrow \mathbb{F}_{\mathbb{N}}B.$$

To see that δ is a natural transformation, consider a morphism $f : A \rightarrow B$. Then

$$\begin{aligned} \delta \circ \mathbb{M}Nf &= \widehat{\mathbb{M}}v \circ \eta \circ \mathbb{M}Nf \\ &= \widehat{\mathbb{M}}v \circ \mathbb{N}Mf \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}Nf \circ \eta \\ &= \widehat{\mathbb{M}}(v \circ \mathbb{N}Nf) \circ \eta \\ &= \widehat{\mathbb{M}}(\mathbb{N}f \circ v) \circ \eta \\ &= \widehat{\mathbb{M}}Nf \circ \widehat{\mathbb{M}}v \circ \eta = \mathbb{N}Mf \circ \delta. \end{aligned}$$

Furthermore, δ is a distributive law since

$$\begin{aligned} \delta \circ \mu &= \widehat{\mathbb{M}}v \circ \eta \circ \mu \\ &= \widehat{\mathbb{M}}v \circ \mathbb{N}\mu \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \hat{\mu} \circ \eta \\ &= \hat{\mu} \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}v \circ \eta \\ &= \mathbb{N}\mu \circ \widehat{\mathbb{M}}(\widehat{\mathbb{M}}v \circ v \circ \mathbb{N}\eta) \circ \eta \\ &= \mathbb{N}\mu \circ \widehat{\mathbb{M}}(v \circ \mathbb{N}\widehat{\mathbb{M}}v \circ \mathbb{N}\eta) \circ \eta \\ &= \mathbb{N}\mu \circ \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}N(\widehat{\mathbb{M}}v \circ \eta) \circ \eta \\ &= \mathbb{N}\mu \circ \widehat{\mathbb{M}}v \circ \mathbb{N}M(\widehat{\mathbb{M}}v \circ \eta) \circ \eta \\ &= \mathbb{N}\mu \circ \widehat{\mathbb{M}}v \circ \eta \circ M(\widehat{\mathbb{M}}v \circ \eta) = \mathbb{N}\mu \circ \delta \circ M\delta \end{aligned}$$

$$\begin{aligned} \delta \circ \mathbb{M}v &= \widehat{\mathbb{M}}v \circ \eta \circ \mathbb{M}v \\ &= \widehat{\mathbb{M}}v \circ \mathbb{N}Mv \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}Nv \circ \eta \\ &= \widehat{\mathbb{M}}(v \circ \mathbb{N}v) \circ \eta \\ &= \widehat{\mathbb{M}}(v \circ v) \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}v \circ \eta \\ &= \widehat{\mathbb{M}}v \circ v \circ \mathbb{N}\eta \circ \widehat{\mathbb{M}}v \circ \eta \\ &= v \circ \mathbb{N}\widehat{\mathbb{M}}v \circ \mathbb{N}\eta \circ \widehat{\mathbb{M}}v \circ \eta \\ &= v \circ \mathbb{N}(\widehat{\mathbb{M}}v \circ \eta) \circ \widehat{\mathbb{M}}v \circ \eta = v \circ \mathbb{N}\delta \circ \delta, \end{aligned}$$

$$\begin{aligned} \delta \circ \varepsilon &= \widehat{\mathbb{M}}v \circ \eta \circ \varepsilon \\ &= \widehat{\mathbb{M}}v \circ \mathbb{N}\varepsilon \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \hat{\varepsilon} \circ \eta = \hat{\varepsilon} \circ v \circ \eta = \mathbb{N}\varepsilon, \end{aligned}$$

$$\begin{aligned} \delta \circ \mathbb{M}\eta &= \widehat{\mathbb{M}}v \circ \eta \circ \mathbb{M}\eta \\ &= \widehat{\mathbb{M}}v \circ \mathbb{N}M\eta \circ \eta \\ &= \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}N\eta \circ \eta \\ &= \widehat{\mathbb{M}}(v \circ \mathbb{N}\eta) \circ \eta = \widehat{\mathbb{M}}\text{id} \circ \eta = \eta. \end{aligned}$$

(1) \Rightarrow (2) Let $\delta : \mathbb{M}N \Rightarrow \mathbb{N}M$ be a distributive law. Given an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$, we set $\hat{\mathfrak{N}}\mathfrak{A} := \langle \mathbb{N}A, \hat{\pi} \rangle$ with $\hat{\pi} := \mathbb{N}\pi \circ \delta : \mathbb{M}\mathbb{N}A \rightarrow \mathbb{N}A$. For a morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$, we set $\hat{\mathfrak{N}}\varphi := \mathbb{N}\varphi$.

We start by showing that this defines a functor $\text{Alg}(\mathbb{M}) \rightarrow \text{Alg}(\mathbb{M})$. If

$\mathcal{A} = \langle A, \pi \rangle$ is an \mathbb{M} -algebra, then so is $\hat{\mathbb{N}}\mathcal{A}$ since

$$\begin{aligned} \hat{\pi} \circ \mathbb{M}\hat{\pi} &= \mathbb{N}\pi \circ \delta \circ \mathbb{M}(\mathbb{N}\pi \circ \delta) & \hat{\pi} \circ \epsilon &= \mathbb{N}\pi \circ \delta \circ \epsilon \\ &= \mathbb{N}\pi \circ \mathbb{N}\mathbb{M}\pi \circ \delta \circ \mathbb{M}\delta & &= \mathbb{N}\pi \circ \mathbb{N}\epsilon \\ &= \mathbb{N}\pi \circ \mathbb{N}\mu \circ \delta \circ \mathbb{M}\delta & &= \text{id}. \\ &= \mathbb{N}\pi \circ \delta \circ \mu & & \\ &= \hat{\pi} \circ \mu, & & \end{aligned}$$

Furthermore, for morphism $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ of \mathbb{M} -algebras, we have

$$\begin{aligned} \hat{\mathbb{N}}\varphi \circ \mathbb{M}\hat{\pi} &= \mathbb{N}\varphi \circ \mathbb{N}\pi \circ \delta \\ &= \mathbb{N}(\pi \circ \mathbb{M}\varphi) \circ \delta \\ &= \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi = \hat{\pi} \circ \mathbb{M}\hat{\mathbb{N}}\varphi. \end{aligned}$$

It now follows that $\hat{\mathbb{N}}$ forms a monad with multiplication $\hat{v} := \hat{\mathbb{N}}v$ and unit map $\hat{\eta} := \hat{\mathbb{N}}\eta$: we have just shown that \hat{v} and $\hat{\eta}$ are morphisms of \mathbb{M} -algebras and the equations for the monad laws immediately follow from those for v and η .

Finally, to show that $\hat{\mathbb{N}}$ is a lift of \mathbb{N} it is sufficient to note that, by definition,

$$\mathbb{U} \circ \hat{\mathbb{N}} = \mathbb{N} \circ \mathbb{U}, \quad \mathbb{U}\hat{v} = v, \quad \mathbb{U}\hat{\eta} = \eta.$$

(2) \Rightarrow (4) Given a set A , we lift the free algebra $\mathcal{F}_A = \langle \mathbb{M}A, \mu \rangle$ to $\hat{\mathbb{N}}\mathcal{F}_A$ and $\hat{\mathbb{N}}\hat{\mathbb{N}}\mathcal{F}_A$. The products of these algebras are morphisms

$$\hat{\mu} : \mathbb{M}\mathbb{N}\mathbb{M}A \rightarrow \mathbb{N}\mathbb{M}A \quad \text{and} \quad \hat{\hat{\mu}} : \mathbb{M}\mathbb{N}\mathbb{N}\mathbb{M}A \rightarrow \mathbb{N}\mathbb{N}\mathbb{M}A.$$

We claim that

$$\kappa := v \circ \mathbb{N}\hat{\mu}$$

is the desired morphism.

First, note that, $\hat{\mu}$ being the product of an \mathbb{M} -algebra, we have

$$\hat{\mu} \circ \mathbb{M}\hat{\mu} = \hat{\mu} \circ \mu \quad \text{and} \quad \hat{\mu} \circ \epsilon = \text{id}.$$

The first of these two equations implies that $\hat{\mu}$ is a morphism of \mathbb{M} -algebras. Consequently, so is $\hat{\mathbb{N}}\hat{\mu} = \mathbb{N}\hat{\mu}$ and we have

$$\mathbb{N}\hat{\mu} \circ \hat{\mu} = \hat{\mu} \circ \mathbb{M}\mathbb{N}\hat{\mu}.$$

Furthermore, by definition of a lift, the underlying morphisms of \hat{v} and $\hat{\eta}$ are, respectively, v and η . As $\hat{v} : \hat{\mathbb{N}}\hat{\mathbb{N}}\mathbb{M}A \rightarrow \hat{\mathbb{N}}\mathbb{M}A$ and $\hat{\eta} : \mathbb{M}A \rightarrow \hat{\mathbb{N}}\mathbb{M}A$ are morphisms of \mathbb{M} -algebras, it follows that

$$\hat{\mu} \circ \mathbb{M}v = v \circ \hat{\mu} \quad \text{and} \quad \hat{\mu} \circ \mathbb{M}\eta = \eta \circ \mu.$$

For (M1), we have to show that κ is a natural transformation $\mathbb{M}\mathbb{N}\mathbb{M} \Rightarrow \mathbb{N}\mathbb{M}$ and that κ and $\eta \circ \epsilon$ satisfy the three monad laws. For the former, let $\varphi : A \rightarrow B$ be a morphism. Then $\mathbb{M}\varphi$ is a morphism $\mathcal{F}_A \rightarrow \mathcal{F}_B$ and, therefore, $\hat{\mathbb{N}}\mathbb{M}\varphi = \mathbb{N}\mathbb{M}\varphi$ is one of $\mathcal{F}_A \rightarrow \mathcal{F}_B$. This implies that $\hat{\mu} \circ \mathbb{M}\mathbb{N}\mathbb{M}\varphi = \mathbb{N}\mathbb{M}\varphi \circ \hat{\mu}$. Thus, $\hat{\mu}$ is a natural transformation $\mathbb{M}\mathbb{N}\mathbb{M} \Rightarrow \mathbb{N}\mathbb{M}$. Since $v : \mathbb{N}\mathbb{N} \Rightarrow \mathbb{N}$ is also a natural transformation it follows that so is the composition $\kappa = v \circ \mathbb{N}\hat{\mu}$.

It remains to check the monad laws.

$$\begin{aligned} \kappa \circ \kappa &= v \circ \mathbb{N}\hat{\mu} \circ v \circ \mathbb{N}\hat{\mu} \\ &= v \circ v \circ \mathbb{N}\mathbb{N}\hat{\mu} \circ \mathbb{N}\hat{\mu} \\ &= v \circ \mathbb{N}v \circ \mathbb{N}\hat{\mu} \circ \mathbb{M}\mathbb{N}\hat{\mu} \\ &= v \circ \mathbb{N}\hat{\mu} \circ \mathbb{N}\mathbb{M}v \circ \mathbb{M}\mathbb{N}\hat{\mu} \\ &= \kappa \circ \mathbb{N}\mathbb{M}\kappa, \end{aligned}$$

$$\kappa \circ \eta \circ \epsilon = v \circ \mathbb{N}\hat{\mu} \circ \mathbb{N}\epsilon \circ \eta = v \circ \eta = \text{id},$$

$$\begin{aligned} \kappa \circ \mathbb{N}\mathbb{M}(\eta \circ \epsilon) &= v \circ \mathbb{N}\hat{\mu} \circ \mathbb{N}\mathbb{M}(\eta \circ \epsilon) \\ &= v \circ \mathbb{N}(\hat{\mu} \circ \mathbb{M}\eta) \circ \mathbb{N}\mathbb{M}\epsilon \\ &= v \circ \mathbb{N}(\eta \circ \mu) \circ \mathbb{N}\mathbb{M}\epsilon \\ &= \text{id}. \end{aligned}$$

For (M2), we have

$$\begin{aligned}\kappa \circ \mathbb{N}\varepsilon \circ \mathbb{N}\mathbb{N}\varepsilon &= \nu \circ \mathbb{N}\hat{\mu} \circ \mathbb{N}\varepsilon \circ \mathbb{N}\mathbb{N}\varepsilon = \nu \circ \mathbb{N}\mathbb{N}\varepsilon = \mathbb{N}\varepsilon \circ \nu, \\ \kappa \circ \eta \circ \mathbb{M}\eta &= \nu \circ \mathbb{N}\hat{\mu} \circ \eta \circ \mathbb{M}\eta = \nu \circ \eta \circ \hat{\mu} \circ \mathbb{M}\eta = \hat{\mu} \circ \mathbb{M}\eta = \eta \circ \mu.\end{aligned}$$

For (M3), we have

$$\kappa \circ \mathbb{N}(\varepsilon \circ \eta) = \nu \circ \mathbb{N}\hat{\mu} \circ \mathbb{N}(\varepsilon \circ \eta) = \nu \circ \mathbb{N}\eta = \text{id}.$$

(4) \Rightarrow (5) Suppose that κ satisfies (M1)–(M3). We claim that it also satisfies (C1)–(C5). (C1) is the same as (M1). For (C3), we have

$$\begin{aligned}\kappa \circ \mathbb{N}\varepsilon &= \kappa \circ \mathbb{N}(\varepsilon \circ \kappa \circ \mathbb{N}(\varepsilon \circ \eta)) \\ &= \kappa \circ \mathbb{N}\mathbb{M}\kappa \circ \mathbb{N}(\varepsilon \circ \mathbb{N}\varepsilon \circ \mathbb{N}\eta) \\ &= \kappa \circ \kappa \circ \mathbb{N}\varepsilon \circ \mathbb{N}\mathbb{N}\varepsilon \circ \mathbb{N}\mathbb{N}\eta \\ &= \kappa \circ \mathbb{N}\varepsilon \circ \nu \circ \mathbb{N}\mathbb{N}\eta \\ &= \kappa \circ \mathbb{N}(\varepsilon \circ \eta) \circ \nu \\ &= \nu,\end{aligned}$$

where the first and last steps follow from (M3) and the fourth step from (M2). (C4) now follows by (C3).

$$\begin{aligned}\nu \circ \mathbb{N}\kappa &= \kappa \circ \mathbb{N}\varepsilon \circ \mathbb{N}\kappa \\ &= \kappa \circ \mathbb{N}\mathbb{M}\kappa \circ \mathbb{N}\varepsilon \\ &= \kappa \circ \kappa \circ \mathbb{N}\varepsilon \\ &= \kappa \circ \nu.\end{aligned}$$

For (C2), we have

$$\begin{aligned}\kappa \circ \mathbb{N}\mathbb{M}\eta &= \kappa \circ \nu \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}\eta \\ &= \nu \circ \mathbb{N}\kappa \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}\eta \\ &= \nu \circ \mathbb{N}(\eta \circ \mu) \\ &= \nu \circ \mathbb{N}\eta \circ \mathbb{N}\mu \\ &= \mathbb{N}\mu,\end{aligned}$$

where the second step follows by (C4) and the third one by (M2). Finally, (C6) follows by (C5).

$$\begin{aligned}\mathbb{N}\mu \circ \kappa &= \kappa \circ \mathbb{N}\mathbb{M}\eta \circ \kappa \\ &= \kappa \circ \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{M}\eta \\ &= \kappa \circ \mathbb{N}\mathbb{M}\kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{M}\eta \\ &= \kappa \circ \mathbb{N}\mathbb{M}(\kappa \circ \mathbb{N}\mathbb{M}\eta) \\ &= \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mu.\end{aligned}$$

(5) \Rightarrow (1) Given κ satisfying (C1)–(C5), we set

$$\delta := \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon.$$

Then δ is a natural transformation since so are κ , η , and ε . Furthermore,

$$\begin{aligned}\delta \circ \mu &= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mu \\ &= \kappa \circ \eta \circ \mu \circ \mathbb{M}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ \mathbb{N}\mu \circ \eta \circ \mathbb{M}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ (\kappa \circ \mathbb{N}\mathbb{M}\eta) \circ \eta \circ \mathbb{M}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ \mathbb{N}\mathbb{M}\kappa \circ \mathbb{N}\mathbb{M}\eta \circ \eta \circ \mathbb{M}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ \eta \circ \mathbb{M}\kappa \circ \mathbb{M}\eta \circ \mathbb{M}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}(\mu \circ \varepsilon) \circ \mathbb{M}\delta \\ &= \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mu \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{M}\delta \\ &= \mathbb{N}\mu \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{M}\delta \\ &= \mathbb{N}\mu \circ \delta \circ \mathbb{M}\delta\end{aligned}$$

$$\begin{aligned}
\delta \circ \mathbb{M}v &= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{M}v \\
&= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{M}(\kappa \circ \mathbb{N}\varepsilon) \\
&= \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{N}\mathbb{M}\kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= \kappa \circ \mathbb{N}\mathbb{M}\kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon, \\
&= \kappa \circ \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon, \\
&= \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon, \\
&= \kappa \circ v \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon, \\
&= v \circ \mathbb{N}(\kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon) \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon, \\
&= v \circ \mathbb{N}\delta \circ \delta, \\
\delta \circ \varepsilon &= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \varepsilon \\
&= \kappa \circ \eta \circ \varepsilon \circ \mathbb{N}\varepsilon \\
&= \kappa \circ \zeta \circ \mathbb{N}\varepsilon \\
&= \mathbb{N}\varepsilon, \\
\delta \circ \mathbb{M}\eta &= \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \circ \mathbb{M}\eta \\
&= \kappa \circ \eta \circ \mathbb{M}(\eta \circ \varepsilon) \\
&= \kappa \circ \mathbb{N}\mathbb{M}\zeta \circ \eta \\
&= \eta.
\end{aligned}$$

It remains to prove that the above translations are bijective. We start by showing that (1) \Rightarrow (3) and (3) \Rightarrow (1) are inverse to each other. We map a given distributive law δ , to the functor $\widehat{\mathbb{M}}$ with

$$\widehat{\mathbb{M}}\varphi = v \circ \delta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta).$$

This functor is then mapped back to

$$\begin{aligned}
\delta' &= \widehat{\mathbb{M}}v \circ \eta \\
&= v \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(v \circ \eta) \circ \eta \\
&= v \circ \mathbb{N}\delta \circ \eta \\
&= v \circ \mathbb{N}\delta \circ \delta \circ \mathbb{M}\eta \\
&= \delta \circ \mathbb{M}v \circ \mathbb{M}\eta \\
&= \delta.
\end{aligned}$$

Conversely, given an extension $\widehat{\mathbb{M}}$, we construct the distributive law

$$\delta := \widehat{\mathbb{M}}v \circ \eta,$$

which in turn produces the functor $\widehat{\mathbb{M}}'$ with

$$\begin{aligned}
\widehat{\mathbb{M}}'\langle \mathbb{N}C, v \rangle &= \langle \mathbb{N}MC, v \rangle = \widehat{\mathbb{M}}\langle \mathbb{N}C, v \rangle, \\
\widehat{\mathbb{M}}'\varphi &= v \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \\
&= v \circ \mathbb{N}(\widehat{\mathbb{M}}v \circ \eta) \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \\
&= v \circ \mathbb{N}\widehat{\mathbb{M}}v \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \\
&= \widehat{\mathbb{M}}v \circ v \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \\
&= \widehat{\mathbb{M}}v \circ \mathbb{N}\mathbb{M}(\varphi \circ \eta) \\
&= \widehat{\mathbb{M}}v \circ \widehat{\mathbb{M}}\mathbb{N}(\varphi \circ \eta) \\
&= \widehat{\mathbb{M}}(v \circ \mathbb{N}(\varphi \circ \eta)) \\
&= \widehat{\mathbb{M}}(\varphi \circ v \circ \mathbb{N}\eta) \\
&= \widehat{\mathbb{M}}\varphi.
\end{aligned}$$

For the remaining translations (1) \Rightarrow (2) \Rightarrow (4) \Rightarrow (5) \Rightarrow (1), we prove bijectivity by showing that, starting with one of the four kinds of objects and applying all translations in order, we obtain the original object back.

Given a distributive law δ , the first translation maps it to a monad $\hat{\mathbb{N}}$ with $\hat{\mu} = \mathbb{N}\mu \circ \delta$. The second step, maps this to the morphism $\kappa = v \circ \mathbb{N}\hat{\mu}$. The

third step is the identity, and the last step maps κ to $\kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon$. Composing these steps we obtain

$$\begin{aligned} \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon &= \nu \circ \mathbb{N}\hat{\mu} \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \nu \circ \mathbb{N}(\mathbb{N}\mu \circ \delta) \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\mu \circ \nu \circ \mathbb{N}\delta \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\mu \circ \nu \circ \eta \circ \delta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\mu \circ \mathbb{N}\mathbb{M}\varepsilon \circ \delta \\ &= \delta, \end{aligned}$$

as desired.

Similarly, if we start with κ as in (4) or (5), we translate it into $\delta = \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon$ and then into the functor mapping $\langle A, \pi \rangle$ to $\langle \mathbb{N}A, \mathbb{N}\pi \circ \delta \rangle$. Finally, we obtain the morphism

$$\begin{aligned} \nu \circ \mathbb{N}\hat{\mu} &= \nu \circ \mathbb{N}(\mathbb{N}\mu \circ \delta) \\ &= \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}(\kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon) \\ &= \mathbb{N}\mu \circ \nu \circ \mathbb{N}\kappa \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\mu \circ \kappa \circ \nu \circ \mathbb{N}\eta \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\mu \circ \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa \circ \mathbb{N}\mathbb{M}\mathbb{N}\mu \circ \mathbb{N}\mathbb{M}\mathbb{N}\varepsilon \\ &= \kappa. \end{aligned}$$

Finally, consider the case where we start with a lifting $\hat{\mathbb{N}}$. Then we obtain $\kappa = \nu \circ \mathbb{N}\hat{\mu}$, $\delta = \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon$, and finally the functor mapping $\langle A, \pi \rangle$ to $\langle A, \mathbb{N}\pi \circ \delta \rangle$. We have to prove that the resulting algebra is equal to $\hat{\mathbb{N}}\langle A, \pi \rangle = \langle \mathbb{N}A, \hat{\pi} \rangle$. Note that the associative law

$$\pi \circ \mu = \pi \circ \mathbb{M}\pi$$

implies that the morphism $\pi : \mathbb{M}A \rightarrow A$ is a morphism of \mathbb{M} -algebras $\langle \mathbb{M}A, \mu \rangle \rightarrow \langle A, \pi \rangle$. Consequently, its image under $\hat{\mathbb{N}}$ is also a morphism of

\mathbb{M} -algebras and we have

$$\hat{\mathbb{N}}\pi \circ \hat{\mu} = \hat{\pi} \circ \mathbb{M}\hat{\mathbb{N}}\pi.$$

Since $\hat{\mathbb{N}}\pi = \mathbb{N}\pi$ it follows that

$$\begin{aligned} \mathbb{N}\pi \circ \delta &= \mathbb{N}\pi \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\pi \circ \nu \circ \mathbb{N}\hat{\mu} \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\pi \circ \nu \circ \eta \circ \hat{\mu} \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \mathbb{N}\pi \circ \hat{\mu} \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \hat{\pi} \circ \mathbb{M}\mathbb{N}\pi \circ \mathbb{M}\mathbb{N}\varepsilon \\ &= \hat{\pi}, \end{aligned}$$

as desired. \square

Corollary 6.5. Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$.

(a) The composition $\mathbb{N}\mathbb{M}$ forms a monad where multiplication and singleton operation are given by the morphisms

$$\nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\delta : \mathbb{N}\mathbb{M}\mathbb{N}\mathbb{M} \Rightarrow \mathbb{N}\mathbb{M} \quad \text{and} \quad \eta \circ \varepsilon : \text{Id} \Rightarrow \mathbb{N}\mathbb{M}.$$

(b) One can lift \mathbb{N} to a functor on \mathbb{M} -algebras that maps an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the \mathbb{M} -algebra $\mathbb{N}\mathfrak{A}$ with product

$$\mathbb{N}\pi \circ \delta : \mathbb{M}\mathbb{N}A \rightarrow \mathbb{N}A.$$

(c) Every lift $\mathbb{N}\mathfrak{A}$ of an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$ carries a canonical $\mathbb{N}\mathbb{M}$ -algebra structure with product

$$\nu \circ \mathbb{N}\mathbb{N}\pi \circ \mathbb{N}\delta : \mathbb{N}\mathbb{M}\mathbb{N}A \rightarrow \mathbb{N}A.$$

Remark. What do we do if we can't combine two monads \mathbb{M} and \mathbb{N} but there is no distributive law between them? In this case we can use the free monad construction and use the more complicated monad $(\mathbb{N}\mathbb{M})^*$.

Exercise 6.1. Let \mathbb{M} and \mathbb{N} be two monads.

- (a) Prove that $(\mathbb{M} + \mathbb{N})^* \cong (\mathbb{NM})^*$.
- (b) Show that every distributive law $\delta : \mathbb{MN} \Rightarrow \mathbb{NM}$ induces an isomorphism $(\mathbb{NM})^* \cong \mathbb{NM}$.

We conclude this section with a simple observation of how to transfer a distributive law along a morphism of monads. A similar result holds for monomorphic ρ .

Lemma 6.6. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$, $\langle \mathbb{M}', \mu', \varepsilon' \rangle$, and $\langle \mathbb{N}, \nu, \eta \rangle$ be monads, $\rho : \mathbb{M} \Rightarrow \mathbb{M}'$ a morphism of monads, $\delta : \mathbb{MN} \Rightarrow \mathbb{NM}$ a distributive law, and $\delta' = (\delta'_A)_A$ a family of functions such that

$$\delta' \circ \rho = \mathbb{N}\rho \circ \delta.$$

If ρ and $\mathbb{M}\rho$ are epimorphisms, then δ' is a distributive law $\mathbb{M}'\mathbb{N} \Rightarrow \mathbb{NM}'$.

Proof. We start by checking that δ' is natural. Given a morphism $f : A \rightarrow B$, we have

$$\begin{aligned} \mathbb{NM}'f \circ \delta' \circ \rho &= \mathbb{NM}'f \circ \mathbb{N}\rho \circ \delta \\ &= \mathbb{N}\rho \circ \mathbb{NM}f \circ \delta \\ &= \mathbb{N}\rho \circ \delta \circ \mathbb{MN}f = \delta' \circ \mathbb{M}'\mathbb{N}f. \end{aligned}$$

It remains to prove the four equations for a distributive law. Note that

$$\begin{aligned} \delta' \circ \varepsilon' &= \delta' \circ \rho \circ \varepsilon \\ &= \mathbb{N}\rho \circ \delta \circ \varepsilon \\ &= \mathbb{N}\rho \circ \mathbb{N}\varepsilon \\ &= \mathbb{N}\varepsilon', \\ \delta' \circ \mathbb{M}'\eta \circ \rho &= \delta' \circ \rho \circ \mathbb{M}\eta \\ &= \mathbb{N}\rho \circ \delta \circ \mathbb{M}\eta \\ &= \mathbb{N}\rho \circ \eta \\ &= \eta \circ \rho, \end{aligned}$$

$$\begin{aligned} \delta' \circ \mu' \circ (\rho \circ \mathbb{M}\rho) &= \delta' \circ \rho \circ \mu \\ &= \mathbb{N}\rho \circ \delta \circ \mu \\ &= \mathbb{N}\rho \circ \mathbb{N}\mu \circ \delta \circ \mathbb{M}\delta \\ &= \mathbb{N}(\mu' \circ \rho \circ \mathbb{M}\rho) \circ \delta \circ \mathbb{M}\delta \\ &= \mathbb{N}\mu' \circ \mathbb{N}\rho \circ \delta \circ \mathbb{MN}\rho \circ \mathbb{M}\delta \\ &= \mathbb{N}\mu' \circ \delta' \circ \rho \circ \mathbb{M}(\delta' \circ \rho) \\ &= \mathbb{N}\mu' \circ \delta' \circ \mathbb{M}'\delta' \circ (\rho \circ \mathbb{M}\rho), \\ \delta' \circ \mathbb{M}'\nu \circ \rho &= \delta' \circ \rho \circ \mathbb{M}\nu \\ &= \mathbb{N}\rho \circ \delta \circ \mathbb{M}\nu \\ &= \mathbb{N}\rho \circ \nu \circ \mathbb{N}\delta \circ \delta \\ &= \nu \circ \mathbb{NN}\rho \circ \mathbb{N}\delta \circ \delta \\ &= \nu \circ \mathbb{N}(\delta' \circ \rho) \circ \delta \\ &= \nu \circ \mathbb{N}\delta' \circ \delta' \circ \rho. \end{aligned}$$

As ρ and $\rho \circ \mathbb{M}\rho$ are epimorphisms, it follows that

$$\begin{aligned} \delta' \circ \varepsilon' &= \mathbb{N}\varepsilon, & \delta' \circ \mu' &= \mathbb{N}\mu' \circ \delta' \circ \mathbb{M}'\delta', \\ \delta' \circ \mathbb{M}'\eta &= \eta, & \delta' \circ \mathbb{M}'\nu &= \nu \circ \mathbb{N}\delta' \circ \delta'. \end{aligned} \quad \square$$

7. Bialgebras

The theory of distributive laws also provides a criterion on when two algebra structures $\pi : \mathbb{M}A \rightarrow A$ and $\rho : \mathbb{N}A \rightarrow A$ on the same set A can be combined into a single \mathbb{NM} -algebra structure. The case where we are particularly interested in, consists in adding a meet operation $\inf : \mathbb{U}A \rightarrow A$ to an \mathbb{M} -algebra $\pi : \mathbb{M}A \rightarrow A$.

Definition 7.1. Let $\delta : \mathbb{MN} \Rightarrow \mathbb{NM}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$.

(a) A δ -bialgebra $\mathfrak{A} = \langle A, \pi, \rho \rangle$ consists of an object A and two morphisms $\pi : \mathbb{M}A \rightarrow A$ and $\rho : \mathbb{N}A \rightarrow A$ such that

- $\langle A, \pi \rangle$ forms an \mathbb{M} -algebra,
- $\langle A, \rho \rangle$ forms an \mathbb{N} -algebra,
- $\pi \circ \mathbb{M}\rho = \rho \circ \mathbb{N}\pi \circ \delta$.

(b) A *morphism of δ -bialgebras* is a morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ that is both an \mathbb{M} -morphism and an \mathbb{N} -morphism.

(c) We denote the category of all δ -bialgebras and their morphisms by $\text{Bialg}(\delta)$. \lrcorner

Example. Let $\delta : \mathbb{M}\mathbb{A} \Rightarrow \mathbb{A}\mathbb{M}$ be the distributive law from the example on page 66. A δ -bialgebra is a structure $\langle S, \pi, \rho \rangle$ where $\langle S, \pi \rangle$ is an arbitrary monoid and $\langle S, \rho \rangle$ a commutative. If we write the former multiplicatively and the latter additively, the bialgebra axiom

$$\pi \circ \mathbb{M}\rho = \rho \circ \mathbb{N}\pi \circ \delta$$

reduces to the familiar distributive law

$$(a + b) \cdot (c + d) = ac + ad + bc + bd.$$

Hence, δ -bialgebras are the same as semirings. \lrcorner

It turns out that a δ -bialgebra is just another formalisation for an \mathbb{NM} -algebra.

Theorem 7.2. Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$.

- (a) Every δ -bialgebra $\mathfrak{A} = \langle A, \pi, \rho \rangle$ induces an \mathbb{NM} -algebra $\langle A, \sigma \rangle$ with product

$$\sigma = \rho \circ \mathbb{N}\pi.$$

- (b) Every \mathbb{NM} -algebra $\mathfrak{A} = \langle A, \sigma \rangle$ induces a δ -bialgebra $\langle A, \pi, \rho \rangle$ with products

$$\pi = \sigma \circ \eta \quad \text{and} \quad \rho = \sigma \circ \mathbb{N}\varepsilon.$$

- (c) This correspondence between \mathbb{NM} -algebras and δ -bialgebras is bijective.
- (d) A morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ is an \mathbb{NM} -morphism if, and only if, it is a δ -bialgebra morphism.

Proof. Remember from Theorem 6.4 that the monad \mathbb{NM} has the unit map $\eta \circ \varepsilon$ and the product

$$\kappa := \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\delta.$$

- (a) Let $\mathfrak{A} = \langle A, \pi, \rho \rangle$ be a δ -bialgebra. Then

$$\begin{aligned} \sigma \circ (\eta \circ \varepsilon) &= \rho \circ \mathbb{N}\pi \circ \eta \circ \varepsilon \\ &= \rho \circ \eta \circ \pi \circ \varepsilon \\ &= \text{id} \circ \text{id}, \\ \sigma \circ \mathbb{NM}\sigma &= \rho \circ \mathbb{N}\pi \circ \mathbb{NM}(\rho \circ \mathbb{N}\pi) \\ &= \rho \circ \mathbb{N}(\pi \circ \mathbb{M}\rho) \circ \mathbb{NM}\mathbb{N}\pi \\ &= \rho \circ \mathbb{N}(\rho \circ \mathbb{N}\pi \circ \delta) \circ \mathbb{NM}\mathbb{N}\pi \\ &= \rho \circ \mathbb{N}\rho \circ \mathbb{NN}\pi \circ \mathbb{NNM}\pi \circ \mathbb{N}\delta \\ &= \rho \circ \nu \circ \mathbb{NN}(\pi \circ \mu) \circ \mathbb{N}\delta \\ &= (\rho \circ \mathbb{N}\pi) \circ (\nu \circ \mathbb{NN}\mu \circ \mathbb{N}\delta) \\ &= \sigma \circ \kappa. \end{aligned}$$

- (b) Let $\mathfrak{A} = \langle A, \sigma \rangle$ be an \mathbb{NM} -algebra. Then $\langle A, \sigma \circ \eta \rangle$ is the η -reduct of \mathfrak{A} . By Lemma II.3.2, it is therefore also an \mathbb{M} -algebra. Furthermore, $\langle A, \sigma \circ \mathbb{N}\varepsilon \rangle$ is an \mathbb{N} -algebra since

$$\begin{aligned} \rho \circ \eta &= \sigma \circ \mathbb{N}\varepsilon \circ \eta \\ &= \sigma \circ \eta \circ \varepsilon \\ &= \text{id}, \end{aligned}$$

$$\begin{aligned}
\rho \circ \mathbb{N}\rho &= \sigma \circ \mathbb{N}\varepsilon \circ \mathbb{N}(\sigma \circ \mathbb{N}\varepsilon) \\
&= \sigma \circ \mathbb{N}\mathbb{M}(\sigma \circ \mathbb{N}\varepsilon) \circ \mathbb{N}\varepsilon \\
&= \sigma \circ \kappa \circ \mathbb{N}(\mathbb{M}\mathbb{N}\varepsilon \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\delta \circ \mathbb{N}(\mathbb{M}\mathbb{N}\varepsilon \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{N}(\mathbb{N}\mu \circ \mathbb{N}\mathbb{M}\varepsilon \circ \delta \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{N}(\delta \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{N}\mathbb{N}\varepsilon \\
&= \sigma \circ \mathbb{N}\varepsilon \circ \nu \\
&= \rho \circ \nu.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
\rho \circ \mathbb{N}\pi \circ \delta &= (\sigma \circ \mathbb{N}\varepsilon) \circ \mathbb{N}(\sigma \circ \eta) \circ \delta \\
&= \sigma \circ \mathbb{N}\mathbb{M}\sigma \circ \mathbb{N}\varepsilon \circ \mathbb{N}\eta \circ \delta \\
&= \sigma \circ \kappa \circ \mathbb{N}\varepsilon \circ \mathbb{N}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\delta \circ \mathbb{N}\varepsilon \circ \mathbb{N}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\mathbb{N}\varepsilon \circ \mathbb{N}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{N}\eta \circ \delta \\
&= \sigma \circ \delta \\
&= \sigma \circ \mathbb{N}(\mu \circ \mathbb{M}\varepsilon) \circ \delta \\
&= \sigma \circ \mathbb{N}\mu \circ \delta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= \sigma \circ \mathbb{N}\mu \circ (\nu \circ \eta) \circ \delta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= \sigma \circ \nu \circ \mathbb{N}\mathbb{N}\mu \circ \mathbb{N}\delta \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= \sigma \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= \sigma \circ \mathbb{N}\mathbb{M}\sigma \circ \eta \circ \mathbb{M}\mathbb{N}\varepsilon \\
&= (\sigma \circ \eta) \circ \mathbb{M}(\sigma \circ \mathbb{N}\varepsilon) \\
&= \pi \circ \mathbb{M}\rho.
\end{aligned}$$

(c) We have to show that the mappings from (a) and (b) are inverse to each other. First, let $\langle A, \sigma \rangle$ be an $\mathbb{N}\mathbb{M}$ -algebra. The corresponding δ -bialgebra

has products $\pi = \sigma \circ \eta$ and $\rho = \sigma \circ \mathbb{N}\varepsilon$. Mapping it back, we obtain the product

$$\begin{aligned}
\sigma' &= \rho \circ \mathbb{N}\pi \\
&= (\sigma \circ \mathbb{N}\varepsilon) \circ \mathbb{N}(\sigma \circ \eta) \\
&= \sigma \circ \mathbb{N}(\mathbb{M}\sigma \circ \varepsilon \circ \eta) \\
&= \sigma \circ \kappa \circ \mathbb{N}(\varepsilon \circ \eta) \\
&= \sigma,
\end{aligned}$$

where the last step follows by Theorem 6.4 (4).

Conversely, let $\langle A, \pi, \rho \rangle$ be a δ -bialgebra. The corresponding $\mathbb{N}\mathbb{M}$ -algebra has the product $\sigma = \rho \circ \mathbb{N}\pi$. Mapping it back, we obtain the products

$$\begin{aligned}
\pi' &= \sigma \circ \eta \\
&= \rho \circ \mathbb{N}\pi \circ \eta \\
&= \rho \circ \eta \circ \pi \\
&= \pi, \\
\rho' &= \sigma \circ \mathbb{N}\varepsilon \\
&= \rho \circ \mathbb{N}\pi \circ \mathbb{N}\varepsilon \\
&= \rho.
\end{aligned}$$

(d) Let $\langle A, \sigma \rangle$ and $\mathfrak{B} = \langle B, \sigma' \rangle$ be $\mathbb{N}\mathbb{M}$ -algebras, and let $\langle A, \pi, \rho \rangle$ and $\langle B, \pi', \rho' \rangle$ be the corresponding δ -algebras. If $\varphi : A \rightarrow B$ is an $\mathbb{N}\mathbb{M}$ -morphism, we have

$$\begin{aligned}
\varphi \circ \pi &= \varphi \circ \sigma \circ \eta \\
&= \sigma' \circ \mathbb{N}\mathbb{M}\varphi \circ \eta \\
&= \sigma' \circ \eta \circ \mathbb{M}\varphi \\
&= \pi' \circ \mathbb{M}\varphi,
\end{aligned}$$

$$\begin{aligned}
\varphi \circ \rho &= \varphi \circ \sigma \circ \mathbb{N}\varepsilon \\
&= \sigma' \circ \mathbb{N}\mathbb{M}\varphi \circ \mathbb{N}\varepsilon \\
&= \sigma' \circ \mathbb{N}\varepsilon \circ \mathbb{N}\varphi \\
&= \rho' \circ \mathbb{N}\varphi.
\end{aligned}$$

Conversely, if $\varphi : A \rightarrow B$ is a δ -bialgebra morphism, we have

$$\begin{aligned}
\varphi \circ \sigma &= \varphi \circ \rho \circ \mathbb{N}\pi \\
&= \rho' \circ \mathbb{N}\varphi \circ \mathbb{N}\pi \\
&= \rho' \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{M}\varphi \\
&= \sigma' \circ \mathbb{N}\mathbb{M}\varphi.
\end{aligned}$$

□

As a consequence the functor $\text{Alg}(\mathbb{M}) \rightarrow \text{Alg}(\mathbb{N}\mathbb{M})$ from Corollary 6.5 can be used to turn every \mathbb{M} -algebra into a δ -bialgebra.

Theorem 7.3. *Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law. The forgetful functor $\text{Bialg}(\delta) \rightarrow \text{Alg}(\mathbb{M})$ has a left adjoint, which maps an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the δ -bialgebra $\langle \mathbb{N}A, \mathbb{N}\pi \circ \delta, \nu \rangle$. The unit of this adjunction is given by the unit morphism η of the monad \mathbb{N} .*

Proof. Let $\mathbb{W} : \text{Bialg}(\delta) \rightarrow \text{Alg}(\mathbb{M})$ be the forgetful functor and let $\widehat{\mathbb{N}} : \text{Alg}(\mathbb{M}) \rightarrow \text{Bialg}(\delta)$ be its supposed left adjoint.

First, note that $\widehat{\mathbb{N}}$ is the functor $\text{Alg}(\mathbb{M}) \rightarrow \text{Alg}(\mathbb{N}\mathbb{M}) \rightarrow \text{Bialg}(\delta)$ that is induced by \mathbb{N} according to Corollary 6.5 (c) and Theorem 7.2 (b). To see that, note that it maps an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$ first to the $\mathbb{N}\mathbb{M}$ -algebra $\langle \mathbb{N}A, \sigma \rangle$ with

$$\sigma = \nu \circ \mathbb{N}\mathbb{N}\pi \circ \mathbb{N}\delta,$$

and then to the δ -bialgebra $\langle \mathbb{N}A, \hat{\pi}, \hat{\rho} \rangle$ where

$$\begin{aligned}
\hat{\pi} &= \sigma \circ \eta \\
&= \nu \circ \mathbb{N}\mathbb{N}\pi \circ \mathbb{N}\delta \circ \eta \\
&= \mathbb{N}\pi \circ \nu \circ \mathbb{N}\delta \circ \delta \circ \mathbb{M}\eta \\
&= \mathbb{N}\pi \circ \delta \circ \mathbb{M}\nu \circ \mathbb{M}\eta \\
&= \mathbb{N}\pi \circ \delta, \\
\hat{\rho} &= \sigma \circ \mathbb{N}\varepsilon \\
&= \nu \circ \mathbb{N}\mathbb{N}\pi \circ \mathbb{N}\delta \circ \mathbb{N}\varepsilon \\
&= \nu \circ \mathbb{N}\mathbb{N}\pi \circ \mathbb{N}\mathbb{N}\varepsilon \\
&= \nu.
\end{aligned}$$

To prove that $\widehat{\mathbb{N}} \dashv \mathbb{W}$, we construct the unit and counit. Let $\mathfrak{A} = \langle A, \pi \rangle$ be an \mathbb{M} -algebra and $\mathfrak{B} = \langle B, \pi', \rho' \rangle$ a δ -bialgebra. Then

$$\begin{aligned}
\mathbb{W}\widehat{\mathbb{N}}\mathfrak{A} &= \langle \mathbb{N}A, \mathbb{N}\pi \circ \delta \rangle = \mathbb{N}\mathfrak{A}, \\
\widehat{\mathbb{N}}\mathbb{W}\mathfrak{B} &= \langle \mathbb{N}B, \mathbb{N}\pi' \circ \delta, \nu \rangle.
\end{aligned}$$

We claim that the unit of the adjunction is given by $\eta : \text{Id} \rightarrow \mathbb{W}\widehat{\mathbb{N}}$ and its counit by $\tau : \widehat{\mathbb{N}}\mathbb{W} \rightarrow \text{Id}$ where $\tau_{\mathfrak{B}} := \rho'$ is the \mathbb{N} -product of \mathfrak{B} . Note that they satisfy the identities

$$\tau \circ \mathbb{N}\eta = \nu \circ \mathbb{N}\eta = \text{id} \quad \text{and} \quad \mathbb{W}\rho \circ \eta = \rho \circ \eta = \text{id}$$

characterising the unit and counit. Hence, it remains to prove that η and τ are natural transformations of the correct form.

First, note that η is indeed an \mathbb{M} -morphism since

$$(\mathbb{N}\pi \circ \delta) \circ \mathbb{M}\eta = \mathbb{N}\pi \circ \eta = \eta \circ \pi,$$

while τ is a morphism of δ -bialgebras since

$$\begin{aligned}\tau \circ (\mathbb{N}\pi' \circ \delta) &= \rho' \circ \mathbb{N}\pi' \circ \delta \\ &= \pi' \circ \mathbb{M}\rho' \\ &= \pi' \circ \mathbb{M}\tau, \\ \tau \circ \nu &= \rho' \circ \nu \\ &= \rho' \circ \mathbb{M}\rho' \\ &= \rho' \circ \mathbb{M}\tau.\end{aligned}$$

For naturality, we consider two morphisms $\varphi : \langle A_o, \pi_o \rangle \rightarrow \langle A_i, \pi_i \rangle$ and $\psi : \langle B_o, \pi'_o, \rho'_o \rangle \rightarrow \langle B_i, \pi'_i, \rho'_i \rangle$. Then

$$\eta \circ \varphi = \mathbb{N}\varphi \circ \eta \quad \text{and} \quad \rho'_i \circ \mathbb{N}\psi = \psi \circ \rho'_o$$

follow from naturality of η and the fact that ψ is an \mathbb{N} -morphism. \square

Corollary 7.4. Let $\mathfrak{A} = \langle A, \pi \rangle$ be an \mathbb{M} -algebra, $\mathfrak{B} = \langle B, \pi', \rho' \rangle$ a δ -bialgebra, and $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ an \mathbb{M} -morphism. Then

$$\hat{\varphi} := \pi' \circ \mathbb{N}\varphi$$

is the unique δ -bialgebra morphism $\hat{\varphi} : \mathbb{N}\mathfrak{A} \rightarrow \mathfrak{B}$ such that

$$\varphi = \hat{\varphi} \circ \eta.$$

For our second construction of δ -bialgebras, we need to generalise the notion of a δ -bialgebra as follows.

Definition 7.5. Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \epsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$. A function $j : \mathfrak{A} \rightarrow \mathfrak{B}$ from an \mathbb{M} -algebra \mathfrak{A} to an \mathbb{N} -algebra \mathfrak{B} is δ -distributive if there exists a function $\sigma : \mathbb{M}B \rightarrow B$ such that

$$\sigma \circ \mathbb{M}(\nu \circ \mathbb{N}j) = \nu \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}j.$$

Lemma 7.6. Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \epsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$. Let \mathfrak{A} be both an \mathbb{M} -algebra and an \mathbb{N} -algebra. Then \mathfrak{A} is a δ -bialgebra if, and only if, the identity $\text{id} : A \rightarrow A$ is δ -distributive.

Proof. (\Rightarrow) In the definition of meet-distributivity, we can take $\sigma := \pi$.

(\Leftarrow) Let σ be the function from the definition of meet-distributivity. Then

$$\begin{aligned}\sigma &= \sigma \circ \mathbb{M}(\nu \circ \eta) \\ &= \nu \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\eta \\ &= \nu \circ \mathbb{N}\pi \circ \eta \\ &= \nu \circ \eta \circ \pi \\ &= \pi.\end{aligned}$$

Hence,

$$\begin{aligned}\pi \circ \mathbb{M}\nu &= \sigma \circ \mathbb{M}(\nu \circ \text{Nid}) \\ &= \nu \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\text{Nid} = \nu \circ \mathbb{N}\pi \circ \delta,\end{aligned}$$

and \mathfrak{A} is a δ -bialgebra. \square

Proposition 7.7. Let $\delta : \mathbb{M}\mathbb{N} \Rightarrow \mathbb{N}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \epsilon \rangle$ and $\langle \mathbb{N}, \nu, \eta \rangle$ and suppose that E is a class of epimorphisms preserved by both \mathbb{M} and \mathbb{N} . Let $\varphi : \mathfrak{C} \rightarrow \mathfrak{A}$ be a δ -distributive function such that $\rho \circ \mathbb{N}\varphi \in E$ (i.e., such that $\text{rng } \varphi$ generates $\mathfrak{A} = \langle A, \rho \rangle$ as an \mathbb{N} -algebra). There exists a unique function $\sigma : \mathbb{M}A \rightarrow A$ such that $\langle A, \sigma, \rho \rangle$ forms a δ -bialgebra and φ an \mathbb{M} -morphism.

Proof. Let $\sigma : \mathbb{M}A \rightarrow A$ be the function from the definition of δ -distributivity.

ity. We start by proving that $\langle A, \sigma \rangle$ is an \mathbb{M} -algebra. Note that

$$\begin{aligned}
\sigma \circ \varepsilon \circ (\rho \circ \mathbb{N}\varphi) &= \sigma \circ \mathbb{M}(\rho \circ \eta) \circ \varepsilon \circ (\rho \circ \mathbb{N}\varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \eta) \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi) \circ \varepsilon \\
&= \sigma \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi) \circ \mathbb{M}\eta \circ \varepsilon \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi) \circ \mathbb{M}\eta \circ \varepsilon \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \varepsilon \circ \mathbb{N}(\rho \circ \mathbb{N}\varphi) \circ \eta \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \varepsilon \circ \eta \circ (\rho \circ \mathbb{N}\varphi) \\
&= \rho \circ \mathbb{N}\pi \circ \mathbb{N}\varepsilon \circ \eta \circ (\rho \circ \mathbb{N}\varphi) \\
&= \rho \circ \eta \circ (\rho \circ \mathbb{N}\varphi) \\
&= \rho \circ \mathbb{N}\varphi, \\
\sigma \circ \mathbb{M}\sigma \circ \mathbb{M}\mathbb{M}(\rho \circ \mathbb{N}\varphi) &= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{M}\varphi \circ \delta) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi \circ \mathbb{N}\pi \circ \delta) \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi \circ \mathbb{M}(\mathbb{N}\pi \circ \delta) \\
&= \rho \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{M}(\varphi \circ \pi) \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{N}(\pi \circ \mathbb{M}(\pi \circ \mathbb{M}\varphi)) \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{N}(\pi \circ \mu \circ \mathbb{M}\mathbb{M}\varphi) \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{N}(\pi \circ \mathbb{M}\varphi \circ \mu) \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{N}(\pi \circ \mathbb{M}\varphi) \circ \delta \circ \mu \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi \circ \mu \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mu \circ \mathbb{M}\mathbb{M}\mathbb{N}\varphi \\
&= \sigma \circ \mathbb{M}\rho \circ \mu \circ \mathbb{M}\mathbb{M}\mathbb{N}\varphi \\
&= \sigma \circ \mu \circ \mathbb{M}\mathbb{M}(\rho \circ \mathbb{N}\varphi).
\end{aligned}$$

Since $\rho \circ \mathbb{N}\varphi \in E$ and $\mathbb{M}\mathbb{M}(\rho \circ \mathbb{N}\varphi) \in E$ are epimorphisms, it follows that

$$\sigma \circ \varepsilon = \text{id} \quad \text{and} \quad \sigma \circ \mathbb{M}\sigma = \sigma \circ \mu.$$

To check that \mathfrak{A} is even a δ -bialgebra, note that

$$\begin{aligned}
&\sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\text{id}) \circ \mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\rho \circ \mathbb{N}\mathbb{N}\varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \nu \circ \mathbb{N}\mathbb{N}\varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi \circ \nu) \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi \circ \mathbb{M}\nu \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}(\nu \circ \mathbb{N}\mathbb{N}\varphi) \\
&= \rho \circ \mathbb{N}\pi \circ \nu \circ \mathbb{N}\delta \circ \delta \circ \mathbb{M}\mathbb{N}\mathbb{N}\varphi \\
&= \rho \circ \mathbb{N}\pi \circ \nu \circ \mathbb{N}\delta \circ \mathbb{N}\mathbb{M}\mathbb{N}\varphi \circ \delta \\
&= \rho \circ \nu \circ \mathbb{N}(\mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi) \circ \delta \\
&= \rho \circ \mathbb{N}\rho \circ \mathbb{N}(\mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi) \circ \delta \\
&= \rho \circ \mathbb{N}(\sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi)) \circ \delta \\
&= \rho \circ \mathbb{N}\sigma \circ \delta \circ \mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi) \\
&= \rho \circ \mathbb{N}\sigma \circ \delta \circ \mathbb{M}\mathbb{N}\text{id} \circ \mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi).
\end{aligned}$$

By surjectivity of $\mathbb{M}\mathbb{N}(\rho \circ \mathbb{N}\varphi) \in E$, this implies that

$$\sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\text{id}) = \rho \circ \mathbb{N}\sigma \circ \delta \circ \mathbb{M}\mathbb{N}\text{id}.$$

To see that φ is an \mathbb{M} -morphism, note that

$$\begin{aligned}
\sigma \circ \mathbb{M}\varphi &= \sigma \circ \mathbb{M}(\rho \circ \eta \circ \varphi) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi \circ \eta) \\
&= \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi \circ \mathbb{M}\eta \\
&= \rho \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{M}\varphi \circ \delta \circ \mathbb{M}\eta \\
&= \rho \circ \mathbb{N}(\varphi \circ \pi) \circ \delta \circ \mathbb{M}\eta \\
&= \rho \circ \mathbb{N}(\varphi \circ \pi) \circ \eta \\
&= \rho \circ \eta \circ \varphi \circ \pi \\
&= \varphi \circ \pi.
\end{aligned}$$

Finally, for uniqueness suppose that $\sigma' : \mathbb{M}A \rightarrow A$ is another function like this. Then we have

$$\sigma \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi) = \rho \circ \mathbb{N}\pi \circ \delta \circ \mathbb{M}\mathbb{N}\varphi = \sigma' \circ \mathbb{M}(\rho \circ \mathbb{N}\varphi),$$

and surjectivity of $\mathbb{M}(\rho \circ \mathbb{N}\varphi) \in E$ implies that $\sigma = \sigma'$. \square

Notes

Much of the material in this chapter is standard and can be found in various accounts on category theory. The definition of discrete categories seems to be new. For a treatment of polynomial functors in a general category-theoretical setting, see [14].

II. Algebra

1. Factorisation Systems

To further develop our algebraic machinery we need the notions of a sub-algebra and a quotient. We start with the observation that, in **Set**, every function can be uniquely factorised into a surjective function followed by an injective one. Let us axiomatise this property.

Definition 1.1. Let \mathcal{C} be a category and E and M two sets of morphisms.

(a) We call the morphisms in M *embeddings*, or *M-morphisms*, and those in E *projections*, or *E-morphisms*.

(b) An *EM-factorisation* of a morphism $\varphi : A \rightarrow B$ consists of two morphisms $\varepsilon \in E$ and $\mu \in M$ such that $\varphi = \mu \circ \varepsilon$.

(c) A *factorisation system* for \mathcal{C} is a pair $\langle E, M \rangle$ of classes of morphisms satisfying the following conditions.

(FS1) E consists of epimorphisms and it contains all isomorphisms.


(FS2) M consists of monomorphisms and it contains all isomorphisms.

(FS3) E and M are both closed under composition.

(FS4) Every morphism $\varphi : A \rightarrow B$ has an *EM-factorisation*.

(FS5) For every choice of morphisms $\varphi, \psi, \mu, \varepsilon$ with $\mu \in M$, $\varepsilon \in E$, and $\psi \circ \varepsilon = \mu \circ \varphi$, there exists a unique morphism δ such that

$$\begin{aligned} \mu \circ \delta &= \psi, \\ \delta \circ \varepsilon &= \varphi. \end{aligned}$$




(FS5) is often referred to as the *diagonal fill-in property*.

(d) If $\varphi = \mu \circ \varepsilon$ with $\mu \in M$ and $\varepsilon \in E$, we call μ the *image* of φ and ε its *kernel*. We denote them by $\text{im } \varphi$ and $\ker \varphi$, respectively. (We will see below that this factorisation is unique, up to isomorphism.)

Remark. Axiom (FS5) is the category-theoretical analogue of the Factorisation Lemma from Universal Algebra. The following reformulation makes this more apparent: given morphisms $\varepsilon : A \rightarrow B$, $\varphi : A \rightarrow C$, $\psi : B \rightarrow C$ with $\varepsilon \in E$ satisfying $\varphi = \psi \circ \varepsilon$, it follows that

$$\ker \varphi = (\ker \psi) \circ \varepsilon \quad \text{and} \quad \text{im } \varphi = \text{im } \psi.$$



Example. (a) The only factorisation system for **Set** is $\langle E, M \rangle$ where E consists of all surjective functions and M of all injective ones. Most axioms follow immediately, only (FS5) requires a bit of thought.

Hence, suppose that $\psi \circ \varepsilon = \mu \circ \varphi$ where ε is surjective and μ injective. We define δ as follows. Given an element b , we choose some $a \in \varepsilon^{-1}(b)$ and we set $\delta(b) := \varphi(a)$. This definition immediately implies that

$$\delta(\varepsilon(a)) = \varphi(a), \quad \text{for all } a.$$

Hence, it remains to check that the definition of δ does not depend on the choice of a . Suppose that $\varepsilon(a) = \varepsilon(a')$. Then

$$\mu(\varphi(a)) = \psi(\varepsilon(a)) = \psi(\varepsilon(a')) = \mu(\varphi(a')),$$

which by injectivity of μ implies that $\varphi(a) = \varphi(a')$, as desired.

(b) We obtain a factorisation system for **Pos** by taking for M all injective functions and for E all surjective functions $\varepsilon : A \rightarrow B$ satisfying

$$\varepsilon(a) \leq \varepsilon(b) \iff a' \leq b', \quad \text{for some } a' \in \varepsilon^{-1}(a) \text{ and } b' \in \varepsilon^{-1}(b).$$

Again most of the axioms are trivial. For (FS5) we can proceed as in the case of sets above. It only remains to check that the function δ is monotone. Hence, suppose that $b \leq b'$. This implies that $a \leq a'$ for all $a \in \varepsilon^{-1}(b)$ and $a' \in \varepsilon^{-1}(b')$. Consequently, $\delta(b) = \varphi(a) \leq \varphi(a') = \delta(b')$.

(c) A second factorisation system for **Pos** consists of all surjective monotone functions for E all and all embeddings for M , i.e., all functions $\mu : A \rightarrow B$ satisfying

$$a \leq b \quad \text{iff} \quad \mu(a) \leq \mu(b), \quad \text{for all } a, b \in A.$$

To check (FS5), consider morphisms $\psi \circ \varepsilon = \mu \circ \varphi$ where ε is surjective and μ is an embedding. Again, we define

$$\delta(b) := \varphi(a), \quad \text{for some } a \in \varepsilon^{-1}(b).$$

To show that this is well-defined and that δ is monotone, note that

$$\begin{aligned} \varepsilon(a) \leq \varepsilon(a') &\Rightarrow \psi(\varepsilon(a)) \leq \psi(\varepsilon(a')) \\ &\Rightarrow \mu(\varphi(a)) \leq \mu(\varphi(a')) \\ &\Rightarrow \varphi(a) \leq \varphi(a'), \end{aligned}$$

where the last step follows since μ is an embedding.

(d) For **Top** there are again two canonical factorisation systems: we can either take all surjective continuous maps and all embeddings, or all quotients and all injective maps.

(e) In **Θ-Set**, the only factorisation system consists of the surjective maps and the injective ones.

Let us collect a few useful properties of factorisation systems. We start with two simple remarks that save us some work. The first one tells us that it is sufficient to prove certain properties only for one of the two sets E and M . The corresponding statement for the other set then follows by duality. By the second remark, it is sufficient to define a factorisation system for \mathcal{D} . Then we can lift it to \mathcal{D}^Ξ . The proofs are straightforward.


Lemma 1.2. Let $\langle E, M \rangle$ be a factorisation system for \mathcal{C} .

- (a) $\langle M, E \rangle$ is a factorisation system for \mathcal{C}^{op} .
- (b) $\langle E^{\perp}, M^{\perp} \rangle$ is a factorisation system for \mathcal{C}^{\perp} .

Let us collect a few basic consequences of the axioms.

Lemma 1.3. Let $\langle E, M \rangle$ be a factorisation system for \mathcal{C} .

- (a) If $\mu \circ \varepsilon = \mu' \circ \varepsilon'$ for morphisms $\mu, \mu' \in M$ and $\varepsilon, \varepsilon' \in E$, then there exists a unique isomorphism σ such that $\varepsilon' = \sigma \circ \varepsilon$ and $\mu = \mu' \circ \sigma$.



- (b) A morphism belongs to $E \cap M$ if, and only if, it is an isomorphism.
- (c) $q \circ f \in E$ implies $q \in E$.
- (d) $f \circ e \in M$ implies $e \in M$.

Proof. (a) Suppose that $\mu \circ \varepsilon = \mu' \circ \varepsilon'$. By (FS5), there exists a unique morphism δ such that

$$\mu' \circ \delta = \mu \quad \text{and} \quad \delta \circ \varepsilon = \varepsilon'.$$

In the same way, $\mu' \circ \varepsilon' = \mu \circ \varepsilon$ implies that there exists a unique morphism δ' with

$$\mu \circ \delta' = \mu' \quad \text{and} \quad \delta' \circ \varepsilon' = \varepsilon.$$

We claim that δ and δ' are inverse to each other. Since $\mu \circ \varepsilon = \mu \circ \varepsilon$, there exists a unique morphism γ such that $\mu \circ \gamma = \mu$ and $\gamma \circ \varepsilon = \varepsilon$. As both id and $\delta' \circ \delta$ satisfy these two equations it follows that $\delta' \circ \delta = \text{id}$. In the same way, one can show that $\delta \circ \delta' = \text{id}$.

(b) By (FS1) and (FS2), every isomorphism belongs to both E and M . Conversely, let $\sigma \in E \cap M$. Then $\text{id} \circ \sigma = \sigma \circ \text{id}$ and, by (FS5), there exists a unique morphism δ such that $\sigma \circ \delta = \text{id}$ and $\delta \circ \sigma = \text{id}$. Hence, σ has the inverse δ .

(c) Let $q \circ f \in E$ and fix EM -factorisations $q = \mu \circ \varepsilon$ and $\varepsilon \circ f = \mu' \circ \varepsilon'$ of, respectively, q and $\varepsilon \circ f$. Then we have

$$\text{id} \circ (q \circ f) = \mu \circ \varepsilon \circ f = (\mu \circ \mu') \circ \varepsilon'.$$

By (FS5), we obtain a unique morphism δ such that

$$\delta \circ q \circ f = \varepsilon' \quad \text{and} \quad \mu \circ \mu' \circ \delta = \text{id}.$$

Hence, μ is a monomorphism with a right inverse, which implies that μ is an isomorphism. In particular, $\mu \in E$ and it follows that $q = \mu \circ \varepsilon \in E$.

(d) follows from (c) by duality. \square

It turns out that each class E and M in a factorisation system uniquely determines the other one via the diagonal fill-in property. Let us start with introducing notation for this property.

Definition 1.4. Let $\varepsilon : A \rightarrow B$ and $\mu : C \rightarrow D$ be morphisms. We write $\varepsilon \perp \mu$ if, for all morphisms $\varphi : A \rightarrow C$ and $\psi : B \rightarrow D$ with $\psi \circ \varepsilon = \mu \circ \varphi$, there exists a unique morphism $\delta : B \rightarrow C$ such that $\delta \circ \varepsilon = \varphi$ and $\mu \circ \delta = \psi$.

Lemma 1.5. Let $\langle E, M \rangle$ be a factorisation system.

- (a) $E = {}^{\perp}M := \{ \varepsilon \mid \varepsilon \perp \mu \text{ for all } \mu \in M \}$
- (b) $M = E^{\perp} := \{ \mu \mid \varepsilon \perp \mu \text{ for all } \varepsilon \in E \}$

Proof. We only need to prove (a), since (b) then follows by duality. We have $E \subseteq {}^{\perp}M$ by (FS5). Conversely, suppose that $\eta \in {}^{\perp}M$ and let $\eta = \mu \circ \varepsilon$ be its EM -factorisation. Since $\eta \perp \mu$ and $\text{id} \circ \eta = \mu \circ \varepsilon$, there exists a unique morphism δ with

$$\delta \circ \eta = \varepsilon \quad \text{and} \quad \mu \circ \delta = \text{id}.$$

Furthermore, $\varepsilon \perp \mu$ and $\mu \circ \varepsilon = \mu \circ \varepsilon$ implies that there exists a unique morphism ζ with

$$\zeta \circ \varepsilon = \varepsilon \quad \text{and} \quad \mu \circ \zeta = \mu.$$

Since $\zeta = \text{id}$ and $\zeta = \delta \circ \mu$ both satisfy these equations, it follows that $\delta \circ \mu = \text{id}$. Hence, μ is an isomorphism with inverse δ . This implies that $\mu \in E$ and, thus, $\eta = \mu \circ \varepsilon \in E$. \square

It follows in particular from this lemma that each of the two sets of a factorisation system is determined by the other one.

Exercise 1.1. Prove that $\langle E, M \rangle$ is a factorisation system if, and only if, $E = {}^\perp M$ and $M = E^\perp$. \lrcorner

Exercise 1.2. Prove that $\text{id} \perp \mu$ implies that μ is a monomorphism, and $\varepsilon \perp \text{id}$ implies that ε is an epimorphism. \lrcorner

The following remark saves us some work when checking the diagonal fill-in property.

Lemma 1.6. Suppose that $\psi \circ \varepsilon = \mu \circ \varphi$.

- (a) If ε is an epimorphism, there exists at most one morphism δ satisfying $\delta \circ \varepsilon = \varphi$ and this morphism automatically satisfies $\mu \circ \delta = \psi$.
- (b) If μ is a monomorphism, there exists at most one morphism δ satisfying $\mu \circ \delta = \psi$ and this morphism automatically satisfies $\delta \circ \varepsilon = \varphi$.

Proof. (a) As ε is an epimorphism,

$$\mu \circ \delta \circ \varepsilon = \mu \circ \varphi = \psi \circ \varepsilon \quad \text{implies} \quad \mu \circ \delta = \psi.$$

Furthermore, uniqueness of δ follows since

$$\delta' \circ \varepsilon = \varphi = \delta \circ \varepsilon \quad \text{implies} \quad \delta' = \delta.$$

(b) follows by (a) and duality. \square

Let us take a look at closure properties of E and M . Our first observation concerns morphisms with inverses. The proof is based on the relation \perp .

Lemma 1.7. Let $\langle E, M \rangle$ be a factorisation system.

- (a) Every coequaliser belongs to E .

(b) Every equaliser belongs to M .

(c) Every morphism with a right inverse belongs to E .

(d) Every morphism with a left inverse belongs to M .

Proof. By duality it is sufficient to prove (b) and (d).

(b) Let μ be the equaliser of f and g . We claim $\mu \in E^\perp$, i.e., that $\varepsilon \perp \mu$, for all $\varepsilon \in E$. Hence, let $\varepsilon \in E$ and suppose that $\mu \circ \varphi = \psi \circ \varepsilon$. Then

$$f \circ \psi \circ \varepsilon = f \circ \mu \circ \varphi = g \circ \mu \circ \varphi = g \circ \psi \circ \varepsilon$$

implies that $f \circ \psi = g \circ \psi$ since $\varepsilon \in E$ is an epimorphism. By universality of μ , we can therefore find a unique morphism δ such that $\psi = \mu \circ \delta$. Since equalisers are monomorphisms, it therefore follows by Lemma 1.6 (b) that $\varphi = \delta \circ \varepsilon$.

(d) By (b) it is sufficient to prove that every morphism $\mu : A \rightarrow B$ with a left inverse $\varepsilon : B \rightarrow A$ is an equaliser. We claim that μ is the equaliser of $\mu \circ \varepsilon$ and id_B . Clearly,

$$\mu \circ \varepsilon \circ \mu = \mu \circ \text{id} = \text{id} \circ \mu.$$

For universality, suppose that $\varphi : C \rightarrow B$ is a morphism with $\mu \circ \varepsilon \circ \varphi = \text{id} \circ \varphi$. Then $\varphi = \mu \circ (\varepsilon \circ \varphi)$ and this factorisation is unique since $\varphi = \mu \circ \psi$ implies that $\psi = \varepsilon \circ \mu \circ \psi = \varepsilon \circ \varphi$. \square

Next, we consider closure under limits.

Definition 1.8. Let F be a class of morphisms.

(a) We say that F is *closed under limits* (of a certain kind) if, for every natural transformation $\tau : D \Rightarrow E$ between two diagrams (of the given kind) with $\tau_i : D(i) \rightarrow E(i) \in F$, for all i , the canonical map $\lim D \rightarrow \lim E$ also belongs to F . Similarly, we say that F is *closed under products* if this is true for all products. Finally, F is *closed under colimits* if, for every such $\tau : D \Rightarrow E$, the canonical map $\text{colim } D \rightarrow \text{colim } E$ belongs to F .

(b) We say that F is *closed under pushouts or pullbacks* if, for every $f \in F$ and every morphism g , the pushout/pullback of f along g belongs to F . \lrcorner

Note that our terminology is not entirely consistent: closure under pullbacks is not the same as closure under limits that are pullbacks.

Lemma 1.9. *Let $\langle E, M \rangle$ be a factorisation system.*

- (a) *E is closed under colimits.*
- (b) *M is closed under limits.*
- (c) *E is closed under pushouts.*
- (d) *M is closed under pullbacks.*

Proof. Again we only prove (a) and (c), since the other two follow by duality.

(a) Let $(\kappa_i)_i$ be the limiting cocone for D and $(\lambda_i)_i$ the one for F . Suppose that $\varepsilon_i : D(i) \rightarrow F(i)$ is an E -morphism, for each index i , and let $\hat{\varepsilon} : \text{colim } D \rightarrow \text{colim } F$ be the corresponding morphism between the colimits. Consider a commuting square $\psi \circ \hat{\varepsilon} = \mu \circ \varphi$ with $\mu : A \rightarrow B$ in M . We have to find a diagonal morphism $\hat{\delta} : \text{colim } F \rightarrow A$. Applying (FS5) to the square $(\psi \circ \lambda_i) \circ \varepsilon_i = \mu \circ (\varphi \circ \kappa_i)$, we obtain a diagonal morphism $\delta_i : F(i) \rightarrow A$. Let $\hat{\delta} : \text{colim } F \rightarrow A$ be the unique morphism with $\hat{\delta} \circ \lambda_i = \delta_i$. Then

$$\hat{\delta} \circ \hat{\varepsilon} \circ \kappa_i = \hat{\delta} \circ \lambda_i \circ \varepsilon_i = \delta_i \circ \varepsilon_i = \varphi \circ \kappa_i,$$

$$\text{and } \mu \circ \hat{\delta} \circ \lambda_i = \mu \circ \delta_i = \psi \circ \lambda_i.$$

Since this holds for all i and $(\kappa_i)_i$ and $(\lambda_i)_i$ are limiting cocones, it follows that $\hat{\delta} \circ \hat{\varepsilon} = \varphi$ and $\mu \circ \hat{\delta} = \psi$.

(c) Let $\varepsilon : A \rightarrow B$ be an E -morphism, let $g : C \rightarrow D$ be the pushout of ε along $f : A \rightarrow C$, and let $f' : B \rightarrow D$ be the other morphism of the pushout square. We claim that $g \in {}^\perp M$. Consider a square $\psi \circ g = \mu \circ \varphi$ with $\mu \in M$. Applying (FS5) to the square $\psi \circ f' \circ \varepsilon = \mu \circ \varphi \circ f$, we obtain a unique morphism ζ with

$$\zeta \circ \varepsilon = \varphi \circ f \quad \text{and} \quad \mu \circ \zeta = \psi \circ f'.$$

The first equation shows that ζ and φ form a cocone for the pushout diagram. Consequently, there exists a unique morphism δ with

$$\zeta = \delta \circ f' \quad \text{and} \quad \varphi = \delta \circ g.$$

This implies that $\mu \circ \delta \circ g = \mu \circ \varphi = \psi \circ g$. Since g is an epimorphism, we obtain $\mu \circ \delta = \psi$. Furthermore, it follows by Lemma 1.6 (b) that δ is unique and that $\delta \circ \varepsilon = \varphi$. \square

Corollary 1.10. *Let $\langle E, M \rangle$ be a factorisation system where E is closed under products and M is closed under coproducts. Then M and E are closed under all polynomial functors.*

Proof. This follows from the preceding lemma since every polynomial functor is composed out of products and a coproduct. \square

We have already seen above that a factorisation system $\langle E, M \rangle$ on \mathcal{D} induces one on \mathcal{D}^Ξ . Let us show next that this factorisation system in turn induces one for \mathbb{M} -algebras.

Lemma 1.11. *Let $\langle E, M \rangle$ be a factorisation system on \mathcal{D} where E^Ξ and M^Ξ are closed under the monad \mathbb{M} . Then the sets*

$$E^{\mathbb{M}} := \{ \varepsilon \in E^\Xi \mid \varepsilon \text{ an } \mathbb{M}\text{-morphism} \},$$

$$M^{\mathbb{M}} := \{ \mu \in M^\Xi \mid \mu \text{ an } \mathbb{M}\text{-morphism} \}$$

form a factorisation system $\langle E^{\mathbb{M}}, M^{\mathbb{M}} \rangle$ on $\text{Alg}(\mathbb{M})$.

Proof. (FS1) If ε is an epimorphism of \mathcal{D}^Ξ and an \mathbb{M} -morphism, it is also an epimorphism of $\text{Alg}(\mathbb{M})$. Furthermore, every isomorphism of $\text{Alg}(\mathbb{M})$ is also an isomorphism of \mathcal{D}^Ξ .

(FS2) analogous.

(FS3) E^Ξ and M^Ξ are closed under composition.

(FS4) Let $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ be an \mathbb{M} -morphism and let $\varphi = \mu \circ \varepsilon$ be its factorisation with respect to $\langle E^\Xi, M^\Xi \rangle$. Let C be the codomain of ε . We start by equipping C with the structure of an \mathbb{M} -algebra such that ε and μ become \mathbb{M} -morphisms. Since $\mathbb{M}\varepsilon \in E^\Xi$ and

$$(\pi \circ \mathbb{M}\mu) \circ \mathbb{M}\varepsilon = \pi \circ \mathbb{M}\varphi = \varphi \circ \pi = \mu \circ (\varepsilon \circ \pi),$$

it follows by (FS5) that there exists a unique map $\sigma : \mathbb{M}C \rightarrow C$ such that

$$\mu \circ \sigma = \pi \circ \mathbb{M}\mu \quad \text{and} \quad \sigma \circ \mathbb{M}\varepsilon = \varepsilon \circ \pi.$$

Hence, Lemma I.5.5, implies that $\langle C, \sigma \rangle$ forms an \mathbb{M} -algebra and $\varepsilon : A \rightarrow C$ is a morphism of \mathbb{M} -algebras.

For μ , it now follows that

$$\pi \circ \mathbb{M}\mu \circ \mathbb{M}\varepsilon = \pi \circ \mathbb{M}\varphi = \varphi \circ \pi = \mu \circ \varepsilon \circ \pi = \mu \circ \pi \circ \mathbb{M}\varepsilon.$$

Since $\mathbb{M}\varepsilon$ is an epimorphism, this implies that $\pi \circ \mathbb{M}\mu = \mu \circ \pi$.

(FS5) Suppose that $\psi \circ \varepsilon = \mu \circ \varphi$ where $\mu \in M^{\mathbb{M}}$ and $\varepsilon \in E^{\mathbb{M}}$. Let δ be the unique function with

$$\mu \circ \delta = \psi \quad \text{and} \quad \delta \circ \varepsilon = \varphi.$$

Since ε and φ are \mathbb{M} -morphisms, it follows by Lemma I.5.6 that so is δ . \square

The Standard Factorisation System

The following factorisation system is the one we will use in the rest of this book.

Definition 1.12. Let \mathcal{D} be a category.

(a) The *standard factorisation system* $\langle E, M \rangle$ for \mathcal{D} consists of

$$E := \{ \varepsilon \mid \mathbb{V}\varepsilon \text{ surjective} \} \quad \text{and} \quad M := E^\perp.$$

(b) We say that the standard factorisation system of \mathcal{D} is *well-defined* if the standard factorisation system is indeed a factorisation system. \lrcorner

In general, the standard factorisation system does not need to be a factorisation system, but in the examples we are interested in it is.

Remark. For each of the categories Set, Pos, Top, Met, and \mathfrak{G} -Set, we have already seen factorisation systems where E consists of all surjective morphisms. Consequently, in every of these categories the standard factorisation system is well-defined. \lrcorner

The reason why we are interested in the standard factorisation system is that it has several special properties. Of particular importance will be the closure of E under limits and pullbacks.

Lemma 1.13. Let $\langle E, M \rangle$ be the standard factorisation system on \mathcal{D} .

- (a) E is closed under limits.
- (b) E is closed under pullbacks.

Proof. (a) Let $\tau : C \Rightarrow D$ be a natural transformation between two diagrams where $\tau_i \in E$, for all indices i . Let $(\lambda_i)_i$ and $(\mu_i)_i$ be the limiting cones of, respectively, C and D , and let $\varepsilon : \lim C \rightarrow \lim D$ be the unique morphism with $\tau_i \circ \lambda_i = \mu_i \circ \varepsilon$. Since \mathbb{V} preserves limits, it follows that $(\mathbb{V}\lambda_i)_i$ and $(\mathbb{V}\mu_i)_i$ are the limiting cones of $\mathbb{V} \circ C$ and $\mathbb{V} \circ D$, and that $\mathbb{V}\varepsilon : \lim (\mathbb{V} \circ C) \rightarrow \lim (\mathbb{V} \circ D)$ satisfies

$$\mathbb{V}\tau_i \circ \mathbb{V}\lambda_i = \mathbb{V}\mu_i \circ \mathbb{V}\varepsilon.$$

We claim that $\varepsilon \in E$. By Lemma I.2.3, it is sufficient to prove that $\mathbb{V}\varepsilon$ is surjective.

Hence, let b be an element of $\lim D$. Set $b_i := \mu_i(b) \in D(i)$. By surjectivity of τ_i , there is some $a_i \in C(i)$ with $\tau_i(a_i) = b_i$. For every morphism $f : i \rightarrow j$ of the index category, it follows that

$$\begin{aligned} (\tau_j \circ Cf)(a_i) &= (Df \circ \tau_i)(a_i) \\ &= Df(b_i) \\ &= (Df \circ \mu_i)(b) = \mu_j(b) = b_j = \tau_j(a_j). \end{aligned}$$

As τ_j is surjective, this implies that $Cf(a_i) = a_j$. Hence, there is some $a \in \lim C$ with $\lambda_i(a) = a_i$, for all i . It follows that

$$\mu_i(\varepsilon(a)) = \tau_i(\lambda_i(a)) = \tau_i(a_i) = b_i, \quad \text{for all } i.$$

This implies that $\varepsilon(a) = b$. Hence, $b \in \text{rng } \varepsilon$.

(b) Let $\varepsilon \in E$ and let ε' be the pullback of ε along some morphism f . Since \mathbb{V} preserves limits, it follows that $\mathbb{V}\varepsilon'$ is the pullback of $\mathbb{V}\varepsilon$ along $\mathbb{V}f$. In Set pullbacks of surjective functions are surjective. Hence, $\mathbb{V}\varepsilon'$ is surjective, which implies by Lemma I.2.3 that ε' is a surjective epimorphism. Hence, $\varepsilon' \in E$. \square

2. Subalgebras

Using a factorisation system we can now define subobjects and quotients. We start with the former. In general, we can represent a subobject of A by an M -morphism $\mu : C \rightarrow A$. For discrete categories, we can be a bit more concrete.

Definition 2.1. Let $A \in \mathcal{D}$, $X \subseteq \mathbb{V}A$, and let $j : \mathbb{J}X \rightarrow A$ be the morphism corresponding to the inclusion map $X \rightarrow \mathbb{V}A$ via the adjunction.

(a) The *subobject* of A generated by X is the domain of the image

$$\langle\langle X \rangle\rangle_A := \text{dom}(\text{im } j).$$


(b) Similarly, for an \mathbb{M} -algebra \mathfrak{A} and $X \subseteq A$, we define the *subalgebra* of \mathfrak{A} generated by X by

$$\langle\langle X \rangle\rangle_{\mathfrak{A}} := \langle C, \pi_o \rangle \quad \text{where} \quad C := \text{dom}(\text{im}(\pi \circ \mathbb{M}j))$$

and the product π_o is chosen according to the following lemma.

Lemma 2.2. Let \mathfrak{A} be an \mathbb{M} -algebra, $X \subseteq A$, and let C be the domain of $\langle\langle X \rangle\rangle_{\mathfrak{A}}$. There exists a function $\pi_o : \mathbb{M}C \rightarrow C$ turning $\langle\langle X \rangle\rangle_{\mathfrak{A}} = \langle C, \pi_o \rangle$ into an \mathbb{M} -algebra and the inclusion map $\langle\langle X \rangle\rangle_{\mathfrak{A}} \rightarrow \mathfrak{A}$ into a morphism of \mathbb{M} -algebras. Furthermore, $\mathbb{V}C = \pi[\mathbb{V}X]$.

Proof. Let $\mu \circ \varepsilon$ be the EM -factorisation of $\pi \circ \mathbb{M}j : \mathbb{M}X \rightarrow A$ and let C be the domain of μ . Similarly, let $\mu' \circ \varepsilon'$ be the EM -factorisation of $\pi \circ \mathbb{M}\mu : \mathbb{M}C \rightarrow A$ and let C' be the domain of μ' .



Set $i := \varepsilon \circ \text{sing}$ and $i' := \varepsilon' \circ \text{sing}$. Then

$$\begin{aligned} \mu \circ i &= \mu \circ \varepsilon \circ \text{sing} = \pi \circ \mathbb{M}j \circ \text{sing} = \pi \circ \text{sing} \circ j = j, \\ \mu' \circ i' &= \mu' \circ \varepsilon' \circ \text{sing} = \pi \circ \mathbb{M}\mu \circ \text{sing} = \pi \circ \text{sing} \circ \mu = \mu. \end{aligned}$$

We claim that i' is an isomorphism. For the proof, note that the two EM -factorisations

$$\mu \circ (\varepsilon \circ \text{flat}) = \mu' \circ (\varepsilon' \circ \mathbb{M}\varepsilon)$$

must be isomorphic. Hence, there exists a morphism $k : C' \rightarrow C$ such that

$$\mu \circ k = \mu' \quad \text{and} \quad \varepsilon \circ \text{flat} = k \circ (\varepsilon' \circ \mathbb{M}\varepsilon).$$

In particular, it follows that

$$\mu \circ k \circ i' = \mu' \circ i' = \mu \quad \text{and} \quad \mu' \circ i' \circ k = \mu \circ k = \mu'.$$

Since μ and μ' are monomorphisms this implies that k and i' are inverses of each other. We define the product of $\langle\langle X \rangle\rangle_{\mathfrak{A}}$ by $\pi_o := k \circ \varepsilon'$.

We start by checking that $\langle\langle X \rangle\rangle_{\mathfrak{A}}$ is an \mathbb{M} -algebra. For the unit law, we have

$$\pi_o \circ \text{sing} = k \circ \varepsilon' \circ \text{sing} = k \circ i' = \text{id}.$$

For the associative law, note that

$$\mu \circ \pi_o = \mu \circ k \circ \varepsilon' = \mu' \circ \varepsilon' = \pi \circ \mathbb{M}\mu$$

implies that

$$\begin{aligned} \mu \circ \pi_o \circ \text{flat} &= \pi \circ \mathbb{M}\mu \circ \text{flat} \\ &= \pi \circ \text{flat} \circ \mathbb{M}\mu \\ &= \pi \circ \mathbb{M}\pi \circ \mathbb{M}\mu \\ &= \pi \circ \mathbb{M}(\mu \circ \pi_o) = \mu \circ \pi_o \circ \mathbb{M}\pi_o. \end{aligned}$$

Since $\mu \in M$, it follows that $\pi_o \circ \text{flat} = \pi_o \circ \mathbb{M}\pi_o$. Finally, to see that the inclusion $\mu : C \rightarrow A$ is a morphism of \mathbb{M} -algebras, note that

$$\mu \circ \pi_o = \mu \circ k \circ \varepsilon' = \mu' \circ \varepsilon' = \pi \circ \mathbb{M}\mu. \quad \square$$

The concept of a subobject generated by some set immediately provides a notion of cardinality for objects: the minimal cardinality of a generating set. The aim of the rest of this section is to link this cardinality to purely category theoretical properties. The usual approach is in terms of so-called *accessible categories* and *κ -presentable objects*. Here, we will use a slightly non-standard formalisation that better fits with the categories we are interested in.

We start by studying how objects are built up from their subobjects. It turns out that the axioms of a discrete category are not quite strong enough, to make this notion as well-behaved as we would like.

Lemma 2.3. *For every $A \in \mathcal{D}$ and every regular cardinal κ , there exists a κ -directed diagram $D : \mathcal{I} \rightarrow \mathcal{D}$ and a bijective morphism $\sigma : \text{colim } D \rightarrow A$.*

Proof. Let \mathcal{I} consist of all subsets $X \subseteq \mathbb{V}A$ of size $|X| < \kappa$ ordered by inclusion and set $D(X) := \langle\langle X \rangle\rangle_A$. Then \mathcal{I} is κ -directed because κ is regular. To find the desired morphism σ , let $(\lambda_X)_X$ be the limiting cocone of D and let $j_X : \mathbb{J}X \rightarrow A$ be the inclusion morphism. Then $(j_X)_X$ forms a cocone

from D to A . Since $(\lambda_X)_X$ is limiting, there exists a unique morphism $\sigma : \text{colim } D \rightarrow A$ with

$$\sigma \circ \lambda_X = j_X, \quad \text{for all } X \in \mathcal{I}.$$

It remains to show that $\mathbb{V}\sigma$ is bijective. Surjectivity follows from the fact that $\text{colim}(\mathbb{V} \circ D) \subseteq \bigcup_X X = \mathbb{V}A$. For injectivity, fix two elements $a, b \in \mathbb{V} \text{colim } D$ with $a \neq b$. As \mathcal{I} is κ -filtered, we can find some index $X \in \mathcal{I}$ with $a, b \in \text{rng } \mathbb{V}\lambda_X = \mathbb{V}\lambda_X[\mathbb{V}\langle\langle X \rangle\rangle_A]$. Consequently, there are $a', b' \in \langle\langle X \rangle\rangle_A$ with $a = \lambda_X(a')$ and $b = \lambda_X(b')$. Hence,

$$a \neq b \quad \text{implies} \quad \sigma(a) = a' \neq b' = \sigma(b). \quad \square$$

To continue, we have to assume that σ is actually an isomorphism. This leads to the following definition.

Definition 2.4. Let \mathcal{D} be a discrete category that has a standard factorisation system (E, M) and let κ a regular cardinal.

(a) The *canonical κ -subobject diagram* $D : \mathcal{I} \rightarrow \mathcal{D}$ of an object $A \in \mathcal{D}$ is defined as follows. The index category \mathcal{I} is the subcategory of $(\mathcal{D} \downarrow A)$ consisting of all morphisms $i : C \rightarrow A$ with $i \in M$ and $C = \langle\langle X \rangle\rangle_C$, for some set X of size $|X| < \kappa$. The functor D maps an object $i : C \rightarrow A$ to its domain C and a morphism $\gamma : i \rightarrow j$ to itself (now regarded as a morphism between the corresponding domains).

(b) We say that \mathcal{D} *has canonical subobject diagrams* if, for every regular cardinal κ and every $A \in \mathcal{D}$ with canonical κ -subobject diagram $D : \mathcal{I} \rightarrow \mathcal{D}$, the morphisms $i \in \mathcal{I}$ form a limiting cocone from D to A . \lrcorner

Example. The categories **Set**, **Pos**, **Met**, and **\mathfrak{G} -Set** all have canonical subobject diagrams. **Top** does not. A counterexample is the space \mathfrak{X} where $X = \kappa$ and a set $U \subseteq X$ is closed if, and only if, $|U| < \kappa$. Every subspace of \mathfrak{X} of size less than κ is discrete. Hence, the colimit of the canonical κ -subobject diagram has the discrete topology. \lrcorner

In categories with canonical subobject diagrams, we can characterise the cardinality of an object in the following way.

Definition 2.5. Let κ a regular cardinal and \mathcal{C} a category with a factorisation system $\langle E, M \rangle$.

(a) Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram and $(\mu_i)_i$ a cocone from D to some object B . We say that a morphism $\varphi : A \rightarrow B$ *factorises through* $(\mu_i)_i$ if there exists an index $i \in \mathcal{I}$ such that

$$\varphi = \mu_i \circ \psi, \quad \text{for some morphism } \psi : A \rightarrow D(i).$$

We say that such a factorisation is *essentially unique* if, given two morphisms $\psi, \psi' : A \rightarrow D(i)$ with $\mu_i \circ \psi = \varphi = \mu_i \circ \psi'$, there exist a morphism $f : i \rightarrow k$ of \mathcal{I} such that $Df \circ \psi = Df \circ \psi'$.

(b) A diagram $D : \mathcal{I} \rightarrow \mathcal{C}$ is κ -*M-generating* if it has a limiting cocone $(\lambda_i)_i$ such that

$$\bigoplus_{i \in \mathcal{I}} \lambda_i \text{ is surjective} \quad \text{and} \quad \lambda_i \in M, \quad \text{for all } i \in \mathcal{I}.$$

(c) An object $A \in \mathcal{C}$ is κ -*M-generated* if, for every diagram $D : \mathcal{I} \rightarrow \mathcal{C}$ that is κ -filtered and κ -*M-generating*, every morphism $f : A \rightarrow \text{colim } D$ factorises essentially uniquely through $(\lambda_i)_i$. \lrcorner

Remark. Note that the canonical κ -subobject diagram is κ -*M-generating* and κ -directed. \lrcorner

Proposition 2.6. Let κ be a regular cardinal and let \mathcal{D} be a discrete category that has canonical subobject diagrams. An object $A \in \mathcal{D}$ is κ -*M-generated* if, and only if, it is of the form $\langle\langle X \rangle\rangle_A$, for some set X of size $|X| < \kappa$.

Proof. Suppose that A is κ -*M-generated*. By assumption there exists a κ -filtered diagram $D : \mathcal{I} \rightarrow \mathcal{D}$ such that $\text{colim } D = A$ and $D(i) = \langle\langle X_i \rangle\rangle_A$ with $|X_i| < \kappa$, for all i . This diagram is κ -*M-generating*. Let $(\lambda_i)_i$ be the corresponding limiting cocone. By assumption, the identity map $\text{id} : A \rightarrow A$ factorises through $(\lambda_i)_i$. Hence, there is some $i \in \mathcal{I}$ and some function $\gamma : A \rightarrow D(i)$ such that

$$\text{id} = \lambda_i \circ \gamma.$$

Thus, g has a left inverse, which means that $\mathbb{V}g$ is injective. Fix a set $X \subseteq A$ of size $|X| < \kappa$ such that $D(i) = \langle\langle X \rangle\rangle_A$. Let $j : X \rightarrow \mathbb{V}D(i)$ be the inclusion map and $j_* : \mathbb{J}X \rightarrow D(i)$ the morphism associated with it via the adjunction $\mathbb{J} \dashv \mathbb{V}$. Since $D(i) = \langle\langle X \rangle\rangle_A$, we have $j_* \in E$. As λ_i has a right inverse, it also belongs to E . Consequently, $\lambda_i \circ j_* \in E$, which implies that $A = \langle\langle \mathbb{V}\lambda_i[X] \rangle\rangle_A$. Since $|\mathbb{V}\lambda_i[X]| \leq |X| < \kappa$, the claim follows.

(\Leftarrow) Suppose that $A = \langle\langle X \rangle\rangle_A$, for some set of size $|X| < \kappa$. Let $j : X \rightarrow \mathbb{V}A$ be the corresponding inclusion morphism and $j_* : \mathbb{J}X \rightarrow A$ its associate via the adjunction $\mathbb{J} \dashv \mathbb{V}$. By assumption, j_* is an isomorphism. To show that A is κ -*M-generated*, we consider a diagram $D : \mathcal{I} \rightarrow \text{Set}$ that is κ -filtered and κ -*M-generating*. Let B be its colimit and $(\lambda_i)_i$ the corresponding limiting cocone. Given a morphism $f : A \rightarrow B$, we set $f_o := \mathbb{V}f \circ j : X \rightarrow \mathbb{V}B$. To show that f factorises through $(\lambda_i)_i$ we pick, for every $x \in X$, some index $i_x \in \mathcal{I}$ such that $f_o(x) \in \text{rng } \mathbb{V}\lambda_{i_x}$. (This is possible since surjectivity of $\bigoplus_i \lambda_i$ implies that $\mathbb{V}B = \bigcup_i \text{rng } \mathbb{V}\lambda_i$.) Since \mathcal{I} is κ -filtered, we can find some index $k \in \mathcal{I}$ and morphisms $g_x : i_x \rightarrow k$. Then

$$f_o(x) \in \text{rng } \mathbb{V}\lambda_{i_x} = \text{rng } \mathbb{V}(\lambda_k \circ Dg_x) \subseteq \text{rng } \mathbb{V}\lambda_k, \quad \text{for all } x \in X.$$

Consequently, there exists a function $f' : X \rightarrow \mathbb{V}D(k)$ such that

$$\mathbb{V}\lambda_k(f'(x)) = f_o(x), \quad \text{for all } x \in X.$$

Via the adjunction $\mathbb{J} \dashv \mathbb{V}$ this function induces a morphism $f_* : \mathbb{J}X \rightarrow D(k)$. It follows that

$$\mathbb{V}\lambda_k \circ f' = f_o = \mathbb{V}f \circ j \quad \text{is associated with} \quad \lambda_k \circ f_* = f \circ j_*.$$

As j_* is an isomorphism, we have $j_* \in E$. Since $\lambda_k \in M$, we can therefore use the diagonal fill-in property to find a morphism $\delta : A \rightarrow D(k)$ such that

$$\delta \circ j_* = f_* \quad \text{and} \quad \lambda_k \circ \delta = f.$$

Hence, f factorises through $(\lambda_i)_i$ and it remains to show that this factorisation is essentially unique.

Suppose that $f = \lambda_k \circ g$, for some $g : A \rightarrow D(k)$. For every $x \in X$, we have

$$(\mathbb{V}\lambda_k \circ \mathbb{V}\delta \circ j)(x) = (\mathbb{V}f \circ j)(x) = (\mathbb{V}\lambda_k \circ \mathbb{V}g \circ j)(x).$$

As $\lambda_k \in M$, it follows that $\mathbb{V}\delta \circ j = \mathbb{V}g \circ j$, which, via the adjunction, implies that $\delta \circ j_* = g \circ j_*$. As j_* is an isomorphism, we obtain $\delta = g$. \square

This results applies to Set, Pos, Met, and \mathfrak{G} -Set, but not to Top. For topological spaces, we obtain the following characterisation.

Example. In Top, a space \mathfrak{X} is κ -presentable if, and only if, $|X| < \kappa$ and \mathfrak{X} is discrete. The implication (\Leftarrow) follows using the same proof as in the preceding proposition. For (\Rightarrow) , suppose that \mathfrak{X} is κ -presentable. We only have to show that \mathfrak{X} is discrete. Then it follows that $|X| < \kappa$ using the same proof as for Set. Let \mathfrak{X} be a topological space of size that is not discrete and fix a subset $P \subseteq X$ that is not closed. Let $D : \kappa \rightarrow \text{Top}$ be the diagram with $D(i) := X + \kappa$ where a set U is closed if, and only if,

$$U \cap X \text{ is closed in } \mathfrak{X} \quad \text{or} \quad \downarrow i \subseteq U.$$

For $i < k < \kappa$, we choose for $D(i, k)$ the identity map. Then $\text{colim } D = X + \kappa$ where the closed sets are exactly those which are closed in every $D(i)$. Hence, the only closed sets are $X + \kappa$ and those which are closed in \mathfrak{X} . In particular, the inclusion $X \rightarrow X + \kappa$ is continuous, but it does not factorise through any $D(i)$ since the set P is closed in $D(i)$ but not in \mathfrak{X} . \lrcorner

Since we have used κ -filtered colimits to characterise cardinality, it is of interest to know whether a given functor preserves such colimits. Our aim is to prove that this is the case for all polynomial functors. Unfortunately, in the general case, we only obtain a slightly weaker statement involving a bijective morphism instead of an isomorphism.

Definition 2.7. (a) We say that a functor $\mathbb{F} : \mathcal{C} \rightarrow \mathcal{D}$ preserves limits of a certain kind if, whenever $D : \mathcal{I} \rightarrow \mathcal{C}$ is a diagram of this type and $(\lambda_i)_i$ is a limiting cocone of D , then $(\mathbb{F}\lambda_i)_i$ is a limiting cocone of $\mathbb{F} \circ D$.

(b) Similarly, we say that \mathbb{F} reflects limits of this kind if, whenever $D : \mathcal{I} \rightarrow \mathcal{C}$ is a diagram of this type and $(\mathbb{F}\lambda_i)_i$ is a limiting cocone of $\mathbb{F} \circ D$, then $(\lambda_i)_i$ is a limiting cocone of D .

(c) We use the same terminology for preserving and reflecting colimits. \lrcorner

Example. The forgetful functors from Pos, Top, Met, and \mathfrak{G} -Set to Set preserve and reflect κ -filtered colimits. \lrcorner

The case for arbitrary polynomial functors follows once we can prove the statement for power operators. For a general discrete category, we have the following statement.

Proposition 2.8. Let \mathcal{D} be a discrete category, $D : \mathcal{I} \rightarrow \mathcal{D}^\Xi$ a κ -filtered diagram, and X a set of size $|X| < \kappa$. There exists a unique bijective morphism

$$\sigma : \text{colim}(D^X) \rightarrow (\text{colim } D)^X$$

satisfying

$$\sigma \circ \mu_i = \lambda_i^X, \quad \text{for all } i \in \mathcal{I},$$

where $(\lambda_i)_i$ is the limiting cocone of D and $(\mu_i)_i$ the one of D^X .

Proof. Set $A := \text{colim } D$ and $B := \text{colim } D^X$. Then $(\lambda_i^X)_i$ is a cocone from D^X to A^X . As $(\mu_i)_i$ is limiting, there exists a unique morphism $\sigma : B \rightarrow A^X$ with

$$\sigma \circ \mu_i = \lambda_i^X, \quad \text{for all } i \in \mathcal{I}.$$

It remains to show that $\mathbb{V}\sigma$ is bijective.

For surjectivity, consider a sequence $(a_x)_{x \in X} \in A^X$. Since $a_x \in A = \text{colim } D$, we can find, for every $x \in X$, some index $i \in \mathcal{I}$ and an element $b \in D(i)$ with $a_x = \lambda_i(b)$. As $|X| < \kappa$ and \mathcal{I} is κ -directed, there is therefore some index $k \in \mathcal{I}$ with $\{a_x \mid x \in X\} \subseteq \lambda_k[D(k)]$. Consequently, we have $(a_x)_x = \lambda_k^X(b)$ for some $b \in D(k)^X$. It follows that

$$(a_x)_x = \lambda_k^X(b) = (\sigma \circ \mu_k)(b) \in \text{rng } \sigma.$$

For injectivity, consider elements $b, b' \in B$ with $\sigma(b) = \sigma(b')$. Since $b, b' \in \text{colim } D^X$, we can find indices $i, i' \in \mathcal{I}$ with $b \in \mu_i[D(i)^X]$ and $b' \in \mu_{i'}[D(i')^X]$. As \mathcal{I} is κ -directed, we may assume that $i = i'$. Fix $a, a' \in D(i)^X$ with $b = \mu_i(a)$ and $b' = \mu_i(a')$. Then

$$\lambda_i^X(a) = (\sigma \circ \mu_i)(a) = \sigma(b) = \sigma(b') = (\sigma \circ \mu_i)(a') = \lambda_i^X(a').$$

For $x \in X$, it follows that

$$\lambda_i(p_x(a)) = p_x(\lambda_i^X(a)) = p_x(\lambda_i^X(a')) = \lambda_i(p_x(a')).$$

Consequently, we can find, for every $x \in X$, some \mathcal{I} -morphism $f_x : i \rightarrow k_x$ with

$$Df_x(p_x(a)) = Df_x(p_x(a')).$$

As \mathcal{I} is κ -filtered and $|X| < \kappa$, we may assume that $k_x = k_y$ and $f_x = f_y$, for all $x, y \in X$. Let us denote this morphism by $f : i \rightarrow k$. Then $D^X f(a) = D^X f(a')$, which implies that

$$b = \mu_i(a) = \mu_k(D^X f(a)) = \mu_k(D^X f(a')) = \mu_i(a') = b'. \quad \square$$

As above, if our category has canonical subobject diagrams, we can improve the result as follows.

Proposition 2.9. *Let \mathcal{D} be a discrete category with canonical subobject diagrams and let X be a set of size $|X| < \kappa$. The functor $(-)^X : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ preserves κ -filtered colimits.*

Proof. Let $D : \mathcal{I} \rightarrow \mathcal{D}^\Xi$ be a κ -filtered diagram with colimit $A := \text{colim } D$ and limiting cocone $(\lambda_i)_i$. Set $B := \text{colim } D^X$ and let $(\mu_i)_i$ be the limiting cocone from D^X to B . We have to show that $B \cong A^X$.

As $(\lambda_i^X)_i$ is a cocone from D^X to A^X , there exists a unique morphism $\sigma : B \rightarrow A^X$ such that

$$\sigma \circ \mu_i = \lambda_i^X, \quad \text{for all } i \in \mathcal{I}.$$

We claim that σ is the desired isomorphism. We construct its inverse as follows.

Note that $\mathbb{V}\sigma$ is bijective by Proposition 2.8. If we can find a right inverse $\tau : A^X \rightarrow B$ of σ , then

$$\mathbb{V}\sigma \circ \mathbb{V}\tau = \mathbb{V}(\sigma \circ \tau) = \text{id}$$

implies by bijectivity of $\mathbb{V}\sigma$ that $\mathbb{V}\tau = (\mathbb{V}\sigma)^{-1}$. Hence,

$$\mathbb{V}(\tau \circ \sigma) = \mathbb{V}\tau \circ \mathbb{V}\sigma = \text{id}.$$

As \mathbb{V} is faithful, we obtain $\tau \circ \sigma = \text{id}$. Hence, σ is an isomorphism with inverse τ .

It remains to construct τ . Let $E : \mathcal{K} \rightarrow \mathcal{D}$ be the canonical κ -subobject diagram of A^X . By assumption, we have $\text{colim } E = A^X$. Let $(\kappa_k)_k$ be the corresponding limiting cocone and let $p_x : A^X \rightarrow A$ be the projection to the x -th component. We obtain morphisms $p_x \circ \kappa_k : E(k) \rightarrow A$, for all $k \in \mathcal{K}$ and $x \in X$. By Proposition 2.6, each $E(k)$ is κ - M -generated. Hence, these morphisms factorise essentially uniquely as $p_x \circ \kappa_k = \lambda_{i_{k,x}}^x \circ \eta_k^x$, for some $i_{k,x} \in \mathcal{I}$ and some morphism $\eta_k^x : E(k) \rightarrow D(i_{k,x})$.

$$\begin{array}{ccccc} D^X(i(k)) & \xleftarrow{\eta_k} & E(k) & \xrightarrow{\eta_k^x} & D(i(k)) \\ \mu_{i(k)} \downarrow & \searrow \lambda_{i(k)}^X & \downarrow \kappa_k & & \downarrow \lambda_{i(k)} \\ B & \xrightarrow{\sigma} & A^X & \xrightarrow{p_x} & A \end{array}$$

Since \mathcal{K} is κ -directed, there exists some index $i(k) \geq i_{k,x}$ for all $x \in X$. It follows that, for every $k \in \mathcal{K}$, we find factorisations of the form

$$p_x \circ \kappa_k = \lambda_{i(k)}^x \circ \eta_k^x, \quad \text{for all } x \in X.$$

Set $\eta_k := \langle \eta_k^x \rangle_{x \in X} : E(k) \rightarrow D(i(k))^X$. Then

$$\begin{aligned} \lambda_{i(k)}^X \circ \eta_k &= \lambda_{i(k)}^X \circ \langle \eta_k^x \rangle_x \\ &= \langle \lambda_{i(k)}^X \circ \eta_k^x \rangle_x \\ &= \langle p_x \circ \kappa_k \rangle_x \\ &= \langle p_x \rangle_x \circ \kappa_k = \text{id} \circ \kappa_k = \kappa_k. \end{aligned}$$

We obtain morphisms $\mu_{i(k)} \circ \eta_k : E(k) \rightarrow B$. To see that these form a cocone from E to B , consider indices $k \leq l$ in \mathcal{K} . As \mathcal{I} is κ -filtered, we can fix morphisms $f : i(k) \rightarrow m$ and $g : i(l) \rightarrow m$ in \mathcal{I} for some m . Then

$$\begin{aligned} \lambda_m \circ Dg \circ \eta_l^x \circ E(k, l) &= \lambda_{i(l)} \circ \eta_l^x \circ E(k, l) \\ \text{and } \lambda_m \circ Df \circ \eta_k^x &= \lambda_{i(k)} \circ \eta_k^x \end{aligned}$$

are two factorisations of $p_x \circ \kappa_k$. By essential uniqueness, it follows that there is some morphism $h : m \rightarrow n$ in \mathcal{I} such that

$$Dh \circ Dg \circ \eta_l^x \circ E(k, l) = Dh \circ Df \circ \eta_k^x.$$

Hence,

$$\begin{aligned} \mu_{i(l)} \circ \eta_l \circ E(k, l) &= \mu_n \circ D^X(h \circ g) \circ \eta_l \circ E(k, l) \\ &= \mu_n \circ D^X(h \circ f) \circ \eta_k \\ &= \mu_{i(k)} \circ \eta_k, \end{aligned}$$

as desired.

Since A^X is the colimit of E , we therefore obtain a unique morphism $\tau : A^X \rightarrow B$ with

$$\tau \circ \kappa_k = \mu_{i(k)} \circ \eta_k, \quad \text{for all } k.$$

It follows that

$$\sigma \circ \tau \circ \kappa_k = \sigma \circ \mu_{i(k)} \circ \eta_k = \lambda_{i(k)}^X \circ \eta_k = \kappa_k.$$

As limiting cocones are jointly epimorphic, it follows that $\sigma \circ \tau = \text{id}$. Hence, τ is indeed the right inverse of σ . \square

We obtain the following corollaries for polynomial functors.

Corollary 2.10. *Let κ be a regular cardinal and \mathcal{D} a discrete category with canonical κ -subobject diagrams. Every polynomial functor $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ of arity at most κ preserves κ -filtered colimits.*

Proof. Let $D : \mathcal{I} \rightarrow \mathcal{D}^\Xi$ be κ -filtered and suppose that $\mathbb{F}X = \sum_{k \in K} X^{\text{dom}(k)}$. Then

$$\begin{aligned} \text{colim}_i(\mathbb{F}D(i)) &= \text{colim}_i \sum_{k \in K} D(i)^{\text{dom}(k)} \\ &= \sum_{k \in K} \text{colim}_i D(i)^{\text{dom}(k)} \\ &= \sum_{k \in K} (\text{colim } D)^{\text{dom}(k)} = \mathbb{F}(\text{colim } D), \end{aligned}$$

where the third step follows by Proposition 2.9. \square

Without canonical subobject diagrams, we only obtain the following statements.

Corollary 2.11. *Let \mathcal{D} be a discrete category, κ a regular cardinal, and $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ a polynomial functor of arity at most κ . For every κ -filtered diagram $D : \mathcal{I} \rightarrow \mathcal{D}^\Xi$, there exists a unique bijective morphism*

$$\sigma : \text{colim}(\mathbb{F} \circ D) \rightarrow \mathbb{F}(\text{colim } D).$$

Proof. Let $D : \mathcal{I} \rightarrow \mathcal{D}^\Xi$ be κ -filtered and suppose that $\mathbb{F}X = \sum_{k \in K} X^{\text{dom}(k)}$. Then

$$\begin{aligned} \text{colim}_i(\mathbb{F}D(i)) &= \text{colim}_i \sum_{k \in K} D(i)^{\text{dom}(k)} = \sum_{k \in K} \text{colim}_i D(i)^{\text{dom}(k)} \\ \text{and } \mathbb{F}(\text{colim } D) &= \sum_{k \in K} (\text{colim } D)^{\text{dom}(k)}. \end{aligned}$$

We can therefore obtain the desired morphism from the former to the latter by Proposition 2.8. \square

A bijective morphism is sufficient to prove the following corollary about the elements of $\mathbb{F}A$.

Corollary 2.12. *Let \mathcal{D} be a discrete category, κ a regular cardinal, and $\mathbb{F} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ a polynomial functor of arity at most κ . Then*

$$\mathbb{V}\mathbb{F}A = \bigcup \{ \mathbb{F}C \mid C \subseteq \mathbb{V}A, |C| < \kappa \}.$$

3. Reducts

Besides adding or removing elements of an algebra we can also add or remove operations or sorts. We start by taking a look at the former. Removing operations corresponds to replacing the product $\pi : \mathbb{M}A \rightarrow A$ by a restriction $\pi \upharpoonright \mathbb{M}^\circ A : \mathbb{M}^\circ A \rightarrow A$ where $\mathbb{M}^\circ A \subseteq \mathbb{M}A$.

Definition 3.1. Let $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ be a morphism of monads and $\mathfrak{A} = \langle A, \pi \rangle$ an \mathbb{M} -algebra.

(a) The ρ -*reduct* of \mathfrak{A} is the \mathbb{M}° -algebra $\mathfrak{A}|_\rho := \langle A, \pi \circ \rho \rangle$. If ρ is understood, we also speak of the \mathbb{M}° -*reduct* of \mathfrak{A} .

(b) For an \mathbb{M} -morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$, we define $\varphi|_\rho : \mathfrak{A}|_\rho \rightarrow \mathfrak{B}|_\rho$ by

$$\varphi|_\rho := \varphi.$$

Lemma 3.2. *Let $\rho : \langle \mathbb{M}^\circ, \mu^\circ, \varepsilon^\circ \rangle \rightarrow \langle \mathbb{M}, \mu, \varepsilon \rangle$ be morphism of monads. Then $|_\rho : \text{Alg}(\mathbb{M}) \rightarrow \text{Alg}(\mathbb{M}^\circ)$ is a functor.*

Proof. Let \mathfrak{A} be an \mathbb{M} -algebra. Then $\mathfrak{A}|_\rho$ is an \mathbb{M}° -algebra since, by the axioms of a morphism of monads,

$$(\pi \circ \rho) \circ \varepsilon^\circ = \pi \circ \varepsilon = \text{id},$$

$$\begin{aligned} (\pi \circ \rho) \circ \mu^\circ &= \pi \circ \mu \circ \rho \circ \mathbb{M}^\circ \rho \\ &= \pi \circ \mathbb{M} \pi \circ \rho \circ \mathbb{M}^\circ \rho = (\pi \circ \rho) \circ \mathbb{M}^\circ (\pi \circ \rho). \end{aligned}$$

Similarly, let $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ be an \mathbb{M} -morphism. Then $\varphi|_\rho$ is an \mathbb{M}° -morphism since

$$\varphi \circ (\pi \circ \rho) = \pi \circ \mathbb{M} \varphi \circ \rho = \pi \circ (\rho \circ \mathbb{M}^\circ \varphi).$$

Finally, the fact that $(\varphi \circ \psi)|_\rho = \varphi|_\rho \circ \psi|_\rho$ follows immediately from the definition. \square

Instead of operations, we can also add or remove sorts to an algebra.

Definition 3.3. Let $\Delta \subseteq \Xi$ be sets of sorts.

(a) The *expansion* of $A = (A_\xi)_{\xi \in \Delta} \in \mathcal{D}^\Delta$ to \mathcal{D}^Ξ is the set $A \upharpoonright \Xi \in \mathcal{D}^\Xi$ defined by

$$(A \upharpoonright \Xi)_\xi := \begin{cases} A_\xi & \text{if } \xi \in \Delta, \\ 0 & \text{otherwise,} \end{cases}$$

where 0 denotes the initial object of \mathcal{D} . Similarly, we define the expansion $f \upharpoonright \Xi : A \upharpoonright \Xi \rightarrow B \upharpoonright \Xi$ of a morphism $f = (f_\xi)_{\xi \in \Delta} : A \rightarrow B$ by

$$(f \upharpoonright \Xi)_\xi := \begin{cases} f_\xi & \text{if } \xi \in \Delta, \\ \text{id}_0 & \text{otherwise.} \end{cases}$$

(b) The Δ -*reduct* of $A = (A_\xi)_{\xi \in \Xi} \in \mathcal{D}^\Xi$ is the set

$$A|_\Delta := (A_\xi)_{\xi \in \Delta} \in \mathcal{D}^\Delta.$$

Similarly, for a function $f : A \rightarrow B$ in \mathcal{D}^Ξ , we denote the induced function $A|_\Delta \rightarrow B|_\Delta$ by $f|_\Delta$.

(c) The Δ -reduct of a functor $\mathbb{M} : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$ is defined by

$$\mathbb{M}|_\Delta A := (\mathbb{M}(A \upharpoonright \Xi))|_\Delta \quad \text{and} \quad \mathbb{M}|_\Delta f := (\mathbb{M}(f \upharpoonright \Xi))|_\Delta.$$

(d) The Δ -reduct of an \mathbb{M} -algebra $\mathfrak{A} = \langle A, \pi \rangle$ is the $\mathbb{M}|_\Delta$ -algebra $\mathfrak{A}|_\Delta$ with universe $A|_\Delta$ and product $\pi|_\Delta \upharpoonright \mathbb{M}|_\Delta(A|_\Delta)$. For a class \mathcal{C} of \mathbb{M} -algebras we set $\mathcal{C}|_\Delta := \{ \mathfrak{A}|_\Delta \mid \mathfrak{A} \in \mathcal{C} \}$. \square

Let us check that these notions are well-behaved.

Lemma 3.4. *Let $\Delta \subseteq \Xi$ be sets of sorts and \mathbb{M} a monad on \mathcal{D}^Ξ .*

- (a) The reduct functor $|_{\Delta} : \mathcal{D}^{\Xi} \rightarrow \mathcal{D}^{\Delta}$ has the left adjoint $-\uparrow \Xi : \mathcal{D}^{\Delta} \rightarrow \mathcal{D}^{\Xi}$.
- (b) $(\mathbb{M}|_{\Delta}, (\text{flat}|_{\Delta} \uparrow \mathbb{M}|_{\Delta} \mathbb{M}|_{\Delta} A|_{\Delta}), \text{sing}|_{\Delta})$ is a monad on \mathcal{D}^{Δ} .
- (c) The reduct functor induces a functor $|_{\Delta} : \text{Alg}(\mathbb{M}) \Rightarrow \text{Alg}(\mathbb{M}|_{\Delta})$.
- (d) If $\text{Alg}(\mathbb{M})$ has coequalisers, the functor $|_{\Delta}$ from (c) has a left adjoint $-\uparrow : \text{Alg}(\mathbb{M}|_{\Delta}) \Rightarrow \text{Alg}(\mathbb{M})$.
- (e) If $\text{Alg}(\mathbb{M})$ has coequalisers and \mathbb{M} preserves E -morphisms, the functor $-\uparrow$ from (d) preserves E -morphisms.

Proof. (a) It is straightforward to check that $|_{\Delta}$ and $-\uparrow \Xi$ are functors: we have to show that

- ♦ $A|_{\Delta} \in \mathcal{D}^{\Delta}$, for $A \in \mathcal{D}^{\Xi}$,
- ♦ $\varphi|_{\Delta} : A|_{\Delta} \rightarrow B|_{\Delta}$, for $\varphi : A \rightarrow B$, and
- ♦ $(\varphi \circ \psi)|_{\Delta} = \varphi|_{\Delta} \circ \psi|_{\Delta}$,
- ♦ the corresponding statements for $-\uparrow \Xi$.

All of these statements follow immediately from the definition.

It remains to prove that $-\uparrow \Xi$ and $|_{\Delta}$ form an adjunction. We define a bijection

$$\mathcal{D}^{\Xi}(A \uparrow \Xi, B) \cong \mathcal{D}^{\Delta}(A, B|_{\Delta})$$

by mapping $f : A \uparrow \Xi \rightarrow B$ to $f|_{\Delta} : A \rightarrow B|_{\Delta}$. This function is indeed bijective since its inverse is given by $f \mapsto f \uparrow \Xi$. To see that the bijection is natural in A and B note that, for $f : A \uparrow \Xi \rightarrow B$, $g : A \rightarrow A'$, and $h : B \rightarrow B'$, we have

$$(f \circ g \uparrow \Xi)|_{\Delta} = f|_{\Delta} \circ g \quad \text{and} \quad (h \circ f)|_{\Delta} = g|_{\Delta} \circ f|_{\Delta}.$$

Before proving the other statements, let us quickly derive the unit and the counit of the adjunction. Since $(A \uparrow \Xi)|_{\Delta} = A$, the identity morphisms induce a natural isomorphism $e := \text{id} : \text{Id} \Rightarrow (-\uparrow \Xi)|_{\Delta}$ which forms the unit. For the counit, we define morphisms $i : A|_{\Delta} \uparrow \Xi \rightarrow A$ by choosing $i_{\xi} := \text{id}_{A_{\xi}}$, for $\xi \in \Delta$, while, for $\xi \in \Xi \setminus \Delta$, $i_{\xi} : 0 \rightarrow A_{\xi}$ is the unique

morphism from the initial object. Then i forms a natural transformation $i : -|_{\Delta} \uparrow \Xi \Rightarrow \text{Id}$, which is the counit of the adjunction.

(b), (c), (d) follow immediately by Proposition I.6.1.

(e) Let $\varepsilon : A \rightarrow B$ be an E -morphism. By Lemma I.7, the coequalisers $\rho_A : \mathbb{M}A \rightarrow \mathfrak{A}^{\uparrow}$ and $\rho_B : \mathbb{M}B \rightarrow \mathfrak{B}^{\uparrow}$ belong to E . By assumption, we also have $\mathbb{M}\varepsilon \in E$. Hence, $\rho_B \circ \mathbb{M}\varepsilon \in E$ and $\rho_A \in E$. Since $f^{\uparrow} \circ \rho_A = \rho_B \circ \mathbb{M}\varepsilon$, it therefore follows by Lemma I.3 that $f^{\uparrow} \in E$. \square

Open Question. Under which conditions is the unit $v : \mathfrak{A} \rightarrow \mathfrak{A}^{\uparrow}|_{\Delta}$ of the adjunction in (d) an isomorphism?

4. Congruences

Finally, we turn to quotients and congruences. For Σ -algebras and partial orders, there is a one-to-one correspondence between quotients and congruences. Thus, a congruence is simply a concise encoding of a quotient. As there is no general definition of a congruence, we will use quotients throughout but adapting the terminology of congruences. One of the most important general facts about congruences is that they form a complete partial order.

Definition 4.1. Let \mathcal{C} be a category with a factorisation system (E, M) .

(a) Let CPO_{κ} be the category of all partial orders where every set of size less than κ has an infimum. The morphisms of CPO_{κ} are all monotone functions preserving such infima.

(b) Given an object $A \in \mathcal{C}$, we define a preorder \leq on the class of all E -morphisms with domain A by setting

$$\varepsilon \leq \varepsilon' \quad \text{iff} \quad \varepsilon' = \rho \circ \varepsilon, \quad \text{for some } \rho.$$

The equivalence classes associated with \leq are called *congruences* of A . Usually, we identify a congruence with its representatives. Given a congruence $\theta : A \rightarrow B$ on A , we denote the *quotient* by $A/\theta := B$.

(c) We define a contravariant functor $\text{Cong} : \mathcal{C} \rightarrow \text{Pos}$ as follows. For an object $A \in \mathcal{C}$, $\text{Cong}(A)$ is the partial order of congruences of A . For a morphism $\varphi : A \rightarrow B$ and a congruence $\theta : B \rightarrow C$, we define

$$\text{Cong}(\varphi)(\theta) := \ker(\theta \circ \varphi).$$

Since we will frequently be considering congruences of an \mathbb{M} -algebra \mathfrak{A} and of the underlying universe A at the same time, we will use the notation $\text{Cong}_{\mathbb{M}}(\mathfrak{A})$ for the former and $\text{Cong}(A)$ for the latter.

(d) A category \mathcal{C} has *congruences* if the quotients of every object form a partial order in CPO_{κ} .

Let us quickly check that the functor Cong is well-defined.

Lemma 4.2. *Let $\varphi : A \rightarrow B$, $q : B \rightarrow C$, and $q' : B \rightarrow C'$ be morphisms. Then*

$$\ker q \leq \ker q' \quad \text{implies} \quad \ker(q \circ \varphi) \leq \ker(q' \circ \varphi).$$

Proof. Without loss of generality, we may assume that q and q' belong to E . Suppose that $q \leq q'$ and fix a morphism ρ such that $q' = \rho \circ q$. Let $\mu \circ \varepsilon$, $\mu' \circ \varepsilon'$, and $\hat{\mu} \circ \hat{\varepsilon}$ be the *EM*-factorisations of, respectively, $q \circ \varphi$, $q' \circ \varphi$, and $\rho \circ \mu$. Then $\hat{\mu} \circ (\hat{\varepsilon} \circ \varepsilon)$ is a second *EM*-factorisation of $q' \circ \varphi$. By uniqueness of factorisations, we obtain an isomorphism σ such that

$$\sigma \circ (\hat{\varepsilon} \circ \varepsilon) = \varepsilon' \quad \text{and} \quad \mu' \circ \sigma = \hat{\mu}.$$

Consequently, $\ker(q \circ \varphi) = \varepsilon \leq \varepsilon' = \ker(q' \circ \varphi)$. \square

Computing infima and suprema of congruences is fortunately straightforward.

Lemma 4.3. *Suppose that \mathcal{C} has products of size κ and let $\theta_i : A \rightarrow B_i$, for $i < \kappa$, be a family of congruences of A . Then*

$$\inf_{i < \kappa} \theta_i = \ker \langle \theta_i \rangle_{i < \kappa}.$$

Proof. Let $p_k : \prod_i B_i \rightarrow B_k$ be the projection. Then

$$p_k \circ \langle \theta_i \rangle_i = \theta_k \quad \text{implies} \quad \ker \langle \theta_i \rangle_i \leq \theta_k.$$

Conversely, suppose that $q : A \rightarrow C$ is some epimorphism with

$$q \leq \theta_i, \quad \text{for all } i < \kappa.$$

Then there exist morphisms $f_i : C \rightarrow B_i$ such that $f_i \circ q = \theta_i$. Hence, $(f_i)_i$ forms a cone from C to $(B_i)_i$. Since $(p_i)_i$ is limiting, there exists a unique morphism $g : C \rightarrow \prod_i B_i$ with $p_i \circ g = f_i$, for all i . It follows that

$$p_i \circ g \circ q = f_i \circ q = \theta_i.$$

Hence, $g \circ q = \langle \theta_i \rangle_i$, which implies that $q \leq \langle \theta_i \rangle_i$, as desired. \square

Exercise 4.1. Show that $\sup_{i < \kappa} \theta_i = \ker \bigoplus_i \theta_i$, if the coproduct is defined.

Lemma 4.4. *Let \mathcal{C} be a category with products of size less than κ . Then $\text{Cong} : \mathcal{C} \rightarrow \text{CPO}_{\kappa}$ is a contravariant functor.*

Proof. Let $f : A \rightarrow B$ be a morphism and let H be a set of congruences of B of size $|H| < \kappa$. Then

$$\begin{aligned} \inf \{ \text{Cong}(f)(\eta) \mid \eta \in H \} &= \langle \text{Cong}(f)(\eta) \rangle_{\eta \in H} \\ &= \ker \langle \eta \circ f \rangle_{\eta \in H} \\ &= \ker \langle \langle \eta \rangle_{\eta \in H} \circ f \rangle \\ &= \ker((\inf H) \circ f) \\ &= \text{Cong}(f)(\inf H). \end{aligned} \quad \square$$

Congruences in Set

The rather abstract definition of the congruence lattice of some object A above is not very useful when one is trying to understand $\text{Cong}(A)$ since it

requires considering all possible morphisms from A to some other object of the category. A definition that only depends on A itself would be much more convenient. Of course, such a definition requires knowledge of the internal structure of A , which is not accessible in an abstract category. Consequently, we will have to work in concrete categories. We start with the category **Set**.

Definition 4.5. Let A be a set.

(a) For an equivalence relation \sim on A , we denote by $[a]_{\sim}$ the \sim -class of $a \in A$.

(b) The *kernel* of a function $f : A \rightarrow B$ is the relation

$$\ker f := \{ \langle a, b \rangle \in A \times A \mid f(a) = f(b) \}.$$

(We use the same notation as above. It should always be obvious from the context which version we are referring to.)

Proposition 4.6. Let $A \in \mathbf{Set}$ be a set. $\text{Cong}(A)$ is isomorphic to the lattice of all equivalence relations on A ordered by inclusion. This isomorphism maps a surjective function $\varepsilon : A \rightarrow C$ to its kernel $\ker \varepsilon$.

Proof. Clearly, each kernel $\ker \varepsilon$ is an equivalence relation. Consequently, the map $\sigma : \varepsilon \mapsto \ker \varepsilon$ is well-defined. To show that it is monotone, suppose that $\varepsilon \leq \varepsilon'$. Then $\varepsilon' = \rho \circ \varepsilon$ for some function ρ . Consequently,

$$\begin{aligned} \langle a, b \rangle \in \ker \varepsilon &\Rightarrow \varepsilon(a) = \varepsilon(b) \\ &\Rightarrow \varepsilon'(a) = \varepsilon'(b) \Rightarrow \langle a, b \rangle \in \ker \varepsilon'. \end{aligned}$$

Conversely, If $\ker \varepsilon \subseteq \ker \varepsilon'$, then

$$\varepsilon(a) = \varepsilon(b) \text{ implies } \varepsilon'(a) = \varepsilon'(b).$$

As ε is surjective, this implies that there exists some function ρ with $\varepsilon' = \rho \circ \varepsilon$. We have shown that

$$\varepsilon \leq \varepsilon' \quad \text{iff} \quad \sigma(\varepsilon) \leq \sigma(\varepsilon').$$

Hence, it remains to show that σ is bijective. Injectivity follows immediately from the fact that

$$\varepsilon \leq \varepsilon' \quad \text{iff} \quad \sigma(\varepsilon) \leq \sigma(\varepsilon').$$

For surjectivity, let \sim be an equivalence relation on A . Then $\sim = \ker p_{\sim}$, where $p_{\sim} : A \rightarrow A/\sim$ is the projection mapping each $a \in A$ to its \sim -class $[a]_{\sim}$. \square

Congruences in Pos

Next, let us take a look at **Pos**. As seen above there are two natural factorisation systems for this categories. We start with the one consisting of quotients and injective functions. In that case, we can again use equivalence relations for congruences. Given a partial order $\langle A, \leq \rangle$ and an equivalence relation \sim on A , we can define the quotient as the set A/\sim equipped with the ordering

$$[a]_{\sim} \leq [b]_{\sim} \quad : \text{iff} \quad a' \leq b' \quad \text{for some } a' \in [a]_{\sim} \text{ and } b' \in [b]_{\sim}.$$

If we define the kernel of a function $\varphi : A \rightarrow B$ in the same way as for sets, we again obtain an isomorphism between kernels and quotients.

For the standard factorisation system, which consists of surjective functions and embeddings, a different approach is needed based on preorders instead of equivalence relations.

Definition 4.7. Let $\langle A, \leq \rangle$ be a partial order.

(a) A *congruence ordering* on A is a preorder $\sqsubseteq \subseteq A \times A$ with $\leq \sqsubseteq \sqsubseteq$.

(b) The *kernel* of a function $f : A \rightarrow B$ in **Pos** is the relation

$$\ker f := \{ \langle a, a' \rangle \in A \times A \mid f(a) \leq f(a') \}.$$

(c) The set of \sqsubseteq -classes is

$$A/\sqsubseteq := \{ [a]_{\sqsubseteq} \mid a \in A \} \quad \text{where} \quad [a]_{\sqsubseteq} := \{ b \in A \mid b \sqsubseteq a \text{ and } a \sqsubseteq b \}.$$

We equip it with the ordering

$$[a]_{\sqsubseteq} \leq [b]_{\sqsubseteq} \quad : \text{iff} \quad a \sqsubseteq b.$$

(d) The *quotient map* $q : A \rightarrow A/\sqsubseteq$ maps $a \in A$ to $[a]_{\sqsubseteq}$.

As above, we obtain the following correspondence. The proof is similar to that for Set.

Proposition 4.8. *Let $A \in \text{Pos}$ be a partial order. $\text{Cong}(A)$ is isomorphic to the lattice of all congruence orderings on A ordered by inclusion. This isomorphism maps a surjective monotone function $\varepsilon : A \rightarrow C$ to its kernel $\ker \varepsilon$.*

In this terminology, the diagonal fill-in property takes the following form.

Lemma 4.9 (Factorisation Lemma). *Let $f : A \rightarrow B$ and $g : A \rightarrow C$ be functions in Pos and assume that f is surjective. Then*

$$\ker f \subseteq \ker g \quad \text{iff} \quad g = h \circ f, \quad \text{for some } h : B \rightarrow C.$$

Moreover, the function h is unique, if it exists.

Proof. The uniqueness of h follows from the surjectivity of f , since surjective functions are epimorphisms: $h \circ f = g = h' \circ f$ implies $h = h'$. Hence, it remains to consider existence.

(\Rightarrow) If $g = h \circ f$, then

$$f(a) \leq f(b) \quad \text{implies} \quad g(a) = h(f(a)) \leq h(f(b)) = g(b).$$

(\Leftarrow) Suppose that $\ker f \subseteq \ker g$. As f is surjective, it has a right inverse r (in Set, r might not be monotone). We claim that $h := g \circ r$ is the desired function.

For monotonicity, suppose that $a \leq b$ in B . Then

$$f(r(a)) = a \leq b = f(r(b)) \quad \text{implies} \quad \langle r(a), r(b) \rangle \in \ker f \subseteq \ker g.$$

Consequently,

$$h(a) = g(r(a)) \leq g(r(b)) = h(b).$$

To show that $g = h \circ f$, set $e := r \circ f$. For $a \in A$, it follows that

$$f(e(a)) = (f \circ r \circ f)(a) = f(a).$$

Hence, $\langle a, e(a) \rangle, \langle e(a), a \rangle \in \ker f \subseteq \ker g$, which implies that $g(a) = g(e(a))$. Thus $g = g \circ e = g \circ r \circ f = h \circ f$. \square

Congruences of Algebras

We have defined congruences for objects in \mathcal{D} . But what we are really interested in are congruences of \mathbb{M} -algebras. Before presenting the general definition let us first take a look at the case of Σ -algebras.

Example. Let Σ be a signature and \mathcal{A} a Σ -algebra. A congruence of \mathcal{A} is an equivalence relation $\sim \subseteq A \times A$ such that, for every function $f \in \Sigma$,

$$a_i \sim b_i, \text{ for all } i, \quad \text{implies} \quad f(\bar{a}) \sim f(\bar{b}).$$

This is equivalent to \sim inducing a subalgebra of $\mathcal{A} \times \mathcal{A}$.

The definition of congruences for \mathbb{M} -algebras is similar. We have already seen in Lemma 1.11 how to lift a factorisation system from \mathcal{D} to $\text{Alg}(\mathbb{M})$. This immediately gives us the notion of a congruence for \mathbb{M} -algebras.

Definition 4.10. An \mathbb{M} -congruence is a congruence in the category $\text{Alg}(\mathbb{M})$.

We start with a criterion for an congruence $\varepsilon \in \mathcal{D}^\Xi$ to be an \mathbb{M} -congruence.

Lemma 4.11. *Let $\langle E, M \rangle$ be a factorisation system where E and M are closed under the monad \mathbb{M} . Let \mathcal{A} be an \mathbb{M} -algebra and $f : A \rightarrow B$ a morphism of \mathcal{D}^Ξ . Then $\ker f$ is an \mathbb{M} -algebra morphism if, and only if,*

$$\ker \mathbb{M}f \leq \ker(f \circ \pi).$$

Proof. (\Rightarrow) Let $f = \mu \circ \varepsilon$ be the EM -factorisation of f and suppose that $\varepsilon : \mathcal{A} \rightarrow \mathcal{C}$ is an \mathbb{M} -algebra morphism. Then

$$\pi \circ \mathbb{M}\varepsilon = \varepsilon \circ \pi \quad \text{implies} \quad \mathbb{M}\varepsilon \leq \varepsilon \circ \pi.$$

Note that $\mathbb{M}f = \mathbb{M}\mu \circ \mathbb{M}\varepsilon$ is the EM -factorisation of $\mathbb{M}f$. Furthermore, by Lemma 1.7, $\pi \circ \text{sing} = \text{id}$ implies that $\pi \in E$. Hence, $f \circ \pi = \mu \circ (\varepsilon \circ \pi)$ is the EM -factorisation of $f \circ \pi$. Consequently, we have

$$\ker \mathbb{M}f = \mathbb{M}\varepsilon \leq \varepsilon \circ \pi = \ker(f \circ \pi).$$

(\Leftarrow) Let $\varepsilon := \ker f : A \rightarrow C$. Since $\mathbb{M}\varepsilon = \ker \mathbb{M}f \leq \ker(f \circ \pi) = \varepsilon \circ \pi$, there exists a morphism $\rho : \mathbb{M}C \rightarrow A$ with $\rho \circ \mathbb{M}\varepsilon = \varepsilon \circ \pi$. Note that

$$\begin{aligned} \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}\varepsilon &= \rho \circ \mathbb{M}(\varepsilon \circ \pi) \\ &= \varepsilon \circ \pi \circ \mathbb{M}\pi \\ &= \varepsilon \circ \pi \circ \text{flat} \\ &= \rho \circ \mathbb{M}\varepsilon \circ \text{flat} = \rho \circ \text{flat} \circ \mathbb{M}\mathbb{M}\varepsilon. \end{aligned}$$

As $\mathbb{M}\mathbb{M}\varepsilon$ is an epimorphism, it therefore follows that $\rho \circ \mathbb{M}\rho = \rho \circ \text{flat}$. Consequently, $\mathfrak{C} := \langle C, \rho \rangle$ is an \mathbb{M} -algebra and $\varepsilon = \ker f : \mathfrak{A} \rightarrow \mathfrak{C}$ a morphism of \mathbb{M} -algebras. \square

Lemma 4.12. $\text{Cong}_{\mathbb{M}}(\mathfrak{A})$ is closed (as a subset of $\text{Cong}(A)$) under all infima that exists in $\text{Cong}(A)$.

Proof. Let $\Theta \subseteq \text{Cong}_{\mathbb{M}}(\mathfrak{A})$ be a set with infimum $\theta := \inf \Theta$ (in $\text{Cong}(A)$). By Lemma 4.3, we have

$$\theta = \langle \varepsilon \rangle_{\varepsilon \in \Theta} : \mathfrak{A} \rightarrow \prod_{\varepsilon \in \Theta} \mathfrak{A}/\varepsilon.$$

It remains to show that θ is a morphism of \mathbb{M} -algebras. Let π_ε be the product of \mathfrak{A}/ε , π_Π the one of $\prod_\varepsilon \mathfrak{A}/\varepsilon$, and let $p_\eta : \prod_\varepsilon A/\varepsilon \rightarrow A/\eta$ and $q_\eta : \prod_\varepsilon \mathbb{M}A/\varepsilon \rightarrow \mathbb{M}A/\eta$ be the projections. For every $\eta \in \Theta$, we have

$$\begin{aligned} q_\eta \circ \langle \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon &= \mathbb{M}p_\eta \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon \\ &= \mathbb{M}(p_\eta \circ \langle \varepsilon \rangle_\varepsilon) = \mathbb{M}q_\eta = q_\eta \circ \langle \mathbb{M}\varepsilon \rangle_\varepsilon. \end{aligned}$$

Consequently, $\langle \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon = \langle \mathbb{M}\varepsilon \rangle_\varepsilon$, which implies that

$$\begin{aligned} \theta \circ \pi &= \langle \varepsilon \rangle_\varepsilon \circ \pi \\ &= \langle \varepsilon \circ \pi \rangle_\varepsilon \\ &= \langle \pi_\varepsilon \circ \mathbb{M}\varepsilon \rangle_\varepsilon \\ &= \prod_\varepsilon \pi_\varepsilon \circ \langle \mathbb{M}\varepsilon \rangle_\varepsilon \end{aligned}$$

$$\begin{aligned} &= \prod_\varepsilon \pi_\varepsilon \circ \langle \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon \\ &= \langle \pi_\varepsilon \circ \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon \\ &= \langle p_\varepsilon \circ \pi_\Pi \rangle_\varepsilon \circ \mathbb{M}\theta \\ &= \pi_\Pi \circ \mathbb{M}\theta. \end{aligned}$$

\square


5. Varieties

The topic of this section concerns classes of objects that are closed under certain algebraic operations like subobjects, quotients, products, and the like. The aim is to find a simple description of such a class. But before doing so, we introduce the following technical property.

Projective Objects

One important property of sets of the form $\mathbb{J}X$ is that they satisfy the following property.

Definition 5.1. Let \mathcal{C} be a category. An object $P \in \mathcal{C}$ is *projective* with respect to a morphism $\varepsilon : A \rightarrow B$ if, for every $\varphi : P \rightarrow B$, there exists a morphism $\hat{\varphi} : P \rightarrow A$ such that $\varphi = \varepsilon \circ \hat{\varphi}$.



Let us check that the notion of projectivity in fact generalises freeness. In order to define what a free object is, we need additional structure on the category. Therefore, in the following statement we will work in a concrete category (i.e., one equipped with a faithful functor to Set).

Lemma 5.2. Let $\mathbb{J} \dashv \mathbb{V}$ be an adjunction between Set and \mathcal{C} . An object of the form $\mathbb{J}X$ is projective with respect to a morphism $\varepsilon : A \rightarrow B$ if, and only if, $X = \emptyset$ or $\mathbb{V}\varepsilon$ is surjective.

Proof. (\Leftarrow) If $X = \emptyset$, then X is initial in \mathbf{Set} . Since left adjoints preserve colimits, it follows that $\mathbb{J}X$ is initial in \mathcal{C} . Consequently, if $f : \mathbb{J}X \rightarrow B$ is the unique morphism to B , then $\varepsilon \circ g = f$, where $g : \mathbb{J}X \rightarrow A$ is the unique morphism to A .

Now suppose that $\mathbb{V}\varepsilon$ is surjective. Then it has a right inverse $r : \mathbb{V}B \rightarrow \mathbb{V}A$. To show projectivity of $\mathbb{J}X$, consider a morphism $f : \mathbb{J}X \rightarrow B$. Let $f_o : X \rightarrow \mathbb{V}B$ be its image under the adjunction and let $g : \mathbb{J}X \rightarrow A$ be the image of $g_o := r \circ f_o$. As the image of $\varepsilon \circ g$ under the adjunction is

$$\mathbb{V}\varepsilon \circ g_o = \mathbb{V}\varepsilon \circ r \circ f_o = f_o,$$

it follows that $\varepsilon \circ g = f$.

(\Rightarrow) Suppose that $X \neq \emptyset$ and that $\mathbb{V}\varepsilon$ is not surjective. Then there exists a function $f_o : X \rightarrow \mathbb{V}B$ with $\text{rng } f_o \cap \text{rng } \mathbb{V}\varepsilon = \emptyset$. Let $f : \mathbb{J}X \rightarrow B$ be the image of f_o under the adjunction. We claim that there is no morphism $g : \mathbb{J}X \rightarrow A$ with $\varepsilon \circ g = f$. For a contradiction, suppose otherwise. Let $g_o : X \rightarrow \mathbb{V}A$ be the image of g under the adjunction. Then $\mathbb{V}\varepsilon \circ g_o = f_o$, which contradicts the fact that the ranges of these two functions are non-empty and disjoint. \square

This situation is typical: very few objects are projective with respect to a non-surjective morphism ε .

Exercise 5.1. An object A in a category \mathcal{C} is called a *generator* if, for every object $B \in \mathcal{C}$, the set of all morphisms $A \rightarrow B$ is jointly epimorphic. (For the dual concept, see Definition III.1.9.)

- (a) Show that, in the categories \mathbf{Set} , \mathbf{Pos} , or \mathbf{Top} , every non-empty object is a generator. (*Hint.* It is sufficient to prove that the terminal object is a generator.)
- (b) Show that, if a generator P is projective with respect to a morphism $\varepsilon : A \rightarrow B$, then ε is an epimorphism. \lrcorner

More generally, we can lift projectiveness along every adjunction, but only in one direction.

Lemma 5.3. Let $\mathbb{J} \dashv \mathbb{V}$ be an adjunction where \mathbb{V} preserves epimorphisms and let $\varepsilon : A \rightarrow B$ be an epimorphism. If X is projective with respect to $\mathbb{V}\varepsilon$, then $\mathbb{J}X$ is projective with respect to ε .

Proof. Suppose that $\varepsilon : A \rightarrow B$ is an epimorphism and let $\psi : \mathbb{J}X \rightarrow B$. Then ψ corresponds to a morphism $\psi' : X \rightarrow \mathbb{V}B$. As X is projective with respect to $\mathbb{V}\varepsilon$, we can find a morphism $\rho' : X \rightarrow \mathbb{V}A$ with $\mathbb{V}\varepsilon \circ \rho' = \psi'$. This morphism corresponds to a morphism $\rho : \mathbb{J}X \rightarrow A$. We claim that $\varepsilon \circ \rho = \psi$.

Note that $\psi \in \mathcal{D}(\mathbb{J}X, B)$ corresponds to $\psi' \in \mathbf{Set}(X, \mathbb{V}B)$, and $\rho \in \mathcal{D}(\mathbb{J}X, A)$ corresponds to $\rho' \in \mathbf{Set}(X, \mathbb{V}A)$. As this correspondence is natural in A , it follows that $\varepsilon \circ \rho$ corresponds to $\mathbb{V}\varepsilon \circ \rho' = \psi'$. Hence, ψ and $\varepsilon \circ \rho$ both correspond to ψ' . As the correspondence is bijective, it follows that $\psi = \varepsilon \circ \rho$. \square

Corollary 5.4. Let \mathcal{D} be a discrete category.

- (a) Every object of the form $\mathbb{J}X$ is projective with respect to all surjective morphisms.
- (b) Every free \mathbb{M} -algebra of the form $\mathbb{M}\mathbb{J}X$ is projective with respect to all surjective morphisms of \mathbb{M} -algebras.

Proof. (a) In \mathbf{Set} every set $X \in \mathbf{Set}$ is projective with respect to all surjective functions.

(b) follows by applying the lemma to (a) and the adjunction between $\mathcal{D}^{\mathbb{M}}$ and $\mathbf{Alg}(\mathbb{M})$ that is canonically associated with the monad \mathbb{M} . \square

Definition 5.5. For a class \mathcal{P} of objects, we set

$$E_{\mathcal{P}} := \{ \varepsilon \in E \mid \text{every } P \in \mathcal{P} \text{ is projective w.r.t. } \varepsilon \}.$$

Many of the properties in Lemma 1.3 carry over to the set $E_{\mathcal{P}}$.

Lemma 5.6. Let \mathcal{A} be a category with factorisation system $\langle E, M \rangle$ and let $\mathcal{P} \subseteq \mathcal{A}$ be a subcategory.

- (a) $E_{\mathcal{P}}$ contains all isomorphisms.

- (b) $E_{\mathcal{P}}$ is closed under composition.
- (c) $p \in E$ and $q \circ p \in E_{\mathcal{P}}$ implies $q \in E_{\mathcal{P}}$.
- (d) If E is closed under products of a certain size, so is $E_{\mathcal{P}}$.

Proof. (a) By (FS1), E contains all isomorphisms. Since clearly every object is projective with respect to an isomorphism, the claim follows.

(b) Let $p, q \in E_{\mathcal{P}}$ be morphisms such that $q \circ p$ is defined. As E is closed under composition, we have $q \circ p \in E$. Consequently, the claim follows from the straightforward fact that, if an object P is projective with respect to p and q , it is also projective with respect to $p \circ q$.

(c) Suppose that $p : A \rightarrow B$ and $q : B \rightarrow C$. By Lemma 1.3, we have $q \in E$. Hence, it remains to show that every $P \in \mathcal{P}$ is projective with respect to q . Let $P \in \mathcal{P}$ and $\varphi : P \rightarrow C$. Since $q \circ p \in E_{\mathcal{P}}$, we can find a morphism $\psi : P \rightarrow A$ such that $q \circ p \circ \psi = \varphi$. Consequently, the morphism $\hat{\varphi} := p \circ \psi$ satisfies $q \circ \hat{\varphi} = \varphi$.

(d) Suppose that $\varepsilon_i : A_i \rightarrow B_i$, for $i < \alpha$, are $E_{\mathcal{P}}$ -morphisms with $\prod_i \varepsilon_i \in E$. We have to show that every $P \in \mathcal{P}$ is projective with respect to $\prod_i \varepsilon_i$.

Fix $P \in \mathcal{P}$ and a morphism $\varphi : P \rightarrow \prod_i B_i$. Let $q_i : \prod_k B_k \rightarrow B_i$ be the projection. Then $q_i \circ \varphi : P \rightarrow B_i$ and $\varepsilon_i \in E_{\mathcal{P}}$ implies that there is some morphism $\psi_i : P \rightarrow A_i$ with $\varepsilon_i \circ \psi_i = q_i \circ \varphi$. It follows that

$$\prod_i \varepsilon_i \circ \langle \psi_i \rangle_i = \langle q_i \circ \varphi \rangle_i = \varphi. \quad \square$$

Example. Let \mathcal{D} be a discrete category and $\mathcal{P} := \{ \mathbb{J}X \mid X \in \text{Set} \}$. We have shown in Corollary 5.4 that $E_{\mathcal{P}}$ contains all surjective E -morphisms. \square

Varieties

Let us start simply by considering classes closed under subobjects and quotients.

Definition 5.7. Let \mathcal{C} be a category, \mathcal{P}, \mathcal{K} classes of objects, and $E_* \subseteq E_{\mathcal{P}}$ a set containing all isomorphisms.

(a) An object $A \in \mathcal{C}$ is \mathcal{P} -generated if there exists an E -morphism $P \rightarrow A$, for some $P \in \mathcal{P}$.

(b) We denote the closure of \mathcal{K} under \mathcal{P} -generated M -subobjects by

$$S(\mathcal{K}) := \{ A \mid e : A \rightarrow K, e \in M, K \in \mathcal{K}, A \text{ is } \mathcal{P}\text{-generated} \}.$$

(c) We denote the closure of \mathcal{K} under all E_* -quotients by

$$H(\mathcal{K}) := \{ A \mid q : K \rightarrow A, q \in E_*, K \in \mathcal{K} \}. \quad \square$$

Example. For $\mathcal{C} = \text{Alg}(\mathbb{M})$ with the standard factorisation system and for $\mathcal{P} = \{ \mathbb{M}\mathbb{J}X \mid X \text{ finite} \}$, an algebra \mathfrak{A} is \mathcal{P} -generated if, and only if, it is finitely generated. \square

To compute the closure of a class \mathcal{K} under subobjects and quotients, we can take alternating sequences of these two operations. It turns out that this is not really necessary. Sequences of length two are sufficient.

Lemma 5.8. If $E_* = E$, then

$$S \circ H \subseteq H \circ S.$$

Proof. Suppose that $A \in SH(\mathcal{K})$. Then there exist some E_* -morphism $q : C \rightarrow B$, an M -morphism $e : A \rightarrow B$, and an E -morphism $\rho : P \rightarrow A$, for some $C \in \mathcal{K}$ and $P \in \mathcal{P}$. Since P is projective with respect to q , we can find a morphism $f : P \rightarrow C$ such that $q \circ f = e \circ \rho$.

Let D be the pullback of e and q , and let e' and q' be the other morphisms of the pullback square. As D is a limit, there exists a morphism $g : P \rightarrow D$ such that

$$q' \circ g = \rho \quad \text{and} \quad e' \circ g = f.$$

By Lemma 1.9, we have $e' \in M$. Let $g = i \circ h$ be the EM -factorisation of g and let K be the codomain of h . Then $e' \circ i : K \rightarrow C$ belongs to M and $h : P \rightarrow K$ to E . Furthermore, by Lemma 1.3, $\rho, h \in E$ implies that $q' \circ i : K \rightarrow A$ also belongs to E . Consequently, $K \in S(\mathcal{K})$ and $A \in HS(\mathcal{K})$. \square

Proposition 5.9. *Let \mathcal{C} be a category, \mathcal{P}, \mathcal{K} two classes of objects, and $E_* = E$. Then $\text{HS}(\mathcal{K})$ is the closure of class \mathcal{K} under E -quotients and \mathcal{P} -generated M -subobjects.*

Proof. $\text{HS}(\mathcal{K})$ is closed under both operations since $\text{HHS}(\mathcal{K}) = \text{HS}(\mathcal{K})$ and $\text{SHS}(\mathcal{K}) \subseteq \text{HSS}(\mathcal{K}) = \text{HS}(\mathcal{K})$. \square

Next let us add products to our closure operations. We are both interested in finite products and in arbitrary ones.

Definition 5.10. Let \mathcal{C} be a category, \mathcal{P}, \mathcal{K} classes of objects, $E_* \subseteq E_{\mathcal{P}}$, and κ a cardinal or $\kappa = \infty$.

(a) We denote the closure of \mathcal{K} under products of size less than κ by

$$\text{P}_{\kappa}(\mathcal{K}) := \{ \prod_{i \in I} A_i \mid A_i \in \mathcal{K}, |I| < \kappa \}.$$

(b) A *variety* is a class \mathcal{K} that is closed under (i) E_* -quotients and (ii) \mathcal{P} -generated M -subobjects of arbitrary products.

(c) A *pseudo-variety* is a class \mathcal{K} that is closed under (i) E_* -quotients and (ii) \mathcal{P} -generated M -subobjects of finite products. \lrcorner

Examples. (a) The classes of (i) all groups, (ii) all modules over a fixed semiring, (iii) all aperiodic monoids, and (iv) all distributive lattices form varieties.

(b) The subclasses of the classes in (a) consisting of all (i) finitely generated algebras or (ii) all finite algebras form pseudo-varieties. \lrcorner

The reason we have combined the operations S and P_{κ} in the above definition is the fact that a product of \mathcal{P} -generated objects does not need to be \mathcal{P} -generated.

Lemma 5.11. *If E_* is closed under products of size less than κ , then*

$$\text{P}_{\kappa}H \subseteq H\text{P}_{\kappa} \quad \text{and} \quad \text{SP}_{\kappa}S \subseteq \text{SP}_{\kappa}.$$

Proof. For the first inclusion, let $A := \prod_{i \in I} B_i$ be a product with $B_i \in H(\mathcal{K})$ and fix E_* -morphisms $\varepsilon_i : B_i \rightarrow C_i$ with $C_i \in \mathcal{K}$. By Lemma 5.6, the product

$\prod_i \varepsilon_i : \prod_i B_i \rightarrow \prod_i C_i$ belongs to E_* . Consequently, A is an E_* -quotient of $\prod_i C_i \in \text{P}_{\kappa}(\mathcal{K})$.

For the second inclusion, let A be a \mathcal{P} -generated M -subobject of a product $B := \prod_{i \in I} C_i$ with $C_i \in S(\mathcal{K})$. Fix an M -morphism $\mu : A \rightarrow B$ and M -morphisms $e_i : C_i \rightarrow D_i$ with $D_i \in \mathcal{K}$. By Lemma 1.9, the product $\prod_i e_i : \prod_i C_i \rightarrow \prod_i D_i$ belongs to M . Hence, so does $\prod_i e_i \circ \mu$. Consequently, A is a \mathcal{P} -generated M -subobject of $\prod_i D_i \in \text{P}_{\kappa}(\mathcal{K})$. \square

Theorem 5.12. *Suppose that $E_* = E$ and E is closed under products (arbitrary ones for (a) and finite ones for (b)).*

(a) *A class \mathcal{V} is a variety if, and only if, $\mathcal{V} = \text{HSP}_{\infty}(\mathcal{V})$.*

(b) *A class \mathcal{V} is a pseudo-variety if, and only if, $\mathcal{V} = \text{HSP}_{\aleph_0}(\mathcal{V})$.*

Proof. Let κ be either \aleph_0 or ∞ . The implication (\Leftarrow) is trivial. For (\Rightarrow) , it is sufficient to note that, by the above lemmas,

$$\begin{aligned} \text{HSP}_{\kappa} \circ \text{HSP}_{\kappa} &\subseteq \text{HSHP}_{\kappa} \text{SP}_{\kappa} \\ &\subseteq \text{HHSP}_{\kappa} \text{SP}_{\kappa} \\ &= \text{HSP}_{\kappa} \text{SP}_{\kappa} \subseteq \text{HSP}_{\kappa} \text{P}_{\kappa} = \text{HSP}_{\kappa}. \end{aligned} \quad \square$$

Notes

Again, most of the material in this chapter is standard and can be extracted from the literature, although I know of no systematic account. The abstract variety theorem in Section 5 is based on [16].

III. Languages

1. Weights

Our main point of interest is to determine which languages are definable in a given logic. As it does not make the presentation significantly more complicated, we will state our results in the more general setting of *weighted* languages. Thus, in the following we fix some set $\Omega \in \mathcal{D}$ of weights. The standard choice for Ω is of course the two element set $\{0, 1\}$. In this case, we will speak of an *unweighted* language. As this is the case we are most interested in, we will sometimes also present the unweighted version of statements and definitions, in particular, if it is easier to understand than the more general version. Other common choices for Ω include

- ♦ the interval $\langle [0, 1], \leq \rangle$,
- ♦ the *tropical semiring* $\langle \mathbb{N} \cup \{\infty\}, \min, +, \infty, 0 \rangle$,
- ♦ the *arctic semiring* $\langle \mathbb{N} \cup \{-\infty\}, \max, +, -\infty, 0 \rangle$,
- ♦ the fields $\langle \mathbb{Q}, +, \cdot, 0, 1 \rangle$ and $\langle \mathbb{R}, +, \cdot, 0, 1 \rangle$,
- ♦ the *language semiring* $\langle \mathcal{P}(\Sigma^*), \cup, \cdot, \emptyset, \{\langle \rangle\} \rangle$,
- ♦ matrices $\langle S^{n \times n}, +, \cdot, 0, I \rangle$ over some semiring S .

A language is then a function $\mathbb{M}\Sigma \rightarrow \Omega$ for some alphabet Σ , or, more generally, a function $A \rightarrow \Omega$ for an arbitrary set A . The set of all such functions can be canonically equipped with operations of the following form.

Definition 1.1. For functions $\omega : \Omega^n \rightarrow \Omega$ and $\kappa_0, \dots, \kappa_{n-1} : X \rightarrow \Omega$, we set

$$\omega[\tilde{\kappa}] := \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle, \quad \text{for } \kappa_0, \dots, \kappa_{n-1} : X \rightarrow \Omega.$$

An Ω -operation of arity n is a function of the form

$$\hat{\omega} : \mathcal{D}(X, \Omega)^n \rightarrow \mathcal{D}(X, \Omega) : \bar{\kappa} \mapsto \omega[\bar{\kappa}], \quad \text{for some } \omega : \Omega^n \rightarrow \Omega.$$

Usually, we will not distinguish notationally between the functions ω and $\hat{\omega}$.

Remark. In the unweighted case, a $\{0, 1\}$ -operation of the form

$$\text{Pos}(X, \{0, 1\})^n \rightarrow \text{Pos}(X, \{0, 1\})$$

is simply a positive Boolean combination of n elements.

For technical reasons, we frequently have to make one assumption on the set Ω : that every partial function $A \rightarrow \Omega$ can be extended to a total one. The formal definition is as follows.

Definition 1.2. Let \mathcal{D} be a category with a factorisation system $\langle E, M \rangle$. An object $S \in \mathcal{D}$ is *injective* if, for every M -morphism $\mu : A \rightarrow B$ and every morphism $\varphi : A \rightarrow S$, there exists a morphism $\hat{\varphi} : B \rightarrow S$ with $\hat{\varphi} \circ \mu = \varphi$.

Examples. (a) In Set and $\mathfrak{G}\text{-Set}$ every object is injective.

(b) In Pos a partial order $\langle A, \leq \rangle$ is injective if, and only if, it is complete, that is, every subset of A has an infimum and a supremum. Given a complete order $\langle S, \leq \rangle$ and an arbitrary order $\langle B, \leq \rangle$, we can extend every monotone function $f : A \rightarrow S$ with $A \subseteq B$ to a function $\hat{f} : B \rightarrow S$ by setting

$$\hat{f}(b) := \sup\{f(a) \mid a \in A, a \leq b\}.$$

Conversely, if $\langle S, \leq \rangle$ is not complete, then the identity map $\text{id} : S \rightarrow S$ has no extension to $C \rightarrow S$, where $\langle C, \leq \rangle$ is the completion of S .

Remark. For most applications, the requirement of injectivity is not a restriction as we can replace Ω by some (possibly larger) injective set. Categories where this is the case are said to *have enough injectives*. For instance, Pos has enough injectives since every partial order can be embedded into its completion.

Finite Sets Conditions

We also need a suitable notion of what it means for a set to be ‘finite’. There are several possible definitions, depending on the set Ω of weights.

Definition 1.3. (a) A set $A \in \mathcal{D}$ is *finite* if $\forall A$ is finite.

(b) A set $A \in \mathcal{D}$ is *strongly finite-dimensional* if there exists an M -morphism $e : A \rightarrow \Omega^d$, for some $d < \omega$.

(c) A set $A \in \mathcal{D}$ is *weakly finite-dimensional* if there exists a surjective E -morphism $\varepsilon : \hat{A} \rightarrow A$ and an M -morphism $\mu : \hat{A} \rightarrow \Omega^d$ with $d < \omega$.

(d) For some property P , we say that $A \in \mathcal{D}^\Xi$ is *sort-wise* P if A_ξ has property P , for every $\xi \in \Xi$.

(e) An \mathbb{M} -algebra \mathcal{A} is *finitary* if it is finitely generated and sort-wise finite.

Remark. (a) Clearly, the class of finite sets is closed under subsets, quotients, and finite products. The class of finitary algebras is closed under quotients and finitely generated subalgebras of finite products.

(b) It follows immediately from the definition that the class of strongly finite-dimensional sets is closed under subsets and finite products.

(c) If Ω is finite, so are all weakly finite-dimensional sets. The converse statement only holds under additional assumptions. We will prove one version below for $\mathcal{D} = \text{Pos}$.

While the class of sort-wise finite sets has the nicest closure properties, it is too restrictive when working with weighted languages since syntactic algebras are not always sort-wise finite. As a substitute, we can frequently use the class of sort-wise weakly finite-dimensional sets.

Lemma 1.4. *The class of weakly finite-dimensional sets is closed under surjective E -quotients, M -subobjects, and finite products.*

Proof. Set $\mathcal{P} := \{\mathbb{J}X \mid X \in \text{Set}\}$ and $E_* := \{\varepsilon \in E \mid \varepsilon \text{ surjective}\}$. Then every object A is \mathcal{P} -generated since the morphisms of the counit $\iota : \mathbb{J}V \Rightarrow \text{Id}$ belong to E . Consequently, it follows by Theorem II.5.12 that the closure of Ω under the above operations is given by $\text{HSP}_{\mathfrak{K}_0}(\{\Omega\})$, which is precisely the class of weakly finite-dimensional sets. \square

Noetherian Sets

As a technical tool let us introduce the notion of a noetherian set, which is a rather mild finiteness condition.

Definition 1.5. An object $A \in \mathcal{D}$ is *noetherian* if its congruence lattice $\text{Cong}(A)$ satisfies the *descending chain condition*, that is, if every strictly descending sequence $\varepsilon_0 > \varepsilon_1 > \dots$ of congruences of A is finite. \lrcorner

Examples. (a) In the categories Set , Pos , and Top an object A is noetherian if, and only if, it is finite.

(b) In Met an object is noetherian if, and only if, it has at most one element.

(c) In the category of modules over some fixed ring \mathfrak{R} , a module \mathfrak{M} is noetherian if, and only if, every ideal $I \subseteq \mathfrak{M}$ is finitely generated. In particular, if \mathfrak{R} is a field, the noetherian modules are precisely the finite-dimensional (both in our sense for $\Omega = \mathfrak{R}$, and in the sense of linear algebra) vector spaces. \lrcorner

Exercise 1.1. We consider the following eight partial orders in Pos .

- ◆ an infinite increasing chain of length ω , an increasing decreasing chain of length ω , a countably infinite antichain,
- ◆ the orders obtained from the three above ones by adding a new top element,
- ◆ the orders obtained from the three above ones by adding a new bottom element,
- ◆ the orders obtained from the three above ones by adding both a top and bottom element.

- (a) Prove that none of the eight above orders is noetherian.
- (b) Prove that, for every infinite $A \in \text{Pos}$, there exists a surjective map from A to at least one of the above orders.
- (c) Prove that an order $A \in \text{Pos}$ is noetherian if, and only if, it is finite. \lrcorner

The closure properties of noetherian sets are quite weak.

Lemma 1.6. If A is noetherian, so is every quotient of A .

Proof. Let $q : A \rightarrow B$ be an E -morphism. For a contradiction, suppose that there exists an infinite strictly descending sequence $\varepsilon_0 > \varepsilon_1 > \dots$ of quotients of B . Then $\varepsilon_0 \circ q \geq \varepsilon_1 \circ q \geq \dots$ forms a descending sequence of quotients of A . Since A is noetherian, we can find indices $i < k$ with $\varepsilon_i \circ q \leq \varepsilon_k \circ q$. Thus, there exists a morphism ρ with $\varepsilon_k \circ q = \rho \circ \varepsilon_i \circ q$. As q is an epimorphism, it follows that $\varepsilon_k = \rho \circ \varepsilon_i$, that is, $\varepsilon_i \leq \varepsilon_k$. A contradiction. \square

Concerning the existence of noetherian sets, we have the following observation.

Lemma 1.7. Let \mathcal{D} be a category such that, for every finite set $X \in \text{Set}$, there are, up to isomorphism, only finitely many $A \in \mathcal{D}$ with $\mathbb{V}A = X$. Then every finite $A \in \mathcal{D}$ is noetherian.

Proof. Let A be finite and $\varepsilon_0 > \varepsilon_1 > \dots$ an infinite strictly descending sequence of quotients of A . For $i < j$, let $\rho_{ij} : A/\varepsilon_i \rightarrow A/\varepsilon_j$ be the function with $\varepsilon_j = \rho_{ij} \circ \varepsilon_i$. We obtain a descending sequence

$$\ker \mathbb{V}\varepsilon_0 \supseteq \ker \mathbb{V}\varepsilon_1 \supseteq \dots$$

of equivalence relations on $\mathbb{V}A$. As A is finite, this sequence must stabilise at some k . Consequently,

$$\mathbb{V}\rho_{ij} : \mathbb{V}(A/\varepsilon_i) \rightarrow \mathbb{V}(A/\varepsilon_j) \text{ is bijective, for all } i, j \geq k.$$

As $\mathbb{V}(A/\varepsilon_i)$ is finite, it follows by assumption on \mathcal{D} that there exists an infinite set I such that

$$\rho_{ij} : A/\varepsilon_i \cong A/\varepsilon_j \text{ is an isomorphism, for all } i, j \in I.$$

Hence, $\varepsilon_j \leq \varepsilon_i$. A contradiction. \square

Finite-Dimensional Sets

To prove weak or strong finite-dimensionality, we have to consider embeddings into powers of Ω . We start by showing how to replace infinite powers by finite ones.

Lemma 1.8. *Let $e : A \rightarrow \Omega^\kappa$ be an M -morphism where κ is some cardinal. If A is noetherian, then A is strongly finite-dimensional.*

Proof. For $w \subseteq \kappa$, let $p_w : \Omega^\kappa \rightarrow \Omega^w$ be the projection to the components in w . Note that

$$u \subseteq w \quad \text{implies} \quad p_u \geq p_w.$$

Since A is noetherian, we can therefore find some finite set $w \subseteq \kappa$ such that $\ker(p_w \circ e) = \ker e = \text{id}$. Consequently, $p_w \circ e \in M$ and $p_w \circ e : A \rightarrow \Omega^w$ is the desired embedding. \square

In light of the preceding lemma, it is of interest to know which sets can be embedded into some power of Ω . The following general notion is useful.

Definition 1.9. A set $C \in \mathcal{D}$ is a *cogenerator* if, for every pair of functions $f, g : A \rightarrow B$,

$$k \circ f = k \circ g, \quad \text{for all } k : B \rightarrow C, \quad \text{implies} \quad f = g.$$

Given a factorisation system $\langle E, M \rangle$, we call C an *EM-cogenerator* if, for every $A \in \mathcal{D}$,

$$\left(\inf_{k:A \rightarrow C} \ker k \right) \leq \text{id}_A,$$

where the infimum is computed in $\text{Cong}(A)$. \lrcorner

Lemma 1.10. *Let \mathcal{D} be a category with arbitrary products and let $\langle E, M \rangle$ be a factorisation system.*

- (a) $\Omega \in \mathcal{D}$ is a cogenerator of \mathcal{D} if, and only if, for every set $A \in \mathcal{D}$, there exists a monomorphism $A \rightarrow \Omega^\kappa$, for some cardinal κ .

- (b) If Ω is an EM-cogenerator, we can choose the monomorphism to belong to M .

Proof. (a) follows from the fact that Ω is a cogenerator if, and only if, for every set $A \in \mathcal{D}$, the morphism

$$\varphi := \langle k \rangle_{k:A \rightarrow \Omega} : A \rightarrow \Omega^{\mathcal{D}(A, \Omega)}$$

is a monomorphism.

(b) Let φ be the morphism from (a). By Lemma II.4.3, we have $\ker \varphi = \inf_k \ker k$. As Ω is an EM-cogenerator, it follows that $\ker \varphi = \text{id}$. Hence, $\varphi \in M$. \square

Corollary 1.11. *If Ω is an EM-cogenerator, every noetherian set A is strongly finite-dimensional.*

Proof. By Lemma 1.10, for every set A , we can find an M -morphism $A \rightarrow \Omega^\kappa$, for some κ . If A is noetherian, we can therefore use Lemma 1.8 to show that it is finite-dimensional. \square

The existence of EM-cogenerators depends on the category in question. We start with Pos.

Lemma 1.12. *If $\Omega \in \text{Pos}$ is not an antichain, then every partial order A can be embedded into Ω^κ , for some κ .*

Proof. Fix two values $u, v \in \Omega$ with $u < v$. For each $c \in A$, we define a function $\mu_c : A \rightarrow \Omega$ by

$$\mu_c(a) := \begin{cases} v & \text{if } a \geq c, \\ u & \text{otherwise.} \end{cases}$$

Note that μ_c is monotone since

$$\begin{aligned} \mu_c(a) \not\leq \mu_c(b) &\Rightarrow \mu_c(a) = u \quad \text{and} \quad \mu_c(b) = v \\ &\Rightarrow a \geq c \quad \text{and} \quad b \not\geq c \\ &\Rightarrow a \not\leq b. \end{aligned}$$

To see that $\mu := \langle \mu_c \rangle_{c \in A}$ is an embedding $A \rightarrow \Omega^A$, we have to check that

$$a \leq b \quad \text{iff} \quad \mu_c(a) \leq \mu_c(b), \quad \text{for all } c \in A.$$

(\Rightarrow) follows by monotonicity of μ_c . For (\Leftarrow), we have

$$\begin{aligned} a \not\leq b &\Rightarrow \mu_a(a) = v \quad \text{and} \quad \mu_a(b) = u \\ &\Rightarrow \mu_a(a) \not\leq \mu_a(b). \end{aligned} \quad \square$$

Corollary 1.13. *Suppose that $\mathcal{D} = \text{Pos}$ and that Ω is not an antichain. Every finite set A is strongly finite-dimensional.*

Remark. Note that no antichain is injective in Pos . Hence, the assumptions of the preceding two results are automatically satisfied in our setting.

Proposition 1.14. *Let $\mathfrak{S} = \{0, 1\} \in \text{Top}$ be Sierpiński space (where the closed sets are \emptyset , $\{0\}$, and $\{0, 1\}$). For a sober space $\Omega \in \text{Top}$ the following statements are equivalent.*

- (1) *Every sober space $\mathfrak{X} \in \text{Top}$ can be embedded (via an injective continuous map) into Ω^κ , for some κ .*
- (2) *There exists an injective continuous map $\mathfrak{S} \rightarrow \Omega$.*

Proof. (1) \Rightarrow (2) Since \mathfrak{S} is sober, we can use (1) to find an embedding $e : \mathfrak{S} \rightarrow \Omega^\kappa$, for some κ . Set $x := e(0)$ and $y := e(1)$, and let $p_i : \Omega^\kappa \rightarrow \Omega$ be the projection to the i -th component. Since $x \neq y$, there is some $i < \kappa$ with $p_i(x) \neq p_i(y)$. Ω being sober, we can find a closed set C containing one of $p_i(x)$ and $p_i(y)$, but not both. By symmetry, we may assume that $x \in C$. We claim that the function $f : \mathfrak{S} \rightarrow \Omega$ with $f(0) := x$ and $f(1) := y$ is the desired embedding. For continuity, let $D \subseteq \Omega$ be closed. If $f^{-1}[D]$ is one of \emptyset , $\{0\}$, or S , we are done. Hence, suppose that $f^{-1}[D] = \{1\}$, i.e., $p_i(y) \in D$ and $p_i(x) \notin D$. Then $U := p_i^{-1}[C]$ and $V := p_i^{-1}[D]$ are closed sets in Ω^κ with $x \in U \setminus V$ and $y \in V \setminus U$. Since e is continuous, it follows that $e^{-1}[V] = e^{-1}(y) = \{1\}$ is closed in \mathfrak{S} . A contradiction.

(2) \Rightarrow (1) The embedding $\mathfrak{S} \rightarrow \Omega$ induces an embedding $\mathfrak{S}^\kappa \rightarrow \Omega^\kappa$. Hence, it is sufficient to prove that every sober space \mathfrak{X} can be embedded

into \mathfrak{S}^κ , for some κ . Fix \mathfrak{X} . For a closed set $C \subseteq X$, let $\mu_C : X \rightarrow S$ be the map defined by

$$\mu_C(x) := \begin{cases} 0 & \text{if } x \in C, \\ 1 & \text{otherwise.} \end{cases}$$

Then μ_C is continuous since $\mu_C^{-1}[\emptyset] = \emptyset$, $\mu_C^{-1}[\{0\}] = C$, and $\mu_C^{-1}[S] = X$ are all closed. Setting $e := \langle \mu_C \rangle_C$ we obtain a continuous map $\mathfrak{X} \rightarrow \mathfrak{S}^T$, where T is the class of all closed sets in \mathfrak{X} .

It remains to see that e is injective. As \mathfrak{X} is sober we can find, for every pair of distinct points $x \neq y$, some closed set C containing one of them, but not the other. Consequently, $\mu_C(x) \neq \mu_C(y)$, which implies that $e(x) \neq e(y)$. \square

Lemma 1.15. *Let $\Omega = [0, 1] \in \text{Met}$ with the metric $d(x, y) := |x - y|$. Every metric space \mathfrak{X} can be embedded into Ω^X .*

Proof. Let $\mathfrak{X} \in \text{Met}$. For every $z \in X$, we define a function $f_z : \mathfrak{X} \rightarrow \Omega$ via $f_z(x) := d(z, x)$. We claim that f_z is non-expansive and that the function $e := \langle f_z \rangle_{z \in X} : \mathfrak{X} \rightarrow \Omega^X$ is the desired embedding. For the first claim, note that

$$d(f_z(x), f_z(y)) = |d(z, x) - d(z, y)| \leq d(x, y),$$

for all $x, y, z \in X$. Hence,

$$d(e(x), e(y)) \leq \sup_z d(f_z(x), f_z(y)) \leq d(x, y), \quad \text{for all } x, y \in X.$$

For the second one, we have to show that e is an isometry. For $x, y \in X$, we have

$$\begin{aligned} d(x, y) &\leq |d(x, x) - d(x, y)| \\ &= |f_x(x) - f_x(y)| \\ &\leq d(f_x(x), f_x(y)) \\ &\leq \sup_z d(f_z(x), f_z(y)) = d(e(x), e(y)) \leq d(x, y). \end{aligned} \quad \square$$

Quasi-Finite Sets

Finally, let us take a quick look at the following generalisation of the notion of an orbit finite set from nominal set theory.

Definition 1.16. An object $A \in \mathcal{D}$ is *quasi-finite* if there exists a surjective morphism $q : \mathbb{J}X \rightarrow A$, for some finite set X . \lrcorner

Lemma 1.17. Every finite set A is quasi-finite. If the unit $\varepsilon : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ of the adjunction is an isomorphism, the converse also holds.

Proof. By Lemma I.2.3, the counit morphism $\iota_A : \mathbb{J}\mathbb{V}A \rightarrow A$ is surjective.

For the second statement, suppose that $\varepsilon_X : X \rightarrow \mathbb{V}\mathbb{J}X$ is an isomorphism and let $f : \mathbb{J}X \rightarrow A$ be a surjective function with finite X . Then $\mathbb{V}f \circ \varepsilon_X : X \rightarrow \mathbb{V}A$ is surjective and $|\mathbb{V}A| \leq |X|$ is finite. \square

Remark. This lemma implies that in the categories Set , Pos , Top , and Met , quasi-finiteness is the same notion as finiteness and therefore not that interesting. In $\mathfrak{G}\text{-Set}$ a set is quasi-finite if, and only if, it has only finitely many orbits under \mathfrak{G} . Thus, quasi-finiteness is strictly weaker than finiteness and it is frequently the more useful concept. \lrcorner

Lemma 1.18. The class of quasi-finite sets is closed under subsets, finite products, and images of surjective morphisms.

Proof. For images under surjective morphisms note that, if $q : A \rightarrow B$ and $f : \mathbb{J}X \rightarrow A$ are surjective, so is $q \circ f : \mathbb{J}X \rightarrow B$.

For subsets, suppose that $f : \mathbb{J}X \rightarrow A$ is surjective and let $i : C \rightarrow A$ be a monomorphism. Then $\mathbb{V}i : \mathbb{V}C \rightarrow \mathbb{V}A$ is injective and so is its pullback $j : Z \rightarrow \mathbb{V}\mathbb{J}X$ along the function $\hat{f} : X \rightarrow \mathbb{V}A$ corresponding to f via the adjunction. Let $q : Z \rightarrow \mathbb{V}C$ be the other morphism of the pullback square and let $\hat{q} : \mathbb{J}Z \rightarrow C$ be the morphism corresponding to it via the adjunction. Note that q is surjective since, in Set , the pullback of an epimorphism is an epimorphism. Furthermore, we have seen in Lemma I.2.3 (e) that the counit $\iota : \mathbb{J}\mathbb{V}C \rightarrow C$ is surjective. Hence, so is $\hat{q} = \iota \circ \mathbb{J}q$.

For the empty product, note that \mathbb{V} preserves products. Hence, $\mathbb{V}1_{\mathcal{D}} = 1_{\text{Set}}$ and the unique map $u : 1_{\text{Set}} \rightarrow \mathbb{V}1_{\mathcal{D}}$ corresponds to a morphism $\hat{u} : \mathbb{J}1_{\text{Set}} \rightarrow 1_{\mathcal{D}}$, which is clearly surjective.

For binary products, suppose that $f : \mathbb{J}X \rightarrow A$ and $g : \mathbb{J}Y \rightarrow B$ are surjective. Let $p : X \times Y \rightarrow X$ and $q : X \times Y \rightarrow Y$ be the two projections. Note that p and q are surjective since \mathbb{V} preserves limits. Hence, $\mathbb{V}(X \times Y) = \mathbb{V}X \times \mathbb{V}Y$ and $\mathbb{V}p$ and $\mathbb{V}q$ are the corresponding projections. Consequently, $f \circ p$ and $g \circ q$ are surjective and, hence, so is $(f \circ p, g \circ q) : \mathbb{J}(X \times Y) \rightarrow A \times B$. \square

Corollary 1.19. If Ω is quasi-finite, then so is every weakly finite-dimensional set.

Proof. Suppose that Ω is quasi-finite and let A be weakly finite-dimensional. Fix an M -morphism $\mu : \tilde{A} \rightarrow \Omega^d$ and a surjective E -morphism $\varepsilon : \tilde{A} \rightarrow A$ with finite d . Then it follows by the preceding lemma that Ω^d , \tilde{A} , and A are also quasi-finite. \square

2. Languages

These preparations out of the way, we can finally start to develop formal language theory. Our abstract setting is as follows.

Conventions.

- (I) \mathcal{D} is a discrete category with arbitrary limits and colimits.
- (II) $\Omega \in \mathcal{D}$ is injective.
- (III) $\langle \mathbb{M}, \text{flat}, \text{sing} \rangle$ is a polynomial monad on \mathcal{D}^{Ξ} such that \mathbb{M} preserves epimorphisms.
- (IV) $\langle E, M \rangle$ is the standard factorisation system on \mathcal{D} , which is well-defined and M is closed under coproducts.

Definition 2.1. (a) An *alphabet* is an object $\Sigma \in \mathcal{D}^{\Xi}$ that is isomorphic to one of the form $\mathbb{J}X$, for some finite set X . We denote by Alph the category

of all alphabets whose morphisms are all functions of the form $\mathbb{J}f$, for some function f on finite sets.

(b) An Ω -language over the alphabet Σ is a function $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$, for some sort ξ .

(c) A family of Ω -languages is a function \mathcal{K} mapping each alphabet Σ to a class $\mathcal{K}[\Sigma]$ of languages over Σ .

(d) A function $f : A \rightarrow B$ recognises a function $\kappa : A_\xi \rightarrow \Omega$ if

$$\kappa = \mu \circ f, \quad \text{for some function } \mu : B_\xi \rightarrow \Omega.$$

(e) Let $f : A \rightarrow B$ and $\kappa : B_\xi \rightarrow \Omega$ be functions. The inverse image of κ is

$$f^{-1}[\kappa] := \kappa \circ f : A_\xi \rightarrow \Omega.$$

There are two common special cases. Functions of the form $f = \mathbb{M}g$, for some morphism $f : \Sigma \rightarrow \Gamma$ of \mathbf{Alph} , are called *relabellings*. In this case, we call $f^{-1}[\kappa]$ an *inverse relabelling* of κ . Similarly, if $f : \mathbb{M}\Sigma \rightarrow \mathbb{M}\Gamma$ is a morphism of \mathbb{M} -algebras, we call $f^{-1}[\kappa]$ the image of κ under an *inverse morphism*.

Remark. In the unweighted case for $\mathcal{D} = \mathbf{Pos}$, a language $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \{0, 1\}$ of course corresponds to a subset $K := \kappa^{-1}(1) \subseteq \mathbb{M}_\xi \Sigma$. Such a language is recognised by $f : \mathbb{M}\Sigma \rightarrow A$ if there is an upwards closed set $P \subseteq A$ such that $K = f^{-1}[P]$.

The inverse image of K under a function $g : \mathbb{M}\Sigma \rightarrow C$ is simply the preimage

$$g^{-1}[K] := \{c \in C \mid g(c) \in K\}.$$

Note that we always assume alphabets to be of the form $\mathbb{J}X$. This is required for the variety theorem in the next section. But sometimes it is useful to also work with languages over other ‘alphabets’ C . We do so by simply replacing C by $\mathbb{J}VC$. This leads to the following extension of the notion of a family of languages.

Definition 2.2. Let \mathcal{K} be a family of languages. For an arbitrary $C \in \mathcal{D}^\Xi$, we define

$$\mathcal{K}[C] := \{ \kappa : \mathbb{M}_\xi C \rightarrow \Omega \mid \xi \in \Xi, \kappa \circ \mathbb{M}\iota \in \mathcal{K}[\mathbb{J}VC] \},$$

where $\iota : \mathbb{J}V \Rightarrow \text{Id}$ is the counit of the adjunction $\mathbb{J} \dashv V$.

Example. Let Ω be a semiring and Σ an alphabet. A *weighted automaton* $\mathcal{A} = \langle \Omega, \Sigma, \delta, i, t \rangle$ (for words) over Ω consists of a finitely generated Ω -module Ω , an *initial state* $i \in Q$, a *final form* $t : Q \rightarrow \Omega$, and a *transition function* $\delta : \Sigma^* \rightarrow \text{End}(\Omega)$, which is a monoid homomorphism into the monoid $\text{End}(\Omega)$ of all linear maps $\Omega \rightarrow \Omega$. Such an automaton recognises the language $\kappa : \Sigma^* \rightarrow \Omega$ defined by

$$\kappa(w) := t(\delta(w)(i)).$$

Note that this language is recognised by the morphism $\delta : \Sigma^* \rightarrow \text{End}(\Omega)$ since

$$\kappa = \mu \circ \delta \quad \text{where} \quad \mu(f) := t(f(i)).$$

Language Algebras

While we are mostly interested in Ω -languages (and the special case of unweighted languages), some of the results below will be developed in a more abstract setting which hopefully simplifies the exposition. To this end, let us introduce a (contravariant) functor \mathbb{L}° that maps each set A to a *language algebra* $\mathbb{L}^\circ A$, whose elements are just the Ω -languages $\kappa : A \rightarrow \Omega$.

Definition 2.3. (a) Let $A \in \mathcal{D}$. We denote by $\mathbb{L}^\circ A := \mathcal{D}(A, \Omega)$ the algebra whose elements are all functions $A \rightarrow \Omega$ and whose operations are all Ω -operations.

(b) For a function $f : A \rightarrow B$ in \mathcal{D} , we denote by $\mathbb{L}^\circ f : \mathbb{L}^\circ B \rightarrow \mathbb{L}^\circ A$ the function $\mathbb{L}^\circ f(\kappa) := \kappa \circ f$. We call $\mathbb{L}^\circ f(\kappa)$ the *inverse image* of κ under f and we usually use the more suggestive notation $f^{-1}[\kappa] := \mathbb{L}^\circ f(\kappa)$.

Remark. (a) For unweighted languages we can again simplify this definition. If we work in an unordered setting, i.e., with algebras in Set , we can use power-set algebras

$$\mathbb{L}^\circ A := \langle \wp(A), \cup, \cap, \neg \rangle.$$

For unweighted languages in Pos , we use

$$\mathbb{L}^\circ A := \langle \mathbb{U}A, \cup, \cap \rangle$$

instead, where $\mathbb{U}A$ denotes the set of upwards-closed subsets of A . In both cases, $f^{-1}[K]$ is just the preimage of $K \subseteq B$ under $f : A \rightarrow B$.

To see that we only need the above operations, note that every function $[2]^n \rightarrow [2]$ can be expressed using disjunction, conjunction, and negation. Similarly, every monotone function $[2]^n \rightarrow [2]$ can be expressed using disjunction and conjunction only.

(b) Using the above notation, we can say that an Ω -language $\kappa \in \mathbb{L}^\circ \mathbb{M}_\xi \Sigma$ is recognised by a morphism $\varphi : \mathbb{M}\Sigma \rightarrow \mathfrak{A}$ if, and only if, $\kappa = \varphi_\xi^{-1}[\mu]$, for some $\mu \in \mathbb{L}^\circ A_\xi$.

(c) Note that, for $f : A \rightarrow B$, the function $\mathbb{L}^\circ f$ is indeed a homomorphism $\mathbb{L}^\circ B \rightarrow \mathbb{L}^\circ A$ since, for each Ω -operation $\omega : \Omega^n \rightarrow \Omega$,

$$\begin{aligned} \omega[\mathbb{L}^\circ f(\kappa_0), \dots, \mathbb{L}^\circ f(\kappa_{n-1})] &= \omega \circ \langle \kappa_0 \circ f, \dots, \kappa_{n-1} \circ f \rangle \\ &= \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle \circ f \\ &= \mathbb{L}^\circ f(\omega[\kappa_0, \dots, \kappa_{n-1}]). \end{aligned}$$

Lemma 2.4. *If $\mu : A \rightarrow B$ is an M -morphism, then $\mathbb{L}^\circ \mu : \mathbb{L}^\circ B \rightarrow \mathbb{L}^\circ A$ is surjective.*

Proof. Let $\kappa : A \rightarrow \Omega$ be long to $\mathbb{L}^\circ A$. Since Ω is injective, there exists some function $\lambda : B \rightarrow \Omega$ with $\lambda \circ \mu = \kappa$. Thus, $\kappa = \mathbb{L}^\circ \mu(\lambda)$. \square

Language algebras $\mathbb{L}^\circ A$ over finite-dimensional sets A are particularly simple. Let us mention two properties.

Lemma 2.5. *A set $A \in \mathcal{D}$ is strongly finite-dimensional if, and only if, the language algebra $\mathbb{L}^\circ A$ is finitely generated.*

Proof. (\Rightarrow) Fix an M -morphism $\mu : A \rightarrow \Omega^d$ and let $\kappa : A \rightarrow \Omega$ be an element of $\mathbb{L}^\circ A$. Since Ω is injective, there exists a function $\omega : \Omega^d \rightarrow \Omega$ with $\kappa = \omega \circ \mu$. Set $\lambda_i := p_i \circ \mu$ where $p_i : \Omega^d \rightarrow \Omega$ is the projection to the i -th component. Then

$$\begin{aligned} \omega[\bar{\lambda}] &= \omega \circ \langle \lambda_0, \dots, \lambda_{d-1} \rangle \\ &= \omega \circ \langle p_0 \circ \mu, \dots, p_{d-1} \circ \mu \rangle = \omega \circ \mu = \kappa. \end{aligned}$$

(\Leftarrow) Fix generators μ_0, \dots, μ_{n-1} of $\mathbb{L}^\circ A$. By assumption, we can find an Ω -operation ω with $\text{id} = \omega[\bar{\mu}']$ for some subset $\bar{\mu}'$ of the generators. Suppose that $\bar{\mu}' = \mu_0, \dots, \mu_{d-1}$ and set

$$e := \langle \mu_0, \dots, \mu_{d-1} \rangle : A \rightarrow \Omega^d.$$

Since $\omega \circ e = \omega[\bar{\mu}'] = \text{id}$, the function e has a left-inverse. This implies that $e \in M$. \square

Corollary 2.6. *A set $A \in \mathcal{D}$ is weakly finite-dimensional if, and only if, the language algebra $\mathbb{L}^\circ A$ is a subalgebra of a finitely generated one.*

Lemma 2.7. *Let $\varepsilon : A \rightarrow B$ and $\eta : A \rightarrow C$ be E -morphisms in \mathcal{D} where C is strongly finite-dimensional. Then*

$$\varepsilon \leq \eta \quad \text{iff} \quad \mathbb{L}^\circ \varepsilon[\mathbb{L}^\circ B] \supseteq \mathbb{L}^\circ \eta[\mathbb{L}^\circ C].$$

Proof. (\Rightarrow) Suppose that there exists a function $\rho : B \rightarrow C$ with $\eta = \rho \circ \varepsilon$. For $\kappa \in \mathbb{L}^\circ C$ it follows that

$$\mathbb{L}^\circ \eta(\kappa) = \kappa \circ \eta = \kappa \circ \rho \circ \varepsilon = \mathbb{L}^\circ \varepsilon(\kappa \circ \rho) \in \mathbb{L}^\circ \varepsilon[\mathbb{L}^\circ B].$$

(\Leftarrow) Since C is strongly finite-dimensional, there exists an M -morphism $e : C \rightarrow \Omega^d$, for some $d < \omega$. Let $p_i : \Omega^d \rightarrow \Omega$ be the projection to the i -th component. Note that $p_i \circ e : C \rightarrow \Omega$ belongs to $\mathbb{L}^\circ C$. Hence, $\mathbb{L}^\circ \eta[\mathbb{L}^\circ C] \subseteq \mathbb{L}^\circ \varepsilon[\mathbb{L}^\circ B]$ implies that there exist $\kappa_i \in \mathbb{L}^\circ B$ with

$$\kappa_i \circ \varepsilon = p_i \circ e \circ \eta.$$

It follows that

$$\langle \kappa_i \rangle_i \circ \varepsilon = \langle p_i \circ e \circ \eta \rangle_i = \langle p_i \rangle_i \circ e \circ \eta = e \circ \eta.$$

By the diagonal fill-in property, we can therefore find a function $\delta : B \rightarrow C$ with $\delta \circ \varepsilon = \eta$ and $e \circ \delta = \langle \kappa_i \rangle_i$. In particular, $\varepsilon \leq \eta$. \square

For finite-dimensional sets, the language algebra of a finite product is generated by languages over the various projections.

Lemma 2.8. *Let $A, B \in \mathcal{D}$ be strongly finite-dimensional and let $p : A \times B \rightarrow A$ and $q : A \times B \rightarrow B$ be the two projections. Then*

$$\mathbb{L}^\circ(A \times B) = \langle\langle p^{-1}[\mathbb{L}^\circ A] \cup q^{-1}[\mathbb{L}^\circ B] \rangle\rangle,$$

where $\langle\langle X \rangle\rangle$ denotes the subalgebra of $\mathbb{L}^\circ(A \times B)$ generated by X .

Proof. (\supseteq) is trivial. For (\subseteq) , fix $\kappa \in \mathbb{L}^\circ(A \times B)$. As A and B are finite-dimensional, there exist M -embeddings $e : A \rightarrow \Omega^m$ and $f : B \rightarrow \Omega^n$. Then $e \times f \in M$ and, since Ω is injective, we can find a function $\omega : \Omega^{m+n} \rightarrow \Omega$ with $\kappa = \omega \circ (e \times f)$. Consequently,

$$\begin{aligned} \kappa &= \omega \circ (e \times f) \\ &= \omega[(u_0 \circ e \circ p), \dots, (u_{m-1} \circ e \circ p), \\ &\quad (v_0 \circ f \circ q), \dots, (v_{n-1} \circ f \circ q)], \end{aligned}$$

where $u_i : \Omega^m \rightarrow \Omega$ and $v_i : \Omega^n \rightarrow \Omega$ are the projections. Since $u_i \circ e \in p^{-1}[\mathbb{L}^\circ A]$ and $v_i \circ f \in q^{-1}[\mathbb{L}^\circ B]$, it follows that

$$\kappa \in \langle\langle p^{-1}[\mathbb{L}^\circ A] \cup q^{-1}[\mathbb{L}^\circ B] \rangle\rangle. \quad \square$$

Corollary 2.9. *If $A, B \in \mathcal{D}$ are strongly finite-dimensional, then*

$$\mathbb{L}^\circ(A \times B) \cong \mathbb{L}^\circ A + \mathbb{L}^\circ B.$$

3. Minimal Algebras

Let us start with the observation that the recognisability of a language $\kappa : A_\xi \rightarrow \Omega$ over an algebra \mathfrak{A} depends on which \mathbb{M} -congruences of \mathfrak{A} are contained in $\ker \kappa$. (Note that $\ker \kappa$ itself is usually not an \mathbb{M} -congruence.)

Lemma 3.1. *Let \mathfrak{A} be an \mathbb{M} -algebra and $\kappa : A_\xi \rightarrow \Omega$ a language.*

- (a) *An E -morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ of \mathbb{M} -algebras recognises κ if, and only if, $\ker \varphi \leq \ker \kappa$.*
- (b) *If $\varepsilon \leq \ker \kappa$ is an \mathbb{M} -congruence, then $\varepsilon : \mathfrak{A} \rightarrow \mathfrak{A}/\varepsilon$ recognises κ .*
- (c) *κ is recognised by a morphism into an \mathbb{M} -algebra with some property P if, and only if, there exists an \mathbb{M} -congruence $\varepsilon \leq \ker \kappa$ such that the quotient \mathfrak{A}/ε has property P .*

Proof. Note that (b) follows by (a) and (c) by (b). Hence, we only need to prove (a). If φ recognises κ , there exists some function $\mu : B_\xi \rightarrow \Omega$ such that $\kappa = \mu \circ \varphi$. This implies that $\ker \varphi \leq \ker \kappa$. The converse follows immediately by definition of the order on quotients:

$$\ker \varphi \leq \ker \kappa \quad \text{iff} \quad \rho \circ \ker \varphi = \ker \kappa, \quad \text{for some } \rho.$$

Since $\varphi \in E$, we have $\ker \varphi = \varphi$. Furthermore, $\kappa = i \circ \ker \kappa$, for some $i \in M$. Hence, we have

$$\rho \circ \ker \varphi = \ker \kappa \quad \text{iff} \quad (i \circ \rho) \circ \varphi = \kappa,$$

and the claim follows for $\mu := i \circ \rho$. \square

It follows that classifying the algebras recognising κ amounts to classifying all \mathbb{M} -congruences contained in $\ker \kappa$. In particular, if there is a coarsest such congruence, then there exists a smallest algebra recognising κ .

Definition 3.2. Let $\kappa \in A_\xi \rightarrow \Omega$ be a language and $\varphi : \mathfrak{A} \rightarrow \mathfrak{M}$ an E -morphism of \mathbb{M} -algebras recognising κ . We call φ the *minimal morphism* of κ if, for every surjective morphism $\psi : \mathfrak{A} \rightarrow \mathfrak{B}$ recognising κ , there exists a unique morphism $\rho : \mathfrak{B} \rightarrow \mathfrak{M}$ such that $\varphi = \rho \circ \psi$. In this case we call \mathfrak{M} the *minimal algebra* of κ . \lrcorner

Remark. (a) In category-theoretic language, the minimal morphism of κ is the terminal object in the category of all surjective morphisms recognising κ . In particular, it is unique up to isomorphism.

(b) In terms of congruences, it follows that, if φ is the minimal morphism of κ , then $\ker \varphi$ is the maximal \mathbb{M} -congruence with $\ker \varphi \leq \ker \kappa$.

Let us start with an example showing that minimal morphisms do not need to exist.

Example. Recall the monad $\mathbb{C} = \mathbb{C}_{\aleph_0}$ for countable chains from the example on page 35. We consider the language

$$K := \{ w \in \mathbb{C}\{a\} \mid \text{for every } n < \omega, w \text{ has a factorisation } w = ua^n v \}.$$

Every morphism $\varphi : \mathbb{C}\{a\} \rightarrow \mathfrak{B}$ recognising K satisfies

$$\varphi(a^m) \neq \varphi(a^n), \quad \text{for } m \neq n.$$

To prove that K does not have a minimal morphism, consider the words

$$w_n := \eta a^n \eta, \quad \text{for } n < \omega$$

where η denotes a dense order labelled with a , and let $h_n : \mathbb{C}\{a\} \rightarrow \mathbb{C}\{a\}$ be the function with

$$h_n(x) := \begin{cases} w_k & \text{if } x = w_{k+1} \text{ for } k \geq n, \\ x & \text{otherwise.} \end{cases}$$

Note that there exists a function $\sigma_n : \mathbb{C}\{a\} \rightarrow \mathbb{C}\{a\}$ such that

$$h_n \circ \text{flat} = \sigma_n \circ \mathbb{C}\sigma.$$

As h_n is surjective, it therefore follows by Lemma I.5.5 that $\mathfrak{A}_n := \langle \mathbb{C}\{a\}, \sigma_n \rangle$ forms an \mathbb{C} -algebra and $h_n : \mathbb{C}\{a\} \rightarrow \mathfrak{A}_n$ a morphism.

For a contradiction, suppose that K has a minimal morphism $\varphi : \mathbb{C}\{a\} \rightarrow \mathfrak{M}$. Then there are morphisms $\rho_n : \mathfrak{A}_n \rightarrow \mathfrak{M}$ with $\rho_n \circ h_n = \varphi$. Hence,

$$\varphi(w_{n+1}) = \rho_n(h_n(w_{n+1})) = \rho_n(w_n) = \rho_n(h_n(w_n)) = \varphi(w_n),$$

for all $n < \omega$. This implies that

$$\varphi(w_1 w_1 w_1 \dots) = \varphi(w_1 w_2 w_3 \dots).$$

But $w_1 w_1 w_1 \dots \notin K$ while $w_1 w_2 w_3 \dots \in K$. A contradiction.

Note that, by Lemma 3.1, the existence of a minimal algebra for a language κ only depends on the kernel $\ker \kappa$.

Lemma 3.3. *Let \mathfrak{A} be an \mathbb{M} -algebra. A language $\kappa \in A_\xi \rightarrow \Omega$ has a minimal algebra if, and only if, there exists a greatest \mathbb{M} -congruence $\varepsilon \leq \ker \kappa$.*

Proof. (\Leftarrow) Let $\eta : \mathfrak{A} \rightarrow \mathfrak{M}$ be the minimal morphism of κ . We claim that $\ker \eta$ is the greatest \mathbb{M} -congruence contained in $\ker \kappa$. First, note that Lemma 3.1 implies that $\ker \eta \leq \ker \kappa$ since η recognises κ .

For maximality, consider an \mathbb{M} -congruence $\varepsilon \leq \ker \kappa$. By Lemma 3.1, $\varepsilon : \mathfrak{A} \rightarrow \mathfrak{A}/\varepsilon$ recognises κ . By minimality of η , we can therefore find a morphism $\rho : \mathfrak{A}/\varepsilon \rightarrow \mathfrak{M}$ with $\eta = \rho \circ \varepsilon$. This implies that $\varepsilon \leq \ker(\rho \circ \varepsilon) = \ker \eta$.

(\Rightarrow) We claim that $\varepsilon : \mathfrak{A} \rightarrow \mathfrak{A}/\varepsilon$ is the minimal morphism of κ . Hence, suppose that $\psi : \mathfrak{A} \rightarrow \mathfrak{C}$ is surjective and recognises κ . By Lemma 3.1, it follows that $\ker \psi \leq \ker \kappa$. As ψ is a morphism of \mathbb{M} -algebras, its kernel $\ker \psi$ is an \mathbb{M} -congruence. By choice of ε , it follows that $\ker \psi \leq \varepsilon$. Consequently, there exists a morphism $\rho : \mathfrak{C} \rightarrow \mathfrak{A}/\varepsilon$ with $\rho \circ \psi = \varepsilon$. Uniqueness of ρ follows from the fact that ψ is surjective. \square

We will prove the existence of minimal algebras in the case $\mathcal{D} = \text{Pos}$ where we can work with congruence orderings. (The exact same proofs work for Set and $\mathfrak{G}\text{-Set}$.) To use the preceding lemma, we need to construct maximal congruences.

Definition 3.4. We denote the *transitive closure* of a relation $\theta \subseteq A \times A$ by $\text{TC}(\theta)$.

Lemma 3.5. *Let $\mathcal{D} = \text{Pos}$ and let \mathfrak{A} be an \mathbb{M} -algebra.*

(a) *$\text{Cong}(A)$ forms a complete lattice where*

$$\inf \Theta = \bigcap \Theta \quad \text{and} \quad \sup \Theta = \text{TC}(\bigcup \Theta), \quad \text{for } \Theta \subseteq \text{Cong}(A).$$

(b) $\text{Cong}_{\mathbb{M}}(\mathfrak{A})$ forms a complete lattice where

$$\inf \Theta = \bigcap \Theta, \quad \text{for } \Theta \subseteq \text{Cong}_{\mathbb{M}}(\mathfrak{A}).$$

Proof. Part (a) is obvious and (b) follows by Lemma II.4.12. \square

Unfortunately, joins of congruences are much harder to compute. Let us collect a few lemmas to do so. We start with a technical result showing how to compute transitive closures of relations on $\mathbb{M}A$.

Lemma 3.6. *Let Θ be a set of reflexive binary relations on some set $A \in \text{Pos}^{\Xi}$, and let $\sigma \subseteq \text{TC}(\bigcup \Theta)$ be finite. Then*

$$s \sigma^{\mathbb{M}} t \quad \text{implies} \quad s = r_0 \theta_0^{\mathbb{M}} \dots \theta_{n-1}^{\mathbb{M}} r_n = t, \\ \text{for some } \theta_0, \dots, \theta_{n-1} \in \Theta \text{ and } r_0, \dots, r_n \in \mathbb{M}A.$$

Proof. Suppose that $s \sigma^{\mathbb{M}} t$. By definition, there exists some $u \in \mathbb{M}\sigma$ such that $s = \mathbb{M}p(u)$ and $t = \mathbb{M}q(u)$, where $p, q : A \times A \rightarrow A$ are the two projections. We fix an enumeration $\langle a_0, b_0 \rangle, \dots, \langle a_{n-1}, b_{n-1} \rangle$ of σ and we define $g_k : \sigma \rightarrow A$, for $k \leq n$, by

$$g_k(\langle a_i, b_i \rangle) := \begin{cases} a_i & \text{if } i \geq k, \\ b_i & \text{if } i < k. \end{cases}$$

Then $s = \mathbb{M}g_0(u)$ and $t = \mathbb{M}g_n(u)$. Consequently, it is sufficient to find, for every $k < n$, relations $\theta_0, \dots, \theta_{n-1} \in \Theta$ and elements $r_0, \dots, r_n \in \mathbb{M}A$ such that

$$\mathbb{M}g_k(u) = r_0 \theta_0^{\mathbb{M}} \dots \theta_{n-1}^{\mathbb{M}} r_n = \mathbb{M}g_{k+1}(u).$$

Thus, fix $k < n$. Since $a_k \hat{\theta} b_k$, we can find elements $c_0, \dots, c_m \in A$ and relations $\theta_0, \dots, \theta_{m-1} \in \Theta$ with

$$a_k = c_0 \theta_0 \dots \theta_{m-1} c_m = b_k.$$

Define $h_l : \sigma \rightarrow A \times A$, for $l < m$ by

$$h_l(\langle a_i, b_i \rangle) := \begin{cases} \langle c_l, c_{l+1} \rangle & \text{if } i = k, \\ \langle a_i, a_i \rangle & \text{if } i > k, \\ \langle b_i, b_i \rangle & \text{if } i < k. \end{cases}$$

Then $v_l := \mathbb{M}h_l(u) \in \mathbb{M}\theta_l$ implies that

$$\begin{aligned} \mathbb{M}g_k(u) &= \mathbb{M}p(v_0) \theta_0 \mathbb{M}q(v_0) = \mathbb{M}p(v_1) \\ \theta_1 \mathbb{M}q(v_1) &= \dots \\ \theta_{m-2} \mathbb{M}q(v_{m-2}) &= \mathbb{M}p(v_{m-1}) \\ \theta_{m-1} \mathbb{M}q(v_{m-1}) &= \mathbb{M}g_{k+1}(u), \end{aligned}$$

as desired. \square

Note that the join of a set of congruences can be computed by first taking their union and then the transitive closure of the resulting relation.

Proposition 3.7. *Let \mathbb{M} be a monad on Pos^{Ξ} where Ξ is finite, let \mathfrak{A} be an \mathbb{M} -algebra, and let $\Xi_0, \Xi_1 \in \text{Cong}_{\mathbb{M}}(\mathfrak{A})$. If Ξ_0 has finitary index, then $\text{TC}(\Xi_0 \cup \Xi_1) \in \text{Cong}_{\mathbb{M}}(\mathfrak{A})$.*

Proof. Let $\Xi := \text{TC}(\Xi_0 \cup \Xi_1)$. Clearly, Ξ is reflexive, transitive, and it contains \leq . Hence, we only need to prove the congruence property, that is, we have to show that

$$s \in \mathbb{M}\Xi \quad \text{implies} \quad \pi(\mathbb{M}p(s)) \subseteq \pi(\mathbb{M}q(s)),$$

where $p, q : A \times A \rightarrow A$ are the two projections.

First, consider the case where $s \in \mathbb{M}\sigma$, for some finite $\sigma \subseteq \Xi$. Then we can use Lemma 3.6 to find a sequence

$$\mathbb{M}p(s) = r_0 \Xi_{i_0}^{\mathbb{M}} \dots \Xi_{i_{n-1}}^{\mathbb{M}} r_n = \mathbb{M}q(s),$$

where $i_0, \dots, i_{n-1} \in \{0, 1\}$ and $r_0, \dots, r_n \in \mathbb{M}A$. As \sqsubseteq_0 and \sqsubseteq_1 are congruence orderings, it follows that

$$\pi(\mathbb{M}p(s)) = \pi(r_0) \sqsubseteq_{i_0} \dots \sqsubseteq_{i_{n-1}} \pi(r_n) = \pi(\mathbb{M}q(s)).$$

Consequently, $\pi(\mathbb{M}p(s)) \sqsubseteq \pi(\mathbb{M}q(s))$.

For the general case, let $s \in \mathbb{M}\sqsubseteq$. Fix a function $\gamma_0 : A/\sqsubseteq_0 \rightarrow A$ (not necessarily monotone) and let $\gamma : A \rightarrow A$ be the composition of the quotient map $A \rightarrow A/\sqsubseteq_0$ with γ_0 . Then

$$\mathbb{M}p(s) \sqsubseteq_0^{\mathbb{M}} \mathbb{M}(\gamma \circ p)(s) \quad \text{and} \quad \mathbb{M}(\gamma \circ q)(s) \sqsubseteq_0^{\mathbb{M}} \mathbb{M}p(s).$$

Furthermore, the set $\sigma := \sqsubseteq \cap (\text{rng } \gamma \times \text{rng } \gamma)$ is finite and $\mathbb{M}(\gamma \times \gamma)(s) \in \mathbb{M}\sigma$. By the special case we have proved above and the fact that \sqsubseteq_0 is a congruence ordering, it therefore follows that

$$\pi(\mathbb{M}p(s)) \sqsubseteq_0 \pi(\mathbb{M}(\gamma \circ p)(s)) \sqsubseteq \pi(\mathbb{M}(\gamma \circ q)(s)) \sqsubseteq_0 \pi(\mathbb{M}q(s)) \quad \square$$

To obtain a stronger statement, we need to make additional assumptions on the monad \mathbb{M} .

Definition 3.8. A functor $\mathbb{M} : \text{Pos}^{\sqsubseteq} \rightarrow \text{Pos}^{\sqsubseteq}$ is *finitary* if it preserves directed colimits, that is, if

$$\mathbb{M}(\text{colim } D) \cong \text{colim}(\mathbb{M} \circ D),$$

for every directed diagram $D : I \rightarrow \text{Pos}^{\sqsubseteq}$.

Remark. (a) In particular, if \mathbb{M} is finitary then $\mathbb{M}A$ is equal to the directed colimit of the diagram consisting of $\mathbb{M}C$, for all finite $C \subseteq A$.

(b) The word functor $\mathbb{M}A := A^+$ is finitary as every finite word uses only finitely many labels. The functor

$$\mathbb{M}\langle A_1, A_\infty \rangle := \langle A_1^+, A_1^+ A_\infty \cup A_1^\omega \rangle$$

for infinite words, on the other hand, is not finitary as an infinite word can contain infinitely many different labels. Thus, in general

$$A^\omega \neq \bigcup \{ C^\omega \mid C \subseteq A \text{ finite} \}.$$

(c) More generally, a polynomial functor in Pos^{\sqsubseteq} is finitary if, and only if, its arity is at most \aleph_0 . (We have proved (\Leftarrow) in Corollary II.2.10.)

For finitary monads we can prove the following more general version of Proposition 3.7. In this case, we can form joins of arbitrary many congruence orderings and we can drop the restriction on the number of sorts.

Proposition 3.9. Let \mathbb{M} be a finitary monad on Pos^{\sqsubseteq} and let Θ be a set of \mathbb{M} -congruences on an \mathbb{M} -algebra \mathfrak{A} . Then $\text{TC}(\bigcup \Theta)$ is also an \mathbb{M} -congruence on \mathfrak{A} .

Proof. Set $\sqsubseteq := \text{TC}(\bigcup \Theta)$. Clearly, \sqsubseteq is a preorder containing \leq . Hence, we only need to prove the congruence property, that is, we have to show that

$$s \in \mathbb{M}\sqsubseteq \quad \text{implies} \quad \pi(\mathbb{M}p(s)) \sqsubseteq \pi(\mathbb{M}q(s)),$$

where $p, q : A \times A \rightarrow A$ are the two projections. Since \mathbb{M} is finitary, there exists a finite subset $\sigma \subseteq \sqsubseteq$ such that $s \in \mathbb{M}\sigma$. Hence, we can use Lemma 3.6 to find a sequence

$$\mathbb{M}p(s) = r_0 \theta_0^{\mathbb{M}} \dots \theta_{n-1}^{\mathbb{M}} r_n = \mathbb{M}q(s),$$

where $\theta_0, \dots, \theta_{n-1} \in \Theta$ and $r_0, \dots, r_n \in \mathbb{M}A$. As the θ_i are congruence orderings, it follows that

$$\pi(\mathbb{M}p(s)) = \pi(r_0) \theta_0 \dots \theta_{n-1} \pi(r_n) = \pi(\mathbb{M}q(s)).$$

This implies that $\pi(\mathbb{M}p(s)) \sqsubseteq \pi(\mathbb{M}q(s))$. \square

After these preparations, we can present the following two conditions on the existence of minimal algebras.

Theorem 3.10. Let \mathbb{M} be a monad on Pos^{\sqsubseteq} where \sqsubseteq is finite. Every language $\kappa \in \mathbb{M}_{\text{f}}A \rightarrow \Omega$ that is recognised by a surjective morphism $\varphi : \mathbb{M}A \rightarrow \mathfrak{B}$ into a finitary \mathbb{M} -algebra has a minimal morphism.

Proof. By Lemma 3.3, it is sufficient to prove that the set

$$\Theta := \{ \theta \in \text{Cong}_{\mathbb{M}}(\mathcal{A}) \mid \theta \subseteq \ker \kappa \}$$

has a greatest element. First, note that Θ contains the congruence $\sigma := \ker \varphi$, which has finitary index. Furthermore, the set

$$\Theta_o := \{ \theta \in \Theta \mid \sigma \subseteq \theta \}.$$

is finite since every congruence in Θ_o can be obtained from σ by merging some of the classes and making some of the other classes comparable. Therefore, we can use Proposition 3.7 to find a greatest element $\rho \in \Theta_o$ (the transitive closure of the union of all congruences in Θ_o).

We claim that ρ is also the greatest element of Θ . Hence, let $\theta \in \Theta$. Then Proposition 3.7 implies that $\hat{\theta} := \text{TC}(\theta \cup \sigma) \in \Theta_o$. Consequently, $\theta \subseteq \hat{\theta} \subseteq \rho$. \square

Note that there are examples of (i) non-recognisable languages and of (ii) recognisable languages over infinitely many sorts that do not have a minimal algebra. For finitary monads on the other-hand, minimal algebras always exist.

Theorem 3.11. *Let \mathbb{M} be a monad on Pos^{Ξ} . The following statements are equivalent.*

- (1) *For every \mathbb{M} -algebra \mathcal{A} , every set Ω , and every $\xi \in \Xi$, every Ω -language $\kappa : A_{\xi} \rightarrow \Omega$ has a minimal morphism.*
- (2) *\mathbb{M} is finitary.*
- (3) *For every \mathbb{M} -algebra \mathcal{A} , the set $\text{Cong}_{\mathbb{M}}(\mathcal{A})$ of all \mathbb{M} -congruences is closed, as a subset of $\text{Cong}(A)$, under arbitrary joins.*

Proof. (2) \Rightarrow (3) follows by Proposition 3.9.

(3) \Rightarrow (1) Let θ be the supremum (in $\text{Cong}(A)$) of all \mathbb{M} -congruences below $\ker \kappa$. By assumption, θ is an \mathbb{M} -congruence. Consequently, the claim follows by Lemma 3.3.

(1) \Rightarrow (2) Fix $A \in \text{Set}$ and let A' be a disjoint copy of A . For every finite $C \subseteq A$ with copy $C' \subseteq A'$, we consider the function $\sigma_C : A + A' \rightarrow A + A'$ that acts as the identity on $A + (A' \setminus C)$ and that maps elements $c' \in C'$ to their copy $c \in C$. For $s, t \in \mathbb{M}(A + A')$, define

$$s \sqsubseteq t \quad : \text{iff} \quad \mathbb{M}\sigma_C(s) \leq \mathbb{M}\sigma_C(t), \quad \text{for some finite } C \subseteq A.$$

We start by showing that \sqsubseteq is a congruence order on $\mathbb{M}(A + A')$. Reflexivity is obvious. For transitivity, note that

$$\mathbb{M}\sigma_C(r) \leq \mathbb{M}\sigma_C(s) \quad \text{and} \quad \mathbb{M}\sigma_D(s) \leq \mathbb{M}\sigma_D(t)$$

implies that $\mathbb{M}\sigma_{C \cup D}(r) \leq \mathbb{M}\sigma_{C \cup D}(t)$. Hence, \sqsubseteq is a preorder. Furthermore, monotonicity of σ_C implies that $\leq \subseteq \sqsubseteq$. Hence, it remains to prove the congruence property. Let $q : \mathbb{M}(A + A') \rightarrow \mathbb{M}(A + A')/\sqsubseteq$ be the quotient map. We can regard q as an Ω -language for $\Omega := \mathbb{M}(A + A')/\sqsubseteq$. By assumption, it therefore follows that q has a minimal morphism $\eta : \mathbb{M}(A + A') \rightarrow \mathfrak{M}$. As $\ker \eta$ is a congruence relation, it is sufficient to prove that $\sqsubseteq = \ker \eta$.

(\supseteq) As η recognises q , there exists some function $g : M \rightarrow \mathbb{M}(A + A')/\sqsubseteq$ such that $q = g \circ \eta$. This implies that $\ker \eta \subseteq \ker q = \sqsubseteq$.

(\subseteq) Suppose that $s \sqsubseteq t$, that is, $\mathbb{M}\sigma_C(s) \leq \mathbb{M}\sigma_C(t)$, for some finite $C \subseteq A$. Since $\ker \mathbb{M}\sigma_C \subseteq \sqsubseteq = \ker q$, we know that $\mathbb{M}\sigma_C$ recognises q . As $\mathbb{M}\sigma_C$ is a morphism of \mathbb{M} -algebras, it follows by the definition of a minimal morphism that there exists some morphism $\rho : \text{rng } \mathbb{M}\sigma_C \rightarrow \mathfrak{M}$ with $\rho \circ \mathbb{M}\sigma_C = \eta$. Consequently,

$$\eta(s) = \rho(\mathbb{M}\sigma_C(s)) \leq \rho(\mathbb{M}\sigma_C(t)) = \eta(t),$$

that is, $\langle s, t \rangle \in \ker \eta$, as desired.

To conclude the proof, let \equiv be the equivalence relation associated with \sqsubseteq and let $\iota : A \rightarrow A'$ be the functions mapping $a \in A$ to $a' \in A'$. For $a \in A$, it follows that

$$\mathbb{M}\sigma_{\{a\}}(\text{sing}(a)) = \text{sing}(a) = \mathbb{M}\sigma_{\{a\}}(\mathbb{M}\iota(\text{sing}(a))).$$

Consequently, $\text{sing}(a) \equiv \mathbb{M}l'(\text{sing}(a))$. Since \sqsubseteq is a congruence ordering and $\mathbb{M}A$ is generated by rng sing , it follows that

$$s \equiv \mathbb{M}l(s), \quad \text{for all } s \in \mathbb{M}A.$$

By definition of \sqsubseteq it follows that, for every $s \in \mathbb{M}A$, there exists some finite set $C_s \subseteq A$ such that

$$\mathbb{M}\sigma_{C_s}(s) = \mathbb{M}\sigma_{C_s}(\mathbb{M}l(s)).$$

To see that \mathbb{M} is finitary, let $s \in \mathbb{M}A$ and fix a finite set $C \subseteq A$ with

$$\mathbb{M}\sigma_C(s) = \mathbb{M}(\sigma_C \circ \iota)(s).$$

Fix $c \in C$ and let $\tau : A + A' \rightarrow A$ be the function acting as the identity on A while mapping every element of A' to c . Then

$$s = \mathbb{M}\tau(s) = \mathbb{M}(\tau \circ \sigma_C)(s) = \mathbb{M}(\tau \circ \sigma_C \circ \iota)(s).$$

The claim now follows since $\text{rng}(\tau \circ \sigma_C \circ \iota) = C$ is finite. \square

Example. The word monad $\mathbb{M}X = X^+$ has arity \aleph_0 and is therefore finitary. Consequently, every Ω -language $\kappa : \Sigma^+ \rightarrow \Omega$ has a minimal algebra. \lrcorner

This result seems to be bad news for languages over infinite objects like ω -words or infinite trees. Fortunately, we do not need minimal algebras for *all* languages. We are only interested in languages of the form $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ where Σ is a finite alphabet. Furthermore, there are usually only a few choices for Ω that are of interest.

4. Syntactic Algebras

The definition of a minimal algebra is rather abstract. In many cases, we can give a more concrete description, which is the subject of this section. To motivate the general constructions that follow, let us take a look at the

case of finite words. With a language $K \subseteq \Sigma^*$, we can associate its *syntactic congruence*

$$u \approx_K v \quad : \text{iff} \quad xuy \in K \Leftrightarrow xvy \in K, \quad \text{for all } x, y \in \Sigma^*.$$

This relation is a congruence for monoid multiplication, which means we can define a monoid

$$\text{Syn}(K) := \Sigma^* / \approx_K,$$

called the *syntactic monoid* of K . It turns out that this monoid is finite if, and only if, K is a regular language. Furthermore, the quotient map $\Sigma^* \rightarrow \Sigma^* / \approx_K$ is a monoid homomorphism recognising K .

In the general case the construction is analogous. We start by generalising operations of the form $u \mapsto xuy$.

Definition 4.1. Let \mathfrak{A} be an \mathbb{M} -algebra.

(a) A *context* is an element of $\mathbb{M}(A + \square_\zeta)$, where \square_ζ is considered as some special symbol of sort $\zeta \in \Xi$ called a *hole*. For a context $p \in \mathbb{M}_\xi(A + \square_\zeta)$ and an element $a \in A_\zeta$, we define

$$p[a] := \sigma_a(p) \in A_\xi$$

where $\sigma_a : \mathbb{M}(A + \square_\zeta) \rightarrow \mathfrak{A}$ is the unique morphism that extends the function $s_a : A + \square_\zeta \rightarrow A$ given by

$$s_a(\square_\zeta) := a \quad \text{and} \quad s_a(c) := c, \quad \text{for } c \in A.$$

In the case where $\mathfrak{A} = \mathbb{M}\Sigma$ is a free \mathbb{M} -algebra, we will also consider elements $p \in \mathbb{M}(\Sigma + \square_\zeta)$ as contexts, by identifying them with their image under $\mathbb{M}(\text{sing} + \text{id})$.

(b) An *internal operation* of \mathfrak{A} is a function $A_\zeta \rightarrow A_\xi$ of the form

$$a \mapsto \pi(p[a]), \quad \text{for some context } p \in \mathbb{M}_\xi(A + \square_\zeta).$$

In the following, we will often identify a context with the associated internal operation.

(c) The composition of two contexts $p \in \mathbb{M}_\xi(A + \square_\zeta)$ and $q \in \mathbb{M}_\zeta(A + \square_\eta)$ is the context obtained from p by replacing every occurrence of \square_ζ by a copy of q . Formally,

$$pq := \hat{p}[q] \in \mathbb{M}_\xi(A + \square_\eta),$$

where $\hat{p} := \mathbb{M}(\text{sing} + \text{id})(p)$ and the expression $\hat{p}[q]$ is evaluated in the \mathbb{M} -algebra $\mathbb{M}(A + \square_\eta)$.

(d) A derivative of $\kappa : A_\xi \rightarrow \Omega$ is a function of the form

$$p^{-1}[\kappa] := \kappa \circ p : A_\zeta \rightarrow \Omega, \quad \text{where } p \in \mathbb{M}_\xi(A + \square_\zeta) \text{ is a context.}$$

We can now lift the language functor \mathbb{L}° from \mathcal{D} to $\text{Alg}(\mathbb{M})$ as follows.

Definition 4.2. (a) For an \mathbb{M} -algebra \mathfrak{A} , we denote by $\mathbb{L}\mathfrak{A}$ the language algebra with domains $(\mathbb{L}^\circ A_\xi)_{\xi \in \Xi} \in \text{Set}^\Xi$ and the following operations. We retain all operations of each algebra $\mathbb{L}^\circ A_\xi$ (that is, all Ω -operations on $\mathcal{D}(A_\xi, \Omega)$) and we add all internal operations $\mathbb{L}^\circ p : \mathbb{L}^\circ A_\xi \rightarrow \mathbb{L}^\circ A_\zeta$, for contexts $p \in \mathbb{M}_\xi(A + \square_\zeta)$, $\xi, \zeta \in \Xi$.

(b) For a morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ of \mathbb{M} -algebras, we set

$$\mathbb{L}\varphi(\kappa) := \mathbb{L}^\circ \varphi_\xi(\kappa), \quad \text{for } \kappa \in \mathbb{L}^\circ A_\xi.$$

(c) A morphism of language algebras is a function $\varphi : \mathbb{L}\mathfrak{A} \rightarrow \mathbb{L}\mathfrak{B}$ such that, for every internal operation p of $\mathbb{L}\mathfrak{B}$, there is some operation q of $\mathbb{L}\mathfrak{A}$ with

$$q \circ \varphi = \varphi \circ p.$$

(d) For a subset $C \subseteq \mathbb{L}\mathfrak{A}$, we denote by $\langle\langle C \rangle\rangle_{\mathbb{L}}$ the subalgebra of $\mathbb{L}\mathfrak{A}$ generated by C .

Example. For unweighted languages in Pos , the language algebra $\mathbb{L}\mathfrak{A}$ consists of all upwards-closed subsets of A . The operations are generated by (i) intersection \cup ; (ii) union \cap ; (iii) the constants \emptyset and A ; and (iv) derivatives $a^{-1}K$ and Ka^{-1} , for $a \in A$.

Lemma 4.3. Let $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ be a morphism of \mathbb{M} -algebras and $p \in \mathbb{M}_\xi(A + \square_\zeta)$ a context.

(a) The image $q := \mathbb{M}(\varphi + \text{id})(p) \in \mathbb{M}_\xi(B + \square_\zeta)$ is a context satisfying

$$\varphi_\xi \circ \hat{p} = \hat{q} \circ \varphi_\zeta,$$

where \hat{p} and \hat{q} are the associated internal operations.

(b) If φ is an E -morphism and \mathbb{M} preserves E -morphisms, the induced map $\mathbb{M}(\varphi + \text{id})$ on contexts is also an E -morphism.

Proof. (a) Note that

$$\begin{aligned} (\hat{q} \circ \varphi)(s) &= \hat{q}(\varphi(s)) \\ &= \pi(\mathbb{M}(\varphi + \text{id})(p)[\varphi(s)]) \\ &= \pi(\mathbb{M}\varphi(p[s])) \\ &= \varphi(\text{flat}(p[s])) \\ &= (\varphi \circ \hat{p})(s). \end{aligned}$$

(b) follows by Lemma II.1.9 (a) and our assumption on \mathbb{M} . \square

Lemma 4.4. In an algebra $\mathbb{L}\mathfrak{A}$, Ω -operations commute with internal operations in the sense that

$$p^{-1}[\omega[\bar{\kappa}]] = \omega[p^{-1}[\kappa_0], \dots, p^{-1}[\kappa_{n-1}]],$$

for every Ω -operation $\omega : \Omega^n \rightarrow \Omega$, every context p , and all $\kappa_0, \dots, \kappa_{n-1} \in \mathbb{L}A$.

Proof. By definition,

$$\begin{aligned} p^{-1}[\omega[\bar{\kappa}]] &= \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle \circ p \\ &= \omega \circ \langle \kappa_0 \circ p, \dots, \kappa_{n-1} \circ p \rangle \\ &= \omega[p^{-1}[\kappa_0], \dots, p^{-1}[\kappa_{n-1}]]. \end{aligned}$$

\square

Definition 4.5. Let \mathfrak{A} be an \mathbb{M} -algebra and $\kappa : A_\xi \rightarrow \Omega$.

(a) The syntactic morphism of κ is

$$\text{syn}_\kappa := \ker \rho,$$

where $\rho = (\rho_\zeta)_{\zeta \in \Xi}$ is the morphism with components

$$\rho_\zeta := \langle \kappa \circ p \rangle_{p \in \mathbb{M}_\zeta(A + \square)} : A_\zeta \rightarrow \prod_p \Omega.$$

(b) We call $\ker \text{syn}_\kappa$ the *syntactic congruence* of κ .

(c) The *syntactic residue* is the unique function $\text{res}_\kappa : \text{Syn}_\xi(\kappa) \rightarrow \Omega$ satisfying $\kappa = \text{res}_\kappa \circ \text{syn}_{\kappa, \xi}$.

(d) We say that κ has a *syntactic algebra* if the syntactic congruence is an \mathbb{M} -congruence. In this case the quotient codom syn_κ has the structure of an \mathbb{M} -algebra. This algebra is called the *syntactic algebra* of κ . We denote it by $\text{Syn}(\kappa)$. \lrcorner

Remark. (a) Using the general definition of the kernel of a morphism, we have $\ker \text{syn}_\kappa = \text{syn}_\kappa$. Hence, having a separate notion of a syntactic congruence is mainly useful in categories \mathcal{D} where we have a more concrete notion of a congruence.

(b) If $\text{Syn}(\kappa)$ exists, the syntactic morphism $\text{syn}_\kappa : \mathfrak{A} \rightarrow \text{Syn}(\kappa)$ is a morphism of \mathbb{M} -algebras.

(c) Uniqueness of res_κ follows from the fact that syn_κ is an epimorphism. \lrcorner

In the category Pos the definitions simplify as follows.

Lemma 4.6. *Suppose that $\mathcal{D} = \text{Pos}$ and let \leq_κ be the syntactic congruence of $\kappa : A_\xi \rightarrow \Omega$. Then*

$$\begin{aligned} a \leq_\kappa b & \quad \text{iff} \quad \kappa(p[a]) \leq \kappa(p[b]), \quad \text{for all } p \in \mathbb{M}_\xi(A + \square_\zeta), \\ & \quad \text{iff} \quad \mu(a) \leq \mu(b), \quad \text{for all } \mu \in \langle \kappa \rangle_{\mathbb{L}}. \end{aligned}$$

Proof. The first equivalence follows immediately by unravelling the definitions. Since $p^{-1}[\kappa] \in \langle \kappa \rangle_{\mathbb{L}}$, it therefore remains to prove that

$$\kappa(p[a]) \leq \kappa(p[b]), \quad \text{for all } p \in \mathbb{M}_\xi(A + \square_\zeta),$$

implies that

$$\mu(a) \leq \mu(b), \quad \text{for all } \mu \in \langle \kappa \rangle_{\mathbb{L}}.$$

Hence, fix $\mu \in \langle \kappa \rangle_{\mathbb{L}}$. Then $\mu = s(\kappa)$, where s is a composition of Ω -operations and internal operations. By Lemma 4.4, we may assume that $s = \omega \circ p^{-1}$ for an Ω -operation $\omega : \Omega \rightarrow \Omega$ and some context p . Since ω is monotone,

$$\kappa(p[a]) \leq \kappa(p[b]) \quad \text{implies} \quad \omega[\kappa](p[a]) \leq \omega[\kappa](p[b]). \quad \square$$

Remark. In the unweighted case for $\mathcal{D} = \text{Pos}$, we obtain for $K \subseteq A$


$$a \leq_K b \quad \text{iff} \quad p[a] \in K \Rightarrow p[b] \in K, \quad \text{for all contexts } p.$$

For $\mathcal{D} = \text{Set}$, we obtain the familiar definition

$$a \approx_K b \quad \text{iff} \quad p[a] \in K \Leftrightarrow p[b] \in K, \quad \text{for all contexts } p. \quad \lrcorner$$

Examples. Let $\mathbb{M}\Sigma := \Sigma^*$ be the word monad on Pos .

(a) We consider the language $K := a^*b^* \subseteq \mathbb{M}\{a, b\}$. Its syntactic congruence \leq_K has the following 5 classes.



(b) Let us consider the \mathbb{N} -language $\kappa : \Sigma^* \rightarrow \mathbb{N}$ defined by

$$\kappa(w) := |w|_a.$$

That is, κ maps a word $w \in \{a, b\}^*$ to the number of letters a it contains. Then \leq_κ is a linear order with

$$w \leq_\kappa w' \quad \text{iff} \quad |w|_a \leq |w'|_a. \quad \lrcorner$$

Example. Syntactic congruences are also used in the theory of programming languages. Let ΛX denote the set of all λ -expressions with basic operations

from the set X . Then Λ forms a monad on Set . Two expressions $s, t \in \Lambda X$ are called *observational equivalent* if

$$us \text{ terminates} \quad \text{iff} \quad ut \text{ terminates}, \quad \text{for every } u \in \Lambda X.$$

This is just the syntactic congruence associated with the language

$$K := \{s \in \Lambda X \mid s \text{ terminates}\}.$$

A *model* is a morphism $\varphi : \Lambda X \rightarrow \mathfrak{A}$ of Λ -algebras. Such a model is *fully abstract* if the kernel of φ coincides with observational equivalence, that is, if φ factorises through the syntactic morphism as $\varphi = i \circ \text{syn}_\kappa$, for some embedding $i : \text{Syn}(K) \rightarrow \mathfrak{A}$.

Let us collect a few basic properties of the syntactic congruence. We start by showing that our definitions make sense.

Lemma 4.7. *Let $\kappa : A_\xi \rightarrow \Omega$.*

- (a) $\text{syn}_\kappa : A \rightarrow \text{Syn}(\kappa)$ recognises κ .
- (b) $\text{res}_\kappa : \text{Syn}_\xi(\kappa) \rightarrow \Omega$ exists.
- (c) $\ker \text{syn}_\kappa \leq \ker \kappa$.
- (d) $\ker \text{syn}_\kappa = \inf_p \ker(\kappa \circ p)$

Proof. (d) Let $\rho_\xi := \langle \kappa \circ p \rangle_p : A_\xi \rightarrow \prod_p \Omega$ be the morphisms used to define syn_κ . Then

$$\ker \text{syn}_\kappa = \ker \rho = \ker \langle \kappa \circ p \rangle_p = \inf_p \ker(\kappa \circ p).$$

This proves (d).

(c) We have $\kappa \circ p_o = \kappa$, for the trivial context $p_o := \text{sing}(\square)$. Hence, (d) implies (c).

(a), (b) By (c) we can find some function μ with $\ker \kappa = \mu \circ \ker \text{syn}_\kappa$. This implies that

$$\kappa = \text{im } \kappa \circ \ker \kappa = \text{im } \kappa \circ \mu \circ \text{syn}_\kappa.$$

Hence, syn_κ recognises κ and $\text{res}_\kappa := \text{im } \kappa \circ \mu$ exists. \square

The following lemma contains the key property of the syntactic congruence.

Lemma 4.8. *Let $\kappa : A_\xi \rightarrow \Omega$ be a language such that syn_κ is an \mathbb{M} -congruence, and let $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ be an E -morphism of \mathbb{M} -algebras. The following statements are equivalent.*

- (1) φ recognises κ .
- (2) $\ker \varphi \leq \ker \text{syn}_\kappa$
- (3) $\ker \varphi \leq \ker \kappa$

Proof. (1) \Leftrightarrow (3) has already been proved in Lemma 3.1 (a).

(1) \Rightarrow (2) If $\kappa = \mu \circ \varphi$, for some $\mu : B_\xi \rightarrow \Omega$, then

$$\varphi = \ker \varphi \leq \ker(\mu \circ \varphi) = \ker \kappa = \ker \text{syn}_\kappa.$$

(2) \Rightarrow (1) Since $\varphi = \ker \varphi \leq \ker \text{syn}_\kappa = \text{syn}_\kappa$, there exists a function $\rho : B \rightarrow \text{Syn}(\kappa)$ such that $\text{syn}_\kappa = \rho \circ \varphi$. Hence,

$$\kappa = \text{res}_\kappa \circ \text{syn}_\kappa = \text{res}_\kappa \circ \rho \circ \varphi.$$

which implies that φ recognises κ . \square

A noteworthy consequence of this lemma is the fact that, if it exists, the syntactic algebra of a language κ is equal to its minimal algebra.

Theorem 4.9. *Let \mathfrak{A} be an \mathbb{M} -algebra and $\kappa : A_\xi \rightarrow \Omega$ be a language that has a syntactic algebra. For every E -morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ recognising κ , there exists a unique morphism $\rho : \mathfrak{B} \rightarrow \text{Syn}(\kappa)$ such that $\text{syn}_\kappa = \rho \circ \varphi$.*

Proof. Suppose that φ recognises κ . By Lemma 4.8 we have $\ker \varphi \leq \ker \text{syn}_\kappa$. Therefore, there exists some function $\rho : B \rightarrow \text{Syn}(\kappa)$ with $\text{syn}_\kappa = \rho \circ \varphi$. It is unique as φ is an epimorphism. Furthermore, ρ is a morphism of \mathbb{M} -algebras by Lemma I.5.6. \square

Corollary 4.10. *Let \mathfrak{A} be an \mathbb{M} -algebra and $\kappa : A_\xi \rightarrow \Omega$ a language with a syntactic algebra. Then $\text{Syn}(\kappa)$ is equal to the minimal algebra of κ and $\text{syn}_\kappa : \mathfrak{A} \rightarrow \text{Syn}(\kappa)$ is equal to its minimal morphism.*

Consequently, the syntactic congruence provides a concrete way to compute the minimal algebra of a language. Note that there do exist languages that have a minimal algebra, but not a syntactic one. The preceding corollary justifies the following notation.

Definition 4.11. Let $\kappa : A_\xi \rightarrow \Omega$ be a language with a minimal algebra. Even if κ does not have a syntactic algebra, we denote the minimal algebra, the minimal morphism, and its residue by, respectively, $\text{Syn}(\kappa)$, syn_κ , and res_κ .

We are particularly interested in languages whose syntactic algebra is finite-dimensional.

Lemma 4.12. *If $\text{Syn}(\kappa)$ is sort-wise noetherian, it is sort-wise strongly finite-dimensional.*

Proof. Let $\rho_\xi : \mathbb{M}_\xi A \rightarrow \Omega^X$ be the morphism used to define $\text{Syn}(\kappa)$. Then $\text{im } \rho_\xi : \text{Syn}_\xi(\kappa) \rightarrow \Omega^X$ is an M -morphism. Since $\text{Syn}_\xi(\kappa)$ is noetherian, it follows by Lemma 1.8 that it is strongly finite-dimensional. \square

Languages Recognised by Syntactic Algebras

Let us take a look at what other languages are recognised by a syntactic algebra, besides the language the algebra was originally constructed from. We start with a technical lemma about closure properties of the syntactic congruence.

Lemma 4.13. *Let \mathfrak{A} be an \mathbb{M} -algebra, $\kappa : A_\xi \rightarrow \Omega$ a function, and $p \in \mathbb{M}_\xi(A + \square_\xi)$ a context.*

- (a) $\ker \text{syn}_{\kappa, \zeta} \leq \ker (\text{syn}_{\kappa, \xi} \circ p)$.
- (b) $\ker \text{syn}_\kappa \leq \ker \text{syn}_{p^{-1}[\kappa]}$.

Proof. Let $\rho_\eta := \langle \kappa \circ q \rangle_q : A_\eta \rightarrow \prod_q \Omega$ be the morphisms used to define syn_κ , and let $u_r : \prod_q \Omega \rightarrow \Omega$ be the projection to the r -th component.

(a) Note that $\rho = \mu \circ \text{syn}_\kappa$, for some $\mu \in M$. Consequently, we have

$$\ker \text{syn}_\kappa = \ker \rho \quad \text{and} \quad \ker (\text{syn}_\kappa \circ p) = \ker (\rho \circ p).$$

Setting $g := \langle u_{pq} \rangle_q$, it follows that

$$\rho_\xi \circ p = \langle \kappa \circ q \rangle_q \circ p = \langle \kappa \circ q \circ p \rangle_q = g \circ \langle \kappa \circ q \rangle_q = g \circ \rho_\zeta.$$

This implies that

$$\ker \text{syn}_{\kappa, \zeta} = \ker \rho_\zeta \leq \ker (\rho_\xi \circ p) = \ker (\text{syn}_{\kappa, \xi} \circ p).$$

(b) Note that the morphism $\text{syn}_{p^{-1}[\kappa]}$ is defined via

$$\sigma_\eta := \langle \kappa \circ p \circ q \rangle_q : A_\eta \rightarrow \prod_q \Omega.$$

Setting $g := \langle u_{pq} \rangle_q$, it follows that $\sigma_\eta = g \circ \rho_\eta$. Consequently, we have

$$\ker \text{syn}_\kappa = \ker \rho \leq \ker (g \circ \rho) = \ker \sigma = \ker \text{syn}_{p^{-1}[\kappa]}.$$

\square

The language algebras of a syntactic algebra are particularly simple: they are generated by a single element.

Proposition 4.14. *Let \mathfrak{A} be an \mathbb{M} -algebra and $\kappa : A_\xi \rightarrow \Omega$ a language with a sort-wise noetherian syntactic algebra. Then*

$$\mathbb{L}\text{Syn}(\kappa) = \langle\langle \text{res}_\kappa \rangle\rangle_{\mathbb{L}}.$$

Proof. Let $\mu \in \mathbb{L}_\zeta \text{Syn}(\kappa)$ and set $\lambda := \mu \circ \text{syn}_\kappa$. As $\text{Syn}_\zeta(\kappa)$ is noetherian, there exists a finite set $H \subseteq \mathbb{M}_\zeta(A + \square_\xi)$ of contexts such that

$$\ker (q_H \circ \text{im } \rho)|_{\text{Syn}_\zeta(\kappa)} = \ker (\text{im } \rho)|_{\text{Syn}_\zeta(\kappa)},$$

where $\rho_\xi := \langle \kappa \circ p \rangle_p : A_\xi \rightarrow \prod_p \Omega$ are the morphisms used to define syn_κ and $q_H : \prod_p \Omega \rightarrow \prod_{p \in H} \Omega$ is the projection to the components in H . It follows that

$$\ker (q_H \circ \rho)|_{A_\zeta} = \ker \rho|_{A_\zeta} = \text{syn}_\kappa|_{A_\zeta} \leq (\mu \circ \text{syn}_\kappa)|_{A_\zeta} = \ker \lambda.$$

As Ω is injective, there therefore exists a function $\omega : \Omega^H \rightarrow \Omega$ such that

$$\omega \circ (q_H \circ \rho)|_{A_\zeta} = \lambda.$$

By definition of ρ , we have $(q_H \circ \rho)|_{A_\zeta} = \langle \kappa \circ p \rangle_{p \in H}$. Hence,

$$\begin{aligned} \mu \circ \text{syn}_\kappa &= \lambda = \omega \circ \langle \kappa \circ p \rangle_{p \in H} \\ &= \omega \circ \langle \text{res}_\kappa \circ \text{syn}_\kappa \circ p \rangle_{p \in H} \\ &= \omega \circ \langle \text{res}_\kappa \circ \mathbb{M}(\text{syn}_\kappa + \text{id})(p) \circ \text{syn}_\kappa \rangle_{p \in H} \\ &= \omega \circ \langle \text{res}_\kappa \circ p \rangle_{p \in \mathbb{M}(\text{syn}_\kappa + \text{id})[H]} \circ \text{syn}_\kappa, \end{aligned}$$

where the fourth step follows by Lemma 4.3. As syn_κ is an epimorphism, it follows that

$$\mu = \omega[\langle p^{-1}[\text{res}_\kappa] \rangle_{p \in \mathbb{M}(\text{syn}_\kappa + \text{id})[H]}] \in \langle\langle \text{res}_\kappa \rangle\rangle_{\mathbb{L}}. \quad \square$$

We obtain the following preliminary characterisation of which languages are recognised by $\text{Syn}(\kappa)$. A more general statement where κ and λ are allowed to have different domains will be derived in Proposition 4.16 below.

Proposition 4.15. *Let \mathfrak{A} be an \mathbb{M} -algebra and let $\kappa : A_\xi \rightarrow \Phi$ and $\lambda : A_\zeta \rightarrow \Omega$ be languages such that κ has a syntactic algebra. The following statements are equivalent.*

- (1) $\ker \text{syn}_\kappa \leq \ker \text{syn}_\lambda$
- (2) $\text{syn}_\kappa : \mathfrak{A} \rightarrow \text{Syn}(\kappa)$ recognises λ .
- (3) Every morphism of \mathbb{M} -algebras recognising κ also recognises λ .

If $\text{Syn}(\kappa)$ is sort-wise noetherian, the following statement is also equivalent to those ones.

- (4) $\lambda = \omega[p_o^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]]$, for some operation $\omega : \Phi^n \rightarrow \Omega$ and finitely many contexts p_o, \dots, p_{n-1} .

If in addition $\Phi = \Omega$, the following statement is also equivalent to (1)–(4).

- (5) $\lambda \in \langle\langle \kappa \rangle\rangle_{\mathbb{L}}$

Proof. (1) \Leftrightarrow (2) follows directly by Lemma 4.8.

(3) \Rightarrow (2) is trivial as we have seen in Lemma 4.7 that syn_κ recognises κ .

(1) \Rightarrow (3) Suppose that $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ recognises κ . By Lemma 4.8, it follows that $\ker \varphi \leq \ker \text{syn}_\kappa \leq \ker \text{syn}_\lambda$, which, by the same lemma, implies that φ recognises λ .

(2) \Rightarrow (5) Suppose that $\lambda = \mu \circ \text{syn}_\kappa$ for some $\mu : \text{Syn}_\zeta(\kappa) \rightarrow \Omega$. By Proposition 4.14, we have $\mu \in \langle\langle \text{res}_\kappa \rangle\rangle_{\mathbb{L}}$. Hence, there is some operation ω of $\mathbb{L}\text{Syn}(\kappa)$ with $\mu = \omega \circ \text{res}_\kappa$. It follows that

$$\lambda = \mu \circ \text{syn}_\kappa = \omega \circ \text{res}_\kappa \circ \text{syn}_\kappa = \omega \circ \kappa.$$

This implies that $\lambda \in \langle\langle \kappa \rangle\rangle_{\mathbb{L}}$.

(5) \Rightarrow (4) Suppose that $\lambda \in \langle\langle \kappa \rangle\rangle_{\mathbb{L}}$. Then $\lambda = s[\kappa]$, for some term s . By Lemma 4.4, it follows that we can replace s by a term consisting of derivatives followed by Ω -operations. As derivatives and Ω -operations are both closed under composition, it follows that

$$\lambda = \omega[p_o^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]],$$

for some Ω -operation ω and contexts p_o, \dots, p_{n-1} .

(4) \Rightarrow (2) We have

$$\begin{aligned} \lambda &= \omega[p_o^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]] \\ &= \omega \circ \langle \kappa \circ p_o, \dots, \kappa \circ p_{n-1} \rangle \\ &= \omega \circ \langle \text{res}_\kappa \circ \text{syn}_\kappa \circ p_o, \dots, \text{res}_\kappa \circ \text{syn}_\kappa \circ p_{n-1} \rangle \\ &= \omega \circ \langle \text{res}_\kappa \circ \mathbb{M}(\text{syn}_\kappa + \text{id})(p_o) \circ \text{syn}_\kappa, \dots, \\ &\quad \text{res}_\kappa \circ \mathbb{M}(\text{syn}_\kappa + \text{id})(p_{n-1}) \circ \text{syn}_\kappa \rangle \\ &= \omega \circ \langle \text{res}_\kappa \circ \mathbb{M}(\text{syn}_\kappa + \text{id})(p_o), \dots, \\ &\quad \text{res}_\kappa \circ \mathbb{M}(\text{syn}_\kappa + \text{id})(p_{n-1}) \rangle \\ &\quad \circ \text{syn}_\kappa, \end{aligned}$$

where the fourth step follows by Lemma 4.3. Hence, $\lambda = \mu \circ \text{syn}_\kappa$, for some μ , and syn_κ recognises λ . \square

Remark. In the unweighted case, the condition in (4) reads

$$L = \bigcup_{i < m} \bigcap_{k < n_i} p_{ik}^{-1}[K].$$

To see this, note that we can express $\omega : \{0, 1\}^n \rightarrow \{0, 1\}$ as a finite boolean combination in disjunctive normal form. As ω is monotone, we can omit every negated term from the resulting expression without changing the result. \lrcorner

The next proposition describes languages recognised by syntactic algebras via arbitrary morphisms.

Proposition 4.16. *Let \mathfrak{A} and \mathfrak{B} be \mathbb{M} -algebras and let $\kappa : A_\xi \rightarrow \Phi$ be a Φ -language with a sort-wise noetherian syntactic algebra. An Ω -language $\lambda : B_\zeta \rightarrow \Omega$ is recognised by $\text{Syn}(\kappa)$ if, and only if, it is of the form*

$$\lambda = \omega[\varphi^{-1}[p_0^{-1}[\kappa]], \dots, \varphi^{-1}[p_{n-1}^{-1}[\kappa]]],$$

for some a morphism $\varphi : \mathfrak{B} \rightarrow \mathfrak{A}$, contexts $p_0, \dots, p_{n-1} \in \mathbb{M}(A + \square)$, and an operation $\omega : \Phi^n \rightarrow \Omega$.

Proof. Note that we can write

$$\omega[\varphi^{-1}[p_0^{-1}[\kappa]], \dots, \varphi^{-1}[p_{n-1}^{-1}[\kappa]]] = \omega \circ \langle p_i^{-1}[\kappa] \rangle_{i < n} \circ \varphi.$$

(\Leftarrow) By Proposition 4.15, the morphism $\text{syn}_\kappa : \mathfrak{A} \rightarrow \text{Syn}(\kappa)$ recognises the language

$$\lambda' := \omega \circ \langle p_i^{-1}[\kappa] \rangle_{i < n}.$$

Consequently, $\text{syn}_\kappa \circ \varphi : \mathbb{M}\Gamma \rightarrow \text{Syn}(\kappa)$ recognises $\lambda' \circ \varphi = \lambda$.

(\Rightarrow) Fix a morphism $\psi : \mathfrak{B} \rightarrow \text{Syn}(\kappa)$ and a language $\mu : \text{Syn}_\zeta(\kappa) \rightarrow \Omega$ such that $\lambda = \psi^{-1}[\mu]$. By Corollary II.5.4, there exists a morphism $\varphi : \mathfrak{B} \rightarrow \mathfrak{A}$ such that $\text{syn}_\kappa \circ \varphi = \psi$. The language $\lambda' := \text{syn}_\kappa^{-1}[\mu]$ is recognised by syn_κ . Consequently, we can use Proposition 4.15 to find contexts p_0, \dots, p_{n-1} and a Φ -operation $\omega : \Phi^n \rightarrow \Phi$ such that

$$\lambda' = \omega[p_0^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]].$$

It follows that

$$\lambda = \mu \circ \psi = \mu \circ \text{syn}_\kappa \circ \varphi = \omega \circ \langle p_i^{-1}[\kappa] \rangle_{i < n} \circ \varphi. \quad \square$$

Remark. In the unweighted case, the condition reads

$$L = \varphi^{-1}\left[\bigcup_{i < m} \bigcap_{k < n_i} p_{ik}^{-1}[K]\right]. \quad \lrcorner$$

Existence of Syntactic Algebras

In general, there is no reason why the syntactic congruence should be an \mathbb{M} -congruence. For finitary monads, we have already shown that minimal algebras always exist. Let us check that the same is true for syntactic algebras. It turns out that the following technical property is of fundamental importance.

Lemma 4.17. *Let $\varepsilon : A \rightarrow C$ be an E -morphism of \mathbb{M} -algebras. Then*

$$\varepsilon \leq \varepsilon \circ p, \quad \text{for all contexts } p \in \mathbb{M}(A + \square).$$

Proof. Given a context p , we can use Lemma 4.3 to find some context q with $q \circ \varepsilon = \varepsilon \circ p$. This implies that $\varepsilon \leq \varepsilon \circ p$. \square

For finitary monads on Pos^Ξ , congruences have the following simple description.

Lemma 4.18. *Let \mathbb{M} be a finitary monad on Pos^Ξ and let \mathfrak{A} and \mathfrak{C} be \mathbb{M} -algebras. An E -morphism $\varepsilon : A \rightarrow C$ is a morphism of \mathbb{M} -algebras if, and only if,*

$$\varepsilon \leq \varepsilon \circ p, \quad \text{for all contexts } p \in \mathbb{M}(A + \square).$$

Proof. (\Rightarrow) was already proved in Lemma 4.17.

(\Leftarrow) It is sufficient to prove that $\Xi := \ker \varepsilon$ is an \mathbb{M} -congruence. Hence, consider a term $u \in \mathbb{M}\Xi$. We have to show that

$$\pi(\mathbb{M}q(u)) \subseteq \pi(\mathbb{M}q'(u)),$$

where $q, q' : A \times A \rightarrow A$ are the two projections. Since \mathbb{M} is finitary, we have $u \in \mathbb{M}\sigma$ for some finite subset $\sigma \subseteq \Xi$. Let $\langle a_o, b_o \rangle, \dots, \langle a_{m-1}, b_{m-1} \rangle$ be an enumeration of σ . For $i < m$, set

$$\theta_i := \leq \cup \{ \langle a_i, b_i \rangle \}.$$

Then $\sigma \subseteq \text{TC}(\theta_o \cup \dots \cup \theta_{m-1})$ and we can use Lemma 3.6 to find $r_o, \dots, r_n \in \mathbb{M}A$ and indices $k_o, \dots, k_{n-1} < m$ such that

$$\mathbb{M}q(u) = r_o \theta_{k_o}^{\mathbb{M}} r_1 \theta_{k_1}^{\mathbb{M}} \dots \theta_{k_{m-1}}^{\mathbb{M}} r_m = \mathbb{M}q'(u).$$

By definition of θ_{k_i} , we can find, for each index $i < m$, some context p_i such that

$$r_i = p_i[a_{k_i}] \quad \text{and} \quad r_{i+1} = p_i[b_{k_i}].$$

Hence,

$$a_{k_i} \subseteq b_{k_i} \quad \text{implies} \quad \pi(r_i) = p_i(a_{k_i}) \subseteq p_i(b_{k_i}) = \pi(r_{i+1}).$$

Consequently, $\pi(\mathbb{M}q(u)) \subseteq \dots \subseteq \pi(\mathbb{M}q'(u))$. □

Theorem 4.19. *Let \mathbb{M} be a finitary monad on Pos^Ξ and \mathfrak{A} an \mathbb{M} -algebra. Every language $\kappa : A_\Xi \rightarrow \Omega$ has a syntactic algebra.*

Proof. By Lemma 4.13 (a), we have

$$\text{syn}_\kappa \leq \text{syn}_\kappa \circ p, \quad \text{for all contexts } p.$$

Hence, it follows by Lemma 4.18 that syn_κ is an \mathbb{M} -congruence. □

Unfortunately, not all the monads \mathbb{M} used in applications are finitary. In particular those needed for languages of infinite words or infinite trees are not. Therefore, we have to extend the preceding proposition to a larger class of functors. It turns out that, in all the known examples of a non-finitary functors where syntactic algebras exists, the functor in question is ‘governed’ in a certain sense by a subfunctor which is finitary. The precise definitions are as follows.

Definition 4.20. (a) A morphism $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ of monads is *dense* over a class \mathcal{C} of \mathbb{M} -algebras if, for all $\mathfrak{A} \in \mathcal{C}$, $C \subseteq A$, and $s \in \mathbb{M}C$, there exists $s^\circ \in \mathbb{M}^\circ C$ with $\pi(\rho(s^\circ)) = \pi(s)$.

(b) We say that a monad \mathbb{M} is *essentially finitary* over a class \mathcal{C} if there exists a morphism $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ such that \mathbb{M}° is finitary and ρ is dense over the closure of \mathcal{C} under binary products. □

Example. Let us again consider the functor

$$\mathbb{M}\langle A_1, A_\infty \rangle := \langle A_1^+, A_1^+ A_\infty \cup A_1^\omega \rangle$$

for infinite words and let

$$\mathbb{M}^\circ \langle A_1, A_\infty \rangle := \langle A_1^+, A_1^+ A_\infty \cup A_1^{\text{up}} \rangle,$$

where A_1^{up} denotes the set of all ultimately periodic words in A_1^ω . One can use a straightforward Ramsey argument to show that the infinite product of a finite ω -semigroup is completely determined by its restriction to all ultimately periodic words. This implies that the inclusion map $\mathbb{M}^\circ \Rightarrow \mathbb{M}$ is dense over the class of all finite ω -semigroups. The case of infinite trees is similar and will be treated in detail in Section V.4. □

Lemma 4.21. *A morphism $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ is dense over a class \mathcal{C} if, and only if, for every algebra $\mathfrak{A} \in \mathcal{C}$ and every set $C \subseteq A$, the subalgebra $\langle\langle C \rangle\rangle_{\mathfrak{A}}$ generated by C in \mathfrak{A} coincides with the subalgebra $\langle\langle C \rangle\rangle_{\mathfrak{A}|_\rho}$ generated by C in the ρ -reduct $\mathfrak{A}|_\rho$.*

Proof. (\Rightarrow) If ρ is dense over \mathcal{C} , we have

$$\langle\langle C \rangle\rangle_{\mathfrak{A}} = \{ \pi(s) \mid s \in \mathbb{M}C \} = \{ \pi(\rho(s^\circ)) \mid s^\circ \in \mathbb{M}^\circ C \} = \langle\langle C \rangle\rangle_{\mathfrak{A}|_\rho},$$

where the first and third step follows from Lemma II.2.2.

(\Leftarrow) Given $s \in \mathbb{M}C$, we have $\pi(s) \in \langle\langle C \rangle\rangle_{\mathfrak{A}} = \langle\langle C \rangle\rangle_{\mathfrak{A}|_\rho}$. Consequently, there exists some $s^\circ \in \mathbb{M}^\circ C$ with $\pi(\rho(s^\circ)) = \pi(s)$. □

If $\mathbb{M}^\circ \Rightarrow \mathbb{M}$ is dense over \mathcal{C} , every \mathbb{M} -algebra in \mathcal{C} is uniquely determined by its \mathbb{M}° -reduct. This will be used below to prove the existence of syntactic algebras for essentially finitary monads.

Lemma 4.22. *Let $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ be dense over a class \mathcal{C} that is closed under binary products.*

- (a) *Any two algebras in \mathcal{C} with the same \mathbb{M}° -reduct are isomorphic.*
- (b) *Let $\varphi : \mathfrak{A}^\circ \rightarrow \mathfrak{B}^\circ$ be a morphism of \mathbb{M}° -algebras and assume that \mathfrak{A}° and \mathfrak{B}° are the \mathbb{M}° -reducts of two \mathbb{M} -algebras $\mathfrak{A}, \mathfrak{B} \in \mathcal{C}$. Then φ is also a morphism $\mathfrak{A} \rightarrow \mathfrak{B}$ of \mathbb{M} -algebras.*
- (c) *An E -morphism ε is a congruence of an \mathbb{M} -algebra $\mathfrak{A} \in \mathcal{C}$ if, and only if, it is a congruence of the \mathbb{M}° -reduct \mathfrak{A}° of \mathfrak{A} .*

Proof. (a) Suppose that \mathcal{C} contains two \mathbb{M} -algebras $\mathfrak{A} = \langle A, \pi \rangle$ and $\mathfrak{A}' = \langle A, \pi' \rangle$ with the same \mathbb{M}° -reduct $\mathfrak{A}^\circ = \langle A, \pi^\circ \rangle$. To show that $\pi = \pi'$, fix an element $s \in \mathbb{M}A$. Set $t := \mathbb{M}d(s) \in \mathbb{M}\Delta$ where $\Delta := \{ \langle a, a \rangle \mid a \in A \}$ is the diagonal of $A \times A$ and $d : A \rightarrow \Delta$ is the diagonal map. By assumption, the product $\mathfrak{A} \times \mathfrak{A}'$ belongs to \mathcal{C} . As ρ is dense, we can find some $t^\circ \in \mathbb{M}^\circ \Delta$ with $\pi^\circ(t^\circ) = \pi(t)$. Note that $t^\circ \in \mathbb{M}^\circ \Delta$ implies that $\mathbb{M}^\circ p(t^\circ) = \mathbb{M}^\circ q(t^\circ)$ where $p, q : A \times A \rightarrow A$ are the two projections. Consequently,

$$\begin{aligned} \pi(s) &= \pi(\mathbb{M}p(t)) = p(\pi(t)) \\ &= p(\pi^\circ(t^\circ)) \\ &= \pi^\circ(\mathbb{M}^\circ p(t^\circ)) \\ &= \pi^\circ(\mathbb{M}^\circ q(t^\circ)) \\ &= q(\pi^\circ(t^\circ)) \\ &= q(\pi(t)) = \pi'(\mathbb{M}q(t)) = \pi'(s). \end{aligned}$$

(b) Fix $s \in \mathbb{M}A$. To show that $\pi(\mathbb{M}\varphi(s)) = \varphi(\pi(s))$, we consider the graph

$$G := \{ \langle a, \varphi(a) \rangle \mid a \in A \}$$

of φ . Let $i := \langle \text{id}, \varphi \rangle : A \rightarrow G$ be the natural bijection and set $t := \mathbb{M}i(s) \in \mathbb{M}G$. Since $\mathfrak{A} \times \mathfrak{B} \in \mathcal{C}$ and ρ is dense, we can find some $t^\circ \in \mathbb{M}^\circ G$ with $\pi(t^\circ) = \pi(t)$. Let $p : A \times B \rightarrow A$ and $q : A \times B \rightarrow B$ be the two projections. Note that

$$\varphi = q \circ i \quad \text{and} \quad q(g) = \varphi(p(g)), \quad \text{for } g \in G,$$

which implies that $\mathbb{M}^\circ q(t^\circ) = \mathbb{M}^\circ(\varphi \circ p)(t^\circ)$. Therefore,

$$\begin{aligned} \pi(\mathbb{M}\varphi(s)) &= \pi(\mathbb{M}q(t)) \\ &= q(\pi(t)) \\ &= q(\pi(t^\circ)) \\ &= \pi(\mathbb{M}^\circ q(t^\circ)) \\ &= \pi(\mathbb{M}^\circ(\varphi \circ p)(t^\circ)) \\ &= \varphi(p(\pi(t^\circ))) \\ &= \varphi(p(\pi(t))) \\ &= \varphi(\pi(\mathbb{M}p(t))) = \varphi(\pi(s)). \end{aligned}$$

(c) Clearly, every \mathbb{M} -algebra morphism is also an \mathbb{M}° -algebra morphism. The converse follows from (b). \square

Let us also note the following property.

Proposition 4.23. *Suppose that \mathbb{M} is essentially finitary over \mathcal{C} and that \mathcal{D} has canonical subobject diagrams. Every algebra in \mathcal{C} is the colimit (in $\text{Alg}(\mathbb{M})$) of its canonical \aleph_0 -subobject diagram.*

Proof. Let $\mathfrak{A} \in \mathcal{C}$, let D be the canonical \aleph_0 -subobject diagram of A , and let $\iota_X : \langle\langle X \rangle\rangle_{\mathfrak{A}} \rightarrow A$ be the inclusion map. Since \mathcal{D} has canonical subobject diagrams, it follows that A is the limit of D in \mathcal{D}^Ξ and that $(\iota_X)_X$ is the corresponding limiting cone.

It remains to show that it is also limiting in $\text{Alg}(\mathbb{M})$. Hence, consider a cocone $(\mu_X)_X$ from D to some \mathbb{M} -algebra \mathfrak{B} . We obtain a unique morphism $\varphi : A \rightarrow B$ satisfying

$$\varphi \circ \iota_X = \mu_X, \quad \text{for all } X.$$

It remains to show that $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ is a \mathbb{M} -morphism. By assumption, there exists a morphism of monads $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ such that \mathbb{M}° is finitary and ρ is dense over \mathcal{C} . By Lemma 4.22, it is sufficient to prove that φ is a \mathbb{M}° -morphism $\mathfrak{A}^\circ \rightarrow \mathfrak{B}^\circ$ between the respective ρ -reducts. Let $s \in \mathbb{M}^\circ A$.

As \mathbb{M}° is finitary, there is some finite set $X \subseteq A$ with $s \in \mathbb{M}^\circ X$. It follows that

$$\pi(\mathbb{M}^\circ \varphi(s)) = \pi(\mathbb{M}^\circ \mu_X(s)) = \mu_X(\pi(s)) = \varphi(\pi(s)),$$

where the first step and the last one follow from the fact that $\varphi \upharpoonright \langle\langle X \rangle\rangle_{\mathfrak{A}} = \varphi \circ \iota_X = \mu_X$. \square

Using the existence result for the category Pos , we obtain the following consequence.

Theorem 4.24. *Let \mathbb{M} be a monad on Pos^Ξ that is essentially finitary over \mathcal{C} and let \mathfrak{A} be an \mathbb{M} -algebra. If $\kappa : A_\xi \rightarrow \Omega$ is recognised by some E -morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{C}$ with $\mathfrak{C} \in \mathcal{C}$, then $\ker \text{syn}_\kappa$ is an \mathbb{M} -congruence.*

Proof. Suppose that $\kappa = \varphi^{-1}[\mu]$, for some $\mu \in \mathbb{L}^\circ C_\xi$. Let $\mathfrak{B} \subseteq \mathfrak{C}$ be the subalgebra induced by $\text{rng } \varphi$. By Theorem 4.19, syn_μ is a morphism of \mathbb{M}° -algebras. Hence, Lemma 4.22 (c) implies that it is also a \mathbb{M} -algebra morphism. Consequently, so is its restriction to \mathfrak{B} . To prove that syn_κ is a morphism of \mathbb{M} -algebras it is therefore sufficient to show that

$$\text{syn}_\kappa = \text{syn}_\mu \circ \varphi.$$

Let

$$\begin{aligned} \rho &:= \langle \kappa \circ p \rangle_p : A \rightarrow \prod_p \Omega, \\ \sigma &:= \langle \mu \circ p \rangle_p : C \rightarrow \prod_p \Omega \end{aligned}$$

be the morphisms such that $\text{syn}_\kappa := \ker \rho$ and $\text{syn}_\mu := \ker \sigma$. Lemma 4.3 implies that

$$\rho = \langle \kappa \circ p \rangle_p = \langle \mu \circ \varphi \circ p \rangle_p = \langle \mu \circ q \circ \varphi \rangle_q = \sigma \circ \varphi.$$

Since $\varphi \in E$ it follows that

$$\text{syn}_\kappa = \ker \rho = \ker (\sigma \circ \varphi) = \ker \sigma \circ \varphi = \text{syn}_\mu \circ \varphi. \quad \square$$

The Algebraic Structure on the Weights

Usually, the set Ω of weights is not a plain set, but carries additional algebraic structure (e.g., a field, a boolean algebra, or a semiring). In this section, we explain how to make use of this additional structure. Assume that Ω is equipped with an \mathbb{O} -algebra product $\pi : \mathbb{O}\Omega \rightarrow \Omega$, for some a monad $\langle \mathbb{O}, \lambda, \varepsilon \rangle$ on \mathcal{D} . The three most common choices for \mathbb{O} are as follows.

- If we do not need any algebraic structure on Ω , we can take the identity monad $\mathbb{I}X = X$.
- If Ω is a semiring, we usually take the monad Lin mapping a set X to the Ω -semimodule $\text{Lin}(X)$ generated by X , i.e., the set of all finite linear combinations of elements of X with coefficients in Ω .
- If Ω is a lattice, we can take the monad DL mapping a set X to the free distributive lattice $\text{DL}(X)$ generated by X , i.e., the set of all finite positive boolean combinations of elements of X .

We denote the lift of \mathbb{O} to \mathcal{D}^Ξ by $\mathbb{O}^\dagger : \mathcal{D}^\Xi \rightarrow \mathcal{D}^\Xi$. To make sure that this additional algebraic structure plays well with the given monad \mathbb{M} , we additionally assume that there exists a distributive law $\delta : \mathbb{M}\mathbb{O}^\dagger \rightarrow \mathbb{O}^\dagger\mathbb{M}$. This distributive law in particular induces a monad structure on the composition $\mathbb{O}^\dagger\mathbb{M}$.

As $\mathbb{O}\mathbb{M}_\xi\Sigma$ is the free \mathbb{O} -algebra over $\mathbb{M}_\xi\Sigma$, we can lift each Ω -language $\kappa : \mathbb{M}_\xi\Sigma \rightarrow \Omega$ to a unique morphism of \mathbb{O} -algebras $\hat{\kappa} : \mathbb{O}\mathbb{M}_\xi\Sigma \rightarrow \Omega$ with $\hat{\kappa} \circ \varepsilon = \kappa$. Furthermore, if $\varphi : \mathbb{O}^\dagger\mathbb{M}\Sigma \rightarrow \mathfrak{A}$ is a surjective morphism of $\mathbb{O}^\dagger\mathbb{M}$ -algebras recognising $\hat{\kappa}$, i.e., $\hat{\kappa} = \mu \circ \varphi$, it follows by Lemma I.5.6 that $\mu : A_\xi \rightarrow \Omega$ is a morphism of \mathbb{O} -algebras.

We can now apply the framework developed so far to the functor $\mathbb{O}^\dagger\mathbb{M}$ instead of \mathbb{M} . In particular, the results concerning the existence of syntactic algebras still apply in this setting, as well as the Variety Theorem and the Reiterman theorem which we will prove below.

What we have gained by this translation is that we may assume that all languages $\kappa : \mathbb{O}\mathbb{M}_\xi\Sigma \rightarrow \Omega$ are morphisms of \mathbb{O} -algebras, that they are recognised by morphisms of $\mathbb{O}^\dagger\mathbb{M}$ -algebras, and, in particular, that the corresponding function $\mu : A_\xi \rightarrow \Omega$ is an \mathbb{O} -algebra morphism.

Let us take a look at how the syntactic $\mathbb{O}^+\mathbb{M}$ -algebra of an Ω -language κ relates to its syntactic \mathbb{M} -algebra. It turns out that, as \mathbb{M} -algebras, the latter is a retract of the former.

Proposition 4.25. *Let $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ be an Ω -language and $\hat{\kappa} : \mathbb{O}\mathbb{M}_\xi \Sigma \rightarrow \Omega$ its lifting to $\mathbb{O}^+\mathbb{M}\Sigma$. There exists an embedding $e : \text{Syn}(\kappa) \rightarrow \text{Syn}(\hat{\kappa})$ of \mathbb{M} -algebras such that $\text{Syn}(\hat{\kappa})$ is generated (as an \mathbb{O}^+ -algebra) by $\text{rng } e$ and e satisfies*

$$\hat{\mu} \circ e = \mu \quad \text{and} \quad e \circ \text{syn}_\kappa = \text{syn}_{\hat{\kappa}} \circ \varepsilon,$$

where $\mu : \text{Syn}_\xi(\kappa) \rightarrow \Omega$ and $\hat{\mu} : \text{Syn}_\xi(\hat{\kappa}) \rightarrow \Omega$ are the functions such that $\kappa = \mu \circ \text{syn}_\kappa$ and $\hat{\kappa} = \hat{\mu} \circ \text{syn}_{\hat{\kappa}}$.

Proof. Let $\pi : \mathbb{O}\Omega \rightarrow \Omega$ be the \mathbb{O} -algebra product of Ω . Since

$$\pi \circ \mathbb{O}\mu \circ \mathbb{O}^+\text{syn}_\kappa = \pi \circ \mathbb{O}\kappa = \pi \circ \mathbb{O}(\hat{\kappa} \circ \varepsilon) = \hat{\kappa} \circ \pi \circ \mathbb{O}\varepsilon = \hat{\kappa},$$

the morphism $\mathbb{O}^+\text{syn}_\kappa$ recognises $\hat{\kappa}$. Consequently, there exists an $\mathbb{O}^+\mathbb{M}$ -algebra morphism $\varphi : \mathbb{O}^+\text{Syn}(\kappa) \rightarrow \text{Syn}(\hat{\kappa})$ with

$$\varphi \circ \mathbb{O}^+\text{syn}_\kappa = \text{syn}_{\hat{\kappa}}.$$

Setting $e := \varphi \circ \varepsilon$ it follows that

$$e \circ \text{syn}_\kappa = \varphi \circ \varepsilon \circ \text{syn}_\kappa = \varphi \circ \mathbb{O}\text{syn}_\kappa \circ \varepsilon = \text{syn}_{\hat{\kappa}} \circ \varepsilon.$$


As syn_κ is an epimorphism, this implies that

$$\text{rng } e = \text{rng } (\varphi \circ \varepsilon) = \text{rng } (\varphi \circ \varepsilon \circ \text{syn}_\kappa) = \text{rng } (\text{syn}_{\hat{\kappa}} \circ \varepsilon).$$

Hence, e factorises through \mathfrak{S} and we obtain a morphism $e_o : \text{Syn}(\kappa) \rightarrow \mathfrak{S}$ with

$$i \circ e_o = e \quad \text{and} \quad e_o \circ \text{syn}_\kappa = \eta.$$

Furthermore, $\text{Syn}(\hat{\kappa})$ is generated by $\text{rng } e = \text{rng } (\text{syn}_{\hat{\kappa}} \circ \varepsilon)$ since the morphism $\text{syn}_{\hat{\kappa}}$ is an epimorphism and its domain $\mathbb{O}^+\mathbb{M}\Sigma$ is generated by $\text{rng } \varepsilon$.



To prove that e is an embedding of \mathbb{M} -algebras, we show that e_o is an isomorphism by finding an inverse. Note that $\mathbb{O}^+\mathbb{M}\Sigma$ is a \mathbb{M} -algebra with product $\mathbb{O}^+\text{flat} \circ \delta$. Since

$$(\mathbb{O}^+\text{flat} \circ \delta) \circ \mathbb{M}\varepsilon = \mathbb{O}^+\text{flat} \circ \varepsilon = \varepsilon \circ \text{flat},$$

it therefore follows that $\varepsilon : \mathbb{M}\Sigma \rightarrow \mathbb{O}\mathbb{M}\Sigma$ is a morphism of \mathbb{M} -algebras. Hence, so is $\eta := \text{syn}_{\hat{\kappa}} \circ \varepsilon$. Let $\mathfrak{S} \subseteq \text{Syn}(\hat{\kappa})$ be the \mathbb{M} -subalgebra of $\text{Syn}(\hat{\kappa})$ induces by $\text{rng } \eta$ and let $i : \mathfrak{S} \rightarrow \text{Syn}(\hat{\kappa})$ be the corresponding inclusion map. Since

$$\hat{\mu} \circ \text{syn}_{\hat{\kappa}} \circ \varepsilon = \hat{\kappa} \circ \varepsilon = \kappa,$$

the morphism $\eta : \mathbb{M}\Sigma \rightarrow \mathfrak{S}$ recognises κ . Therefore, there exists an \mathbb{M} -algebra morphism $f : \mathfrak{S} \rightarrow \text{Syn}(\kappa)$ with

$$f \circ \eta = \text{syn}_\kappa.$$

To see that f is the desired inverse of e_o , note that

$$f \circ e_o \circ \text{syn}_\kappa = f \circ \eta = \text{syn}_\kappa \quad \text{and} \quad e_o \circ f \circ \eta = e_o \circ \text{syn}_\kappa = \eta.$$

By surjectivity of syn_κ and η , we obtain $f \circ e_o = \text{id}$ and $e_o \circ f = \text{id}$. Consequently, $e_o : \text{Syn}(\kappa) \cong \mathfrak{S}$ is an isomorphism and $e = i \circ e_o : \text{Syn}(\kappa) \rightarrow \text{Syn}(\hat{\kappa})$ an embedding (of \mathbb{M} -algebras).

Finally, note that

$$\hat{\mu} \circ e \circ \text{syn}_\kappa = \hat{\mu} \circ \text{syn}_{\hat{\kappa}} \circ \varepsilon = \hat{\kappa} \circ \varepsilon = \kappa = \mu \circ \text{syn}_\kappa.$$

By surjectivity of syn_κ , this implies that $\hat{\mu} \circ e = \mu$. \square

5. Varieties

After these preparations, we come to the first of the central theorems of algebraic language theory: the Variety Theorem. This theorem characterises which kind of language families are amenable to our algebraic tools by establishing a correspondence between language families and the classes of algebras recognising them.

We fix a class \mathcal{R} that serves as recognisers for the languages we are interested in.

Convention. In this section, we adopt the following assumptions.

- We denote by \mathcal{T} the class of all free \mathbb{M} -algebras of the form $\mathbb{M}\Sigma$, for $\Sigma \in \text{Alph}$.
- We fix a class \mathcal{R} of \mathbb{M} -algebras such that
 - every $\mathfrak{A} \in \mathcal{R}$ is \mathcal{T} -generated, sort-wise strongly finite-dimensional, and sort-wise noetherian;
 - \mathcal{R} is closed under \mathcal{T} -generated M -subalgebras of finite products.

We call the objects in \mathcal{R} recognisers and those in \mathcal{T} term algebras.

Remark. The requirement that the recognisers are \mathcal{T} -generated is no restriction: if $\varphi : \mathfrak{T} \rightarrow \mathfrak{A}$ recognises a language κ with $\mathfrak{A} \in \mathcal{R}$, then so does $\ker \varphi : \mathfrak{T} \rightarrow \mathfrak{R}$, and \mathfrak{R} is \mathcal{T} -generated. Hence, $\mathfrak{R} \in \mathcal{R}$. \lrcorner

In Section II.5, we defined a variety to be a class $\mathcal{V} \subseteq \text{Alg}(\mathbb{M})$ closed under E_* -quotients and \mathcal{P} -generated M -subobjects of arbitrary products, while a pseudo-variety is a class $\mathcal{V} \subseteq \text{Alg}(\mathbb{M})$ closed under E_* -quotients

and \mathcal{P} -generated M -subobjects of finite products. For our present purpose, we need a slightly different definition: to support infinite sets of sorts Ξ we have to add one more closure property to the classes in question.

Definition 5.1. An \mathbb{M} -algebra \mathfrak{A} is a *sort-accumulation point* of a class \mathcal{C} of \mathbb{M} -algebras if, for every finite subset $\Delta \subseteq \Xi$, there is some $\mathfrak{C} \in \mathcal{C}$ and an E -morphism $\varepsilon : \mathfrak{C}|_\Delta \rightarrow \mathfrak{A}|_\Delta$. \lrcorner

We will show below that there is a precise correspondence between the following families of languages and classes of algebras.

Definition 5.2. (a) A *variety of languages* is a contravariant functor \mathcal{K} mapping the full subcategory of $\text{Alg}(\mathbb{M})$ induced by all algebras of the form $\mathbb{M}\Sigma$ with $\Sigma \in \text{Alph}$ to the category of language algebras such that

- (i) each $\mathcal{K}[\mathfrak{T}]$ forms a subalgebra of $\mathbb{L}\mathfrak{T}$,
- (ii) for $\varphi : \mathfrak{S} \rightarrow \mathfrak{T}$, the morphism $\mathcal{K}[\varphi] : \mathcal{K}[\mathfrak{T}] \rightarrow \mathcal{K}[\mathfrak{S}]$ is the corresponding restriction of $\mathbb{L}\varphi : \mathbb{L}\mathfrak{T} \rightarrow \mathbb{L}\mathfrak{S}$.
- (b) A class $\mathcal{V} \subseteq \mathcal{R}$ is an \mathcal{R} -variety if it is closed under
 - (i) images under surjective E -morphisms that belong to \mathcal{R} ,
 - (ii) \mathcal{T} -generated M -subalgebras of finite products,
 - (iii) \mathcal{T} -generated sort-accumulation points that belong to \mathcal{R} . \lrcorner

Remark. (a) Note that every variety of languages \mathcal{K} is uniquely determined by the domains of the algebras $\mathcal{K}[\mathfrak{T}]$ the algebraic structure being induced by that of $\mathbb{L}\mathfrak{T}$. In the following we will not strictly distinguish between the view of a variety as a functor or as a set of languages.

(b) We obtain the three usual closure properties of a variety of languages as follows. (i) The fact that $\mathcal{K}[T_\xi]$ is a language algebra means that $\mathcal{K}[T_\xi]$ is closed under language operations (Boolean operations, Ω -operations, etc.) and under derivatives. (ii) The fact that $\mathcal{K}[\varphi] = \mathbb{L}\varphi$ is a morphism $\mathcal{K}[\mathfrak{S}] \rightarrow \mathcal{K}[\mathfrak{T}]$ implies that \mathcal{K} is closed under inverse morphisms.

(c) In the definition of an \mathcal{R} -variety of algebras, closure under quotients is superfluous as it is implied by closure under sort-accumulation points.

We have left it as a requirement in the definition to emphasise the analogy to the usual definition in the setting with finitely many sorts.

(d) The reason why we combine the operations of taking subalgebras and forming products into a single one is that, in general, the product of two finitary algebras need not be finitely generated (see Section V.4 for a counterexample).

(e) If the set of sorts Ξ is finite and \mathcal{R} is the class of all finite \mathbb{M} -algebras, an \mathcal{R} -variety is just a pseudo-variety. \lrcorner

We start with a characterisation of varieties in terms of closure operations like in Theorem II.5.12.

Definition 5.3. For a class \mathcal{C} of \mathbb{M} -algebras we set

$$\begin{aligned} \mathcal{C}(\mathcal{C}) &:= \{ \mathfrak{A} \mid \mathfrak{A} \text{ a sort-accumulation point of some } \mathfrak{B} \in \mathcal{C} \}, \\ \mathcal{H}_o(\mathcal{C}) &:= \{ \mathfrak{A} \mid \mathfrak{A} \text{ image of an } E\text{-morphisms of some } \mathfrak{B} \in \mathcal{C} \}, \\ \mathcal{H}(\mathcal{C}) &:= \mathcal{H}_o(\mathcal{C}) \cap \mathcal{R}, \\ \mathcal{S}(\mathcal{C}) &:= \{ \mathfrak{A} \mid \mathfrak{A} \text{ a } \mathcal{T}\text{-generated } M\text{-subalgebra of some } \mathfrak{B} \in \mathcal{C} \}, \\ \mathcal{P}(\mathcal{C}) &:= \{ \mathfrak{A} \mid \mathfrak{A} \text{ a finite product of algebras in } \mathcal{C} \}. \end{aligned} \quad \lrcorner$$

Proposition 5.4. Suppose that $\text{Alg}(\mathbb{M})$ has coequalisers and let $\mathcal{V} \subseteq \mathcal{R}$ be a class of algebras.

- (a) $(\text{HSP})^2(\mathcal{V}) = \text{HSP}(\mathcal{V})$.
- (b) The following conditions are equivalent.
 - (1) \mathcal{V} is an \mathcal{R} -variety.
 - (2) $\mathcal{V} = \text{HSP}(\mathcal{V})$ and $\mathcal{V} = \mathcal{C}(\mathcal{V})$
 - (3) \mathcal{V} satisfies the following two statements.
 - ◆ The reduct $\mathcal{V}|_\Delta$ is an $\mathcal{R}|_\Delta$ -variety, for every finite $\Delta \subseteq \Xi$.
 - ◆ \mathcal{V} is the closure of the reducts $\mathcal{V}|_\Delta$ in the sense that

$$\mathfrak{A} \in \mathcal{V} \quad \text{iff} \quad \mathfrak{A}|_\Delta \in \mathcal{V}|_\Delta, \quad \text{for all finite } \Delta \subseteq \Xi.$$

Proof. (a) We have shown in Theorem II.5.12 that (in the notation of the current section) $(\text{H}_o\text{SP})^2 = \text{H}_o\text{SP}$. Since \mathcal{R} is closed under SP, it follows that

$$(\text{HSP})^2 = (\text{H}_o\text{SP} \cap \mathcal{R})^2 = (\text{H}_o\text{SP})^2 \cap \mathcal{R} = \text{H}_o\text{SP} \cap \mathcal{R} = \text{HSP}.$$

(b) (1) \Rightarrow (2) immediately follows from the closure properties of an \mathcal{R} -variety and (2) \Rightarrow (1) follows by (a).

(3) \Rightarrow (2) It is sufficient to show that $\text{CHSP}(\mathcal{V}) = \mathcal{V}$. Hence, let $\mathfrak{A} \in \text{CHSP}(\mathcal{V})$. Then $\mathfrak{A}|_\Delta \in \text{HSP}(\mathcal{V}|_\Delta) = \mathcal{V}|_\Delta$, for all finite $\Delta \subseteq \Xi$. Consequently, $\mathfrak{A} \in \mathcal{V}$.

(2) \Rightarrow (3) If \mathfrak{A} is an algebra with $\mathfrak{A}|_\Delta \in \mathcal{V}|_\Delta$, for all finite $\Delta \subseteq \Xi$. Then $\mathcal{C}(\mathcal{V}) = \mathcal{V}$ implies that $\mathfrak{A} \in \mathcal{V}$. For the other claim, note that

$$\mathcal{V} = \text{HSP}(\mathcal{V}) \quad \text{implies} \quad \mathcal{V}|_\Delta = \text{HSP}(\mathcal{V}|_\Delta),$$

since the reduct operation $|_\Delta$ commutes with H, S, and P. By (a), this means that $\mathcal{V}|_\Delta$ is an $\mathcal{R}|_\Delta$ -variety. \square

The aim of the rest of this section is to establish a one-to-one correspondence between varieties of languages and \mathcal{R} -varieties of \mathbb{M} -algebras. The arguments are mostly standard, except for some adjustments needed to support infinitely many sorts. We start with the following observation.

Lemma 5.5. Let \mathcal{V} be an \mathcal{R} -variety and $\kappa : T_\Xi \rightarrow \Omega$ a language where $\Xi \in \mathcal{T}$ such that κ has a minimal algebra that belongs to \mathcal{R} . Then κ is recognised by some algebra $\mathfrak{A} \in \mathcal{V}$ if, and only if, $\text{Syn}(\kappa) \in \mathcal{V}$.

Proof. (\Leftarrow) is trivial since $\text{syn}_\kappa : \mathfrak{T} \rightarrow \text{Syn}(\kappa)$ recognises κ . For (\Rightarrow), consider a morphism $\varphi : \mathfrak{T} \rightarrow \mathfrak{A}$ recognising κ with $\mathfrak{A} \in \mathcal{V}$. Set $\mathfrak{B} := \mathfrak{T} / \ker \varphi$. Since $\text{im } \varphi : \mathfrak{B} \rightarrow \mathfrak{A}$ is an M -morphism and $\ker \varphi : \mathfrak{T} \rightarrow \mathfrak{B}$ an E -morphism, it follows that $\mathfrak{B} \in \mathcal{V}$. Furthermore, $\ker \varphi$ recognises κ . As $\text{Syn}(\kappa)$ is minimal, we can find a morphism $\rho : \mathfrak{A} \rightarrow \text{Syn}(\kappa)$ with $\text{syn}_\kappa = \rho \circ \ker \varphi$. Since syn_κ and $\ker \varphi$ are E -morphisms, it follows by Lemma II.1.3 (c) that $\rho \in E$. Hence, $\text{Syn}(\kappa) \in \text{H}(\mathcal{V}) = \mathcal{V}$. \square

The first step in correlating varieties of languages and \mathcal{R} -varieties of algebras consists in the following fact.

Proposition 5.6. *Let \mathcal{V} be an \mathcal{R} -variety. The family of languages \mathcal{K} recognised by the algebras in \mathcal{V} forms a variety of languages.*

Proof. We have to prove two closure properties.

(i) We start by proving that $\mathcal{K}[\mathfrak{Z}]$ forms a subalgebra of $\mathbb{L}\mathfrak{Z}$, that is, that

$$\langle\langle \mathcal{K}[\mathfrak{Z}] \rangle\rangle_{\mathbb{L}} = \mathcal{K}[\mathfrak{Z}].$$

Since all operations of $\mathbb{L}\mathfrak{Z}$ have finite arity, it is sufficient to show that

$$\langle\langle C \rangle\rangle_{\mathbb{L}} \subseteq \mathcal{K}[\mathfrak{Z}], \quad \text{for all finite } C \subseteq \mathcal{K}[\mathfrak{Z}].$$

Hence, consider languages $\kappa_0, \dots, \kappa_{n-1} \in \mathcal{K}[\mathfrak{Z}]$ for $\mathfrak{Z} \in \mathcal{T}$. Then there are morphisms $\varphi_i : \mathfrak{Z} \rightarrow \mathfrak{A}^i$ with $\mathfrak{A}^i \in \mathcal{R}$ and languages $\mu_i \in \mathbb{L}A^i$ such that $\kappa_i = \mathbb{L}\varphi_i(\mu_i)$. Fix $\lambda \in \langle\langle \kappa_0, \dots, \kappa_{n-1} \rangle\rangle_{\mathbb{L}}$. We have to show that $\lambda \in \mathcal{K}[\mathfrak{Z}]$.

Let $\mathfrak{Z} \xrightarrow{e} \mathfrak{B} \xrightarrow{m} \prod_{i < n} \mathfrak{A}^i$ be the EM -factorisation of $\langle \varphi_0, \dots, \varphi_{n-1} \rangle$. Then \mathfrak{B} is a \mathcal{T} -generated M -subalgebra of a finite product of algebras in \mathcal{R} . This implies that $\mathfrak{B} \in \mathcal{R}$. To prove that $\lambda \in \mathcal{K}[\mathfrak{Z}]$ it is therefore sufficient to show that the morphism $e : \mathfrak{Z} \rightarrow \mathfrak{B}$ recognises λ . Note that, being a morphism language algebras, $\mathbb{L}e$ commutes with all operations. Consequently, we have

$$\begin{aligned} \lambda \in \langle\langle \kappa_0, \dots, \kappa_{n-1} \rangle\rangle_{\mathbb{L}} &= \langle\langle \mathbb{L}\varphi_0(\mu_0), \dots, \mathbb{L}\varphi_{n-1}(\mu_{n-1}) \rangle\rangle_{\mathbb{L}} \\ &= \langle\langle \mathbb{L}(\varphi_0 \times \dots \times \varphi_{n-1})(\bar{\mu}) \rangle\rangle_{\mathbb{L}} \\ &= \langle\langle \mathbb{L}(m \circ e)(\bar{\mu}) \rangle\rangle_{\mathbb{L}} \\ &= \langle\langle \mathbb{L}e\mathbb{L}m(\bar{\mu}) \rangle\rangle_{\mathbb{L}} \\ &= \mathbb{L}e\langle\langle \mathbb{L}m(\bar{\mu}) \rangle\rangle_{\mathbb{L}}, \end{aligned}$$

which implies that $\lambda = \mathbb{L}e(v)$, for some $v \in \mathbb{L}B$.

(ii) It remains to prove that, for every morphism $\psi : \mathfrak{S} \rightarrow \mathfrak{Z}$ between algebras in \mathcal{T} , the morphism $\mathbb{L}\psi : \mathbb{L}\mathfrak{Z} \rightarrow \mathbb{L}\mathfrak{S}$ restricts to a morphism

$\mathcal{K}[\mathfrak{Z}] \rightarrow \mathcal{K}[\mathfrak{S}]$. Hence, fix $\kappa \in \mathcal{K}[\mathfrak{Z}]$. By assumption, there is some algebra $\mathfrak{A} \in \mathcal{V}$ recognising κ . Suppose that $\kappa = \mathbb{L}\varphi(\mu)$ where $\varphi : \mathfrak{Z} \rightarrow \mathfrak{A}$ and $\mu \in \mathbb{L}_{\xi}A$. Then

$$\mathbb{L}\psi(\kappa) = \mathbb{L}\psi\mathbb{L}\varphi(\mu) = \mathbb{L}(\varphi \circ \psi)(\mu),$$

which means that $\varphi \circ \psi : \mathfrak{S} \rightarrow \mathfrak{A}$ recognises $\mathbb{L}\psi(\kappa)$. This implies that $\mathbb{L}\psi(\kappa) \in \mathcal{K}_{\xi}[\mathfrak{S}]$. \square

It remains to prove the converse direction of the correspondence. We start with two lemmas.

Lemma 5.7. *Let $q : \mathfrak{A} \rightarrow \mathfrak{B}$ be an E -morphism and let $\mathfrak{Z} \in \mathcal{T}$. Every language $\kappa : T_{\xi} \rightarrow \Omega$ recognised by \mathfrak{B} is also recognised by \mathfrak{A} .*

Proof. Suppose that $\kappa = \psi^{-1}[\mu]$ where $\psi : \mathfrak{Z} \rightarrow \mathfrak{B}$ and $\mu : B_{\xi} \rightarrow \Omega$. Since \mathfrak{Z} is projective with respect to E -morphisms, we can find a morphism $\varphi : \mathfrak{Z} \rightarrow \mathfrak{A}$ such that $q \circ \varphi = \psi$. It follows that φ also recognises κ since $\mathbb{L}\varphi\mathbb{L}q(\mu) = \mathbb{L}\psi(\mu) = \kappa$ and $\mathbb{L}q(\mu) \in \mathbb{L}_{\xi}A$. \square

Lemma 5.8. *Let $\mathfrak{Z} \in \mathcal{T}$ and suppose that $\kappa : T_{\xi} \rightarrow \Omega$ is recognised by an algebra \mathfrak{C} such that there exists an M -morphism $\sigma : \mathfrak{C} \rightarrow \prod_{i < n} \mathfrak{A}^i$ with $\mathfrak{A}^i \in \mathcal{R}$ and $n < \omega$. Then*

$$\kappa \in \langle\langle \lambda_0, \dots, \lambda_{n-1} \rangle\rangle_{\mathbb{L}},$$

where each $\lambda_j : T_{\xi} \rightarrow \Omega$ is recognised by some factor \mathfrak{A}^{k_j} .

Proof. Suppose that $\kappa = \mathbb{L}\varphi(\mu)$ for $\varphi : \mathfrak{Z} \rightarrow \mathfrak{C}$ and $\mu : C_{\xi} \rightarrow \Omega$. By Lemma 2.4, the function $\mathbb{L}\sigma$ is surjective. Fix some $v \in \mathbb{L}_{\xi} \prod_i A^i$ with $\mathbb{L}\sigma(v) = \mu$. Since $\mathfrak{A}^i \in \mathcal{R}$, the sets A^i are strongly finite-dimensional. Hence, it follows by Lemma 2.8 that there exists an ω -operation ω and languages $\lambda_i \in \mathbb{L}_{\xi}A^{k_i}$ such that

$$v = \omega[\mathbb{L}p_0(\lambda_0), \dots, \mathbb{L}p_{n-1}(\lambda_{n-1})],$$

where $p_i : \prod_j A^j \rightarrow A^i$ are the projections. Consequently,

$$\begin{aligned} \kappa &= \mathbb{L}\varphi(\mu) \\ &= \mathbb{L}\varphi\mathbb{L}\sigma(\nu) \\ &= \mathbb{L}(\sigma \circ \varphi)(\omega[\mathbb{L}p_o(\lambda_o), \dots, \mathbb{L}p_{n-1}(\lambda_{n-1})]) \\ &= \omega[\mathbb{L}(\sigma \circ \varphi \circ p_o)(\lambda_o), \dots, \mathbb{L}(\sigma \circ \varphi \circ p_{n-1})(\lambda_{n-1})], \end{aligned}$$

where each language $\mathbb{L}(\sigma \circ \varphi \circ p_j)(\lambda_j)$ is recognised by some factor \mathfrak{A}^{k_j} . \square

Theorem 5.9. *Let \mathcal{K} be a variety of languages such that every $\kappa \in \mathcal{K}$ has a minimal algebra $\text{Syn}(\kappa) \in \mathcal{R}$. A language κ belongs to \mathcal{K} if, and only if, it is recognised by some algebra from the \mathcal{R} -variety \mathcal{V} generated by the set $\mathcal{S} := \{\text{Syn}(\kappa) \mid \kappa \in \mathcal{K}\}$.*

Proof. (\Rightarrow) Every language $\kappa \in \mathcal{K}$ is recognised by $\text{Syn}(\kappa)$, which belongs to \mathcal{V} .

(\Leftarrow) Set $\mathcal{V}_o := \mathcal{S}$ and $\mathcal{V}_{n+1} := \text{CHSP}(\mathcal{V}_n)$, for $n < \omega$. It follows from Proposition 5.4 that $\mathcal{V} = \bigcup_{n < \omega} \mathcal{V}_n$. By induction on n , we show that every language recognised by an algebra $\mathfrak{A} \in \mathcal{V}_n$ belongs to \mathcal{K} . For $n = o$, suppose that $\lambda \in \mathbb{L}_\zeta \mathcal{S}$ is recognised by $\text{Syn}(\kappa)$ where $\kappa \in \mathcal{K}_\xi[\mathfrak{T}]$ and $\mathfrak{S}, \mathfrak{T} \in \mathcal{T}$. Then we have $\lambda = \mathbb{L}\varphi(\mu)$ for some morphism $\varphi : \mathfrak{S} \rightarrow \text{Syn}(\kappa)$ and a language $\mu \in \mathbb{L}_\zeta \text{Syn}(\kappa)$. As $\text{Syn}(\kappa)$ is sort-wise noetherian, it follows by Proposition 4.14 that $\mu \in \langle\langle \text{res}_\kappa \rangle\rangle_{\mathbb{L}}$. Thus,

$$\lambda = \mathbb{L}\varphi(\mu) \in \mathbb{L}\varphi[\langle\langle \text{res}_\kappa \rangle\rangle_{\mathbb{L}}] \subseteq \mathbb{L}\varphi[\mathcal{K}_\zeta[\mathfrak{T}]] \subseteq \mathcal{K}_\zeta[\mathfrak{S}].$$

For the inductive step, suppose that we have already proved the claim algebras in \mathcal{V}_n and consider an algebra $\mathfrak{B} \in \mathcal{V}_{n+1}$. Let $\varphi : \mathfrak{T} \rightarrow \mathfrak{B}$ be a morphism recognising $\kappa = \mathbb{L}\varphi(\mu)$ with $\mu \in \mathbb{L}_\xi \mathcal{B}$. Suppose that $\mathfrak{T} = \mathbb{M}\Sigma$ for $\Sigma \in \text{Alph}$ and let $\Delta \subseteq \Xi$ be the set consisting of ξ and all sorts appearing in Σ . As \mathfrak{B} is a sort-accumulation point of $\text{HSP}(\mathcal{V}_n)$, we can find an algebra $\mathfrak{A} \in \text{HSP}(\mathcal{S})$ and an E -morphism $\varepsilon : \mathfrak{A}|_\Delta \rightarrow \mathfrak{B}|_\Delta$. Since $\mathfrak{T}|_\Delta$ is projective with respect to E -morphisms, there exists a unique morphism $\psi_o : \mathfrak{T}|_\Delta \rightarrow$

$\mathfrak{A}|_\Delta$ with $\varepsilon \circ \psi_o = \varphi|_\Delta$. We turn ψ_o into a morphism $\psi : \mathfrak{T} \rightarrow \mathfrak{A}$ as follows. For $c \in \Sigma$, we set

$$\psi(\text{sing}(c)) := \psi_o(\text{sing}(c)).$$

As \mathfrak{T} is freely generated by rng sing , this uniquely determines a morphism $\psi : \mathfrak{T} \rightarrow \mathfrak{A}$ of \mathbb{M} -algebras. For $s \in \mathbb{M}_\xi \Sigma$, it follows that

$$\psi(s) = (\pi \circ \mathbb{M}(\psi_o \circ \text{sing}))(s) = (\psi_o \circ \pi \circ \mathbb{M}\text{sing})(s) = \psi_o(s),$$

where the second step follows from the fact that ψ_o is a morphism of $\mathbb{M}|_\Delta$ -algebras and $\xi \in \Delta$. Consequently,

$$(\mu \circ \varepsilon \circ \psi)(s) = (\mu \circ \varepsilon \circ \psi_o)(s) = (\mu \circ \varphi)(s) = \kappa(s).$$

Thus, κ is recognised by $\mathfrak{A} \in \text{HSP}(\mathcal{V}_n)$. Using Lemmas 5.7 and 5.8 and the closure properties of \mathcal{K} , it follows that $\kappa \in \mathcal{K}$. \square

As we have just seen, every \mathcal{R} -variety of algebras is associated with a variety of languages and every variety of languages is associated with an \mathcal{R} -variety of algebras. We conclude this section by proving that this correspondence is bijective. As usual we start with a lemma.

Lemma 5.10. *Let $\mathfrak{A} \in \mathcal{R}$ be an algebra such that every language recognised by \mathfrak{A} has a minimal algebra in \mathcal{R} . Then \mathfrak{A} belongs to an \mathcal{R} -variety \mathcal{V} if, and only if, $\text{Syn}(\kappa) \in \mathcal{V}$, for every language κ recognised by \mathfrak{A} .*

Proof. (\Rightarrow) If κ is recognised by $\mathfrak{A} \in \mathcal{V}$, it follows by Lemma 5.5 that $\text{Syn}(\kappa) \in \mathcal{V}$.

(\Leftarrow) Suppose that $\text{Syn}(\kappa) \in \mathcal{V}$, for every language κ recognised by \mathfrak{A} . As $\mathfrak{A} \in \mathcal{A}$ is \mathcal{T} -generated, there exists an E -morphism $\varepsilon : \mathfrak{T} \rightarrow \mathfrak{A}$ with $\mathfrak{T} \in \mathcal{T}$. Since \mathcal{V} is closed under sort-accumulation points, it is sufficient to show that, for every finite $\Delta \subseteq \Xi$, there is some E -morphism $\mathfrak{B}|_\Delta \rightarrow \mathfrak{A}|_\Delta$ with $\mathfrak{B} \in \mathcal{V}$.

Hence, fix $\Delta \subseteq \Xi$. Since \mathfrak{T} is finitely generated, so is \mathfrak{A} . By enlarging Δ , if necessary, we may therefore assume that $\mathfrak{A}|_\Delta$ generates \mathfrak{A} . Furthermore,

by Lemma 2.5 we can find finite sets H_ξ generating $\mathbb{L}_\xi A$, for $\xi \in \mathcal{E}$. For every $h \in H_\xi$ with $\xi \in \Delta$, we consider the language $\kappa_h := \mathbb{L}\varepsilon(h)$ and the morphism

$$\rho := \langle \text{syn}_{\kappa_h} \rangle_h : \mathfrak{Z} \rightarrow \prod_{\xi \in \Delta} \prod_{h \in H_\xi} \text{Syn}(\kappa_h).$$

Let $\mathfrak{Z} \xrightarrow{\eta} \mathfrak{B} \xrightarrow{\mu} \prod_h \text{Syn}(\kappa_h)$ be the EM -factorisation of ρ . Since \mathcal{V} is closed under \mathcal{T} -generated M -subalgebras of finite products, we have $\mathfrak{B} \in \mathcal{V}$.

For $h \in H_\xi$ with $\xi \in \Delta$, let $p_h : \prod_{i \in H_\xi} \text{Syn}(\kappa_i) \rightarrow \text{Syn}(\kappa_h)$ be the projection. Then we have

$$\begin{aligned} \mathbb{L}\varepsilon(h) &= \kappa_h = \mathbb{L}\text{syn}_{\kappa_h}(\text{res}_{\kappa_h}) \\ &= \mathbb{L}(p_h \circ \mu \circ \eta)(\text{res}_{\kappa_h}) = \mathbb{L}\eta(\mathbb{L}(p_h \circ \mu)(\text{res}_{\kappa_h})), \end{aligned}$$

where $\mathbb{L}(p_h \circ \mu)(\text{res}_{\kappa_h}) \in \mathbb{L}_\xi B$. Consequently,

$$\mathbb{L}\varepsilon[\mathbb{L}_\xi A] = \mathbb{L}\varepsilon[\langle H_\xi \rangle_{\mathbb{L}}] = \langle \mathbb{L}\varepsilon[H_\xi] \rangle_{\mathbb{L}} \subseteq \mathbb{L}\eta[\mathbb{L}_\xi B].$$

Since A_ξ is strongly finite-dimensional, we can therefore use Lemma 2.7 to find a function $q_\xi : B_\xi \rightarrow A_\xi$ with $\varepsilon_\xi = q_\xi \circ \eta_\xi$. Combining these into a single function $q : B|_\Delta \rightarrow A|_\Delta$ we obtain $\varepsilon|_\Delta = q \circ \eta|_\Delta$. By Lemma I.5.6, this implies that q is in fact a morphism $\mathfrak{B}|_\Delta \rightarrow \mathfrak{A}|_\Delta$ of $\mathbb{M}|_\Delta$ -algebras. Furthermore, it follows by Lemma II.1.3 that $q \in E$. \square

Theorem 5.11 (Variety Theorem). *Let \mathcal{V} be an \mathcal{R} -variety of \mathbb{M} -algebras such that every language recognised by an algebra in \mathcal{V} has a minimal algebra that also belongs to \mathcal{R} , and let \mathcal{K} be a variety of Ω -languages such that every language in \mathcal{K} has a minimal algebra that belongs to \mathcal{R} . The following statements are equivalent.*

- (1) \mathcal{K} consists of those languages that are recognised by some algebra in \mathcal{V} .
- (2) \mathcal{K} consists of all languages κ with $\text{Syn}(\kappa) \in \mathcal{V}$.
- (3) \mathcal{V} consists of those algebras that only recognise languages in \mathcal{K} .
- (4) \mathcal{V} is the \mathcal{R} -variety generated by the set $\{\text{Syn}(\kappa) \mid \kappa \in \mathcal{K}\}$.

Proof. (1) \Leftrightarrow (2) follows by Lemma 5.5, and (4) \Rightarrow (1) by Theorem 5.9.

(2) \Rightarrow (3) If $\mathfrak{A} \in \mathcal{V}$ and κ is recognised by \mathfrak{A} , it follows by Lemma 5.10 that $\text{Syn}(\kappa) \in \mathcal{V}$. By (2), this implies that $\kappa \in \mathcal{K}$. Conversely, if \mathfrak{A} only recognises languages in \mathcal{K} , (2) implies that $\text{Syn}(\kappa) \in \mathcal{V}$ for all languages κ recognised by \mathfrak{A} . By Lemma 5.10 it follows that $\mathfrak{A} \in \mathcal{V}$.

(3) \Rightarrow (4) Let \mathcal{V}_\circ be the \mathcal{R} -variety generated by $\{\text{Syn}(\kappa) \mid \kappa \in \mathcal{K}\}$. First, suppose that $\mathfrak{A} \in \mathcal{V}$. Then every language recognised by \mathfrak{A} belongs to \mathcal{K} and Lemma 5.10 implies that $\mathfrak{A} \in \mathcal{V}_\circ$. Hence, $\mathcal{V} \subseteq \mathcal{V}_\circ$.

For the converse inclusion it is sufficient to prove that $\text{Syn}(\kappa) \in \mathcal{V}$, for all $\kappa \in \mathcal{K}$. Hence, fix a language $\kappa \in \mathcal{K}[\mathfrak{Z}]$. It follows by Proposition 4.16 that every language $\lambda \in \mathbb{L}\mathfrak{S}$ recognised by $\text{Syn}(\kappa)$ is of the form

$$\lambda = \mathbb{L}\varphi(\mu), \quad \text{for some } \mu \in \langle \kappa \rangle_{\mathbb{L}} \text{ and } \varphi : \mathfrak{S} \rightarrow \mathfrak{Z}.$$

Since $\mathcal{K}[\mathfrak{Z}]$ is a language algebra, we have $\langle \kappa \rangle_{\mathbb{L}} \subseteq \mathcal{K}[\mathfrak{Z}]$. Consequently,

$$\lambda \in \mathbb{L}\varphi[\langle \kappa \rangle_{\mathbb{L}}] \subseteq \mathbb{L}\varphi[\mathcal{K}[\mathfrak{Z}]] \subseteq \mathcal{K}[\mathfrak{S}].$$

We have shown that all languages recognised by $\text{Syn}(\kappa)$ belong to \mathcal{K} . By assumption, this implies that $\text{Syn}(\kappa) \in \mathcal{V}$. \square

6. The Profinitary Term Monad

The goal of this section and the next one is to derive an axiomatisation of pseudo-varieties in terms of systems of inequalities. We start by defining the kind of terms allowed in our axioms. The actual axiomatisation will then be presented in Section 7 below. A natural choice for the terms would be to take the elements of $\mathbb{M}X$, for some set X of ‘variables’. But it turns out that this does not work. To capture the restriction to \mathbb{M} -algebras from the class \mathcal{R} , we have to use a more general notion of a term. The classic result by Reiterman characterises the pseudo-varieties of finite semigroups as exactly those axiomatisable by a set of *profinite* equations. Analogously, we have to define *pro- \mathcal{R} \mathbb{M} -terms* for our version of this theorem. While the general definition below works for an arbitrary category \mathcal{D} , the subsequent development is based on duality arguments which are specific to the underlying category. We will therefore work mostly in $\mathcal{D} = \text{Pos}$.

The Monad $\widehat{\mathbb{M}}_{\mathcal{A}}$

In this section we will make the following additional assumptions.

Conventions.

- $\mathcal{T} := \{ \mathbb{M} \mathbb{J} X \mid X \text{ finite} \}$
- \mathcal{R} is a class of \mathbb{M} -algebras that is closed under \mathcal{T} -generated \mathbb{M} -subalgebras of finite products and such that, up to isomorphism, \mathcal{R} forms a set.

To explain how we arrive at the definition below, let us collect our requirements on this set of terms. We are looking for a functor $\widehat{\mathbb{M}}$ mapping an (unordered) set X of ‘variables’ to some set $\widehat{\mathbb{M}}X$ of ‘terms’. These terms should generalise the ordinary terms from $\mathbb{M}X$, i.e., we need an embedding $\iota : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}X$. Furthermore, we should be able to ‘evaluate’ a term $t \in \widehat{\mathbb{M}}X$ in a given \mathbb{M} -algebra $\mathfrak{A} \in \mathcal{R}$ with respect to a given ‘variable assignment’ $\beta : X \rightarrow \mathfrak{A}$. Let us denote the resulting value by $\text{val}(t; \beta)$. For ordinary terms $t \in \mathbb{M}X$, this value should of course correspond to the value of t in \mathfrak{A} . Thus,

$$\text{val}(\iota(t); \beta) = \pi(\mathbb{M}\beta(t)),$$

where $\pi(\mathbb{M}\beta(t))$ is the canonical extension of $\beta : X \rightarrow \mathfrak{A}$ to $\mathbb{M}X \rightarrow \mathfrak{A}$. Furthermore, $\text{val}(t; \beta)$ should be compatible with morphisms of \mathbb{M} -algebras. That is,

$$\text{val}(t; \varphi \circ \beta) = \varphi(\text{val}(t; \beta)), \quad \text{for every morphism } \varphi : \mathfrak{A} \rightarrow \mathfrak{B}.$$

This leads to the following construction. We work in the category of all morphisms $\mathbb{M}X \rightarrow \mathfrak{A}$. In this category we consider the diagram of all $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ where \mathfrak{A} belongs to a given class \mathcal{R} of ‘recognisers’. Then we take for $\iota : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}X$ the limit. The morphisms $\widehat{\mathbb{M}}X \rightarrow \mathfrak{A}$ of the corresponding limiting cone can then be taken as our evaluation maps. The formal construction is as follows.

Definition 6.1. Let $\mathcal{R} \subseteq \text{Alg}(\mathbb{M})$ be a subcategory of \mathbb{M} -algebras and $X \in \mathcal{D}$. We denote the comma category $(\mathbb{M}X \downarrow \text{Alg}(\mathbb{M}))$ by \mathcal{C} , the subcategory $(\mathbb{M}X \downarrow \mathcal{R})$ by \mathcal{C}_o , and the inclusion diagram by $D : \mathcal{C}_o \rightarrow \mathcal{C}$.

(a) We denote by $\iota_{\mathcal{R}} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}X$ the limit $\iota_{\mathcal{R}} := \lim D$ of D , and the limiting cone by $(\text{val}_{\mathcal{R}}(-; \beta))_{\beta \in \mathcal{C}_o}$. If \mathcal{R} is the category of all finitary \mathbb{M} -algebras, we drop the subscript and simply write $\widehat{\mathbb{M}}$, ι , and $\text{val}(-; \beta)$.

(b) We turn $\widehat{\mathbb{M}}_{\mathcal{R}}$ into a functor as follows. Given $f : X \rightarrow Y$, the family $(\text{val}(-; \beta \circ \mathbb{M}f))_{\beta}$ (where β ranges over all morphisms $\beta : \mathbb{M}Y \rightarrow \mathfrak{A} \in \mathcal{R}$) forms a cone from $\widehat{\mathbb{M}}X$ to D . As the cone $(\text{val}(-; \beta))_{\beta}$ is limiting, there exists a unique function $f' : \widehat{\mathbb{M}}X \rightarrow \widehat{\mathbb{M}}Y$ such that

$$\text{val}(-; \beta \circ \mathbb{M}f) = \text{val}(-; \beta) \circ f', \quad \text{for all } \beta : \mathbb{M}Y \rightarrow \mathfrak{A} \in \mathcal{R}.$$

We set $\widehat{\mathbb{M}}f := f'$.

Remark. A more concise way to define $\widehat{\mathbb{M}}$ is as the so-called ‘codensity monad’ associated with the forgetful functor $\mathbb{U} : \mathcal{R} \rightarrow \mathcal{D}^{\mathbb{E}}$ mapping an \mathbb{M} -algebra to its universe. By definition, this monad is the right Kan extension of \mathbb{U} along itself. Unravelling all the definitions leads to the explicit definition above.

Let us start by checking that $\widehat{\mathbb{M}}_{\mathcal{R}}$ is well-defined and reasonably behaved.

Lemma 6.2. If \mathcal{D} is complete, the limit $\iota_{\mathcal{R}} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}X$ exists.

Proof. By assumption, \mathcal{D} has arbitrary limits. It follows by Proposition I.5.7 that so has $\text{Alg}(\mathbb{M})$. Now, let $D : \mathcal{C}_o \rightarrow \mathcal{C}$ be the diagram defining $\iota_{\mathcal{R}} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}X$ and let $U : \mathcal{C} \rightarrow \text{Alg}(\mathbb{M})$ be the forgetful functor mapping $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ to the codomain \mathfrak{A} . Since \mathcal{R} forms a set (up to isomorphism), so does the index category \mathcal{C}_o . As $\text{Alg}(\mathbb{M})$ is complete, $U \circ D$ therefore has a limit \mathfrak{Z} . Let $(\lambda_{\beta})_{\beta}$ be the corresponding limiting cone. As $(\beta)_{\beta}$ forms a cone from $\mathbb{M}X$ to $U \circ D$, we obtain a unique morphism $\varphi : \mathbb{M}X \rightarrow \mathfrak{Z}$ such that $\lambda_{\beta} \circ \varphi = \beta$, for all β . It is now straightforward to check that $\varphi : \mathbb{M}X \rightarrow \mathfrak{Z}$ is the limit of D and $(\lambda_{\beta})_{\beta}$ is the corresponding limiting cone. \square

We collect a few basic facts about the evaluation morphisms that will be useful in the proofs below.

Lemma 6.3. Let \mathcal{R} be a class of \mathbb{M} -algebras such that $\iota_{\mathcal{R}}$ exists, $\mathfrak{A}, \mathfrak{B} \in \mathcal{R}$ algebras, $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$, $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$, and $f : Y \rightarrow X$ morphisms, and $s, t \in \widehat{\mathbb{M}}_{\mathcal{R}}X$.

- (a) $\text{val}_{\mathcal{R}}(-; \beta) \circ \iota_{\mathcal{R}} = \beta$
- (b) $\varphi \circ \text{val}_{\mathcal{R}}(-; \beta) = \text{val}_{\mathcal{R}}(-; \varphi \circ \beta)$
- (c) $\text{val}_{\mathcal{R}}(-; \beta) \circ \widehat{\mathbb{M}}_{\mathcal{R}}f = \text{val}_{\mathcal{R}}(-; \beta \circ \mathbb{M}f)$
- (d) $\langle \text{val}_{\mathcal{R}}(-; \beta) \rangle_{\beta} = \text{id}$.
- (e) If X is finite and \mathcal{R} is closed under \mathcal{T} -generated M -subalgebras then, for every $\hat{s} \in \widehat{\mathbb{M}}_{\mathcal{R}}X$, there is some $s \in \mathbb{M}X$ with $\text{val}_{\mathcal{R}}(\hat{s}; \beta) = \beta(s)$.

Proof. (a) By the definition of a cone, $\text{val}_{\mathcal{R}}(-; \beta)$ is a morphism from $\iota_{\mathcal{R}} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}X$ to $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$. This is equivalent to (a).

(b) In the comma category, $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ corresponds to a morphism from $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ to $\varphi \circ \beta : \mathbb{M}X \rightarrow \mathfrak{B}$. Hence, (b) holds again by definition of a cone.

(c) holds by definition of $\widehat{\mathbb{M}}_{\mathcal{R}}f$.

(d) One explicit way to define the limit $\widehat{\mathbb{M}}_{\mathcal{R}}X$ is to take all sequences $(a_{\beta})_{\beta}$ indexed by morphisms $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ satisfying

$$a_{\gamma} = \varphi(a_{\beta}), \quad \text{for all } \varphi : \mathfrak{A} \rightarrow \mathfrak{B} \text{ with } \gamma = \varphi \circ \beta.$$

Then the function $\text{val}_{\mathcal{R}}(-; \beta)$ is simply the projection to the component a_{β} . Consequently,

$$\langle \text{val}_{\mathcal{R}}(-; \beta) \rangle_{\beta} = \text{id}.$$

(e) Let $\beta = i \circ \beta_0$ be the EM -factorisation of β and let \mathfrak{A}_0 be the codomain of β_0 . Note that $\mathfrak{A}_0 \in \mathcal{R}$ since \mathcal{R} is closed under \mathcal{T} -generated M -subalgebras. Fix $\hat{s} \in \widehat{\mathbb{M}}_{\mathcal{R}}X$. By (a), we have $\text{rng } \text{val}_{\mathcal{R}}(-; \beta_0) \supseteq \text{rng } \beta_0$ which, by surjectivity of β_0 , implies that the two ranges are in fact equal. Hence, there is some $s \in \mathbb{M}X$ with $\beta_0(s) = \text{val}_{\mathcal{R}}(\hat{s}; \beta_0)$. By (b), it follows that

$$\beta(s) = i(\beta_0(s)) = i(\text{val}_{\mathcal{R}}(\hat{s}; \beta_0)) = \text{val}_{\mathcal{R}}(\hat{s}; i \circ \beta_0) = \text{val}_{\mathcal{R}}(\hat{s}; \beta). \quad \square$$

Corollary 6.4. Let X be a set and $f, g : C \rightarrow \widehat{\mathbb{M}}X$ functions.

$$f = g \quad \text{iff} \quad \text{val}(-; \beta) \circ f = \text{val}(-; \beta) \circ g, \quad \text{for all } \beta : \mathbb{M}X \rightarrow \mathfrak{A}.$$

Proof. This statement holds generally for all limits. For our special case, we can give a simple proof using Lemma 6.3 (d). By this lemma it follows that

$$\begin{aligned} f = g & \quad \text{iff} \quad \langle \text{val}_{\mathcal{R}}(-; \beta) \rangle_{\beta} \circ f = \langle \text{val}_{\mathcal{R}}(-; \beta) \rangle_{\beta} \circ g \\ & \quad \text{iff} \quad \langle \text{val}_{\mathcal{R}}(-; \beta) \circ f \rangle_{\beta} = \langle \text{val}_{\mathcal{R}}(-; \beta) \circ g \rangle_{\beta}. \quad \square \end{aligned}$$

Proposition 6.5. Let \mathcal{R} be a class such that $\iota_{\mathcal{R}}$ is defined. $\widehat{\mathbb{M}}_{\mathcal{R}}$ forms a monad and $\iota_{\mathcal{R}} : \mathbb{M} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}$ a natural transformation. The unit map of $\widehat{\mathbb{M}}_{\mathcal{R}}$ is $\varepsilon := \iota_{\mathcal{R}} \circ \text{sing}$ and the multiplication $\mu : \widehat{\mathbb{M}}_{\mathcal{R}} \circ \widehat{\mathbb{M}}_{\mathcal{R}} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}$ is uniquely determined by the equations

$$\text{val}(-; \beta) \circ \mu = \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)), \quad \text{for all } \beta.$$

Proof. To simplify notation, let us drop the subscript \mathcal{R} . To see that $\widehat{\mathbb{M}}$ is a functor, note that the uniqueness of the function f' in the definition of $\widehat{\mathbb{M}}f$ implies that $\widehat{\mathbb{M}}(f \circ g) = \widehat{\mathbb{M}}f \circ \widehat{\mathbb{M}}g$.

To show that ι is a natural transformation, consider a function $f : X \rightarrow Y$. For every $\beta : \mathbb{M}Y \rightarrow \mathfrak{A} \in \mathcal{R}$, Lemma 6.3 (c) implies that

$$\begin{aligned} \text{val}(-; \beta) \circ \widehat{\mathbb{M}}f \circ \iota &= \text{val}(-; \beta \circ \mathbb{M}f) \circ \iota \\ &= \beta \circ \mathbb{M}f = \text{val}(-; \beta) \circ \iota \circ \mathbb{M}f. \end{aligned}$$

Consequently, it follows by Corollary 6.4 that $\widehat{\mathbb{M}}f \circ \iota = \iota \circ \mathbb{M}f$.

We define the multiplication $\mu : \widehat{\mathbb{M}} \circ \widehat{\mathbb{M}} \Rightarrow \widehat{\mathbb{M}}$ as follows. For every morphism $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ with $\mathfrak{A} \in \mathcal{R}$, we have

$$\begin{aligned} \beta &= \beta \circ \pi \circ \text{sing} \\ &= \pi \circ \mathbb{M}\beta \circ \text{sing} \\ &= \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \mathbb{M}\iota \circ \text{sing} \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \iota \circ \mathbb{M}\iota \circ \text{sing}. \end{aligned}$$

Furthermore, for two such morphisms $\alpha : \mathbb{M}X \rightarrow \mathfrak{A}$ and $\beta : \mathbb{M}X \rightarrow \mathfrak{B}$ and a morphism $\varphi : \mathfrak{A} \rightarrow \mathfrak{B}$ with $\beta = \varphi \circ \alpha$, we have

$$\begin{aligned} \varphi \circ \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \alpha)) &= \text{val}(-; \varphi \circ \pi \circ \mathbb{M}\text{val}(-; \alpha)) \\ &= \text{val}(-; \pi \circ \mathbb{M}\varphi \circ \mathbb{M}\text{val}(-; \alpha)) \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \varphi \circ \alpha)) \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)). \end{aligned}$$

Consequently, the morphisms $(\text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)))_\beta$ form a cone from

$$\iota \circ \mathbb{M}\iota \circ \text{sing} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}\mathbb{M}X}$$

to the diagram $(\mathbb{M}X \downarrow \mathcal{R})$. As $\iota : \mathbb{M}X \rightarrow \widehat{\mathbb{M}X}$ is the limit of this cone, there exists a unique map $\mu : \widehat{\mathbb{M}\mathbb{M}X} \rightarrow \widehat{\mathbb{M}X}$ such that

$$\mu \circ \iota \circ \mathbb{M}\iota \circ \text{sing} = \iota$$

and $\text{val}(-; \beta) \circ \mu = \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta))$, for all β .

Note that the first of these equations follows from the second one since, for every β ,

$$\begin{aligned} &\text{val}(-; \beta) \circ \mu \circ \iota \circ \mathbb{M}\iota \circ \text{sing} \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \iota \circ \mathbb{M}\iota \circ \text{sing} \\ &= \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \mathbb{M}\iota \circ \text{sing} \\ &= \pi \circ \mathbb{M}\beta \circ \text{sing} \\ &= \beta \circ \pi \circ \text{sing} \\ &= \beta \\ &= \text{val}(-; \beta) \circ \iota, \end{aligned}$$

which, by Corollary 6.4, implies that $\mu \circ \iota \circ \mathbb{M}\iota \circ \text{sing} = \iota$.

Let us start by showing that these morphisms μ form a natural transformation. Hence, fix a function $f : X \rightarrow Y$. For every $\beta : \mathbb{M}Y \rightarrow \mathfrak{A}$, we

have

$$\begin{aligned} \text{val}(-; \beta) \circ \mu \circ \widehat{\mathbb{M}\mathbb{M}}f &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \widehat{\mathbb{M}\mathbb{M}}f \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \mathbb{M}\widehat{\mathbb{M}}f) \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta \circ \mathbb{M}f)) \\ &= \text{val}(-; \beta \circ \mathbb{M}f) \circ \mu \\ &= \text{val}(-; \beta) \circ \widehat{\mathbb{M}}f \circ \mu. \end{aligned}$$

By Corollary 6.4, this implies that $\mu \circ \widehat{\mathbb{M}\mathbb{M}}f = \widehat{\mathbb{M}}f \circ \mu$.

The fact that $\varepsilon := \iota \circ \text{sing}$ is a natural transformation follows immediately from the facts that ι and sing are natural transformations. It therefore remains to check the three axioms of a monad. For every $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$, we have

$$\begin{aligned} \text{val}(-; \beta) \circ \mu \circ \varepsilon &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \iota \circ \text{sing} \\ &= \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \text{sing} \\ &= \text{val}(-; \beta) \circ \pi \circ \text{sing} \\ &= \text{val}(-; \beta), \\ \text{val}(-; \beta) \circ \mu \circ \widehat{\mathbb{M}}\varepsilon &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \widehat{\mathbb{M}}\varepsilon \\ &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \mathbb{M}\varepsilon) \\ &= \text{val}(-; \pi \circ \mathbb{M}(\text{val}(-; \beta) \circ \iota \circ \text{sing})) \\ &= \text{val}(-; \pi \circ \mathbb{M}(\beta \circ \text{sing})) \\ &= \text{val}(-; \beta \circ \pi \circ \mathbb{M}\text{sing}) \\ &= \text{val}(-; \beta), \end{aligned}$$

$$\begin{aligned}
\text{and } \text{val}(-; \beta) \circ \mu \circ \widehat{\mathbb{M}}\mu &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \widehat{\mathbb{M}}\mu \\
&= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta) \circ \mathbb{M}\mu) \\
&= \text{val}(-; \pi \circ \mathbb{M}(\text{val}(-; \beta) \circ \mu)) \\
&= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta))) \\
&= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \mu \\
&= \text{val}(-; \beta) \circ \mu \circ \mu.
\end{aligned}$$

By Corollary 6.4, this implies that

$$\mu \circ \varepsilon = \text{id}, \quad \mu \circ \widehat{\mathbb{M}}\varepsilon = \text{id}, \quad \text{and} \quad \mu \circ \widehat{\mathbb{M}}\mu = \mu \circ \mu. \quad \square$$

The next lemma states that, without loss of generality, we may assume that the morphisms $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ are all surjective. This will be convenient in some situations.

Lemma 6.6. *Let X be a finite set and \mathcal{R} a class of \mathbb{M} -algebras that is closed under \mathcal{T} -generated M -subalgebras of finite products.*

- (a) $\mathcal{C}_o = (\mathbb{M}X \downarrow \mathcal{R})$ is cofiltered.
- (b) In the definition of $\widehat{\mathbb{M}}_{\mathcal{R}}X$, we can restrict the category \mathcal{C}_o to the surjective morphisms without changing the result.

Proof. (a) There are two axioms to check. First, let $\alpha : \mathbb{M}X \rightarrow \mathfrak{A}$ and $\beta : \mathbb{M}X \rightarrow \mathfrak{B}$ be two objects of \mathcal{C}_o . We have to find some $\gamma : \mathbb{M}X \rightarrow \mathfrak{C}$ and morphisms $\varphi : \gamma \rightarrow \alpha$ and $\psi : \gamma \rightarrow \beta$. Set $\gamma := \langle \alpha, \beta \rangle : \mathbb{M}X \rightarrow \mathfrak{A} \times \mathfrak{B}$ and let \mathfrak{C} be the codomain of $\ker \gamma$. Then \mathfrak{C} is a \mathcal{T} -generated M -subalgebra of $\mathfrak{A} \times \mathfrak{B}$. Hence, $\mathfrak{C} \in \mathcal{R}$, $\text{im } \gamma \in \mathcal{C}_o$, and we have morphisms $p : \text{im } \gamma \rightarrow \alpha$ and $q : \text{im } \gamma \rightarrow \beta$, where $p : C \rightarrow A$ and $q : C \rightarrow B$ are the two projections.

For the second axiom, consider two morphisms $\varphi, \psi : \alpha \rightarrow \beta$ with $\alpha : \mathbb{M}X \rightarrow \mathfrak{A}$ and $\beta : \mathbb{M}X \rightarrow \mathfrak{B}$ in \mathcal{C}_o . The set

$$C := \{a \in A \mid \varphi(a) = \psi(a)\}$$

induces a subalgebra of \mathfrak{A} since, for $s \in \mathbb{M}C$, we have

$$\varphi(\pi(s)) = \pi(\mathbb{M}\varphi(s)) = \pi(\mathbb{M}\psi(s)) = \psi(\pi(s)).$$

For $x \in X$, we have

$$\varphi(\alpha(x)) = \beta(x) = \psi(\alpha(x)),$$

which implies that $\alpha[X] \subseteq C$. Hence, $\text{rng } \alpha \subseteq C$. Let $\alpha = i \circ \alpha_o$ be the EM -factorisation of α and let \mathfrak{D} be the codomain of α_o . Then \mathfrak{D} is a \mathcal{T} -generated M -subalgebra of \mathfrak{A} , which implies that $\mathfrak{D} \in \mathcal{R}$. Furthermore, $i : \alpha_o \rightarrow \alpha$ satisfies $\varphi \circ i = \psi \circ i$.

(b) Let \mathcal{C}_{oo} be the full subcategory of $\mathcal{C}_o = (\mathbb{M}X \downarrow \mathcal{R})$ consisting of all morphisms $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ that are surjective. It is sufficient to show that the inclusion $\mathcal{C}_{oo} \rightarrow \mathcal{C}_o$ is final, which means, we have to establish the following two properties.

- (i) Every $\beta \in \mathcal{C}_o$ factorises through some $\beta_o \in \mathcal{C}_{oo}$.
- (ii) For all $\alpha, \alpha' \in \mathcal{C}_{oo}$, $\beta \in \mathcal{C}_o$, and all morphisms $\varphi : \alpha \rightarrow \beta$ and $\varphi' : \alpha' \rightarrow \beta$, there is some $\gamma \in \mathcal{C}_{oo}$ with morphisms $\psi : \gamma \rightarrow \alpha$ and $\psi' : \gamma \rightarrow \alpha'$ such that $\varphi \circ \psi = \varphi' \circ \psi'$.

(i) Given $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$, let \mathfrak{A}_o be the subalgebra of \mathfrak{A} induced by $\text{rng } \beta$, let $i : \mathfrak{A}_o \rightarrow \mathfrak{A}$ be the inclusion function, and $\beta_o : \mathbb{M}X \rightarrow \mathfrak{A}_o$ be the corestriction of β . Then $\beta = i \circ \beta_o$. Since \mathfrak{A}_o is finitely generated by $\beta_o[X]$ and \mathcal{R} is closed under finitely generated subalgebras, we have $\mathfrak{A}_o \in \mathcal{R}$ and $\beta_o \in \mathcal{C}_o$.

(ii) Consider $\alpha : \mathbb{M}X \rightarrow \mathfrak{A}$, $\alpha' : \mathbb{M}X \rightarrow \mathfrak{A}'$ in \mathcal{C}_{oo} , $\beta : \mathbb{M}X \rightarrow \mathfrak{B}$ in \mathcal{C}_o , and $\varphi : \alpha \rightarrow \beta$ and $\varphi' : \alpha' \rightarrow \beta$. Let \mathfrak{C} be the subalgebra of $\mathfrak{A} \times \mathfrak{A}'$ induced by the range of $\gamma := \langle \alpha, \alpha' \rangle : \mathbb{M}X \rightarrow \mathfrak{A} \times \mathfrak{A}'$. As above, it follows that $\mathfrak{C} \in \mathcal{R}$ and $\gamma \in \mathcal{C}_{oo}$. The two projections $p : \mathfrak{C} \rightarrow \mathfrak{A}$ and $p' : \mathfrak{C} \rightarrow \mathfrak{A}'$ are morphisms of \mathcal{C}_{oo} satisfying $\varphi \circ p = \varphi' \circ p'$. \square

Let us show that (in a certain sense) the transformation $\iota_{\mathcal{R}} : \mathbb{M} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}$ is the terminal object of the category of all morphisms of monads $\rho : \mathbb{M} \Rightarrow \mathbb{N}$.

Proposition 6.7. *Let $\rho : \mathbb{M} \Rightarrow \mathbb{N}$ be a morphism of monads, \mathcal{R} a class of \mathbb{N} -algebras, and \mathcal{R}_p the class of their p -reducts. There exists a unique morphism $\varphi : \mathbb{N} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}_p}$ of monads such that $\iota_{\mathcal{R}_p} = \varphi \circ \rho$.*

Proof. Fix a set X . For every morphism $\beta : \mathbb{M}X \rightarrow \mathfrak{A}|_\rho$ with $\mathfrak{A} \in \mathcal{R}$, we define $\mu_\beta : \mathbb{N}X \rightarrow \mathfrak{A}$ by

$$\mu_\beta := \pi \circ \mathbb{N}(\beta \circ \text{sing}) .$$

To see that μ_β is a morphism of \mathbb{N} -algebras, note that

$$\begin{aligned} \pi \circ \mathbb{N}\mu_\beta &= \pi \circ \mathbb{N}(\pi \circ \mathbb{N}(\beta \circ \text{sing})) \\ &= \pi \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{N}(\beta \circ \text{sing}) \\ &= \pi \circ \text{flat} \circ \mathbb{N}\mathbb{N}(\beta \circ \text{sing}) \\ &= \pi \circ \mathbb{N}(\beta \circ \text{sing}) \circ \text{flat} = \mu_\beta \circ \text{flat} . \end{aligned}$$

We claim that $(\mu_\beta)_\beta$ is a cone from $\rho : \mathbb{M}X \rightarrow \mathbb{N}X$ to $(\mathbb{M}X \downarrow \mathcal{R}_\rho)$. First, we have

$$\begin{aligned} \mu_\beta \circ \rho &= \pi \circ \mathbb{N}(\beta \circ \text{sing}) \circ \rho \\ &= (\pi \circ \rho) \circ \mathbb{M}(\beta \circ \text{sing}) \\ &= \beta \circ \text{flat} \circ \mathbb{M}\text{sing} \\ &= \beta \end{aligned}$$

where the third step follows from the fact that $\pi \circ \rho$ is the product of the \mathbb{M} -algebra $\mathfrak{A}|_\rho$, while $\text{flat} : \mathbb{M}\mathbb{M}X \rightarrow \mathbb{M}X$ is the product of $\mathbb{M}X$. Consequently, μ_β is a morphism $\rho \rightarrow \beta$ of the comma category.

Furthermore, for a morphism $\psi : \beta \rightarrow \gamma$ of the comma category, we have

$$\begin{aligned} \psi \circ \mu_\beta &= \psi \circ \pi \circ \mathbb{N}(\beta \circ \text{sing}) \\ &= \pi \circ \mathbb{N}(\psi \circ \beta \circ \text{sing}) \\ &= \pi \circ \mathbb{N}(\gamma \circ \text{sing}) \\ &= \mu_\gamma , \end{aligned}$$

where the second step follows since the morphism $\psi : \beta \rightarrow \gamma$ is induced by a morphism $\psi : \text{codom}(\beta) \rightarrow \text{codom}(\gamma)$ of \mathbb{N} -algebras.

It follows that $(\mu_\beta)_\beta$ is a cone. As $(\text{val}(-; \beta))_\beta$ is the limiting cone, we obtain a unique morphism $\varphi : \mathbb{N}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}_\rho} X$ satisfying $\varphi \circ \rho = \iota_{\mathcal{R}_\rho}$. We claim that φ is a morphism of monads.

We start by showing that φ is natural in X . Hence, consider a function $f : X \rightarrow Y$. Note that, by definition, φ is the unique function such that

$$\text{val}(-; \beta) \circ \varphi = \mu_\beta .$$

Consequently, for each $\beta : \mathbb{M}X \rightarrow \mathfrak{A}|_\rho$ with $\mathfrak{A} \in \mathcal{R}$, we have

$$\begin{aligned} \text{val}(-; \beta) \circ \varphi \circ \mathbb{N}f &= \mu_\beta \circ \mathbb{N}f \\ &= \pi \circ \mathbb{N}(\beta \circ \text{sing}) \circ \mathbb{N}f \\ &= \pi \circ \mathbb{N}(\beta \circ \text{sing} \circ f) \\ &= \pi \circ \mathbb{N}(\beta \circ \mathbb{M}f \circ \text{sing}) \\ &= \mu_{\beta \circ \mathbb{M}f} \\ &= \text{val}(-; \beta \circ \mathbb{M}f) \circ \varphi = \text{val}(-; \beta) \circ \widehat{\mathbb{M}}_{\mathcal{R}_\rho} f \circ \varphi . \end{aligned}$$

By Corollary 6.4, it follows that $\varphi \circ \mathbb{N}f = \widehat{\mathbb{M}}_{\mathcal{R}_\rho} f \circ \varphi$. Hence, φ is a natural transformation $\mathbb{N} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}_\rho}$.

It remains to check the two axioms for a morphism of monads. For the first equation, let $\varepsilon = \iota_{\mathcal{R}_\rho} \circ \text{sing}$ be the unit map of $\widehat{\mathbb{M}}_{\mathcal{R}_\rho}$. Then

$$\varepsilon = \iota_{\mathcal{R}_\rho} \circ \text{sing} = \varphi \circ \rho \circ \text{sing} = \varphi \circ \text{sing} ,$$

where the last step follows from the fact that ρ is a morphism of monads.

For the second equation, consider a morphism $\beta : \mathbb{M}X \rightarrow \mathfrak{A}|_\rho$ with $\mathfrak{A} \in \mathcal{R}$.

Recall that the multiplication $\mu : \widehat{\mathbb{M}}_{\mathcal{R}_\rho} \circ \widehat{\mathbb{M}}_{\mathcal{R}_\rho} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}_\rho}$ satisfies

$$\text{val}(-; \beta) \circ \mu = \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) .$$

Hence,

$$\begin{aligned} \text{val}(-; \beta) \circ \mu \circ \varphi \circ \mathbb{N}\varphi &= \text{val}(-; \pi \circ \mathbb{M}\text{val}(-; \beta)) \circ \varphi \circ \mathbb{N}\varphi \\ &= \mu_{\pi \circ \mathbb{M}\text{val}(-; \beta)} \circ \mathbb{N}\varphi \\ &= \pi \circ \mathbb{N}(\pi \circ \mathbb{M}\text{val}(-; \beta) \circ \text{sing}) \circ \mathbb{N}\varphi \\ &= \pi \circ \mathbb{N}(\pi \circ \text{sing} \circ \text{val}(-; \beta) \circ \varphi) \end{aligned}$$

$$\begin{aligned}
&= \pi \circ \mathbb{N}\mu_\beta \\
&= \pi \circ \mathbb{N}(\pi \circ \mathbb{N}(\beta \circ \text{sing})) \\
&= \pi \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{N}(\beta \circ \text{sing}) \\
&= \pi \circ \text{flat} \circ \mathbb{N}\mathbb{N}(\beta \circ \text{sing}) \\
&= \pi \circ \mathbb{N}(\beta \circ \text{sing}) \circ \text{flat} \\
&= \mu_\beta \circ \text{flat} \\
&= \text{val}(-; \beta) \circ \varphi \circ \text{flat}.
\end{aligned}$$

By Corollary 6.4, it therefore follows that $\mu \circ \varphi \circ \mathbb{N}\varphi = \varphi \circ \text{flat}$. \square

Priestley Spaces

To continue our investigation of the monad $\widehat{\mathbb{M}}_{\mathcal{R}}$, we require some tools from topology. As these are specific to the underlying category, we will work exclusively with $\mathcal{D} = \text{Pos}$ throughout this section. We start with a variant of Stone duality for ordered topological spaces.

Definition 6.8. (a) A *Priestley space* consists of an ordered set $A \in \text{Pos}$ equipped with a topology that is compact and has the following separation property: for every pair of elements $a, b \in A$ with $a \not\leq b$, there exists a clopen set $C \subseteq A$ which is upwards-closed and contains a , but not b . A *morphism of Priestley spaces* is a function $f : A \rightarrow B$ that is monotone and continuous. We denote the category of all Priestley spaces and their morphisms by PSP .

(b) We denote by Dist the category of all distributive lattices (with top and bottom elements) and all lattice homomorphisms (preserving top and bottom). \lrcorner

Remark. Every Priestley space is a Stone space, i.e., compact, Hausdorff, and totally disconnected. \lrcorner

Theorem 6.9 (Priestley). *The category PSP is equivalent to Dist^{op} .*

To translate between these two categories we can map a Priestley space to the lattices of its upwards-closed clopen subsets, and a distributive lattice to the

set of its prime filters (with a suitable topology). We start our investigation of Priestley spaces by showing how to compute limits in PSP^{Ξ} .

Definition 6.10. (a) Let $(\mu_i)_{i \in I}$ be a cone where $\mu_i : A \rightarrow B_i$ and each B_i is a topological space. The *cone topology* induced by $(\mu_i)_i$ is the topology on A which has a closed subbasis consisting of all sets of the form $\mu_i^{-1}[K]$ with $i \in I$ and $K \subseteq B_i$ closed. If A is the limit of a diagram $D : I \rightarrow \text{Pos}^{\Xi}$ and we do not specify a cone explicitly, we will always consider the cone topology induced by the corresponding limiting cone.

(b) For a functor $\mathbb{M} : \text{Pos}^{\Xi} \rightarrow \text{Pos}^{\Xi}$ for which we have defined a lifting to $\text{PSP}^{\Xi} \rightarrow \text{PSP}^{\Xi}$, we write $\text{PAlg}(\mathbb{M})$ for the category of \mathbb{M} -algebras in PSP^{Ξ} . \lrcorner

Remark. Let X be a finite set and \mathcal{R} an \mathcal{R} -variety. When we equip each $\mathfrak{A} \in \mathcal{R}$ with the discrete topology, we can turn $\mathbb{M}X$ and $\widehat{\mathbb{M}}_{\mathcal{R}}X$ into topological spaces where the topology is induced by the cones $(\beta)_{\beta}$ and $(\text{val}(-; \beta))_{\beta}$, respectively. Then it follows by Lemma 6.3 (e) that the embedding $\iota_{\mathcal{R}} : \mathbb{M}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{R}}X$ is dense with respect to these topologies. In fact, the space $\widehat{\mathbb{M}}_{\mathcal{R}}X$ can be seen as the topological completion of $\mathbb{M}X$. In particular, every element of $\widehat{\mathbb{M}}_{\mathcal{R}}X$ is the limit of a suitable sequence in $\mathbb{M}X$. In the semigroup case, for instance, $\widehat{\mathbb{M}}_{\mathcal{R}}X$ contains the *idempotent power* x^{π} which is the limit of the sequence $(x^{n!})_{n < \omega}$. \lrcorner

Lemma 6.11. *The forgetful functor $\mathbb{U} : \text{PSP}^{\Xi} \rightarrow \text{Pos}^{\Xi}$ reflects limits. More precisely, the limit $\lim D$ of a diagram $D : I \rightarrow \text{PSP}^{\Xi}$ is the space obtained by equipping the set $\lim (\mathbb{U} \circ D)$ with the cone topology.*

Proof. Let $A := \lim D$ and $B := \lim (\mathbb{U} \circ D)$ and let $(\lambda_i)_i$ and $(\mu_i)_i$ be the corresponding limiting cones. We start by showing that the cone topology on B is sort-wise Priestley. Note that B_{ξ} is the subset of $\prod_{i \in I} D_{\xi}(i)$ consisting of all families $(a_i)_i$ such that $a_i = Df(a_k)$, for all I -morphisms $f : k \rightarrow i$. Hence, $B_{\xi} = \bigcap_f H_f$ where

$$H_f := \{ (a_i)_i \in \prod_i D_{\xi}(i) \mid Df(a_k) = a_i \}, \quad \text{for } f : k \rightarrow i.$$

Since, for distinct $a, b \in D_\xi(k)$, we can always find a clopen set C with $a \in C$ and $b \notin C$, we can express H_f as the intersection of all sets of the form

$$(\mu_k^{-1}[(Df)^{-1}[C]] \cap \mu_l^{-1}[C]) \cup (\mu_k^{-1}[(Df)^{-1}[C']] \cap \mu_l^{-1}[C']),$$

where C, C' range over all partitions of $D_\xi(k)$ into two clopen classes. It follows that the sets H_f are all closed. By the Theorem of Tychonoff, the product $\prod_i D_\xi(i)$ is compact. Consequently, $B_\xi = \bigcap_f H_f$ is a closed subset of a compact space and, therefore, also compact.

To show that the topology is Priestley, consider two distinct elements $a \not\leq b$ in B . By the definition of the ordering of a limit in Pos^Ξ , there exists an index $i \in I$ with $\mu_i(a) \not\leq \mu_i(b)$. Therefore we can find a clopen, upwards-closed set $C \subseteq D(i)$ such that $\mu_i(a) \in C$ and $\mu_i(b) \notin C$. The preimage $C' := \mu_i^{-1}[C]$ is clopen in B and satisfies $a \in C'$ and $b \notin C'$. Suppose that C' is not upwards-closed. Then there are elements $c \leq d$ with $c \in C'$ and $d \notin C'$. Consequently, $\mu_i(c) \leq \mu_i(d)$ and $\mu_i(c) \in C$ and $\mu_i(d) \notin C$. This contradicts the fact that C is upwards-closed.

We have shown that B with the cone topology belongs to PSP^Ξ . Since B is the limit in Pos^Ξ , there exists a unique map $f : A \rightarrow B$ (in Pos^Ξ) such that $\lambda_i = \mu_i \circ f$, for all i . Similarly, there exists a unique morphism $g : B \rightarrow A$ of PSP^Ξ such that $\mu_i = \lambda_i \circ g$. We can see that the function f is continuous as follows. Let $C = \mu_i^{-1}[K]$ for a basic closed set $K \subseteq B$. Then $f^{-1}[C] = (\mu_i \circ f)^{-1}[K] = (\lambda_i)^{-1}[K]$. Hence, continuity of λ_i implies that the preimage $f^{-1}[C]$ is closed.

Consequently, we can apply the same universality argument two more times to obtain $f \circ g = \text{id}$ and $g \circ f = \text{id}$. Therefore, B and A with the cone topology are isomorphic as topological spaces. \square

The following version of compactness and its corollary below contain our key topology-based argument.

Proposition 6.12. *Let $E : I \rightarrow \text{PSP}$ be a cofiltered diagram. If all spaces $E(i)$, $i \in I$, are non-empty, so is the limit $\lim E$.*

Proof. Let $\mu_i : \lim E \rightarrow E(i)$, $i \in I$, be the morphisms of the limiting cone. Note that, by the Theorem of Tychonoff, the product $\prod_{i \in I} E_\xi(i)$ is compact and Hausdorff.

For each morphism $f : k \rightarrow l$ of I , consider the set

$$H_f := \{ (a_i)_i \in \prod_i E(i) \mid Ef(a_k) = a_l \}.$$

In the proof of Lemma 6.11, we have shown that the sets H_f are closed. For a contradiction, suppose that $\lim E = \bigcap_f H_f$ is empty. By compactness, we can then find finitely many morphisms f_0, \dots, f_n with

$$H_{f_0} \cap \dots \cap H_{f_n} = \emptyset.$$

Suppose that $f_i : k_i \rightarrow l_i$. Since I is cofiltered, there exists some $m \in I$ and morphisms $g_i : m \rightarrow k_i$ such that

- ♦ $l_i = l_j$ implies $f_i \circ g_i = f_j \circ g_j$,
- ♦ $k_i = k_j$ implies $g_i = g_j$,
- ♦ $k_i = l_j$ implies $g_i = f_j \circ g_j$.

(It follows by induction on n that such morphisms exist.) Fixing some element $a_m \in E(m)$, we set

$$a_{k_i} := Eg_i(a_m) \quad \text{and} \quad a_{l_i} := E(f_i \circ g_i)(a_m).$$

By choice of the g_i , this is well-defined. Let $(b_i)_i \in \prod_i E(i)$ be any family with $b_{k_i} = a_{k_i}$ and $b_{l_i} = a_{l_i}$. Then $(b_i)_i \in H_{f_0} \cap \dots \cap H_{f_n} \neq \emptyset$. A contradiction. \square

The following consequence is what we will need below.

Lemma 6.13. *Let $D : I \rightarrow \text{PSP}^\Xi$ be a cofiltered diagram and $(\mu_i)_i$ a cone from $A \in \text{PSP}^\Xi$ to D where each $\mu_i : A \rightarrow D(i)$ is surjective. The induced morphism $\varphi : A \rightarrow \lim D$ is surjective.*

Proof. Fix an element $c \in \lim D$ and let $(\lambda_i)_i$ be the limiting cone. To show that $\varphi^{-1}(c) \neq \emptyset$, we define a diagram $E : I \rightarrow \text{PSp}$ as follows. For $i \in I$, we set

$$E(i) := \mu_i^{-1}(\lambda_i(c)).$$

As singleton sets are closed and μ_i is continuous, $E(i)$ is a closed subset of A_ξ , where ξ is the sort of c . Consequently, $E(i) \in \text{PSp}$. For a morphism $f : i \rightarrow j$ of I , we let $Ef : E(i) \rightarrow E(j)$ be the inclusion map. This is well-defined, since

$$\begin{aligned} a \in E(i) &\Rightarrow \mu_i(a) = \lambda_i(c) \\ &\Rightarrow \mu_j(a) = Df(\mu_i(a)) = Df(\lambda_i(c)) = \lambda_j(c) \\ &\Rightarrow a \in E(j). \end{aligned}$$

By Proposition 6.12, the limit $B := \lim E \neq \emptyset$ is non-empty. Fix $b \in B$. Note that the inclusion maps $E(i) \rightarrow A_\xi$ constitute a natural transformation $\sigma : E \Rightarrow A_\xi^I$, where A_ξ^I denotes the constant diagram $I \rightarrow \text{PSp}$ with value A_ξ and identity maps everywhere. Let $\psi : B \rightarrow A_\xi$ be the induced map between the correpsoring limits. Since

$$\lambda_i(\varphi(\tau(b))) = \mu_i(\tau(b)) = \lambda_i(c), \quad \text{for all } i \in I,$$

it follows that $\varphi(\tau(b)) = c$, as desired. \square

Our main technical tool in the next section is the following natural transformation relating the functors $\widehat{\mathbb{M}}_{\mathcal{V}}$ and $\widehat{\mathbb{M}}_{\mathcal{R}}$, for different classes \mathcal{V} and \mathcal{R} . The important case below will be where \mathcal{V} is the \mathcal{R} -variety under consideration and \mathcal{R} the class of all finitary \mathbb{M} -algebras.

Theorem 6.14. *Let $\mathcal{V} \subseteq \mathcal{R} \subseteq \text{Alg}(\mathbb{M})$.*

- (a) *There exists a unique morphism $\rho : \widehat{\mathbb{M}}_{\mathcal{R}} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{V}}$ of monads that makes the following diagram commute, for all morphisms $\beta : \mathbb{M}X \rightarrow \mathcal{A}$ where $\mathcal{A} \in \mathcal{V}$ and $X \in \mathcal{D}^\Xi$.*

$$\begin{array}{ccc} \widehat{\mathbb{M}}_{\mathcal{R}}X & \xrightarrow{\rho} & \widehat{\mathbb{M}}_{\mathcal{V}}X \\ \wr_{\mathcal{R}} \swarrow & \text{val}_{\mathcal{R}}(-; \beta) \searrow & \swarrow \text{val}_{\mathcal{V}}(-; \beta) \\ \mathbb{M}_{\mathcal{R}}X & \xrightarrow{\beta} & \mathcal{A} \end{array}$$

- (b) *Suppose that $X = \mathbb{J}Z$, for some finite set Z , that every algebra in \mathcal{R} is finitary, and that \mathcal{V} and \mathcal{R} are both closed under finitely generated subalgebras of finite products. Then the induced morphism $\rho_X : \widehat{\mathbb{M}}_{\mathcal{R}}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{V}}X$ is surjective.*

Proof. (a) Fix a set $X \in \mathcal{D}^\Xi$. The family $(\text{val}_{\mathcal{R}}(-; \beta))_{\beta \in (\mathbb{M}X \downarrow \mathcal{V})}$ forms a cone from $\widehat{\mathbb{M}}_{\mathcal{R}}X$ to the diagram defining $\widehat{\mathbb{M}}_{\mathcal{V}}X$. As $(\text{val}_{\mathcal{V}}(-; \beta))_{\beta \in (\mathbb{M}X \downarrow \mathcal{V})}$ is limiting, there exists a unique map $\rho_X : \widehat{\mathbb{M}}_{\mathcal{R}}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{V}}X$ such that

$$\text{val}_{\mathcal{V}}(-; \beta) \circ \rho_X = \text{val}_{\mathcal{R}}(-; \beta), \quad \text{for all } \beta : \mathbb{M}X \rightarrow \mathcal{A}.$$

As the equation $\text{val}_{\mathcal{V}}(-; \beta) \circ \iota_{\mathcal{V}} = \beta$ was already proved in Lemma 6.3 (a), it therefore remains to prove that the family $\rho := (\rho_X)_X$ forms a morphism of monads. To see that it is a natural transformation, consider a function $f : X \rightarrow Y$. Then

$$\begin{aligned} \text{val}_{\mathcal{R}}(-; \beta) \circ \widehat{\mathbb{M}}_{\mathcal{R}}f \circ \rho &= \text{val}_{\mathcal{R}}(-; \beta \circ \mathbb{M}f) \circ \rho \\ &= \text{val}_{\mathcal{V}}(-; \beta \circ \mathbb{M}f) \\ &= \text{val}_{\mathcal{V}}(-; \beta) \circ \widehat{\mathbb{M}}_{\mathcal{V}}f \\ &= \text{val}_{\mathcal{R}}(-; \beta) \circ \rho \circ \widehat{\mathbb{M}}_{\mathcal{V}}f. \end{aligned}$$

By Corollary 6.4, it follows that $\widehat{\mathbb{M}}_{\mathcal{R}}f \circ \rho = \rho \circ \widehat{\mathbb{M}}_{\mathcal{V}}f$, as desired.

To check the two axioms of a morphism of monads, let $\mu_{\mathcal{V}}$ and $\varepsilon_{\mathcal{V}}$ be the multiplication and unit map of $\widehat{\mathbb{M}}_{\mathcal{V}}$, and $\mu_{\mathcal{R}}$ and $\varepsilon_{\mathcal{R}}$ those of $\widehat{\mathbb{M}}_{\mathcal{R}}$. For every

$\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ with $\mathfrak{A} \in \mathcal{V}$, we have

$$\begin{aligned}
 \text{val}_{\mathcal{V}}(-; \beta) \circ \rho \circ \mu_{\mathcal{R}} &= \text{val}_{\mathcal{R}}(-; \beta) \circ \mu_{\mathcal{R}} \\
 &= \text{val}_{\mathcal{R}}(-; \pi \circ \mathbb{M}\text{val}_{\mathcal{R}}(-; \beta)) \\
 &= \text{val}_{\mathcal{V}}(-; \pi \circ \mathbb{M}\text{val}_{\mathcal{R}}(-; \beta)) \circ \rho \\
 &= \text{val}_{\mathcal{V}}(-; \pi \circ \mathbb{M}\text{val}_{\mathcal{V}}(-; \beta) \circ \mathbb{M}\rho) \circ \rho \\
 &= \text{val}_{\mathcal{V}}(-; \pi \circ \mathbb{M}\text{val}_{\mathcal{V}}(-; \beta)) \circ \widehat{\mathbb{M}}\rho \circ \rho \\
 &= \text{val}_{\mathcal{V}}(-; \beta) \circ \mu_{\mathcal{V}} \circ \widehat{\mathbb{M}}\rho \circ \rho
 \end{aligned}$$

$$\begin{aligned}
 \text{and } \text{val}_{\mathcal{V}}(-; \beta) \circ \rho \circ \varepsilon_{\mathcal{R}} &= \text{val}_{\mathcal{V}}(-; \beta) \circ \rho \circ \iota_{\mathcal{R}} \circ \text{sing} \\
 &= \text{val}_{\mathcal{R}}(-; \beta) \circ \iota_{\mathcal{R}} \circ \text{sing} \\
 &= \beta \circ \text{sing} \\
 &= \text{val}_{\mathcal{V}}(-; \beta) \circ \iota_{\mathcal{V}} \circ \text{sing} = \text{val}_{\mathcal{V}}(-; \beta) \circ \varepsilon_{\mathcal{V}}.
 \end{aligned}$$

By Corollary 6.4, it follows that $\rho \circ \mu_{\mathcal{R}} = \mu_{\mathcal{V}} \circ \widehat{\mathbb{M}}\rho \circ \rho$ and $\rho \circ \varepsilon_{\mathcal{R}} = \varepsilon_{\mathcal{V}}$.

(b) To apply the topological machinery we have just set up, we translate the problem into the category of Priestley spaces. We equip each algebra $\mathfrak{A} \in \mathcal{R}$ with the discrete topology, which is Priestley since \mathfrak{A} is finitary. According to Lemma 6.6 (b), we can define the limits $\widehat{\mathbb{M}}_{\mathcal{V}}X$ and $\widehat{\mathbb{M}}_{\mathcal{R}}X$ in terms of only the surjective morphisms $\beta : \mathbb{M}X \rightarrow \mathfrak{A}$ with \mathfrak{A} in \mathcal{V} or \mathcal{R} . Furthermore, it follows by Lemma 6.11 that $\widehat{\mathbb{M}}_{\mathcal{V}}X$ and $\widehat{\mathbb{M}}_{\mathcal{R}}X$ are also sort-wise Priestley spaces when equipped with the cone topology. In addition, the limits in the category PSP^{Ξ} coincide with $\widehat{\mathbb{M}}_{\mathcal{V}}X$ and $\widehat{\mathbb{M}}_{\mathcal{R}}X$.

Let $\Xi_0 \subseteq \Xi$ be the set of all sorts ξ such that $\mathbb{M}_{\xi}X \neq \emptyset$. By Lemma 6.3 (e), it follows that these are exactly the same sorts ξ with $\widehat{\mathbb{M}}_{\mathcal{V}, \xi}X \neq \emptyset$, $\widehat{\mathbb{M}}_{\mathcal{R}, \xi}X \neq \emptyset$, and with $A_{\xi} \neq \emptyset$, for $\mathfrak{A} \in \mathcal{V}$. Consequently, we can perform the rest of the proof in the category Pos^{Ξ_0} . By the definition of the cone topology, all the maps $\text{val}_{\mathcal{V}}(-; \beta)$ and $\text{val}_{\mathcal{R}}(-; \beta)$ are continuous. Furthermore, since we restricted the diagram to surjective maps β , $\text{val}_{\mathcal{R}}(-; \beta) \circ \iota = \beta$ implies that the value maps $\text{val}_{\mathcal{R}}(-; \beta)$ are also surjective. By Lemma 6.6 (a), $\widehat{\mathbb{M}}_{\mathcal{V}}X$ is a cofiltered limit. Consequently, we can use Lemma 6.13, to show that $\rho : \widehat{\mathbb{M}}_{\mathcal{R}}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{V}}X$ is surjective. \square

Finally Copresentable Algebras

In the remainder of this section we will prove that algebras of the form $\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta}X$ are what is called *finitely copresentable* (at least if X and Δ are finite). This is another result requiring us to work with Priestly spaces. Already the next proposition fails in Set or Pos . Unfortunately, it also does only hold for finitely many sorts.

Definition 6.15. An object A of a category \mathcal{C} is *finitely copresentable* if, for every cofiltered diagram $D : I \rightarrow \mathcal{C}$ with limit C and limiting cone $(\lambda_i)_{i \in I}$, and for every morphism $f : C \rightarrow A$, there exists an index $k \in I$ and an essentially unique morphism $g : D(k) \rightarrow A$ such that $f = g \circ \lambda_k$. Essentially uniqueness here means that, if $g' : D(k) \rightarrow A$ is another morphism with $f = g' \circ \lambda_k$, then there exists an I -morphisms $h : l \rightarrow k$ with $g \circ Dh = g' \circ Dh$. \square

Let us note the following standard fact from category theory.

Lemma 6.16. An object $A \in \mathcal{C}$ is *finitely copresentable* if, and only if, the hom-functor $\mathcal{C}(-, A)$ preserves cofiltered surjective limits.

The following results makes essential use of duality. It is one of the main reasons why we work with Priestly spaces.

Proposition 6.17. Let Ξ be a finite set of sorts. Every finite Priestley space is *finitely copresentable* in PSP^{Ξ} .

Proof. First, note that the duality theorem implies that PSP^{Ξ} is equivalent to $(\text{Dist}^{\Xi})^{\text{op}}$. Furthermore, the corresponding translation maps finite spaces to finite lattices. Consequently it is sufficient to show that every finite lattice is finitely presentable in Dist^{Ξ} .

Hence, let \mathfrak{L} be a Ξ -sorted finite lattice, let $D : I \rightarrow \text{Dist}^{\Xi}$ be a directed diagram with colimit \mathfrak{N} and limiting cocone $(\lambda_i)_{i \in I}$, and let $\varphi : \mathfrak{L} \rightarrow \mathfrak{N}$ be a lattice homomorphism. Since L is finite there exists some index $i \in I$ such that $\text{rng } \varphi \subseteq \text{rng } \lambda_i$. For every $a \in L$, fix an element $f(a) \in \lambda_i^{-1}(\varphi(a))$. This defines a function $f : L \rightarrow D(i)$ with $\lambda_i \circ f = \varphi$. But note that, in general,

f is neither monotone nor a lattice homomorphism. For $a, b \in L$, it follows that

$$\begin{aligned}\lambda_i(f(a) \sqcap f(b)) &= \lambda_i(f(a)) \sqcap \lambda_i(f(b)) \\ &= \varphi(a) \sqcap \varphi(b) \\ &= \varphi(a \sqcap b) \\ &= \lambda_i(f(a \sqcap b)).\end{aligned}$$

By the definition of a colimit, this implies that there is some index $k \geq i$ such that

$$D(i, k)(f(a) \sqcap f(b)) = D(i, k)(f(a \sqcap b)).$$

The same argument provides an index $k \geq i$ with

$$D(i, k)(f(a) \sqcup f(b)) = D(i, k)(f(a \sqcup b)).$$

Since I is directed and there are only finitely many pairs $a, b \in L$, it follows that we can find some index $l \geq i$ such that, for all $a, b \in L$,

$$D(i, l)(f(a) \sqcap f(b)) = D(i, l)(f(a \sqcap b))$$

and $D(i, l)(f(a) \sqcup f(b)) = D(i, l)(f(a \sqcup b)).$

It follows that the function $\mu := D(i, l) \circ f$ is a lattice homomorphism satisfying

$$\lambda_l \circ \mu = \lambda_l \circ D(i, l) \circ f = \lambda_i \circ f = \varphi.$$

It remains to show that μ is essentially unique. Hence, suppose that there is a second homomorphism $\mu' : \mathcal{L} \rightarrow D(I')$ with $\lambda_{l'} \circ \mu' = \varphi$. Fixing some index $l'' \geq l, l'$ and replacing μ and μ' by, respectively, $D(l, l'') \circ \mu$ and $D(l', l'') \circ \mu'$ we may assume that $l = l'$. For every element $a \in L$, the fact that $\lambda_l(\mu(a)) = \varphi(a) = \lambda_l(\mu'(a))$ implies that there is some index $k \geq l$ with $D(l, k)(\mu(a)) = D(l, k)(\mu'(a))$. As I is directed and L is finite, it follows that we can find an index $k \geq l$ such that

$$D(l, k)(\mu(a)) = D(l, k)(\mu'(a)), \quad \text{for all } a \in L.$$

Thus $D(l, k) \circ \mu = D(l, k) \circ \mu'$, as desired. \square

It remains to transfer this result from \mathbf{PSp}^{Ξ} to $\mathbf{PAlg}(\widehat{\mathbb{M}})$.

Proposition 6.18. *Let Ξ be a finite set of sorts and \mathcal{R} a class of finite \mathbb{M} -algebras. The functor $\widehat{\mathbb{M}}_{\mathcal{R}}$ preserves cofiltered limits.*

Proof. We obtain a very concise proof if we employ a bit of category-theoretical machinery. I have tried to present the proof in a way that it should be intelligible without knowledge of the actual definitions of the terms involved.

As already noted above one can define $\widehat{\mathbb{M}}_{\mathcal{R}}$ as the *codensity monad* associated with the forgetful functor $I : \mathcal{R} \rightarrow \mathbf{PSp}^{\Xi}$ mapping an \mathbb{M} -algebra $\mathfrak{A} \in \mathcal{R}$ to its universe A (equipped with the discrete topology). By definition, this means that

$$\widehat{\mathbb{M}}_{\mathcal{R}} = \mathbf{Ran}_I I$$

is the *right Kan extension* of I along itself. Furthermore, we can compute such a Kan extension as

$$(\mathbf{Ran}_I I)(X) = \int_{\mathfrak{A} \in \mathcal{R}} (I\mathfrak{A})^{\mathbf{PSp}^{\Xi}(X, I\mathfrak{A})} = \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathbf{PSp}^{\Xi}(X, A)},$$

where the integral sign is a certain kind of limit for binary functors called an *end*. Note that we have seen in Lemma I.2.9 that there exists a natural isomorphism

$$\mathbf{PSp}^{\Xi}(B, A^X) \cong \mathbf{Set}^{\Xi}(X, \mathbf{PSp}^{\Xi}(B, A)),$$

for sets $X \in \mathbf{Set}^{\Xi}$ and spaces $A, B \in \mathbf{PSp}^{\Xi}$. For a fixed space A , it follows that $A^{(-)} : (\mathbf{Set}^{\Xi})^{\text{op}} \rightarrow \mathbf{PSp}^{\Xi}$ is the right adjoint of the hom-functor $\mathbf{PSp}^{\Xi}(-, A) : \mathbf{PSp}^{\Xi} \rightarrow (\mathbf{Set}^{\Xi})^{\text{op}}$. This in particular implies that $A^{(-)}$ preserves all limits. Furthermore we have seen in Lemma 6.16, that a space $A \in \mathbf{PSp}^{\Xi}$ is finitely copresentable if, and only if, the hom-functor $\mathbf{PSp}^{\Xi}(-, A)$ preserves cofiltered surjective limits. As we have seen in Proposition 6.17 that the universe of a finite \mathbb{M} -algebra is finitely copresentable in \mathbf{PSp}^{Ξ} , it follows that the composition $A^{\mathbf{PSp}^{\Xi}(-, A)}$ preserves cofiltered limits, for every $\mathfrak{A} \in \mathcal{R}$.

Given a cofiltered diagram $D : J \rightarrow \mathbf{PSp}^{\Xi}$, we therefore have


$$\begin{aligned} \widehat{\mathbb{M}}_{\mathcal{R}}(\lim D) &= (\text{Ran}_I I)(\lim D) \\ &= \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathbf{PSp}^{\Xi}(\lim_{j \in J} D(j), A)} \\ &= \int_{\mathfrak{A} \in \mathcal{R}} \lim_{j \in J} A^{\mathbf{PSp}^{\Xi}(D(j), A)} \\ &= \lim_{j \in J} \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathbf{PSp}^{\Xi}(D(j), A)} = \lim_{j \in J} \widehat{\mathbb{M}}_{\mathcal{R}} D(j), \end{aligned}$$

where the fourth step follows by the fact that an end is a limit and limits commute. \square

Lemma 6.19. *Let \mathcal{C} be a category, $\mathbb{M} : \mathcal{C} \rightarrow \mathcal{C}$ a monad preserving cofiltered limits, and \mathfrak{A} an \mathbb{M} -algebra with finitely copresentable domain A . Then \mathfrak{A} is finitely copresentable in $\mathbf{Alg}(\mathbb{M})$.*

Proof. Fix a cofiltered diagram $D : I \rightarrow \mathbf{Alg}(\mathbb{M})$ with limit \mathfrak{C} and a limiting cone $(\lambda_i)_{i \in I}$, and let $\varphi : \mathfrak{C} \rightarrow \mathfrak{A}$ be a morphism of \mathbb{M} -algebras. As the domain A is finitely copresentable, there exists an index $k \in I$ and a \mathcal{C} -morphism $f : D(k) \rightarrow A$ such that $\varphi = f \circ \lambda_k$. Since φ and λ_k are morphisms of \mathbb{M} -algebras, we have

$$f \circ \pi \circ \mathbb{M}\lambda_k = f \circ \lambda_k \circ \pi = \varphi \circ \pi = \pi \circ \mathbb{M}\varphi = \pi \circ \mathbb{M}f \circ \mathbb{M}\lambda_k.$$



As A is finitely copresentable and $\mathbb{M}C$ is the limit of $\mathbb{M} \circ D$ with limiting cone $(\mathbb{M}\lambda_i)_i$, we can find an index $m \in I$ and an essentially unique morphism

$f' : \mathbb{M}D(m) \rightarrow A$ such that

$$\pi \circ \mathbb{M}\varphi = f' \circ \mathbb{M}\lambda_m.$$

Since I is cofiltered, there exists an index $m' \in I$ and I -morphisms $g : m' \rightarrow m$ and $h : m' \rightarrow k$. Replacing m by m' and f' by $f' \circ Dg$, we may therefore assume w.l.o.g. that there exists an I -morphism $h : m \rightarrow k$. By essential uniqueness of f' , we can find a morphism $h' : m' \rightarrow m$

$$f \circ \pi \circ \mathbb{M}Dh \circ \mathbb{M}Dh' = f' \circ \mathbb{M}Dh' = \pi \circ \mathbb{M}f \circ \mathbb{M}Dh \circ \mathbb{M}Dh'.$$

Consequently, $\psi := f \circ D(h \circ h')$ satisfies

$$\begin{aligned} \psi \circ \pi &= f \circ D(h \circ h') \circ \pi \\ &= f \circ \pi \circ \mathbb{M}D(h \circ h') = \pi \circ \mathbb{M}f \circ \mathbb{M}D(h \circ h') = \pi \circ \mathbb{M}\psi, \end{aligned}$$

and $\varphi = f \circ \lambda_k = f \circ Dh \circ \lambda_m = \psi \circ \lambda_m$.

Hence, ψ is the desired morphism of \mathbb{M} -algebras. For essential uniqueness of ψ , note that A is finitely copresentable. Consequently, ψ is even essentially unique as a \mathcal{C} -morphism. \square

Corollary 6.20. *Let $\Delta \subseteq \Xi$ be a finite set of sorts and \mathcal{R} a class of finitary \mathbb{M} -algebras. For every finite set $X \in \mathbf{PSp}^{\Delta}$, the $\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta}$ -algebra $\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta} X$ is finitely copresentable in $\mathbf{PAlg}(\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta})$.*

Proof. By Proposition 6.17, the set X (with the discrete topology) is finitely copresentable in \mathbf{PSp}^{Δ} . As we have shown in Proposition 6.18 that $\widehat{\mathbb{M}}|_{\Delta}$ preserves cofiltered limits, the claim therefore follows by Lemma 6.19. \square

7. Axiomatisations

After the preparations in the previous section we are now able to define the type of inequalities we use to axiomatise \mathcal{R} -varieties and to prove the characterisation theorem. To apply the results of the preceding section, we again work in $\mathcal{D} = \mathbf{Pos}$.

Convention. In this section, we denote by \mathcal{R} the class of all finitary \mathbb{M} -algebras.

Definition 7.1. Let $X \in \text{Set}^\Xi$ be finite and $\mathcal{V} \subseteq \mathcal{R}$.

(a) An \mathbb{M} -inequality over X is a statement of the form $s \leq t$ with $s, t \in \widehat{\mathbb{M}}\mathbb{J}X$.

(b) An algebra $\mathfrak{A} \in \mathcal{R}$ satisfies an \mathbb{M} -inequality $s \leq t$ over X if

$$\text{val}_{\mathcal{R}}(s; \beta) \leq \text{val}_{\mathcal{R}}(t; \beta), \quad \text{for all } \beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A}.$$

We write $\mathfrak{A} \models s \leq t$ to denote this fact.

(c) The \mathbb{M} -theory $\text{Th}(\mathcal{V})$ of \mathcal{V} is the set of all \mathbb{M} -inequalities $s \leq t$ satisfied by every algebra in \mathcal{V} . (We do *not* fix the set X these inequalities are over.)

(d) A set Φ of \mathbb{M} -inequalities (possibly over several different sets X) axiomatises the following subclass of \mathcal{V} .

$$\text{Mod}_{\mathcal{V}}(\Phi) := \{ \mathfrak{A} \in \mathcal{V} \mid \mathfrak{A} \models s \leq t \text{ for all } s \leq t \in \Phi \}.$$

Let us start with the following important property connecting the theory of a class \mathcal{V} to the morphism $\rho_{\mathcal{V}}$ from Theorem 6.14.

Lemma 7.2. Let \mathcal{V} be a class of \mathbb{M} -algebras, X a finite set, and $s \leq t$ an \mathbb{M} -inequality over X . Then

$$s \leq t \in \text{Th}(\mathcal{V}) \quad \text{iff} \quad \rho_{\mathcal{V}}(s) \leq \rho_{\mathcal{V}}(t),$$

where $\rho_{\mathcal{V}} : \widehat{\mathbb{M}} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{V}}$ is the morphism from Theorem 6.14.

Proof. By Lemma 6.3 (d), we have

$$\begin{aligned} & \mathfrak{A} \models s \leq t, & & \text{for all } \mathfrak{A} \in \mathcal{V} \\ \text{iff } & \text{val}_{\mathcal{R}}(s; \beta) \leq \text{val}_{\mathcal{R}}(t; \beta), & & \text{for all } \beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A} \in \mathcal{V} \\ \text{iff } & \text{val}_{\mathcal{V}}(\rho_{\mathcal{V}}(s); \beta) \leq \text{val}_{\mathcal{V}}(\rho_{\mathcal{V}}(t); \beta), & & \text{for all } \beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A} \in \mathcal{V} \\ \text{iff } & \rho_{\mathcal{V}}(s) \leq \rho_{\mathcal{V}}(t). & & \square \end{aligned}$$

The easier direction is to show that every axiomatisable class is an \mathcal{R} -variety.

Proposition 7.3. Let \mathcal{V} be an \mathcal{R} -variety and Φ a set of \mathbb{M} -inequalities. Then $\text{Mod}_{\mathcal{V}}(\Phi)$ is an \mathcal{R} -variety.

Proof. We have to check three closure properties. First, consider a finitary subalgebra \mathfrak{A} of a product $\prod_{i \in I} \mathfrak{B}^i$ with $\mathfrak{B}^i \in \text{Mod}_{\mathcal{V}}(\Phi)$. Let $p_k : \prod_i B^i \rightarrow B^k$ be the projection. For $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A}$ it follows that

$$p_k(\text{val}(s; \beta)) = \text{val}(s; p_k \circ \beta) \leq \text{val}(t; p_k \circ \beta) = p_k(\text{val}(t; \beta)),$$

where the second step follows from the fact that $\mathfrak{B}^k \models s \leq t$. As the ordering of the product is defined component-wise, this implies that $\text{val}(s; \beta) \leq \text{val}(t; \beta)$. Consequently, $\mathfrak{A} \in \text{Mod}_{\mathcal{V}}(\Phi)$.

Next, consider a quotient $q : \mathfrak{B} \rightarrow \mathfrak{A}$ with $\mathfrak{B} \in \text{Mod}_{\mathcal{V}}(\Phi)$. Fix $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A}$. Since q is surjective, we can use Corollary II.5.4 to find some $\gamma : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{B}$ with $\beta = q \circ \gamma$. Then

$$\begin{aligned} \text{val}(s; \beta) &= \text{val}(s; q \circ \gamma) = q(\text{val}(s; \gamma)) \\ &\leq q(\text{val}(t; \gamma)) = \text{val}(t; q \circ \gamma) = \text{val}(t; \beta), \end{aligned}$$

where the third step follows by monotonicity of q and the fact that $\mathfrak{B} \models s \leq t$. Consequently, $\mathfrak{A} \in \text{Mod}_{\mathcal{V}}(\Phi)$.

Finally, suppose that \mathfrak{A} is a sort-accumulation point of $\text{Mod}_{\mathcal{V}}(\Phi)$. Fix $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \rightarrow \mathfrak{A}$. We have to show that

$$\text{val}_{\mathcal{V}}(s; \beta) \leq \text{val}_{\mathcal{V}}(t; \beta).$$

Suppose that $s, t \in \widehat{\mathbb{M}}_{\xi}\mathbb{J}X$ and let $\Delta \subseteq \Xi$ be a finite set of sorts containing ξ and all sorts in X . By assumption, there is some algebra $\mathfrak{B} \in \text{Mod}_{\mathcal{V}}(\Phi)$ and a surjective morphism $\mu : \mathfrak{B}|_{\Delta} \rightarrow \mathfrak{A}|_{\Delta}$. By Corollary II.5.4, we can find a morphism $\gamma : \mathbb{M}|_{\Delta}\mathbb{J}X \rightarrow \mathfrak{B}|_{\Delta}$ with $\beta|_{\Delta} = \mu \circ \gamma$. Since $\mathfrak{B} \models s \leq t$ and $s, t \in \widehat{\mathbb{M}}|_{\Delta}\mathbb{J}X$, we have (working in the category Pos^{Δ})

$$\begin{aligned} \text{val}_{\mathcal{V}}(s; \beta|_{\Delta}) &= \text{val}_{\mathcal{V}}(s; \mu \circ \gamma) \\ &= \mu(\text{val}_{\mathcal{V}}(s; \gamma)) \\ &\leq \mu(\text{val}_{\mathcal{V}}(t; \gamma)) \\ &= \text{val}_{\mathcal{V}}(t; \mu \circ \gamma) = \text{val}_{\mathcal{V}}(t; \beta|_{\Delta}). \end{aligned}$$

Since $\text{val}_{\mathcal{V}}(-; \beta|_{\Delta}) = \text{val}_{\mathcal{V}}(-; \beta) \upharpoonright \widehat{\mathbb{M}}|_{\Delta} \mathbb{J}X$, it follows that $\mathfrak{A} \models s \leq t$. \square

For the converse statement – that every \mathcal{R} -variety is axiomatisable – we start with a proposition.

Proposition 7.4. *Let \mathcal{V} be an \mathcal{R} -variety. Then*

$$\mathcal{V} = \{ \mathfrak{A} \mid \mathfrak{A} \text{ a finitary quotient of } \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X \text{ for some finite set } X \}.$$

Proof. (\subseteq) Let $\mathfrak{A} \in \mathcal{V}$. As \mathfrak{A} is finitely generated, there exists a surjective morphism $\beta : \mathbb{M} \mathbb{J}X \rightarrow \mathfrak{A}$, for some finite set X . The claim follows since $\text{val}(-; \beta) \circ \iota = \beta$ implies that $\text{val}(-; \beta) : \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X \rightarrow \mathfrak{A}$ is also surjective.

(\supseteq) Let \mathfrak{A} be finitary and $\varphi : \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X \rightarrow \mathfrak{A}$ surjective. We have to show that $\mathfrak{A} \in \mathcal{V}$. As \mathcal{V} is closed under sort-accumulation points, it is sufficient to show that, for every finite set $\Delta \subseteq \mathcal{E}$ there is some algebra $\mathfrak{B} \in \mathcal{V}$ and a surjective morphism $\mathfrak{B}|_{\Delta} \rightarrow \mathfrak{A}|_{\Delta}$. Hence, fix $\Delta \subseteq \mathcal{E}$. Note that, according to Lemma 6.11 we can define the set $\widehat{\mathbb{M}}_{\mathcal{V}}|_{\Delta} \mathbb{J}X$ as the limit of a cofiltered diagram in PSP^{Δ} . Furthermore, we have seen in Corollary 6.20 that the $\widehat{\mathbb{M}}_{\mathcal{V}}|_{\Delta}$ -algebra $\widehat{\mathbb{M}}_{\mathcal{V}}|_{\Delta} \mathbb{J}X$ is finitely copresentable in $\text{PAlg}(\widehat{\mathbb{M}}_{\mathcal{V}}|_{\Delta})$. Therefore, there exists an algebra $\mathfrak{B} \in \mathcal{V}$ and morphisms $\beta : \mathbb{M} \mathbb{J}X \rightarrow \mathfrak{B}$ and $\mu : \mathfrak{B}|_{\Delta} \rightarrow \mathfrak{A}|_{\Delta}$ such that $\varphi|_{\Delta} = \mu \circ \text{val}(-; \beta)|_{\Delta}$. Since $\varphi|_{\Delta}$ is surjective, so is μ . Consequently, $\mathfrak{A}|_{\Delta}$ is a quotient of $\mathfrak{B}|_{\Delta}$ and $\mathfrak{B} \in \mathcal{V}$. \square

Corollary 7.5. *Let \mathcal{V} and \mathcal{W} be \mathcal{R} -varieties.*

- (a) $\mathcal{V} \subseteq \mathcal{W}$ iff $\text{Th}(\mathcal{V}) \supseteq \text{Th}(\mathcal{W})$.
- (b) $\text{Mod}(\text{Th}(\mathcal{V})) = \mathcal{V}$.

Proof. (a) (\Rightarrow) follows immediately by definition. For (\Leftarrow), let $\rho_{\mathcal{V}, X} : \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X$ and $\rho_{\mathcal{W}, X} : \widehat{\mathbb{M}}_{\mathcal{W}} \mathbb{J}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{W}} \mathbb{J}X$ be the morphisms from Theorem 6.14. It follows by Lemma 7.2 that

$$\text{Th}(\mathcal{W}) \subseteq \text{Th}(\mathcal{V}) \text{ implies } \ker \rho_{\mathcal{W}, X} \subseteq \ker \rho_{\mathcal{V}, X}.$$

Hence, we can use the Factorisation Lemma to find a morphism $q_X : \widehat{\mathbb{M}}_{\mathcal{W}} \mathbb{J}X \rightarrow \widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X$ such that $\rho_{\mathcal{V}, X} = q_X \circ \rho_{\mathcal{W}, X}$. By Theorem 6.14, the

morphism $\rho_{\mathcal{V}, X}$ is surjective. Hence, so is q_X . That means that $\widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X$ is a quotient of $\widehat{\mathbb{M}}_{\mathcal{W}} \mathbb{J}X$. Consequently, every quotient of $\widehat{\mathbb{M}}_{\mathcal{V}} \mathbb{J}X$ is also a quotient of $\widehat{\mathbb{M}}_{\mathcal{W}} \mathbb{J}X$ and it follows by Proposition 7.4 that $\mathcal{V} \subseteq \mathcal{W}$.

(b) We have seen in Proposition 7.3 that the class $\mathcal{W} := \text{Mod}(\text{Th}(\mathcal{V}))$ is an \mathcal{R} -variety. We have to show that $\mathcal{V} = \mathcal{W}$.

(\subseteq) Let $\mathfrak{A} \in \mathcal{V}$. Then we have $\mathfrak{A} \models s \leq t$, for every $s \leq t$ in $\text{Th}(\mathcal{V})$. This implies that $\mathfrak{A} \in \text{Mod}(\text{Th}(\mathcal{V})) = \mathcal{W}$.

(\supseteq) By (a) it is sufficient to prove that $\text{Th}(\mathcal{W}) \supseteq \text{Th}(\mathcal{V})$. Hence, let $s \leq t$ be in $\text{Th}(\mathcal{V})$. Then $\mathfrak{A} \models s \leq t$, for all $\mathfrak{A} \in \text{Mod}(\text{Th}(\mathcal{V})) = \mathcal{W}$, which implies that $s \leq t$ belongs to $\text{Th}(\mathcal{W})$. \square

We are finally able to state our Reiterman theorem for pseudo-varieties of \mathbb{M} -algebras.

Theorem 7.6. *Let \mathcal{R} be the class of all finitary \mathbb{M} -algebras. A class \mathcal{V} is an \mathcal{R} -variety if, and only if, it is of the form $\mathcal{V} = \text{Mod}_{\mathcal{R}}(\Phi)$, for some set Φ of \mathbb{M} -inequalities.*

Proof. (\Leftarrow) was already proved in Proposition 7.3, and (\Rightarrow) follows by Corollary 7.5 since $\mathcal{V} = \text{Mod}_{\mathcal{R}}(\text{Th}(\mathcal{V}))$. \square

Notes

Monadic frameworks for formal language theory were initially put forward by Bojańczyk [5]. Later on they have been generalised and fleshed out in [21, 23, 16, 6, 4].

The material on minimal algebras in Section 3 is taken from lecture notes by Bojańczyk [6]. Theorem 3.11 is unpublished work of Bojańczyk and Plotkin.

The sections on pro-finitary terms and Reiterman's Theorem are based on [11, 23]. An introduction to Priestley spaces can be found, for instance, in Chapter 11 of [12]. For a thorough introduction to profinite groups, see [20].

IV. Logic

1. Abstract Logics

A major application of algebraic language theory consists in deriving criteria for when a language is definable in a given logic. In this section we will introduce an abstract framework covering a large number of the logics used in practice. Our focus will be on isolating some abstract properties of a logic ensuring that the corresponding language family forms a variety and, thus, fits into our framework. In the next section we will then investigate what it means for a language to be definable in a given logic.

Definition 1.1. Let $\Omega \in \mathcal{D}$ be a set of weights and let $\tilde{\Omega} \in \mathcal{D}^\Xi$ be the set with $\tilde{\Omega}_\xi = \Omega$, for all $\xi \in \Xi$.

(a) An Ω -valued logic is a triple $\langle L, \mathcal{M}, \text{Mod} \rangle$ consisting of a set $L \in \text{Set}^\Xi$ of formulae, a class $\mathcal{M} \in \mathcal{D}^\Xi$ of models, and a model function

$$\text{Mod} : L \rightarrow \mathcal{D}^\Xi(\mathcal{M}, \tilde{\Omega}).$$

We call $\text{Mod}(\varphi) : \mathcal{M}_\xi \rightarrow \Omega$ the *class of models* of $\varphi \in L_\xi$. To keep notation light, we usually identify a logic with its set of formulae L .

(b) For a logic $\langle L, \mathcal{M}, \text{Mod} \rangle$, we define the *satisfaction function*

$$\langle \cdot ; \cdot \rangle : \mathbb{V}\mathcal{M} \times L \rightarrow \mathbb{V}\tilde{\Omega} \quad \text{by} \quad \langle M; \varphi \rangle := \text{Mod}(\varphi)(M),$$

and the *theory function*

$$\text{Th}_L := \langle \text{Mod}(\varphi) \rangle_{\varphi \in L} : \mathcal{M} \rightarrow \tilde{\Omega}^L.$$

We call $\text{Th}_L(M)$ the L -theory of M .

(c) The congruence

$$\text{eq}(L) := \ker \text{Th}_L : \mathcal{M} \rightarrow \mathcal{M}/(\ker \text{Th}_L) \subseteq \tilde{\Omega}^L.$$

is called *L-equivalence*.

Remark. (a) If \mathcal{D} is one of Pos or Set and $\Omega = \{0, 1\}$, we usually replace the satisfaction function $\langle \cdot; \cdot \rangle : \mathcal{M} \times L \rightarrow \{0, 1\}$ by the relation $\models \subseteq \mathcal{M} \times L$ which is defined by

$$M \models \varphi \quad \text{iff} \quad \langle M; \varphi \rangle = 1.$$

Then we have

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

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

Note that $\text{Mod}(\varphi)$ is a set, while $\text{Th}_L(M) \in \mathcal{D}$. For $\mathcal{D} = \text{Pos}$, the ordering on $\text{Th}_L(M)$ is given by

$$\begin{aligned} \varphi \leq \psi & \quad \text{iff} \quad N \models \varphi \Rightarrow N \models \psi, \quad \text{for all } N \in \mathcal{M}, \\ & \quad \text{iff} \quad \varphi \text{ implies } \psi. \end{aligned}$$

(b) In Set, *L-equivalence* $\text{eq}(L)$ reduces to the well-known relation \equiv_L given by

$$M \equiv_L N \quad \text{iff} \quad \langle M; \varphi \rangle = \langle N; \varphi \rangle, \quad \text{for all } \varphi \in L.$$

In the unweighted case, this reads

$$M \equiv_L N \quad \text{iff} \quad M \models \varphi \Leftrightarrow N \models \varphi, \quad \text{for all } \varphi \in L.$$

In Pos, we obtain a preorder \sqsubseteq_L instead, which is given by

$$M \sqsubseteq_L N \quad \text{iff} \quad \langle M; \varphi \rangle \leq \langle N; \varphi \rangle, \quad \text{for all } \varphi \in L,$$

$$\text{or } M \sqsubseteq_L N \quad \text{iff} \quad M \models \varphi \Rightarrow N \models \varphi, \quad \text{for all } \varphi \in L.$$

Examples. (a) Let L be the set of all finite (word) automata over a fixed alphabet Σ and set $\mathcal{M} := \Sigma^*$. For $\mathcal{A} \in L$ and $w \in \mathcal{M}$, we define

$$w \models \mathcal{A} \quad \text{iff} \quad \mathcal{A} \text{ accepts the input } w.$$

(b) For a given signature Σ , a set X_1 of first-order variables, and a set X_2 of set variables, we can define *monadic second-order logic* as

$$\langle \text{MSO}[\Sigma, X_1, X_2], \text{Alg}[\Sigma, X_1, X_2], \text{Mod} \rangle$$

where $\text{MSO}[\Sigma, X_1, X_2]$ is the set of all monadic second-order formulae over the signature Σ with free first-order variables in X_1 and free monadic second-order variables in X_2 ; and $\text{Alg}[\Sigma, X_1, X_2]$ is the set of all triples $\langle \mathcal{A}, \beta_1, \beta_2 \rangle$ where \mathcal{A} is a Σ -structure and $\beta_1 : X \rightarrow A$ and $\beta_2 : X \rightarrow \wp(A)$ are variable assignments. We use $\Omega = \{0, 1\}$ as truth values and the satisfaction function maps a pair $\langle M, \varphi \rangle$ to 1 if $M \models \varphi$ and to 0 otherwise.

For multi-valued logics we have to distinguish two different notions of definability. *Strong* definability is the obvious generalisation from the 2-sorted case, whereas *weak* definability may appear less natural at first sight, but turns out to be much better behaved if the logic in question is not closed under all Ω -operations. For this reason, we will mainly use weak definability below.

Definition 1.2. Let $\mathcal{L} = \langle L, \mathcal{M}, \text{Mod} \rangle$ and $\mathcal{L}' = \langle L', \mathcal{M}', \text{Mod} \rangle$ be Ω -valued logics.

(a) The *extension of \mathcal{L} by Ω -operations* is the logic

$$\mathcal{L}[\Omega] := \langle L[\Omega], \mathcal{M}, \text{Mod} \rangle,$$

where $L[\Omega]$ is the set of all pairs $\langle \omega, \tilde{\varphi} \rangle$ such that $\omega : \Omega^n \rightarrow \Omega$ is an Ω -operation of some arity $n < \omega$ and $\tilde{\varphi} \in L^n$ an n -tuple of formulae. We use the syntax $\omega[\tilde{\varphi}]$ for such a tuple. The satisfaction function is defined by

$$\text{Mod}(\omega[\tilde{\varphi}]) := \omega[\text{Mod}(\varphi_0), \dots, \text{Mod}(\varphi_{n-1})].$$

(b) A class $C : \mathcal{M}_\xi \rightarrow \Omega$ is *strongly L-definable* if $C = \text{Mod}(\varphi)$, for some $\varphi \in L_\xi$. It is *(weakly) L-definable* if it is strongly $L[\Omega]$ -definable.

(c) A *strong morphism* $\langle \lambda, \mu \rangle : \mathcal{L} \rightarrow \mathcal{L}'$ consists of functions $\lambda : L \rightarrow L'$ and $\mu : \mathcal{M}' \rightarrow \mathcal{M}$ such that

$$\langle M'; \lambda(\varphi) \rangle = \langle \mu(M'); \varphi \rangle, \quad \text{for all } \varphi \in L \text{ and } M' \in \mathcal{M}'.$$

(d) A *(weak) morphism* $\langle \lambda, \mu \rangle : \mathcal{L} \rightarrow \mathcal{L}'$ is a strong morphism $\mathcal{L} \rightarrow \mathcal{L}'[\Omega]$. We denote the category of all logics and their (weak) morphisms by wLog . \lrcorner

Note that, for logics L that are closed under all Ω -operations, the logics L and $L[\Omega]$ are equi-expressive. Hence, there is no distinction between weak and strong definability, and none between weak and strong morphisms. For instance, for $\mathcal{D} = \text{Pos}$ and $\Omega = \{0, 1\}$, this is the case if the logic in question is closed under finite meets and joins. Similarly, for $\mathcal{D} = \text{Set}$ and $\Omega = \{0, 1\}$, we need closure under all finite boolean operations. Unfortunately, if the set Ω of truth values is infinite, strong definability and weak definability are usually different, and the latter notion is often better behaved than the more standard notion of strong definability. For this reason, we will focus on weak definability below.

Example. Every MSO-interpretation τ (from the signature Σ to Γ) gives rise to a morphism $\text{MSO}[\Gamma, \emptyset, \emptyset] \rightarrow \text{MSO}[\Sigma, \emptyset, \emptyset]$ since we can construct, for every formula $\varphi \in \text{MSO}[\Gamma, \emptyset, \emptyset]$, some formula $\varphi^\tau \in \text{MSO}[\Sigma, \emptyset, \emptyset]$ with

$$\tau(\mathfrak{A}) \models \varphi \quad \text{iff} \quad \mathfrak{A} \models \varphi^\tau, \quad \text{for all } \Sigma\text{-structures } \mathfrak{A}.$$

(Just replace in φ every atomic formula containing a relation symbol $R \in \Gamma$ by the formula from τ defining R .) \lrcorner

Let us isolate a few simple conditions for when a class of models is definable.

Lemma 1.3. *Let $\langle L, \mathcal{M}, \text{Mod} \rangle$ and $\langle L', \mathcal{M}', \text{Mod} \rangle$ be Ω -valued logics.*

(a) *A class $\mathcal{C} : \mathcal{M}_\xi \rightarrow \Omega$ is L -definable if, and only if,*

$$\text{eq}_\xi(\Delta) \leq \mathcal{C}, \quad \text{for some finite } \Delta \subseteq L_\xi.$$

(b) *For sort-wise finite sets $\Delta \subseteq L$ and $\Delta' \subseteq L'$, and a function $f : \mathcal{M} \rightarrow \mathcal{M}'$ the following two statements are equivalent:*

$$(1) \quad \text{eq}_\xi(\Delta) \leq \ker(\text{eq}_\xi(\Delta') \circ f)$$

(2) *If $\mathcal{C} : \mathcal{M}'_\xi \rightarrow \Omega$ is Δ' -definable, then $f^{-1}[\mathcal{C}] : \mathcal{M}_\xi \rightarrow \Omega$ is Δ -definable.*

Proof. (a) (\Rightarrow) Suppose that

$$\mathcal{C} = \omega[\text{Mod}(\varphi_0), \dots, \text{Mod}(\varphi_{n-1})].$$

For $\Delta := \{\varphi_0, \dots, \varphi_{n-1}\}$, it follows that

$$\begin{aligned} \mathcal{C} &= \omega \circ \langle \text{Mod}(\varphi_0), \dots, \text{Mod}(\varphi_{n-1}) \rangle \\ &= \omega \circ \text{Th}_\Delta \\ &= (\omega \circ \text{im Th}_\Delta) \circ \ker \text{Th}_\Delta = (\omega \circ \text{im Th}_\Delta) \circ \text{eq}_\xi(\Delta). \end{aligned}$$

Consequently, $\text{eq}_\xi(\Delta) \leq \mathcal{C}$.

(\Leftarrow) Set $\Theta := \mathcal{M}/\text{eq}_\xi(\Delta)$ and fix a morphism $\omega_\circ : \Theta_\xi \rightarrow \Omega$ such that $\mathcal{C} = \omega_\circ \circ \text{eq}_\xi(\Delta)$. Since $\text{im Th}_{\Delta, \xi} : \Theta_\xi \rightarrow \Omega^\Delta$ is an M -morphism and Ω is injective, we can extend ω_\circ to a function $\omega : \Omega^\Delta \rightarrow \Omega$ with $\omega \circ \text{im Th}_{\Delta, \xi} = \omega_\circ$. Consequently,

$$\begin{aligned} \mathcal{C} &= \omega_\circ \circ \text{eq}_\xi(\Delta) \\ &= \omega \circ \text{im Th}_{\Delta, \xi} \circ \ker \text{Th}_{\Delta, \xi} \\ &= \omega \circ \text{Th}_{\Delta, \xi} \\ &= \omega[\text{Mod}(\varphi_0), \dots, \text{Mod}(\varphi_{n-1})], \end{aligned}$$

where $\varphi_0, \dots, \varphi_{n-1}$ is an enumeration of Δ_ξ .

(b) (1) \Rightarrow (2) Suppose that $\mathcal{C} : \mathcal{M}'_\xi \rightarrow \Omega$ is Δ' -definable. By (1) and (a), it then follows that

$$\text{eq}_\xi(\Delta) \leq \ker(\text{eq}_\xi(\Delta') \circ f) \leq \ker(\mathcal{C} \circ f).$$

Hence, (a) (applied to the logic Δ instead of L) implies that $f^{-1}[\mathcal{C}]$ is Δ -definable.

(2) \Rightarrow (1) Suppose that $\text{eq}_\xi(\Delta) \leq \ker(\text{eq}_\xi(\Delta') \circ f)$. Note that every class $\text{Mod}(\varphi)$ with $\varphi \in \Delta'$ is trivially Δ' -definable. By assumption, it therefore follows that $f^{-1}[\text{Mod}(\varphi)]$ is Δ -definable. Hence, (a) implies that

$$\text{eq}_\xi(\Delta) \leq \ker(\text{Mod}(\varphi) \circ f), \quad \text{for all } \varphi \in \Delta'_\xi.$$

Consequently,

$$\begin{aligned} \text{eq}_\xi(\Delta) &\leq \prod_{\varphi \in \Delta'_\xi} \ker(\text{Mod}(\varphi) \circ f) \\ &= \ker(\langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta'_\xi} \circ f) \\ &= \ker(\text{Th}_{\Delta', \xi} \circ f) \\ &= \ker(\text{eq}_\xi(\Delta') \circ f). \end{aligned} \quad \square$$

Remark. In the category Pos , the above conditions simplify to the following ones.

- (a) $M \sqsubseteq_\Delta N$ implies $\mathcal{C}(M) \leq \mathcal{C}(N)$.
- (b) $M \sqsubseteq_\Delta N$ implies $f(M) \sqsubseteq_{\Delta'} f(N)$.

For Set, we just have to replace \sqsubseteq_Δ by \sqsubseteq_Δ , and \leq by $=$.

The next result provides a characterisation of when a map on models is part of a morphism of logics.

Lemma 1.4. Let $\mathcal{L} = \langle L, \mathcal{M}, \text{Mod} \rangle$ and $\mathcal{L}' = \langle L', \mathcal{M}', \text{Mod} \rangle$ be Ω -valued logics and $\mu : \mathcal{M}' \rightarrow \mathcal{M}$ a function. The following statements are equivalent.

- (1) There exists a function $\lambda : L \rightarrow L'$ such that $\langle \lambda, \mu \rangle : \mathcal{L} \rightarrow \mathcal{L}'$ is a morphism of logics.
- (2) If $\mathcal{C} : \mathcal{M}_\xi \rightarrow \Omega$ is L -definable, then $\mu^{-1}[\mathcal{C}] : \mathcal{M}'_\xi \rightarrow \Omega$ is L' -definable.
- (3) For every sort-wise finite $\Delta \subseteq L$, there exists a sort-wise finite $\Delta' \subseteq L'$ such that

$$\text{eq}(\Delta') \leq \ker(\text{eq}(\Delta) \circ \mu).$$

Proof. (3) \Rightarrow (2) Suppose that $\mathcal{C} : \mathcal{M}_\xi \rightarrow \Omega$ is L -definable. Then it is Δ -definable, for some finite set $\Delta \subseteq L_\xi$. By assumption, we can therefore find some sort-wise finite set $\Delta' \subseteq L'$ such that

$$\text{eq}(\Delta') \leq \ker(\text{eq}(\Delta) \circ \mu).$$

By Lemma 1.3 (b) this implies that $\mathcal{C} \circ \mu$ is Δ' -definable. In particular, it is L' -definable.

(2) \Rightarrow (1) We define $\lambda : L \rightarrow L'$ as follows. For each $\varphi \in L$, the class $\text{Mod}(\varphi)$ is obviously L -definable. By assumption it follows that the preimage $\text{Mod}(\varphi) \circ \mu$ is L' -definable, i.e., there are formulae $\psi'_0, \dots, \psi'_n \in L'$ and an Ω -operation $\omega : \Omega^n \rightarrow \Omega$ such that

$$\text{Mod}(\varphi) \circ \mu = \omega[\text{Mod}(\psi'_0), \dots, \text{Mod}(\psi'_n)].$$

We can therefore set $\lambda(\varphi) := \langle \omega, \psi'_0, \dots, \psi'_n \rangle$.

(1) \Rightarrow (3) Given $\Delta \subseteq L$, we set

$$\Delta' := \bigcup \{ \{ \psi'_0, \dots, \psi'_n \} \mid \varphi \in \Delta, \lambda(\varphi) = \langle \omega, \psi'_0, \dots, \psi'_n \rangle \}.$$

Then

$$\text{Mod}(\varphi) \circ \mu = \omega[\text{Mod}(\psi'_0), \dots, \text{Mod}(\psi'_n)], \quad \text{for every } \varphi \in \Delta.$$

Consequently, there exists a function $\chi : \Omega^{\Delta'} \rightarrow \Omega^\Delta$ such that

$$\langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ \mu = \chi \circ \langle \text{Mod}(\psi) \rangle_{\psi \in \Delta'}.$$

Hence,

$$\text{Th}_{\Delta'} = \langle \text{Mod}(\psi) \rangle_{\psi \in \Delta'} \leq \langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ \mu = \text{Th}_\Delta \circ \mu,$$

which implies that

$$\text{eq}(\Delta') = \ker \text{Th}_{\Delta'} \leq \ker(\text{Th}_\Delta \circ \mu) = \text{eq}(\Delta) \circ \mu. \quad \square$$

Here, we are mainly interested in logics whose class of models is of the form $\mathcal{M} = \mathbb{M}\Sigma$ with $\Sigma \in \text{Alph}$, as these can be used to define languages. As with families of languages, we also need to consider families of logics indexed by the alphabet used.

Definition 1.5. (a) A logic L is *over* an alphabet Σ if its class of models is equal to $\mathbb{M}\Sigma$.

(b) A *family of logics* is a functor $L : \text{Alph} \rightarrow \text{wLog}$ such that

- ♦ for every alphabet Σ , the image $L[\Sigma]$ is a logic over Σ ,
- ♦ for every function $f : \Sigma \rightarrow \Gamma$, the image $L[f]$ is a morphism $\langle \lambda, \mu \rangle : L[\Sigma] \rightarrow L[\Gamma]$ with $\mu = \mathbb{M}f$.

(c) Let L be a family of logics. A family of languages \mathcal{K} is *L -definable* if, for all alphabets Σ and all sorts ξ , every $\kappa \in \mathcal{K}_\xi[\Sigma]$ is $L[\Sigma]$ -definable.

(d) Let L be a family of logics and $A \in \mathcal{D}$ finite. We call a function $\kappa : \mathbb{M}A \rightarrow \Omega$ *L -definable*, if the language $\kappa \circ \mathbb{M}l : \mathbb{M}\mathbb{J}\forall A \rightarrow \Omega$ is $L[\mathbb{J}\forall A]$ -definable.

(e) A family L of logics is *varietal* if the class of all L -definable languages forms a variety of languages.

(f) We call a family of logics L (*sort-wise*) *finite* if, for every alphabet Σ , the set of formulae $L[\Sigma]$ is (sort-wise) finite.

(g) To keep notation light we will drop the signature in cases where it is understood. Thus, we will usually write L instead of $L[\Sigma]$.

Example. For the word monad $\mathbb{M}A := A^*$ and monadic second-order logic, we can define a family MSO that maps an alphabet Σ to the logic $\text{MSO}[\hat{\Sigma}, \emptyset, \emptyset]$ where

$$\hat{\Sigma} := \{E, \leq\} \cup \{P_a \mid a \in \Sigma\}$$

is the signature consisting of the successor relation E , the ordering \leq , and predicates P_a for all letters in Σ .

2. Compositionality

As the notion is very general, there is not much one can prove for an arbitrary logic. To get non-trivial statements we need some kind of restriction. As languages come equipped with a monadic composition operation, it is natural to require our logics to be well-behaved under this form of composition. This leads to the following definition.

Definition 2.1. A family L of logics is *\mathbb{M} -compositional* if, for every finite subfamily $\Phi \subseteq L$, there exists some sort-wise finite subfamily $\Phi \subseteq \Delta \subseteq L$ such that

$$\text{eq}(\Delta[\Sigma]) \text{ is an } \mathbb{M}\text{-congruence on } \mathbb{M}\Sigma, \quad \text{for all alphabets } \Sigma.$$

Example. For words $u_0, \dots, u_{n-1}, v_0, \dots, v_{n-1} \in \Sigma^*$ we have

$$u_i \equiv_{\text{MSO}_m} v_i, \quad \text{for all } i, \quad \text{implies} \quad u_0 \cdots u_{n-1} \equiv_{\text{MSO}_m} v_0 \cdots v_{n-1},$$

where MSO_m denotes the set of MSO-formulae of quantifier rank at most m . Consequently, MSO is \mathbb{M} -compositional for the word monad $\mathbb{M}A = A^*$.

The importance of \mathbb{M} -compositionality stems from the fact that the set of theories of such a logic forms an \mathbb{M} -algebra.

Proposition 2.2. A family of logics L is *\mathbb{M} -compositional* if, and only if, for every finite subfamily $\Phi \subseteq L$, there exist

- ♦ a sort-wise finite subfamily $\Phi \subseteq \Delta \subseteq L$,
- ♦ a functor $\Theta_\Delta : \text{Alph} \rightarrow \text{Alg}(\mathbb{M})$, and
- ♦ an epimorphic natural transformation $\theta_\Delta : (\mathbb{M} \upharpoonright \text{Alph}) \Rightarrow \Theta_\Delta$

such that

$$(\theta_\Delta)_\Sigma = \text{eq}(\Delta[\Sigma]), \quad \text{for every } \Sigma.$$

Proof. (\Leftarrow) Given $\Phi \subseteq L$, choose $\Phi \subseteq \Delta \subseteq L$ such that $\text{eq}(\Delta[\Sigma]) = \theta_\Delta$. Since θ_Δ is a morphism of \mathbb{M} -algebras, it follows that $\text{eq}(\Delta[\Sigma])$ is an \mathbb{M} -congruence.

(\Rightarrow) Given $\Phi \subseteq L$, choose $\Phi \subseteq \Delta \subseteq L$ such that $\text{eq}(\Delta[\Sigma])$ is an \mathbb{M} -congruence, for all Σ . Set

$$\Theta_\Delta \Sigma := \mathbb{M}\Sigma / \text{eq}(\Delta[\Sigma]) \quad \text{and} \quad (\theta_\Delta)_\Sigma := \text{eq}(\Delta[\Sigma]) : \mathbb{M}\Sigma \rightarrow \Theta_\Delta \Sigma.$$

Given a function $f : \Sigma \rightarrow \Gamma$, we define the morphism $\Theta_\Delta f : \Theta_\Delta \Sigma \rightarrow \Theta_\Delta \Gamma$ as follows. By definition of a family of logics, $L[f] = \langle \lambda, \mathbb{M}f \rangle$ is a morphism of logics. Hence, Lemma 1.3 implies that

$$\text{eq}(\Delta[\Sigma]) \leq \ker(\text{eq}(\Delta[\Gamma]) \circ \mathbb{M}f).$$

Consequently, there exists some function $\psi : \Theta_\Delta \Sigma \rightarrow \Theta_\Delta \Gamma$ with

$$\psi \circ \theta_\Delta = \psi \circ \text{eq}(\Delta[\Sigma]) = \text{eq}(\Delta[\Gamma]) = \theta_\Delta \circ \mathbb{M}f.$$

We set $\Theta_\Delta f := \psi$.

From this definition it immediately follows that θ_Δ is a natural transformation $\mathbb{M} \Rightarrow \Theta_\Delta$ since

$$\Theta_\Delta f \circ \theta_\Delta = \theta_\Delta \circ \mathbb{M}f.$$

Hence, it remains to show that Θ_Δ is a functor. Consider two functions $f : \Sigma \rightarrow \Gamma$ and $g : \Gamma \rightarrow \Upsilon$. By the equation we have just established, we have

$$\begin{aligned} \Theta_\Delta(g \circ f) \circ \theta_\Delta &= \theta_\Delta \circ \mathbb{M}(g \circ f) \\ &= \theta_\Delta \circ \mathbb{M}g \circ \mathbb{M}f \\ &= \Theta_\Delta g \circ \theta_\Delta \circ \mathbb{M}f = \Theta_\Delta g \circ \Theta_\Delta f \circ \theta_\Delta. \end{aligned}$$

As θ_Δ is an epimorphism, this implies that $\Theta_\Delta(g \circ f) = \Theta_\Delta g \circ \Theta_\Delta f$. \square

It follows immediately from the definition that the algebras $\Theta_\Delta \Sigma$ are sort-wise finite-dimensional.

Lemma 2.3. *Let Δ be a sort-wise finite set such that $\Theta_\Delta \Sigma$ exists. Then $\Theta_\Delta \Sigma$ is sort-wise strongly finite-dimensional.*

Proof. By definition we have

$$\ker \theta_\Delta = \text{eq}(\Delta) \leq \ker \text{Mod}(\varphi), \quad \text{for all } \varphi \in \Delta.$$

For every $\varphi \in \Delta$, we can therefore find a function $\mu_\varphi : \Theta_\Delta \Sigma \rightarrow \Omega$ such that

$$\mu_\varphi \circ \theta_\Delta = \text{Mod}(\varphi).$$

It is sufficient to prove that $e := \langle \mu_\varphi \rangle_{\varphi \in \Delta} : \Theta_\Delta \Sigma \rightarrow \Omega^\Delta$ belongs to M . Note that

$$e \circ \theta_\Delta = \langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta} = \text{Th}_\Delta$$

implies that

$$\ker(e \circ \theta_\Delta) = \ker \text{Th}_\Delta = \text{eq}(\Delta) = \theta_\Delta.$$

Consequently, $e \in M$. \square

It follows immediately from the definition that the theory algebra $\Theta_\Delta \Sigma$ recognises every Δ -definable language.

Lemma 2.4. *The morphism $\theta_\Delta : \mathbb{M}\Sigma \rightarrow \Theta_\Delta \Sigma$ recognises every Δ -definable language $\kappa : \mathbb{M}\Sigma \rightarrow \Omega$.*

Proof. Suppose that $\kappa = \omega[\text{Mod}(\varphi_0), \dots, \text{Mod}(\varphi_{n-1})]$ for some formulae $\varphi_0, \dots, \varphi_{n-1} \in \Delta$ and an Ω -operation $\omega : \Omega^n \rightarrow \Omega$. Since

$$\theta_\Delta = \text{eq}(\Delta[\Sigma]) \leq \text{eq}(\{\varphi_i\}), \quad \text{for all } i < n,$$

we have

$$\begin{aligned} \theta_\Delta &\leq \inf_{i < n} \text{eq}(\{\varphi_i\}) \\ &= \inf_{i < n} \ker \text{Mod}(\varphi_i) \\ &= \ker \langle \text{Mod}(\varphi_i) \rangle_{i < n} \\ &\leq \ker(\omega \circ \langle \text{Mod}(\varphi_i) \rangle_{i < n}) = \ker \kappa. \end{aligned}$$

Consequently, we can use Lemma III.3.1 (a) to show that θ_Δ recognises κ . \square

Corollary 2.5. *Suppose the L is \mathbb{M} -compositional, $\Delta \subseteq L$ a sort-wise finite set such that θ_Δ exists, and let $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ be a language with a syntactic algebra. Then κ is Δ -definable if, and only if,*

$$\text{eq}(\Delta) \leq \ker \text{syn}_\kappa.$$

Proof. (\Rightarrow) As θ_Δ recognises κ , we can use Theorem III.4.9 to find a function $\rho : \Theta_\Delta \Sigma \rightarrow \text{Syn}(\kappa)$ with $\rho \circ \theta_\Delta = \text{syn}_\kappa$. It follows that

$$\text{eq}(\Delta) = \ker \theta_\Delta \leq \ker (\rho \circ \theta_\Delta) = \ker \text{syn}_\kappa.$$

(\Leftarrow) By Lemma III.4.7, we have

$$\text{eq}(\Delta) \leq \ker \text{syn}_\kappa \leq \ker \kappa \leq \kappa.$$

Therefore Lemma 1.3 implies that κ is Δ -definable. \square

Corollary 2.6. *Suppose the L is \mathbb{M} -compositional, $\Delta \subseteq L$ a sort-wise finite set such that θ_Δ exists, and let $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ be Δ -definable. If the minimal algebra $\text{Syn}(\kappa)$ exists, it is sort-wise weakly finite-dimensional.*

Proof. As θ_Δ recognises κ , we can use Theorem III.4.9 to find a function $\rho : \Theta_\Delta \Sigma \rightarrow \text{Syn}(\kappa)$ with $\rho \circ \theta_\Delta = \text{syn}_\kappa$. Since syn_κ and θ_Δ are E -morphisms, so is ρ . Furthermore, we have shown in Lemma 2.3 that $\Theta_\Delta \Sigma$ is strongly finite-dimensional. \square

Next, let us take a look at the closure properties of definable languages. Our first observation concerns closure under inverse relabellings, which holds for every logic L . Then we show that \mathbb{M} -compositionality implies, but is slightly stronger than, closure under derivatives.

Lemma 2.7. *Let L be a family of logics. The class of L -definable languages is closed under inverse relabellings.*

Proof. If $f : \Sigma \rightarrow \Gamma$ is a morphism of Alph , it follows by the definition of a family of logics that there is some function λ such that $L[f] = \langle \lambda, \mathbb{M}f \rangle$ is a morphism of logics. Consequently, it follows by Lemma 1.4 that $(\mathbb{M}f)^{-1}[\kappa]$ is L -definable, for every L -definable language $\kappa : \mathbb{M}_\xi \Gamma \rightarrow \Omega$. \square

Lemma 2.8. *Let L be an \mathbb{M} -compositional family of logics, and let $\Delta \subseteq L$ be a subfamily such that $\text{eq}(\Delta)$ is an \mathbb{M} -congruence. Then*

$$\text{eq}(\Delta) \leq \text{eq}(\Delta) \circ p, \quad \text{for all contexts } p \in \mathbb{M}(\Sigma + \square).$$

Proof. By Lemma III.4.3, there exists some context q with

$$q \circ \theta_\Delta = \theta_\Delta \circ p.$$

Hence, $\theta_\Delta \leq \ker (\theta_\Delta \circ p)$, and it follows that

$$\text{eq}(\Delta) = \theta_\Delta \leq \ker (\theta_\Delta \circ p) = \ker (\text{eq}(\Delta) \circ p). \quad \square$$

Usually, the theory algebras $\Theta_\Delta \Sigma$ from Proposition 2.2 are not very well understood. (Otherwise, we would not need to introduce a special algebraic framework to study definability questions.) To shed a bit more light on what these algebras look like, we present an alternative construction for the theory functor Θ .

Definition 2.9. Let L be a family of logics such that every L -definable language has a syntactic algebra. The *syntactic theory morphism* (for an alphabet Σ) is

$$\tilde{\theta}_L := \langle \text{syn}_{\text{Mod}(\varphi)} \rangle_{\varphi \in L[\Sigma]} : \mathbb{M}\Sigma \rightarrow \prod_{\varphi \in L[\Sigma]} \text{Syn}(\text{Mod}(\varphi)).$$

Lemma 2.10. *Let L be a family of logics such that every L -definable language has a syntactic algebra, and let $\Delta \subseteq L$ be sort-wise finite. The following statements are equivalent.*

- (1) *The class of Δ -definable languages is closed under derivatives.*
- (2) $\text{eq}(\Delta) = \ker \tilde{\theta}_\Delta(t).$
- (3) $\text{eq}(\Delta) \leq \ker (\text{eq}(\Delta) \circ p), \quad \text{for every context } p.$

Proof. (1) \Leftrightarrow (3) We have shown in Lemma 1.3 (b) that

$$\text{eq}(\Delta[\Sigma]) \leq \ker (\Delta[\Sigma] \circ p),$$

if, and only if, $p^{-1}[\kappa] : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ is Δ -definable, for every Δ -definable $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$.

(2) \Leftrightarrow (3) First, note that $\text{eq}(\Delta) \circ p = \text{eq}(\Delta)$ for the empty context $p = \text{sing}(\square)$. Consequently, (3) is equivalent to

$$\text{eq}(\Delta) = \inf \{ \ker(\text{eq}(\Delta) \circ p) \mid p \text{ a context} \}.$$

The equivalence of (2) and (3) therefore follows from the fact that

$$\begin{aligned} \ker \tilde{\theta}_\Delta &= \inf \{ \ker \text{syn}_{\text{Mod}(\varphi)} \mid \varphi \in \Delta \} \\ &= \inf \{ \ker(\text{Mod}(\varphi) \circ p) \mid \varphi \in \Delta, p \text{ context} \} \\ &= \inf \{ \ker(\langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ p) \mid p \text{ context} \} \\ &= \inf \{ \ker(\text{Th}_\Delta \circ p) \mid p \text{ context} \} \\ &= \inf_p \ker(\text{eq}(\Delta) \circ p), \end{aligned}$$

where the second step follows by Lemma III.4.7. \square

Theorem 2.11. *Let L be a family of logics such that every L -definable language has a syntactic algebra. The following statements are equivalent.*

- (1) L is \mathbb{M} -compositional.
- (2) For every finite $\Phi \subseteq L$, there exists a sort-wise finite $\Phi \subseteq \Delta \subseteq L$ such that the class of Δ -definable languages is closed under derivatives.

Proof. (1) \Rightarrow (2) This follows immediately from Lemma 1.3 (b) together with Lemma 2.8.

(2) \Rightarrow (1) Given a subfamily $\Delta \subseteq L$ with the above closure properties, it follows by Lemma 2.10 that $\text{eq}(\Delta) = \ker \tilde{\theta}_\Delta$. In particular, $\text{eq}(\Delta)$ is an \mathbb{M} -congruence. \square

Apart from a criterion for \mathbb{M} -compositionality, this theorem also gives us an explicit construction of the theory algebra $\Theta_\Delta \Sigma$ in language-theoretic terms. It therefore provides a more direct link between properties of a logic L and properties of the class of L -definable languages.

3. Definable Algebras

We have finally arrived at the central part of this chapter where we make the connection between algebra and logic. It follows from Theorem III.5.11 that, to every varietal logic L , there corresponds a unique \mathcal{R} -variety \mathcal{V} of \mathbb{M} -algebras recognising the family of L -definable languages. We would like to use these \mathbb{M} -algebras to study the expressive power of our logic L . To do so, we need to know as much as possible about how the algebras in \mathcal{V} look like. Unfortunately, Theorem III.5.11 does not tell us very much about that. The following definition provides a slightly more concrete description.

Definition 3.1. Let \mathfrak{A} be an \mathbb{M} -algebra and L a family of logics.

(a) A finite subset $C \subseteq \mathbb{V}A$ is *L -definably embedded* in \mathfrak{A} if, for every sort $\xi \in \Xi$ and every function $\mu : A_\xi \rightarrow \Omega$, the composition $\mu \circ \pi \upharpoonright \mathbb{M}_\xi \mathbb{J}C$ is L -definable.

(b) \mathfrak{A} is *locally L -definable* if every finite subset $C \subseteq \mathbb{V}A$ is L -definably embedded in \mathfrak{A} .

(c) \mathfrak{A} is *L -definable* if it is finitely generated, sort-wise weakly finite-dimensional, and locally L -definable. \lrcorner

Remark. Suppose that $\mathcal{D} = \text{Pos}$ and $\Omega = \{0, 1\}$. A subset C of \mathfrak{A} is L -definably embedded if, for every subset $P \subseteq A_\xi$, the preimage

$$\pi^{-1}[\upharpoonright P] \cap \mathbb{M}C \text{ is } L\text{-definable.}$$

Furthermore, being weakly finite-dimensional is the same as being finite. \lrcorner

Example. For the word functor $\mathbb{M}A = A^+$, every finite algebra (i.e., every finite semigroup) is MSO-definable since we can evaluate products in MSO. (Just guess a labelling that associates with every position the product of the corresponding prefix.)

The same is true for the functor $\mathbb{M}\langle A_1, A_\infty \rangle = \langle A_1^+, A_1^+ A_\infty \cup A_1^\omega \rangle$ for infinite words, and for the functor for finite trees. (For the former, one can use a reduction to the semigroup case via a simple application of the Theorem of Ramsey; for the latter, one can compute the product of a tree bottom-up similarly to the semigroup case.)

For infinite trees the situation is more complicated: there exist finitary algebras that are not MSO-definable. One such example will be presented in Section V.4.

If our logic L is sufficiently well-behaved, it immediately follows from this definition that L -definable algebras only recognise L -definable languages. (The converse, that every L -definable language is recognised by some L -definable algebra, is harder to prove. We will do so in the next section.) Note that this correspondence – besides being trivial – is also not that useful for understanding the expressive power of L as the definition makes essential use of L -definability. But the above definition can serve as a starting point for deriving more useful descriptions – that of course will be specific to the logic in question.

Before proving that the L -definable algebras are exactly those that only recognise L -definable languages, let us start by looking at definably embedded sets.

Lemma 3.2. *Let L be a family of Ω -valued logics and \mathfrak{A} a sort-wise weakly finite-dimensional \mathbb{M} -algebra. A finite set $C \subseteq \mathbb{V}\mathfrak{A}$ is L -definably embedded in \mathfrak{A} if, and only if, there exists a sort-wise finite set $\Delta \subseteq L[C]$ such that*

$$\text{eq}(\Delta) \upharpoonright \mathbb{M}\mathbb{J}C \leq \ker(\pi \upharpoonright \mathbb{M}\mathbb{J}C).$$

Proof. (\Leftarrow) Let Δ be a sort-wise finite set such that

$$\text{eq}(\Delta) \upharpoonright \mathbb{M}\mathbb{J}C \leq \ker(\pi \upharpoonright \mathbb{M}\mathbb{J}C).$$

For $\mu : A_\xi \rightarrow \Omega$, it follows that

$$\text{eq}(\Delta) \upharpoonright \mathbb{M}\mathbb{J}C \leq \ker(\pi \upharpoonright \mathbb{M}\mathbb{J}C) \leq \ker(\mu \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C).$$

By Lemma 1.3 (a) this means that $\mu \circ \pi \upharpoonright \mathbb{M}_\xi C$ is L -definable.

(\Rightarrow) Since A_ξ is weakly finite-dimensional, there exist morphisms $q_\xi : \tilde{A}_\xi \rightarrow A$ and $e_\xi : \tilde{A}_\xi \rightarrow \Omega^{d_\xi}$, for some $d_\xi < \omega$, with $q_\xi \in E$ and $e_\xi \in M$. Set $q := (q_\xi)_\xi$, $e := (e_\xi)_\xi$, let $j : C \rightarrow \mathbb{V}\mathfrak{A}$ be the inclusion, and set

$$\pi_o := \pi \circ \mathbb{M}(q \circ \iota \circ \mathbb{J}j) : \mathbb{M}\mathbb{J}C \rightarrow A$$

be the restriction of the product. Since $\mathbb{M}\mathbb{J}C$ is projective with respect to E -morphisms and $q \in E$, we can find morphisms $\tilde{\pi}_\xi : \mathbb{M}_\xi \mathbb{J}C \rightarrow \tilde{A}_\xi$ with

$$q_\xi \circ \tilde{\pi}_\xi = \pi_o, \quad \text{for all } \xi.$$

Let $p_i : \Omega^{d_\xi} \rightarrow \Omega$ be the projection to the i -th component. As C is L -definably embedded in \mathfrak{A} , we can fix, for every $i < d_\xi$, a finite set $\Delta_{\xi,i} \subseteq L$ of formulae such that the composition $p_i \circ e_\xi \circ \tilde{\pi}_\xi$ is $\Delta_{\xi,i}$ -definable. The union $\Delta := \bigcup_{\xi \in \Xi} \bigcup_{i < d_\xi} \Delta_{\xi,i}$ is sort-wise finite and satisfies

$$\text{eq}_\xi(\Delta) \leq \text{eq}_\xi(\Delta_{\xi,i}) \leq \ker(p_i \circ e_\xi \circ \tilde{\pi}_\xi),$$

where the second inequality follows by Lemma 1.3 (a). Hence,

$$\begin{aligned} \text{eq}_\xi(\Delta) &\leq \inf_i \ker(p_i \circ e_\xi \circ \tilde{\pi}_\xi) \\ &= \ker(e_\xi \circ \tilde{\pi}_\xi) \\ &= \ker \tilde{\pi}_\xi \\ &\leq \ker(q_\xi \circ \tilde{\pi}_\xi) = \ker(\pi_o)_\xi, \end{aligned}$$

where the third step follows since $e_\xi \in M$. \square

In general, the closure properties of definably embedded sets are rather weak. To make them better behaved we have to impose some restriction on the logic L .

Lemma 3.3. *Let \mathfrak{A} be an \mathbb{M} -algebra, L a family of logics, and $C \subseteq \mathbb{V}\mathfrak{A}$ a finite set that is L -definably embedded in \mathfrak{A} .*

- (a) *Every subset of C is L -definably embedded in \mathfrak{A} .*
- (b) *If the class of L -definable languages is closed under inverse morphisms, every finite subset $D \subseteq \langle\langle C \rangle\rangle_{\mathfrak{A}}$ is L -definably embedded in \mathfrak{A} .*

Proof. (a) Fix $D \subseteq C$ and let $i : D \rightarrow C$ and $j : \mathbb{J}C \rightarrow A$ be the respective inclusion maps. For $\mu : A_\xi \rightarrow \Omega$ it follows that

$$\mu \circ \pi \upharpoonright \mathbb{M}_\xi D = \mu \circ \pi \circ \mathbb{M}j \circ \mathbb{M}i = (\mu \circ \pi \circ \mathbb{M}j) \circ \mathbb{M}i.$$

By assumption, $\mu \circ \pi \circ \mathbb{M}j$ is L -definable. Consequently, Lemma 2.7 implies that so is $\mu \circ \pi \circ \mathbb{M}j \circ \mathbb{M}ji$.

(b) By (a) it is sufficient to consider the case where $D = \mathbb{V}\langle C \rangle_{\mathfrak{A}}$. Let $i : C \rightarrow D$ and $j : \mathbb{J}D \rightarrow A$ be the inclusion maps. For every $d \in D$, we can find an element $f(d) \in \mathbb{M}C$ such that $\pi(f(d)) = d$. This defines a function $f : D \rightarrow \mathbb{V}\mathbb{M}jC$ with

$$\mathbb{V}(\pi \circ \mathbb{M}(j \circ \mathbb{J}i)) \circ f = \text{id}_D.$$

Via the adjunction, we obtain a function $f^* : \mathbb{J}D \rightarrow \mathbb{M}jC$. Let $\varphi : \mathbb{M}jD \rightarrow \mathbb{M}jC$ be the (unique) extension of f^* to a morphism of \mathbb{M} -algebras. Then

$$\begin{aligned} \pi \circ \mathbb{M}j \circ \mathbb{M}ji \circ \varphi \circ \text{sing} &= \pi \circ \mathbb{M}j \circ \mathbb{M}ji \circ f \\ &= \text{id} \\ &= \pi \circ \mathbb{M}j \circ \text{sing}. \end{aligned}$$

Since morphisms of \mathbb{M} -algebras are uniquely determined by their restriction to rng sing , it follows that

$$\pi \circ \mathbb{M}j \circ \mathbb{M}ji \circ \varphi = \pi \circ \mathbb{M}j.$$

To show that D is L -definably embedded, consider a function $\mu : A_\xi \rightarrow \Omega$. We have to show that $\mu \circ \pi \circ \mathbb{M}j : \mathbb{M}jD \rightarrow \Omega$ is L -definable. As C is L -definably embedded, we know that $\mu \circ \pi \circ \mathbb{M}(j \circ \mathbb{J}i) : \mathbb{M}_\xi \mathbb{J}C \rightarrow \Omega$ is L -definable. Furthermore, by assumption, L -definable languages are closed under inverse morphisms. Hence,

$$\mu \circ \pi \circ \mathbb{M}(j \circ \mathbb{J}i) \circ \varphi = \mu \circ \pi \circ \mathbb{M}j$$

is also L -definable. \square

It follows immediately from the definition that an L -definable algebra only recognises L -definable languages. We start with a slightly more precise statement.

Theorem 3.4. *Let L be a family of logics such that the L -definable languages are closed under inverse morphisms. An \mathbb{M} -algebra \mathfrak{A} is locally L -definable if, and only if, every language recognised by \mathfrak{A} is L -definable.*

Proof. (\Leftarrow) If some finite subset $C \subseteq \mathbb{V}A$ is not L -definably embedded, we can find a partial function $\mu : A_\xi \rightarrow \Omega$ such that the composition $\kappa := \mu \circ \pi \upharpoonright \mathbb{M}jC$ is not L -definable. Thus, the restriction $\pi \upharpoonright \mathbb{M}jC : \mathbb{M}jC \rightarrow \mathfrak{A}$ of the product is a morphism recognising the non- L -definable language κ .

(\Rightarrow) Let $\varphi : \mathbb{M}\Sigma \rightarrow \mathfrak{A}$ be a morphism and $\mu : A_\xi \rightarrow \Omega$ a function. We have to show that $\kappa := \mu \circ \varphi$ is L -definable. By assumption, the set $C := \text{rng}(\varphi \circ \text{sing})$ is L -definably embedded in \mathfrak{A} . By definition, this implies that the function $\mu \circ \pi \upharpoonright \mathbb{M}jC : \mathbb{M}jC \rightarrow \Omega$ is L -definable. As Σ is of the form $\mathbb{J}X$, we can use Corollary II.5.4 to find a function $\hat{\varphi} : \mathbb{M}\Sigma \rightarrow \mathbb{M}jC$ satisfying $\varphi = \pi \circ \hat{\varphi}$. Hence,

$$\kappa = \mu \circ \varphi = \mu \circ \pi \circ \hat{\varphi}.$$

Since $\mu \circ \pi \upharpoonright \mathbb{M}jC$ is L -definable and since we have assumed above that the class of L -definable languages is closed under inverse morphisms, it follows that κ is also L -definable. \square

Next, let us take a look at the closure properties of L -definable algebras.

Proposition 3.5. *The class of locally L -definable \mathbb{M} -algebras is closed under M -subalgebras and finite products.*

Proof. For subalgebras, suppose that $\mathfrak{A} \subseteq \mathfrak{B}$ where \mathfrak{B} is locally L -definable. Let $\pi_A : \mathbb{M}A \rightarrow A$ and $\pi_B : \mathbb{M}B \rightarrow B$ be the respective products. Given a finite set $C \subseteq \mathbb{V}A$ and a function $\mu : A_\xi \rightarrow \Omega$, consider the language $\kappa := \mu \circ \pi_A \upharpoonright \mathbb{M}jC$. Since $\pi_A \upharpoonright \mathbb{M}jC = \pi_B \upharpoonright \mathbb{M}jC$, we have $\kappa = \mu \circ \pi_B \upharpoonright \mathbb{M}jC$ and it follows by our assumption on \mathfrak{B} that κ is L -definable.

It remains to consider products. First, note that, according to Proposition I.5.7, the empty product \mathfrak{A} has exactly one element 1_ξ of each sort ξ . For a function $\mu : A_\xi \rightarrow \Omega$, it follows that $\kappa := \mu \circ \pi \upharpoonright \mathbb{M}_\xi C$ is the constant map with value $\mu(1_\xi) \in \Omega$. Hence, we can write

$$\kappa = \omega,$$

for the Ω -operation $\omega : \Omega^\circ \rightarrow \Omega : \langle \rangle \mapsto \mu(1_\xi)$. Thus, κ is L -definable.

It remains to consider the case of a non-empty, finite product $\mathfrak{A} = \prod_{i < n} \mathfrak{B}^i$ where each \mathfrak{B}^i is locally L -definable. Given a finite set $C \subseteq \mathbb{V}A$, we choose finite sets $D^i \subseteq \mathbb{V}B^i$, for $i < n$, such that $C \subseteq \prod_i D^i$. Let $p_k : \prod_i B^i \rightarrow B^k$ be the projections.

To show that C is L -definably embedded in \mathfrak{A} , we use the characterisation from Lemma 3.2. Since D^k is L -definably embedded in \mathfrak{B}^k , there exist finite sets $\Delta_k \subseteq L$ such that

$$\text{eq}(\Delta_k) \upharpoonright \mathbb{M}J D^k \leq \ker \pi \upharpoonright \mathbb{M}J D^k.$$

As the L -definable languages are closed under inverse relabellings, we can use Lemma 1.3 (b) to find finite sets $\Delta'_k \subseteq L$ such that

$$\text{eq}(\Delta'_k) \leq \ker(\text{eq}(\Delta_k) \circ \mathbb{M}p_k).$$

Setting $\Delta := \Delta'_0 \cup \dots \cup \Delta'_{n-1}$, it follows that

$$\begin{aligned} \text{eq}(\Delta) \upharpoonright \mathbb{M}J C &\leq \ker(\text{eq}(\Delta_k) \circ \mathbb{M}p_k) \upharpoonright \mathbb{M}J C \\ &\leq \ker(\pi \circ \mathbb{M}p_k) \upharpoonright \mathbb{M}J C \\ &= \ker(p_k \circ \pi) \upharpoonright \mathbb{M}J C. \end{aligned}$$

which implies that

$$\begin{aligned} \text{eq}(\Delta) \upharpoonright \mathbb{M}J C &\leq \inf_k \ker(p_k \circ \pi) \upharpoonright \mathbb{M}J C \\ &= \ker(\langle p_k \rangle_k \circ \pi) \upharpoonright \mathbb{M}J C = \ker \pi \upharpoonright \mathbb{M}J C \end{aligned}$$

as desired. \square

Theorem 3.6. *Let \mathcal{R} be the class of all finitely generated, sort-wise weakly finite-dimensional \mathbb{M} -algebras. The class of L -definable \mathbb{M} -algebras forms an \mathcal{R} -variety.*

Proof. We start by proving closure under finitely generated subalgebras of finite products. Let \mathfrak{A} be a finitely generated subalgebra of $\prod_{i < n} \mathfrak{B}^i$ where

each \mathfrak{B}^i is L -definable. Then \mathfrak{A} is finitely generated and, according to Proposition 3.5, it is locally L -definable. Finally, the fact that \mathfrak{A} is sort-wise weakly finite-dimensional follows from Lemma III.1.4.

It therefore remains to prove closure under sort-accumulation points. Let \mathfrak{A} be a finitely generated sort-accumulation point of the class of L -definable algebras. We have to show that \mathfrak{A} is L -definable.

To check that \mathfrak{A} is sort-wise weakly finite-dimensional, we fix a sort $\xi \in \Xi$. By assumption, there exists some L -definable algebra \mathfrak{B} and an E -morphism $q : \mathfrak{B}|_{\{\xi\}} \rightarrow \mathfrak{A}|_{\{\xi\}}$. In particular, there exists an E -morphism $B_\xi \rightarrow A_\xi$. Since B_ξ is weakly finite-dimensional, so is therefore A_ξ .

It remains to show that \mathfrak{A} is locally L -definable. Consider a finite set $C \subseteq \mathbb{V}A$ and a function $\mu : A_\xi \rightarrow \Omega$. Let $\Delta \subseteq \Xi$ be a finite set of sorts such that $\xi \in \Delta$ and $C \subseteq A|_\Delta$. By assumption, we can find an L -definable algebra \mathfrak{B} such that $\mathfrak{A}|_\Delta$ is a quotient of $\mathfrak{B}|_\Delta$. Let $q : \mathfrak{B}|_\Delta \rightarrow \mathfrak{A}|_\Delta$ be an E -morphism. Then $\mathbb{V}q$ is surjective and has a right-inverse. Via the adjunction we obtain a function $f : \mathbb{J}\mathbb{V}A|_\Delta \rightarrow B|_\Delta$ such that $q \circ f = \iota$, where $\iota : \mathbb{J}\mathbb{V} \Rightarrow \text{Id}$ is the counit of the adjunction. Setting $D := f[\mathbb{J}C]$, it follows that

$$\begin{aligned} \mu \circ \pi \upharpoonright \mathbb{M}J C &= \mu \circ \pi \circ \mathbb{M}\iota \upharpoonright \mathbb{M}J C \\ &= \mu \circ \pi \circ \mathbb{M}q \circ \mathbb{M}f \upharpoonright \mathbb{M}J C \\ &= \mu \circ q \circ \pi \circ \mathbb{M}f \upharpoonright \mathbb{M}J C. \end{aligned}$$

This language is L -definable since $(\mu \circ q) \circ \pi$ is L -definable and the class of L -definable languages is closed under inverse relabellings. Consequently, $\mu \circ \pi \upharpoonright \mathbb{M}J C$ is also L -definable. \square

4. Definable Languages

Finally, we are ready to make the connection between definable languages and definable algebras. We start by proving that syntactic algebras and theory algebras are (locally) L -definable.

Lemma 4.1. *Let L be a varietal family of logics. If $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ is an L -definable language with a syntactic algebra, then $\text{Syn}(\kappa)$ is locally L -definable.*

Proof. Let $C \subseteq \mathbb{V}\text{Syn}(\kappa)$ be finite and fix $\mu : \text{Syn}_\xi(\kappa) \rightarrow \Omega$. As $\lambda := \mu \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C : \mathbb{M}\mathbb{J}C \rightarrow \Omega$ is recognised by the restriction $\pi \upharpoonright \mathbb{M}\mathbb{J}C : \mathbb{M}\mathbb{J}C \rightarrow \text{Syn}(\kappa)$, it follows by Proposition III.4.16 that λ is of the form

$$\lambda = \varphi^{-1}[\omega[p_o^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]]],$$

for some morphism $\varphi : \mathbb{M}C \rightarrow \mathbb{M}\Sigma$, contexts p_o, \dots, p_{n-1} , and an Ω -operation $\omega : \Omega^n \rightarrow \Omega$. By the assumed closure properties, languages of this form are L -definable. Consequently, it follows by Theorem 3.4 that C is L -definably embedded in $\text{Syn}(\kappa)$. \square

Theorem 4.2. *Let L be a family of logics such that every L -definable language has a syntactic algebra. The following statements are equivalent.*

- (1) L is varietal.
- (2) For every L -definable language $\kappa : \mathbb{M}\Sigma \rightarrow \Omega$, the syntactic algebra $\text{Syn}(\kappa)$ is locally L -definable.

Proof. (1) \Rightarrow (2) follows by Lemma 4.1. For (2) \Rightarrow (1), fix an L -definable language $\kappa : \mathbb{M}_\xi \Gamma \rightarrow \Omega$. Then $\kappa = \mu \circ \text{syn}_\kappa$, for some $\mu : \text{Syn}_\xi(\kappa) \rightarrow \Omega$. For closure under inverse morphisms, consider $\varphi : \mathbb{M}\Sigma \rightarrow \mathbb{M}\Gamma$. Then

$$\kappa \circ \varphi = \mu \circ \text{syn}_\kappa \circ \varphi,$$

is recognised by $\text{syn}_\kappa \circ \varphi : \mathbb{M}\Sigma \rightarrow \text{Syn}(\kappa)$. Hence, it is L -definable by Theorem 3.4.

For closure under derivatives, consider a context $p \in \mathbb{M}(\Gamma + \square)$. By Proposition III.4.15, $p^{-1}[\kappa]$ is recognised by $\text{syn}_\kappa : \mathbb{M}\Gamma \rightarrow \text{Syn}(\kappa)$. Hence, L -definability follows again by Theorem 3.4. \square

In general there is no reason why a syntactic algebra should be weakly L -definable. Hence, we only obtain local definability. For theory algebras, the situation better. Below we will characterise under which conditions these are indeed L -definable. The proof rests on the following technical result.

Lemma 4.3. *Let L be an \mathbb{M} -compositional family of logics. For every sort-wise finite set $\Delta \subseteq L$ such that $\text{eq}(\Delta)$ is an \mathbb{M} -congruence, the set $\text{rng } \mathbb{V}(\theta_\Delta \circ \text{sing})$ is L -definably embedded in $\Theta_\Delta \Sigma$.*

Proof. Let $\varepsilon : \text{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \text{Id}$ be the unit and the counit of the adjunction. Suppose that $\Sigma = \mathbb{J}X$ and let

$$f := \mathbb{V}(\theta_\Delta \circ \text{sing}) \circ \varepsilon : X \rightarrow \mathbb{V}\Theta_\Delta \Sigma$$

be the morphism corresponding to $\theta_\Delta \circ \text{sing} : \Sigma \rightarrow \Theta_\Delta \Sigma$ via the adjunction. Let $C := \text{rng } f$ and choose a right-inverse $g : C \rightarrow X$ of $f : X \rightarrow C$. Finally, let $i : C \rightarrow \mathbb{V}\Theta_\Delta \Sigma$ be the inclusion and $\pi_o := \pi \circ \mathbb{M}(\iota \circ \mathbb{J}i) : \mathbb{M}\mathbb{J}C \rightarrow \Theta_\Delta \Sigma$ the restriction of the product to $\mathbb{M}\mathbb{J}C$. Then

$$\begin{aligned} \theta_\Delta \circ \mathbb{M}\mathbb{J}g &= \theta_\Delta \circ (\text{flat} \circ \mathbb{M}\text{sing}) \circ \mathbb{M}\mathbb{J}g \\ &= \pi \circ \mathbb{M}\theta_\Delta \circ \mathbb{M}\text{sing} \circ \mathbb{M}(\iota \circ \mathbb{J}\varepsilon \circ \mathbb{J}g) \\ &= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}\mathbb{V}(\theta_\Delta \circ \text{sing}) \circ \mathbb{J}(\varepsilon \circ g)) \\ &= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}(f \circ g)) \\ &= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}i) \\ &= \pi_o. \end{aligned}$$

To show that C is L -definably embedded in $\Theta_\Delta \Sigma$, let $\mu : \Theta_\Delta \Sigma \rightarrow \Omega$. By Lemma 1.3,

$$\text{eq}(\Delta) = \ker \theta_\Delta \leq \ker (\mu \circ \theta_\Delta)$$

implies that the language $\kappa := \mu \circ \theta_\Delta$ is L -definable. Furthermore, we have shown in Lemma 2.7 that the class of L -definable languages is closed under inverse relabellings. Consequently, the language

$$\mu \circ \pi_o = \mu \circ \theta_\Delta \circ \mathbb{M}\mathbb{J}g = \kappa \circ \mathbb{M}\mathbb{J}g,$$

is also L -definable. \square

Theorem 4.4. *Let L be an \mathbb{M} -compositional family of logics such that every L -definable language has a syntactic algebra. The following statements are equivalent.*

- (1) L is varietal.

- (2) The class of L -definable languages is closed under inverse morphisms.
- (3) Every algebra of the form $\Theta_\Delta \Sigma$ is L -definable.
- (4) For every L -definable language $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$, the syntactic algebra $\text{Syn}(\kappa)$ is L -definable.
- (5) For every L -definable language $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$, the syntactic algebra $\text{Syn}(\kappa)$ is locally L -definable.

Proof. (1) \Rightarrow (2) is trivial.

(2) \Rightarrow (1) Closure under inverse morphisms holds by assumption; closure under Ω -operations follows trivially from the definition of weak definability; and closure under derivatives was proved in Theorem 2.11.

(1) \Leftrightarrow (5) has already been established in Theorem 4.2.

(4) \Rightarrow (5) is trivial.

(5) \Rightarrow (4) It remains to show that $\text{Syn}(\kappa)$ is finitely generated and weakly finite-dimensional. For the former, note that $\text{Syn}(\kappa)$ is generated by the finite set $\text{rng}(\text{syn}_\kappa \circ \text{sing})$. For the latter, we fix a sort-wise finite set $\Delta \subseteq L$ such that κ is Δ -definable and $\Theta_\Delta \Sigma$ is defined. Then we can use Theorem III.4.9 to find an E -morphism $\rho : \Theta_\Delta \Sigma \rightarrow \text{Syn}(\kappa)$. Since $\Theta_\Delta \Sigma$ is strongly finite-dimensional, it follows that $\text{Syn}(\kappa)$ is weakly finite-dimensional.

(2) \Rightarrow (3) First, note that $\Theta_\Delta \Sigma$ is generated by the finite set

$$C := \text{rng } \mathbb{V}(\theta_\Delta \circ \text{sing}).$$

Furthermore, we have shown in Lemma 2.3 that it is strongly finite-dimensional. It therefore remains to prove L -definability. By Lemma 4.3, the set C is L -definably embedded in $\Theta_\Delta \Sigma$. As every finite $D \subseteq \Theta_\Delta \Sigma$ is contained in $\langle\langle C \rangle\rangle_{\Theta_\Delta \Sigma} = \Theta_\Delta \Sigma$, we can therefore use Lemma 3.3 (b) to show that every finite subset is L -definably embedded.

(3) \Rightarrow (2) Fix a morphism $\varphi : \mathbb{M} \Sigma \rightarrow \mathbb{M} \Gamma$ and an L -definable language $\kappa : \mathbb{M}_\xi \Gamma \rightarrow \Omega$. We will construct two sort-wise finite sets $\Delta, \Delta' \subseteq L$ such that κ is $\Delta[\Gamma]$ -definable and

$$\text{eq}(\Delta'[\Sigma]) \leq \ker(\text{eq}(\Delta[\Gamma]) \circ \varphi).$$

Then it follows by Lemma 1.3 (b) that $\varphi^{-1}[\kappa]$ is L -definable.

Hence, it remains to find the sets Δ and Δ' . As L is \mathbb{M} -compositional, we can choose a sort-wise finite subset $\Delta \subseteq L$ such that κ is $\Delta[\Gamma]$ -definable and $\text{eq}(\Delta)$ is an \mathbb{M} -congruence. Set

$$f := \theta_\Delta \circ \varphi \circ \text{sing} : \Sigma \rightarrow \Theta_\Delta \Gamma \quad \text{and} \quad C := \text{rng } \mathbb{V} f.$$

Then we have

$$\begin{aligned} \pi \circ \mathbb{M} f &= \pi \circ \mathbb{M}(\theta_\Delta \circ \varphi \circ \text{sing}) \\ &= \theta_\Delta \circ \varphi \circ \text{flat} \circ \mathbb{M} \text{sing} = \theta_\Delta \circ \varphi. \end{aligned}$$

By assumption, C is L -definably embedded in $\Theta_\Delta \Gamma$. We can therefore use Lemma 3.2 to find a sort-wise finite subset $\Psi \subseteq L$ such that

$$\text{eq}(\Psi) \upharpoonright \mathbb{M} \mathbb{J} C \leq \ker(\pi \upharpoonright \mathbb{M} \mathbb{J} C).$$

Let $\Delta_0 \subseteq \Delta$ be the (finite) subset of all formulae whose sort is equal to the sort of some element of C . We have shown in Lemma 2.7 that L -definable languages are closed under inverse relabellings. Therefore, we can use Lemma 1.4 to find a sort-wise finite set $\Psi_\xi \cup \Delta_0 \subseteq \Delta' \subseteq L$ such that

$$\text{eq}(\Delta'[\Sigma]) \leq \ker(\text{eq}(\Psi) \circ \mathbb{M} f).$$

Consequently,

$$\begin{aligned} \text{eq}_\xi(\Delta'[\Sigma]) &\leq \ker(\text{eq}(\Psi) \circ \mathbb{M} f) \\ &\leq \ker(\pi \circ \mathbb{M} f) \\ &= \ker(\theta_\Delta \circ \varphi) = \ker(\text{eq}(\Delta[\Gamma]) \circ \varphi). \quad \square \end{aligned}$$

As a consequence we obtain the following counterpart to Theorem 3.4.

Corollary 4.5. *Let L be an \mathbb{M} -compositional, varietal family of logics such that every L -definable language has a syntactic algebra, and let $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ be a language. The following statements are equivalent.*

- (1) κ is L -definable

- (2) κ is recognised by an L -definable algebra.
 (3) κ is recognised by a sort-wise strongly finite-dimensional L -definable algebra.

Proof. (3) \Rightarrow (2) is trivial and (2) \Rightarrow (1) has already been proved in Theorem 3.4. For (1) \Rightarrow (3), we fix some sort-wise finite $\Delta \subseteq L$ such that κ is Δ -definable and $\Theta_\Delta \Sigma$ exists. By the preceding theorem, the algebra $\Theta_\Delta \Sigma$ is L -definable. According to Lemma 2.3 it is even sort-wise strongly finite-dimensional. Finally, we have seen in Lemma 2.4 that the morphism $\theta_\Delta : \mathbb{M}\Sigma \rightarrow \Theta_\Delta \Sigma$ recognises κ . \square

Under some additional assumptions, the non-trivial implication in the preceding corollary actually characterises \mathbb{M} -compositionality.

Theorem 4.6. *Suppose that M is closed under products with arbitrary morphisms and let L be a family of logics such that every L -definable language has a syntactic algebra and the class of L -definable algebras is closed under \mathcal{T} -generated M -subalgebras of finite products. A logic L is \mathbb{M} -compositional if, and only if, every L -definable language is recognised by a sort-wise strongly finite-dimensional L -definable algebra.*

Proof. (\Rightarrow) follows by Corollary 4.5.

(\Leftarrow) Let $\Phi \subseteq L$ be finite. For every formula $\varphi \in \Phi$, we fix a morphism $\eta_\varphi : \mathbb{M}\Sigma \rightarrow \mathfrak{A}_\varphi$ to a sort-wise strongly finite-dimensional L -definable algebra \mathfrak{A}_φ that recognises the language $\text{Mod}(\varphi)$. Set $\eta := \ker \langle \eta_\varphi \rangle_{\varphi \in \Phi}$. The codomain \mathfrak{A} of η is a finitely generated M -subalgebra of $\prod_\varphi \mathfrak{A}_\varphi$. By Theorem 3.6, \mathfrak{A} is L -definable. Clearly, it is also sort-wise strongly finite-dimensional. Hence, there exist embeddings $e_\xi : A_\xi \rightarrow \Omega^{d_\xi}$ with $d_\xi < \omega$, for every $\xi \in \Xi$. Furthermore, we can find formulae $\psi'_{\xi,i}$, $i < d_\xi$, defining the languages $p_i \circ e_\xi$, where $p_i : \Omega^{d_\xi} \rightarrow \Omega$ is the projection to the i -th component. Since L is closed under inverse relabellings, there exist formulae $\psi_{\xi,i}$ such that

$$s \models \psi_{\xi,i} \quad \text{iff} \quad \mathbb{M}(\eta \circ \text{sing})(s) \models \psi'_{\xi,i}.$$

Set $\Delta := \Phi \cup \{ \psi_{\xi,i} \mid \xi \in \Xi, i < d_\xi \}$. Then Δ is sort-wise finite. By choice of η , there exist functions $\mu_\varphi : A_\xi \rightarrow \Omega$, for $\varphi \in \Phi$, such that

$$\text{Mod}(\varphi) = \mu_\varphi \circ \eta.$$

Furthermore, set $\mu_{\psi_{\xi,i}} := p_i \circ e_\xi$ and $\rho := \langle \mu_\varphi \rangle_{\varphi \in \Delta}$. Note that, for $\xi \in \Xi$,

$$\langle \mu_{\psi_{\xi,i}} \rangle_{i < d_\xi} = \langle p_i \circ e_\xi \rangle_i = \langle p_i \rangle_i \circ e_\xi = e_\xi \in M.$$

By our assumption on M , it follows that $\rho = \langle \langle \mu_\varphi \rangle_\varphi, \langle \mu_{\psi_{\xi,i}} \rangle_{\xi,i} \rangle \in M$. Consequently,

$$\begin{aligned} \text{eq}(\Delta) &= \ker \text{Th}_\Delta \\ &= \ker \langle \text{Mod}(\varphi) \rangle_{\varphi \in \Delta} \\ &= \ker \langle \mu_\varphi \circ \eta \rangle_{\varphi \in \Delta} \\ &= \ker (\langle \mu_\varphi \rangle_{\varphi \in \Delta} \circ \eta) = \ker \eta, \end{aligned}$$

which implies that $\text{eq}(\Delta)$ is an \mathbb{M} -congruence. \square

Remark. The condition on M in the preceding theorem is satisfied in Set (the product of an injective function and an arbitrary one is injective), but not in Pos or Top , for example (where the product is injective, but not always an embedding).

We have seen above that the class of L -definable algebras forms an \mathcal{R} -variety. The next theorem provides a more concrete description of this \mathcal{R} -variety: it is generated by the theory algebras $\Theta_\Delta \Sigma$.

Theorem 4.7. *Let L be a varietal \mathbb{M} -compositional logic and let \mathcal{R} be the class of all finitely generated sort-wise weakly finite-dimensional \mathbb{M} -algebras.*

A sort-wise strongly finite-dimensional \mathbb{M} -algebra \mathfrak{A} is L -definable if, and only if, it belongs to the \mathcal{R} -variety \mathcal{V} generated by all theory algebras of the form $\Theta_\Delta \mathbb{J}X$ where X is some finite set and $\Delta \subseteq L$ a sort-wise finite subfamily such that $\text{eq}(\Delta)$ is an \mathbb{M} -congruence.

Proof. (\Leftarrow) We have seen in Theorem 3.6 that the class of all L -definable algebras forms an \mathcal{R} -variety \mathcal{W} , and in Theorem 4.4 that \mathcal{W} contains all theory algebras. Consequently, $\mathcal{V} \subseteq \mathcal{W}$.

(\Rightarrow) Let \mathfrak{A} be L -definable and fix a finite set $C \subseteq \mathbb{V}A$ of generators. It is sufficient to prove that \mathfrak{A} is a sort-accumulation point of theory algebras. Hence, fix a finite set $X \subseteq \mathcal{E}$ of sorts. W.l.o.g. we may assume that X contains all the sorts in C . Since A_ξ is sort-wise strongly finite-dimensional, there exists an M -morphism $e_\xi : A_\xi \rightarrow \Omega^{d_\xi}$, for some $d_\xi < \omega$. Let $p_i : \Omega^{d_\xi} \rightarrow \Omega$ be the projection to the i -th component and let $\Delta \subseteq L$ be a sort-wise finite set such that Θ_Δ is defined and every language $p_i \circ e_\xi \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C$ with $\xi \in X$ and $i < d_\xi$ is Δ -definable. Then

$$\ker \theta_\Delta \upharpoonright \mathbb{M}\mathbb{J}C = \text{eq}(\Delta) \upharpoonright \mathbb{M}\mathbb{J}C \leq \ker (p_i \circ e_\xi \circ \pi) \upharpoonright \mathbb{M}\mathbb{J}C$$

implies that

$$\begin{aligned} \ker \theta_\Delta \upharpoonright \mathbb{M}\mathbb{J}C &\leq \inf_{\xi, i} \ker (p_i \circ e_\xi \circ \pi) \upharpoonright \mathbb{M}\mathbb{J}C \\ &= \inf_{\xi} \ker (e_\xi \circ \pi) \upharpoonright \mathbb{M}\mathbb{J}C \\ &= \inf_{\xi} \ker \pi \upharpoonright \mathbb{M}\mathbb{J}C \\ &= \ker \pi \upharpoonright \mathbb{M}\mathbb{J}C, \end{aligned}$$

where the second step follows since $\langle p_i \rangle_i = \text{id}$ and the third one since $e_\xi \in M$. Consequently, $\ker \theta_\Delta|_X \leq \ker \pi|_X$ and there exist a morphism $q : \Theta_\Delta \mathbb{J}C|_X \rightarrow \mathfrak{A}|_X$ such that $\pi = q \circ \theta_\Delta$. In particular, $\mathfrak{A}|_X$ is a quotient of $\Theta_\Delta \mathbb{J}C|_X$ and $\Theta_\Delta \mathbb{J}C \in \mathcal{V}$. \square

Corollary 4.8. *Suppose that $\mathcal{D} = \text{Set}$ or $\mathcal{D} = \text{Pos}$, and let $\Omega \in \mathcal{D}$ be a finite set. Let L be a varietal \mathbb{M} -compositional logic, let \mathcal{S} be the class of all theory algebras $\Theta_\Delta \Sigma$, and \mathcal{R} the class of all finitary \mathbb{M} -algebras. An algebra $\mathfrak{A} \in \mathcal{R}$ is L -definable if, and only if, it satisfies every \mathbb{M} -inequality in $\text{Th}(\mathcal{S})$.*

Proof. We distinguish three cases. First assume that $\mathcal{D} = \text{Pos}$ and Ω is not an antichain. Then being weakly finite-dimensional is the same as being

finite. Let \mathcal{V} be the \mathcal{R} -variety of all L -definable algebras. By Theorem 4.7, \mathcal{V} is the smallest \mathcal{R} -variety containing \mathcal{S} . By Proposition III.7.3, the class $\mathcal{W} := \text{Mod}(\text{Th}(\mathcal{S}))$ is also a \mathcal{R} -variety containing \mathcal{S} . Consequently, $\mathcal{V} \subseteq \mathcal{W}$. Furthermore, $\mathcal{S} \subseteq \mathcal{V}$ implies $\text{Th}(\mathcal{S}) \supseteq \text{Th}(\mathcal{V})$. Hence, it follows by Corollary III.7.5 that

$$\mathcal{W} = \text{Mod}(\text{Th}(\mathcal{S})) \subseteq \text{Mod}(\text{Th}(\mathcal{V})) = \mathcal{V}.$$

If $\mathcal{D} = \text{Set}$, we can use the same proof. It therefore remains to consider the case where $\mathcal{D} = \text{Pos}$ and Ω is an antichain. Since \mathbb{M} is polynomial, it follows that every set of the form $\mathbb{M}\Sigma$, for $\Sigma \in \text{Alph}$, is also an antichain. Furthermore, every algebra \mathfrak{A} recognising some language $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ must be an antichain as, otherwise, there can be no functions $A \rightarrow \Omega$. Consequently, the problem reduces to the case where $\mathcal{D} = \text{Set}$. \square

The following theorem summarises our various characterisations of when a language is definable in a given logic. It can be considered the main result of this article. Of these characterisations, (8) and (10) are the most useful; (8) mainly when trying to prove that a language is L -definable and (10) when devising a decision procedure for L -definability. Of course, for the latter one has to first determine the set of inequalities in question. Depending on the logic L this can be a highly non-trivial task.

Theorem 4.9. *Let L be an \mathbb{M} -compositional varietal family of logics such that every L -definable language has a syntactic algebra, and let $\kappa : \mathbb{M}_\xi \Sigma \rightarrow \Omega$ be an Ω -language that has a syntactic algebra. The following statements are equivalent.*

- (1) κ is L -definable.
- (2) κ is recognised by some L -definable algebra.
- (3) $\text{Syn}(\kappa)$ is L -definable.
- (4) $\text{Syn}(\kappa)$ is a quotient of $\Theta_\Delta \Gamma$, for some Δ and Γ .
- (5) $\text{syn}_\kappa = \rho \circ \theta_\Delta$, for some Δ and a surjective morphism $\rho : \Theta_\Delta \Sigma \rightarrow \text{Syn}(\kappa)$.
- (6) κ is recognised by $\Theta_\Delta \Gamma$, for some Δ and Γ .
- (7) $\theta_\Delta : \mathbb{M}\Sigma \rightarrow \Theta_\Delta \Sigma$ recognises κ , for some Δ .

(8) $\text{eq}(\Delta) \leq \ker \kappa$, for some Δ .

(9) $\text{eq}(\Delta) \leq \ker \text{syn}_\kappa$, for some Δ .

If $\mathcal{D} = \text{Set}$ or $\mathcal{D} = \text{Pos}$, and if $\Omega \in \mathcal{D}$ is finite, the above statements are equivalent to the following one.

(10) $\text{Syn}(\kappa)$ satisfies all \mathbb{M} -inequalities $s \leq t$ that hold in every theory algebra $\Theta_\Delta \Gamma$.

(Here Δ ranges over sort-wise finite subsets of L and Γ ranges over alphabets.)

Proof. (1) \Leftrightarrow (2) has been proved in Corollary 4.5; (1) \Leftrightarrow (8) in Lemma 1.3; (3) \Leftrightarrow (10) follows by Corollary 4.8.

(1) \Rightarrow (7) was proved in Lemma 2.4.

(7) \Rightarrow (8) Suppose that $\kappa = \mu \circ \theta_\Delta$ for some function $\mu : \Theta_\Delta \Sigma \rightarrow \Omega$. Then

$$\text{eq}(\Delta) = \ker \theta_\Delta \leq \ker (\mu \circ \theta_\Delta) = \ker \kappa.$$

(1) \Rightarrow (9) follows by Corollary 2.5.

(9) \Leftrightarrow (5) Note that $\text{eq}(\Delta) = \ker \theta_\Delta$. For a finite set $\Delta \subseteq L$, it therefore follows that

$$\begin{aligned} \text{eq}(\Delta) \leq \ker \text{syn}_\kappa & \quad \text{iff} \quad \ker \theta_\Delta \leq \ker \text{syn}_\kappa \\ & \quad \text{iff} \quad \text{syn}_\kappa = \rho \circ \theta_\Delta, \quad \text{for some morphism } \rho, \end{aligned}$$

where the fact that ρ is a morphism of \mathbb{M} -algebras follows by Lemma I.5.6.

(5) \Rightarrow (4) is trivial.

(4) \Rightarrow (6) Since $\text{syn}_\kappa : \mathbb{M}\Sigma \rightarrow \text{Syn}(\kappa)$ recognises κ , the claim follows by Lemma III.5.7.

(6) \Rightarrow (2) follows by Theorem 4.4 (3).

(1) \Rightarrow (3) follows by Theorem 4.4 (4).

(3) \Rightarrow (2) holds as $\text{syn}_\kappa : \mathbb{M}\Sigma \rightarrow \text{Syn}(\kappa)$ recognises κ . \square

Our algebraic framework works for logics that are varietal and \mathbb{M} -compositional. In order to simplify proofs of these two properties, let us introduce a strong form of compositionality that implies both of them.

Definition 4.10. A logic L is *strongly \mathbb{M} -compositional* if, for every finite set $\Phi \subseteq L$, there exists a sort-wise finite set $\Delta \supseteq \Phi \subseteq L$ with the following property. For every $\varphi \in \Delta$, there exists a finite subset $\Delta_\varphi \subseteq \Delta$ and a formula $\psi \in L[\Omega]$ satisfying

$$\text{Mod}(\varphi) \circ \text{flat} = \text{Mod}(\psi) \circ \mathbb{M}\text{Th}_{\Delta_\varphi}.$$

Remark. Note that \mathbb{M} -compositionality implies that

$$\mathbb{M}\text{Th}_\Delta(s) \quad \text{determines} \quad \text{Th}_\Delta(\text{flat}(s)).$$

Strong \mathbb{M} -compositionality strengthens this implication by requiring that the theory of $\text{flat}(s)$ can be computed by an L -formula.

Proposition 4.11. Every strongly \mathbb{M} -compositional logic L is varietal and \mathbb{M} -compositional.

Proof. We start by proving that L is \mathbb{M} -compositional. Given a finite set $\Phi \subseteq L$, fix a sort-wise finite set $\Delta \supseteq \Phi$ as in the definition of strong \mathbb{M} -compositionality. We claim that $\text{eq}(\Delta)$ is an \mathbb{M} -congruence, i.e., that

$$\ker \mathbb{M}\text{eq}(\Delta) \leq \ker (\text{eq}(\Delta) \circ \text{flat}).$$

By strong \mathbb{M} -compositionality we can choose, for every formula $\varphi \in \Delta$, a formula ψ_φ and a finite set $\Delta_\varphi \subseteq \Delta$ such that

$$\text{Mod}(\varphi) \circ \text{flat} = \text{Mod}(\psi_\varphi) \circ \mathbb{M}\text{Th}_{\Delta_\varphi}.$$

It follows that

$$\text{eq}(\Delta) \circ \text{flat} = \left\langle \text{Mod}(\varphi) \right\rangle_{\varphi \in \Delta} \circ \text{flat} = \left\langle \text{Mod}(\psi_\varphi) \circ \mathbb{M}\text{Th}_{\Delta_\varphi} \right\rangle_{\varphi \in \Delta}.$$

Since $\Delta_\varphi \subseteq \Delta$, there are functions p_φ with $\text{Th}_{\Delta_\varphi} = p_\varphi \circ \text{eq}(\Delta)$. Consequently,

$$\begin{aligned} \text{eq}(\Delta) \circ \text{flat} &= \left\langle \text{Mod}(\psi_\varphi) \circ \mathbb{M}(p_\varphi \circ \text{eq}(\Delta)) \right\rangle_{\varphi \in \Delta} \\ &= \left\langle \text{Mod}(\psi_\varphi) \circ \mathbb{M}p_\varphi \right\rangle_{\varphi \in \Delta} \circ \mathbb{M}\text{eq}(\Delta), \end{aligned}$$

which implies that $\ker \mathbb{M}eq(\Delta) \leq \ker (eq(\Delta) \circ \text{flat})$.

It remains to prove that L is varietal. By Theorem 4.4, it is sufficient to show that every theory algebra $\Theta_\Delta \Sigma$ is L -definable. First, note that $\Theta_\Delta \Sigma$ is finitely generated by $C := \text{rng}(\theta_\Delta \circ \text{sing})$. Let $i : C \rightarrow \Theta_\Delta \Sigma$ be the inclusion map and let $q : \Sigma \rightarrow C$ be the function such that $i \circ q = \theta_\Delta \circ \text{sing}$. We have to show that, for every function $\mu : \Theta_\Delta \Sigma \rightarrow \Omega$, the composition $\mu \circ \pi \circ \mathbb{M}i : \mathbb{M}C \rightarrow \Omega$ is L -definable.

Hence, fix $\mu : \Theta_\Delta \Sigma \rightarrow \Omega$. Then $\mu \circ \theta_\Delta : \mathbb{M}\Sigma \rightarrow \Omega$ is Δ -definable. Let $\varphi \in \Delta[\Omega]$ be a formula with $\text{Mod}(\varphi) = \mu \circ \theta_\Delta$. By strong \mathbb{M} -compositionality, there is some formula ψ and a set $\Delta_o \subseteq \Delta$ such that

$$\text{Mod}(\varphi) \circ \text{flat} = \text{Mod}(\psi) \circ \mathbb{M}\text{Th}_{\Delta_o}.$$

Since $\Delta_o \subseteq \Delta$, we can find a function p with $\text{Th}_{\Delta_o} = p \circ \theta_\Delta$. It follows that

$$\begin{aligned} (\mu \circ \pi \circ \mathbb{M}i) \circ \mathbb{M}q &= \mu \circ \pi \circ \mathbb{M}(\theta_\Delta \circ \text{sing}) \\ &= \mu \circ \theta_\Delta \circ \text{flat} \circ \mathbb{M}\text{sing} \\ &= \text{Mod}(\varphi) \circ \text{flat} \circ \mathbb{M}\text{sing} \\ &= \text{Mod}(\psi) \circ \mathbb{M}\text{Th}_{\Delta_o} \circ \mathbb{M}\text{sing} \\ &= \text{Mod}(\psi) \circ \mathbb{M}(p \circ \theta_\Delta) \circ \mathbb{M}\text{sing} \\ &= \text{Mod}(\psi) \circ \mathbb{M}(p \circ i \circ q). \end{aligned}$$

Since q (and hence $\mathbb{M}q$) is surjective, it follows that

$$\mu \circ \pi \circ \mathbb{M}i = \text{Mod}(\psi) \circ \mathbb{M}(p \circ i).$$

Thus $\mu \circ \pi \circ \mathbb{M}i$ is an inverse relabelling of an L -definable language and, therefore, also L -definable. \square

Notes

Over the years several abstract logical frameworks have been in use, most of them not developed enough to be ever published. Among the major ones are the framework for abstract model theory proposed by Barwise (see,

e.g., [2]), the notion of an *abstract elementary class* introduced by Shelah (see, e.g., [1]), and the theory of *institutions* developed by Goguen and Burstall (see, e.g., [13]). The framework presented here is somewhat similar to the latter, the main difference being that we do not equip our class of models with the structure of a category.

V. Trees

1. Monads and Logics for Trees and Forests

It is time to apply the abstract theory we have developed so far to a concrete setting. Let us take an in-depth look at languages of infinite trees, which exhibit many of the complications not arising in simpler settings. We start by introducing the corresponding monad. It turns out that in different situations slightly different versions of this monad will be needed. Therefore, we define several variants starting with the most general one and then deriving the others from it.

Convention. *In this chapter we will work exclusively in the category $\mathcal{D} = \text{Pos}$.*

Definition 1.1. Let X be a countably infinite set of variables.

(a) For sorts we will be using the set $\Xi := \wp_\omega(X)$ of all finite sets of variables. Sometimes we also take the larger set $\Xi_+ := \wp(X)$. The *arity* of an element a of sort ξ is the cardinality $|\xi|$.

(b) Let $A \in \mathcal{D}^\Xi$. An A -labelled (nondeterministic) rooted graph is a countable directed graph g where the vertices are labelled by elements of A and the edges by elements of X in such a way that, if a vertex v has a label a of sort ξ , then the labels of all outgoing edges belong to ξ . We assume that all outgoing edges with the same label are ordered from left-to-right, but there is no ordering between edges whose label is different. In addition, some of the vertices of g are marked as *roots*, which are also assumed to be ordered left-to-right, and we require that every vertex can be reached by a path from some root.

(c) We denote by \mathbb{G} the polynomial functor where $\mathbb{G}_\xi A$ is the set of all $(A + \xi)$ -labelled rooted graphs where the labels in ξ are assumed to have sort \emptyset and every variable $x \in \xi$ appears as the label of at least one vertex.

(d) We denote the set of vertices of a graph $g \in \mathbb{G}_\xi A$ by $\text{dom}_+(g)$ and by $\text{dom}(g)$ the subset of vertices with a label in A . Vertices with a label in ξ are called *holes*. An x -edge is an edge with label x . If there is an x -edge from v to u , we call u an x -successor of v and v an x -predecessor of u .

In addition, we use the following functors derived from \mathbb{G} .


Definition 1.2. Let $A \in \mathcal{D}^\Xi$.

(a) We denote by $\mathbb{R}A$ the subset of $\mathbb{G}A$ consisting of all graphs with a single root such that the root is not labelled by a variable and such that, for every vertex v of sort ξ and every $x \in \xi$, there is exactly one outgoing edge with label x . We call such graphs *deterministic*.


(b) $\mathbb{T}^\times A \subseteq \mathbb{R}A$ is the set of deterministic graphs that form trees, and $\mathbb{T}_\xi A \subseteq \mathbb{T}^\times A$ is the set of trees where every variable $x \in \xi$ appears exactly once. We call the elements of $\mathbb{T}A$ *linear ranked trees* and those of $\mathbb{T}^\times A$ *non-linear ranked trees*.

(c) $\mathbb{F}^\times A \subseteq \mathbb{G}A$ is the set of nondeterministic graphs that form forests, and $\mathbb{F}_\xi A \subseteq \mathbb{F}^\times A$ is the set of forests where every variable $x \in \xi$ appears exactly once. We call the elements of $\mathbb{F}A$ *linear (unranked) forests* and those of $\mathbb{F}^\times A$ *non-linear (unranked) forests*. An *unranked tree* is an unranked forests with a single connected component.

Remark. (a) To denote finite trees we will frequently use term notation. For instance, $a(x, b(y))$ is the tree



(We assumed a fixed ordering on the variables, so that we can speak of the first successor, the second one, etc.) Similarly, we write forests using an addition operation $+$ for disjoint union and a constant o for the empty forest. For instance, $a(x + c, o, b(y)) + b(x)$ denotes the forest



(b) For unranked forests, the edge labellings are usually only necessary when we consider forests in $s \in \mathbb{F}FA$ where the component forests $s(v)$ might contain several different variables. For forests $s \in \mathbb{F}\Sigma$ labelled by some alphabet Σ on the other hand, we could do without labelled edges. For this reason, we frequently work with alphabets where every letter has arity 1. In particular, we will frequently write $\mathbb{F}\Sigma$, for one-sorted sets $\Sigma \in \mathcal{D}$, which we implicitly regard as Ξ -sorted ones by assigning to each letter $c \in \Sigma$ the sort $\{x\}$, for some fixed variable x .

To turn the functors we have just introduced into monads we use the following operations on graphs.

Definition 1.3. Let $A \in \mathcal{D}^\Xi$.

(a) The *unravelling* of a graph $g \in \mathbb{G}A$ is the graph $\text{un}(g)$ whose vertices correspond to all finite paths in g that start at some root. The label of such a path is the label of its end-point in g . A path q is the x -successor of a path p if q can be obtained from p by appending a single x -edge. The roots of $\text{un}(g)$ are all the paths of length 0 (see Figure 1 for an example).

(b) For $a \in A_\xi$, we denote by $\text{sing}(a)$ the graph g with a single root r labelled a which has one x -successor u_x , for every $x \in \xi$, with label $g(u_x) = x$. In term notation, $\text{sing}(a) = a(x_0, \dots, x_{n-1})$ for $\xi = \{x_0, \dots, x_{n-1}\}$.




Figure 1.: A graph and its unravelling.

(c) The *flattening* $\text{flat}(g)$ of a graph $g \in \mathbb{G}\mathbb{G}A$ is defined as follows (see Figure 2 for an example). Let h be the disjoint union of all graphs $g(v)$, for $v \in \text{dom}_+(g)$, and let R be the binary relation on $\text{dom}_+(h)$ containing all pairs $\langle u, v \rangle$ such that

- $h(u) = x$ for some variable x ,
- u is a leaf of some component $g(w)$ with $w \in \text{dom}(g)$,
- v is a root of some component $g(w')$ where w' is the x -successor of v in g .

Then $\text{flat}(g)$ is the graph obtained from h by

- deleting every vertex corresponding to a hole in $g(v)$ with $v \in \text{dom}(g)$ and
- for every x -predecessor u' of a deleted vertex u , adding x -edges from u' to all vertices v with $\langle u, v \rangle \in R$.

(d) We denote the restrictions of $\text{flat} : \mathbb{G}\mathbb{G}A \rightarrow \mathbb{G}A$ and $\text{sing} : A \rightarrow \mathbb{G}A$ to the subfunctors \mathbb{R} , \mathbb{T} , and \mathbb{F} also by flat and sing . In cases where we want to distinguish between these versions, we add the functor as a superscript: $\text{flat}^{\mathbb{G}}$, $\text{flat}^{\mathbb{R}}$, $\text{flat}^{\mathbb{T}}$, etc.

(e) Finally, for the functors \mathbb{T}^\times and \mathbb{F}^\times we set $\text{flat}^\times := \text{un} \circ \text{flat}$ (restricted to $\mathbb{T}^\times \mathbb{T}^\times A$ or $\mathbb{F}^\times \mathbb{F}^\times A$, respectively).

It is now straightforward (but a bit tedious) to check the monad laws.

Proposition 1.4. \mathbb{G} , \mathbb{R} , \mathbb{T} , and \mathbb{F} form monads on Pos^Ξ with multiplication flat and unit morphism sing .

Proof. It is sufficient to prove the claim for \mathbb{G} . The other three cases then follow from Lemma ?? (a) by considering the corresponding inclusion morphisms. Clearly, we have $\text{flat} \circ \text{sing} = \text{id}$ and $\text{flat} \circ \mathbb{G}\text{sing} = \text{id}$. For associativity, let $G \in \mathbb{G}\mathbb{G}\mathbb{G}A$. By Proposition 1.3.7 and the Remark on page 38, we obtain




Figure 2.: A graph $g \in \mathbb{G}\mathbb{G}\{a, b, c\}$ and its flattening. The sorts of a, b, c are $\{u, v\}$, $\{u\}$, and \emptyset , respectively. The roots are marked by an incoming arrow.

labelling-preserving functions

$$\begin{aligned}
\mu &: \text{dom}(\text{flat}(G)) \rightarrow \sum_{v \in \text{dom}(G)} \text{dom}(G(v)), \\
\mu' &: \text{dom}(\text{flat}(\text{flat}(G))) \rightarrow \sum_{w \in \text{dom}(\text{flat}(G))} \text{dom}(\text{flat}(G)(w)), \\
\mu_v &: \text{dom}(\text{flat}(G(v))) \rightarrow \sum_{u \in \text{dom}(G(v))} \text{dom}(G(v)(u)), \\
\mu'' &: \text{dom}(\text{flat}(\mathbb{G}\text{flat}(G))) \rightarrow \sum_{v \in \text{dom}(\mathbb{G}\text{flat}(G))} \text{dom}(\text{flat}(G(v))),
\end{aligned}$$

for $v \in \text{dom}(G)$. As the flattening function for graphs does not erase or duplicated vertices (this is different for \mathbb{T}^\times and \mathbb{F}^\times), these functions are all bijective. Furthermore, it is straightforward to check that they preserve the edge relation. Composing these we obtain a bijection

$$\begin{aligned}
\text{dom}(\text{flat}(\text{flat}(G))) &\rightarrow \sum_{w \in \text{dom}(\text{flat}(G))} \text{dom}(\text{flat}(G)(w)) \\
&\rightarrow \sum_{(v,u) \in \sum_{v \in \text{dom}(G)} \text{dom}(G(v))} \text{dom}(G(v)(u)) \\
&\rightarrow \sum_{v \in \text{dom}(G)} \sum_{u \in \text{dom}(G(v))} \text{dom}(G(v)(u)) \\
&\rightarrow \sum_{v \in \text{dom}(\mathbb{G}\text{flat}(G))} \sum_{u \in \text{dom}(G(u))} \text{dom}(G(v)(u)) \\
&\rightarrow \sum_{v \in \text{dom}(\mathbb{G}\text{flat}(G))} \text{dom}(\text{flat}(G(v))) \\
&\rightarrow \text{dom}(\text{flat}(\mathbb{G}\text{flat}(G)))
\end{aligned}$$

that preserves the labelling and the edge relation. This implies that the graphs $\text{flat}(\text{flat}(G))$ and $\text{flat}(\mathbb{G}\text{flat}(G))$ are in fact equal. \square

The non-linear cases \mathbb{T}^\times and \mathbb{F}^\times have to be dealt with separately since, in contrast to their linear counterparts, they do *not* form a submonad of \mathbb{R} and \mathbb{G} , respectively. Instead they are *quotients*.

Lemma 1.5. $\langle \mathbb{T}^\times, \text{flat}^\times, \text{sing} \rangle$ and $\langle \mathbb{F}^\times, \text{flat}^\times, \text{sing} \rangle$ form monads and the unravelling maps $\text{un} : \mathbb{R} \Rightarrow \mathbb{T}^\times$ and $\text{un} : \mathbb{G} \Rightarrow \mathbb{F}^\times$ are morphisms of monads.

Proof. By Lemma ??, it is sufficient to check that

$$\text{sing} = \text{un} \circ \text{sing} \quad \text{and} \quad \text{flat}^\times \circ \text{un} \circ \mathbb{R}\text{un} = \text{un} \circ \text{flat}$$

(and similarly for \mathbb{F}^\times and \mathbb{G} .) The first equation immediately follows from the fact that $\text{un}(\text{sing}(a)) = \text{sing}(a)$. For the second one, note that the vertices of $\text{un}(\text{flat}(g))$ correspond to the finite paths of $\text{flat}(g)$, while those of $\text{un}(\text{flat}(\text{un}(\mathbb{R}\text{un}(g))))$ correspond to those of $\text{flat}(\text{un}(\mathbb{R}\text{un}(g)))$. Furthermore, every path α in a graph of the form $\text{flat}(h)$ corresponds to a path $(v_n)_n$ of h and a family of paths β_n of $h(v_n)$ such that α can be identified with the concatenation $\beta_0\beta_1\dots$. Finally, a path in $\text{un}(h)$ is the same as a path in h . Consequently, each path of $\text{flat}(\text{un}(\mathbb{R}\text{un}(g)))$ corresponds to (i) a path of g together with (ii) a family of paths in some components $g(v)$ as above. This correspondence induces a bijection between

$$\text{dom}(\text{un}(\text{flat}(g))) \quad \text{and} \quad \text{dom}(\text{un}(\text{flat}(\text{un}(\mathbb{R}\text{un}(g)))).$$

As this bijection preserves the labelling it follows that

$$\text{un}(\text{flat}(g)) = \text{un}(\text{flat}(\text{un}(\mathbb{R}\text{un}(g))). \quad \square$$

Logics for Forests

The main logics we are interested in are *first-order logic* FO and *monadic second-order logic* MSO. To define the satisfaction relation between formulae of these logics and elements of $\mathbb{G}\Sigma$, we encode a graph $g \in \mathbb{G}_\xi\Sigma$ as the structure

$$\mathfrak{G} = \langle V, \leq_{\text{pf}}, \leq_{\text{so}}, (\text{suc}_x)_{x \in X}, (P_c)_{c \in \Sigma + \xi}, R \rangle$$

where

- ♦ $V := \text{dom}_+(g)$ is the set of vertices,
- ♦ \leq_{pf} the reflexive transitive closure of the edge relation,

- \leq_{so} is the left-to-right ordering on the successors of each vertex (where successors of distinct vertices are incomparable, and the same for x -successors and y -successors with $x \neq y$),
- suc_x is the x -successor relation,
- $P_c := g^{-1}(c)$ the set of all vertices with label $c \in \Sigma + \xi$, and
- R is a unary relation that marks the roots.

Depending on which of the above relations we allow in a formula, we obtain several variants of first-order or monadic second-order logic.

Definition 1.6. (a) $\text{MSO}[\leq_{so}]$, monadic second-order logic with *successor ordering*, is MSO over structures of the above form where every relation can be used.

(b) $\text{MSO}[\leq_{pf}]$, monadic second-order logic with *forest ordering*, or *prefix ordering*, is MSO over the above structures, but without the relation \leq_{so} .

(c) $\text{MSO}[\text{suc}]$, monadic second-order logic with *successor*, is MSO over the above structures, but without the relations \leq_{pf} and \leq_{so} .

(d) We define the same variants for first-order logic FO and *counting monadic second-order logic* CMSO. The latter is the extension of MSO by atomic formulae of the form

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

for set variables X and numbers $0 < k < m < \omega$. By definition, a formula of the first form hold if the set X is finite, and a formula of the second form holds if X is finite and its cardinality is congruent k modulo m .

(e) For a number $m < \omega$, we denote by $\text{MSO}_m[\leq_{so}]$, $\text{MSO}_m[\leq_{pf}]$, ... the restrictions of the corresponding logics to formulae of quantifier-rank at most m (counting both first-order and second-order quantifiers).

Let us give a quick overview over several well-known tools that help us study the expressive power of these logics on forests. We use the following notation.

Definition 1.7. Let L be one of the above logics. For two structures \mathcal{A} and \mathcal{B} and tuples $\bar{a} \in A^k$, $\bar{b} \in B^k$, $\bar{P} \in \wp(A)^l$, $\bar{Q} \in \wp(B)^l$, we write

$$\mathcal{A}, \bar{P}\bar{a} \equiv_L \mathcal{B}, \bar{Q}\bar{b} \quad : \text{iff} \quad [\mathcal{A} \models \varphi(\bar{P}, \bar{a}) \Leftrightarrow \mathcal{B} \models \varphi(\bar{Q}, \bar{b})],$$

for all $\varphi \in L$.

We can compute this relation by induction on m as follows.

Proposition 1.8. Let \mathcal{A} and \mathcal{B} be two structures over a finite, relational signature, and let $\bar{a} \in A^k$ and $\bar{b} \in B^k$ be tuples of elements. Then

$$\mathcal{A}, \bar{a} \equiv_{\text{FO}_{m+1}} \mathcal{B}, \bar{b}$$

if, and only if,

- for every $a' \in A$, there is some $b' \in B$ with $\mathcal{A}, \bar{a}a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{b}b'$,
- for every $b' \in B$, there is some $a' \in A$ with $\mathcal{A}, \bar{a}a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{b}b'$.

Proposition 1.9. Let L be one of the above variants of MSO or CMSO, and let L_m be the corresponding fragment of restricted quantifier-rank. Let \mathcal{A} and \mathcal{B} be two structures over a finite, relational signature, and let $\bar{a} \in A^k$, $\bar{b} \in B^k$, $\bar{P} \in \wp(A)^l$, $\bar{Q} \in \wp(B)^l$ be tuples of elements and sets. Then

$$\mathcal{A}, \bar{P}\bar{a} \equiv_{\text{FO}_{m+1}} \mathcal{B}, \bar{Q}\bar{b}$$

if, and only if,

- for every $a' \in A$, there is some $b' \in B$ with $\mathcal{A}, \bar{P}\bar{a}a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{Q}\bar{b}b'$,
- for every $b' \in B$, there is some $a' \in A$ with $\mathcal{A}, \bar{P}\bar{a}a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{Q}\bar{b}b'$,
- for every $P' \subseteq A$, there is some $Q' \subseteq B$ with $\mathcal{A}, \bar{P}P'a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{Q}Q'b'$,
- for every $Q' \subseteq B$, there is some $P' \subseteq A$ with $\mathcal{A}, \bar{P}P'a' \equiv_{\text{FO}_m} \mathcal{B}, \bar{Q}Q'b'$.

This method is called a *back-and-forth argument*. As the readers can verify in the following sections for themselves, proofs containing such arguments are often quite tedious. When applicable, so-called composition theorems provide a much simpler way to compute the relation \equiv_L . Let us present two such theorems. They are based on the following two operations.

Definition 1.10. Let L be one of the above logics and let Σ and Γ be finite relational signatures.

(a) A *simple L -interpretation* is a unary operation that is defined by a tuple

$$\tau = \langle \delta(x), (\varphi_R(\bar{x}))_{R \in \Gamma} \rangle$$

of L -formulae as follows. Given a Σ -structure \mathfrak{A} it produces the Γ -structure

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

with universe

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

and relations

$$\varphi_R^{\mathfrak{A}} := \{ \bar{a} \in A^n \mid \mathfrak{A} \models \varphi_R(\bar{a}) \}, \quad \text{for } R \in \Gamma$$

(where n is the arity of R).

(b) Let \mathfrak{I} be a Γ -structure and $(\mathfrak{A}^i)_{i \in I}$ a family of Σ -structures indexed by the elements of \mathfrak{I} . The *generalised sum*

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

is the $(\Sigma + \Gamma + \{\sim\})$ -structure with universe

$$\sum_{i \in I} A^i := \{ \langle i, a \rangle \mid i \in I, a \in A^i \}$$

and relations

$$\begin{aligned} R &:= \{ \langle \langle i, a_0 \rangle, \dots, \langle i, a_{n-1} \rangle \rangle \mid i \in I, \langle a_0, \dots, a_{n-1} \rangle \in R^{\mathfrak{A}^i} \}, \\ S &:= \{ \langle \langle i_0, a_0 \rangle, \dots, \langle i_{n-1}, a_{n-1} \rangle \rangle \mid \langle i_0, \dots, i_{n-1} \rangle \in S^{\mathfrak{I}}, a_j \in A^{i_j} \}, \\ \sim &:= \{ \langle \langle i, a \rangle, \langle i, b \rangle \rangle \mid i \in I, a, b \in A^i \}, \end{aligned}$$

for $R \in \Sigma$ and $S \in \Gamma$.

The corresponding composition results can both be proved by a (mostly) straightforward induction on the quantifier-rank.

Proposition 1.11. Let L be one of the above logics, let L_m be the corresponding fragment of restricted quantifier-rank, and let $\tau = \langle \delta(x), (\varphi_R(\bar{x}))_{R \in \Gamma} \rangle$ be a simple L_k -interpretation. For every formula $\varphi \in L_m$, there exists a formula $\varphi^\tau \in L_{m+k}$ such that

$$\tau(\mathfrak{A}) \models \varphi \quad \text{iff} \quad \mathfrak{A} \models \varphi^\tau, \quad \text{for all structures } \mathfrak{A}.$$

Theorem 1.12. Let L be one of the above variants of MSO or CMSO, and let L_m be the corresponding fragment of restricted quantifier-rank. For every formula $\varphi \in L_m$, there exist finitely many formulae $\chi_0, \dots, \chi_{n-1} \in L_m$ and $\psi \in L$ such that

$$\sum_{i \in \mathfrak{I}} \mathfrak{A}^i \models \varphi \quad \text{iff} \quad \langle \mathfrak{I}, \llbracket \chi_0 \rrbracket, \dots, \llbracket \chi_{n-1} \rrbracket \rangle \models \psi,$$

where $\llbracket \chi \rrbracket := \{ i \in I \mid \mathfrak{A}^i \models \chi \}$.

2. Finite Forests

Before turning to the general case, let us take a brief look at languages of *finite forests*.

Definition 2.1. We denote by $\mathbb{F}^{\text{fin}} A$ the set of all *finite forests* in $\mathbb{F}^\times A$.

Definition 2.2. Let \mathfrak{A} be an \mathbb{F}^{fin} -algebra.

(a) The *horizontal semigroup* associated with \mathfrak{A} is the structure $\langle A_\emptyset, + \rangle$ where the *horizontal product* $+$ is defined by

$$a + b := \pi(s), \quad \text{where } s \text{ is the forest consisting of the two trees } \text{sing}(a) \text{ and } \text{sing}(b) \text{ (in this order).}$$

(b) The *vertical semigroup* associated with \mathfrak{A} is the structure $\langle A_{\{x\}}, \cdot \rangle$ (for some fixed variable $x \in X$) where the *vertical product* \cdot is defined by

$$a \cdot b := \pi(s), \quad \text{where } s \text{ is the tree consisting of a root labelled } a \text{ and a leaf labelled } b.$$

Proposition 2.3.

- (a) Every finitary \mathbb{F}^{fin} -algebra is $\text{MSO}[\leq_{\text{so}}]$ -definable.
- (b) Every finitary \mathbb{F}^{fin} -algebra with a commutative horizontal semigroup is $\text{CMSO}[\leq_{\text{pf}}]$ -definable.
- (c) Every finitary \mathbb{F}^{fin} -algebra whose horizontal semigroup is commutative and aperiodic is $\text{MSO}[\leq_{\text{pf}}]$ -definable.

Proof. (a) Let \mathfrak{A} be finitary, $C \subseteq A$ a finite set of generators, and $c \in A_\xi$. We have to construct a formula φ_c such that

$$s \models \varphi_c \quad \text{iff} \quad \pi(s) = c, \quad \text{for } s \in \mathbb{F}_\xi^{\text{fin}} C.$$

The formula φ_c guesses sets $(Z_a)_a$ such that

$$Z_a := \{v \in \text{dom}(s) \mid \eta(s|_v) = a\}, \quad \text{for } a \in \bigcup_{\zeta \subseteq \xi} A_\zeta,$$

and then it checks that these guesses are correct. For a finite forest s we can do so bottom-up starting at the leaves. Let v be a vertex with successors u_0, \dots, u_{m-1} and label $s(v) = a$. Suppose that $u_i \in Z_{b_i}$. Then we have to check that $v \in Z_{a(b_0 + \dots + b_{m-1})}$. To do so, we have to be able to evaluate the sum $b_0 + \dots + b_{m-1}$. Since $\langle \bigcup_{\zeta \subseteq \xi} A_\zeta, + \rangle$ is a finite semigroup, there exist MSO -formula $(\vartheta_a)_a$ such that

$$\langle \{u_0, \dots, u_{m-1}\}, \leq_{\text{so}}, \bar{Z} \rangle \models \vartheta_a \quad \text{iff} \quad b_0 + \dots + b_{m-1} = a.$$

Consequently, we can set

$$\varphi_c := \exists \bar{Z} \left[Z_c(\text{root}) \wedge \forall x \left[\bigwedge_a Z_a x \leftrightarrow \bigvee_{b, d: b(d)=a} (P_b x \wedge \vartheta_d^{(x)}) \right] \right],$$

where $\vartheta^{(x)}$ denotes the relativisation of ϑ to the set of successors of x .

(b) We can use the same construction as in (a), we only have to modify the formulae ϑ_a . Since horizontal composition is commutative, there exist constants k and p such that every sum

$$b_0 + \dots + b_{m-1}$$

is uniquely determined by the values

$$\min \{n_a, k\} \quad \text{and} \quad n_a \bmod p, \quad \text{for every } a \in \bigcup_{\zeta \subseteq \xi} A_\zeta,$$

where n_a is the number of terms equal to a . Both of these values can be computed in CMSO .

(c) To compute a sum in a horizontal semigroup that is commutative and aperiodic, it is sufficient to know only the values $\min \{n_a, k\}$, which can be computed in MSO \square

Exercise 2.1. *Anti-chain logic* is the variant of MSO where quantification is restricted to anti-chains. *Anti-chain logic with counting* is the corresponding restriction of CMSO .

(a) Let Σ be an alphabet where every symbol has arity 0 or 2. Prove that every regular language $K \subseteq \mathbb{T}^{\text{fin}} \Sigma$ can be defined in anti-chain logic.

(b) Let $K \subseteq \mathbb{F}^{\text{fin}} \Sigma$ be a regular language of forests where no vertex has exactly one successor. Prove that K can be defined in anti-chain logic with counting. \square

3. Countable Chains

Another simple example are countable linear orders, which can be seen as forests of height 0, i.e., forests where all labels have arity 0.

Definition 3.1. (a) We denote by $\mathbb{C} : \mathcal{D} \rightarrow \mathcal{D}$ the functor mapping a set A to the set $\mathbb{C}A$ of all countable A -labelled linear orders. Formally, \mathbb{C} is the polynomial functor

$$\mathbb{C}X := \sum_I X^{\text{dom}(I)},$$

where I ranges over all countable linear orders.

(b) The flattening operation $\text{flat} : \mathbb{C}\mathbb{C} \Rightarrow \mathbb{C}$ is just the *ordered sum*, which is defined as follows. The flattening of $S \in \mathbb{C}\mathbb{C}A$ is the linear order with

domain

$$\begin{aligned} \text{dom}(\text{flat}(S)) &:= \sum_{i \in \text{dom}(S)} \text{dom}(S(i)) \\ &= \{ \langle i, k \rangle \mid i \in \text{dom}(S), k \in \text{dom}(S(i)) \} \end{aligned}$$

and ordering

$$\langle i, k \rangle \leq \langle j, l \rangle \quad : \text{iff} \quad i < j \quad \text{or} \quad i = j \text{ and } k \leq l.$$

(c) The singleton operation $\text{sing} : \mathbb{C} \Rightarrow \mathbb{C}$ maps an element $a \in A$ to the linear order $\text{sing}(a)$ with a single element that is labelled a .

Remark. As already remarked above, we can regard \mathbb{C} as a submonad of \mathbb{F} if we identify Set with $\text{Set}^{\{\emptyset\}}$. Then an alphabet $\Sigma \in \text{Set}$ corresponds to a $\{\emptyset\}$ -sorted set Σ' with $\Sigma'_\emptyset = \Sigma$. Furthermore, every vertex of a forest $s \in \mathbb{F}\Sigma'$ has sort \emptyset and must therefore be a leaf. Consequently, s is just the linear order formed by its countably many roots. Thus, we obtain an embedding $\mathbb{C} \Rightarrow \mathbb{F}|_{\{\emptyset\}}$. Furthermore, this embedding preserves the flattening and singleton operations, that is, it is a morphism of monads.

In light of this correspondence, we regard the product of a \mathbb{C} -algebra as the horizontal sum of an $\mathbb{F}|_{\{\emptyset\}}$ -algebra. Consequently, we will use additive notation for this product and call it a *sum*.

Example. The set $W \subseteq \mathbb{C}A$ of all A -labelled well-orders is recognisable by the finite \mathbb{C} -algebra with elements $\{0, 1\}$ and the infimum operation as multiplication. Then $W = \varphi^{-1}(1)$ where φ maps well-orders to 1 and all other orders to 0.

Exercise 3.1. Prove that $\langle \mathbb{C}, \text{flat}, \text{sing} \rangle$ forms a monad.

To show that monadic second-order logic over countable linear orders fits into our algebraic framework, we have to check that it is varietal, \mathbb{C} -compositional, and that syntactic algebras exist.

Definition 3.2. For the monad \mathbb{C} , we use the obvious variant of MSO that encodes a linear order $s \in \mathbb{C}A$ as a structure of the form $\langle \text{dom}(s), \leq, (P_a)_{a \in A} \rangle$

where the unary predicates $P_a := s^{-1}(a)$ contain the positions labelled by the elements of A .

Remark. Using the above identification of \mathbb{C} with a submonad of $\mathbb{F}|_{\{\emptyset\}}$, this version of MSO corresponds to the variant $\text{MSO}[\leq_{\text{so}}]$ defined above.

Compositionality is straightforward.

Theorem 3.3. MSO is strongly \mathbb{C} -compositional.

Proof. First note that, for $s \in \mathbb{C}\Sigma$, we can write

$$\text{flat}(s) \cong \tau \left(\sum_{i \in \text{dom}(s)} s(i) \right)$$

as a generalised sum follows by a quantifier-free interpretation τ . By the corresponding composition theorems (Proposition 1.11 and Theorem 1.12) it therefore follows that, given an MSO-formula φ of quantifier-rank m , we can compute an MSO-formula ψ such that

$$\sum_{i \in \text{dom}(s)} s(i) \models \varphi \quad \text{iff} \quad \mathbb{C}\theta_{\Delta_m}(s) \models \psi,$$

where Δ_m is the set of MSO-formulae of quantifier-rank at most m . \square

Corollary 3.4. MSO is varietal and \mathbb{C} -compositional.

Proof. By Proposition IV.4.11. \square

Next we turn to syntactic algebras. The proof utilises the submonad \mathbb{C}^{reg} of all linear orders that are ‘regular’ in the following sense.

Definition 3.5. Let A be a set.

(a) The *shuffle* of A is the linear order $s \in \mathbb{C}A$ with domain \mathbb{Q} such that

$$\text{for all } x < y \text{ and all } a \in A \text{ there is some } x < z < y \text{ with } s(z) = a.$$

For finite sets $A = \{a_0, \dots, a_{n-1}\}$, we denote the shuffle by $a_0 \sqcup \dots \sqcup a_{n-1}$.

(b) A *Läuchli-Leonard operation* over the set A is an operation on $\mathbb{C}A$ of one of the following forms:

- ♦ constants a , for $a \in A$.
- ♦ a binary operation $s + t$, for $s, t \in \mathbb{C}A$,
- ♦ two unary operations $s \times \omega$ and $s \times \omega^{\text{op}}$, for $s \in \mathbb{C}A$,
- ♦ an n -ary operation $s_0 \sqcup \dots \sqcup s_{n-1}$, for $s_0, \dots, s_{n-1} \in \mathbb{C}A$ and $n < \omega$.

$a \in A$ denotes the singleton order $\text{sing}(a)$; $s + t$ is the *ordered sum* of s and t ; $s \times \omega$ the ordered sum of ω copies of s arranged in an infinite increasing chain; and $s \times \omega^{\text{op}}$ the ordered sum of ω copies of s arranged in an infinite decreasing chain.

(c) A countable linear order s is *regular* if it is the value of a finite term of L\"auchli-Leonard operations. We denote the set of all such orders by $\mathbb{C}^{\text{reg}}A$.

Remark. The shuffle of a countable set A exists and it is unique up to isomorphism. Uniqueness follows by a straightforward back-and-forth argument. For existence, fix an enumeration $(a_n)_{n < \omega}$ (possibly with duplicates) of A and set

$$Q := (o + 1)^* 1.$$

Note that Q equipped with the lexicographic ordering is a dense linear order. We define a labelling $s : Q \rightarrow A$ by

$$s(o^n 1) := a_n \quad \text{and} \quad s(w 1 o^n 1) := a_n, \quad \text{for } n < \omega \text{ and } w \in \{o, 1\}^*.$$

Then s is the shuffle of A .

Lemma 3.6. *Let \mathfrak{S} be a finite \mathbb{C} -algebra, $C \subseteq S$, and let $K \subseteq \mathbb{C}C$ the least set satisfying the following conditions.*

- ♦ $C \subseteq K$.
- ♦ $u, v \in K$ implies $u + v \in K$.
- ♦ If $w_0, w_1, \dots \in K$ with $\pi(w_0) = \pi(w_1) = \dots$, then $w_0 + w_1 + \dots \in K$ and $\dots + w_1 + w_0 \in K$.
- ♦ If $s \in \mathbb{C}K$ such that $\mathbb{C}\pi(s)$ is the shuffle of a finite set, then $\text{flat}(s) \in K$.

Then $\mathbb{C}C \subseteq K$.

Proof. Let us call a term $s \in \mathbb{C}C$ *decomposable* if every non-empty factor of s belongs to K . We will show that every $s \in \mathbb{C}C$ is decomposable. Consider the following relation on $\text{dom}(s)$.

$$x \sim y \quad \text{iff} \quad x = y \text{ or the factor corresponding to the interval } (x, y] \text{ or } (y, x] \text{ is decomposable.}$$

Clearly, \sim is reflexive and symmetric. For transitivity, suppose that $x \sim y \sim z$ where w.l.o.g. $x < y < z$. Every factor w in the interval $(x, z]$ can be factorised as $w = uv$ where u lies in the interval $(x, y]$ and v in $(y, z]$. By assumption, u and v belong to K (if they are non-empty). This implies that $w = uv \in K$, as desired.

Thus \sim is an equivalence relation. Let H be a \sim -class and let w be the corresponding factor of s . We claim that w is decomposable. If $H = \emptyset$, this is trivial. Otherwise, fix some index $x \in H$ and set

$$H_0 := \{y \in H \mid y \leq x\} \quad \text{and} \quad H_1 := \{y \in H \mid y > x\}.$$

Let w_0 and w_1 be the factors corresponding to these two sets. As we have already proved above that decomposable factors are closed under binary concatenation, it is sufficient to show that w_0 and w_1 are decomposable.

If H_1 has a greatest element z , then $x \sim z$ implies that w_1 is decomposable. Otherwise, since H is countable, we can fix an increasing sequence $x = z_0 < z_1 < \dots$ of positions in H that is not bounded in H . Then $z_i \sim z_j$ implies that every factor u_{ij} corresponding to an interval $(z_i, z_j]$ with $i < j$ is decomposable. By the Theorem of Ramsey, there exist an infinite set $I \subseteq \omega$ such that $\pi(u_{ij}) = \pi(u_{i'j'})$ for all indices $i < j$ and $i' < j'$ in I . It follows that every non-empty factor v of w_1 can be factorised as $v'v''$ where v' corresponds to $(y, z_i]$ and v'' to $(z_i, y']$, for some index $i \in I$ and for some positions y and y' . By the closure properties of K , this implies that $v', v'' \in K$ (if they are non-empty). Hence, $v \in K$ as well.

The proof for w_0 is analogous.

To conclude the proof, let \hat{s} be the order obtained from s by replacing every \sim -class I with corresponding factor w by the single element $\pi(w)$. If \hat{s} is a singleton, it follows that s is decomposable and we are done.

Hence, suppose otherwise. If \hat{s} would contain two positions $x < y$ with no position in between, the concatenation of the corresponding \sim -classes would correspond to a decomposable factor (since decomposable factors are closed under binary concatenation). Consequently, \hat{s} is an infinite dense linear order. To obtain the desired contradiction it is sufficient to show that some interval of \hat{s} is a shuffle (since then the union of the corresponding \sim -classes would be decomposable again).

We prove the claim by the number of elements $a \in S$ appearing as a label in \hat{s} . If there is such an element a that does not appear in some infinite factor w of \hat{s} , the claim follows by applying the inductive hypothesis to w . Hence, we may assume that every label that appears in \hat{s} appears in every infinite factor of \hat{s} . Consequently, \hat{s} is the shuffle of these labels. \square

Theorem 3.7. *The inclusion $\mathbb{C}^{\text{reg}} \Rightarrow \mathbb{C}$ is dense over the class of all finite \mathbb{C} -algebras.*

Proof. Let $s \in \mathbb{C}\mathbb{C}$. Then $s \in K$, where K is the set from Lemma 3.6. By the inductive definition of K , we can construct a Lauchli-Leonard term s° with $\pi(s^\circ) = \pi(s)$. \square

Corollary 3.8. *Every language $K \subseteq \mathbb{C}\Sigma$ over a finite alphabet has a syntactic algebra.*

The goal of the remainder of this section is to prove that the MSO-definable languages are exactly those whose syntactic algebra is finite. The hard part of the proof consists in showing that the product of every finite \mathbb{C} -algebra is MSO-definable. To do so, we use a variant of regular expressions for languages over \mathbb{C} . These expressions are obtained by generalising the Lauchli-Leonard operations from trees to sets of trees.

Definition 3.9. (a) For an n -ary Lauchli-Leonard operation σ and languages $K_0, \dots, K_{n-1} \subseteq \mathbb{C}\Sigma$ we set

$$\sigma(\tilde{K}) := \{ \sigma(\tilde{s}) \mid s_i \in K_i \}.$$

(b) The iteration of a language $K \subseteq \mathbb{C}\Sigma$ is the set

$$K^+ := \{ s_0 + \dots + s_{n-1} \mid 0 < n < \omega, s_0, \dots, s_{n-1} \in K \}.$$

(c) A *regular expression* is a finite term using

- ♦ constants for all singleton languages $\{\text{sing}(a)\}$ with $a \in \Sigma$,
- ♦ boolean operations $K \cup L, K \cap L, \mathbb{C}\Sigma \setminus K$,
- ♦ the Lauchli-Leonard operations $K+L, K \times \omega, K \times \omega^{\text{op}}, K_0 \sqcup \dots \sqcup K_{n-1}$,
- ♦ iteration K^+ .

We start by noting that the above operations are all MSO-definable.

Lemma 3.10. *The class of MSO-definable languages is closed under all operations from the preceding definition.*

Proof. Closure under boolean operations is trivial and every language of the form $\{\text{sing}(a)\}$ is clearly MSO-definable.

For concatenation, note that we can check whether a given chain s belongs to $K+L$ by guessing a partition of $P \cup Q = \text{dom}(s)$ such that P is downwards closed, Q is upwards closed, and the restriction of s to P belongs to K , while the restriction to Q belongs to L . Each of these properties can be expressed in MSO, provided that K and L are MSO-definable.

For the remaining operations, we need to encode an arbitrary factorisation of a chain s by a single set $P \subseteq \text{dom}(s)$. We use the fact that every subset $P \subseteq \text{dom}(s)$ induces an equivalence relation on $\text{dom}(s)$ by

$$x \sim_P y \quad : \text{iff} \quad x \in P \Leftrightarrow z \in P, \quad \text{for all } x \leq z \leq y.$$

Hence, we can check whether $s \in K^+$ by guessing a subset $P \subseteq \text{dom}(s)$ such that there are only finitely many \sim_P -classes and every \sim_P -class belongs to K . Similarly, $s \in K \times \omega$, if the \sim_P -classes form an infinite increasing chain of length ω and each class belongs to K . Similarly, for $K \times \omega^{\text{op}}$. Finally, we have $s \in K_0 \sqcup \dots \sqcup K_{n-1}$ if there is a set P such that, for every pair of \sim_P -classes $I < J$ and every $i < n$, there is some \sim_P -class H with $I < H < J$ that belongs to K_i . \square

The main part of the proof relies on algebraic tools from semigroup theory. The corresponding adaptation to \mathbb{C} -algebras looks as follows.

Definition 3.11. Let \mathfrak{S} be a \mathbb{C} -algebra. For $a, b \in S$, we define the following Green's relations.

$$\begin{aligned} a \leq_L b & \quad \text{iff} \quad a = s + b, & \text{for some } s \in S, \\ a \leq_R b & \quad \text{iff} \quad a = b + s, & \text{for some } s \in S, \\ a \leq_J b & \quad \text{iff} \quad a = s + b + t, & \text{for some } s, t \in S. \end{aligned}$$

Let \equiv_L, \equiv_R , and \equiv_J be the corresponding equivalence relations and set

$$\equiv_H := \equiv_L \cap \equiv_R.$$

We call the corresponding equivalence classes the *L-class*, *R-class*, *J-class*, and *H-class* of the given element.

(b) A chain $s \in \mathbb{C}S$ is *J-homogeneous* if there exists a J-class J such that $s \in \mathbb{C}J$ and $\pi(s) \in J$.

Below we need the following standard facts about these relations.

Lemma 3.12. Let \mathfrak{S} be a finite \mathbb{C} -algebra.

- (a) $(\equiv_L \circ \equiv_R) = \equiv_J = (\equiv_R \circ \equiv_L)$
- (b) $a \equiv_J ab \equiv_J b$ implies $a \equiv_R ab$ and $ab \equiv_L b$.
- (c) Every H-class H containing an idempotent element forms a group (with respect to the semigroup-multiplication induced by the product of \mathfrak{S}).

We can use Green's relations to compute products in a \mathbb{C} -algebra as follows.

Lemma 3.13. Let \mathfrak{S} be a finite \mathbb{C} -algebra, $J \subseteq S$ a J-class containing an idempotent $e \in J$, and let $s \in \mathbb{C}J$ be a J-homogeneous chain of order type ω . Then

$$\pi(s) = a + e \times \omega,$$

where the element $a \in S$ can be computed given e and the R-class of the first letter of s .

Proof. Let $b := s(o)$ be the first letter of s and let s' be the suffix such that $s = b + s'$. Since $b \equiv_J e$, we have $b = a + e + c$ for some $a, c \in J$. We claim that we can compute a from e and the R-class of b . Hence, suppose that $b \equiv_R b'$, for some other element $b' \equiv_J e$. Then $b' = b + x$, for some $x \in S$, and we have

$$b' = b + x = a + e + c + x,$$

as desired. To conclude the proof, it is therefore sufficient to show that $e + c + s' = e^\omega$.

Set $t := c + s'$. By the Theorem of Ramsey, there exists a factorisation $t = u_0 + u_1 + \dots$ and an element $f \in S$ such that

$$\pi(u_i + \dots + u_j) = f, \quad \text{for all } 0 < i \leq j < \omega.$$

In particular, we have

$$f + f = \pi(u_1) + \pi(u_2) = \pi(u_1 + u_2) = f.$$

Hence, f is idempotent. Since f is the product of a factor of a J-homogeneous chain, we furthermore have $f \in J$. Setting $b := \pi(u_0)$, it follows that

$$\pi(t) = b + f + f + \dots = b + f \times \omega.$$

As $f \equiv_J e$, we have $f = x + e + y = x + e + e + y$, for some $x, y \in S$. Consequently,

$$\begin{aligned} \pi(e + t) &= e + b + f \times \omega \\ &= e + b + (x + e + e + y) \times \omega \\ &= e + b + x + e + (e + y + x + e) \times \omega. \end{aligned}$$

Setting $g := e + b + x + e$ and $h := e + y + x + e$, we have $g \equiv_H e \equiv_H h$. In particular, the H-class H of g and h contains an idempotent, which implies by Lemma 3.12 (c) that H forms a group with neutral element e . Choose $m, n < \omega$ such that $g \times m = e$ and $h \times n = e$. Then

$$g \times \omega = (g \times m) \times \omega = e \times \omega \quad \text{and} \quad h \times \omega = (h \times n) \times \omega = e \times \omega.$$

Consequently,

$$\pi(e + t) = g + h \times \omega = g + g \times \omega = g \times \omega = e \times \omega. \quad \square$$

Lemma 3.14. *Let \mathfrak{S} be a finite \mathbb{C} -algebra and let $F \subseteq S$ be upwards closed with respect to the order \leq_J .*

- (a) *The function mapping a chain $s \in \pi^{-1}[F]$ to the R-class of $\pi(s)$ is MSO-definable.*
- (b) *The function mapping a chain $s \in \pi^{-1}[F]$ to $\pi(s)$ is MSO-definable.*
- (c) *The language $\pi^{-1}[F]$ is MSO-definable.*

Proof. We prove all three claims by a simultaneous induction on the number of J-classes contained in F . Fix a minimal J-class $J \subseteq F$, set

$$F_0 := F \setminus J, \quad K := \pi^{-1}[F_0], \quad \text{and} \quad L := \pi^{-1}[J].$$

Furthermore, we define

$$K_+ := K \cup K \times \omega \cup K \times \omega^{\text{op}} \cup ((K \cup J) + (K \cup J)).$$

Note that K , K_+ , and $K \cup L$ are all closed under factors. Furthermore, given an increasing sequence $(u_i)_i$ of prefixes of some chain s with $u_i \in K$, the union of all u_i belongs to $K \times \omega \subseteq K_+$.

We call a prefix u of $s \in \pi^{-1}[F]$ an ω -prefix if u is either empty or an infinite sum of the form $u_0 + u_1 + \dots$ with $u_i \in L$. Similarly, an ω^{op} -suffix of s is a suffix of the form $\dots + v_1 + v_0$ with $v_i \in L$.

Our proof consists of a simultaneous induction establishing not only (a), (b), and (c), but also the following three claims.

- (I) The function mapping $s \in K_+$ to $\pi(s)$ is MSO-definable.
- (II) A union of ω -prefixes of some chain in $\pi^{-1}[F]$ is again an ω -prefix.
- (III) We can factorise every $s \in \pi^{-1}[F]$ as

$$s = u + s_0 + \dots + s_{n-1} + v,$$

where $s_0, \dots, s_{n-1} \in K_+$, u is an ω -prefix of s , and v is an ω^{op} -suffix.

(I) Given $s \in K_+$, we can use part (c) of the inductive hypothesis to define a factorisation of s belonging to one of the languages

$$K, \quad K \times \omega, \quad K \times \omega^{\text{op}}, \quad \text{or} \quad ((K \cup J) + (K \cup J)).$$

(Such a factorisation can be encoded as a single set variable consisting of the union of every other interval in the decomposition.) By part (b) of the inductive hypothesis, we can define the product of each factor. Consequently, we can compute $\pi(s)$ with the help of Lemma 3.10.

(II) Let w be a union of ω -prefixes of $s \in \pi^{-1}[F]$. For each position $x \in \text{dom}(w)$, we can find an ω -prefix $u_0 + u_1 + \dots$ containing x . If $x \in \text{dom}(u_i)$, then u_{i+1} is a chain in L that lies entirely to the right of x . Consequently, we can find a cofinal sequence v_0, v_1, \dots of disjoint factors of w that all belong to L and such that v_i lies to the left of v_{i+1} . Enlarging these factors, if necessary, we obtain a decomposition $w = v_0 + v_1 + \dots \in L \times \omega$.

(III) Let $s \in \pi^{-1}[F]$. By (II), s has a longest ω -prefix u . Similarly, it has a longest ω^{op} -suffix v that is disjoint from u . Thus, we can write $s = u + t_0 + v$ and it remains to show that we can decompose t_0 into a finite sum of chains in K_+ . To do so, we construct a sequence $s_0, s_1, \dots \in K_+$ and chains t_1, t_2, \dots inductively as follows.

Suppose that we have ready defined t_i . If t_i has a prefix $s_i \in K_+ \setminus K$, we fix one such prefix and we choose for t_{i+1} the suffix of t_i such that $t_i = s_i + t_{i+1}$. Otherwise, let s_i be the union of all prefixes of t_i that belong to K and let t_{i+1} be the remainder of t_i .

To prove that our construction terminates, we will show below that every chain s_i is non-empty. Then it follows that, after finitely many steps, we must arrive at a situation where $t_i = 0$ since, otherwise, t_0 had a prefix $(s_0 + s_1) + (s_2 + s_3) + \dots$ with $s_i + s_{i+1} \in L$. Hence, s would have an ω -prefix $u + s_0 + s_1 + \dots$ that is longer than u . A contradiction.

It remains to prove the claim. If $s_i \in K_+ \setminus K$, we are done as the empty chain belongs to K . Otherwise, t_{i+1} cannot have a least element since by adding such an element to s_i we would obtain a word in $K_+ \setminus K$. A contradiction to our choice of s_i . Hence, we can fix a factorisation $t_{i+1} = \dots + w_2 + w_1 + w_0$ of order type ω^{op} . Only finitely many of the factors w_j belong to L since,

otherwise, t_{i+1} had an ω^{op} -suffix v' . Hence, $v' + v$ would be an ω^{op} -suffix of s disjoint from u . A contradiction to our choice of v . Consequently, t_{i+1} has a prefix that belongs to $K \times \omega^{\text{op}} \subseteq K_+$. In particular, it follows that s_i is non-empty since, otherwise, we could have chosen $s_i = \cdots + w_2 + w_1$.

(a) Given $s \in \pi^{-1}[F]$, we can compute the R-class of $\pi(s)$ as follows. If $s \in K$, we can use the inductive hypothesis. Otherwise, let s_o be the union of all prefixes p of s that belong to K , and let s_1 be the suffix of s such that $s = s_o + s_1$. Then $s_o \in K_+$. Set $a_o := \pi(s_o)$ and $a_1 := \pi(s_1)$. If $a_o \in J$ then $a_o + a_1 \in J$ implies, by Lemma 3.12 (b), that the R-classes of a_o and $a_o + a_1 = \pi(s)$ coincide. By (i), a_o is definable. Hence, so is its R-class.

It remains to consider the case where $a_o \in F_o$. First, suppose that s_1 has a least element. Let c be its label. By maximality of s_o , we have $a_o + c \in J$. As above, this implies that the R-classes of $a_o + c$ and $\pi(s)$ coincide. Since $a_o + c$ is definable, the claim follows.

Hence, we may assume that s_1 does not have a least element. By the Theorem of Ramsey, we can factorise s_1 as $s_1 = \cdots + w_2 + w_1 + w_o$ where $\pi(w_i) = \pi(w_j)$, for $i, j \neq o$. Let c be the value of this product. It follows that s_1 has a prefix with product $c \times \omega^{\text{op}}$. As above, this implies that the R-class of $\pi(s)$ is equal to that of $a_o + c \times \omega^{\text{op}}$. It remains to define c . If $c \in F_o$, it is MSO-definable by inductive hypothesis. Otherwise, Lemma 3.13 implies that $c \times \omega^{\text{op}} = e \times \omega^{\text{op}}$, for any idempotent $e \in J$.

(b) Given $s \in \pi^{-1}[F]$, let $s = u + s_o + \cdots + s_{n-1} + v$ be the factorisation from (11). Note that, by (a), we know the R-class of (the first factor of) $\pi(u)$ and we can use the dual statement to define the L-class of $\pi(v)$. By Lemma 3.13 (a), this is sufficient to compute $\pi(u)$ and $\pi(v)$. Furthermore, we can compute the product $\pi(s_o) + \cdots + \pi(s_{n-1})$ in MSO since we can obtain $\pi(s_i)$ via (i) and MSO allows us to evaluate finite products in a finite semigroup.

(c) Note that

$$\pi^{-1}[F] = \mathbb{CS} \setminus \pi^{-1}[I] \quad \text{where} \quad I := \pi^{-1}[S \setminus F].$$

Hence, it is sufficient to show that $\pi^{-1}[I]$ is MSO-definable. To do so we introduce the following concept. Let $N = (N_a)_{a \in F}$ be a family of sets

$N_a \subseteq \mathbb{CS}$. We call a chain $w \in \mathbb{CS}$ *N-critical* if it satisfies one of the following conditions.

- ♦ $w = \text{sing}(a)$, for some $a \notin F$.
- ♦ $w = u + v$ with $u \in N_a$, $v \in N_b$ with $a + b \notin F$.
- ♦ $w = u_o + u_1 + \cdots$ with $u_o, u_1, \dots \in N_a$ and $a \times \omega \notin F$.
- ♦ $w = \cdots + u_1 + u_o$ with $u_o, u_1, \dots \in N_a$ and $a \times \omega^{\text{op}} \notin F$.
- ♦ $w \in N_{a_o} \sqcup \cdots \sqcup N_{a_{n-1}}$ with $a_o \sqcup \cdots \sqcup a_{n-1} \notin F$.

Below we will prove the following claim: if N is a family of sets with

$$\pi^{-1}(a) \subseteq N_a \subseteq \pi^{-1}[I \cup \{a\}], \quad \text{for all } a \in F,$$

then

$$s \in \pi^{-1}[I] \quad \text{iff} \quad s \text{ has an } N\text{-critical factor.}$$

Before doing so, let us show how this implies (c). For every $a \in F$, we obtain from (b) an MSO-formula φ_a such that

$$s \models \varphi_a \quad \text{iff} \quad \pi(s) = a, \quad \text{for } s \in \pi^{-1}[F].$$

Let N_a be the language defined by φ_a . Then we can use Lemma 3.10 to construct an MSO-formula expressing the existence of an *N-critical* factor. By the above claim, this is equivalent to $s \in \pi^{-1}[I]$ since we have

$$\pi^{-1}(a) \subseteq N_a \subseteq \pi^{-1}[\{a\} \cup I].$$

It therefore remains to prove the claim.

(\Leftarrow) Note that I is downwards closed with respect to \leq_J . Hence, if s has an *N-critical* factor w , then $w \in \pi^{-1}[I]$ implies $s \in \pi^{-1}[I]$.

(\Rightarrow) It is sufficient to show that the language

$$M := \{s \in \mathbb{CS} \mid \pi(s) \in F \text{ or } s \text{ has an } N\text{-critical factor}\}$$

coincides with \mathbb{CS} . We will do so using Lemma 3.6. To apply the lemma, we have to check several closure properties of M . Since the arguments involved are all identical, we will only prove closure under binary sums.

For a contradiction, suppose that $s, t \in M$ but $s + t \notin M$. Then $s + t$ does not have an N -critical factor. Hence, neither do s and t . It follows that $s, t \in \pi^{-1}[F]$. Hence, $s \in N_a$ and $t \in N_b$, for some $a, b \in F$. But $s + t \notin M$ implies that $a + b = \pi(s + t) \in S \setminus F = I$. Hence, it follows by the second condition above that $s + t$ has an N -critical factor. A contradiction. \square

Theorem 3.15. *A \mathbb{C} -algebra \mathfrak{S} is MSO-definable if, and only if, it is finite.*

Proof. (\Rightarrow) is trivial since, by definition, every MSO-definable algebra is finitary. As \mathbb{C} -algebras have only one sort, this is the same as being finite.

(\Leftarrow) Fix $a \in S$ and let $J := [a]_J$ be the J -class of a . By Lemma 3.14 (c), there exists an MSO-formula ψ_J defining the language $K_J := \pi^{-1}[J]$. Furthermore, by part (b) of that lemma, there exists an MSO-formula φ_a defining $\pi(s)$, for $s \in K_J$. It follows that the conjunction $\psi_J \wedge \varphi_a$ defines $\pi^{-1}(a)$. \square

Corollary 3.16. *A language $K \subseteq \mathbb{C}\Sigma$ is MSO-definable if, and only if, it has a finite syntactic algebra.*

4. Counterexamples

Before continuing, let us give several counterexamples showing that, for infinite trees, MSO-definable algebras are more complicated than one might hope. We start with an example showing that not every finitary \mathbb{T}^\times -algebra is MSO-definable.

Definition 4.1. A tree $t \in \mathbb{T}^\times \Sigma$ is *antiregular* if it is infinite and no two subtrees of t are isomorphic. We call t *densely antiregular* if every subtree of t has an antiregular subtree. \dashv

First, note that antiregular trees do exist.

Lemma 4.2. *Let Σ be an alphabet containing two elements a and b of arity 2. There exists an antiregular tree in $\mathbb{T}_\emptyset \Sigma$.*

Proof. Let $P \subseteq \{0, 1\}^*$ be the language (of finite words) of all palindromes. We claim that the tree $t \in \mathbb{T}\{a, b\}$ defined by

$$\text{dom}(t) = \{0, 1\}^* \quad \text{and} \quad t(v) := \begin{cases} a & \text{if } v \in P, \\ b & \text{otherwise,} \end{cases}$$

is antiregular.

Let s and s' be two subtrees of t rooted at u and u' , respectively. Let \bar{u} be the mirror image of u . Then $u\bar{u} \in L$ and $u'\bar{u} \notin L$, which implies that $s(\bar{u}) = a$ and $s'(\bar{u}) = b$. In particular, s and s' are not isomorphic. \square

Theorem 4.3. *Suppose that $\mathcal{D} = \text{Pos}$ and let $\Sigma := \{a, b\}$ be an alphabet with two elements, both of arity 2. The language L of all densely antiregular trees over Σ is not regular, but it can be recognised by a finitary \mathbb{T}^\times -algebra.*

Proof. To show that L is not regular, note that L is non-empty by Lemma 4.2. If L were regular, it would therefore also contain some regular tree. As regular trees cannot be densely antiregular, this is not possible.

Hence, it remains to construct a finitary \mathbb{T}^\times -algebra recognising L . For $\xi \in \Xi$, let Δ_ξ be the set of all finite trees in $\mathbb{T}_\xi \Sigma$. As Σ contains only binary elements, every leaf of a tree $t \in \Delta_\xi$ must be labelled by a variable. Hence, t has at most $|\xi|$ leaves and, therefore, at most $|\xi| - 1$ internal vertices. This implies that Δ_ξ is a finite set.

The domains of the desired algebra \mathfrak{A} are

$$A_\xi := \Delta_\xi \cup \{\perp, *\}, \quad \text{for } \xi \in \Xi,$$

which we order such that \perp is the least element and all other elements are incomparable. To give a bit of intuition for the definition of the product, let us first describe the morphism $\eta : \mathbb{T}^\times \Sigma \rightarrow A$ that will recognise L . For

$t \in \mathbb{T}_\xi^\times \Sigma$, we set

$$\eta(t) := \begin{cases} t & \text{if } t \in \Delta_\xi, \\ \perp & \text{if } t \text{ has a subtree without variables that is not densely} \\ & \text{antiregular,} \\ ? & \text{every subtree of } t \text{ has a variable and some variable} \\ & \text{occurs at least twice,} \\ * & t \text{ has a subtree without variables and every such subtree} \\ & \text{is densely antiregular.} \end{cases}$$

Then $L = \eta^{-1}[*]$. Hence, it remains to equip A with a product that makes η into a morphism.

We define the product as follows. Let us call a tree $s \in \mathbb{T}^\times A$ *good* if every subtree of s contains some vertex v such that $t(v) = *$, or such that $t|_v \in \mathbb{T}^\times \Delta$ and $\text{flat}(t|_v)$ is densely antiregular. For $t \in \mathbb{T}_\xi^\times A$, we then set

$$\pi(t) := \begin{cases} \perp & \text{if } t \text{ contains the label } \perp, \\ \perp & \text{if } t \text{ has a subtree without variables that is not} \\ & \text{good,} \\ * & \text{if } t \text{ has a subtree without variables and every such} \\ & \text{subtree is good,} \\ * & \text{if } t \text{ contains the label } * \text{ and every subtree of } t \text{ has} \\ & \text{a variable,} \\ ? & \text{if } t \in \mathbb{T}^\times (\Delta + \{?\}), t \text{ contains the label } ?, \text{ and} \\ & \text{every subtree of } t \text{ has a variable,} \\ ? & \text{if } t \in \mathbb{T}^\times \Delta, \text{ every subtree of } t \text{ has a variable, and} \\ & \text{some variable appears more than once in } t, \\ \text{flat}(t) & \text{if } t \in \mathbb{T}^\times \Delta \text{ is finite and every variable appears} \\ & \text{exactly once.} \end{cases}$$

It is straightforward (but tedious) to check that these cases cover all possibilities, that the function π is monotone, and that it satisfies

$$\eta \circ \text{flat} = \pi \circ \mathbb{T}^\times \eta.$$

Consequently, it follows by Lemma I.5.5 that $\mathfrak{A} := \langle A, \pi \rangle$ is a \mathbb{T}^\times -algebra and $\eta : \mathbb{T}^\times \Sigma \rightarrow \mathfrak{A}$ an \mathbb{T}^\times -morphism. \square

Definition 4.4. We call the algebra \mathfrak{A} from the above proof the *Bojańczyk-Klin algebra*. \lrcorner

The next two examples concern closure properties of MSO-definable algebras. Our definition of \mathcal{R} -varieties was complicated by the fact that the usual algebraic operations might produce algebras that are not finitely generated. Here we present two examples showing that a subalgebra or a finite product of an MSO-definable \mathbb{T}^\times -algebra need not be finitely generated.

Let us start with subalgebras. We use a result about so-called *clones*. A clone \mathfrak{C} is a set of functions (of various arities) over some fixed set X that contains all projections and that is closed under composition, i.e., if \mathfrak{C} contains $f : X^n \rightarrow X$ and $g_0, \dots, g_{n-1} : X^m \rightarrow X$, it also contains the m -ary function

$$\bar{x} \mapsto f(g_0(\bar{x}), \dots, g_{n-1}(\bar{x})).$$

Note that this composition also makes sense if the functions g_0, \dots, g_{n-1} have different arities since we can make their arities equal by composing them by suitable projections (which are in \mathfrak{C} by assumption).

Theorem 4.5 (Yanov, Muchnik). *There are uncountably many clones on a three element set.*

As there are only countably many finitely generated clones, it follows in particular that there exists some clone \mathfrak{C} that is not finitely generated. We will use it to construct the desired \mathbb{T}^\times -algebra.

Example. Let $[3] = \{0, 1, 2\}$ be a three element set and let A_ξ be the set of all functions $[3]^\xi \rightarrow [3]$ together with a special error value \perp . We turn

$A := (A_\xi)_\xi$ into a \mathbb{T}^\times -algebra by defining the following multiplication $\pi : \mathbb{T}^\times A \rightarrow A$. For a finite tree $t \in \mathbb{T}^\times A$ that does not contain the symbol \perp , we compute the product $\pi(t)$ by composing all the functions that label the vertices of t . For all other trees, we set $\pi(t) := \perp$. The resulting structure $\mathfrak{A} = \langle A, \pi \rangle$ forms a \mathbb{T}^\times -algebra which is finitely generated. (To see the latter, one can, e.g., represent every 3-valued function in a similar way as boolean functions can be written in disjunctive normal form.) Furthermore, \mathfrak{A} is MSO-definable since, when evaluating a tree t an automaton is able to first check that t is finite and does not contain \perp , and then evaluate t bottom up by remembering where each (of the bounded number) of the input arguments is mapped to.

To conclude the construction recall that we have seen above that there exists a clone on $[3]$ that is not finitely generated. Let $\mathfrak{C} \subseteq \mathfrak{A}$ be the subalgebra of \mathfrak{A} consisting of the elements of that clone. Then \mathfrak{C} is not finitely generated.

Example. Our counterexample for products looks as follows. For simplicity, we will use \mathbb{T} -algebras instead of \mathbb{T}^\times -algebras. We start with a \mathbb{T} -algebra \mathfrak{B} where the elements of sort ξ are all finite sequences in ξ^* that contain every variable at most once. We define the product as follows. Suppose we have sequences $\alpha \in B_\xi$ and $\beta \in (B_\zeta)^\xi$ where the β_x are disjoint. If $\alpha = \langle x_{i_0}, \dots, x_{i_{k-1}} \rangle$, we set

$$\alpha(\beta) := \beta_{x_{i_0}} \dots \beta_{x_{i_{k-1}}},$$

i.e., we substitute β_x for x in α . For a finite tree $t \in \mathbb{T}B_\xi$, we can now inductively define

$$\pi(t) = \alpha(\pi(s_0), \dots, \pi(s_{m-1})),$$

where $\alpha := t(\langle \rangle)$ is the label at the root and s_0, \dots, s_{m-1} are the attached subtrees. (With the convention that $\pi(s_i) = \langle x \rangle$ in case that $s_i = x$ is a single variable.)

We can extend this definition to infinite trees as follows. If t does not contain variables, we set $\pi(t) = \langle \rangle$. Otherwise, we choose a finite prefix s of t

that contains all the variables (here we need the fact that $t \in \mathbb{T}B$), separately compute the products of s and of the attached subtrees, and then multiply the results as above. Note that this definition ensures that $\pi(t)$ is the sequence of all variables appearing in t , but not necessarily in the order they appear in.

Again it is straightforward to check that \mathfrak{B} is a \mathbb{T} -algebra. Furthermore, note that by suitably choosing the ordering of the variables of t we can write every sequence $\alpha \in B_\xi$ as the product of a tree t where all internal vertices are labelled by $\langle x \rangle$ or $\langle x, y \rangle$, for fixed $x, y \in X$. Hence, \mathfrak{B} is finitely generated by three elements $\langle \rangle, \langle x \rangle, \langle x, y \rangle$.

Furthermore, \mathfrak{B} is MSO-definable since, given an element $\alpha \in B_\xi$ and a finite set of generators, an automaton can determine whether an input tree evaluates to α since all intermediate results are sequences of length at most $|\xi|$.

We claim that the product $\mathfrak{B} \times \mathfrak{B}$ is not finitely generated. For a contradiction suppose otherwise and fix a finite set C of generators. Choose a number m that is greater than the arity of all elements in C and set $\xi := \{x_0, \dots, x_{2m-1}\}$. We consider the element $\langle \alpha, \beta \rangle \in B_\xi \times B_\xi$ where

$$\begin{aligned} \alpha &:= \langle x_0, \dots, x_{2m-1} \rangle \\ \beta &:= \langle x_m, x_0, x_{m+1}, x_1, \dots, x_{m+i}, x_i, \dots, x_{2m-1}, x_{m-1} \rangle. \end{aligned}$$

By assumption, there is a tree $t \in \mathbb{T}C$ with product $\langle \alpha, \beta \rangle$. Let $\langle \gamma, \delta \rangle$ be the label at the root of t and let s_0, \dots, s_{n-1} be the subtrees attached to it. (For simplicity, we assume that $n > 1$. Otherwise our proof needs to be slightly modified.) By choice of m , there is some subtree s_i that contains at least two variables. Let $\sigma, \tau : [n] \rightarrow [n]$ be the permutations such that

$$\gamma = \langle x_{\sigma(0)}, \dots, x_{\sigma(n-1)} \rangle \quad \text{and} \quad \delta = \langle x_{\tau(0)}, \dots, x_{\tau(n-1)} \rangle,$$

and let $p : \mathfrak{B} \times \mathfrak{B} \rightarrow \mathfrak{B}$ be the projection to the first component. By looking

at the first components, we see that

$$\begin{aligned} & \pi(\mathbb{T}p(s_{\sigma(o)})) \dots \pi(\mathbb{T}p(s_{\sigma(n-1)})) \\ &= \gamma(\pi(\mathbb{T}p(s_o)), \dots, \pi(\mathbb{T}p(s_{n-1}))) \\ &= \alpha \\ &= \langle x_o, \dots, x_{2m-1} \rangle. \end{aligned}$$

Consequently, there exist numbers $k < l$ such that the term s_i contains the variables $x_k, x_{k+1}, \dots, x_{l-1}$. By choice of i , we have $l \geq k + 2$.

Looking at the second components, we see that β must have some segment of length $l - k \geq 2$ which contains the variables $x_k, x_{k+1}, \dots, x_{l-1}$ (in any order). But the only segments of β of this form are those of length 1 and the one of length $2m$. A contradiction. \perp

5. MSO-definable Algebras

The traditional tool to study the expressive power of monadic second-order logic over forests are automata. As our algebraic theory is not quite developed enough to replace all automata-based techniques, let us briefly recall some material on automata that we will need below.

Definition 5.1. (a) A *forest automaton* is a tuple $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_o, \Omega \rangle$ where

- $\langle Q, \sigma \rangle$ forms a finite \mathbb{C} -algebra whose elements are called *states*,
- Σ is the *input alphabet*,
- $\xi \in \Xi$ the *input sort*,
- $q_o \in Q$ the *initial state*,
- $\Omega : Q \rightarrow \omega$ the *priority function*, and
- $\Delta = (\Delta_\zeta)_\zeta$ the *transition relation* where

$$\Delta_\zeta \subseteq Q \times (\Sigma + \xi)_\zeta \times Q^\zeta.$$

(b) A *run* of an automaton \mathcal{A} on an input forest $s \in \mathbb{F}_\xi^\times \Sigma$ is a function $\rho : \text{dom}_+(s) \rightarrow Q$ such that, for every vertex $v \in \text{dom}_+(s)$,

$$\langle \rho(v), s(v), \bar{q} \rangle \in \Delta,$$

where, for each $x \in \xi$, $q_x := \sigma(w_x)$ and $w_x \in \mathbb{C}Q$ is the linear order obtained by restricting ρ to the set of x -successors of v .

(c) A run ρ is *accepting* if,

- $\sigma(w) = q_o$, where $w \in \mathbb{C}Q$ is the restriction of ρ to the roots of s , and
- for every infinite branch β of s , the following *parity condition* is satisfied:

$$\liminf_{n \rightarrow \infty} \Omega(\rho(\beta_n)) \text{ is even, where } \beta_n \text{ denotes the } n\text{-th vertex on } \beta.$$

(d) The language *recognised* by \mathcal{A} is the set

$$L(\mathcal{A}) := \{ s \in \mathbb{F}_\xi^\times \Sigma \mid \text{there is an accepting run of } \mathcal{A} \text{ on } s \}.$$

Example. Let $\Sigma = \{a, b\}$ where a and b both have arity 1, and let $K \subseteq \mathbb{F}_\emptyset^\times \Sigma$ the set of all forests with at least one occurrence of the letter a . The following automaton $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_o, \Omega \rangle$ recognises K . We use two states $Q := \{o, 1\}$ with initial state $q_o := 1$. The input sort is $\xi := \emptyset$ and the priority function is the identity $\Omega(q) := q$. Finally, the transition relation is defined by

$$\begin{aligned} \Delta &:= \{ \langle 1, a, q \rangle \mid q \in Q \} \\ &\cup \{ \langle q, c, q \rangle \mid q \in Q, c \in \{a, b\} \}. \end{aligned}$$

To see that this automaton recognises K , consider a forest $s \in K$. We obtain an accepting run ρ of \mathcal{A} on s by choosing some vertex v labelled a and setting all states on the path from the corresponding root to v to 1, while all other vertices get the state o . Conversely, given an accepting run ρ on a forest s , note that the set of vertices $P := \rho^{-1}(1)$ is downwards closed (with respect to the forest ordering \leq_{pf}). Furthermore, since ρ is accepting, P is non-empty and it does not contain an infinite branch. Consequently, P must have a maximal element v , which implies that $s(v) = a$. \perp

The behaviour of an automaton $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_o, \Omega \rangle$ on an input forest $s \in \mathbb{F}_\xi^\times \Sigma$ can also be described by a game, called the *Automaton-Pathfinder Game*. This game has two players, *Automaton* and *Pathfinder*, where Automaton tries to prove that the given forest is accepted by \mathcal{A} , while Pathfinder tries to refute this claim.

The game positions for Automaton are of the form $\langle p, q \rangle$ or $\langle v, q \rangle$ where $q \in Q$ is a state, $v \in \text{dom}(s)$ a vertex, and $p \subseteq \text{dom}(s)$ is what we call a *place*. By definition, a place is the set of roots of some subforest of s , that is, p is either the set of roots of s , or the set of all x -successors of some vertex $u \in \text{dom}(s)$. The former is called the *root place* of s , while the latter is the *x -successor place* of u . The positions for Pathfinder are either of the form $\langle p, f \rangle$ where p is a place and $f : p \rightarrow Q$ a function, or of the form $\langle u, \delta \rangle$ where u is a vertex and $\delta \in \Delta$ a transition. The starting position of the game is $\langle r, q_o \rangle$ where r is the root place of s and q_o the initial state.

Each round of the game proceeds in two phases. In the position $\langle p, q \rangle$, Automaton first chooses a function $f : p \rightarrow Q$ with $\sigma(f) = q$ and Pathfinder selects a vertex $u \in p$. Then Automaton picks a transition $\langle f(u), s(u), \tilde{q}' \rangle \in \Delta$ and Pathfinder answers by choosing a variable x . The new position is $\langle p_x, q'_x \rangle$ where p_x is the set of x -successors of u and q'_x is the x -component of \tilde{q}' .

The game continues until either a position is reached where the corresponding player cannot make a move, in which case this player loses, or an infinite number of rounds is played and Automaton wins if the sequence of states occurring in the corresponding game positions satisfies the parity condition. Using this game, we obtain the following description of the behaviour of an automaton.

Lemma 5.2. *Let \mathcal{A} be an automaton and s an input forest. Then \mathcal{A} accepts s if, and only if, Automaton has a winning strategy for the corresponding Automaton-Pathfinder Game.*

Below we will need the following two facts from automata theory. (There are currently no purely algebraic proofs of these results.)

Theorem 5.3. *A language $K \subseteq \mathbb{F}_\xi^\times \Sigma$ is MSO-definable if, and only if, it can be recognised by some automaton \mathcal{A} .*

Definition 5.4. A forest $s \in \mathbb{F}_\xi^\times \Sigma$ is *regular* if, up to isomorphism, s has only finitely many different subtrees. \lrcorner

Theorem 5.5. *Two MSO-definable languages $K, L \subseteq \mathbb{F}^\times \Sigma$ are equal if, and only if, they contain the same regular forests.*

We will also have to deal with runs of automata on forests $\text{flat}(s)$ that are partitioned into several factors. We can decompose such a run into pieces, one for each factor $s(v)$. These pieces are not themselves runs since they do not necessarily start at the starting state and at variables the transition relation does not need to be satisfied.

Definition 5.6. Let $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_o, \Omega \rangle$ be a forest automaton.

(a) A *partial run* of \mathcal{A} on some forest $s \in \mathbb{F}_\xi^\times \Sigma$ (where ξ might be different from ξ) is a function $\rho : \text{dom}_+(s) \rightarrow Q$ such that

- ♦ ρ satisfies the transition relation

$$\langle \rho(v), s(v), \tilde{q} \rangle \in \Delta$$

at every vertex v that is not labelled by a variable, and

- ♦ ρ satisfies the parity condition for every infinite branch of s .

(b) The *starting state* of a partial run ρ is the state $\sigma(w)$, where $w \in \mathbb{C}Q$ is the restriction of ρ to the roots of the input forest s .

(c) The *profile* of a partial run ρ on a forest $s \in \mathbb{F}_\xi^\times \Sigma$ is the tuple $\tau = \langle p, \tilde{U} \rangle$ where p is the starting state of ρ and, for each $z \in \xi$, U_z is the set of all pairs $\langle k, q \rangle$ such that there exists some leaf v of s labelled z with state $q := \rho(v)$ and such that the least priority seen along the path from the corresponding root to v is equal to k . \lrcorner

It follows that, given a forest $s \in \mathbb{F}^\times \mathbb{F}^\times \Sigma$ and a partial run $\rho(v)$ on each factor $s(v)$ such that the states at the holes of $s(v)$ are equal to the σ -product of the results of the corresponding successors of u , we can compose the $\rho(v)$ into a single run ρ on $\text{flat}(s)$. We will prove below that sets of profiles with their natural composition form a \mathbb{F}^\times -algebra and that the function mapping a forest to the set of its profiles is a \mathbb{F}^\times -morphism. This gives an alternative

proof that every language recognised by an automaton is recognised by a finitary \mathbb{F}^\times -algebra.

Syntactic Algebras

Once we have verified that the assumptions of Theorem IV.4.9 are satisfied (compositionality and the existence of syntactic algebras), it follows that the family of MSO-definable languages (for, say, the monad \mathbb{F}^\times) corresponds to the class of MSO-definable \mathbb{F}^\times -algebras. We start by proving the existence of syntactic algebras.

Theorem 5.7. *Let \mathbb{M} be one of the monads \mathbb{T} , \mathbb{T}^\times , \mathbb{F} , or \mathbb{F}^\times . Then \mathbb{M} is essentially finitary over the class of all $\text{MSO}[\leq_{\text{so}}]$ -definable \mathbb{M} -algebras.*

Proof. Let $\mathbb{M}^{\text{reg}}A \subseteq \mathbb{M}^\times A$ be the set of all *regular* forests in $\mathbb{M}^\times A$. Since every regular forest has only finitely many different labels, this functor is finitary. Hence, it remains to prove that the inclusion morphism $\mathbb{M}^{\text{reg}} \Rightarrow \mathbb{M}^\times$ is dense over the class of all finite products of MSO-definable \mathbb{M}^\times -algebras.

Let $\mathfrak{A}_0, \dots, \mathfrak{A}_{n-1}$ be MSO-definable, $B \subseteq A_0 \times \dots \times A_{n-1}$, and $t \in \mathbb{M}^\times B$ a forest with $\pi(t) = \bar{a}$. We have to find a regular forest $t^\circ \in \mathbb{M}^{\text{reg}}B$ with $\pi(t^\circ) = \bar{a}$. Let $C_i \subseteq A_i$ be a finite set of generators of \mathfrak{A}_i . Since \mathfrak{A}_i is MSO-definable, there exists an automaton \mathcal{A}_i recognising the preimage $\pi^{-1}(a_i) \cap \mathbb{M}^\times C_i$. Suppose that Q_i is the set of states of \mathcal{A}_i , $\Omega_i : Q_i \rightarrow \omega$ its priority function, and $K_i := \text{rng } \Omega_i$ the set of priorities used by \mathcal{A}_i . For every $\bar{b} \in B$, we fix forests $\sigma_i(\bar{b}) \in \mathbb{M}^\times C_i$ with $\pi(\sigma_i(\bar{b})) = b_i$, for $i < n$. This defines a function $\sigma_i : \mathbb{V}B \rightarrow \mathbb{M}^\times \mathbb{V}C_i$, which we can extend to a morphism $\hat{\sigma}_i : \mathbb{M}^\times \mathbb{V}B \rightarrow \mathbb{M}^\times \mathbb{V}C_i$.

We construct the desired forest t° by the following variant of the usual Automaton-Pathfinder game. In this game Automaton tries to construct a forest $s \in \mathbb{M}^\times B$ such that, for every $i < n$, $\hat{\sigma}_i(s)$ is accepted by \mathcal{A}_i , while Pathfinder tries to prove that such a forest does not exist. We will define the game in such a way that there is a correspondence between winning strategies for Automaton and such forests s . Note that these are exactly the

forests s with $\pi(s) = \bar{a}$, since

$$\begin{aligned} \pi(\hat{\sigma}_i(s)) &= \pi(\text{flat}(\mathbb{M}^\times \sigma_i(s))) \\ &= \pi(\mathbb{M}^\times \pi(\mathbb{M}^\times \sigma_i(s))) \\ &= \pi(\mathbb{M}^\times p_i(s)) \\ &= p_i(\pi(s)), \end{aligned}$$

where $p_i : A_0 \times \dots \times A_{n-1} \rightarrow A_i$ is the projection to the i -th component. As $\pi(t) = \bar{a}$, it follows that Automaton indeed has a winning strategy for the game. Furthermore, the winning condition of our game is regular. Therefore, it follows by the Büchi-Landweber Theorem that Automaton even has a winning strategy that uses only a finite amount of memory. As the forests s corresponding to finite-memory strategies via the above correspondence are regular, the claim follows.

To conclude the proof, it therefore remains to define a regular game with the above properties. In each round, Automaton picks the label $\bar{b} \in B$ for the next vertex v of s and Pathfinder responds by choosing one of the successors of v . While doing so, we have to keep track of all the states of the various automata from which we want to accept the remaining subforest.

The positions for Automaton are of the form $\bar{U} \in \prod_{i < n} \wp(K_i \times Q_i)$, while those for Pathfinder are finite families of tuples $(\bar{V}_x)_{x \in \xi}$ where $\xi \in \Xi$ and each component \bar{V}_x is a position for Automaton. The initial position belongs to Automaton and consists of the tuple $\{\langle o, q_o^i \rangle\}_{i < n}$, where q_o^i is the initial state of \mathcal{A}_i .

In a position \bar{U} , Automaton chooses

- ♦ an element $\bar{b} \in B$ and,
- ♦ for every $i < n$ and every pair $\langle k, q \rangle \in U_i$, a partial run ρ_q of \mathcal{A}_i on the forest $\sigma_i(\bar{b})$ such that the starting state of ρ_q is q .

(It will turn out that Automaton can choose ρ_q independently of k . So we omit the index k to keep the notation light. We also assume that the sets Q_i are disjoint, so we do not need to specify the index i .) Suppose that $\bar{b} \in B_\xi$ has sort ξ . For $i < n$ and $x \in \xi$, let H_{ix} be the set of all vertices

of $\sigma_i(\bar{b})$ labelled by the variable x . We denote by $W_{ix}(q)$ the set of all pairs $\langle k', q' \rangle \in K_i \times Q_i$ such that there is some $v \in H_{ix}$ with

$$\rho_q(v) = q' \quad \text{and} \quad k' := \min \{ \Omega_i(\rho_q(w)) \mid w \leq v \}.$$

The new position is $(\tilde{V}^x)_{x \in \xi}$ where

$$V_i^x := \bigcup_{\langle k, q \rangle \in U_i} W_{ix}(q).$$

Pathfinder responds by choosing some $x \in \xi$, after which the game proceeds to position \tilde{V}^x .

Automaton wins a play of this game if either the play ends in the position $\langle \rangle$ where Pathfinder cannot make a move, or if the play is infinite and satisfies the following variant of the parity condition. Suppose that the play is $\tilde{U}^0, \tilde{V}^0, \tilde{U}^1, \tilde{V}^2, \dots$ and let $W_{ix}^l(q)$ be the sets used in the l -th turn by Automaton to determine the next position $\tilde{V}^l = (\tilde{V}_x^l)_x$. We call a sequence $k_0, q_0, k_1, q_1, k_2, q_2, \dots$ an i -trace of this play if $\langle k_0, q_0 \rangle \in U_i^0$ and, for all $l < \omega$,

$$\langle k_{l+1}, q_{l+1} \rangle \in W_{ix}^l(q_l), \quad \text{for some } x \text{ with } \tilde{U}^{l+1} = V_x^l.$$

We say that the play satisfies the parity condition if, for all $i < n$,

$$\liminf_{l < \omega} k_l \text{ is even, for all } i\text{-traces } k_0, q_0, k_1, q_1, k_2, q_2, \dots$$

Note that this is a regular winning condition. Furthermore, it is straightforward to check that Automaton wins this game if, and only if, there exists some forest $s \in \mathbb{M}^\times B$ such that, for every $i < n$, the forest $\hat{\sigma}_i(s)$ is accepted by \mathcal{A}_i . \square

Remark. The Bojańczyk-Klin algebra shows that \mathbb{M}^{reg} is not densely embedded in \mathbb{M}^\times over the class of all finitary \mathbb{M}^\times -algebras. \lrcorner

As a consequence we can prove the existence of syntactic algebras (see Theorem III.4.24).

Corollary 5.8. *Let \mathbb{M} be one of the monads \mathbb{T} , \mathbb{T}^\times , \mathbb{F} , or \mathbb{F}^\times . Every MSO-definable language $K \subseteq \mathbb{M}\Sigma$ has a syntactic algebra.*

Compositionality

Our next goal is to show that MSO is varietal and compositional. We start with the latter.

Theorem 5.9. *The logic MSO is \mathbb{F}^\times -compositional and, therefore, also \mathbb{T}^\times -compositional, \mathbb{F} -compositional, and \mathbb{T} -compositional.*

Proof. Because of the translations between formulae and automata, there exists, for every automaton \mathcal{A} and each profile τ of \mathcal{A} , an MSO-formula $\varphi_{\mathcal{A}, \tau}$ stating that there is a partial run of \mathcal{A} on the given forest with profile τ . Furthermore, every MSO-formula is equivalent to a disjunction of formulae of this kind.

For $m < \omega$, let $\text{MSO}_{(m)}$ denote the set of all MSO-formulae equivalent to a formula of the form $\varphi_{\mathcal{A}, \tau}$ where \mathcal{A} is an automaton with at most m states. Since there are only finitely many such automata and each of them has only finitely many profiles of partial runs, it follows that $\text{MSO}_{(m)}$ is finite (up to logical equivalence). Let $\equiv_{(m)}$ be the equivalence relation which holds for two forests if they satisfy the same $\text{MSO}_{(m)}$ -formulae. We claim that $\equiv_{(m)}$ is a congruence relation. This means that, if $S, T \in \mathbb{F}^\times \mathbb{F}^\times \Sigma$ are forests with the same shape, then

$$S(v) \equiv_{(m)} T(v), \quad \text{for all } v, \quad \text{implies} \quad \text{flat}(S) \equiv_{(m)} \text{flat}(T).$$

For the proof, fix a formula $\varphi_{\mathcal{A}, \tau} \in \text{MSO}_{(m)}$ with $\text{flat}(S) \models \varphi_{\mathcal{A}, \tau}$. We have to show that $\text{flat}(T)$ also satisfies $\varphi_{\mathcal{A}, \tau}$, i.e., that there is a partial run of \mathcal{A} on $\text{flat}(T)$ with profile τ . To do so, we introduce the following variant of the Automaton-Pathfinder game. For a given forest $T \in \mathbb{F}^\times \mathbb{F}^\times \Sigma$, Player Automaton tries to prove that there is a partial run of \mathcal{A} on $\text{flat}(T)$ with profile τ , while Pathfinder tries to disprove him. We call a set p of vertices of T a *place* if p is the set of roots of some subforest of T . That is, p is either the set of roots of T , or the set of all x -successors of some vertex $u \in \text{dom}(T)$. The former is called the *root place* of T , while the latter is the *x -successor place* of u . The game starts in the position $\langle r, \tau \rangle$ where r is the root place of T . In a position $\langle p, v \rangle$ where p is a place and v a profile, Automaton tries to show that there exists a partial run ρ on the subforest rooted at p

with profile v . He starts by choosing a family $(\kappa_u)_{u \in p}$ of profiles such that the σ -product of the starting states of the κ_u evaluates to the starting state of v . Pathfinder answers by picking some vertex $u \in p$. Next, Automaton picks a partial run ρ of \mathcal{A} on the forest $T(u)$ whose starting state is the one given by κ_u . Then he has to choose profiles $\tilde{\lambda}$ for all the subforests attached to the copy of $T(u)$ in $\text{flat}(T)$ such that the ‘composition’ of the profile of ρ and $\tilde{\lambda}$ is equal to κ_u . This is done as follows.

Let $\mu = \langle p, \tilde{U} \rangle$ be the profile of ρ . For each component U_x , Automaton chooses a set W_x of triples $\langle k, q, \lambda \rangle$ where k is a priority, q a state, and λ a profile. These sets must satisfy the following conditions.

- U_x is the projection of W_x to the first two components.
- For each $\langle k, q, \lambda \rangle \in W_x$, the state q is equal to the starting state of λ .
- $v = \langle p, \tilde{V} \rangle$ is the composition of μ and the profiles λ . Formally,

$$V_x = \{ \langle l, q' \rangle \mid \langle k, q, \lambda \rangle \in W_x, \lambda = \langle q, \tilde{L} \rangle, \langle k', q' \rangle \in L_x, \\ l = \min \{ k, k' \} \}.$$

Given \tilde{W} , Pathfinder responds by choosing a variable x , an x -successor v_x of u , and a triple $\langle k, q, \lambda \rangle \in W_x$. Then the game continues in the position $\langle v_x, \lambda \rangle$.

If the game reaches a leaf of T , it ends with a win for one of the players. If the leaf is labelled by a variable x and the current position is $\langle v, v \rangle$, then Automaton wins if, and only if, v is of the form $\langle q, \tilde{U} \rangle$ with $U_x = \{q\}$ and $U_z = \emptyset$, for $z \neq x$. Otherwise, Pathfinder wins. If the leaf is not labelled by a variable, then Automaton wins if he can choose $\mu = \langle p, \tilde{U} \rangle$ such that $U_x = \emptyset$, for all x .

In the case where the game is infinite, Automaton wins if the sequence of pairs $\langle k_o, q_o, \lambda_o \rangle, \langle k_1, q_1, \lambda_o \rangle, \dots$ chosen by Pathfinder satisfies the parity condition

$$\liminf_{i < \omega} k_i \text{ is even.}$$

It is straightforward to check that Automaton wins the game on a given forest T if, and only if, there exists a partial run of \mathcal{A} on $\text{flat}(T)$ with profile τ .

(Every partial run of \mathcal{A} on $\text{flat}(T)$ with this profile gives rise to a winning strategy in the game and, conversely, every winning strategy can be used to construct a partial run with the desired profile.)

To conclude the proof we have to show that, if T is a forest with $S(v) \equiv_{(m)} T(v)$, for all v , then Automaton has a winning strategy in the game on T . By construction, Automaton has a winning strategy σ in the game on S . We use it to define a winning strategy σ' in the game on T as follows. If σ tells Automaton to choose a partial run ρ on $S(v)$, σ' returns some partial run ρ' on $T(v)$ with the same profile as ρ . (This is possible since $S(v) \equiv_{(m)} T(v)$.) As only the profile of the chosen run is used by the game and σ is winning, it follows that the resulting strategy σ' is also winning. \square

Remark. Note that in the above proof we have chosen a rather strange stratification of MSO. It might be nice if we could use the usual stratification in terms of the quantifier-rank instead, but this does not seem to work for \mathbb{F}^\times and \mathbb{T}^\times . For the monads \mathbb{F} and \mathbb{T} on the other hand, there is an alternative proof consisting of a simple inductive back-and-forth argument based on the quantifier-rank. \downarrow

According to Theorem IV.4.4, to prove that MSO is varietal it suffices to show that all theory algebras are MSO-definable,

Proposition 5.10. *Let Σ be an alphabet and $\Delta_m := \text{MSO}_{(m)}$ the fragment of MSO used in the proof of Theorem 5.9. The theory algebra $\Theta_{\Delta_m} \Sigma$ is MSO-definable.*

Proof. The set $C := \theta_{\Delta_m}[\Sigma]$ is a finite set of generators of $\Theta_{\Delta_m} \Sigma$. Given a Δ_m -theory $\sigma \in \Theta_{\Delta_m} \Sigma$, we have to find an MSO-formula φ defining the set

$$\pi^{-1}(\sigma) \cap \mathbb{MC}.$$

Every formula in σ is of the form: ‘there exists a partial run of the automaton \mathcal{A} with profile τ ’. Let us write $\chi_{\mathcal{A}, \tau}$ for such a statement. For $t \in \mathbb{MC}$, it follows that $\pi(t) = \sigma$ if, and only if, for every forest $s \in \mathbb{MS}$ with $\mathbb{M}\theta_{\Delta_m}(s) = t$ and every $\chi_{\mathcal{A}, \tau} \in \sigma$, there exists a partial run of \mathcal{A} on s with profile τ . Consequently, to define the above preimage it is sufficient to

express, for a given automaton \mathcal{A} and a profile τ , that every preimage of the given forest t under $\mathbb{M}\theta_{\Delta_m}$ has a partial run of \mathcal{A} with profile τ . This can be done by saying that, for every vertex v , the theory $t(v)$ contains a formula of the form $\chi_{\mathcal{A}, v_v}$, for some profile v_v , such that the ‘composition’ of the profiles v_v yields τ . For this composition, we have to check that the states at the borders match and to compute the minimal priorities on each branch. All of this can easily be expressed in MSO. \square

Corollary 5.11. *MSO is varietal with respect to the functors \mathbb{F}^\times , \mathbb{T}^\times , \mathbb{F} , and \mathbb{T} .*

By the above theorems it follows that the framework we have set up in Chapters III and IV applies to the logic MSO: (I) MSO-definable languages have syntactic algebras which, furthermore, are MSO-definable; (II) the class of all such languages forms a variety of languages; (III) every subvariety can be axiomatised by a set of inequalities. In particular, we can use Theorem IV.4.9 to study the expressive power of monadic second-order logic.

6. First-Order Logic

Let us turn to the logic FO next. Again, we start with compositionality.

Theorem 6.1. *The logic FO is \mathbb{F}^\times -compositional and, therefore, also \mathbb{T}^\times -compositional, \mathbb{F} -compositional, and \mathbb{T} -compositional.*

Proof. Let FO_m denote the set of all first-order formulae of quantifier-rank at most m and denote by \equiv_m equivalence with respect to such formulae. We claim that \equiv_m is a congruence on $\mathbb{F}^\times \Sigma$. For the proof, consider two forests $S, T \in \mathbb{F}^\times \mathbb{F}^\times \Sigma$ with the same shape satisfying

$$S(v) \equiv_m T(v), \quad \text{for all vertices } v.$$

We have to show that $\text{flat}(S) \equiv_m \text{flat}(T)$.

The proof is by induction on m . To make the inductive step go through we have to prove a slightly stronger statement involving parameters. Given a

tuple \bar{a} of vertices of $\text{flat}(S)$ and a copy s of $S(v)$ in $\text{flat}(S)$, we denote by \bar{a}^s the tuple

$$a_i^s := \begin{cases} a_i & \text{if } a_i \in \text{dom}(s), \\ v & \text{if } v \text{ is a hole of } s \text{ and } v \leq_{\text{pf}} a_i \text{ in } \text{flat}(S), \\ v & \text{if } v \text{ is the root of } s \text{ and } v \not\leq_{\text{pf}} a_i \text{ in } \text{flat}(S). \end{cases}$$

We use the same notation for parameters in $\text{flat}(T)$. For a tuple \bar{a} of vertices of some forest s , we write $\langle s, \bar{a} \rangle$ for the expansion of s by constants for the vertices \bar{a} . The claim we prove is that, for forests $S, T \in \mathbb{F}^\times \mathbb{F}^\times \Sigma$ with the same ‘shape’ and with parameters \bar{a} in $\text{flat}(S)$ and \bar{b} in $\text{flat}(T)$,

$$\langle s, \bar{a}^s \rangle \equiv_m \langle t, \bar{b}^t \rangle, \quad \text{for all } v, \text{ copies } s \text{ of } S(v), \text{ and copies } t \text{ of } T(v),$$

implies that

$$\langle \text{flat}(S), \bar{a} \rangle \equiv_m \langle \text{flat}(T), \bar{b} \rangle.$$

For $m = 0$, the proof is straightforward: to see whether $\langle \text{flat}(T), \bar{b} \rangle$ satisfies an atomic formula we only have to check atomic formulae in some of the factors. For the inductive step, suppose that

$$\langle s, \bar{a}^s \rangle \equiv_{m+1} \langle t, \bar{b}^t \rangle, \quad \text{for all } v, s, \text{ and } t.$$

We use a back-and-forth argument to show that

$$\langle \text{flat}(S), \bar{a} \rangle \equiv_{m+1} \langle \text{flat}(T), \bar{b} \rangle.$$

Hence, let $c \in \text{dom}(\text{flat}(S))$ be a new parameter and suppose that c belongs to a copy s of the forest $S(v)$. When we want to apply the inductive hypothesis, we now face the problem that, if $\text{flat}(S)$ contains several copies of $S(v)$, only one of them contains the new parameter. To solve this issue, we have to modify the forests S and T to make sure this does not happen.

Let v_o, \dots, v_n be the path from the root v_o of S to $v = v_n$ and let s_i be the copy of $S(v_i)$ in $\text{flat}(S)$ such that c is a descendant of the root of s_i .

We construct new forests S_0, \dots, S_n and T_0, \dots, T_n as follows. We start with $S_0 := S$ and $T_0 := T$. For the inductive step, suppose we have already defined S_i and T_i for some $i < n$ and that there is a unique copy t_i of $T_i(v_i)$ in $\text{flat}(T_i)$. We choose a vertex d_i of t_i such that

$$\langle s_i, \bar{a}^{s_i} c^{s_i} \rangle \equiv_m \langle t_i, \bar{b}^{t_i} d_i \rangle.$$

Note that the vertex c^{s_i} is a leaf labelled by some variable x . Hence, so is d_i . If there is no other occurrence of x in s_i , we set $S_{i+1} := S_i$. Otherwise, we choose some variable y that does not appear in s_i and we replace every occurrence of x in s_i by y , except for the one at c^{s_i} . Let S_{i+1} be the forest obtained from S_i by

- changing $S(v_i) = s_i$ in this way and
- duplicating the subforest whose root are the x -successors of v_i in such a way that the roots of the new copy are y -successors.

This ensures that $\text{flat}(S_{i+1}) = \text{flat}(S_i)$ and that S_{i+1} contains a unique copy of $S(v_{i+1})$. The forest T_{i+1} is obtained from T_i in exactly the same way.

Having constructed the forests S_n and T_n , we now choose some vertex $d_n \in \text{dom}(T_n(v_n))$ such that

$$\langle s_n, \bar{a}^{s_n} c^{s_n} \rangle \equiv_m \langle t_n, \bar{b}^{t_n} d_n \rangle.$$

Setting $d := d_n$, it follows that $d^{t_i} = d_i$, for all $i \leq n$, which implies that

$$\langle s_i, \bar{a}^{s_i} c^{s_i} \rangle \equiv_m \langle t_i, \bar{b}^{t_i} d^{t_i} \rangle, \quad \text{for all } i \leq n.$$

Note that, if u is a vertex different from v_0, \dots, v_n , s a copy of $S_n(u)$ and t a copy of $T_n(u)$, then c^s is the root of s and d^t is the root of t . Consequently, we also have

$$\langle s, \bar{a}^s c^s \rangle \equiv_m \langle t, \bar{b}^t d^t \rangle.$$

Hence, the forests S_n and T_n together with the parameters \bar{a}, c and \bar{b}, d satisfy our inductive hypothesis and it follows that

$$\langle \text{flat}(S_n), \bar{a}, c \rangle \equiv_m \langle \text{flat}(T_n), \bar{b}, d \rangle.$$

Since $\text{flat}(S_n) = \text{flat}(S)$ and $\text{flat}(T_n) = \text{flat}(T)$, the claim follows.

In the same way we can show that, for every choice of d in $\text{flat}(T)$, we find a matching vertex c in $\text{flat}(S)$. \square

It remains to show that FO is varietal. It turns out that this is only the case for the monads \mathbb{F} and \mathbb{T} , but not for \mathbb{F}^\times or \mathbb{T}^\times .

Proposition 6.2. *FO is closed under inverse \mathbb{F} -morphisms and inverse \mathbb{T} -morphisms.*

Proof. Let $\varphi : \mathbb{F}\Sigma \rightarrow \mathbb{F}\Gamma$ be a morphism of \mathbb{F} -algebras and let $\varphi_o := \varphi \circ \text{sing} : \Sigma \rightarrow \mathbb{F}\Gamma$ be its restriction to Σ . For $s, t \in \mathbb{F}\Sigma$, we will prove that

$$s \equiv_m t \quad \text{implies} \quad \varphi(s) \equiv_m \varphi(t),$$

where \equiv_m denotes equivalence with respect to FO-formulae of quantifier-rank at most m . For the induction step we again need to prove a more general statement involving parameters. We start with setting up a bit of notation.

Note that a forest of the form $\varphi(s) = \text{flat}(\mathbb{F}\varphi_o(s))$ is obtained from s by replacing each vertex u by a forest $\varphi_o(s(u))$. For $s \in \mathbb{F}\Sigma$, we denote by $g_s : \text{dom}(\varphi(s)) \rightarrow \text{dom}(s)$ the function mapping a vertex u of $\varphi(s)$ to the vertex $v := g_s(u)$ such that the copy of the forest $\varphi_o(s(v))$ replacing v in $\varphi(s)$ contains u . (Note that this copy of $\varphi_o(s(v))$ is unique, since we are dealing with the monad \mathbb{F} .)

For an n -tuple \bar{a} of vertices of $\varphi(s)$ and a vertex u of s , we set


$$I_u := \{ i < n \mid g_s(a_i) = u \} \quad \text{and} \quad \bar{a}^u := (a_i)_{i \in I_u},$$

where we consider \bar{a}^u as a tuple of vertices of $\varphi_o(s(u))$.

The statement we will prove by induction on m is the following. Let $s, t \in \mathbb{F}\Sigma$ be forests and \bar{a} and \bar{b} n -tuples of parameters of, respectively, $\varphi(s)$ and $\varphi(t)$. Then

$$\langle s, g_s(\bar{a}) \rangle \equiv_m \langle t, g_t(\bar{b}) \rangle$$

and $\langle \varphi_o(s(u)), \bar{a}^u \rangle \cong \langle \varphi_o(t(v)), \bar{b}^v \rangle$, for all u, v with $I_u = I_v \neq \emptyset$,

Figure 3.: Definition of g_s , I_u , and \tilde{a}^u

implies

$$\langle \varphi(s), \tilde{a} \rangle \equiv_m \langle \varphi(t), \tilde{b} \rangle.$$

For $m = 0$, this is immediate. Hence, suppose that $m > 0$. We have to check the back-and-forth properties. Thus, let $c \in \text{dom}(\varphi(s))$ and set $u := g_s(c)$. Then there is some $v \in \text{dom}(t)$ such that

$$\langle s, g_s(\tilde{a}), g_s(c) \rangle \equiv_{m-1} \langle t, g_t(\tilde{b}), v \rangle.$$

Note that

$$I_u = \{ i \mid g_s(a_i) = u \} = \{ i \mid g_t(b_i) = v \} = I_v.$$

We distinguish two cases. If $I_u = I_v \neq \emptyset$, then there exists an isomorphism

$$\sigma : \langle \varphi_o(s(u)), \tilde{a}^u \rangle \rightarrow \langle \varphi_o(t(v)), \tilde{b}^v \rangle$$

and we can set $d := \sigma(c)$.

Otherwise, $I_u = I_v = \emptyset$ and $s(u) = t(v)$ implies that $\varphi_o(s(u)) \cong \varphi_o(t(v))$. Hence, can choose some element d of $\varphi_o(t(v))$ such that

$$\langle \varphi_o(s(u)), c \rangle \cong \langle \varphi_o(t(v)), d \rangle.$$

In both cases, it now follows that

$$\langle s, g_s(\tilde{a}), g_s(c) \rangle \equiv_m \langle t, g_t(\tilde{b}), g_t(d) \rangle$$

and $\langle \varphi_o(s(u)), \tilde{a}^u c^u \rangle \cong \langle \varphi_o(t(v)), \tilde{b}^v d^v \rangle$, if $I_u = I_v \neq \emptyset$,

which, by inductive hypothesis, implies that

$$\langle \varphi(s), \tilde{a}c \rangle \equiv_{m-1} \langle \varphi(t), \tilde{b}d \rangle.$$

The other direction follows by symmetry. \square

Let us give a counterexample showing that FO is not closed under inverse morphisms of \mathbb{F}^\times -algebras. (The same example works for \mathbb{T}^\times -algebras.) It rests on the following lemma. Recall that a tree is *complete binary* if every non-leaf has exactly two successors.

Lemma 6.3. *There exists a first-order formula φ such that a finite complete binary tree $\mathfrak{T} = \langle T, \text{suc}_o, \text{suc}_i, \leq_{\text{pf}} \rangle$ satisfies φ if, and only if, every leaf of \mathfrak{T} has an even distance from the root.*

Proof. The basic idea is as follows. If every leaf is at an even distance from the root, we can determine whether a vertex x belongs to an even level of the tree by walking a zig-zag path from x downwards until we hit a leaf. For such a path it is trivial to check that its length is even. Hence, our formula only needs to express that the level parities computed in this way are consistent and that the root is on an even level.

To express all this in first-order logic, we first define a few auxiliary formulae.

$$\begin{aligned} \text{suc}(x, y) &:= \text{suc}_o(x, y) \vee \text{suc}_i(x, y) \\ \text{zigzag}(x, y; u, v) &:= [\text{suc}_o(x, y) \wedge \text{suc}_i(u, v)] \\ &\quad \vee [\text{suc}_i(x, y) \wedge \text{suc}_o(u, v)] \\ \text{probe}(x, y) &:= x \leq_{\text{pf}} y \wedge \neg \exists z [\text{suc}(y, z)] \\ &\quad \wedge \forall u \forall v \forall w [x \leq_{\text{pf}} u \wedge \text{suc}(u, v) \wedge \text{suc}(v, w) \\ &\quad \wedge w \leq_{\text{pf}} y \rightarrow \text{zigzag}(u, v; v, w)]. \end{aligned}$$

The first one just states that y is a successor of x ; the second one says that $\langle x, y \rangle$ and $\langle u, v \rangle$ are two edges that go into different directions, one to the

left and one to the right; and the last one states that y is one of the two leaves below x that are reached by a zig-zag path consisting of alternatingly taking left and right successors.

Using these formulae we can express that a vertex x has an even distance from some leaf by

$$\begin{aligned} \text{even}(x) := \exists y [\text{probe}(x, y) \wedge \\ \exists u \exists v [x = y \vee [\text{suc}(x, u) \wedge u \leq_{\text{pf}} v \wedge \text{suc}(v, y) \\ \wedge \text{zigzag}(x, u; v, y)]]] . \end{aligned}$$

Consequently, the desired formula is

$$\begin{aligned} \forall x \forall y [\text{suc}(x, y) \rightarrow [\text{even}(x) \leftrightarrow \neg \text{even}(y)]] \\ \wedge \exists x \forall y [x \leq_{\text{pf}} y \wedge \text{even}(x)] . \end{aligned} \quad \square$$

Corollary 6.4. *FO is not closed under inverse \mathbb{F}^\times -morphisms or inverse \mathbb{T}^\times -morphisms.*

Proof. Let $\Sigma := \{a, c\}$ and $\Gamma := \{b, c\}$ where a is unary, b binary, and c a constant, and let $\varphi := \mathbb{F}^\times \Sigma \rightarrow \mathbb{F}^\times \Gamma$ be the morphism mapping $a(x_o)$ to $b(x_o, x_o)$ and c to c . Let $K \subseteq \mathbb{F}^\times \Gamma$ be the set of all complete binary trees defined by the formula even from Lemma 6.3. Then $\varphi^{-1}[K]$ contains a tree of the form $a^n(c)$ if, and only if, n is even. This is not FO-definable. \square

To summarise, we have obtained the following result.

Corollary 6.5. *FO is varietal with respect to the functors \mathbb{F} and \mathbb{T} , but not with respect to \mathbb{F}^\times or \mathbb{T}^\times .*

Notes

The results about countable chains are originally due to [10]. Our presentation is based on [6]. For an introduction to Green's relations see, e.g., [18, 3]. The Bojańczyk-Klin algebra is from [7] and Theorem 4.5 from [25]. Introductions to automata on infinite trees can be found in [22, 15]. Our model of

a forest automaton is a simplified version of the MSO-automata introduced in [24]. Lemma 6.3 is by Potthoff [19] and its corollary by Bojańczyk and Michalewski [8].

VI. Temporal Logics

1. Temporal Logics

Temporal logics are modal logics talking about transition systems, i.e., labelled directed graphs, and frequently used in verification. Here, we will consider them only over forests instead of arbitrary transition systems. This is not a restriction since each of these logics is closed under a suitable form of bisimulation. In particular, they cannot distinguish between a transition system and its unravelling.

Convention. *In this chapter we will work exclusively in the category $\mathcal{D} = \text{Set}$.*

Well-known temporal logics for forests include basic *modal logic* ML, *computation tree logic* CTL and CTL^* , and *propositional dynamic logic* PDL. We also consider the following variants. We denote by EF the fragment of CTL where we only allow the modal operator EF. *Weak computation tree logic* wCTL* is the variant of CTL^* where the path quantifiers range over *finite* paths only. Finally, we consider several fragments of monadic second-order logic: *Monadic path logic* MPL and *monadic chain logic* MCL are the variants of $\text{MSO}[\leq_{\text{pf}}]$ where quantification is restricted to sets that form, respectively, *paths* and *chains* (i.e., subsets of paths). There are also the *weak* variants of these two logics, WMPL and WMCL, where we can only quantify over *finite* paths and chains.

These logics also have *counting* variants where, instead of the usual modal operators which only check the existence of certain successors or paths, we use counting operators which express, for some constant k , that there are at least k such successors or paths. We call these variants cEF, cCTL^* , cwCTL^* , cPDL, etc.. (Note that the ordinary temporal logics are the bisimulation-invariant fragments of their counting counterparts.)

All of these logics above have the common property that formulae only speak about what happens on a *path* of the given tree. Consequently, if we introduce a quantifier that labels all vertices of a given forest by definable properties of their attached subtrees and then counts the number of certain paths in the resulting forest, we can unify them into a single logic. The definition is as follows.

Definition 1.1. (a) We call a finite set $\Phi = \{\varphi_0, \dots, \varphi_n\}$ of formulae a *syntactic partition* if there are formulae $\vartheta_0, \dots, \vartheta_{n-1}$ such that

$$\varphi_i = \vartheta_i \wedge \bigwedge_{k < i} \neg \vartheta_k, \quad \text{for } i < n, \quad \text{and} \quad \varphi_n = \bigwedge_{k < n} \neg \vartheta_k.$$

(b) Let Φ be a syntactic partition and let \models_t be the relation defined in (c). The Φ -labelling of a path $(v_i)_i$ in a forest s , is the sequence $(\varphi_i)_i$ where φ_i is the unique formula in Φ with $s|_{v_i} \models_t \varphi_i$.

(c) Let \mathcal{K} be a family of ω -word languages. *Counting \mathcal{K} -temporal logic* $c\text{TL}[\mathcal{K}]$ (over a fixed alphabet Σ) has two kinds of formulae: *tree formulae* and *forest formulae*. The set of forest formulae is inductively defined as follows.

- ♦ Forest formulae are closed under finite boolean combinations.
- ♦ For every syntactic partition Φ of tree formulae, every ω -word language $K \in \mathcal{K}[\Phi]$, and every positive integer n , $E_n K$ is a forest formula.

The set of tree formulae is defined as follows.

- ♦ P_a is a tree formula, for every label $a \in \Sigma$.
- ♦ Tree formulae are closed under finite boolean combinations.
- ♦ Every forest formula is a tree formula.

To define the semantics we introduce a satisfaction relation \models_f for forest formulae and one \models_t for tree formulae. In both cases boolean combinations are defined in the usual way. For a tree t , we define

$$\begin{aligned} t \models_t P_a & \quad \text{iff} \quad \text{the root of } t \text{ has label } a, \\ t \models_t \varphi & \quad \text{iff} \quad t' \models_f \varphi, \quad \text{for forest formulae } \varphi, \end{aligned}$$

where t' denotes the forest obtained by removing the root from t . For a forest s we define

$$s \models_f E_n K \quad \text{iff} \quad \text{there exist at least } n \text{ paths each of which has a } \Phi\text{-labelling in } K.$$

We do not require the paths above to be disjoint. They can even be prefixes of each other.

(d) \mathcal{K} -temporal logic $\text{TL}[\mathcal{K}]$ is the fragment of $c\text{TL}[\mathcal{K}]$ consisting of all formulae that do not contain E_n -operators with $n > 1$.

(e) Let \mathcal{S} be a pseudo-variety of ω -semigroups and let \mathcal{K}^∞ , \mathcal{K}^+ , and \mathcal{K}^ω be the classes of, respectively, all ∞ -word languages, all finite word languages, and all ω -word languages recognised by some ω -semigroup in \mathcal{S} . For $x \in \{\infty, +, \omega\}$, we introduce the short hands $c\text{TL}[\mathcal{S}^x] := c\text{TL}[\mathcal{K}^x]$ and $\text{TL}[\mathcal{S}^x] := \text{TL}[\mathcal{K}^x]$.

(f) The *modal rank* of a formula $\varphi \in c\text{TL}[\mathcal{K}]$ is the nesting depth of the modal operators E_n in φ .

The aim of this chapter is to derive algebraic characterisations of various instances of $\text{TL}[\mathcal{K}]$ and $c\text{TL}[\mathcal{K}]$. In many cases we will only consider languages of finitely branching forests to avoid technicalities and to make our task a bit easier.

Definition 1.2. For a set A , we denote by $\mathbb{F}^{\text{fb}} A$ the set of *finitely branching forests* in $\mathbb{F}^* A$ that have only *finitely many roots* and *finitely many holes*.

We start by showing how the logics mentioned above fit into this framework.

Proposition 1.3. *Each of the logics listed in Figure 1 is equivalent to $\text{TL}[\mathcal{K}]$ or $c\text{TL}[\mathcal{K}]$ for the family \mathcal{K} given in the table.*

Most of the equivalences follow by a simple induction on the structure of formulae. The non-trivial cases follow from the equivalences $\text{MPL} = c\text{CTL}^*$, $\text{FO} = \text{WMPL} = \text{cwCTL}^*$, and $\text{cPDL} = \text{WMCL}$, which can be proved using composition arguments.

TL[\mathcal{K}]		cTL[\mathcal{K}]	
logic	\mathcal{K}	logic	\mathcal{K}
EF	$\Sigma^* a$	cEF	$\Sigma^* a$
wCTL	$C^* a$	cwCTL	$C^* a$
CTL	$C^* a$ and C^ω	cCTL	$C^* a$ and C^ω
wCTL*	FO-definable $K \subseteq \Sigma^*$	cwCTL*	FO-definable $K \subseteq \Sigma^*$
CTL*	FO-definable $K \subseteq \Sigma^\infty$	cCTL*	FO-definable $K \subseteq \Sigma^\infty$
PDL	regular $K \subseteq \Sigma^*$	cPDL	regular $K \subseteq \Sigma^*$
		FO	FO-definable $K \subseteq \Sigma^*$
		WMPL	FO-definable $K \subseteq \Sigma^*$
		MPL	FO-definable $K \subseteq \Sigma^\infty$
		WMCL	regular $K \subseteq \Sigma^*$
		MCL	regular $K \subseteq \Sigma^\infty$

Figure 1.: Instances of \mathcal{K} -temporal logic ($a \in \Sigma$ and $C \subseteq \Sigma$)

The rest of this chapter is devoted to the study of logics of the form $\text{TL}[\mathcal{K}]$ and $\text{cTL}[\mathcal{K}]$. We start with a language-theoretic characterisation. The general idea is as follows. Suppose we want to construct a recogniser for a given language, say, an automaton, morphism, or a formula in some logic. One way to make our task easier is to do the recognition in several steps. For instance, we can run a first automaton on the input forest, label the input by the resulting run, and then use a second automaton on the forest expanded in this way. Similarly, we can apply a first morphism to the input forest, label each subtree by its value under this morphism, and then feed the resulting forest to a second morphism. This is the idea behind the following *cascading operation*.

Definition 1.4. Let $K_0, \dots, K_{n-1} \subseteq \mathbb{F}_\emptyset \Sigma$ be forest languages that form a partition of $\mathbb{F}_\emptyset \Sigma$.

(a) For $t \in \mathbb{F}_\emptyset \Sigma$, we denote by $t[\tilde{K}]$ the forest where each vertex gets an additional label encoding to which language K_i the attached subtree belongs.

Formally, we define $t[\tilde{K}] \in \mathbb{F}_\emptyset(\Sigma \times [n])$ by

$$t[\tilde{K}](v) := \langle t(v), i \rangle \quad \text{where } i \text{ is the unique index with } t|_v \in K_i.$$

(b) For $L \subseteq \mathbb{F}_\emptyset(\Sigma \times [n])$, we set

$$L[\tilde{K}] := \{ t \in \mathbb{F}_\emptyset \Sigma \mid t[\tilde{K}] \in L \}.$$

(c) For a family \mathcal{K} of languages, we denote by $\text{Casc}(\mathcal{K})$ the smallest family of languages that contains \mathcal{K} , is closed under finite boolean operations, and satisfies

$$L, \tilde{K} \in \text{Casc}(\mathcal{K}) \implies L[\tilde{K}] \in \text{Casc}(\mathcal{K}).$$

Proposition 1.5. Let \mathcal{K} be a family of ∞ -word languages and \mathcal{M} the family consisting of all forest languages of the form

$$\{ s \in \mathbb{F}_\emptyset \Sigma \mid s \text{ has at least } n \text{ paths with labelling in } K \}$$

and $\{ s \in \mathbb{F}_\emptyset \Sigma \mid \text{the roof of } s \text{ has label } a \},$

for $K \in \mathcal{K}[\Sigma]$, $a \in \Sigma$, and $n < \omega$.

- (a) A language $L \subseteq \mathbb{F}_\emptyset \Sigma$ is $\text{cTL}[\mathcal{K}]$ -definable if, and only if, $L \in \text{Casc}(\mathcal{M})$.
- (b) If we modify the definition of \mathcal{M} to only allow $n = 1$, we obtain an analogous characterisation of $\text{TL}[\mathcal{K}]$ -definable languages.

Proof. We only prove the claim for $\text{cTL}[\mathcal{K}]$. The second statement follows in exactly the same way.

(\Leftarrow) Clearly, every language in \mathcal{M} is $\text{cTL}[\mathcal{K}]$ -definable. Hence, the claim follows by induction if we can show that, if L and \tilde{K} are $\text{cTL}[\mathcal{K}]$ -definable languages, then so is $L[\tilde{K}]$. Suppose that L is defined by the formula φ and K_i by ψ_i . Then we can define $L[\tilde{K}]$ by the formula obtained from φ by replacing every atom of the form $P_{(a,i)}$ by the formula $P_a \wedge \psi_i$.

(\Rightarrow) Given a formula $\varphi \in \text{cTL}[\mathcal{K}]$ we prove by induction on φ that it defines a language in $\text{Casc}(\mathcal{M})$. For $\varphi = P_a$, the defined language is

$$\{ s \in \mathbb{F}_\emptyset \Sigma \mid \text{the roof of } s \text{ has label } a \},$$

which belongs to $\mathcal{M} \subseteq \text{Casc}(\mathcal{M})$ by assumption. As $\text{Casc}(\mathcal{M})$ is closed under finite boolean operations, it remains to consider formulae of the form $\varphi = E_n K$ with $K \in \mathcal{K}[\Phi]$, for some syntactic partition Φ . For every $\psi \in \Phi$, we know by inductive hypothesis that the language L_ψ it defines belongs to $\text{Casc}(\mathcal{M})$. By assumption, so does the language

$$M := \{s \in \mathbb{F}_{\mathcal{O}}\Sigma \mid s \text{ has at least } n \text{ paths with labelling in } K\}.$$

Consequently, $M[\bar{L}] \in \text{Casc}(\mathcal{M})$, which is the language defined by $E_n K$. \square

2. Bisimulation

Before approaching characterisations of various temporal logics, let us start with the simpler case of bisimulation invariance. This example also explains why we have chosen to use algebras with elements of higher arities. As a warm-up, we start with a very simple example: that of sibling-commutative languages.

Definition 2.1. (a) A forest s is a *permutation* of $t \in \mathbb{F}A$ if s is obtained from t by simultaneously rearranging the roots and the successors of every vertex. Formally, we call a function $\sigma : \text{dom}(s) \rightarrow \text{dom}(t)$ a *permutation* if it is bijective and it preserves the successor and sibling relations. Then s is a *permutation* of t if there exists some permutation $\text{dom}(s) \rightarrow \text{dom}(t)$.

(a) A language $K \subseteq \mathbb{T}\Sigma$ is *sibling-commutative* if it is closed under permutations. \lrcorner

Note that being sibling-commutative is not the same as being closed under rearranging the successors of a single vertex (or finitely many of them).

Theorem 2.2. A regular language $K \subseteq \mathbb{F}_{\mathcal{O}}^{\text{fb}}\Sigma$ is *sibling-commutative* if, and only if, its syntactic algebra $\text{Syn}(K)$ satisfies the equations

$$c + d = d + c \quad \text{and} \quad a(x_o, \dots, x_{m-1}) = a(x_{\sigma(o)}, \dots, x_{\sigma(m-1)}),$$

for all sorts $\xi = \{x_o, \dots, x_{m-1}\} \in \Xi$, elements $a \in A_\xi$ and $c, d \in A_{\mathcal{O}}$, and permutations $\sigma : [m] \rightarrow [m]$.

Proof. (\Leftarrow) If s is a permutation of t , we have $\text{syn}_K(s) = \text{syn}_K(t)$. Since syn_K recognises K it follows that $s \in K \Leftrightarrow t \in K$ and that K is sibling-commutative.

(\Rightarrow) Fix an element $a \in A_\xi$ and a permutation $\sigma : [m] \rightarrow [m]$. We have to show that

$$a(x_o, \dots, x_{m-1}) \sim_K a(x_{\sigma(o)}, \dots, x_{\sigma(m-1)})$$

(where \sim_K is the syntactic congruence of K). Hence, let p be a context. Note that the two forests obtained from p by replacing the hole \square by, respectively, $a(x_o, \dots, x_{m-1})$ and $a(x_{\sigma(o)}, \dots, x_{\sigma(m-1)})$ are permutations of each other. As K is sibling-commutative we therefore have

$$p[a(x_o, \dots, x_{m-1})] \in K \Leftrightarrow p[a(x_{\sigma(o)}, \dots, x_{\sigma(m-1)})] \in K.$$

The equation $c + d = d + c$ follows analogously. \square

Remark. It follows in particular that sibling-commutativity is decidable. Given a regular language K , we can compute its syntactic algebra and check whether it satisfies the above equations. (We only need to check them for elements a in a finite set of generators.) \lrcorner

The characterisation of bisimulation is similar. We just have to account for the duplication of successors. Let us recall the definition.

Definition 2.3. (a) A *bisimulation* between two forests s and t is a binary relation $Z \subseteq \text{dom}(s) \times \text{dom}(t)$ such that $\langle u, v \rangle \in Z$ implies that

- ♦ $s(u) = t(v)$ and,
- ♦ for every x -successor u' of u , there is some x -successor v' of v with $\langle u', v' \rangle \in Z$ and vice versa.

(b) Two trees are *bisimilar* if there exists a bisimulation between them that relates their roots. More generally, two forests are bisimilar if every component of one is bisimilar to some component of the other.

(c) A language K of forests is *bisimulation-invariant* if $s \in K$ implies $t \in K$, for every forest t bisimilar to s . \lrcorner

Theorem 2.4. Let Σ be an alphabet where every symbol has arity 1, and let $K \subseteq \mathbb{F}_{\emptyset}\Sigma$ be a forest language with a syntactic algebra. Then K is bisimulation-invariant if, and only if, $\text{Syn}(K)$ satisfies the following equations:

$$\begin{aligned} c + c &= c, & a(x + x) &= a(x), \\ c + d &= d + c, & a(x_0 + x_1 + x_2 + x_3) &= a(x_0 + x_2 + x_1 + x_3), \end{aligned}$$

for all $a \in \text{Syn}_{\{x\}}(K)$ and $c, d \in \text{Syn}_{\emptyset}(K)$.

Proof. We denote the syntactic congruence of K by \sim_K .

(\Rightarrow) Given elements $c, d \in \text{Syn}_{\emptyset}(K)$, we fix forests $s \in \text{syn}_K^{-1}(c)$ and $t \in \text{syn}_K^{-1}(d)$. If K is bisimulation-invariant, we have

$$\begin{aligned} p[s] \in K & \quad \text{iff} \quad p[s + s] \in K, \\ p[s + t] \in K & \quad \text{iff} \quad p[t + s] \in K, \end{aligned}$$

for every context p . Consequently, $s \sim_K s + s$ and $s + t \sim_K t + s$, which implies that $c = c + c$ and $c + d = d + c$.

The remaining two equations are proved similarly. Fix $a \in \text{Syn}_{\{x\}}(K)$ and $s \in \text{syn}_K^{-1}(a)$. Setting $s' := s(x_0 + x_0)$, bisimulation-invariance of K implies that

$$p[s] \in K \quad \text{iff} \quad p[s'] \in K, \quad \text{for every context } p.$$

Consequently $s \sim_K s'$ and $a(x_0) = \text{syn}_K(s) = \text{syn}_K(s') = a(x_0 + x_0)$.

Similarly, for $t := s(x_0 + x_1 + x_2 + x_3)$ and $t' := s(x_0 + x_2 + x_1 + x_3)$, we have

$$p[t] \in K \quad \text{iff} \quad p[t'] \in K, \quad \text{for every context } p.$$

Hence, $t \sim_K t'$ and $a(x_0 + x_1 + x_2 + x_3) = a(x_0 + x_2 + x_1 + x_3)$.

(\Leftarrow) Suppose that $\text{Syn}(K)$ satisfies the four equations above and let s and s' be bisimilar forests. We claim that $\text{syn}_K(s) = \text{syn}_K(s')$, which implies that $s \in K \Leftrightarrow s' \in K$.

Fix a bisimulation relation $Z \subseteq \text{dom}(s) \times \text{dom}(s')$. W.l.o.g. we may assume that Z only relates vertices on the same level of the respective forests




Figure 2.: Transforming bisimilar forests

and that it only relates vertices whose predecessors are also related. (If not, we can always remove the pairs not satisfying this condition without destroying the fact that Z is a bisimulation.) Let \approx be the equivalence relation on $\text{dom}(s) \cup \text{dom}(s')$ generated by Z .

We will transform the forests s and s' in several steps while preserving their value under syn_K until both forests are equal. (Note that each of these steps necessarily modifies the given forest at every vertex.) An example of this process can be found in Figure 2. The first step consists in translating the problem into the algebra $\text{Syn}(K)$. We define two new forests $t_0, t'_0 \in \mathbb{F}_{\emptyset}\text{Syn}(K)$ with the same domains as, respectively, s and s' and the following labelling. If $v \in \text{dom}(s)$ has the successors u_0, \dots, u_{n-1} , we set

$$t_0(v) := \text{syn}_K(s(v))(x_0 + \dots + x_{n-1})$$

and we make u_i an x_i -successor of v in t_o . We obtain t'_o from s' in the same way. By associativity it follows that $\pi(t_o) = \text{syn}_K(s)$ and $\pi(t'_o) = \text{syn}_K(s')$.

Next we make the shapes of the forests t_o and t'_o the same. Let t_1 and t'_1 be the forests with the same domains as t_o and t'_o and the following labelling. For every vertex v of t_o with successors u_o, \dots, u_{n-1} and labelling

$$t_o(v) = a(x_o + \dots + x_{n-1}),$$

we set

$$t_1(v) := a(x_o + \dots + x_o + \dots + x_{n-1} + \dots + x_{n-1}),$$

where each variable x_i is repeated m_i times and the numbers m_i are determined as follows. Let M be some number such that, for every $i < n$, no vertex $v' \approx v$ has at more than M successors u' with $u' \approx u_i$. (Note that there are only finitely many such vertices.) We choose the constants m_i such that

$$\sum_{k \in U_i} m_k = M, \quad \text{where} \quad U_i := \{k < n \mid u_k \approx u_i\}.$$

We obtain the forest t'_1 in the same way from t'_o . By the top right equation in the statement of the theorem, the value of the product is not affected by this modification. Hence, $\pi(t_1) = \pi(t_o)$ and $\pi(t'_1) = \pi(t'_o)$.

Finally, let t_2 and t'_2 be the unravelling of, respectively, t_1 and t'_1 , i.e., the forest where for every vertex v with successors u_o, \dots, u_{n-1} and label

$$t_1(v) = a(x_o + \dots + x_o + \dots + x_{n-1} + \dots + x_{n-1}),$$

we set

$$t_2(v) := a(x_o + \dots + x_k + \dots + x_l + \dots + x_m),$$

where we number the variables from left-to-right, e.g., $a(x_o + x_o + x_1 + x_2 + x_2)$ becomes $a(x_o + x_1 + x_2 + x_3 + x_4)$, and we duplicate each attached subforest a corresponding number of times such that the value of the product does not change. We do the same for t'_2 .

We have arrived at a situation where, for each component r of the forests t_2 , there is some component r' of t'_2 that differs only in the ordering of successors, but not in their number. Consequently, there exists a bijection $\sigma : \text{dom}(t) \rightarrow \text{dom}(r')$ such that, for a vertex v of r with successors u_o, \dots, u_{n-1} ,

$$r'(v) = r(v)(x_{\sigma_v(o)} + \dots + x_{\sigma_v(n-1)}),$$

where the function $\sigma_v : [n] \rightarrow [n]$ is chosen such that $\sigma(u_i)$ is the $x_{\sigma_v(i)}$ -successor of $\sigma(v)$.

Let \hat{r} be the tree obtained from r as follows. For a vertex v with successors u_o, \dots, u_{n-1} and labelling

$$r(v) = a(x_o + \dots + x_{n-1}),$$

we set

$$\hat{r}(v) := a(x_{\sigma_v(o)} + \dots + x_{\sigma_v(n-1)}),$$

and we reorder the attached subtrees accordingly. By associativity and the bottom right equation, this does not change the value of the product. It follows that $\hat{r} = r'$. Consequently, $\pi(r) = \pi(r')$.

We have shown that, for every component of t_o there is some component of t'_o with the same product. Therefore, we can write

$$\pi(t_o) = a_o + \dots + a_{m-1} \quad \text{and} \quad \pi(t'_o) = b_o + \dots + b_{n-1}$$

where the sets $\{a_o, \dots, a_{m-1}\}$ and $\{b_o, \dots, b_{n-1}\}$ coincide. Using the equations $c + c = c$ and $c + d = d + c$ we can therefore transform $\pi(t_o)$ into $\pi(t'_o)$. Consequently,

$$\text{syn}_K(s) = \pi(t_o) = \pi(t'_o) = \text{syn}_K(s').$$

As syn_K recognises K it follows that $s \in K \Leftrightarrow s' \in K$, as desired. \square

Note that we immediately obtain a decision procedure for bisimulation-invariance from this theorem, since we can compute the syntactic algebra and check whether it satisfies the given set of equations.

Corollary 2.5. *It is decidable whether a given regular language K is bisimulation-invariant.*

3. The Logic EF

One of the simplest temporal logics is the logic EF a fragment of CTL where we only allow the modal operator EF . As remarked in Proposition 1.3, we have $EF = TL[\mathcal{K}]$ and $cEF = cTL[\mathcal{K}]$ where \mathcal{K} is the family of all languages of the form $\Sigma^* a$, for some $a \in \Sigma$.

Definition 3.1. For $n, m < \omega$, we denote by cEF_n the fragment of cEF that uses only operators E_l where $l \leq n$, and cEF_n^m is the fragment of cEF_n where the nesting depth of the operators E_l is restricted to m . For $n = 1$, we set $EF := cEF_1$ and $EF^m := cEF_1^m$. \dashv

The following is our main theorem. Before giving the statement a few technical remarks are in order. In the equations below we make use of the ω -power a^ω of an element $a \in A_{\{x\}}$ (which is the infinite vertical product $aaa \dots$), and the *idempotent power* a^π (which is defined as $a^\pi = a^n$ for the minimal number n with $a^n a^n = a^n$). For the horizontal semigroup we use multiplicative notation instead: $n \times a$ for $a + \dots + a$ and $\pi \times a$ for $n \times a$ with n as above.

When writing an ω -power of an element of arity greater than one, we need to specify with respect to which variable we take the power. We use the notation $a^{\omega x}$ to indicate that the variable x should be used. Note that, when using several ω -powers like in $(a(x, (b(x, y))^{\omega y}))^{\omega x}$, the intermediate term after resolving the inner power can be a forest with infinitely many occurrences of the variable x . But after resolving the outer ω -power, we obtain a forest without variables, i.e., a proper element of $\mathbb{F}_\emptyset^{\text{fb}} A$. Consequently, the equations below are all well-defined.

Theorem 3.2. A language $K \subseteq \mathbb{F}_\emptyset^{\text{fb}} \Sigma$ is definable in the logic cEF_n if, and only

if, its syntactic algebra $\text{Syn}(K)$ satisfies the following equations:

$$\begin{aligned} c + d &= d + c & (a(x) + b(x))^\omega &= (ab(x))^\omega \\ (ab)^\pi &= b(ab)^\pi & (a(x) + c)^\omega &= (a(x + c))^\omega \\ a^\omega + a^\omega &= a^\omega & (a(x + c + c))^\omega &= (a(x + c))^\omega \\ (abb')^\omega &= (ab'b)^\omega & [a(b(x, y))^{\omega y}]^{\omega x} &= [ab(x, x)]^{\omega x} \\ (aab)^\omega &= (ab)^\omega & [a(x + bc + c)]^\omega &= [a(x + bc)]^\omega \\ a_\xi(c, \dots, c) + (n - |\xi|) \times c &= a_\xi(c, \dots, c) + (n - |\xi| + 1) \times c, \\ [a(x + (a(n \times x))^\pi(c))]^\omega &= n \times (a(n \times x))^\pi(c) \end{aligned}$$

for all $a, b, b' \in \text{Syn}_{\{x\}}(K)$, $c, d \in \text{Syn}_\emptyset(K)$, $a_\xi \in \text{Syn}_\xi(K)$ with $|\xi| \leq n$.

Before giving the proof of the theorem, let us present some of the consequences.

Corollary 3.3. For fixed n , it is decidable whether a given regular language K is cEF_n -definable.

For the logic cEF, where the value of n is not bounded, a similar result can now be derived as a simple corollary. The basic argument is contained in the following lemma.

Lemma 3.4. Given a finitary \mathbb{F}^{fb} -algebra \mathfrak{A} that is generated by $A_\emptyset \cup A_{\{x\}}$, we can compute a number N such that, if \mathfrak{A} satisfies the equations of Theorem 3.2 for some value of n , it satisfies them for $n = N$.

Proof. Set $N := m_\emptyset^{2m_1} + m_\emptyset$ where $m_\emptyset := |A_\emptyset|$ and $m_1 := |A_{\{x\}}|$. By assumption there is some number n for which \mathfrak{A} satisfies the equations of Theorem 3.2. W.l.o.g. we may assume that $n \geq N$. The only two equations depending on n are

$$\begin{aligned} (1)_n \quad & a_\xi(c, \dots, c) + (n - |\xi|) \times c = a_\xi(c, \dots, c) + (n - |\xi| + 1) \times c \\ (2)_n \quad & [a(x + (a(n \times x))^\pi(c))]^\omega = n \times (a(n \times x))^\pi(c) \end{aligned}$$

We have to show that \mathcal{A} also satisfies $(1)_N$ and $(2)_N$.

For $(2)_N$, note that $n \geq N \geq |A_\emptyset|$ implies that $N \times c = \pi \times c = n \times c$, for all $c \in A_\emptyset$. Consequently,

$$a(N \times x)(c) = a(n \times x)(c)$$

and, therefore,

$$(a(N \times x))^\pi(c) = (a(n \times x))^\pi(c).$$

This implies the claim.

For $(1)_N$, fix $a \in A_\xi$ and $c \in A_\emptyset$. If $|\xi| \leq N - m_o$, then $N - |\xi| \geq m_o = |A_\emptyset|$ implies that $(N - |\xi|) \times c = \pi \times c$. Consequently,

$$\begin{aligned} a(c, \dots, c) + (N - |\xi|) \times c &= a(c, \dots, c) + \pi \times c \\ &= a(c, \dots, c) + \pi \times c + c \end{aligned}$$

and we are done. Thus, we may assume that $|\xi| > N - m_o = m_o^{2m_1}$. As \mathcal{A} is generated by $A_\emptyset \cup A_{\{x\}}$, there exists some forest $s \in \mathbb{F}_\xi(A_\emptyset \cup A_{\{x\}})$ with $\pi(s) = a$. We distinguish several cases.

If some variable $x \in \xi$ does not appear in s , we can use $(1)_n$ to show that

$$\begin{aligned} &a(c, \dots, c, \dots, c) + (N - |\xi|) \times c \\ &= a(c, \dots, c + \dots + c, \dots, c) + (N - |\xi|) \times c \\ &= a(c, \dots, n \times c, \dots, c) + (N - |\xi|) \times c \\ &= a(c, \dots, n \times c, \dots, c) + (N - |\xi|) \times c + c. \end{aligned}$$

Next, suppose that s is highly branching in the sense that it has the form

$$s = r(t_o + \dots + t_{m_o^2-1})$$

where each subtree t_i contains some variable. Then we can find indices $i_o < \dots < i_{m_o-1}$ such that $\pi(t_{i_o}(\bar{c})) = \dots = \pi(t_{i_{m_o-1}}(\bar{c}))$ (where \bar{c} denotes as many copies of c as appear in the respective term). Hence, $(1)_n$ again

implies that

$$\begin{aligned} &a(\bar{c}) + (N - |\xi|) \times c \\ &= \pi(s(\bar{c})) + (N - |\xi|) \times c \\ &= \pi(r(t_o(\bar{c}) + \dots + t_{m_o^2-1}(\bar{c}))) + (N - |\xi|) \times c \\ &= \pi(r(t_o(\bar{c}) + \dots + t_{m_o^2-1}(\bar{c}) + n \times t_{i_o}(\bar{c}))) + (N - |\xi|) \times c \\ &= a(\bar{c}) + (N - |\xi|) \times c + c. \end{aligned}$$

Note that a tree of height $h := m_1$ where every vertex has at most $d := m_o^2$ successors has at most $d^h = m_o^{2m_1}$ leaves. Hence, if s is not highly branching in the sense above, the fact that it contains $|\xi| > m_o^{2m_1}$ variables implies that there must be a chain $v_o < \dots < v_{m_1}$ of vertices such that, for every $i < m_1$, there is some leaf u labelled by a variable with $v_{i-1} < u$ and $v_i \not\leq u$. (For $i = o$, we omit the first condition.) Hence, we can decompose s as

$$s(\bar{c}) = r_o(\bar{c}, r_1(\bar{c}, \dots, r_{m_1}(\bar{c}))).$$

For $i < j$, set

$$r_{ij} := r_i(\bar{c}, r_{i+1}(\bar{c}, \dots, r_{j-1}(\bar{c}, x))).$$

Then there are two indices $i < j$ such that

$$\pi(r_{oi}) = \pi(r_{oj}).$$

Consequently, we can use pumping to obtain a term

$$\pi(s(\bar{c})) = \pi(r_{oi}(r_{ij})^n r_{jm_i})$$

which contains at least n occurrences of c , and the claim follows again by $(1)_n$. \square

According to this lemma, we can check for cEF-definability of a language K , by computing its syntactic algebra $\text{Syn}(K)$, the associated constant N , and then checking the equations for $n = N$.

Corollary 3.5. *It is decidable whether a given regular language K is cEF-definable.*

When taking the special case of $n = 1$ in Theorem 3.2, we obtain the following characterisation of EF-definability.

Theorem 3.6. *A language $K \subseteq \mathbb{F}_{\emptyset}^{\text{fb}} \Sigma$ is definable in the logic EF if, and only if, the syntactic algebra $\text{Syn}(K)$ satisfies the following equations:*

$$\begin{aligned} c + d &= d + c & (a(x) + b(x))^\omega &= (ab(x))^\omega \\ (ab)^\pi &= b(ab)^\pi & (a(x) + c)^\omega &= (a(x + c))^\omega \\ (abb')^\omega &= (ab'b)^\omega & (a(x + c + c))^\omega &= (a(x + c))^\omega \\ (aab)^\omega &= (ab)^\omega & [a(b(x, y))^{\omega y}]^{\omega x} &= [ab(x, x)]^{\omega x} \\ ac &= ac + c & c &= c + c & [a(x + a^\pi c)]^\omega &= a^\pi c, \end{aligned}$$

for all $a, b, b' \in \text{Syn}_{\{x\}}(K)$ and $c, d \in \text{Syn}_{\emptyset}(K)$.

Corollary 3.7. *It is decidable whether a given regular language $K \subseteq \mathbb{F}_{\emptyset}^{\text{fb}} \Sigma$ is EF-definable.*

The Proof

For the proof of Theorem 3.2, we need to set up a bit of machinery. We start by defining the suitable notion of bisimulation for cEF_n . The difference to the standard notion is that we use reachability instead of the edge relation and that we also have to preserve the number of reachable positions.

Definition 3.8. Let $m, n < \omega$.

(a) For trees $s, t \in \mathbb{F}^{\text{fb}} \Sigma$, we define

$$s \approx_n^0 t \quad : \text{iff} \quad \text{the roots of } s \text{ and } t \text{ have the same label,}$$

and, inductively, we set $s \approx_n^{m+1} t$ if,

- the roots of s and t have the same label,

- for every n -tuple \bar{x} in $\text{dom}(s)$ not containing the root, there is some n -tuple \bar{y} in $\text{dom}(t)$ not containing the root such that

$$s|_{x_i} \approx_n^m t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j, \quad \text{for all } i, j < n,$$

- for every n -tuple \bar{y} in $\text{dom}(t)$ not containing the root, there is some n -tuple \bar{x} in $\text{dom}(s)$ not containing the root such that

$$s|_{x_i} \approx_n^m t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j, \quad \text{for all } i, j < n.$$

To simplify notation, we will frequently write $x \approx_n^m y$ for vertices x and y instead of the more cumbersome $s|_x \approx_n^m t|_y$.

(b) For two forests $s, t \in \mathbb{F}^{\text{fb}} \Sigma$ with possibly several components, we set $s \sim_n^{m+1} t$ if

- for every n -tuple \bar{x} in s there is some n -tuple \bar{y} in t such that

$$s|_{x_i} \approx_n^m t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j, \quad \text{for all } i, j < n,$$

- for every n -tuple \bar{y} in t there is some n -tuple \bar{x} in s such that

$$s|_{x_i} \approx_n^m t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j, \quad \text{for all } i, j < n.$$

Let us show that this notion of bisimulation captures the expressive power of cEF. The proof is mostly standard. We start by introducing the following notion of a type.

Definition 3.9. (a) We define the type $\text{tp}_n^m(s)$ of a tree $s \in \mathbb{F} \Sigma$ by

$$\text{tp}_n^0(s) := a \quad \text{and} \quad \text{tp}_n^{m+1}(s) := \langle a, \theta_s \rangle,$$

where a is the label at the root of s and

$$\theta_s := \left\{ \langle l, \sigma \rangle \mid l \leq n, x_0, \dots, x_{l-1} \in \text{dom}(s) \text{ distinct, not equal to the root, } \sigma = \text{tp}_n^m(s|_{x_0}) = \dots = \text{tp}_n^m(s|_{x_{l-1}}) \right\}.$$

(b) For an arbitrary forest $s \in \mathbb{F} \Sigma$, we set

$$\text{Tp}_n^{m+1}(s) := \theta_s,$$

where

$$\theta_s := \{ \langle l, \sigma \rangle \mid l \leq n, x_0, \dots, x_{l-1} \in \text{dom}(s) \text{ distinct}, \\ \sigma = \text{tp}_n^m(s|_{x_0}) = \dots = \text{tp}_n^m(s|_{x_{l-1}}) \}.$$

A standard proof establishes the following equivalences.

Lemma 3.10. *Let $n, m < \omega$.*

(a) *For trees $s, t \in \mathbb{F}_\emptyset \Sigma$, the following statements are equivalent.*

- (1) $s \approx_n^m t$
- (2) $\text{tp}_n^m(s) = \text{tp}_n^m(t)$
- (3) $s \models \varphi \Leftrightarrow t \models \varphi$, for all $\varphi \in \text{cEF}_n^m$.

(b) *For arbitrary forests $s, t \in \mathbb{F}_\emptyset \Sigma$, the following statements are equivalent.*

- (1) $s \sim_n^m t$
- (2) $\text{Tp}_n^m(s) = \text{Tp}_n^m(t)$
- (3) $s \models \varphi \Leftrightarrow t \models \varphi$, for all $\varphi \in \text{cEF}_n^m$.

Proof. (a) (2) \Rightarrow (1) follows by a straightforward induction on m and (1) \Rightarrow (3) by induction on φ . For (3) \Rightarrow (2) it is sufficient to show that, for every type τ , there exists a formula $\chi_\tau \in \text{EF}_n^m$ such that

$$s \models \chi_\tau \quad \text{iff} \quad \text{tp}_n^m(s) = \tau, \quad \text{for every tree } s.$$

We proceed by induction on m . If $m = 0$, the type τ is of the form $a \in \Sigma$. Hence, we can set $\chi_\tau := P_a$. If $m > 0$, then $\tau = \langle a, \theta \rangle$ for some $a \in \Sigma$ and some set θ of types of lower rank. We can set

$$\chi_\tau := P_a \wedge \bigwedge_{\langle l, \sigma \rangle \in \theta} E_l \chi_\sigma \wedge \bigwedge_{\langle l, \sigma \rangle \notin \theta} \neg E_l \chi_\sigma.$$

(b) is proved in the same way. \square

Corollary 3.11. *A language $L \subseteq \mathbb{F}\Sigma$ is cEF_n^m -definable if, and only if, it is regular and satisfies*

$$s \sim_n^m t \quad \text{implies} \quad s \in L \Leftrightarrow t \in L, \quad \text{for all regular forests } s, t \in \mathbb{F}_\emptyset \Sigma.$$

Proof. (\Rightarrow) follows by the implication (1) \Rightarrow (3) of Lemma 3.10.

(\Leftarrow) Set

$$\varphi := \bigvee \{ \chi_\tau \mid \tau = \text{Tp}_n^m(s) \text{ for some regular forest } s \in L \},$$

where χ_τ are the formulae from the proof of Lemma 3.10. For a regular forest $t \in \mathbb{F}_\emptyset \Sigma$, it follows that

$$\begin{aligned} t \models \varphi & \quad \text{iff} \quad \text{Tp}_n^m(t) = \text{Tp}_n^m(s), \quad \text{for some regular forest } s \in L, \\ & \quad \text{iff} \quad t \sim_n^m s, \quad \text{for some regular forest } s \in L, \\ & \quad \text{iff} \quad t \in L. \end{aligned}$$

Let K be the language defined by φ . Since L and K are both regular languages that contain the same regular forests, it follows that $L = K$. Thus, L is cEF_n^m -definable. \square

We want to show that an algebra recognises cEF_n -definable languages if, and only if, it satisfies the following equations.

Definition 3.12. (a) An \mathbb{F} -algebra \mathfrak{A} is an *algebra for cEF_n* if it is finitary, generated by $A_\emptyset \cup A_{\{x\}}$, and satisfies the following equations.

- (G1)_n $a_\xi(c, \dots, c) + (n - |\xi|) \times c = a_\xi(c, \dots, c) + (n - |\xi| + 1) \times c$
- (G2) $(ab)^\pi = b(ab)^\pi$
- (G3) $a^\omega + a^\omega = a^\omega$
- (G4) $c + d = d + c$
- (G5) $(a(x) + b(x))^\omega = (ab(x))^\omega$
- (G6) $(a(x) + c)^\omega = (a(x + c))^\omega$
- (G7) $(a(x + c + c))^\omega = (a(x + c))^\omega$
- (G8) $[a(b(x, y))^{\omega y}]^{\omega x} = [ab(x, x)]^{\omega x}$
- (G9) $(abb')^\omega = (ab'b)^\omega$
- (G10) $(aab)^\omega = (ab)^\omega$
- (G11) $[a(x + bc + c)]^\omega = [a(x + bc)]^\omega$

$$(G12)_n \quad [a(x + (a(n \times x))^\pi(c))]^\omega = n \times (a(n \times x))^\pi(c)$$

where $a, b, b' \in A_{\{x\}}$, $c, d \in A_\emptyset$, $a, \xi \in A_\xi$, and $|\xi| \leq n$.

(b) An \mathbb{F} -algebra \mathfrak{A} is an algebra for cEF if it is an algebra for cEF $_n$, for some $n \geq 1$.

In the proof that algebras for cEF recognise exactly the cEF-definable languages, we use one of the Green's relations (suitably modified for \mathbb{F} -algebras).

Definition 3.13. Let \mathfrak{A} be an \mathbb{F} -algebra. For $a, b \in A_\emptyset$, we define

$$a \leq_L b \quad : \text{iff} \quad a = c(b) \quad \text{or} \quad a = b + d, \\ \text{for some } c \in A_{\{x\}}, d \in A_\emptyset.$$

Lemma 3.14. Let \mathfrak{A} be an algebra for cEF $_k$.

(a) The relation \leq_L is antisymmetric.

(b) For $a \in A_{\{x\}}$, $c \in A_\emptyset$, we have

$$c = c + c \quad \text{implies} \quad ac = ac + c, \\ c = a(c, c) \quad \text{implies} \quad c = c + c.$$

Proof. (a) For a contradiction, suppose that there are elements $a \neq b$ with $a \leq_L b \leq_L a$. By definition, we can find elements c and d such that

$$(1) \ a = c(b) \text{ or } (2) \ a = b + c, \quad \text{and} \quad (i) \ b = d(a) \text{ or } (ii) \ b = a + d.$$

We have thus to consider four cases. In each of them we obtain a contradiction via (G1) $_k$ or (G2).

$$(1, i) \quad a = cb = cda = (cd)^\pi(a) = d(cd)^\pi(a) = da = b. \\ (1, ii) \quad a = cb = c(a + d) = (c(x + d))^\pi(a) \\ = (c(x + d))^\pi(a) + d = a + d = b. \\ (2, i) \quad b = da = d(b + c) = (d(x + c))^\pi(b) \\ = (d(x + c))^\pi(b) + c = b + c = a. \\ (2, ii) \quad a = b + c = a + d + c = a + k \times (d + c) \\ = a + k \times (d + c) + d = a + d = b.$$




Figure 3.: A forest s with a convex set U (in bold) that has three close U -ends (on the left) and five far ones (on the right). The height is $h(s, U) = 2$.

(b) By (G1) $_k$ we have

$$c = c + c \quad \text{implies} \quad ac = a(c + c) = a(k \times c) \\ = a(k \times c) + c = ac + c, \\ c = a(c, c) \quad \text{implies} \quad c = a(c, c) = (a(x, c))^\pi(c) \\ = (a(x, c))^\pi(c) + c = c + c. \quad \square$$

Let us take a look at the following situation (see Figure 3). Let s be a forest and U a set of vertices. We assume that U is convex in the sense that $u \leq_{\text{pf}} v \leq_{\text{pf}} w$ and $u, w \in U$ implies $v \in U$ (where \leq_{pf} denotes the forest order). We call the maximal elements (w.r.t. \leq_{pf}) of U the U -ends. An U -end u is close if $u' \in U$, for all $u' \leq_{\text{pf}} u$. Otherwise, it is far. We would like to know how many of the U -ends are close.

Lemma 3.15. Let $m \geq 0$ and $n \geq 1$, let $s \sim_n^{m+n+2} t$ be two forests, $U \subseteq \text{dom}(s)$ a convex set that is closed under \approx_n^m , and set

$$V := \{v \in \text{dom}(t) \mid u \approx_n^m v \text{ for some } u \in U\}.$$

(a) V is convex and closed under \approx_n^m .

(b) The numbers of ends of U and V are the same, or both numbers are at least n .

- (c) If U has less than n ends, then U is finite if, and only if, V is finite.
- (d) If U is finite and has less than n ends, then U and V have the same numbers of close ends and the same number of far ones.

Proof. (a) If V is not convex, there are vertices $v <_{\text{pf}} v' <_{\text{pf}} v''$ of t with $v, v'' \in V$ and $v' \notin V$. Fix vertices $u <_{\text{pf}} u' <_{\text{pf}} u''$ with $u \approx_n^{m+2} v$, $u' \approx_n^{m+1} v'$, and $u'' \approx_n^m v''$. By definition of V , we have $u, u'' \in U$ and $u' \notin U$. This contradicts the fact that U is convex.

To see that V is closed under \approx_n^m , suppose that $v \in V$ and $v \approx_n^m v'$. By definition of V , there is some $u \in U$ with $u \approx_n^m v$. Hence, $u \approx_n^m v \approx_n^m v'$. As \approx_n^m is transitive, this implies that $v' \in V$.

(b) For a contradiction, suppose that U has $k < n$ ends while V has more than k ends. (By (a), the other case follows by symmetry.) Choose $k + 1$ ends $v_0, \dots, v_k \in V$. Since $s \approx_n^{m+2} t$, there are vertices u_0, \dots, u_k in s with $u_i \approx_n^{m+1} v_i$. By definition of V , we have $u_i \in U$. By assumption, there is some index j such that u_j is not an end. Hence, we can find a vertex $u' >_{\text{pf}} u_j$ with $u' \in U$. Fix a vertex $v' >_{\text{pf}} v_j$ of t with $u' \approx_n^m v'$. Then $v' \in V$ and v_j is not an end. A contradiction.

(c) For a contradiction, suppose that U is finite, but V is not. (The other case follows again by symmetry.) By (b), V has only finitely many ends. Hence, there is some element $v \in V$ such that $v \not\leq_{\text{pf}} v'$ for every end v' of V . Since $s \approx_n^{m+3} t$, we can find a vertex u of s with $u \approx_n^{m+2} v$. This implies that $u \in U$. As U is finite, we can find some end u' of U with $u \leq_{\text{pf}} u'$. Fix some $v' \geq_{\text{pf}} v$ with $u' \approx_n^{m+1} v'$. Then $u' \in U$ implies $v' \in V$. By choice of v , there is some $v'' >_{\text{pf}} v'$ with $v'' \in V$. Choose $u'' >_{\text{pf}} u'$ with $u'' \approx_n^m v''$. By choice of u' , we have $u'' \notin U$. This contradicts the fact that $v'' \in V$.

(d) By (b), we only need to prove that the number of close ends is the same. Let \hat{U} and \hat{V} be the sets of U -ends and V -ends, respectively. We denote by $N(s, U)$ the number of close U -ends and by $F(s, U)$ the set of all proper subforests s' of s that are attached to some vertex v that does not belong to U but where at least one root belongs to U . (A forest s' is a *proper subforest* of s attached at v if s' can be obtained from the subtree $s|_v$ by removing the

root v .) We define the following equivalence relation.

$$\begin{aligned} \langle s, U \rangle \approx_o \langle t, V \rangle & \quad : \text{iff} \quad N(s, U) = N(t, V), \\ \langle s, U \rangle \approx_{i+1} \langle t, V \rangle & \quad : \text{iff} \quad N(s, U) = N(t, V) \text{ and} \\ & \quad \#_\tau(s, U) = \#_\tau(t, V), \\ & \quad \text{for every } \approx_i\text{-class } \tau, \end{aligned}$$

where $\#_\tau(s, U)$ denotes the number of subforests $s' \in F(s, U)$ that belong to the class τ .

We define the U -height of s by

$$h(s, U) := \begin{cases} o & \text{if } F(s, U) = \emptyset \\ 1 + \max \{ h(s', U) \mid s' \in F(s, U) \} & \text{otherwise.} \end{cases}$$

By induction on l , we will prove the following claim:

$$(*) \quad s \sim_n^{m+l+2} t \text{ and } h(s, U) \leq l \text{ implies}$$

$$h(s, U) = h(t, V) \quad \text{and} \quad \langle s, U \rangle \approx_l \langle t, V \rangle.$$

As $h(s, U) \leq |\hat{U}| < n$, it then follows that $\langle s, U \rangle \approx_n \langle t, V \rangle$. In particular, $N(s, U) = N(t, V)$, as desired.

It thus remains to prove $(*)$. First, consider the case where $l = o$. If $h(t, V) > o$, there is some V -end v that is not close. Fix some vertex $v' <_{\text{pf}} v$ with $v' \notin V$. Since $s \sim_n^{m+2} t$, we can find vertices $u' <_{\text{pf}} u$ of s with $u' \approx_n^{m+1} v'$ and $u \approx_n^m v$. By definition of V , it follows that $u' \notin U$ and $u \in U$. As U is finite, we can find some U -end $w \geq_{\text{pf}} u$. But $u' <_{\text{pf}} u \leq_{\text{pf}} w$ implies that w is not close. Hence, $h(s, U) > o$. A contradiction.

For the second part, suppose that $\langle s, U \rangle \not\approx_o \langle t, V \rangle$, that is, $N(s, U) \neq N(t, V)$. By symmetry, we may assume that $m := N(s, U) < N(t, V)$. Pick $m + 1$ distinct close V -ends v_0, \dots, v_m . Since $m + 1 \leq n$ and $s \sim_n^{m+2} t$, there are elements $u_0, \dots, u_m \in \text{dom}(s)$ with $u_i \approx_n^{m+1} v_i$. There must be some index j such that u_j is not a close U -end. As U is closed under \approx_n^m and $u_j \approx_n^m v_j \approx_n^m u$, for some $u \in U$, it follows that $u_j \in U$. Furthermore,

$u_j \approx_n^{m+1} v_j$ and the fact that v_j is a V -end implies that $u' \notin U$, for all $u' \succ_{\text{pt}} u_j$. Thus, u_j is a U -end. But $h(s, U) = 0$ implies that all U -ends of s are close. A contradiction.

For the inductive step, suppose that $s \sim_n^{m+(l+1)+2} t$ holds but we have $h(s, U) \neq h(t, V)$ or $\langle s, U \rangle \not\prec_{l+1} \langle t, V \rangle$. We distinguish several cases.

(i) Suppose that $h(s, U) > h(t, V)$. By definition of h , there is a subforest $s' \in F(s, U)$ with $h(s', U) = h(s, U) - 1$. Then there is some subforest t' of t with $s' \sim_n^{m+l+2} t'$. By inductive hypothesis it follows that

$$h(s, U) = h(s', U) + 1 = h(t', V) + 1 < h(t, V) + 1 \leq h(s, U).$$

A contradiction.

(ii) Suppose that $h(s, U) < h(t, V)$. By definition of h , there is a subforest $t' \in F(t, V)$ with $h(t', V) = h(t, V) - 1$. Fix a subforest s' of s with $s' \sim_n^{m+l+2} t'$. By inductive hypothesis, it follows that

$$h(s, U) > h(s', U) = h(t', V) = h(t, V) - 1 \geq h(s, U).$$

A contradiction.

(iii) Suppose that $N(s, U) \neq N(t, V)$ and there is no \prec_l -class τ with $\#_\tau(s, U) \neq \#_\tau(t, V)$. Then we have $|\hat{U}| - N(s, U) = |\hat{V}| - N(t, V)$. Since $|\hat{U}| = |\hat{V}|$ it follows that $N(s, U) = N(t, V)$. A contradiction.

(iv) Finally, suppose that $\#_\tau(s, U) \neq \#_\tau(t, V)$, for some \prec_l -class τ . By symmetry, we may assume that $m := \#_\tau(s, U) < \#_\tau(t, V)$. We choose $m + 1$ vertices v_0, \dots, v_m of t such that the attached subforests have class τ . Since $s \sim_n^{m+(l+1)+2} t$ and $m + 1 \leq n$, there are vertices u_0, \dots, u_m of s such that $u_i \sim_n^{m+l+2} v_i$, for all $i \leq m$. Let s_i be the subforest of s attached to u_i , and t_i the subforest of t attached to v_i . By inductive hypothesis, it follows that $s_i \prec_l t_i$, for $i \leq m$. Thus, s has at least $m + 1$ different subforest in the class τ . A contradiction. \square

Corollary 3.16. *Let $s \sim_n^{m+n+2} t$ be forests such that, for every $c \in A_{\emptyset}$, the sets*

$$U_c := \{x \in \text{dom}(s) \mid \pi(s|_x) = c\}$$

$$V_c := \{y \in \text{dom}(t) \mid \pi(t|_y) = c\}$$

are convex and closed under \approx_n^m . Then $\pi(s) = \pi(t)$.

Proof. Suppose that $s = s_0 + \dots + s_{l-1}$ and $t = t_0 + \dots + t_{k-1}$, for trees s_i and t_i . It is sufficient to show that, for every $c \in A_{\emptyset}$ such that the number of components s_i with $\pi(s_i) = c$ is different from the number of t_i with $\pi(t_i) = c$, we have

$$\pi(s) = \pi(s) + \pi \times c \quad \text{and} \quad \pi(t) = \pi(t) + \pi \times c.$$

Adding enough terms c to $\pi(s) = \pi(s_0) + \dots + \pi(s_{l-1})$ and $\pi(t) = \pi(t_0) + \dots + \pi(t_{k-1})$ it then follows that $\pi(s) = \pi(t)$.

Hence, fix such an element c . By Lemma 3.15, we obtain one of the following cases.

(i) U and V both have at least n ends. Then they contain an antichain of size n , and we can write s as $r(s'_0, \dots, s'_{n-1})$ with $\pi(s'_i) = c$. Hence, it follows by $(G1)_n$ that

$$\pi(s) = \pi(r)(c, \dots, c) = \pi(r)(c, \dots, c) + \pi \times c = \pi(s) + \pi \times c.$$

For t it follows in the same way that

$$\pi(t) = \pi(t) + \pi \times c.$$

(ii) Both U and V are infinite, but each has less than n ends. Then they contain an infinite chain and we can use Ramsey's Theorem (or the fact that s is regular) to write $\pi(s)$ as $a'e^\omega$ where $ec = c = e^\omega$. By $(G3)$ and $(G1)_n$ it follows that

$$\begin{aligned} \pi(s) &= a'e^\omega = a'(e^\omega + \dots + e^\omega) = a'(c + \dots + c) \\ &= a'(c + \dots + c) + \pi \times c \\ &= \pi(s) + \pi \times c. \end{aligned}$$

For t , we similarly obtain

$$\pi(t) = \pi(t) + \pi \times c.$$

(iii) The last remaining case is where both U and V are finite and they have the same number of close ends. Then the number of indices i with $\pi(s_i) = c$ would be the same as the number of i with $\pi(t_i) = c$, in contradiction to our choice of c . \square

Before presenting the main proof, let us quickly recall how to solve a system of equations using a fixed-point operator. Suppose we are given a system of the form

$$\begin{aligned} x_0 &= r_0(x_0, \dots, x_{n-1}), \\ &\vdots \\ x_{n-1} &= r_{n-1}(x_0, \dots, x_{n-1}), \end{aligned}$$

where $r_0, \dots, r_{n-1} \in \mathbb{F}_\xi A$ and $\xi = \{x_0, \dots, x_{n-1}\}$. Inductively defining

$$s_i(x_0, \dots, x_{i-1}) := (r_i(x_0, \dots, x_i, s_{i+1}, \dots, s_{n-1}))^{\omega_i},$$

we obtain the new system

$$\begin{aligned} x_0 &= s_0, \\ x_1 &= s_1(x_0), \\ &\vdots \\ x_{n-1} &= s_{n-1}(x_0, \dots, x_{n-2}), \end{aligned}$$

which can now be solved by substitution.

Proposition 3.17. *Let \mathcal{A} be an algebra for cEF_n . Then*

$$s \approx_n^{(n+3)(|A_\emptyset|+1)} t \text{ implies } \pi(s) = \pi(t),$$

for all regular trees $s, t \in \mathbb{F}_\emptyset(A_\emptyset \cup A_{\{x\}})$.

Proof. Let m be the number of L-classes above $b := \pi(s)$ (including that of b itself). We will prove by induction on m that

$$s \approx_n^{f(m)} t \text{ implies } \pi(t) = b,$$

where $f(m) := (n+3)(m+1)$. Set

$$S := \{x \in \text{dom}(s) \mid \pi(s|_x) = b\},$$

$$T := \{y \in \text{dom}(t) \mid x \approx_n^{f(m-1)} y \text{ for some } x \in S\}.$$

As t is regular it is the unravelling of some finite graph G . For every $y \in T$, we will prove that $\pi(t|_y) = b$ by induction on the number of strongly connected components of G that are contained in T and that are reachable from y . Hence, fix $y \in T$, let C be the strongly connected component of G containing y , and choose some $x \in S$ with $x \approx_n^{f(m)-1} y$. We distinguish two cases.

(a) Let us begin our induction with the case where C is trivial, i.e., it consists of the single vertex y without self-loop. Then

$$t|_y = a(t_0 + \dots + t_{k-1} + t'_0 + \dots + t'_{q-1})$$

where $a := t(y)$ and the subtrees t_i lie outside of T while the t'_i contain vertices in T . Set $d_i := \pi(t_i)$. By our two inductive hypotheses, we already know that $\pi(t'_i) = b$ and that $b <_L d_i$. Hence,

$$\pi(t|_y) = a(d_0 + \dots + d_{k-1} + q \times b).$$

We have to show that this value is equal to b . Suppose that

$$s|_x = a(s_0 + \dots + s_{l-1} + s'_0 + \dots + s'_{p-1}),$$

where again the trees s_i lie outside of S , while the s'_i contain vertices of S . Setting $c_i := \pi(s_i)$ it follows that

$$\pi(s|_x) = a(c_0 + \dots + c_{l-1} + p \times b).$$

Since $x \in S$, we already know that this value is equal to b . Hence, it remains to show that

$$a(c_0 + \dots + c_{l-1} + p \times b) = a(d_0 + \dots + d_{k-1} + q \times b).$$

For $c \in A_\emptyset$, let U_c be the set of all vertices $u >_{\text{pf}} x$ such that $\pi(s|_u) = c$ and let V_c be the set of vertices $v >_{\text{pf}} y$ with $\pi(t|_v) = c$. As \leq_L is antisymmetric, these sets are convex. Furthermore, by inductive hypothesis on m , they are also closed under $\approx_n^{f(m-1)}$. Since $f(m) - 1 = f(m-1) + n + 2$, it therefore follows by Corollary 3.16 that

$$c_0 + \dots + c_{l-1} = d_0 + \dots + d_{k-1}.$$

If $p = q$, we are done. Hence, we may assume that $p \neq q$. To conclude the proof, we set

$$U := \{ u \in S \mid x <_{\text{pf}} u \} \quad \text{and} \quad V := \{ v \in T \mid y <_{\text{pf}} v \}.$$

If $p > 0$, then $x \approx_n^{f(m)-1} y$ and $U \neq \emptyset$ implies $V \neq \emptyset$. Hence, $q > 0$. In the same way, $q > 0$ implies $p > 0$. Consequently, we have $p, q > 0$. We consider several cases.

(i) If $b + b = b$, then

$$\begin{aligned} a(d_0 + \dots + d_{k-1} + q \times b) &= a(c_0 + \dots + c_{l-1} + q \times b) \\ &= a(c_0 + \dots + c_{l-1} + p \times b) = b, \end{aligned}$$

as desired.

(ii) If U is not a chain, we obtain $b = a'(b, b)$, for some a' . Hence Lemma 3.14 implies that we are in case (i).

(iii) If U contains an infinite chain, we can use Ramsey's Theorem (or the fact that s is regular), to obtain a factorisation $b = e^\omega$, which implies that $b + b = b$ by (G3). Hence, we are in case (i) again.

(iv) If U is a finite chain, then so is V , by Lemma 3.15. Hence, $p = 1 = q$ and we are done.

(b) It remains to consider the case where the component C is not trivial. Then we can factorise

$$t|_y = r(t_0, \dots, t_{k-1}, t'_0, \dots, t'_{q-1}),$$

where $r \in \mathbb{F}A$ is the unravelling of C , the subtrees t_i lie outside of T , while the subtrees t'_i contain vertices in T . Setting $d_i := \pi(t_i)$, it follows by the two inductive hypotheses that $d_i >_L b$ and $\pi(t'_i) = b$. Consequently,

$$\pi(t|_y) = \pi(r)(d_0, \dots, d_{k-1}, b, \dots, b).$$

Let us simplify the term r . Introducing one variable x_v , for every vertex $v \in C$, we can write r as a system of equations

$$x_v = a_v(x_{u_0} + \dots + x_{u_{l-1}} + c_0 + \dots + c_{q-1}), \quad \text{for } v \in C,$$

where u_0, \dots, u_{l-1} are the successors of v that belong to C and c_0, \dots, c_{q-1} are constants from $\{d_0, \dots, d_{k-1}, b\}$ that correspond to successors outside of C . Solving this system of equations in the way we explained above, we obtain a finite term r_0 built up from elements of $A_\emptyset \cup A_{\{x\}}$ using as operations the horizontal product, the vertical product, and the ω -power operation, such that

$$\pi(t|_y) = \pi(r_0)(d_0, \dots, d_{k-1}, b).$$

With the help of the equations (G5)–(G10), we can transform r_0 in several steps (while preserving its product) until it assumes the form

$$\begin{aligned} & [a_0 \cdots a_{j-1}(x + d_0 + \dots + d_{k-1} + b)]^\omega \\ \text{or} \quad & [a_0 \cdots a_{j-1}(x + d_0 + \dots + d_{k-1})]^\omega \end{aligned}$$

where a_0, \dots, a_{j-1} are the labels of the vertices in C .

We distinguish two cases. First suppose that there is no term with value b in the above sum. This means that every subtree attached to C lies entirely outside of the set T . Then $x \approx_n^{f(m)-1} y$ implies that we can factorise $s|_x$ as

$$s|_x = r'(s_0, \dots, s_{l-1})$$

where

- ♦ $\{\pi(s_0), \dots, \pi(s_{l-1})\} = \{d_0, \dots, d_{k-1}\}$,
- ♦ all labels of r' are among a_0, \dots, a_{j-1} ,
- ♦ every vertex of r' has, for every $i < n$, some descendant labelled a_i .

As above we can transform $s|_x$ into

$$[a_0 \cdots a_{j-1}(x + c_0 + \dots + c_{l-1})]^\omega$$

where $c_i := \pi(s_i)$. Since $\{c_0, \dots, c_{l-1}\} = \{d_0, \dots, d_{k-1}\}$ it follows that

$$\begin{aligned} \pi(t|_y) &= (a_0 \cdots a_{j-1}(x + d_0 + \dots + d_{k-1}))^\omega \\ &= (a_0 \cdots a_{j-1}(x + c_0 + \dots + c_{l-1}))^\omega = \pi(s|_x) = b. \end{aligned}$$

It thus remains to consider the case where some term has value b . Using (G7) and (G11) and the fact that $b <_{\mathbb{L}} d_i$, it then follows that

$$\begin{aligned}\pi(t|_y) &= [a_0 \cdots a_{j-1}(x + d_0 + \cdots + d_{k-1} + b)]^\omega \\ &= [a_0 \cdots a_{j-1}(x + b)]^\omega.\end{aligned}$$

For every $i < j$, we fix some $z_i \in S$ with label a_i such that $x <_{\text{pf}} z_i$ and some successor of z_i also belongs to S . Then

$$\pi(s|_{z_i}) = a_i(c_0^i + \cdots + c_{l_i-1}^i + b + \cdots + b),$$

for some $c_0^i, \dots, c_{l_i-1}^i >_{\mathbb{L}} b$. Since

$$\begin{aligned}b = \pi(s|_{z_i}) &= a_i(c_0^i + \cdots + c_{l_i-1}^i + b + \cdots + b) \\ &\leq_{\mathbb{L}} c_0^i + \cdots + c_{l_i-1}^i + b + \cdots + b \leq_{\mathbb{L}} b\end{aligned}$$

it follows by asymmetry of $\leq_{\mathbb{L}}$ that

$$c_0^i + \cdots + c_{l_i-1}^i + b + \cdots + b = b$$

and $a_i(b) = a_i(c_0^i + \cdots + c_{l_i-1}^i + b + \cdots + b) = b$.

Consequently, $a_0 \cdots a_{j-1}b = b$, which implies that $a^\pi b = b$ where $a := a_0 \cdots a_{j-1}$. We claim that $b + b = b$. It then follows that

$$b = a(b) = a(n \times x)(b) = (a(n \times x))^\pi(b),$$

which, by (G12)_n, implies that

$$\begin{aligned}\pi(t|_y) &= [a(x + b)]^\omega = [a(x + a(n \times x)^\pi(b))]^\omega \\ &= n \times a(n \times x)^\pi(b) = n \times b = b,\end{aligned}$$

as desired.

Hence, it remains to prove our claim that $b + b = b$. By our assumption on y and C , there is some vertex $u \in C$ that has some successor $v \notin C$ with

$v \in T$. Since $s|_x \approx_n^{f(m)-1} t|_y$ and $f(m) \geq f(m-1) + n + 1$, there are vertices $x \leq_{\text{pf}} u_0 <_{\text{pf}} \cdots <_{\text{pf}} u_{n-1}$ each of which has some successor $v_i \in S$ with $v_i \not\leq_{\text{pf}} u_{i+1}$. Consequently, we can write

$$\pi(s|_x) = a' a''(b, \dots, b) \quad \text{and} \quad \pi(s|_{u_0}) = a''(b, \dots, b),$$

where $a' \in A_{\{x\}}$ and $a'' \in A_n$. Hence, it follows by (G1)_n that

$$\begin{aligned}b + b &= \pi(s|_{u_0}) + b = a''(b, \dots, b) + b \\ &= a''(b, \dots, b) = \pi(s|_{u_0}) = b.\end{aligned} \quad \square$$

Theorem 3.18. *A regular \mathbb{F} -algebra \mathfrak{A} is an algebra for cEF_n if, and only if, there exists a number $m < \omega$ such that*

$$s \sim_n^m t \quad \text{implies} \quad \pi(s) = \pi(t),$$

for all regular forests $s, t \in \mathbb{F}(A_\emptyset \cup A_{\{x\}})$.

Proof. (\Leftarrow) In each of the equations (G1)_n–(G12)_n, the two terms on both sides are \sim_n^m -equivalent.

(\Rightarrow) By Proposition 3.17, there is some number m such that

$$s \sim_n^m t \quad \text{implies} \quad \pi(s) = \pi(t),$$

for all regular trees $s, t \in \mathbb{F}(A_\emptyset \cup A_{\{x\}})$.

Suppose that $s, t \in \mathbb{F}(A_\emptyset \cup A_{\{x\}})$ are regular forests. We claim that

$$s \sim_n^{m+n+2} t \quad \text{implies} \quad \pi(s) = \pi(t).$$

Suppose that $s = s_0 + \cdots + s_{l-1}$ and $t = t_0 + \cdots + t_{k-1}$, for trees s_i and t_i , and set $c_i := \pi(s_i)$ and $d_i := \pi(t_i)$. As in Part (a) of the proof of Proposition 3.17, we can use Corollary 3.16 to show that $\pi(s) = \pi(t)$. \square

We complete the proof of Theorem 3.2 as follows.

Theorem 3.19. *A regular language $K \subseteq \mathbb{F}_\emptyset \Sigma$ is cEF_n -definable if, and only if, its syntactic algebra $\text{Syn}(K)$ is an algebra for cEF_n .*

Proof. (\Leftarrow) Suppose that $\text{Syn}(K)$ is an algebra for cEF_n . By Theorem 3.18, every language recognised by $\text{Syn}(K)$ is invariant under \sim_n^m , for some m (when considering regular forests only). Consequently, the claim follows by Corollary 3.11.

(\Rightarrow) Suppose that the language K is cEF_n -definable. By Corollary 3.11 it is then \sim_n^m -invariant, for some m . Thus \sim_n^m is contained in the syntactic congruence of K , which means that $\text{syn}_K : \mathbb{F}\Sigma \rightarrow \text{Syn}(K)$ maps \sim_n^m -equivalent forests to the same value. Given forests $s, t \in \mathbb{F}(\text{Syn}_o(K) \cup \text{Syn}_i(K))$ with $s \sim_n^m t$, we can choose forests $s', t' \in \mathbb{F}\Sigma$ with $s' \sim_n^m t'$ and $s = \text{syn}_K(s')$ and $t = \text{syn}_K(t')$. Then

$$s \sim_n^m t \quad \text{implies} \quad \pi(s) = \text{syn}_K(s') = \text{syn}_K(t') = \pi(t).$$

By Theorem 3.18, it follows that $\text{Syn}(K)$ is an algebra for cEF_n . \square

4. Wreath Products

In this section we provide an algebraic analogue to the cascading operation on languages. We start with a bit of useful notation concerning ranked sets and forests.

Definition 4.1. (a) For a set $A \in \mathcal{D}$, we denote by $A^\Delta \in \mathcal{D}^\Xi$ the Ξ -sorted set with $(A^\Delta)_\xi := A^\xi$.

(b) For two sets $A, B \in \mathcal{D}^\Xi$, we set $B^A := \mathcal{D}^\Xi(A, B)$.

(c) For a forest $t \in \mathbb{F}_\xi A$, a vertex $v \in \text{dom}(t)$ of sort ξ , and a variable $x \in \xi$, we denote by $t \downarrow_v^x$ the subforest of t attached to the x -successors of v , i.e.,

$$t|_v = (t(v))((t \downarrow_v^x)_{x \in \xi}).$$

(d) For a forest $t \in \mathbb{F}_\xi A$ and a tuple $\bar{a} \in (A_\emptyset)^\xi$, we denote by $t \Leftarrow \bar{a} \in \mathbb{F}_\emptyset A$ the forest obtained from t by replacing each variable x by the corresponding value a_x , for $x \in \xi$. \lrcorner

First, let us give a simplified algebraic account that does not quite yield an \mathbb{F} -morphism, but that illustrates the main idea of the more complicated construction below.

Definition 4.2. (a) For a forest $t \in \mathbb{F}_\emptyset \Sigma$ and a morphism $\varphi : \mathbb{F}\Sigma \rightarrow \mathfrak{A}$, we define the φ -annotation $t \triangleleft \varphi \in \mathbb{F}_\emptyset(\Sigma \times A_\emptyset^\Delta)$ of t as the forest that adds to each vertex v the φ -images of the subtrees attached to successors of v . Formally,

$$(t \triangleleft \varphi)(v) := \langle t(v), (\varphi(t \downarrow_v^x))_{x \in \xi} \rangle, \quad \text{for } v \in \text{dom}(t) \text{ of arity } \xi.$$

(b) The *cascade composition* of two morphisms $\alpha : \mathbb{F}\Sigma \rightarrow \mathfrak{A}$ and $\beta : \mathbb{F}(\Sigma \times A_\emptyset^\Delta) \rightarrow \mathfrak{B}$ is the function $\alpha \triangleleft \beta : \mathbb{F}_\emptyset \Sigma \rightarrow B_\emptyset$ defined by

$$(\beta \triangleleft \alpha)(t) := \beta(t \triangleleft \alpha), \quad \text{for } t \in \mathbb{F}_\emptyset \Sigma.$$

Thus, to simulate the cascading operation on languages, we can first transform a given input forest $t \in \mathbb{F}_\emptyset \Sigma$ to the forest $t \triangleleft \alpha$ and then compute its image under β . In other words, we can use the cascade composition $\beta \triangleleft \alpha$ to ‘recognise’ a language. But note that these notions are only defined for forests of sort \emptyset . In particular, $\beta \triangleleft \alpha$ is no morphism. To handle higher arities, we need to set up a bit of algebraic machinery.

The problem with generalising the labelling operation $t \triangleleft \alpha$ is that, when computing the images $\alpha(s_i)$ of subforests s_i , we have to deal with the variables appearing in s_i . Each variable indicates a ‘missing’ part of the forest which we have to fill in some way. In order to compute the α -image of the complete subforest we have to know the α -images of these missing parts. For this reason, we will replace in $t \triangleleft \alpha$ every label from A_\emptyset by a function $A^\Delta \rightarrow A_\emptyset$ that maps the missing α -images to the value we are interested in. This is the main idea behind the following definition.

Definition 4.3. The *wreath product* $\mathfrak{A} \circ \mathfrak{B}$ of two \mathbb{F} -algebras \mathfrak{A} and \mathfrak{B} is the \mathbb{F} -algebra with universe

$$C := A \times B^{A_\emptyset^\Delta}$$

and the following product. Let $p : C \rightarrow A$ and $q : C \rightarrow B^{A^\Delta}$ be the respective projections. Given a forest $t \in \mathbb{F}_\xi C$, we introduce two functions λ and ρ where $\lambda(\bar{a})$ computes the product of the (first components of the) subtrees when we substitute the values \bar{a} for the variables, and ρ applies these products to the function stored in the second component of a vertex. Formally, we define $\lambda : A_\emptyset^\xi \rightarrow \mathbb{F}_\emptyset(A_\emptyset^\Delta)$ and $\rho : A_\emptyset^\xi \rightarrow \mathbb{F}B$ by

$$\begin{aligned} \rho(\bar{a})(v) &:= q(t(v))(\lambda(\bar{a})(v)), \\ \lambda(\bar{a})(v) &:= (\pi(s_x))_{x \in \xi}, \quad \text{where } s_x := (\mathbb{F}p(t) \leftarrow \bar{a}) \downarrow_v^x. \end{aligned}$$

Then we set

$$\pi(t) := \langle \pi(\mathbb{F}p(t)), \pi \circ \rho \rangle.$$

Remark. Note that λ is the unique function such that

$$((\mathbb{F}p(t) \leftarrow \bar{a}) \triangleleft \pi)(v) = \langle p(t(v)), \lambda(\bar{a})(v) \rangle.$$

Before proving that $\mathfrak{A} \circ \mathfrak{B}$ really is an \mathbb{F} -algebra, let us see how the wreath product solves our initial problem regarding cascades of morphisms.

Lemma 4.4. *Let Σ be a finite alphabet and \mathfrak{A} and \mathfrak{B} two \mathbb{F} -algebras. There exists a bijection between pairs of morphisms*

$$\alpha : \mathbb{F}\Sigma \rightarrow \mathfrak{A} \quad \text{and} \quad \beta : \mathbb{F}(\Sigma \times A_\emptyset^\Delta) \rightarrow \mathfrak{B}$$

and morphisms $\varphi : \mathbb{F}\Sigma \rightarrow \mathfrak{A} \circ \mathfrak{B}$. This bijection respects the equation

$$p \circ \varphi = \alpha$$

where $p : \mathfrak{A} \circ \mathfrak{B} \rightarrow \mathfrak{A}$ is the projection to the first component.

Proof. Given α and β , we define φ by

$$\varphi(\text{sing}(c)) := \langle \alpha(\text{sing}(c)), \bar{a} \mapsto \beta(\text{sing}(\langle c, \bar{a} \rangle)) \rangle, \quad \text{for } c \in \Sigma.$$

Since $\mathbb{F}\Sigma$ is freely generated by the elements of the form $\text{sing}(c)$ this defines a unique morphism. Conversely, given φ we set

$$\begin{aligned} \alpha(\text{sing}(c)) &:= p(\varphi(\text{sing}(c))) \\ \beta(\text{sing}(\langle c, \bar{a} \rangle)) &:= q(\varphi(\text{sing}(c)))(\bar{a}). \end{aligned}$$

We claim that these two constructions are inverse to each other. To keep the notation simple, we identify $\text{sing}(c)$ with c and drop the sing . For the first direction, suppose that α' and β' are obtained from φ which in turn is obtained from α and β . Then

$$\begin{aligned} \alpha'(c) &= p(\varphi(c)) = \alpha(c), \\ \beta'(c, \bar{a}) &= q(\varphi(c))(\bar{a}) = \beta(c, \bar{a}). \end{aligned}$$

Conversely, suppose that φ' is obtained from α and β which in turn are obtained from φ . Then

$$\begin{aligned} \varphi'(c) &= \langle \alpha(c), \bar{a} \mapsto \beta(c, \bar{a}) \rangle \\ &= \langle p(\varphi(c)), \bar{a} \mapsto q(\varphi(c))(\bar{a}) \rangle \\ &= \langle p(\varphi(c)), q(\varphi(c)) \rangle \\ &= \varphi(c). \end{aligned}$$

□

Proposition 4.5. *If \mathfrak{A} and \mathfrak{B} are \mathbb{F} -algebras, so is $\mathfrak{A} \circ \mathfrak{B}$.*

Proof. Set $\mathfrak{C} := \mathfrak{A} \circ \mathfrak{B}$. For the unit law, let $\langle c, f \rangle \in C_\xi$. Then

$$\pi(\text{sing}(\langle c, f \rangle)) = \langle c, \pi \circ \rho \rangle$$

where

$$\begin{aligned} \mathbb{F}p(t) \leftarrow \bar{a} &= c(\bar{a}), \\ \lambda(\bar{a})(\langle \rangle) &= \bar{a}, \\ \rho(\bar{a}) &= \text{sing}(f(\bar{a})). \end{aligned}$$

Hence, $\pi(\rho(\bar{a})) = \pi(\text{sing}(f(\bar{a}))) = f(\bar{a})$ and $\pi \circ \rho = f$, as desired.

It remains to prove that $\pi(\text{flat}(T)) = \pi(\mathbb{F}\pi(T))$, for $T \in \mathbb{F}_\xi \mathbb{F}C$. Since \mathfrak{A} satisfies the associative law, this equation holds for the first components of the above products and it is sufficient to check the equality of the second components. Let λ_v and ρ_v be the functions used to define $\pi(T(v))$, $\hat{\lambda}$ and $\hat{\rho}$ those defining $\pi(\text{flat}(T))$, and λ_* and ρ_* those for the outer product of $\pi(\mathbb{F}\pi(T))$. Let $R : A_\emptyset^\xi \rightarrow \mathbb{F}_\xi \mathbb{F}B$ be the function defined by

$$R(\bar{a}) \simeq_{\text{sh}} T \quad \text{and} \quad R(\bar{a})(v) := \rho_v(\lambda_*(\bar{a})(v)), \quad \text{for } \bar{a} \in A_\emptyset^\xi.$$

We claim that

$$\hat{\rho} = \text{flat} \circ R \quad \text{and} \quad \mathbb{F}\pi \circ R = \rho_*.$$

It then follows that

$$\begin{aligned} q(\pi(\text{flat}(T))) &= \pi \circ \hat{\rho} = \pi \circ \text{flat} \circ R \\ &= \pi \circ \mathbb{F}\pi \circ R = \pi \circ \rho_* = q(\pi(\mathbb{F}\pi(T))), \end{aligned}$$

as desired.

Hence, it remains to prove the claim. First, note that, for a vertex $v \in \text{dom}(T)$ and a tuple $\bar{a} \in A_\emptyset^\xi$, we have

$$\begin{aligned} \rho_*(\bar{a})(v) &= q(\pi(T(v)))(\lambda_*(\bar{a})(v)) \\ &= \pi(\rho_v(\lambda_*(\bar{a})(v))) = \pi(R(\bar{a})(v)), \end{aligned}$$

which proves the second equation. For the other one, consider two vertices $v \in \text{dom}(T)$ and $x \in \text{dom}(T(v))$, and let x' be the corresponding vertex of $\text{flat}(T)$. Then

$$\begin{aligned} \text{flat}(R(\bar{a}))(x') &= R(\bar{a})(v)(x) \\ &= \rho_v(\lambda_*(\bar{a})(v))(x) \\ &= q(T(v)(x))(\lambda_v(\lambda_*(\bar{a})(v))) \\ &= q(\text{flat}(T)(x'))(\lambda_v(\lambda_*(\bar{a})(v))) \end{aligned}$$

and

$$\hat{\rho}(\bar{a})(x') = q(\text{flat}(T)(x'))(\hat{\lambda}(\bar{a})(x')).$$

Thus, it is sufficient to prove that

$$\lambda_v(\lambda_*(\bar{a})(v))(x) = \hat{\lambda}(\bar{a})(x').$$

Note that

$$\begin{aligned} \lambda_v(\lambda_*(\bar{a})(v))(x) &= (\pi(s_v^z))_z, \\ \lambda_*(\bar{a})(v) &= (\pi(S_*^z))_z, \\ \hat{\lambda}(\bar{a})(x') &= (\pi(\hat{s}^z))_z, \end{aligned}$$

where

$$\begin{aligned} s_v^z &:= (\mathbb{F}p(T(v)) \leftarrow \lambda_*(\bar{a})(v)) \downarrow_x^z = \mathbb{F}p(r_v^z) \leftarrow \lambda_*(\bar{a})(v), \\ S_*^z &:= (\mathbb{F}p(\mathbb{F}\pi(T)) \leftarrow \bar{a}) \downarrow_v^z = \mathbb{F}p(\mathbb{F}\pi(R_*^z)) \leftarrow \bar{a}, \\ \hat{s}^z &:= (\mathbb{F}p(\pi(\text{flat}(T))) \leftarrow \bar{a}) \downarrow_{x'}^z = \mathbb{F}p(\hat{r}^z) \leftarrow \bar{a}, \end{aligned}$$

and

$$r_v^z := T(v) \downarrow_x^z, \quad R_*^z := T \downarrow_v^z, \quad \hat{r}^z := \text{flat}(T) \downarrow_{x'}^z,$$

Note that

$$\hat{r}^z = r_v^z((\text{flat}(R_*^z))_z).$$

Furthermore, it follows that

$$\begin{aligned} \pi(S_*^z) &= \pi(\mathbb{F}p(\mathbb{F}\pi(R_*^z)) \leftarrow \bar{a}) \\ &= \pi(\mathbb{F}\pi(\mathbb{F}\mathbb{F}p(R_*^z) \leftarrow \text{sing}(\bar{a}))) \\ &= \pi(\text{flat}(\mathbb{F}\mathbb{F}p(R_*^z) \leftarrow \text{sing}(\bar{a}))) \\ &= \pi(\text{flat}(\mathbb{F}\mathbb{F}p(R_*^z)) \leftarrow \bar{a}), \\ &= \pi(\mathbb{F}p(\text{flat}(R_*^z)) \leftarrow \bar{a}). \end{aligned}$$

Consequently,

$$\begin{aligned}
\pi(s_v^z) &= \pi(\mathbb{F}p(r_v^z) \Leftarrow \lambda_*(\bar{a})(v)) \\
&= \pi(\mathbb{F}p(r_v^z) \Leftarrow (\pi(\mathbb{F}p(\text{flat}(R_*^y)) \Leftarrow \bar{a}))_y) \\
&= \pi((\mathbb{F}p(r_v^z))(\mathbb{F}p(\text{flat}(R_*^y)) \Leftarrow \bar{a})_y) \\
&= \pi((\mathbb{F}p(r_v^z))(\mathbb{F}p(\text{flat}(R_*^y))_y) \Leftarrow \bar{a}) \\
&= \pi(\mathbb{F}p(r_v^z)((\text{flat}(R_*^y))_y) \Leftarrow \bar{a}) \\
&= \pi(\mathbb{F}p(\hat{r}^z) \Leftarrow \bar{a}) \\
&= \pi(\hat{s}^z),
\end{aligned}$$

as desired. \square

Unfortunately, in general wreath products do not preserve MSO-definable algebras, similar to cartesian products, a wreath product of finitely generated algebras is not necessarily finitely generated. But we can show that every finitely generated subalgebra of a wreath product of MSO-definable algebras is again MSO-definable. For our purposes, this weaker property is sufficient since the images of recognising morphisms are always finitely generated.

Proposition 4.6. *If \mathfrak{A} and \mathfrak{B} are locally MSO-definable \mathbb{F} -algebras, so is $\mathfrak{A} \circ \mathfrak{B}$.*

Proof. Let $C \subseteq A \circ B$ be finite and set

$$D := \{b(\bar{c}) \mid \langle a, b \rangle \in C, \bar{c} \in A_{\emptyset}^A\}.$$

Then D is also finite. By assumption, there exist MSO-formulae φ_a and ψ_b , for $a \in A$ and $b \in B$, such that

$$\begin{aligned}
\pi(t) = a &\quad \text{iff} \quad t \models \varphi_a \quad \text{for } t \in \mathbb{F}p[C], \\
\pi(t) = b &\quad \text{iff} \quad t \models \psi_b \quad \text{for } t \in \mathbb{F}D.
\end{aligned}$$

We have to construct a formula checking that a forest $t \in \mathbb{F}C$ evaluates to a given value.

Let φ'_a be the formula obtained from φ_a by replacing every atomic formula of the form $P_c x$ (where P_c is the predicate checking for the letter c) by the disjunction $\bigvee_{\bar{a}} P_{\langle c, \bar{a} \rangle} x$ (which checks that the first component of the letter is equal to c). Since $p(\pi(t)) = \pi(\mathbb{F}p(t))$ it follows that

$$t \models \varphi'_a \quad \text{iff} \quad p(\pi(t)) = a.$$

For the second component, we proceed as follows. Given an input forest $t \in \mathbb{F}_{\xi}C$, our formula will guess two families $(R_{\bar{a}})_{\bar{a} \in A_{\emptyset}^{\xi}}$ and $(L_{\bar{a}})_{\bar{a} \in A_{\emptyset}^{\xi}}$ of labellings of t where $R_{\bar{a}}$ encodes the forest $\rho(\bar{a})$ and $L_{\bar{a}}$ the forest $\lambda(\bar{a})$. Since $\rho(\bar{a}) \in \mathbb{F}D$, we can then use the formula ψ_b to check whether the labelling $R_{\bar{a}}$ evaluates to b . But first, we have to check that $R_{\bar{a}}$ and $L_{\bar{a}}$ are correct. For $L_{\bar{a}}$, which contains products in \mathfrak{A} , we can use the formulae φ_a . For $R_{\bar{a}}$, note that $\rho(\bar{a})(v)$ is obtained by a local computation from $\lambda(\bar{a})(v)$. Hence, we can check $R_{\bar{a}}$ once we have verified $L_{\bar{a}}$. \square

Let us collect a few more useful algebraic properties of the wreath product. The proofs are straightforward, but rather tedious.

Lemma 4.7. *Let \mathfrak{A} and \mathfrak{B} be \mathbb{F} -algebras.*

- (a) *There exists a canonical embedding $\mathfrak{A} \times \mathfrak{B} \rightarrow \mathfrak{A} \circ \mathfrak{B}$.*
- (b) *The projection $p : \mathfrak{A} \circ \mathfrak{B} \rightarrow \mathfrak{A}$ to the first component is an \mathbb{F} -morphism.*

Proof. (a) Let φ be the function mapping a pair $\langle a, b \rangle$ to $\langle a, f_b \rangle$ where f_b is the constant function with value b . It is easy to check that φ is a morphism.

(b) follows immediately from the definition of the product of $\mathfrak{A} \circ \mathfrak{B}$. \square

Lemma 4.8. *The wreath product is associative.*

Proof. Let \mathfrak{A} , \mathfrak{B} , and \mathfrak{C} be \mathbb{F} -algebras and set

$$\mathfrak{D} := (\mathfrak{A} \circ \mathfrak{B}) \circ \mathfrak{C} \quad \text{and} \quad \mathfrak{E} := \mathfrak{A} \circ (\mathfrak{B} \circ \mathfrak{C}).$$

We claim that \mathfrak{D} and \mathfrak{E} are isomorphic. Note that the domains are

$$D_{\xi} = (A_{\xi} \times B_{\xi}^{(A_{\emptyset})^{\xi}}) \times C_{\xi}^{(A_{\emptyset} \times B_{\emptyset})^{\xi}}$$

$$\text{and } E_\xi = A_\xi \times (B_\xi \times C_\xi^{(B_\emptyset)^\xi})^{(A_\emptyset)^\xi}.$$

There exists a canonical bijection φ between these two sets which maps

$$\langle \langle a, f \rangle, g \rangle \quad \text{to} \quad \langle a, h \rangle \quad \text{where} \quad h(\bar{c}) := \langle f(\bar{c}), \bar{d} \mapsto g(\bar{c}, \bar{d}) \rangle.$$

We claim that this function is an isomorphism. Hence, given a forest $t \in \mathbb{F}D$, we have to show that

$$\varphi(\pi(t)) = \pi(\mathbb{F}\varphi(t)).$$

First, using the projections

$$\begin{aligned} p : D &\rightarrow A \times B^{A_\emptyset^\Delta} & \hat{p} : E &\rightarrow A \\ q : D &\rightarrow C^{(A_\emptyset \times B_\emptyset)^\Delta} & \hat{q} : E &\rightarrow (B \times C^{B_\emptyset^\Delta})^{A_\emptyset^\Delta} \\ p' : A \times B^{A_\emptyset^\Delta} &\rightarrow A & \hat{p}' : B \times C^{B_\emptyset^\Delta} &\rightarrow B \\ q' : A \times B^{A_\emptyset^\Delta} &\rightarrow B^{A_\emptyset^\Delta} & \hat{q}' : B \times C^{B_\emptyset^\Delta} &\rightarrow C^{B_\emptyset^\Delta} \end{aligned}$$

we can recast the definition of φ into the form of three equations:

$$\begin{aligned} p'(p(d)) &= \hat{p}(\varphi(d)), \\ q'(p(d))(\bar{a}) &= \hat{p}'[\hat{q}(\varphi(d))(\bar{a})], \\ q(d)(\bar{a}, \bar{b}) &= \hat{q}'[\hat{q}(\varphi(d))(\bar{a})](\bar{b}), \end{aligned}$$

for $d \in D$, $\bar{a} \in (A_\emptyset)^\xi$, and $\bar{b} \in (B_\emptyset)^\xi$.

We will prove the above equation separately for each of the three components. By Lemma 4.7 (b), we have

$$\begin{aligned} \hat{p} \circ \varphi \circ \pi &= p' \circ p \circ \pi \\ &= \pi \circ \mathbb{F}(p' \circ p) \\ &= \pi \circ \mathbb{F}(\hat{p} \circ \varphi) = \hat{p} \circ \pi \circ \mathbb{F}\varphi. \end{aligned}$$

Thus, it remains to prove that

$$\hat{q} \circ \varphi \circ \pi = \hat{q} \circ \pi \circ \mathbb{F}\varphi.$$

Again we will prove this equation separately for the two components.

Let $\rho, \lambda, \hat{\rho}, \hat{\lambda}$, and ρ', λ' be the functions used to define the products of, respectively, $\mathbb{F}p(t)$, $\mathbb{F}\varphi(t)$, and t . Then

$$\begin{aligned} q'(\pi(\mathbb{F}p(t))) &= \pi \circ \rho, \\ \hat{q}(\pi(\mathbb{F}\varphi(t))) &= \pi \circ \hat{\rho}, \\ q(\pi(t)) &= \pi \circ \rho'. \end{aligned}$$

Further, note that we have

$$\begin{aligned} \mathbb{F}\hat{p}(\mathbb{F}\varphi(t)) &\Leftarrow \bar{a} = \mathbb{F}p'(\mathbb{F}p(t)) \Leftarrow \bar{a}, \\ \mathbb{F}p'(\mathbb{F}p(t)) &\Leftarrow \bar{a} = \mathbb{F}\hat{p}(\mathbb{F}\varphi(t)) \Leftarrow \bar{a}, \end{aligned}$$

which implies that

$$\hat{\lambda}(\bar{a}) = \lambda(\bar{a}) \quad \text{and} \quad p'(\lambda'(\bar{a}\bar{b})) = \hat{\lambda}(\bar{a}).$$

Consequently, we have

$$\begin{aligned} (*_1) \quad \rho(\bar{a})(v) &= q'(p(t(v))) (\lambda(\bar{a})(v)) \\ &= q'(p(t(v))) (\hat{\lambda}(\bar{a})(v)) \\ &= \hat{p}'[\hat{q}(\varphi(t(v))) (\hat{\lambda}(\bar{a})(v))] = \hat{p}'(\hat{\rho}(\bar{a})(v)), \\ (*_2) \quad \rho'(\bar{a}\bar{b})(v) &= q(t(v)) (\lambda'(\bar{a}\bar{b})(v)) \\ &= \hat{q}'[\hat{q}(\varphi(t(v))) (p'(\lambda'(\bar{a}\bar{b})(v)))] (q'(\lambda'(\bar{a}\bar{b})(v))) \\ &= \hat{q}'[\hat{q}(\varphi(t(v))) (\hat{\lambda}(\bar{a})(v))] (q'(\lambda'(\bar{a}\bar{b})(v))) \\ &= \hat{q}'[\hat{\rho}(\bar{a})(v)] (q'(\lambda'(\bar{a}\bar{b})(v))). \end{aligned}$$

We further claim that

$$(*_3) \pi(\rho'(\bar{a}\bar{b})) = \hat{q}'(\pi(\hat{\rho}(\bar{a}))) (\bar{b}).$$

It then follows by $(*_1)$ and $(*_3)$ that

$$\begin{aligned} \hat{p}'[\hat{q}(\varphi(\pi(t)))(\bar{a})] &= q'(\pi(\pi(t)))(\bar{a}) \\ &= q'(\pi(\mathbb{F}p(t)))(\bar{a}) \\ &= \pi(\rho(\bar{a})) \\ &= \pi(\mathbb{F}\hat{p}'(\hat{\rho}(\bar{a}))) \\ &= \hat{p}'(\pi(\hat{\rho}(\bar{a}))) \\ &= \hat{p}'[\hat{q}(\pi(\mathbb{F}\varphi(t)))(\bar{a})], \\ \hat{q}'[\hat{q}(\varphi(\pi(t)))(\bar{a})](\bar{b}) &= q(\pi(t))(\bar{a}, \bar{b}) \\ &= \pi(\rho'(\bar{a}\bar{b})) \\ &= \hat{q}'(\pi(\hat{\rho}(\bar{a}))) (\bar{b}) \\ &= \hat{q}'[\hat{q}(\pi(\mathbb{F}\varphi(t)))(\bar{a})](\bar{b}), \end{aligned}$$

which concludes the proof.

Thus, it remains to prove the above claim. Suppose that

$$\lambda'(\bar{a}\bar{b})(v) = (\pi(s'_x))_x \quad \text{and} \quad \hat{\lambda}(\bar{a})(v) = (\pi(\hat{s}_x))_x,$$

and let $\hat{\rho}_*$, $\hat{\lambda}_*$ and $\bar{\rho}_v^x$, $\bar{\lambda}_v^x$ be the functions used to define the products $\pi(\hat{\rho}(\bar{a}))$ and $\pi(\mathbb{F}p(t) \leftarrow \bar{a}\bar{b}) \downarrow_v^x$, respectively. By definition, we have

$$\begin{aligned} \pi(\hat{\rho}_*(\bar{b})) &= \hat{q}'(\pi(\hat{\rho}(\bar{a}))) (\bar{b}), \\ \pi(\bar{\rho}_v^x(\langle \rangle)) &= q'(\pi(\mathbb{F}p(t) \leftarrow \bar{a}\bar{b}) \downarrow_v^x)(\langle \rangle). \end{aligned}$$

Furthermore, $p' \circ p = \hat{p} \circ \varphi$ implies that

$$\begin{aligned} \bar{\lambda}_v^x(\langle \rangle)(u) &= (\pi((\mathbb{F}p'(\mathbb{F}p(t) \leftarrow \bar{a}\bar{b}) \downarrow_v^x) \leftarrow \langle \rangle) \downarrow_u^y))_y \\ &= (\pi((\mathbb{F}p' \mathbb{F}p(t) \leftarrow \bar{a}) \downarrow_v^x \downarrow_u^y))_y \\ &= (\pi((\mathbb{F}\hat{p}'\varphi(t) \leftarrow \bar{a}) \downarrow_{v \times u}^y))_y \\ &= \hat{\lambda}(\bar{a}) \downarrow_v^x(u). \end{aligned}$$

Below we will show that

$$\bar{\rho}_v^x(\langle \rangle)(u) = \hat{p}'[(\hat{\rho}(\bar{a}) \leftarrow \langle \bar{b}(\langle \rangle), \bar{c} \rangle) \downarrow_v^x(u)].$$

Then it follows that

$$\begin{aligned} q'(\lambda'(\bar{a}\bar{b})(v)) &= (q'(\pi(\mathbb{F}p(t) \leftarrow \bar{a}\bar{b}) \downarrow_v^x)(\langle \rangle))_x \\ &= (\pi(\bar{\rho}_v^x(\langle \rangle)))_x \\ &= (\pi(\mathbb{F}\hat{p}'(\hat{\rho}(\bar{a})) \downarrow_v^x))_x \\ &= \hat{\lambda}_*(\bar{b})(v), \end{aligned}$$

which implies that

$$\begin{aligned} \rho'(\bar{a}\bar{b})(v) &= \hat{q}'[\hat{\rho}(\bar{a})(v)](q'(\lambda'(\bar{a}\bar{b})(v))) \\ &= \hat{q}'[\hat{\rho}(\bar{a})(v)](\hat{\lambda}_*(\bar{b})(v)) \\ &= \hat{\rho}_*(\bar{b})(v). \end{aligned}$$

Consequently, we have

$$\hat{q}'(\pi(\hat{\rho}(\bar{a}))) (\bar{b}) = \pi(\hat{\rho}_*(\bar{b})) = \pi(\rho'(\bar{a}\bar{b})),$$

as desired.

Hence, it remains to prove the above equation. Note that it is sufficient to show that

$$\begin{aligned} \hat{q}(\varphi(t \leftarrow \langle \bar{a}\bar{b}, \bar{c}' \rangle) \downarrow_v^x(u)) &(\hat{\lambda}(\bar{a}) \downarrow_v^x(u)) \\ &= (\hat{\rho}(\bar{a}) \leftarrow \langle \bar{b}(\langle \rangle), \bar{c} \rangle) \downarrow_v^x(u), \end{aligned}$$

provided that $\varphi(\langle \bar{a}\bar{b}, \bar{c}' \rangle) = \langle \bar{a}, \langle \rangle \mapsto \langle \bar{b}(\langle \rangle), \bar{c} \rangle \rangle$, since it then follows that

$$\begin{aligned} \bar{\rho}_v^x(\langle \rangle)(u) &= q(\mathbb{F}p(t) \leftarrow \bar{a}\bar{b}) \downarrow_v^x(u) (\bar{\lambda}_v^x(\langle \rangle)(u)) \\ &= q(p(t \leftarrow \langle \bar{a}\bar{b}, \bar{c} \rangle) \downarrow_v^x(u)) (\hat{\lambda}_v^x(\bar{a}) \downarrow_v^x(u)) \\ &= \hat{p}'[\hat{q}(\varphi(t \leftarrow \langle \bar{a}\bar{b}, \bar{c} \rangle) \downarrow_v^x(u)) (\hat{\lambda}_v^x(\bar{a}) \downarrow_v^x(u))] \\ &= \hat{p}'[(\hat{\rho}(\bar{a}) \leftarrow \langle \bar{b}(\langle \rangle), \bar{c} \rangle) \downarrow_v^x(u)], \end{aligned}$$

as desired. First, consider a vertex u such that $t(vxu) = z$ is a variable. Then

$$\begin{aligned} & \hat{q}(\varphi((t \leftarrow \langle \bar{a}\bar{b}, \bar{c}' \rangle) \downarrow_v^x(u))) (\hat{\lambda}(\bar{a}) \downarrow_v^x(u)) \\ &= \hat{q}(\varphi(\langle a_z b_z, c'_z \rangle)(\langle \rangle)) \\ &= \langle b_z(\langle \rangle), c_z \rangle \\ &= (\hat{p}(\bar{a}) \leftarrow \langle \bar{b}(\langle \rangle), \bar{c} \rangle) \downarrow_v^x(u). \end{aligned}$$

If vxu is not labelled by a variable, we have

$$\begin{aligned} & \hat{q}(\varphi((t \leftarrow \langle \bar{a}\bar{b}, \bar{c}' \rangle) \downarrow_v^x(u))) (\hat{\lambda}(\bar{a}) \downarrow_v^x(u)) \\ &= \hat{q}(\varphi(t(vxu))) (\hat{\lambda}(\bar{a})(vxu)) \\ &= \hat{p}(\bar{a})(vxu) \\ &= (\hat{p}(\bar{a}) \leftarrow \langle \bar{b}(\langle \rangle), \bar{c} \rangle) \downarrow_v^x(u). \quad \square \end{aligned}$$

Finally, let us show that wreath products provide the desired algebraic analogue to the cascading operation.

Proposition 4.9. *Let \mathcal{K} be a family of forest languages and \mathcal{C} a class of \mathbb{F} -algebras such that a language L belongs to \mathcal{K} if, and only if, it is recognised by a morphism into some algebra in \mathcal{C} . Then*

$$L \in \text{Casc}(\mathcal{K}) \quad \text{iff} \quad L \text{ is recognised by an iterated wreath product of algebras in } \mathcal{C}.$$

Proof. (\Rightarrow) We prove the claim by induction on the construction of a language in $\text{Casc}(\mathcal{K})$. By assumption, every language in \mathcal{K} is recognised by such an algebra. As the recognisable languages are closed under finite boolean operations, it therefore remains to consider closure under the cascading operation. Hence, suppose that $L, \bar{K} \in \text{Casc}(\mathcal{K})$. By inductive hypothesis, we can find iterated wreath products \mathfrak{A} and \mathfrak{B} such that L is recognised by a morphism to \mathfrak{A} and \bar{K} by one to \mathfrak{B}_i . It follows that $L[\bar{K}]$ is recognised by a morphism

$$\varphi : \mathbb{F}\Sigma \rightarrow (\mathfrak{B}_0 \times \cdots \times \mathfrak{B}_{n-1}) \circ \mathfrak{A}.$$

We can combine φ with the canonical embedding of $\mathfrak{B}_0 \times \cdots \times \mathfrak{B}_{n-1}$ into $\mathfrak{B}_0 \circ \cdots \circ \mathfrak{B}_{n-1}$ to obtain a morphism

$$\mathbb{F}\Sigma \rightarrow \mathfrak{B}_0 \circ \cdots \circ \mathfrak{B}_{n-1} \circ \mathfrak{A}$$

recognising $L[\bar{K}]$.

(\Leftarrow) We prove the claim by induction on the length of the iteration. Hence, suppose that $L = \varphi^{-1}[P]$ where $\varphi : \mathbb{F}\Sigma \rightarrow \mathfrak{A} \circ \mathfrak{B}$, $P \subseteq A_\emptyset \times B_\emptyset$, $\mathfrak{B} \in \mathcal{C}$, and \mathfrak{A} is an iterated wreath product. Let p and q be the two projections and set

$$K_a := \{s \mid p(\varphi(s)) = a\}, \quad \text{for } a \in A.$$

By inductive hypothesis, we know that $K_a \in \text{Casc}(\mathcal{K})$. By Lemma 4.4, there exists a morphism $\zeta : \mathbb{F}(\Sigma \times A_\emptyset) \rightarrow \mathfrak{B}$ such that

$$q(\varphi(s)) = b \quad \text{iff} \quad \zeta(s \triangleleft (p \circ \varphi)) = b, \quad \text{for } s \in \mathbb{F}_\emptyset \Sigma \text{ and } b \in B_\emptyset.$$

Each language $M_b := \zeta^{-1}(b)$ is recognised by \mathfrak{B} and, hence, belongs to \mathcal{K} . Furthermore, for $s \in \mathbb{F}_\emptyset \Sigma$, we have

$$\begin{aligned} s \in L & \quad \text{iff} \quad \varphi(s) \in P \\ & \quad \text{iff} \quad \text{there is some } \langle a, b \rangle \in P \text{ such that} \\ & \quad \quad p(\varphi(s)) = a \quad \text{and} \quad q(\varphi(s)) = b \\ & \quad \text{iff} \quad \text{there is some } \langle a, b \rangle \in P \text{ such that} \\ & \quad \quad s \in K_a \quad \text{and} \quad \zeta(s \triangleleft (p \circ \varphi)) = b \\ & \quad \text{iff} \quad \text{there is some } \langle a, b \rangle \in P \text{ such that} \\ & \quad \quad s \in K_a \quad \text{and} \quad s \triangleleft (p \circ \varphi) \in M_b \\ & \quad \text{iff} \quad s \in \bigcup_{\langle a, b \rangle \in P} [K_a \cap M_b[\bar{K}]]. \end{aligned}$$

As $\text{Casc}(\mathcal{K})$ is closed under boolean operations, it follows that the latter language belongs to $\text{Casc}(\mathcal{K})$. \square

Open Question. *Does there exist an analogue to the Krohn-Rhodes Theorem for \mathbb{F} -algebras?*

5. Distributive Algebras

In this and the next section article we will use Proposition 4.9 to derive characterisations of the various logics defined in Section 1. We start with the bisimulation-invariant logics $\text{TL}[\mathcal{K}]$.

Definition 5.1. Let \mathfrak{A} be an \mathbb{F} -algebra.

(a) For a path w (starting at some root) of a forest $s \in \mathbb{F}A_{\{x\}}$, we denote by $\text{pr}_w \in \mathbb{F}A_{\{x\}}$ the following path-like forest. If w is infinite, pr_w is simply the restriction of s to the domain

$$\text{dom}(\text{pr}_w) := \{v \in \text{dom}(s) \mid v \text{ lies on } w\}.$$

If w is finite, we take the above restriction and add an additional vertex at the end that is labelled by the variable x . We will not strictly distinguish between the forest $\text{pr}_w \in \mathbb{F}A_{\{x\}}$ and the corresponding word in $A_{\{x\}}^\infty$.

(b) \mathfrak{A} is *distributive* if, for all forests $s, t \in \mathbb{F}A$,

$$\{\pi(\text{pr}_w) \mid w \text{ a path of } s\} = \{\pi(\text{pr}_w) \mid w \text{ a path of } t\}$$

implies $\pi(s) = \pi(t)$.

(b) \mathfrak{A} is *+distributive* if, for all forests $s, t \in \mathbb{F}A$,

$$\{\pi(\text{pr}_w) \mid w \text{ a finite path of } s\} = \{\pi(\text{pr}_w) \mid w \text{ a finite path of } t\}$$

implies $\pi(s) = \pi(t)$. \lrcorner

Remark. Note that the set $\{\pi(\text{pr}_w) \mid w \text{ a path of } s\}$ may contain both elements of arity 0 and of arity 1. \lrcorner

Before using distributive algebras to characterise temporal logics, let us mention two properties that are essential for applications.

Proposition 5.2. Every distributive \mathbb{F} -algebra is locally MSO-definable.

Proof. When deciding whether a product $\pi(s)$ evaluates to a given element a , we only have to see which branches are realised in s . As automata can evaluate products in ω -semigroups, there exists an automaton performing this check. \square

Proposition 5.3. Distributivity of an MSO-definable \mathbb{F} -algebra \mathfrak{A} is decidable.

Proof. By definition, an \mathbb{F} -algebra \mathfrak{A} is not distributive if, and only if, there exist forests $s, t \in \mathbb{F}A$ such that $\pi(s) \neq \pi(t)$, but

$$\{\pi(\text{pr}_w) \mid w \text{ a path of } s\} = \{\pi(\text{pr}_w) \mid w \text{ a path of } t\}.$$

This is the case if, and only if, there exists a forest r with two roots u and v such that the subforests s and t attached to u and v , respectively, satisfy the above condition. As the product of \mathfrak{A} is MSO-definable, we can construct an MSO-formula that checks whether or not a given a forest r is of this form. Since satisfiability of MSO over forests is decidable, the result follows. \square

Theorem 5.4. Let \mathcal{S} be a pseudo-variety of finite ω -semigroups and \mathcal{K} the family of ∞ -word languages recognised by them.

(a) A language $L \subseteq \mathbb{F}\Sigma$ can be defined by a finite boolean combination of formulae of the form EK with $K \in \mathcal{K}$ if, and only if, it is recognised by an MSO-definable \mathbb{F} -algebra \mathfrak{A} that is distributive and whose vertical ω -semigroup $\langle A_{\{x\}}, A_\emptyset, \cdot \rangle$ belongs to \mathcal{S} .

(b) A language $L \subseteq \mathbb{F}\Sigma$ can be defined by a finite boolean combination of formulae of the form EK with $K \in \mathcal{K}$ and $K \subseteq \Sigma^*$ if, and only if, it is recognised by an MSO-definable \mathbb{F} -algebra \mathfrak{A} that is +distributive and whose vertical ω -semigroup $\langle A_{\{x\}}, A_\emptyset, \cdot \rangle$ belongs to \mathcal{S} .

Proof. (a) (\Rightarrow) Let $\varphi : \mathbb{F}\Sigma \rightarrow \mathfrak{A}$ recognise L where \mathfrak{A} is as above. For every $a \in A_\emptyset \cup A_{\{x\}}$, we set

$$H_a := \{\{\pi(\text{pr}_w) \mid w \text{ path in } t\} \mid t \in \pi^{-1}(a)\},$$

$$K_a := \{w \in \Sigma^\infty \mid \pi(\mathbb{F}\varphi(w)) = a\}$$

(where we identify words $w \in \Sigma^\infty$ with forest that have a single branch). It follows that

$$\begin{aligned} \pi(t) = a & \quad \text{iff} \quad \{\pi(\text{pr}_w) \mid w \text{ path in } t\} \in H_a \\ & \quad \text{iff} \quad t \models \bigvee_{I \in H_a} \left[\bigwedge_{c \in I} EK_c \wedge \bigwedge_{c \notin I} \neg EK_c \right]. \end{aligned}$$

(\Leftarrow) Suppose that L is definable by a formula of this form. Then L is regular and recognised by a morphism $\text{syn}_L : \mathbb{F}\Sigma \rightarrow \text{Syn}(L)$ to its syntactic algebra. It is sufficient to prove that $\text{Syn}(L)$ is distributive and that its vertical ω -semigroup belongs to \mathcal{S} .

For distributivity, consider two forests $s, t \in \mathbb{F}S_L$ with

$$\{ \pi(\text{pr}_w) \mid w \text{ a path of } s \} = \{ \pi(\text{pr}_w) \mid w \text{ a path of } t \}.$$

As $\text{syn}_L : \mathbb{F}\Sigma \rightarrow \text{Syn}(L)$ is surjective, it has a right inverse $\iota : \text{Syn}(L) \rightarrow \mathbb{F}\Sigma$. We set

$$s' := \text{flat}(\mathbb{F}\iota(s)) \quad \text{and} \quad t' := \text{flat}(\mathbb{F}\iota(t)).$$

Note that, for every context p , the forest $p[s']$ and $p[t']$ cannot be distinguished by a formula of the above form. This implies that

$$p[s'] \in L \iff p[t'] \in L.$$

Thus $s' \sim_L t'$, which implies that $\text{syn}_L(s') = \text{syn}_L(t')$. Therefore, we have

$$\begin{aligned} \pi(s) &= \pi(\text{syn}_L(\iota(s))) \\ &= \text{syn}_L(\mathbb{F}\text{flat}(\iota(s))) \\ &= \text{syn}_L(s') \\ &= \text{syn}_L(t') \\ &= \text{syn}_L(\mathbb{F}\text{flat}(\iota(t))) \\ &= \pi(\text{syn}_L(\iota(t))) \\ &= \pi(t). \end{aligned}$$

To show that the vertical ω -semigroup belongs to \mathcal{S} consider two ∞ -words u, v of elements of $\langle A_{\{x\}}, A_{\emptyset}, \cdot \rangle$ that should be equal according to the equations for \mathcal{S} . Analogously to above, we can show that the words

$$u' := \text{flat}(\mathbb{F}\iota(u)) \quad \text{and} \quad v' := \text{flat}(\mathbb{F}\iota(v))$$

are syntactically equivalent and, thus, $\pi(u) = \pi(u') = \pi(v') = \pi(v)$.

(b) The above proof goes through with minor modifications. For (\Rightarrow), we change the definition of H_a to only consider finite paths, and the definition of K_a to only contain finite words. For (\Leftarrow), note that the assumption that s and t have the same set of finite paths is sufficient to show that $\langle s', t' \rangle \in \ker \text{syn}_L$. \square

With this theorem we are able to give characterisations for various logics of the form $\text{TL}[\mathcal{S}^\infty]$ and $\text{TL}[\mathcal{S}^+]$. To do so we will make use of the following types of \mathbb{F} -algebras.

We denote by $\mathbb{U}_{(1)}$ the \mathbb{F} -algebra with two elements \circ_ξ and ι_ξ , for each sort ξ , where the product is simply the maximum function (using the order $\circ_\xi < \iota_\xi$):

$$\pi(s) := \max \{ s(v) \mid v \in \text{dom}(s) \}.$$

$\mathbb{U}_{(2)}^*$ is the \mathbb{F} -algebra with domains

$$(U_{(2)}^*)_\xi := \{ \top_\xi \} \cup \wp(\xi), \quad \text{for } \xi \in \Xi.$$

To define the product of a forest $s \in \mathbb{F}U_{(2)}^*$, let us call a vertex $v \in \text{dom}(s)$ *reachable* if, for each $u < v$ and every variable x with $ux \leq v$, we have $x \in s(u)$ (which in particular implies that $s(u) \neq \top$). If some vertex v with label \top is reachable, we set $\pi(s) := \top$. Otherwise,

$$\pi(s) := \{ x \in \xi \mid \text{a vertex with label } x \text{ is reachable} \}.$$

We also define a variant of this algebra called $\mathbb{U}_{(2)}^\omega$. It has the same domains, but the product is slightly modified. We set $\pi(s) := \top$ if some vertex with label \top is reachable, or if there exists an infinite branch all vertices of which are reachable. Otherwise, we define $\pi(s)$ as above as the set of reachable variables.

Finally, we call an algebra *aperiodic* if its vertical ω -semigroup is aperiodic.

Theorem 5.5. *A forest language is definable in one of the logics below if, and only if, it is recognised by an iterated wreath product of algebras from the following table.*

logic	algebras
EF	$\mathbb{U}_{(1)}$
wCTL	$\mathbb{U}_{(2)}^*$
CTL	$\mathbb{U}_{(2)}^*, \mathbb{U}_{(2)}^\omega$
wCTL*	+distributive aperiodic algebras
CTL*	distributive aperiodic algebras
PDL	+distributive algebras

Proof. Note that the only forest languages recognised by a morphism $\varphi : \mathbb{F}\Sigma \rightarrow \mathbb{U}_{(1)}$ are the empty language \emptyset , the full language $\mathbb{F}\Sigma$, the language

$$\{s \in \mathbb{F}\Sigma \mid \text{there is some } v \in \text{dom}(s) \text{ with } s(v) \in \varphi^{-1}(1_\emptyset)\},$$

and its complement. These languages are exactly the languages of the form

$$\{s \in \mathbb{F}\Sigma \mid \text{there is some path with labelling in } \Sigma^* C\}, \quad \text{for } C \subseteq \Sigma,$$

which, by Propositions 1.3 and 1.5, coincide with the EF-definable ones.

Similarly, a morphism $\varphi : \mathbb{F}\Sigma \rightarrow \mathbb{U}_{(2)}^*$ recognises the language of all forests with some path labelled in C^*B where

$$B := \Sigma \cap \varphi^{-1}(\top) \quad \text{and} \quad C := \Sigma \cap \varphi^{-1}(\{x\})$$

(where we consider $\{x\} \in (U_{(2)})_{\{x\}}$ to be an element of arity 1), and the complements of such languages. To see this, note that $(U_{(2)})_{\{x\}}$ contains three elements \emptyset , $\{x\}$, and \top . A path s evaluates to \top if it is of the form $\{x\}^* \top \emptyset^*$, and it evaluates to $\{x\}$ if it is of the form $\{x\}^+$. In all other cases, the product is \emptyset . For higher arities, we need to keep track which variables are reachable by a path labelled $\{x\}^+$ (or more precisely, labelled $\xi_0 \dots \xi_n$ where each ξ_i contains the variable leading to the successor on the path).

For morphisms $\varphi : \mathbb{F}\Sigma \rightarrow \mathbb{U}_{(2)}^\omega$, we can also recognise the language of all forests with an infinite path labelled in C^ω . Now the claim for wCTL and CTL follow as above.

Finally, the last three statements follow by Theorem 5.4. \square

6. Path Algebras

Next, we turn to the logics $\text{cTL}[\mathcal{K}]$, which we can characterise in terms of cascade products of the following class of languages.

Definition 6.1. Let \mathcal{S} be a pseudo-variety of finite ω -semigroups and \mathcal{K} the family of ∞ -word languages recognised by them. A forest language $L \subseteq \mathbb{F}\Sigma$ is an \mathcal{S} -path language if it is a finite boolean combination of languages of the form

$$B_{k,K} := \{t \in \mathbb{F}\Sigma \mid t \text{ has at least } k \text{ paths whose label belongs to } K\},$$

where $k < \omega$ and $K \in \mathcal{K}$. We call languages of the form $B_{k,K}$ a *basic \mathcal{S} -path language*.

If we allow only languages $K \in \mathcal{K}$ with $K \subseteq \Sigma^*$, we speak of an \mathcal{S}^+ -path language.

To introduce the corresponding family of \mathbb{F} -algebras, we need a bit of notation. Recall that a semigroup is aperiodic if it does not contain a group. This is equivalent to the condition that $a^\pi \cdot a = a^\pi$ for all a . Finally, for a semigroup \mathfrak{S} we denote by \mathfrak{S}^1 the extension of \mathfrak{S} by a new element 1 which acts as neutral element, i.e., $a \cdot 1 = a = 1 \cdot a$, for all $a \in \mathfrak{S}$. In case \mathfrak{S} is the vertical semigroup of an \mathbb{F} -algebra, we denote this element by x .

The following is an analogue of one of the Green's relations for \mathbb{F} -algebras.

Definition 6.2. Let \mathfrak{A} be an \mathbb{F} -algebra.

(a) For elements $a, b \in (A_{\{x\}})^1$, we define

$$a \sqsubseteq_{\mathbb{R}} a' \quad \text{iff} \quad a = a'b, \quad \text{for some } b \in (A_{\{x\}})^1.$$

(b) Let $\tilde{\Psi} \subseteq A_{\{x\}}^*$ be the set of all finite words $\langle a_0, \dots, a_{n-1} \rangle$ satisfying

$$a_0 \supseteq_R a_0 a_1 \supseteq_R \dots \supseteq_R a_0 a_1 \dots a_{n-1},$$

and set $\Psi_\xi := \tilde{\Psi} \times A_\xi$, for $\xi \in \Xi$. The *character* $\chi(w)$ of a word $w \in A_{\{x\}}^\infty$ is the word

$$\chi(w) = \langle a_0, \dots, a_{n-1} \rangle \in \Psi$$

obtained from a certain factorisation $w = w_0 \dots w_{n-1}$ of w by setting $a_i := \pi(w_i)$. The factorisation w_0, \dots, w_{n-1} is constructed inductively. If we have already found w_0, \dots, w_{i-1} , we choose the shortest factor w_i of w such that the product $a_i = \pi(w_i)$ satisfies

$$a_0 \dots a_{i-1} a_i \sqsubseteq_R a_0 \dots a_{i-1}.$$

If no such factor exists, we take for w_i the rest of the word.

The *relative character* $\chi(w/\sigma)$ of $w \in A_{\{x\}}^\infty$ with respect to $\sigma \in \Psi$ is defined as follows. Suppose that

$$\chi(\sigma w) = \langle a_0, \dots, a_{n-1}, c \rangle.$$

Then $\sigma = \langle a_0, \dots, a_{k-1}, b \rangle$ and $a_k = bb'$, for some $b, b' \in A_{\{x\}}$, where b' is the product of a prefix of w . We set

$$\chi(w/\sigma) = \langle b', a_{k+1}, \dots, a_{n-1}, c \rangle.$$

Definition 6.3. Let \mathfrak{A} be an \mathbb{F} -algebra and $s \in \mathbb{F}_{\emptyset} A_{\{x\}}$ a forest.

(a) We write sf_v for the subforest of s rooted at the successors of a vertex v , i.e., the forest obtained from the subtree rooted at v by removing the vertex v itself.

(b) Recall the definition of the path pr_w , for paths w in s . For $\sigma \in \Psi$, we define the *reduced path* to w as the forest

$$\tilde{\text{pr}}_{w,\sigma} := \chi(\text{pr}_w/\sigma) \in \mathbb{F}_{\{x\}} A_{\{x\}}.$$

(c) For a given constant $N < \omega$, we denote by

$$\Pi_N(s) : X \rightarrow [N+1]$$

the function mapping a character $\sigma \in \Psi$ to the number of paths w of s such that $\chi(\text{pr}_w) = \sigma$, provided this number is at most N . Otherwise, we set $\Pi_N(s)(\sigma) := N$. We denote by $\Pi_N^+(s)$ the variant that only counts the number of finite paths. For a subset $H \subseteq \text{dom}(s)$, we write $\Pi_N(s/H)$ and $\Pi_N^+(s/H)$ for the corresponding functions that only count paths containing some vertex in H .

After these preparations we can introduce the class of algebras corresponding to path languages.

Definition 6.4. (a) An \mathbb{F} -algebra \mathfrak{A} is an ω -path algebra if it is finitary, its horizontal semigroup $\langle A_\emptyset, + \rangle$ is commutative and aperiodic, and there exists a constant $N < \omega$ such that

$$\Pi_N(s) = \Pi_N(t) \quad \text{implies} \quad \pi(s) = \pi(t), \quad \text{for } s, t \in \mathbb{F}_{\emptyset} A_{\{x\}}.$$

(b) \mathfrak{A} is a $+$ -path algebra if it is an ω -path algebra satisfying the stronger condition that

$$\Pi_N^+(s) = \Pi_N^+(t) \quad \text{implies} \quad \pi(s) = \pi(t), \quad \text{for } s, t \in \mathbb{F}_{\emptyset} A_{\{x\}}.$$

Let us highlight two particular consequences of the axioms which, in the context of finite forests, are in fact equivalent to the axioms above.

Lemma 6.5. Let \mathfrak{A} be an ω -path algebra, $a \in A_{\{x\}}$, and $b, c, d \in A_\emptyset$.

- ♦ $a(b) + a(c + d) = a(b + c) + a(d)$,
- ♦ $aa = a \quad \text{implies} \quad a(a(b + c)) = a(a(b) + c)$.

Proof. Both statements follow from the fact that both forests have the same value under Π_N . For (a), this is immediate since the two forests have the same set of paths with the same multiplicities. For (b), there is exactly one path where the two forests differ: on the left-hand side we have a path with label aac , while on the right-hand side we have ac . But $aa = a$ implies that $a \sqsubseteq_R aa$. Hence, $\chi(aac) = \langle a, c \rangle = \chi(ac)$. \square

It is straightforward to show that every ω -path algebra is locally regular.

Proposition 6.6. *Every ω -path algebra is locally MSO-definable.*

Proof. When deciding whether a product $\pi(s)$ evaluates to a given element a , we only have to count how often each (character of some) branch is realised in s , up to a given constant N . As automata can evaluate products in ω -semigroups, there exists an automaton performing this check. \square

One direction of the characterisation is straightforward.

Proposition 6.7. *Let \mathcal{S} be a pseudo-variety of finite ω -semigroups and $L \subseteq \mathbb{F}\Sigma$ an \mathcal{S} -path language.*

- (a) *The vertical ω -semigroup $\langle \text{Syn}_{\{x\}}(L), \text{Syn}_{\mathcal{O}}(L), \cdot \rangle$ of $\text{Syn}(L)$ belongs to \mathcal{S} .*
- (b) *$\text{Syn}(L)$ is an ω -path algebra.*
- (c) *If L is an \mathcal{S}^+ -path language, $\text{Syn}(L)$ is even a $+$ -path algebra.*

Proof. We use the notation $\sim_K := \ker \text{syn}_K$ for the syntactic congruence of a language K . Suppose that L is a boolean combination of the basic \mathcal{S} -path languages B_{k_i, K_i} , $i < n$.

(a) Fix ω -semigroups $\mathfrak{S}_0, \dots, \mathfrak{S}_{n-1} \in \mathcal{S}$ recognising K_0, \dots, K_{n-1} , respectively. For words $w, w' \in \Sigma^\infty$, we will show that

$$w \sim_{K_i} w', \text{ for all } i < n, \quad \text{implies} \quad w \sim_L w',$$

where in the last equation we can regard w and w' either as forests in $\mathbb{F}_{\mathcal{O}}\Sigma$ or in $\mathbb{F}_{\{x\}}\Sigma$. Then it follows that $\langle \text{Syn}_{\{x\}}(L), \text{Syn}_{\mathcal{O}}(L), \cdot \rangle$ is a quotient of the product $\prod_{i < n} \mathfrak{S}_i$ and, thus, belongs to \mathcal{S} .

For the proof note that, for every context p , $w \sim_{K_i} w'$ implies that $p[w]$ and $p[w']$ have the same number of paths with labelling in K_i . Consequently,

$$p[w] \in L \quad \text{iff} \quad p[w'] \in L.$$

This implies that $w \sim_L w'$ when we regard w and w' as elements of $\mathbb{F}_{\mathcal{O}}\Sigma$.

(b) We start with the horizontal monoid. Consider two elements $a, b \in \text{Syn}_{\mathcal{O}}(L)$ and choose trees $s \in \text{syn}_L^{-1}(a)$ and $t \in \text{syn}_L^{-1}(b)$. For commutativity, we have

$$p[s + t] \in L \Leftrightarrow p[t + s] \in L, \quad \text{for all contexts } p,$$

which implies that $s + t \sim_L t + s$. Consequently, $a + b = b + a$.

For aperiodicity, set $N := \max_i k_i$. Then we have

$$p[N \times s] \in L \Leftrightarrow p[(N+1) \times s] \in L, \quad \text{for all contexts } p.$$

Thus $N \times s \sim_L (N+1) \times s$, which implies that $N \times a = (N+1) \times a$.

For the remaining axiom, suppose that the forests $s, t \in \mathbb{F}_{\mathcal{O}}\text{Syn}_{\{x\}}(L)$ satisfy $\Pi_N(s) = \Pi_N(t)$ where $N := \max_i k_i$. As syn_L is surjective, there exists a function $\varphi : \text{Syn}(L) \rightarrow \mathbb{F}_{\mathcal{O}}\Sigma$ such that $\text{syn}_L \circ \varphi = \text{id}$. We set

$$S := \mathbb{F}\varphi(s) \quad \text{and} \quad T := \mathbb{F}\varphi(t).$$

Then $\Pi_N(\text{flat}(S)) = \Pi_N(\text{flat}(T))$. Consequently, $\text{flat}(S)$ and $\text{flat}(T)$ have the same number of paths in K_i , for $i < n$, (up to the number N) and we have

$$p[\text{flat}(S)] \in L \Leftrightarrow p[\text{flat}(T)] \in L, \quad \text{for all contexts } p.$$

Thus $\text{flat}(S) \sim_L \text{flat}(T)$ and it follows that

$$\begin{aligned} \pi(s) &= \pi(\mathbb{F}\text{syn}_L(S)) \\ &= \text{syn}_L(\text{flat}(S)) \\ &= \text{syn}_L(\text{flat}(T)) \\ &= \pi(\mathbb{F}\text{syn}_L(T)) = \pi(t). \end{aligned}$$

(c) The proof of the stronger condition is analogous to the one above. With s, t, S, T as before, we now only know that $\Pi_N^+(\text{flat}(S)) = \Pi_N^+(\text{flat}(T))$. But since the languages K_i contains only finite words, this is sufficient to imply that

$$p[\text{flat}(S)] \in L \Leftrightarrow p[\text{flat}(T)] \in L, \quad \text{for all contexts } p. \quad \square$$

To show the converse of this statement we start by simplifying the input forest t . Since its value in a given ω -path algebra only depends on the number of paths of each given kind, we can try to write t as a sum of paths. For instance, $a(b + c)$ would become $a(b) + a(c)$. The problem with this is that this changes the number of paths since we have added one copy of the path $a(x)$. To make this idea work we therefore have to be a bit more careful and only allow this operation if this additional path does not change the value of the product. The main technical tool which will enable us to do so is Lemma 6.11 below. But before proving it, let us collect a few technical results about path algebras.

Lemma 6.8. *Let $\langle S, + \rangle$ be a finite semigroup that is commutative and aperiodic. There exists an element $u \in S$ such that $u + a = u$, for all $a \in S$.*

Proof. Suppose that $S = \{c_0, \dots, c_{n-1}\}$ and set $u := c_0 + \dots + c_{n-1}$. For every $a \in S$, there is some index $i < n$ such that $\pi \times a = c_i$. Hence,

$$c_i + a = \pi \times a + a = \pi \times a = c_i,$$

which, by commutativity, implies that $u + a = u$. \square

Lemma 6.9. *Let \mathfrak{A} be a finitary \mathbb{F} -algebra where the horizontal semigroup $\langle A_\emptyset, + \rangle$ is commutative and aperiodic and let $a \in (A_{\{x\}})^1$ be $\subseteq_{\mathbb{R}}$ -minimal. Then*

$$ac = ac', \quad \text{for all } c, c' \in A_\emptyset.$$

Proof. By Lemma 6.8, there exists an element $u \in A_\emptyset$ such that $u + c = u$, for all $c \in A_\emptyset$. The element $\hat{a} := a(x + u)$ satisfies $\hat{a} \subseteq_{\mathbb{R}} a$ and

$$\hat{a}c = a(c + u) = au, \quad \text{for all } c \in A_\emptyset.$$

By $\subseteq_{\mathbb{R}}$ -minimality of a , there are elements b and d such that $a(x) = \hat{a}(b(x) + d)$. Consequently,

$$ac = \hat{a}(bc + d) = \hat{a}u = \hat{a}(bc' + d) = ac', \quad \text{for all } c, c' \in A_\emptyset. \quad \square$$

Lemma 6.10. *Let \mathfrak{A} be a finitary ω -path algebra and $a \in A_{\{x\}}$ an element.*

$$aa = a \quad \text{implies} \quad a(x) = a(x + \pi \times a(o)).$$

Proof. Lemma 6.5 (b) implies that $a(x) = aa(x) = a(x + a(o))$. Iterating this equation, we obtain $a(x) = a(x + \pi \times a(o))$. \square

Lemma 6.11. *Let \mathfrak{A} be a finitary ω -path algebra and $a, b, b' \in (A_{\{x\}})^1$. Then*

$$a = abb' \quad \text{implies} \quad a(x) = a(x + b(o)).$$

Proof. Below we will prove that the element $\hat{b} := (bb')^\pi$ satisfies the following equations.

- (1) $\pi \times \hat{b}(o) = \pi \times \hat{b}(o) + b(o)$
- (2) $a(x) = a(x + \pi \times \hat{b}(o))$

Then it follows that

$$\begin{aligned} a(x) &= a(x + \pi \times \hat{b}(o)) \\ &= a(x + \pi \times \hat{b}(o) + b(o)) \\ &= a((x + b(o)) + \pi \times \hat{b}(o)) \\ &= a(x + b(o)), \end{aligned}$$

where the first and last step follow by (2), the second step by (1), and the third one by commutativity. Hence, it remains to prove the two equations.

(1) We have

$$\begin{aligned} \pi \times \hat{b}(o) &= \pi \times \hat{b}(o) + \pi \times \hat{b}(o) \\ &= \pi \times (\hat{b}(o) + \hat{b}(o)) \\ &= \pi \times (bb'\hat{b}(o) + bb'\hat{b}(o)) \\ &= \pi \times (b(b'\hat{b}(o) + b'\hat{b}(o)) + b(o)) \\ &= \pi \times b(b'\hat{b}(o) + b'\hat{b}(o)) + \pi \times b(o), \end{aligned}$$

where the fifth step follows by Lemma 6.5 (b). Consequently,

$$\begin{aligned}\pi \times \hat{b}(o) + b(o) &= \pi \times b(b'\hat{b}(o) + b'\hat{b}(o)) + \pi \times b(o) + b(o) \\ &= \pi \times b(b'\hat{b}(o) + b'\hat{b}(o)) + \pi \times b(o) \\ &= \pi \times \hat{b}(o).\end{aligned}$$

(2) Set $d := \hat{b}(x) + \pi \times \hat{b}(o)$. Then

$$\begin{aligned}d &= \hat{b}(x) + \pi \times \hat{b}(o) \\ &= \hat{b}(x) + \pi \times \hat{b}(o) + \pi \times \hat{b}(o) \\ &= d(x) + \pi \times \hat{b}(o)\end{aligned}$$

and it follows that

$$d^\pi = dd^\pi = dd^\pi(x) + \pi \times \hat{b}(o) = d^\pi(x) + \pi \times \hat{b}(o).$$

Consequently, we have

$$\begin{aligned}a &= abb' = a\hat{b} = a\hat{b}\hat{b} \\ &= a\hat{b}(\hat{b}(x) + \pi \times \hat{b}(o)) \\ &= a(\hat{b}(x) + \pi \times \hat{b}(o)) \\ &= ad \\ &= ad^\pi d^\pi \\ &= a(d^\pi d^\pi(x) + \pi \times \hat{b}(o)) \\ &= ad^\pi d^\pi(x + \pi \times \hat{b}(o)) \\ &= a(x + \pi \times \hat{b}(o)),\end{aligned}$$

where the forth step holds by Lemma 6.10, the fifth one follows from the equation $a = a\hat{b}$, and the last step from $a = ad^\pi d^\pi$. \square

Lemma 6.12. *Let \mathcal{A} be a finitary ω -path algebra. Then*

$$a = a(b(x) + c) \quad \text{implies} \quad a = ab = a(x + c).$$

Proof. Note that it is sufficient to prove that $a = a(x + c)$, since it then follows that

$$ab = a(b(x) + c) = a.$$

Let $\hat{b} \in A_{\{x\}}$ be the element such that

$$\hat{b}(x) + c = (b(x) + c)^\pi,$$

and set

$$d(x) := \hat{b}(x) + \pi \times \hat{b}(o) + \pi \times c.$$

Note that $d = d(x) + c$, which implies that $d^\pi = d^\pi(x) + c$. Hence, we have

$$d^\pi = d^\pi d^\pi = d^\pi(d^\pi(x) + c) = d^\pi d^\pi(x + c) = d^\pi(x + c),$$

where the third step follows by Lemma 6.5 (b). Consequently,

$$\begin{aligned}a &= a(b(x) + c) \\ &= a(\hat{b}(\hat{b}(x) + c) + c) \\ &= a(\hat{b}(\hat{b}(x) + c + \pi \times (\hat{b}(o) + c)) + c) \\ &= a(\hat{b}(x) + c + \pi \times \hat{b}(o) + \pi \times c) \\ &= a(\hat{b}(x) + \pi \times \hat{b}(o) + \pi \times c) \\ &= ad \\ &= ad^\pi \\ &= ad^\pi(x + c) = a(x + c),\end{aligned}$$

where the third step follows by Lemma 6.10. \square

After these preparations let us proceed to simplifying forests. Since the main argument below is inductive, we have to phrase everything relative to some character σ . What we are aiming for are forests that are σ -separated in the following sense.

Definition 6.13. Let \mathfrak{A} be a ω -path algebra and $\sigma \in \Psi$.

(a) An element $a \in A_{\{x\}}$ is σ -insignificant if $\sigma \sqsubseteq_R \sigma a$. Otherwise, a is called σ -significant.

(b) For forests $s, t \in \mathbb{F}_{\{x\}} A_{\{x\}}$, we set $s \sim_\sigma t$ if there are finite paths $p_i, q_i \in \mathbb{F}_{\{x\}} A_{\{x\}}$ such that

$$\Pi_N(\sigma(s + \sum_{i < m} p_i(o))) = \Pi_N(\sigma(t + \sum_{i < n} q_i(o))),$$

and the products $\pi(p_i)$ and $\pi(q_i)$ are σ -insignificant.

(c) A forest $s \in \mathbb{F}_{\emptyset} A$ is σ -separated if it is a finite horizontal sum of paths and of forests of the form $p(r)$, where r is an arbitrary forest and p is a path whose product $\pi(p)$ is σ -significant. \downarrow

Lemma 6.14. Let \mathfrak{A} be a finitary ω -path algebra and $\sigma \in \Psi$ a character.

(a) $s \sim_\sigma t$ implies $\pi(\sigma(s)) = \pi(\sigma(t))$.

(b) If $a \in A_{\{x\}}$ is σ -significant, then so is $a(x) + c$.

Proof. (a) Suppose that

$$\Pi_N(\sigma(s + \sum_{i < m} p_i(o))) = \Pi_N(\sigma(t + \sum_{i < n} q_i(o))),$$

with $c_i := \pi(p_i)$ and $d_i := \pi(q_i)$ σ -insignificant. Setting $a := \pi(s)$ and $b := \pi(t)$, it follows by Lemma 6.11 that

$$\begin{aligned} \sigma a(x) &= \sigma(a(x) + c_o(o)) \\ &= \sigma(a(x) + c_o(o) + c_1(o)) \\ &= \dots \\ &= \sigma(a(x) + c_o(o) + \dots + c_{m-1}(o)) \\ &= \sigma(b(x) + d_o(o) + \dots + d_{n-1}(o)) \\ &= \dots \\ &= \sigma(b(x)). \end{aligned}$$

(b) Suppose that $a(x) + c$ is σ -insignificant. Then there exists an element b such that $\sigma(x) = \sigma(ab(x) + c)$. By Lemma 6.12, it follows that $\sigma(x) = \sigma(ab(x))$, which means that a is σ -insignificant. \square

We introduce the following operation on forests.

Definition 6.15. Given a forest $s \in \mathbb{F}_{\{x\}} A_{\{x\}}$ and two sets of vertices $H_o \subseteq H \subseteq \text{dom}(s)$ where H_o is finite, we set

$$s[H/H_o] := s' + \sum_{v \in H_o} \tilde{\text{pr}}_{v, \sigma}(\text{sf}_v),$$

where s' is the forest obtained from s by removing all subtrees attached to some vertex in H . \downarrow

Lemma 6.16. Let $s \in \mathbb{F}_{\{x\}} A_{\{x\}}$ be a forest and $\sigma \in \Psi$ a character. Then

$$\Pi_N(\sigma s/H_o) = \Pi_N(\sigma s/H)$$

implies that

$$\Pi_N(\sigma s[H/H_o]) = \Pi_N(\sigma(s + \sum_{v \in H_o} \tilde{\text{pr}}_{v, \sigma}(o))).$$

Proof. Set $t := \sigma(s + \sum_{v \in H_o} \tilde{\text{pr}}_{v, \sigma}(o))$ and $t' := \sigma(s[H/H_o])$. Note that for every path of t' there is a corresponding path of t with the same image under χ . Consequently,

$$\Pi_N(t')(\tau) \leq \Pi_N(t)(\tau), \quad \text{for all } \tau \in \Psi.$$

For the other direction, suppose that there is some sequence $\tau \in \Psi$ with

$$\Pi_N(t')(\tau) < \Pi_N(t)(\tau).$$

By definition of $s[H/H_o]$, it follows that there are vertices $v \leq u$ of s with $v \in H \setminus H_o$ and $\tau = \chi(\sigma s|_u)$. If $\Pi_N(H_o) < N$, this would mean that $\Pi_N(H) > \Pi_N(H_o)$, which contradicts our assumption on H_o . Hence, $\Pi_N(H_o) = N$ and we have

$$N \geq \Pi_N(t)(\tau) > \Pi_N(t')(\tau) \geq \Pi_N(H_o)(\tau) = N.$$

A contradiction. \square

Proposition 6.17. *Let \mathfrak{A} be a finitary ω -path algebra and $\sigma \in \Psi$. For every forest $s \in \mathbb{F}_{\emptyset} A_{\{x\}}$, there exists a σ -separated forest $t \in \mathbb{F}_{\emptyset} A_{\{x\}}$ such that $s \sim_{\sigma} t$.*

Proof. Let us call a vertex $v \in \text{dom}(s)$ *insignificant* if the product $\pi(\text{pr}_v)$ is σ -insignificant. The set I of all insignificant vertices forms an initial subset of $\text{dom}(s)$. The *frontier* H of s consists of the minimal (in the forest ordering) vertices of $\text{dom}(s) \setminus I$. We choose a finite subset $H_o \subseteq H$ such that

$$\Pi_N(\sigma s / H_o) = \Pi_N(\sigma s / H).$$

Set

$$t' := s[H/H_o] \quad \text{and} \quad s' := s + \sum_{v \in H_o} \tilde{\text{pr}}_{v,\sigma}(\circ).$$

Then it follows by Lemma 6.16 that $\Pi_N(\sigma s') = \Pi_N(\sigma t')$, which implies that $s \sim_{\sigma} t'$.

Hence, it is sufficient to find a σ -separated forest $t \sim_{\sigma} t'$. By definition, we have

$$t' = s[H/H_o] = r' + \sum_{v \in H_o} \tilde{\text{pr}}_{v,\sigma}(\text{sf}_v),$$

where r' is the forest obtained from s by removing all subtrees attached to a vertex in H . If we can find a σ -separated forest $r \sim_{\sigma} r'$, we can take

$$t := r + \sum_{v \in H_o} \tilde{\text{pr}}_{v,\sigma}(\text{sf}_v)$$

as the desired σ -separated forest with $t \sim_{\sigma} s$. Set $\mu := \Pi_N(\sigma r')$. Let B be a set containing, for every infinite character $\tau \in A_{\{x\}}^* A_{\emptyset}$, exactly $\mu(\tau)$ infinite branches of r' with character τ . Furthermore, fix a finite initial subset $I \subseteq \text{dom}(r')$ containing, for every finite character $\tau \in A_{\{x\}}^*$, at least $\mu(\tau)$ vertices v such that $\chi(\sigma \text{pr}_v) = \tau$. Let r'' be the forest obtained from r' by deleting every vertex that does not belong to $I \cup B$. Then

$$\Pi_N(\sigma r'') = \mu = \Pi_N(\sigma r').$$

Finally, set

$$r := \sum_{\beta \in B} r''|_{\beta} + \sum_{v \in I} \text{pr}_v(\circ).$$

By construction, we have

$$\Pi_N(\sigma r)(\tau) \geq \mu(\tau) = \Pi_N(\sigma r'')(\tau).$$

Set $\sigma := \Pi_N(\sigma r)$. Then

$$\begin{aligned} & \Pi_N(\sigma(r'' + \sum_{\tau} (\sigma(\eta) - \mu(\eta)) \times \pi(\eta)(\circ)))(\tau) \\ &= \mu(\tau) + \sigma(\tau) - \mu(\tau) \\ &= \sigma(\tau) \\ &= \Pi_N(\sigma r). \end{aligned}$$

Hence, $r \sim_{\sigma} r'' \sim_{\sigma} r'$, as desired. \square

This proposition provides our simplification operation for forests. The next step is to show that the product of σ -separated forests can be defined using path languages. Our main tool will be what we call the *profile* of such a forest.

Definition 6.18. Let \mathfrak{A} be a ω -path algebra and $\sigma \in \Psi$.

(a) For $c, c' \in A_{\emptyset}$, we set

$$c \approx_{\sigma} c' \quad : \text{iff} \quad \sigma a c = \sigma a c', \quad \text{for all } \sigma\text{-significant } a \in A_{\{x\}}.$$

We denote the \approx_{σ} -class of c by $[c]_{\sigma}$.

(b) We denote by $\Psi_{\{x\}}^{(\sigma)}$ the subset of $\Psi_{\{x\}}$ consisting of all characters τ whose product $\pi(\tau)$ is σ -significant.

(c) The σ -profile $\text{pf}(s)$ of a σ -separated forest $s = \sum_{i < m} q_i + \sum_{i < n} p_i(r_i)$ consists of two functions

$$\text{pf}_{\circ}(s) : \Psi_{\emptyset} \rightarrow [N + 1] \quad \text{and} \quad \text{pf}_1(s) : \Psi_{\{x\}}^{(\sigma)} \rightarrow [N + 1] \times A_{\emptyset} / \approx_{\sigma},$$

where the first one maps a character $\tau \in \Psi_{\emptyset}$ to the number

$$|\{i < m \mid \chi(q_i) = \tau\}|$$

and the second one maps a character $\tau \in \Psi_{\{x\}}^{(\sigma)}$ to the pair $\langle k, [c]_{\sigma} \rangle$ where

$$k := |\{i < n \mid \chi(p_i) = \tau\}| \quad \text{and} \quad c := \sum \{ \pi(r_i) \mid \chi(p_i) = \tau \}.$$

Lemma 6.19. *If s and s' are σ -separated forests, then*

$$\text{pf}(s) = \text{pf}(s') \quad \text{implies} \quad \pi(\sigma s) = \pi(\sigma s').$$

Proof. Suppose that

$$\text{pf}_{\circ}(s)(\tau) = l_{\tau} \quad \text{and} \quad \text{pf}_1(s)(\tau) = \langle k_{\tau}, [c_{\tau}]_{\tau} \rangle.$$

Then

$$\begin{aligned} \pi(s) &= b + \sum_{\tau \in \Psi_{\circ}} l_{\tau} \times \tau + \sum_{\tau \in \Psi_1^{(\sigma)}} \sum_{i < k_{\tau}} \tau d_{\tau}^i \\ &= b + \sum_{\tau \in \Psi_{\circ}} l_{\tau} \times \tau + \sum_{\tau \in \Psi_1^{(\sigma)}} (\tau \hat{d}_{\tau} + (k_{\tau} - 1) \times \tau(o)), \\ \pi(s') &= b' + \sum_{\tau \in \Psi_{\circ}} l_{\tau} \times \tau + \sum_{\tau \in \Psi_1^{(\sigma)}} \sum_{i < k_{\tau}} \tau e_{\tau}^i \\ &= b' + \sum_{\tau \in \Psi_{\circ}} l_{\tau} \times \tau + \sum_{\tau \in \Psi_1^{(\sigma)}} (\tau \hat{e}_{\tau} + (k_{\tau} - 1) \times \tau(o)), \end{aligned}$$

where b and b' are sums of σ -insignificant elements, $\hat{d}_{\tau} := \sum_i d_{\tau}^i$, $\hat{e}_{\tau} := \sum_i e_{\tau}^i$, and $\hat{d}_{\tau} \approx_{\sigma} c_{\tau} \approx_{\sigma} \hat{e}_{\tau}$. Note that, according to Lemma 6.14 (b), the element $a(x) + b$ is σ -significant, for every b . By the definition of \approx_{σ} , it therefore follows that

$$\hat{d}_{\tau} \approx_{\sigma} \hat{e}_{\tau} \quad \text{implies} \quad \sigma(\tau \hat{d}_{\tau} + a) = \sigma(\tau \hat{e}_{\tau} + a), \quad \text{for all } a \in A_{\emptyset}.$$

In particular for

$$a := \sum_{\tau \in \Psi_{\circ}} l_{\tau} \times \tau + \sum_{\tau \in \Psi_1} (k_{\tau} - 1) \times \tau(o),$$

we have

$$\begin{aligned} \pi(\sigma(s)) &= \sigma\left(a + b + \sum_{\tau \in \Psi_1} \tau \hat{d}_{\tau}\right) \\ &= \sigma\left(a + \sum_{\tau \in \Psi_1} \tau \hat{d}_{\tau}\right) \\ &= \sigma\left(a + \sum_{\tau \in \Psi_1} \tau \hat{e}_{\tau}\right) \\ &= \sigma\left(a + b' + \sum_{\tau \in \Psi_1} \tau \hat{e}_{\tau}\right) = \pi(\sigma(s')), \end{aligned}$$

where the second and forth step follow by Lemma 6.11. \square

We are finally at the point where we can prove the other part of our characterisation. The main step of the proof is the following proposition.

Proposition 6.20. *Let \mathfrak{A} be an ω -path algebra, \mathcal{S} a pseudo-variety of finite ω -semigroups, and \mathcal{K} the corresponding family of ∞ -word languages. Suppose that \mathcal{K} is closed under concatenation and that the vertical ω -semigroup $\langle A_{\{x\}}, A_{\emptyset}, \cdot \rangle$ belongs to \mathcal{S} . Let $\varphi : \mathbb{T}\Sigma \rightarrow \mathfrak{A}$ be a morphism of \mathbb{F} -algebras, $\sigma \in \Psi_{\{x\}}$ a character, and $c \in A_{\emptyset}$ an element. Then*

$$L_{\sigma, c} := \{s \in \mathbb{F}\Sigma \mid \sigma(\varphi(s)) = c\}$$

is an \mathcal{S} -path language. If \mathfrak{A} is a $+$ -path algebra, $L_{\sigma, c}$ is even an \mathcal{S}^+ -path language.

Proof. We prove the claim by induction on the $\sqsubseteq_{\mathbb{R}}$ -class of σ . If σ is $\sqsubseteq_{\mathbb{R}}$ -minimal, Lemma 6.9 tells us that $L_{\sigma, c}$ is either empty or it contains all forests. It particular, it is an \mathcal{S}^+ -path language for every pseudo-variety \mathcal{S} .

For the inductive step, suppose that σ is not $\subseteq_{\mathbb{R}}$ -minimal. Below we will construct \mathcal{S} -path languages $M_{\tau,k}, M_{\tau,\langle k, [c]_{\sigma} \rangle} \subseteq \mathbb{F}A$ such that, for every σ -separated forest $s \in \mathbb{F}A$,

$$\begin{aligned} s \in M_{\tau,k} & \quad \text{iff} \quad \text{pf}_o(s)(\tau) = k, \\ s \in M_{\tau,\langle k, [c]_{\sigma} \rangle} & \quad \text{iff} \quad \text{pf}_1(s)(\tau) = \langle k, [c]_{\sigma} \rangle. \end{aligned}$$

and, for arbitrary forests $s, s' \in \mathbb{F}_{\emptyset} A_{\{x\}}$, $s \sim_{\sigma} s'$ implies that

$$s \in M_{\tau,k} \Leftrightarrow s' \in M_{\tau,k} \quad \text{and} \quad s \in M_{\tau,\langle k, [c]_{\sigma} \rangle} \Leftrightarrow s' \in M_{\tau,\langle k, [c]_{\sigma} \rangle}.$$

We claim that

$$L_{\sigma,c} = \bigcup \left\{ (\mathbb{F}\varphi_o)^{-1}[M_{\mu}] \mid \mu \text{ a profile such that} \right. \\ \left. \text{pf}(s) = \mu \Rightarrow \pi(\sigma s) = c \right\},$$

where $\varphi_o := \varphi \circ \text{sing}$ is the restriction of φ to the alphabet Σ and

$$M_{\mu} := \bigcap_{\tau \in \Psi_o} M_{\tau, \mu_o(\tau)} \cap \bigcap_{\tau \in \Psi_{\{x\}}^{(o)}} M_{\tau, \mu_1(\tau)}.$$

To prove this equality, we start by noting that, for a σ -separated forest $s \in \mathbb{F}A$,

$$\begin{aligned} s \in M_{\mu} & \quad \text{iff} \quad s \in M_{\tau, \mu_o(\tau)} \text{ and } s \in M_{\tau, \mu_1(\tau)}, \quad \text{for all suitable } \tau, \\ & \quad \text{iff} \quad \text{pf}_o(s)(\tau) = \mu_o(\tau) \text{ and } \text{pf}_1(s)(\tau) = \mu_1(\tau), \\ & \quad \quad \quad \text{for all suitable } \tau, \\ & \quad \text{iff} \quad \text{pf}(s) = \mu. \end{aligned}$$

Now consider a forest $t \in \mathbb{F}\Sigma$ and set $s := \mathbb{F}\varphi_o(t)$. By Proposition 6.17, there exists a σ -separated forest $s' \sim_{\sigma} s$. Furthermore, Lemma 6.14 implies that

$$\pi(\sigma s') = \pi(\sigma s) = \sigma(\varphi(t)).$$

Consequently, we have

$$\begin{aligned} t \in L_{\sigma,c} & \quad \text{iff} \quad s \in M_{\mu}, \quad \text{for some } \mu \text{ such that} \\ & \quad \text{pf}(r) = \mu \Rightarrow \pi(\sigma r) = c \\ & \quad \text{iff} \quad s' \in M_{\mu}, \quad \text{for some } \mu \text{ such that} \\ & \quad \text{pf}(r) = \mu \Rightarrow \pi(\sigma r) = c \\ & \quad \text{iff} \quad \text{pf}(r) = \text{pf}(s') \Rightarrow \pi(\sigma r) = c \\ & \quad \text{iff} \quad \pi(\sigma s') = c \\ & \quad \text{iff} \quad \sigma(\varphi(t)) = c, \end{aligned}$$

where the fourth equivalence follows by Lemma 6.19.

It remains to construct the languages $M_{\sigma,k}$ and $M_{\sigma,\langle k, [c]_{\sigma} \rangle}$. First, note that, since $\langle A_{\{x\}}, A_{\emptyset}, \cdot \rangle \in \mathcal{S}$ and \mathcal{K} is closed under concatenation, the class \mathcal{K} contains all word languages of the form

$$K_{\tau} := \chi^{-1}(\tau), \quad \text{for } \tau \in \Psi.$$

Consequently,

$$M_{\tau,k} = B_{k,K_{\tau}} \setminus B_{k+1,K_{\tau}}$$

is an \mathcal{S} -path language. (Recall that $B_{k,K}$ refers to one of the basic \mathcal{S} -path languages.) Furthermore, τ being σ -significant we clearly have

$$s \sim_{\sigma} s' \quad \text{implies} \quad s \in M_{\tau,k} \Leftrightarrow s' \in M_{\tau,k},$$

for all forests $s, s' \in \mathbb{F}_{\emptyset} A_{\{x\}}$. Before defining $M_{\tau,\langle k, [c]_{\sigma} \rangle}$, we show that the languages

$$N_{[c]_{\sigma}} := \{ s \in \mathbb{F}_{\emptyset} A \mid \pi(s) \approx_{\sigma} c \}$$

are \mathcal{S} -path languages. Let H be the set of all σ -significant elements in $A_{\{x\}}$. Note that

$$\begin{aligned} s \in N_{[c]_\sigma} & \quad \text{iff} \quad \pi(s) \approx_\sigma c \\ & \quad \text{iff} \quad \sigma a(\pi(s)) = \sigma a(c), \quad \text{for all } a \in H, \\ & \quad \text{iff} \quad s \in L_{\sigma a, \sigma a(c)}, \quad \text{for all } a \in H, \\ & \quad \text{iff} \quad s \in \bigcap_{a \in H} L_{\sigma a, \sigma a(c)}. \end{aligned}$$

Using the inductive hypothesis, we see that $N_{[c]_\sigma}$ is an intersection of \mathcal{S} -path languages and, thus, an \mathcal{S} -path language itself.

Let $N'_{\tau, [c]_\sigma}$ be the language obtained from $N_{[c]_\sigma}$ by replacing each basic \mathcal{S} -path language $B_{l, K}$ appearing in the corresponding boolean combination by the language B_{l, K_τ} where

$$K_\tau := \chi^{-1}(\tau), \quad \text{for } \tau \in \Psi_{\{x\}}^{(\sigma)}.$$

For a path p with $\chi(p) = \tau \in \Psi_{\{x\}}^{(\sigma)}$ it then follows

$$ps \in N'_{\tau, [c]_\sigma} \quad \text{iff} \quad \pi(s) \approx_\sigma c,$$

Hence, we can set

$$M_{\tau, \langle k, [c]_\sigma \rangle} := (B_{k, K_\tau} \setminus B_{k+1, K_\tau}) \cap N'_{\tau, [c]_\sigma}.$$

Before showing that $M_{\tau, \langle k, [c]_\sigma \rangle}$ has the desired property, let us note that, since $\pi(\tau)$ is σ -significant, we have

$$s \sim_\sigma s' \quad \text{implies} \quad s \in M_{\tau, \langle k, [c]_\sigma \rangle} \Leftrightarrow s' \in M_{\tau, \langle k, [c]_\sigma \rangle},$$

for arbitrary forests $s, s' \in \mathbb{F}_{\emptyset A_{\{x\}}}$.

Given a σ -separated forest s , the first term of the above intersection expresses that $\text{pf}(s)(\tau) = \langle k, [d]_\sigma \rangle$, for some d . We claim that the second term ensures that $d \approx_\sigma c$. Hence, consider a σ -separated forest s . By the first term of the intersection, we may assume that s has exactly k components of

the form $p(r)$ with $\chi(p) = \tau$. Hence, $s = s_o + \sum_{i < k} p_i(r_i)$ with $\chi(p_i) = \tau$. Set

$$r := \sum_i r_i \quad \text{and} \quad s' := s_o + p_o(r) + \sum_{o < i < k} p_i(o).$$

Then $\Pi_N(s) = \Pi_N(s')$, which implies that

$$\begin{aligned} s \in N'_{\tau, [c]_\sigma} & \quad \text{iff} \quad s' \in N'_{\tau, [c]_\sigma} \\ & \quad \text{iff} \quad p_o(r) \in N'_{\tau, [c]_\sigma} \\ & \quad \text{iff} \quad \pi(r) \approx_\sigma c \\ & \quad \text{iff} \quad \text{pf}_i(s)(\tau) = \langle k, [c]_\sigma \rangle. \end{aligned}$$

To conclude the proof, suppose that \mathfrak{A} is a $+$ -path algebra. We want to show that $L_{\sigma, c}$ is a \mathcal{S}^+ -path language. The above proof goes through with minor modifications: we have to change the definitions of $\text{pf}_o(s)(\tau)$ and $M_{\tau, k}$ to only count finite paths. \square

Combining Proposition 6.7 and 6.20 we obtain the following characterisation result.

Theorem 6.21. *Let \mathcal{S} be a pseudo-variety of finite ω -semigroups such that the corresponding family of ∞ -word languages is closed under concatenation.*

(a) *A forest language L is an \mathcal{S} -path language if, and only if, it is recognised by an ω -path algebra \mathfrak{A} whose vertical semigroup $\langle A_{\{x\}}, A_\emptyset, \cdot \rangle$ belongs to \mathcal{S} .*

(b) *A forest language L is an \mathcal{S}^+ -path language if, and only if, it is recognised by a $+$ -path algebra \mathfrak{A} whose vertical semigroup $\langle A_{\{x\}}, A_\emptyset, \cdot \rangle$ belongs to \mathcal{S} .*

Let us see what this theorem entails for the graded logics we introduced above. We start by defining the needed algebras.

For $n < \omega$, let $\mathfrak{N}_{(1)}^n$ be the \mathbb{F} -algebra with domains

$$(N_{(1)}^n)_\xi := [n+1], \quad \text{for } \xi \in \Xi.$$

The product of a forest s is simply the sum (up to n) of all labels $s(v)$, $v \in \text{dom}(s)$.

For $n < \omega$, let $\mathfrak{N}_{(2)}^{*,n}$ be the \mathbb{F} -algebra with domains

$$(N_{(2)}^{*,n})_{\xi} := [n+1] \times \wp(\xi), \quad \text{for } \xi \in \Xi.$$

To define the product of a forest s , let us call a vertex $v \in \text{dom}(s)$ *reachable* if, for each $u < v$ and every variable x with $ux \leq v$, we have $s(u) = \langle i, I \rangle$ with $x \in I$. We set

$$\pi(s) := \langle j, J \rangle,$$

where j is the sum of all first components of labels of reachable vertices (up to n), and J is the set of all reachable variables.

We also define a variant $\mathfrak{N}_{(2)}^{\omega,n}$ with the domains

$$(N_{(2)}^{\omega,n})_{\xi} := [n+1] \times [n+1] \times \wp(\xi),$$

where the product is given by

$$\pi(s) := \langle j, k, J \rangle,$$

such that j is the sum of all first components of labels of reachable vertices (up to n), k is the sum of all second components of reachable vertices plus the number of infinite branches all vertices of which are reachable, and J is the set of all reachable variables.

The proof of the following theorem is similar to that of Theorem 5.5.

Theorem 6.22. *A forest language is definable in one of the logics below if, and only if, it is recognised by an iterated wreath product of algebras from the following table.*

logic	algebras
cEF	$\mathfrak{N}_{(1)}^n, n < \omega$
cwCTL	$\mathfrak{N}_{(2)}^{*,n}, n < \omega$
cCTL	$\mathfrak{N}_{(2)}^{*,n}, \mathfrak{N}_{(2)}^{\omega,n}, n < \omega$
cPDL	path algebras
cwCTL*	aperiodic path algebras
cCTL*	aperiodic ω -path algebras
FO	aperiodic path algebras
WMPL	aperiodic path algebras
MPL	aperiodic ω -path algebras
WMCL	path algebras
MCL	ω -path algebras

Notes

Much of this chapter is based on [9], where the corresponding material was developed for finite forests. The equivalence between $\text{MPL} = \text{cCTL}^*$ has been proven by Moller and Rabinovitch [17].

Recommended Literature

- 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.
- G. Rozenberg, A. Salomaa, *Handbook of Formal Languages*, Springer 1997.
- J.-É. Pin, *Handbook of Automata Theory*, EMS Press 2021.
- D. Perrin, J.-É. Pin, *Infinite Words*, Birkhäuser 2004.
- J.-É. Pin, *Mathematical Foundations of Automata Theory*, lecture notes.
- J. Berstel, C. Reutenauer, *Noncommutative Rational Series with Applications*, Cambridge Universal Press 2011.
- J. Sakarovitch, *Elements of Automata Theory*, Cambridge Universal Press 2009.

Bibliography

- [1] J. T. BALDWIN, *Categoricity*, American Mathematical Society, 2009.
- [2] J. BARWISE AND S. FEFERMAN, *Model-Theoretic Logics*, Springer, 1985.
- [3] A. BLUMENSATH, *Monadic Second-Order Model Theory*. unpublished lecture notes, <https://www.fi.muni.cz/~blumens/MSO.pdf>.
- [4] ———, *Algebraic Language Theory for Eilenberg–Moore Algebras*, Logical Methods in Computer Science, 17 (2021), pp. 6:1–6:60.
- [5] M. BOJAŃCZYK, *Recognisable languages over monads*. unpublished note, arXiv:1502.04898v1.
- [6] ———, *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.
- [7] M. BOJAŃCZYK AND B. KLIN, *A non-regular language of infinite trees that is recognisable by a finite algebra*, Logical Methods in Computer Science, 15 (2019).
- [8] M. BOJAŃCZYK AND H. MICHAŁEWSKI, *Some connections between universal algebra and logics for trees*. arXiv:1703.04736, unpublished.
- [9] M. BOJAŃCZYK, H. STRAUBING, AND I. WALUKIEWICZ, *Wreath products of forest algebras, with applications to tree logics*, Logical Methods in Computer Science, 8 (2012).
- [10] 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.
- [11] L.-T. CHEN, J. ADÁMEK, S. MILIUS, AND H. URBAT, *Profinite monads, profinite equations and reiterman’s theorem*, in Proc. 19th Int. Conference on Foundations of Software Science and Computation Structures, FoSSaCS, 2016, pp. 531–547. (preprint with proofs: arXiv:1511.02147).

- [12] B. A. DAVEY AND H. A. PRIESTLEY, *Introduction to Lattices and Order*, Cambridge University Press, 2002.
- [13] R. DIACONESCU, *Institution-independent Model Theory*, Birkhäuser, 2008.
- [14] N. GAMBINO AND J. KOCK, *Polynomial Functors and Polynomial Monads*, Math. Proc. Cambridge Phil. Soc., 154 (2013), pp. 153–192.
- [15] C. LÖDING, *Automata on infinite trees*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 265–302.
- [16] S. MILIUS AND H. URBAT, *Equational Axiomatization of Algebras with Structure*, in Proc. 22nd International Conference on Foundations of Software Science and Computation Structures, FOSSACS 2019, 2019, pp. 400–417.
- [17] F. MOLLER AND A. RABINOVITCH, *Counting on CTL*: on the expressive power of monadic path logic*, Information and Computation, 184 (2003), pp. 147–159.
- [18] J.-É. PIN, *Mathematical Foundations of Automata Theory*. unpublished lecture notes, <http://www.liafa.jussieu.fr/~jep/PDF/MPRI/MPRI.pdf>.
- [19] A. POTTHOFF, *Logische Klassifizierung regulärer Baumsprachen*, Ph. D. Thesis, Universität Kiel, Kiel, 1994.
- [20] L. RIBES AND P. ZALESSKII, *Profinite Groups*, Springer-Verlag, 2010.
- [21] J. SALAMANCA, *Unveiling eilenberg-type correspondences: Birkhoff's theorem for (finite) algebras + duality*. unpublished note, arXiv:1702.02822v1.
- [22] W. THOMAS, *Languages, Automata, and Logic*, in Handbook of Formal Languages, G. Rozenberg and A. Salomaa, eds., vol. 3, Springer, New York, 1997, pp. 389–455.
- [23] H. URBAT, J. ADÁMEK, L.-T. CHEN, AND S. MILIUS, *Eilenberg theorems for free*, in 42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21–25, 2017 – Aalborg, Denmark, vol. 83, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, pp. 43:1–43:15.
- [24] I. WALUKIEWICZ, *Monadic second-order logic on tree-like structures*, Theoretical Computer Science, 275 (2002), pp. 311–346.
- [25] Y. I. YANOV AND A. A. MUCHNIK, *Existence of k -valued closed classes without a finite basis*, Dokl. Akad. Nauk., 127 (1959), pp. 44–46.

Symbol Index

Chapter I

$[n]$	$\{0, \dots, n-1\}$, 8
$\uparrow X$	upwards closure, 8
$\downarrow X$	downwards closure, 8
Set	category of sets, 8
Pos	category of posets, 8
Top	category of topological spaces, 8
Met	category of bounded metric spaces, 8
$\mathfrak{S}\text{-Set}$	category of \mathfrak{S} -sets, 9
$\mathbb{J}X$	discrete object, 10
$\mathbb{V}A$	underlying set, 10
A^X	power of A , 19
$\text{dom}(s)$	domain of s , 23
$\Pi(C)$	free product completion, 27

Chapter II

$\text{im } \varphi$	image, 94
$\ker \varphi$	kernel, 94
$\varepsilon \perp \mu$	orthogonal morphisms, 97
$\langle\langle X \rangle\rangle_A$	generated subobject, 104
$\langle\langle X \rangle\rangle_{\mathfrak{A}}$	generated subalgebra, 104
\mathfrak{A}_ρ	ρ -reduct of \mathfrak{A} , 116

$A \uparrow \mathcal{E}$	expansions to sorts in \mathcal{E} , 117
A_Δ	reduct to sorts in Δ , 117
A/θ	quotient, 119
$\text{Cong}(A)$	congruences of A , 120
$\text{Cong}_{\mathbb{M}}(\mathfrak{A})$	\mathbb{M} -congruences of \mathfrak{A} , 120
$[a]_{\sim}$	equivalence class, 122
A/\sqsubseteq	quotient, 123
$[a]_{\sqsubseteq}$	\sqsubseteq -class, 123
$E_{\mathcal{P}}$	epimorphisms making every $P \in \mathcal{P}$ projective, 129
$S(\mathcal{K})$	closure under subobjects, 131
$H(\mathcal{K})$	closure under quotients, 131
$P_{\kappa}(\mathcal{K})$	closure under products, 132

Chapter III

$f^{-1}[\kappa]$	inverse image of κ , 146
$\text{TC}(\theta)$	transitive closure, 153
$\mathbb{L}\mathfrak{A}$	language algebra, 162
$\langle\langle C \rangle\rangle_{\mathbb{L}}$	subalgebra generated by C , 162
syn_{κ}	syntactic morphism, 163
$\text{Syn}(\kappa)$	syntactic algebra, 164
\mathcal{T}	term algebras, 182
\mathcal{R}	recognisers, 182
$\mathfrak{A} \models s \leq t$	satisfaction, 214

$\text{Th}(\mathcal{V})$ theory of \mathcal{V} , 214

Chapter IV

$\text{Mod}(\varphi)$ models of φ , 219
 $\langle\langle M; \varphi \rangle\rangle$ truth value of φ in the
 model M , 219
 $\text{Th}_L(M)$ L -theory, 219
 $\text{eq}(L)$ L -equivalence, 220
 $\mathcal{L}[\Omega]$ Ω -extension, 221

Chapter V

Ξ sorts, 253
 Ξ_+ extended sorts, 253
 \mathbb{G} nondeterministic rooted
 graphs, 253
 $\text{dom}_+(g)$ extended domain, 254
 $\text{dom}(g)$ domain, 254
 \mathbb{R} deterministic rooted graphs,
 254
 $\text{un}(g)$ unravelling, 255
 sing singleton graph, 255
 flat flattening for graphs and
 trees, 256
 $\text{MSO}[\leq_{\text{so}}]$ MSO with successor
 ordering, 260
 $\text{MSO}[\leq_{\text{pf}}]$ MSO with forest ordering,
 260
 $\text{MSO}[\text{suc}]$ MSO with successor
 relations, 260
 CMSO counting MSO, 260
 \equiv_L logical equivalence, 261
 \mathbb{P}^{fin} finite forests, 263

$a + b$ horizontal product, 263
 $a \cdot b$ vertical product, 263
 $s_0 \sqcup \dots \sqcup s_{n-1}$ shuffle, 267
 \mathbb{C}^{reg} regular linear orders, 268
 $a \leq_L b$ L-order, 272
 $a \leq_R b$ R-order, 272
 $a \leq_J b$ J-order, 272
 $a \equiv_L b$ L-relation, 272
 $a \equiv_R b$ R-relation, 272
 $a \equiv_J b$ J-relation, 272

Chapter VI

ML modal logic, 303
 CTL computation tree logic, 303
 PDL propositional dynamic logic,
 303
 EF reachability fragment of
 CTL, 303
 MPL monadic path logic, 303
 MCL monadic chain logic, 303
 cEF counting variant of EF, 303
 cCTL * counting variant of CTL * ,
 303
 cwCTL * counting variant of wCTL * ,
 303
 cPDL counting variant of PDL, 303
 cTL $[\mathcal{K}]$ counting \mathcal{K} -temporal logic,
 304
 \models_{f} satisfiability for forests, 304
 \models_{t} satisfiability for trees, 304
 TL $[\mathcal{K}]$ \mathcal{K} -temporal logic, 305
 \mathbb{P}^{fb} finitely branching forests, 305
 $t[\tilde{K}]$ labelling of t , 307
 $L[\tilde{K}]$ cascade operation, 307

Casc(\mathcal{K}) cascade languages, 307
 a^ω ω -power, 314
 a^π idempotent power, 314
 $\pi \times a$ idempotent multiple, 314
 $a \approx_n^m t$ cEF $_n$ -bisimulation, 318
 $\text{tp}_n^m(s)$ EF $_n$ -type of s , 319
 A^Δ ξ -tuples, 334
 B^A set of functions, 334
 $t \downarrow_v^x$ subforest of x -successors,
 334
 $t \leftrightarrow \bar{a}$ substituting values for
 variables, 334
 $t \triangleleft \varphi$ φ -labelling of t , 335
 $\alpha \triangleleft \beta$ cascade composition, 335
 $\mathfrak{Q} \circ \mathfrak{B}$ wreath product, 335
 pr_w prefix corresponding to w ,
 348

\mathfrak{S}^1 adjoining a unit to a
 semigroup, 353
 Ψ set of characters, 354
 $\chi(w)$ character of w , 354
 $\chi(w/\sigma)$ relative character, 354
 sf_v subforest attached to v , 354
 $\Pi_N(s)$ number of paths, 355
 $\Pi_N^+(s)$ number of finite paths, 355
 $\Pi_N(s/H)$ number of paths through H ,
 355
 $\Pi_N^+(s/H)$ number of finite paths
 through H , 355
 $s \sim_\sigma t$ equivalent w.r.t. a character,
 362
 $c \approx_\sigma d$ equivalent w.r.t. σ , 365
 $[c]_\sigma$ \approx_σ -class, 365
 $\Psi_{\{x\}}^{(\sigma)}$ σ -significant characters, 365
 $\text{pf}(s)$ profile of s , 365

Index

- accepting run, 285
- algebraic signature, 36
- alphabet, 145
- φ -annotation, 335
- anti-chain logic, 265
- antiregular tree, 278
- aperiodic semigroup, 353
- arctic semiring, 135
- arity of a polynomial functor, 22
- arity, 253
- arity of a function symbol, 36
- arity of a signature, 36
- automaton, 284
- Automaton-Pathfinder Game, 286
- axiomatising a class of algebras, 214
- back-and-forth argument, 261
- δ -bialgebra, 81
- bisimilar forests, 309
- bisimulation, 309
- bisimulation-invariance, 309
- Bojańczyk-Klin algebra, 281
- canonical subobject diagram, 107
- cascade composition, 335
- category with congruences, 120
- character of a word, 354
- close end, 323
- closure of morphisms pullback, 99
- closure of morphisms pushouts, 99
- closure of morphisms under (co-)limits, 99
- cogenerator, 140
- comma category, 8
- complete binary tree, 299
- \mathbb{M} -compositional logic, 227
- composition of contexts, 162
- computation tree logic, 303
- congruence, 119
- \mathbb{M} -congruence, 125
- congruence ordering, 123
- context, 161
- convex set, 323
- counting monadic second-order logic, 260
- counting \mathcal{K} -temporal logic, 304
- counting temporal logics, 303
- definable algebra, 233
- definable class, 221
- definable language, 226
- definably embedded subset, 233
- dense morphism of monads, 175
- densely antiregular tree, 278
- derivative of a language, 162

descending chain condition, 138
 deterministic rooted graphs, 254
 diagonal fill-in property, 94
 discrete category, 10
 disjunctive category, 10
 distributive \mathbb{F} -algebra, 348
 δ -distributive function, 88
 distributive law, 66

E-morphism, 93
 x -edge, 254
 EF-bisimulation, 318
 embedding, 93
 U-end, 323
 enough injectives, 136
 L-equivalence, 220
 essentially finitary, 288
 essentially finitary monad, 175
 essentially unique factorisation, 108
 existence of syntactic algebras, 164
 expansion, 117
 Ω -extension of a logic, 221
 extension of a monad, 67

EM-factorisation, 93
 factorisation system, 93
 factorising through a cocone, 108
 family of languages, 146
 family of logics, 226
 far end, 323
 finitary algebra, 137
 finitary functor, 156
 finitary signature, 36
 finite-dimensional, 137
 finitely branching forest, 305
 first-order logic, 259
 flattening, 34

 flattening of a graph, 256
 forest, 254
 forest automaton, 284
 forest formula, 304
 formula, 219
 free algebra, 43
 free functor, 67
 function symbol, 36

 \mathfrak{G} -set, 7
 generalised sum, 262
 \mathcal{P} -generated object, 131
 κ -*M*-generated object, 108
 κ -*M*-generating diagram, 108
 generator, 128
 graph structures, 259
 Green's relations, 272

 H-class, 272
 hole, 161, 254
 horizontal product, 263
 horizontal semigroup of a forest algebra,
 263

 idempotent power, 314
 image of a morphism, 94
 MI-inequality, 214
 injective object, 136
 σ -insignificant element, 362
 internal operation, 161
 interpretation, 262
 inverse image, 146, 147
 inverse morphism, 146
 inverse relabelling, 146
 iteration of a language, 271

 J-class, 272

J-homogeneous word, 272

 kernel of a function in Pos, 123
 kernel of a function in Set, 122
 kernel of a morphism, 94
 Kleisli category, 67

 L-class, 272
 Φ -labelling, 304
 Ω -language, 146
 language algebra, 162
 language recognised by an automaton,
 285
 Luchli-Leonard operation, 267
 lift of a mond, 67
 linear forest, 254
 linear tree, 254
 locally definable algebra, 233
 logic, 219
 logic over an alphabet, 226

M-morphism, 93
 minimal algebra of a language, 151
 minimal morphism of a language, 151
 modal logic, 303
 modal rank, 305
 model, 219
 models of a formula, 219
 monadic chain logic, 303
 monadic path logic, 303
 monadic second-order logic, 259
 monadic second-order logic for linear
 orders, 266
 monadic second-order logic with forest
 ordering, 260
 monadic second-order logic with
 successor, 260

 monadic second-order logic with
 successor ordering, 260
 morphism of δ -bialgebras, 82
 morphism of language algebras, 162
 morphism of logics, 222
 morphism of polynomial functors, 26

 neutral element, 353
 noetherian, 138
 non-linear forest, 254
 non-linear tree, 254
 nondeterministic rooted graph, 253

 Ω -operation, 136
 orbit finite set, 144
 orbit-finite, 7
 origin map, 31

 parity condition, 285
 partial run of an automaton, 287
 +-path algebra, 355
 \mathcal{S} -path language, 353
 ω -path algebra, 355
 permutation of a forest, 308
 place in a forest, 286
 polynomial functor, 22
 power operator, 19
 x -predecessor, 254
 preserving a limit/colimit, 110
 priority function, 284
 σ -profile of a forest, 365
 profile of a partial run, 287
 projection, 93
 projective object, 127
 propositional dynamic logic, 303
 pseudo-variety, 132

quasi-finite set, 144
 quotient, 119

 \mathcal{R} -class, 272
 ranked tree, 254
 recognising a language, 146
 recognising a language by an automaton, 285
 reduced path, 354
 Δ -reduct, 117
 reduct of an \mathbb{M} -algebra, 116
 reflecting a limit/colimit, 111
 regular expression for languages of countable chains, 271
 regular forest, 287
 regular linear order, 268
 relabelling, 146
 relative character, 354
 root, 253
 rooted graph, 253
 run of an automaton, 285

 satisfaction function, 219
 satisfying an \mathbb{M} -inequality, 214
 σ -separated forest, 362
 shape map, 31
 shuffle, 267
 sibling-commutative forest language, 308
 Sierpiński space, 142
 signature, 36
 σ -significant element, 362
 simple interpretation, 262
 singleton operation, 34
 sort, 36
 sort-accumulation point, 183
 sort-wise finite family of logics, 226
 sort-wise properties, 137

 standard factorisation system, 102
 starting state of a partial run, 287
 state of an automaton, 284
 strong morphism of logics, 222
 strongly \mathbb{M} -compositional logic, 249
 strongly definable class, 221
 strongly finite-dimensional, 137
 Σ -structure, 39
 subalgebra, 104
 subobject, 104
 x -successor, 254
 syntactic algebra, 164
 syntactic congruence, 164
 syntactic morphism, 163
 syntactic partition, 304
 syntactic residue, 164

 \mathcal{K} -temporal logic, 305
 term algebra, 43
 theory, 219
 \mathbb{M} -theory of a class, 214
 theory function, 219
 transition relation, 284
 transitive closure, 153
 tree, 254
 tree formula, 304
 tropical semiring, 135
 type of a function symbol, 36

 universe, 39
 unranked tree, 254
 unravelling, 255

 Ω -valued logic, 219
 variable, 253
 varietal logic, 226
 variety, 132

\mathcal{R} -variety, 183
 variety of languages, 183
 vertical product, 263
 vertical semigroup of a forest algebra, 263

 weak coequaliser, 52
 weak morphism of logics, 222
 weakly definable class, 221
 weakly finite-dimensional set, 137
 well-defined factorisation system, 102
 wreath product, 335

The Roman and Fraktur alphabets							
<i>A</i>	<i>a</i>	𝔐	𝔞	<i>N</i>	<i>n</i>	ℛ	ℓ
<i>B</i>	<i>b</i>	𝔖	𝔞	<i>O</i>	<i>o</i>	𝔐	𝔞
<i>C</i>	<i>c</i>	𝔐	𝔠	<i>P</i>	<i>p</i>	ℙ	ℙ
<i>D</i>	<i>d</i>	𝔐	𝔡	<i>Q</i>	<i>q</i>	ℚ	ℚ
<i>E</i>	<i>e</i>	𝔐	𝔢	<i>R</i>	<i>r</i>	ℛ	ℛ
<i>F</i>	<i>f</i>	𝔐	𝔣	<i>S</i>	<i>s</i>	𝔖	𝔖
<i>G</i>	<i>g</i>	𝔐	𝔤	<i>T</i>	<i>t</i>	𝔐	𝔞
<i>H</i>	<i>h</i>	𝔐	𝔥	<i>U</i>	<i>u</i>	ℒ	ℒ
<i>I</i>	<i>i</i>	𝔐	𝔦	<i>V</i>	<i>v</i>	℔	℔
<i>J</i>	<i>j</i>	𝔐	𝔧	<i>W</i>	<i>w</i>	℔	℔
<i>K</i>	<i>k</i>	𝔐	𝔨	<i>X</i>	<i>x</i>	℔	℔
<i>L</i>	<i>l</i>	𝔐	𝔩	<i>Y</i>	<i>y</i>	℔	℔
<i>M</i>	<i>m</i>	ℳ	𝔞	<i>Z</i>	<i>z</i>	𝔖	𝔞

The Greek alphabet					
<i>A</i>	α	alpha	<i>N</i>	ν	nu
<i>B</i>	β	beta	<i>Ξ</i>	ξ	xi
<i>Γ</i>	γ	gamma	<i>O</i>	o	omicron
<i>Δ</i>	δ	delta	<i>Π</i>	π	pi
<i>E</i>	ϵ	epsilon	<i>P</i>	ρ	rho
<i>Z</i>	ζ	zeta	<i>Σ</i>	σ	sigma
<i>H</i>	η	eta	<i>T</i>	τ	tau
<i>Θ</i>	θ	theta	<i>Υ</i>	υ	upsilon
<i>I</i>	ι	iota	<i>Φ</i>	ϕ	phi
<i>K</i>	κ	kappa	<i>X</i>	χ	chi
<i>Λ</i>	λ	lambda	<i>Ψ</i>	ψ	psi
<i>M</i>	μ	mu	<i>Ω</i>	ω	omega