Abstract Algebraic Language Theory

Achim Blumensath

Abstract Algebraic Language Theory

Achim Blumensath

BRNO 2024

Achim Blumensath
blumens@fi.muni.cz

This document was last updated 2024-07-21.
The latest version can be found at

www.fi.muni.cz/~blumens

Contents

A. Algebra 1

I Monads 3

1. An overview . . . 3
2. Discrete categories . . . 10
3. Polynomial functors . . . 23
4. Monads . . . 35
5. Eilenberg-Moore algebras . . . 44
6. Lifting monads . . . 65

II Algebra 93

1. Factorisations . . . 93
2. Subalgebras . . . 110
3. Reducts . . . 122
4. Bialgebras . . . 129
5. Congruences . . . 142
6. Varieties . . . 150

B. Language Theory 157

III Languages 159

1. Weights . . . 159
2. Languages . . . 169
3. Minimal algebras . . . 174

4. Syntactic algebras 184
5. Varieties 205
6. The profinitary term monad 214
7. Axiomatisations 237

IV Logic 243
1. Abstract logics 243
2. Compositionality 251
3. Definable algebras 257
4. Definable languages 263

C. Applications 277

V Trees 279
1. Monads and logics for trees and forests 279
2. Finite forests 289
3. Countable chains 291
4. Counterexamples 304
5. MSO-definable algebras 310
6. First-order logic 324

VI Temporal Logics 331
1. Temporal logics 331
2. Bisimulation 336
3. The logic EF 342
4. Wreath products 362
5. Distributive algebras 376
6. Path algebras 381

D. Advanced Topics 403

VII Power Sets 405

1. Power-set functors 405
2. Linear monads 415
3. Closure under projection 430
4. Non-linear trees 432
5. Substitutions 463
6. Regular expressions 467
7. Distributive lattices 472

VIII Branch-Continuous Algebras 481
1. Sublinear trees 481
2. Semigroup-like algebras 485
3. Meet-distributive algebras 491
4. Trace algebras 493
5. Game algebras 502

Recommended Literature 515
Bibliography 517
Symbol Index 521
Index 525

Part A.

Algebra

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* $\mathfrak{M} = \langle M, \cdot, e\rangle$ is a structure with universe M, a binary operation $\cdot : M \times M \to M$, 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 $f : \Sigma \to M$, there exists a unique homomorphism $\varphi : \Sigma^* \to \mathfrak{M}$ that agrees with f on the elements of Σ.

A *homomorphism* $\varphi : \mathfrak{M} \to \mathfrak{N}$ of monoids is a function $\varphi : M \to N$ 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 $\varphi : \Sigma^* \to \mathfrak{M}$ 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 $\Sigma^* \to \Sigma^*/{\sim_L}$ recognises L. Hence, if $\sim_L$ has only finitely many classes, the quotient $\Sigma^*/{\sim_L}$ 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 $w : \omega \to \Sigma$. 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 $\mathfrak{S} = \langle S_1, S_\infty, \cdot, \pi \rangle$ is a two-sorted structure where S_1 contains the *finite* elements and S_∞ the *infinite* ones. These two sets are equipped with three kinds of multiplication:

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

All products are assumed to the associative. That is,

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

for all $a, b, c, a_0, a_1, \dots \in S_1$, $u \in S_\infty$, 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* $\varphi : \mathfrak{S} \to \mathfrak{T}$ of ω-semigroups consists of a pair of maps $\varphi_1 : S_1 \to T_1$ and $\varphi_\infty : S_\infty \to T_\infty$ that commute with all three products.

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

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

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

and the product is just the maximum operation.

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

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

and the product is defined by

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

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

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* $\mathrm{ar}(a) \in \omega$. 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

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

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 $\kappa : \Sigma^* \to \Omega$, where Ω is some set of *weights*. Frequently, one assumes that Ω forms a semiring. In this case the set of all weighted languages $\Sigma^* \to \Omega$ also forms a semiring, which is usually denoted $\Omega\langle\!\langle \Sigma^* \rangle\!\rangle$ and called the semiring of *formal power-series* over Ω. It is then customary to write a function $\kappa : \Sigma^* \to \Omega$ 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 $\alpha : \mathfrak{G} \to \mathrm{Aut}(X)$. Such sets are called $\mathfrak{G}$*-set*, where $\mathfrak{G}$ is the group in question. Since an action is just a homomorphism $\alpha : \mathfrak{G} \to \mathrm{Aut}(X)$, the category $\mathfrak{G}$-Set of all $\mathfrak{G}$-sets is just the comma category $(\mathfrak{G} \downarrow \mathrm{Aut})$. 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 $\mathfrak{G}$-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

$$\prod_{i \in I} f_i : \prod_{i \in I} A_i \to \prod_{i \in I} B_i\,, \qquad \sum_{i \in I} f_i : \sum_{i \in I} A_i \to \sum_{i \in I} B_i\,,$$
$$\langle g_i \rangle_{i \in I} : A \to \prod_{i \in I} B_i\,, \qquad \bigoplus_{i \in I} h_i : \sum_{i \in I} A_i \to B\,.$$

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 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 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 Set the category of all sets and functions, and by Pos the category of all partial orders and monotone functions. Top denotes the category of topological spaces and continuous maps, and Met is the category of bounded metric spaces and non-expansive functions, that is, metric spaces $\langle X, d \rangle$ satisfying

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

and functions $f : X \to Y$ satisfying

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

Finally, for a group $\mathfrak{G}$, we denote by $\mathfrak{G}$-Set the category of all $\mathfrak{G}$-sets, i.e., sets S equipped with an action $\mathfrak{G} \times S \to S$ 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 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 $X \in \mathsf{Set}$ with the 'weakest possible' $\mathcal{D}$-structure $\mathbb{J}X$. For instance, for $\mathcal{D} = \mathsf{Pos}$, $\mathbb{V}A$ 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 $\mathcal{D} = \mathsf{Top}$, $\mathbb{V}$ 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 $\mathbb{V} : \mathcal{D} \to \mathsf{Set}$ 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 $\mathcal{D}$. (Not every function is monotone/continuous/etc.) Furthermore, $\mathbb{V}$ induces a partial order $\sqsubseteq$ on the objects of $\mathcal{D}$ by

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

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 $\mathcal{D} = \mathsf{Pos}$, we have

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

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

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

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

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

(b) A category $\mathcal{D}$ is *discrete*[1] if it is disjunctive, has arbitrary coproducts, and there exists an adjunction $\mathbb{J} \dashv \mathbb{V}$ between Set and $\mathcal{D}$ such that $\mathbb{V}$ is faithful and preserves coproducts, ⌟

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

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

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

[1] Note that the term 'discrete category' is used differently in the literature.

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

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

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

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

A limit $\lim D \subseteq \prod_i D(i)$ is computed as in 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 $\mathfrak{G}$-Set, we define the adjunction $\mathbb{J} \dashv \mathbb{V}$ by

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

where

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

The corresponding bijection maps $f : X \to Y$ in Set to

$$\langle \mathrm{G} \times X, \gamma\rangle \to \langle Y, \alpha\rangle : \langle g, x\rangle \mapsto \alpha(g)(f(x))$$

in $\mathfrak{G}$-Set. The unit $\varepsilon : X \to G \times X$ and the counit $\iota : \langle G \times X, \gamma \rangle \to \langle X, \alpha \rangle$ 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 disjunctiveness 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 $\varepsilon : X \to X$ and the counit $\iota : \langle X, d_1 \rangle \to \langle X, d \rangle$ 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 disjunctiveness is again similar to that of Set. ⌟

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 $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ be the unit and counit of the adjunction.*

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

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

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

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

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

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

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

(b) By definition, $\iota : \mathbb{J}\mathbb{V}A \to A$ is the image of $\mathrm{id} : \mathbb{V}A \to \mathbb{V}$ 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 $g, h : C \to A$ 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 $s, t : 1 \to \mathbb{V}A$ with $s \neq t$ but $\mathbb{V}f \circ s = \mathbb{V}f \circ t$. Let $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ 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 $\mathbb{J}1 \to A$ 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 $g, h : B \to C$ 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. □

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

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

Let us prove the interesting direction. Suppose that $\operatorname{rng} e$ is dense in $\mathfrak{Y}$ and consider two different continuous maps $g, h : \mathfrak{Y} \to \mathfrak{Z}$. 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 $U \cap \operatorname{rng} e \neq \varnothing$. Consequently, $g \circ e \neq h \circ e$ and e is not an epimorphism. ⌟

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 $X \in \mathsf{Set}$. ⌟

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 $\sigma : A \to \Xi$ 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 $\mathcal{D}^\Xi$ whose objects are families $A = (A_\xi)_{\xi \in \Xi}$ where each A_ξ is an object of $\mathcal{D}$. A morphism $f : A \to B$ between such families is then just a family $f = (f_\xi)_\xi$ of morphisms $f_\xi : A_\xi \to B_\xi$. Similarly, a functor $\mathbb{F} : \mathcal{C} \to \mathcal{D}^\Xi$ is given by a family $\mathbb{F} = (\mathbb{F}_\xi)_\xi$ of functors $\mathbb{F}_\xi : \mathcal{C} \to \mathcal{D}$. 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 $\mathcal{D}^\Xi$. 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 $f : A \to B$ 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 $\mathcal{D}^\Xi$ 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 $X \in \mathsf{Set}$ and $A \in \mathcal{D}$.

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

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

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

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

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

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

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 $\zeta \in \Xi$,

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

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

To see that $\mathbb{F}$ is full, consider a morphism $f : \alpha \to \beta$ 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) $f_\xi : A_\xi \to B_\xi$. Hence, $f = \mathbb{F}(f_\xi)_\xi$.

It remains to show that every object $\alpha : A \to \Xi \cdot 1$ of $(\mathcal{D} \downarrow \Xi \cdot 1)$ is isomorphic to one of the form $\mathbb{F}(A_\xi)_\xi$. Hence, fix $\alpha : A \to \Xi \cdot 1$. As $\mathcal{D}$ is disjunctive, we have $A = \sum_\xi A_\xi$ and $\alpha = \sum_\xi \alpha_\xi$, for suitable A_ξ and $\alpha_\xi : A_\xi \to 1$. 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 $\mathcal{D}^\Xi$. 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 $f : I \to K$ a function in Set. We define functors $\square_f : \mathcal{C}^K \to \mathcal{C}^I$ and $\Sigma_f, \Pi_f : \mathcal{C}^I \to \mathcal{C}^K$ by

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

and similarly for morphisms. ⌟

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 $(\mathcal{C} \downarrow K \cdot 1) \to (\mathcal{C} \downarrow I \cdot 1)$. Prove that this functor maps a morphism $\alpha : C \to \mathbb{J}K$ to its pullback along $\mathbb{J}f : \mathbb{J}I \to \mathbb{J}K$. ⌟

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 $f : I \to K$ be a function in* Set.

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

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

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

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

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

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

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

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

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

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

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

Let us also mention the following commutation relation between them.

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

$$\Box_f \Sigma_g \cong \Sigma_u \Box_v \quad \text{and} \quad \Box_f \Pi_g \cong \Pi_u \Box_v \,. \qquad \lrcorner$$

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 $f : I \to J$ in* Set, *the following two conditions hold.*

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

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

(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 $X \in$ 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 $X \to \mathbb{V}A$ 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 $X \to \mathbb{V}A$ 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 $X \in \mathsf{Set}$, we define the *power operator*

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

Similarly, for $A \in \mathcal{D}^\Xi$ and $X \in \mathsf{Set}^\Xi$, 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 $f : A \to B$ and $g : X \to Y$, we define

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

where $p_y : \prod_{y \in Y} A \to A$ 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 $\mathsf{Set}_\kappa \subseteq \mathsf{Set}$ be the full subcategory of all sets of size less than κ. Then*

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

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

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

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

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

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

where $p_x : \prod_{z \in X} B \to B$ 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 $h : Y \to X$, we have

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

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

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

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

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

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) $\mathbb{V}(A^X) \cong \mathsf{Set}(X, \mathbb{V}A)$.

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

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

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

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

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 $\Xi = \{1, \infty\}$ where sort 1 represents the 'finite' elements and sort ω the 'infinite' ones. The operation $\mathbb{M}$ then maps a Ξ-sorted set $A = \langle A_1, A_\infty \rangle$ to $\mathbb{M}A = \langle \mathbb{M}_1 A, \mathbb{M}_\infty A \rangle$ where

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

(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 $f : A \to B$ induces an operation $\mathbb{M}f : \mathbb{M}A \to \mathbb{M}B$ which applies the function f to each label. This turns $\mathbb{M}$ into a functor. In this section, we introduce the kind of functors $\mathbb{M} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ 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 $\mathbb{F} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ is *polynomial* if it is of the form

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

for fixed sets $I_\xi \in \mathsf{Set}$ and $D^i_\xi \in \mathsf{Set}^\Xi$. Usually, we will use the notation

$$\mathbb{F}X = \sum_{i \in I} X^{\mathrm{dom}(i)},$$

where $I = (I_\xi)_\xi$ and $\mathrm{dom}(i) := D^i_\xi$.

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

(c) If $\mathbb{F}$ is a polynomial functor, we call the elements of $\mathbb{F}A$ *terms*. ⌟

Example. (a) The word functor $\mathbb{F}X := X^+$ is a polynomial functor on 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 Pos^2. ⌟

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 $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ with

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

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

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

and generally

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

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

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

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

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 $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ there exists a polynomial functor $\mathbb{F}^{\circ} : \mathsf{Set}^{\Xi} \to \mathsf{Set}^{\Xi}$ such that*

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

Proof. Suppose that

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

If we use the same expression

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

in the category Set^{Ξ}, we obtain

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

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. □

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

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

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 $f : A \to B$ a morphism. First, suppose that $f : A \to B$ 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 $v \in \mathrm{dom}(s)$. Since f is injective, it follows that $s(v) = t(v)$, for all v. Hence, $s = t$.

Next, suppose that $f : A \to B$ 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 $g : \mathbb{V}B \to \mathbb{V}A$. 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 monomorphisms if, and only if, it is injective. $\square$

Remark. Every functor $\mathbb{F} : \mathsf{Set} \to \mathsf{Set}$ preserves epimorphisms since epimorphisms in Set have right inverses. ⌟

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

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

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

Proof. We have

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

□

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

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

and the domains are

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

⌟

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 $\mathcal{C} \to \mathcal{C}$ and their morphisms by $\mathsf{Poly}(\mathcal{C})$. ⌟

We would like to obtain a more concrete description of such morphisms. To do so, we compare $\mathsf{Poly}(\mathcal{C})$ 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 $(A_i)_{i\in I} \to (B_j)_{j\in J}$ is a pair $\langle f, (\varphi_j)_{j\in J}\rangle$ consisting of a function $f : J \to I$

and a family of morphisms $\varphi_j : A_{f(j)} \to B_j$, for $j \in J$. The composition of two such morphisms is defined by

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

There are obvious maps between $\Pi(\mathsf{Set})$ and $\mathsf{Poly}(\mathcal{C})$ 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*

$$\mathbb{I} : \Pi(\mathsf{Set})^{\mathrm{op}} \to \mathsf{Poly}(\mathcal{D})$$

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 $\Pi(\mathsf{Set})$ and let $\mathbb{F}_A := \mathbb{I}A$ and $\mathbb{F}_B := \mathbb{I}B$ be the associated polynomial functors. We map a morphism $\langle f, (\varphi_i)_i \rangle : B \to A$ to the natural transformation $\tau : \mathbb{F}_A \Rightarrow \mathbb{F}_B$ defined by

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

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 $g : X \to Y$. 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 $\langle f, (\varphi_i)_i \rangle : B \to A$ and $\langle g, (\psi_j)_j \rangle : C \to B$. 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

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

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 $\langle f, (\varphi_i)_i\rangle : B \to A$ and $\langle g, (\psi_i)_i\rangle : B \to A$ 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). □

In general, the embedding functor $\mathbb{I} : \Pi(\mathsf{Set})^{\mathrm{op}} \to \mathsf{Poly}(\mathcal{D})$ is not full. Hence, there can be natural transformations that do not correspond to a morphism of $\Pi(\mathsf{Set})$. To get an idea of how these addional 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 $\alpha : (-)^D \Rightarrow (-)^E$ is of the form*

$$\alpha(s) = \mathbb{V}(\iota \circ \mathbb{J}s) \circ \alpha_\circ\,, \quad \textit{for } s : D \to \mathbb{V}X\,,$$

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

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

$$\mathbb{V}\alpha_{\mathbb{J}D} : \mathbb{V}\big((\mathbb{J}D)^D\big) \to \mathbb{V}\big((\mathbb{J}D)^E\big)$$

induces a function

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

We set

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

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

To show that α is of the required form, let $A \in \mathcal{D}$ and $s \in \mathsf{Set}(D, \mathbb{V}A) \cong \mathbb{V}A^D$. 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

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

This implies that

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

□

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 $\mathfrak{G}$-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 $\tau : (-)^D \Rightarrow (-)^E$ on $\mathfrak{A}$-Set are precisely the functions of the form

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

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

(b) Let $\mathfrak{G}$ be a non-abelian group, fix two elements $g, h \in G$ with $gh \neq hg$. The function $\tau : G^D \to G^D$ 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 $\mathfrak{G}$-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 $\langle f, (\varphi_i)_i \rangle : (E_j)_j \to (\mathbb{V}\mathbb{J}D_i)_i$ of $\Pi(\mathsf{Set})$ such that*

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

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

Proof. Let $\tau : \mathbb{F} \Rightarrow \mathbb{G}$ be a natural transformation. We start by recovering the function $f : I \to J$. Let 1 be the terminal object of $\mathcal{D}^{\Xi}$. 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, $\mathbb{V}\tau_1 : \mathbb{V}\mathbb{F}1 \to \mathbb{V}\mathbb{G}1$ induces a function $f : I \to J$. Given some object A, let $u : A \to 1$ 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

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

It thus remains to construct the functions $f_i' : E_{f(i)} \to \mathbb{V}\mathbb{J}D_i$. 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

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

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

Remark. The above proposition yields a function

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

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

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

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

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

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

Categories satisfying the conditions of this theorem include Set, Pos, Top, and Met, but not $\mathfrak{G}$-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 $f : I \to J$ a function in* Set*. The functor $\Sigma_f : \mathcal{D}^I \to \mathcal{D}^J$ preserves connected limits.*

Proof. Let $D : \mathcal{K} \to \mathcal{D}^I$ 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 $\psi_k : C_k \to D(k)$ such that

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

For every morphism $h : k \to l$ 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

$$C_k = C_l \quad \text{and} \quad Dh \circ \psi^n_k = \psi^n_l \,.$$

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 $\sigma : C \to A$ 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. □

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

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

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

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. ⌟

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

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

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

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

⌟

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 $\mathrm{flat} : \mathbb{M}\mathbb{M}A \to \mathbb{M}A$, that is, it takes an $\mathbb{M}A$-labelled object $s \in \mathbb{M}\mathbb{M}A$ and assembles the appearing labels into a single large object. We call $\mathrm{flat}(s)$ the *flattening* of s. Secondly, there is usually a *singleton operation* $\mathrm{sing} : A \to \mathbb{M}A$ 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 $\mathrm{flat} : (A^*)^* \to A^*$ is simply the concatenation operation and $\mathrm{sing} : A \to A^*$ produces 1-letter words.

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

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 $\mathbb{M} : \mathcal{C} \to \mathcal{C}$ that is equipped with two natural transformations: a *multiplication* $\mathrm{flat} : \mathbb{M} \circ \mathbb{M} \Rightarrow \mathbb{M}$ and a *unit morphism* $\mathrm{sing} : \mathrm{Id} \Rightarrow \mathbb{M}$ (where Id is the identity functor), that satisfying the following equations.

$$\mathrm{flat} \circ \mathrm{sing} = \mathrm{id}\,, \qquad \mathrm{flat} \circ \mathbb{M}\mathrm{sing} = \mathrm{id}\,, \qquad \mathrm{flat} \circ \mathrm{flat} = \mathrm{flat} \circ \mathbb{M}\mathrm{flat}\,.$$

$$\begin{tikzcd} \mathbb{M}A \arrow[r, "\mathrm{sing}"] \arrow[dr, "\mathrm{id}"'] & \mathbb{M}\mathbb{M}A \arrow[d, "\mathrm{flat}"] & \mathbb{M}A \arrow[l, "\mathbb{M}\mathrm{sing}"'] \arrow[dl, "\mathrm{id}"] \\ & \mathbb{M}A & \end{tikzcd} \qquad \begin{tikzcd} \mathbb{M}\mathbb{M}\mathbb{M}A \arrow[r, "\mathrm{flat}"] \arrow[d, "\mathbb{M}\mathrm{flat}"'] & \mathbb{M}\mathbb{M}A \arrow[d, "\mathrm{flat}"] \\ \mathbb{M}\mathbb{M}A \arrow[r, "\mathrm{flat}"'] & \mathbb{M}A \end{tikzcd}$$

⌟

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 $\bigcup : \wp(\wp(X)) \to \wp(X)$ 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 VII.1).

⌟

Examples. (a) Fix an infinite cardinal κ. The *chain monad* is defined as follows. For $A \in \mathsf{Set}$, 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 $\mathrm{flat} : \mathbb{C}_\kappa \mathbb{C}_\kappa A \to \mathbb{C}_\kappa A$ maps a chain $(\mathfrak{A}_i)_{i \in I}$ of linear orders to their ordered sum $\sum_{i \in I} \mathfrak{A}_i$. As $\kappa^2 = \kappa$, the resulting order still has size less than κ. Let us check one of the monad laws.

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

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 $\mathrm{flat}(\mathbb{C}_\kappa \mathrm{flat}(((\mathfrak{A}^i_j)_{j \in J_i})_{i \in I}))$ take the form

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

while those of $\mathrm{flat}(\mathrm{flat}(((\mathfrak{A}^i_j)_{j \in J_i})_{i \in I}))$ have the form

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

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

$$\mathrm{flat}(\mathbb{C}_\kappa \mathrm{flat}(((\mathfrak{A}^i_j)_{j \in J_i})_{i \in I})) = \mathrm{flat}(\mathrm{flat}(((\mathfrak{A}^i_j)_{j \in J_i})_{i \in I}))\,.$$

(b) Every subclass of linear orders that is closed under flat induces a corresponding submonad of $\mathbb{C}_\kappa$. 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 $\mathrm{Lin}_\mathfrak{S}$ mapping a set X to the $\mathfrak{S}$-module $\mathfrak{S}[X]$ generated by X. That is, $\mathrm{Lin}_\mathfrak{S} X$ consists of all finite linear combinations

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

The corresponding flattening operation maps a nested linear combination

$$s(r_0 x_0 + \dots r_{m-1} x_{m-1}) + \dots + t(p_0 y_0 + \dots + p_{n-1} y_{n-1})$$

to

$$s r_0 x_0 + \dots s r_{m-1} x_{m-1} + \dots + t p_0 y_0 + \dots + t p_{n-1} y_{n-1} \,.$$ ⌟

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

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

and multiplication

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

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 $\bar\eta \to \xi$, where $\xi \in \Xi$ and $\bar\eta \in \Xi^\kappa$, 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 $X \in \mathsf{Set}^\Xi$ a set of *variables*. The set $\mathrm{Term}[\Sigma, X] \in \mathsf{Set}^\Xi$ 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 $\bar\eta \to \xi$ and, for each index i, $t_i \in \mathrm{Term}_{\eta_i}[\Sigma, X]$ is a Σ-term of sort η_i, then $f(\bar t)$ is a Σ-term of sort ξ.

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

$$\mathrm{Term}[\Sigma, f] : \mathrm{Term}[\Sigma, X] \to \mathrm{Term}[\Sigma, Y]$$

the function replacing every variable $x \in X$ in a given term $t \in \mathrm{Term}[\Sigma, X]$ 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 $11 \to 1$, the mixed one $\times$ has the type $1\omega \to \omega$, and the infinite product π the type $111\ldots \to \omega$. An example of a term is

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

⌟

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

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

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

Definition 4.3. Let Σ be a signature.

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

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

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

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

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

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

⌟

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

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

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

Remark. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ 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

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

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

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

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

$$\begin{array}{ll} g : 1 \to I & f : K \to I \\ \psi_* : D_{g(*)} \to \mathbb{VJ}1 & \varphi_k : D_{f(k)} \to \mathbb{VJ}E_k\,, \quad \text{for } k \in K\,. \end{array}$$

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

$$\varphi_s : \mathrm{dom}(\mu(s)) \to \sum_{v \in \mathrm{dom}(s)} \mathbb{VJ}\,\mathrm{dom}(s(v))\,, \quad \text{for } s \in \mathbb{MM}A\,.$$

⌟

Let us introduce some useful terminology and notation.

Definition 4.5. Let $\mathbb{M}$ be a polynomial monad. A *factorisation* of a term $t \in \mathbb{M}A$, is a term $T \in \mathbb{MM}A$ such that $t = \mathrm{flat}(T)$. In this case, we call each term $T(v)$, for $v \in \mathrm{dom}(T)$, a *factor* of the factorisation. ⌟

Definition 4.6. Let $\mathbb{F} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ be a functor.

(a) The *lift* of a relation $\theta \subseteq A \times B$ is the relation $\theta^{\mathbb{F}} \subseteq \mathbb{F}A \times \mathbb{F}B$ defined by

$$s\ \theta^{\mathbb{F}}\ t \quad :\text{iff} \quad \mathbb{F}p(u) = s \text{ and } \mathbb{F}q(u) = t\,, \quad \text{for some } u \in \mathbb{F}\theta\,,$$

where $p : A \times B \to A$ and $q : A \times B \to B$ are the two projections.

(b) Two terms $s \in \mathbb{F}A$ and $t \in \mathbb{F}B$ *have the same shape* if they are related via the lift $\simeq_{\mathrm{sh}} := (A \times B)^{\mathbb{F}}$. ⌟

Examples. (a) For a polynomial functor $\mathbb{F}X = \sum_{i \in I} X^{D_i}$ and $s \in \mathbb{F}A$, $t \in \mathbb{F}B$, we have

$$s \simeq_{\mathrm{sh}} t \quad \text{iff} \quad s \in A^{D_i} \quad \text{and} \quad t \in B^{D_i}\,, \quad \text{for the same index } i \in I\,.$$

(This implies that s and t have the same sort, namely that of i.) Then

$$s\ \theta^{\mathbb{F}}\ t \quad \text{iff} \quad s \simeq_{\mathrm{sh}} t \quad \text{and} \quad s(v)\ \theta\ t(v)\,, \quad \text{for all } v \in \mathrm{dom}(s)\,.$$

(b) Two words $u, v \in A^*$ have the same shape if their length is the same.

(c) Given $s \in \mathbb{F}A$ and $t \in \mathbb{F}\wp(A)$, we have

$$s \in^{\mathbb{F}} t \quad \text{iff} \quad s \simeq_{\mathrm{sh}} t \quad \text{and} \quad s(v) \in t(v)\,, \quad \text{for all } v \in \mathrm{dom}(s)\,.$$

This relation will play an important role in Chapter VII. ⌟

Proving that something is a monad is often tedious. The following characterisation reduces the number of things we have to check.

Definition 4.7. Let $\mathcal{C}$ be a category. An *extension system* on $\mathcal{C}$ consisting of

- a function $F : \mathcal{C} \to \mathcal{C}$ on objects,
- a family of morphisms $\eta_A : A \to FA$, for $A \in \mathcal{C}$,
- a family of functions $-^F : \mathcal{C}(A, FB) \to \mathcal{C}(FA, FB)$, for $A, B \in \mathcal{C}$,

that satisfy the following axioms.

$$\begin{aligned} \eta^F &= \mathrm{id}\,, & \\ f^F \circ \eta &= f\,, & &\text{for all } f : A \to FB\,, \\ g^F \circ f^F &= (g^F \circ f)^F\,, & &\text{for all } f : A \to FB \text{ and } g : B \to FC\,. \end{aligned}$$ ⌟

Proposition 4.8. *Let $\mathcal{C}$ be a category. There exists a bijection between extension systems on $\mathcal{C}$ and monads on $\mathcal{C}$. This bijection maps an extension system $\langle F, \eta, -^F \rangle$ to the monad $\langle \mathbb{M}, \mu, \varepsilon \rangle$ where*

$$\begin{aligned} \mu &:= \mathrm{id}^F\,, & \mathbb{M}A &:= FA\,, & &\textit{for } A \in \mathcal{C}\,, \\ \varepsilon &:= \eta\,, & \mathbb{M}f &:= (\eta \circ f)^F\,, & &\textit{for } f : A \to B\,. \end{aligned}$$

Conversely, a monad $\langle \mathbb{M}, \mu, \varepsilon \rangle$ is mapped to the extension system $\langle F, \eta, -^F \rangle$ where

$$FA := \mathbb{M}A\,, \qquad \eta := \varepsilon\,, \qquad f^F := \mu \circ \mathbb{M}f\,.$$

Proof. Let $\langle F, \eta, -^F \rangle$ be an extension system and define $\langle \mathbb{M}, \mu, \varepsilon \rangle$ as above. Then $\mathbb{M}$ is a functor since

$$\begin{aligned}
\mathbb{M}\mathrm{id} &= (\eta \circ \mathrm{id})^F = \eta^F = \mathrm{id}\,, \\
\mathbb{M}(g \circ f) &= (\eta \circ g \circ f)^F \\
&= \left((\eta \circ g)^F \circ \eta \circ f\right)^F \\
&= (\eta \circ g)^F \circ (\eta \circ f)^F = \mathbb{M}g \circ \mathbb{M}f\,.
\end{aligned}$$

Furthermore, μ and ε are natural since

$$\mathbb{M}f \circ \varepsilon = (\eta \circ f)^F \circ \eta = \eta \circ f\,,$$

$$\begin{aligned}
\mathbb{M}F \circ \mu &= (\eta \circ f)^F \circ \mathrm{id}^F \\
&= \left((\eta \circ f)^F \circ \mathrm{id}\right)^F \\
&= \left((\eta \circ f)^F\right)^F \\
&= (\mathbb{M}f)^F \\
&= (\mathrm{id} \circ \mathbb{M}f)^F \\
&= (\mathrm{id}^F \circ \eta \circ \mathbb{M}f)^F \\
&= \mathrm{id}^F \circ (\eta \circ \mathbb{M}f)^F = \mu \circ \mathbb{M}\mathbb{M}f\,.
\end{aligned}$$

Finally, the monad axioms hold since

$$\begin{aligned}
\mu \circ \varepsilon &= \mathrm{id}^F \circ \eta = \mathrm{id}\,, \\
\mu \circ \mathbb{M}\varepsilon &= \mathrm{id}^F \circ (\eta \circ \eta)^F \\
&= (\mathrm{id}^F \circ \eta \circ \eta)^F \\
&= (\mathrm{id} \circ \eta)^F \\
&= \eta^F \\
&= \mathrm{id}\,,
\end{aligned}$$

$$
\begin{aligned}
\mu \circ \mathbb{M}\mu &= \mathrm{id}^F \circ (\eta \circ \mathrm{id}^F)^F \\
&= (\mathrm{id}^F \circ \eta \circ \mathrm{id}^F)^F \\
&= (\mathrm{id} \circ \mathrm{id}^F)^F \\
&= (\mathrm{id}^F)^F \\
&= (\mathrm{id}^F \circ \mathrm{id})^F \\
&= \mathrm{id}^F \circ \mathrm{id}^F \\
&= \mu \circ \mu\,.
\end{aligned}
$$

Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a monad and define $\langle F, \eta, -^F \rangle$ as above. Then

$$
\begin{aligned}
\eta^F &= \mu \circ \mathbb{M}\varepsilon = \mathrm{id}\,, \\
f^F \circ \eta &= \mu \circ \mathbb{M}f \circ \varepsilon = \mu \circ \varepsilon \circ f = f\,, \\
g^F \circ f^F &= \mu \circ \mathbb{M}g \circ \mu \circ \mathbb{M}f \\
&= \mu \circ \mu \circ \mathbb{M}\mathbb{M}g \circ \mathbb{M}f \\
&= \mu \circ \mathbb{M}\mu \circ \mathbb{M}\mathbb{M}g \circ \mathbb{M}f \\
&= \mu \circ \mathbb{M}(\mu \circ \mathbb{M}g \circ f) \\
&= \mu \circ \mathbb{M}(g^F \circ f) \\
&= (g^F \circ f)^F\,.
\end{aligned}
$$

□

Corollary 4.9. *A monad $\langle \mathbb{M}, \mu, \varepsilon \rangle$ is uniquely determined by*

- *the object function $\mathbb{M} : \mathcal{C}^{\mathrm{obj}} \to \mathcal{C}^{\mathrm{obj}}$,*
- *the family of morphisms $\varepsilon = (\varepsilon_A)_A$, and*
- *the operation $-^*$ given by $f^* = \mu \circ \mathbb{M}f$, for $f : A \to \mathbb{M}B$.*

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

$\mathbb{M}\Sigma \to \mathfrak{A}$ 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* $A \in \mathsf{Set}^{\Xi}$ and
- one function $f^{\mathfrak{A}} : \prod_i A_{\eta_i} \to A_{\xi}$, for each function symbol $f \in \Sigma$ of type $\bar{\eta} \to \xi$. ⌟

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 $\cdot^{\mathfrak{S}} : S \times S \to S$. 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, $11 \to 1$, $1\infty \to \infty$, and $111\ldots \to \infty$. 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

$$\mathrm{Term}[\Sigma, X] \times A^X \to A\,,$$

where the second argument $\beta : X \to A$ 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

$$\mathrm{Term}[\Sigma, A] \to A\,.$$

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 $\pi : \mathbb{M}A \to A$. For instance, for words this product takes the form $\pi : A^+ \to A$, 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 $\pi : A^+ \to A$ 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 $\cdot : A \times A \to A$. 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

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

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 $\mathbb{M} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ 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 $\pi : \mathbb{M}A \to A$ satisfying

$$\pi \circ \mathbb{M}\pi = \pi \circ \mathrm{flat},$$
$$\pi \circ \mathrm{sing} = \mathrm{id}.$$

$$\begin{tikzcd}
\mathbb{M}\mathbb{M}A \arrow[r, "\mathrm{flat}"] \arrow[d, "\mathbb{M}\pi"'] & \mathbb{M}A \arrow[d, "\pi"] \\
\mathbb{M}A \arrow[r, "\pi"'] & A
\end{tikzcd}$$

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

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

$$\begin{tikzcd} \mathbb{M}A \arrow[r, "\mathbb{M}\varphi"] \arrow[d, "\pi"'] & \mathbb{M}B \arrow[d, "\pi"] \\ A \arrow[r, "\varphi"'] & B \end{tikzcd}$$

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

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

(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 $s \in \mathbb{M}C$. ⌟

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

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

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

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

The laws of an $\mathbb{M}$-algebra ensure that π_1 corresponds to a semigroup product $A_1 \times A_1 \to A_1$ and π_∞ correspond to the additional products $A_1 \times A_\infty \to A_\infty$ and $A_1^\omega \to A_\infty$ 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

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

and unit function

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

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 $\mathsf{Alg}(\mathbb{M})$ of $\mathbb{M}$-algebras is isomorphic to the category $\mathfrak{M}$-Set of $\mathfrak{M}$-sets.

(d) Given a semiring $\mathfrak{S}$, let $\mathbb{L}$ be the functor mapping a set X to the set of all finite linear combinations of elements in X with coefficients from $\mathfrak{S}$. Algebras for $\mathbb{L}$ are $\mathfrak{S}$-modules. ⌟

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 $\mathrm{flat} : \mathbb{M}\mathbb{M}A \to \mathbb{M}A$ 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 $\langle \mathbb{M}A, \mathrm{flat} \rangle$.*

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

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

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

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

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

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

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 algabra.* of the from $\mathrm{Term}[\Sigma, X]$, for some set X. Its elements are all terms and a function f maps a tuple of terms $\bar{t}$ to the term $f(\bar{t})$. ⌟

Corollary 5.4. *For every monad $\mathbb{M} : \mathcal{C} \to \mathcal{C}$ there exists an adjunction $\mathbb{F} \dashv \mathbb{U}$ between $\mathcal{C}$ and $\mathsf{Alg}(\mathbb{M})$ 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 $\langle \mathbb{M}A, \mathrm{flat} \rangle$, 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

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

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

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 $f : A \to B$ and $p : \mathbb{M}B \to B$ 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 $f : \mathfrak{A} \to \mathfrak{B}$ 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

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

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

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

Lemma 5.6. *Let $\varphi : \mathfrak{A} \to \mathfrak{B}$ and $\varepsilon : \mathfrak{A} \to \mathfrak{C}$ be $\mathbb{M}$-morphisms where $\mathbb{M}\varepsilon$ is an epimorphism. If $f : C \to B$ 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 $\pi \circ \mathbb{M}f = f \circ \pi$. □

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 $\mathsf{Alg}(\mathbb{M}) \to \mathcal{D}^{\Xi}$ reflects limits. Hence, $\mathsf{Alg}(\mathbb{M})$ has all limits that exist in $\mathcal{D}$.*

Proof. Suppose that $\mathcal{D}$ has limits with index category $\mathcal{I}$. Then so does $\mathcal{D}^{\Xi}$. (As limits in $\mathcal{D}^{\Xi}$ are computed componentwise.) Let $D : \mathcal{I} \to \mathsf{Alg}(\mathbb{M})$ be a diagram and let $\mathbb{U} : \mathsf{Alg}(\mathbb{M}) \to \mathcal{D}^{\Xi}$ 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 $\lambda_i : \mathfrak{A} \to D(i)$ is an $\mathbb{M}$-morphism. To define the product of $\mathfrak{A}$, we set

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

where $\pi_i : \mathbb{M}D(i) \to D(i)$ is the product of $D(i)$. For every morphism $f : i \to j$ 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 $\rho : \mathbb{M}A \to A$ 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

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

for all i. Limiting cones being jointly monomorphic, we obtain $\rho \circ \mathrm{sing} = \mathrm{id}$ and $\rho \circ \mathrm{flat} = \rho \circ \mathbb{M}\rho$. Consequently, $(\lambda_i)_i$ forms a cone from $\langle A, \rho \rangle$ to D

in $\mathsf{Alg}(\mathbb{M})$. 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 $\sigma : B \to A$ (in $\mathcal{D}^\Xi$) 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. $\square$

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 $\pi : \mathbb{M}A \to A$ is the coequaliser of* flat *and $\mathbb{M}\pi$ in* $\mathsf{Alg}(\mathbb{M})$.

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

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

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 $\varphi : \mathbb{M}A \to \mathfrak{B}$ is another $\mathbb{M}$-morphism such that $\varphi \circ \mathrm{flat} = \varphi \circ \mathbb{M}\pi$. The function $\sigma := \varphi \circ \mathrm{sing}$ satisfies

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

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

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

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

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

□

The following criterion greatly simplyfies existence proofs.

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

Proof. $(\Rightarrow)$ is trivial. For $(\Leftarrow)$, suppose that $\mathsf{Alg}(\mathbb{M})$ has coequalisers. Since a category is cocomplete if, and only if, it has coproducts and coequalisers, it is sufficient to show that $\mathsf{Alg}(\mathbb{M})$ 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 $s_i : A_i \to \sum_j A_j$ and $t_i : \mathbb{M}A_i \to \sum_j \mathbb{M}A_j$ 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 $w : \sum_j \mathbb{M}A_j \to \mathbb{M}\sum_j A_j$ such that

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

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

$$\mathbb{M}\textstyle\sum_i \pi_j,\ \mathrm{flat} \circ \mathbb{M}w : \mathbb{M}\sum_j \mathbb{M}A_j \to \mathbb{M}\sum_j A_j\,.$$

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

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

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

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

we can use Lemma 5.8 to find unique $\mathbb{M}$-morphisms $\lambda_i : \mathfrak{A}_i \to \mathfrak{C}$ 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 $u : \sum_j A_j \to B$ 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 $v : \sum_j \mathbb{M}A_j \to \mathbb{M}B$ such that

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

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

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

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

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

This implies that

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

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

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

Since limiting cocones are jointly epimorphic, it follows that

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

Consequently,

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

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

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

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

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

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

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

Hence,

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

as desired. □

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 $\mathsf{Alg}(\mathbb{M})$ is cocomplete.

(b) Show that every polynomial functor $\mathsf{Set} \to \mathsf{Set}$ preserves coequalisers. ⌟

Thus, we only have to show that the category $\mathsf{Alg}(\mathbb{M})$ 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 $\varphi, \psi : A \to B$ is a morphism $\rho : B \to C$ such that $\rho \circ \varphi = \rho \circ \psi$ and every morphism $\sigma : B \to D$ with $\sigma \circ \varphi = \sigma \circ \psi$ factorises (not necessarily uniquely) through ρ. ⌟

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 $U : (A \downarrow \mathcal{C}) \to \mathcal{C}$ be the functor mapping an object $f : A \to B$ to B and a morphism $\varphi : f \to g$ to the corresponding morphism between the codomains of f and g. It is sufficient to show that U creates limits. Hence, let $D : \mathcal{I} \to (A \downarrow \mathcal{C})$ 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 $h : i \to k$ of $\mathcal{I}$, the fact that Dh is a morphism $\mu_i \to \mu_k$ 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 $\sigma : A \to B$ such that

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

We can consider λ_i as a morphism $\lambda_i : \sigma \to \mu_i$ 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. □

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 $\varphi, \psi : A \to B$ with a weak coequaliser $q : B \to \mathrm{C}$. 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 $\sigma : B \to \mathrm{C}$ 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 $\mathcal{W} \to \mathcal{E}$ is final (which means that the inclusions $\mathcal{W} \to (B \downarrow \mathcal{C})$ and $\mathcal{E} \to (B \downarrow \mathcal{C})$ have the same limit).

Let $D : \mathcal{W} \to (B \downarrow \mathcal{C})$ be the inclusion functor. Since $\mathcal{W} \subseteq \mathcal{C}(B, \mathrm{C})$ is a set, D has a limit $\tau : B \to \mathrm{C}$. Let $\varepsilon : \rho \to \tau$ 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 $\rho \to \sigma$ we choose an arbitrary $v \in \mathcal{W}$. Let $\lambda : \tau \to v$ be the corresponding component of the limiting cone from τ to D. Since $v \in \mathcal{W}$ there also exists a morphism $f : v \to \sigma$. We therefore obtain a morphism

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

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 $f, g : \rho \to \sigma$. Let $e : v \to \rho$ 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 $h : \rho \to v$. Since $\varepsilon \circ e \circ h$ and id are both morphisms $\tau \to \tau$ and ε is the equaliser of all such morphisms, it follows that

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

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

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

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 $\mathsf{Alg}(\mathbb{M})$ 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 $\varphi, \psi : \mathfrak{A} \to \mathfrak{B}$. By Lemma 5.12 and since $\mathsf{Alg}(\mathbb{M})$ 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 $P, Q : \kappa + 1 \to \mathcal{C}$,
- 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 $\sigma_\kappa : P_* \to Q_*$ will the product morphism, and $q_\kappa : \mathfrak{B} \to Q_*$ will be the coequaliser morphism.

$$\begin{tikzcd}
\mathbb{M}A \arrow[r, shift left, "\mathbb{M}\varphi"] \arrow[r, shift right, "\mathbb{M}\psi"'] \arrow[dd, "\pi"'] & \mathbb{M}B \arrow[r, "\mathbb{M}q_i"] \arrow[dr, "p_i"'] \arrow[dd, "\pi"] & \mathbb{M}Q(i) \arrow[r, "\mathbb{M}Q(i,k)"] & \mathbb{M}Q(k) \\
 & & P(i) \arrow[u, "\tau_i"'] \arrow[r, "P(i,k)"'] \arrow[d, "\sigma_i"] & P(k) \arrow[u, "\tau_k"'] \arrow[d, "\sigma_k"] \\
\mathfrak{A} \arrow[r, shift left, "\varphi"] \arrow[r, shift right, "\psi"'] & \mathfrak{B} \arrow[r, "q_k"'] & Q(i) \arrow[r, "Q(i,k)"'] & Q(k)
\end{tikzcd}$$

Our construction proceeds by induction on $i < \kappa + 1$. For $i = 0$, let $p_i : B \to P(0)$ be the coequaliser (in $\mathcal{C}$) of $\mathbb{M}\varphi$ and $\mathbb{M}\psi$ and let $q_i : B \to Q(0)$ 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 $\sigma_0 : P(0) \to Q(0)$ 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 that there exists a unique morphism $\tau_0 : P(0) \to \mathbb{M}Q(0)$ such that

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

For the inductive step, suppose that $P(i)$, $Q(i)$, σ_i, τ_i, p_i, q_i are defind for all $i < \alpha$. We distinguish two cases. If α is a limit ordinal, we choose for $P(\alpha)$ the colimit of the diagram $P|_\alpha : \alpha \to \mathcal{C}$ and for $Q(\alpha)$ the colimit of $Q|_\alpha : \alpha \to \mathcal{C}$. 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 $\sigma_\alpha : P(\alpha) \to Q(\alpha)$ and $\tau_\alpha : P(\alpha) \to \mathbb{M}Q(\alpha)$ 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 $\sigma_{i+1} : \mathbb{M}Q(i) \to Q(i+1)$ be the coequaliser of

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

We set

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

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

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

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

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

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

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

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

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

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

$$\begin{aligned} q_* \circ \pi &= Q(0, \kappa) \circ q_0 \circ \pi \\ &= Q(0, \kappa) \circ \sigma_0 \circ p_0 \\ &= \sigma_* \circ P(0, \kappa) \circ p_0 \\ &= \sigma_* \circ p_\kappa \\ &= (\sigma_* \circ \tau_*^{-1}) \circ \mathbb{M}q_* \,. \end{aligned}$$

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

$$\begin{tikzcd} \mathbb{M}\mathbb{M}Q(i) & \mathbb{M}P(i) \arrow[l, "\mathbb{M}\tau_i"'] \arrow[r, shift left, "\mathbb{M}\sigma_i"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\tau_i"'] & \mathbb{M}Q(i) \arrow[r, "\sigma_{i+1}"] & Q(i+1) \end{tikzcd}$$

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

$$\begin{tikzcd} \mathbb{M}\mathbb{M}Q(i) \arrow[d, "\mathbb{M}\mathbb{M}Q(i,k)"'] & \mathbb{M}P(i) \arrow[l, "\mathbb{M}\tau_i"'] \arrow[r, shift left, "\mathbb{M}\sigma_i"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\tau_i"'] \arrow[d, "\mathbb{M}P(i,k)"'] & \mathbb{M}Q(i) \arrow[r, "\sigma_{i+1}"] \arrow[d, "\mathbb{M}Q(i,k)"'] & Q(i+1) \arrow[d, "Q(i,k)"'] \\ \mathbb{M}\mathbb{M}Q(k) & \mathbb{M}P(k) \arrow[l, "\mathbb{M}\tau_k"'] \arrow[r, shift left, "\mathbb{M}\sigma_k"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\tau_k"'] & \mathbb{M}Q(k) \arrow[r, "\sigma_{k+1}"] & Q(k+1) \end{tikzcd}$$

This colimit is of the form

$$\begin{tikzcd} \mathbb{M}\mathbb{M}Q(\kappa) & \mathbb{M}P(\kappa) \arrow[l, "\alpha"'] \arrow[r, shift left, "\beta"] \arrow[r, shift right, "\gamma"'] & \mathbb{M}Q(\kappa) \arrow[r, "\delta"] & Q(\kappa) \end{tikzcd}$$

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

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

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

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

as desired.

As $q_* \circ \varphi = Q(0,\kappa) \circ q_0 \circ \varphi = Q(0,\kappa) \circ q_0 \circ \psi = q_* \circ \psi$, it remains to prove weak universality. We will show below that, for every $\mathbb{M}$-morphism $\chi : \mathfrak{B} \to \mathfrak{C}$ 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 $\chi = \mu_\kappa \circ q_*$ and μ_κ is an $\mathbb{M}$-morphism since

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

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 $\mathcal{C}$). Hence, $\chi = \mu_0 \circ q_0$, for some $\mu_0 : Q(0) \to C$. 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

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

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

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

It follows that

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

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 $\mu_\delta : Q(\delta) \to \mathrm{C}$ 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

In this section we present several ways to construct monads by transferring a monad from one category to another one. We start with transfer a monad along a natural transformation.

Definition 6.1. Let $\langle \mathbb{M}, \mu, \varepsilon\rangle$ and $\langle \mathbb{P}, \nu, \eta\rangle$ be monads. A natural transformation $\rho : \mathbb{M} \Rightarrow \mathbb{P}$ is a *morphism of monads* if

$$\eta = \rho \circ \varepsilon \quad \text{and} \quad \nu \circ (\rho \circ \mathbb{M}\rho) = \rho \circ \mu\,.$$

If there exists a morphism $\mathbb{M} \Rightarrow \mathbb{P}$, we say that $\mathbb{M}$ is a *reduct* of $\mathbb{P}$. ⌟

The following lemma is frequently useful to prove that a functor forms a monad.

Lemma 6.2. *Let $\mathbb{M}$ and $\mathbb{P}$ be functors, $\mu = (\mu_A)_A$, $\nu = (\nu_A)_A$, $\varepsilon = (\varepsilon_A)_A$, and $\eta = (\eta_A)_A$ families of morphisms*

$$\begin{aligned} \mu_A &: \mathbb{M}\mathbb{M}A \to \mathbb{M}A\,, \qquad & \varepsilon_A &: A \to \mathbb{M}A\,, \\ \nu_A &: \mathbb{P}\mathbb{P}A \to \mathbb{P}A\,, \qquad & \eta_A &: A \to \mathbb{P}A\,, \end{aligned}$$

and let $\rho : \mathbb{M} \Rightarrow \mathbb{P}$ be a natural transformation satisfying

$$\eta = \rho \circ \varepsilon \quad \textit{and} \quad \nu \circ \rho \circ \mathbb{M}\rho = \rho \circ \mu\,.$$

(a) *Suppose that ρ consists of monomorphisms. If $\langle \mathbb{P}, \nu, \eta \rangle$ is a monad, then so is $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\rho : \mathbb{M} \Rightarrow \mathbb{P}$ is a morphism of monads.*

(b) *Suppose that ρ, $\mathbb{M}\rho$, and $\mathbb{M}\mathbb{M}\rho$ consist of epimorphisms. If $\langle \mathbb{M}, \mu, \varepsilon \rangle$ is a monad, then so is $\langle \mathbb{P}, \nu, \eta \rangle$ and $\rho : \mathbb{M} \Rightarrow \mathbb{P}$ is a morphism of monads.*

Proof. (a) We start by proving that ν and η are natural transformations. Given a morphism $f : A \to B$ we have

$$\begin{aligned} \rho \circ \mu \circ \mathbb{M}\mathbb{M}f &= \nu \circ \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}f \\ &= \nu \circ \rho \circ \mathbb{M}\mathbb{P}f \circ \mathbb{M}\rho \\ &= \nu \circ \mathbb{P}\mathbb{P}f \circ \rho \circ \mathbb{M}\rho \\ &= \mathbb{P}f \circ \nu \circ \rho \circ \mathbb{M}\rho \\ &= \mathbb{P}f \circ \rho \circ \mu \\ &= \rho \circ \mathbb{M}f \circ \mu\,, \end{aligned}$$

$$\rho \circ \varepsilon \circ f = \eta \circ f = \mathbb{P}f \circ \eta = \mathbb{P}f \circ \rho \circ \varepsilon = \rho \circ \mathbb{M}f \circ \varepsilon\,.$$

As ρ is a monomorphism, it follows that $\mathbb{M}f \circ \mu = \mu \circ \mathbb{M}\mathbb{M}f$ and $\mathbb{M}f \circ \varepsilon = \varepsilon \circ f$.

It remains to check that $\langle \mathbb{P}, \nu, \eta \rangle$ is a monad.

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

$$\begin{aligned}
\rho \circ \mu \circ \mathbb{M}\varepsilon &= \nu \circ \rho \circ \mathbb{M}\rho \circ \mathbb{M}\varepsilon \\
&= \nu \circ \rho \circ \mathbb{M}\eta \\
&= \nu \circ \mathbb{P}\eta \circ \rho \\
&= \rho\,,
\end{aligned}$$

$$\begin{aligned}
\rho \circ \mu \circ \mu &= \nu \circ \rho \circ \mathbb{M}\rho \circ \mu \\
&= \nu \circ \mathbb{P}\rho \circ \rho \circ \mu \\
&= \nu \circ \mathbb{P}\rho \circ \nu \circ \rho \circ \mathbb{M}\rho \\
&= \nu \circ \nu \circ \mathbb{P}\mathbb{P}\rho \circ \mathbb{P}\rho \circ \rho \\
&= \nu \circ \mathbb{P}\nu \circ \mathbb{P}\mathbb{P}\rho \circ \mathbb{P}\rho \circ \rho \\
&= \nu \circ \rho \circ \mathbb{M}(\nu \circ \mathbb{P}\rho \circ \rho) \\
&= \nu \circ \rho \circ \mathbb{M}(\nu \circ \rho \circ \mathbb{M}\rho) \\
&= \nu \circ \rho \circ \mathbb{M}(\rho \circ \mu) \\
&= \rho \circ \mu \circ \mathbb{M}\mu\,.
\end{aligned}$$

As ρ is a monomorphism, it follows that $\mu \circ \varepsilon = \mathrm{id} = \mu \circ \mathbb{M}\varepsilon$ and $\mu \circ \mu = \mu \circ \mathbb{M}\mu$.

(b) We start by proving that ν and η are natural transformations. Since ρ and ε are natural transformations, so is $\eta = \rho \circ \varepsilon$. Furthermore, given a morphism $f : A \to B$ we have

$$\begin{aligned}
\mathbb{P}f \circ \nu \circ \rho \circ \mathbb{M}\rho &= \mathbb{P}f \circ \rho \circ \mu \\
&= \rho \circ \mathbb{M}f \circ \mu \\
&= \rho \circ \mu \circ \mathbb{M}\mathbb{M}f \\
&= \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}f \\
&= \rho \circ \mathbb{M}\mathbb{P}f \circ \mathbb{M}\rho \\
&= \mathbb{P}\mathbb{P}f \circ \rho \circ \mathbb{M}\rho\,.
\end{aligned}$$

As ρ and $\mathbb{M}\rho$ are epimorphisms, it follows that $\mathbb{P}f \circ \nu = \nu \circ \mathbb{P}\mathbb{P}f$.

It remains to check that $\langle \mathbb{P}, \nu, \eta \rangle$ is a monad.

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

As ρ is an epimorphism, it follows that $\nu \circ \eta = \mathrm{id} = \nu \circ \mathbb{P}\eta$. Furthermore,

$$\begin{aligned}
\nu \circ \nu \circ \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}\rho &= \nu \circ \rho \circ \mu \circ \mathbb{M}\mathbb{M}\rho \\
&= \nu \circ \rho \circ \mathbb{M}\rho \circ \mu \\
&= \rho \circ \mu \circ \mu \\
&= \rho \circ \mu \circ \mathbb{M}\mu \\
&= \nu \circ \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mu \\
&= \nu \circ \rho \circ \mathbb{M}\nu \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}\rho \\
&= \nu \circ \mathbb{P}\nu \circ \rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}\rho \,.
\end{aligned}$$

As ρ, $\mathbb{M}\rho$, and $\mathbb{M}\mathbb{M}\rho$ are all epimorphisms, so is their composition. Consequently, it follows that $\nu \circ \nu = \nu \circ \mathbb{P}\nu$. ☐

Next, we consider a transfer along an adjunction.

Proposition 6.3. *Let $\mathbb{F} \dashv \mathbb{G}$ be an adjunction between the categories $\mathcal{C}$ and $\mathcal{D}$, let $e : \mathrm{Id} \Rightarrow \mathbb{G}\mathbb{F}$ be its unit, $i : \mathbb{F}\mathbb{G} \Rightarrow \mathrm{Id}$ the counit, and let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a monad on $\mathcal{D}$.*

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

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

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

(c) *If* $\mathsf{Alg}(\mathbb{M})$ *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* $\mathsf{Alg}(\mathbb{M})$*) of*

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

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

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

For the monad axioms, note that

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

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

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

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

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

(c) The idea of the proof is as follows. When constructing $\mathfrak{Q} := \widehat{\mathbb{F}}\mathfrak{A}$ from $\mathfrak{A}$ we have to transfer the product $\pi : \mathbb{P}A \to A$ to a function $\mathbb{M}Q \to Q$. 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 $\mathfrak{A}$ to $\mathfrak{Q}$. But this does not work since not every term $s \in \mathbb{M}\mathbb{F}A$ corresponds to some $t \in \mathbb{P}A$. 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 $\mathbb{M}\mathbb{F}A$ 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 $s, t \in \mathbb{M}\mathbb{F}A$ where the products of the corresponding terms in $\mathbb{P}A$ coincide. Thus, we arrive at our final definiton, we take for $\mathfrak{Q}$ a suitable quotient of $\mathbb{M}\mathbb{F}A$.

The formal definition is as follows. Let $\rho : \mathbb{M}\mathbb{F}A \to \mathfrak{Q}$ with $\mathfrak{Q} = \langle Q, \sigma \rangle$ be the coequaliser (in $\mathsf{Alg}(\mathbb{M})$) of

$$\mathbb{M}\mathbb{F}\pi,\ \mu \circ \mathbb{M}i : \mathbb{M}\mathbb{F}\mathbb{P}A \to \mathbb{M}\mathbb{F}A\,.$$

(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, $\mathbb{M}\mathbb{F}\pi$ and $\mathbb{M}i$ are $\mathbb{M}$-morphisms.) We set $\widehat{\mathbb{F}}\mathfrak{A} := \mathfrak{Q}$.

$$\begin{tikzcd}
\mathbb{M}\mathbb{F}\mathbb{P}A \arrow[r, shift left, "\mathbb{M}\mathbb{F}\pi"] \arrow[r, shift right, "\mu \circ \mathbb{M}i"'] \arrow[dr, "\mathbb{M}i"'] & \mathbb{M}\mathbb{F}A \arrow[r, "\rho"] & Q \\
& \mathbb{M}\mathbb{M}\mathbb{F}A \arrow[u, "\mu"] \arrow[r, "\mathbb{M}\rho"'] & \mathbb{M}Q \arrow[u, "\sigma"']
\end{tikzcd}$$

To define the action of $\widehat{\mathbb{F}}$ on morphisms, let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be a $\mathbb{P}$-morphism. Then φ induces morphisms between the coequaliser diagrams assciated with $\widehat{\mathbb{F}}\mathfrak{A}$ and $\widehat{\mathbb{F}}\mathfrak{B}$. (This is in fact the definition of $\widehat{\mathbb{F}}\varphi$.)

$$\begin{tikzcd}
\mathbb{M}\mathbb{F}\mathbb{P}A \arrow[r, shift left, "\mathbb{M}\mathbb{F}\pi"] \arrow[r, shift right, "\mu \circ \mathbb{M}i"'] \arrow[d, "\mathbb{M}\mathbb{F}\mathbb{P}\varphi"'] & \mathbb{M}\mathbb{F}A \arrow[r, "\rho_A"] \arrow[d, "\mathbb{M}\mathbb{F}\varphi"] & Q_A \arrow[d, "\widehat{\mathbb{F}}\varphi"] \\
\mathbb{M}\mathbb{F}\mathbb{P}B \arrow[r, shift left, "\mathbb{M}\mathbb{F}\pi"] \arrow[r, shift right, "\mu \circ \mathbb{M}i"'] & \mathbb{M}\mathbb{F}B \arrow[r, "\rho_B"'] & Q_B
\end{tikzcd}$$

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

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

Furthermore, every component ρ_A of ρ is an epimorphism (in $\mathsf{Alg}(\mathbb{M})$): given $\mathbb{M}$-morphisms $f, g : \mathfrak{Q} \to \mathfrak{C}$ 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 $h : \mathfrak{Q} \to \mathfrak{C}$ 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 $\tau : \mathrm{Id} \Rightarrow \widehat{\mathbb{G}}\widehat{\mathbb{F}}$ of the adjunction by $\tau := \mathbb{G}\rho \circ \eta$. Note that τ is a $\mathbb{P}$-morphism $\mathfrak{A} \to \widehat{\mathbb{G}}\mathfrak{Q}$ since

$$\begin{aligned}
\tau \circ \pi &= \mathbb{G}\rho \circ \eta \circ \pi \\
&= \mathbb{G}\rho \circ \mathbb{G}\varepsilon \circ e \circ \pi \\
&= \mathbb{G}(\rho \circ \varepsilon \circ \mathbb{F}\pi) \circ e \\
&= \mathbb{G}(\rho \circ \mathbb{MF}\pi \circ \varepsilon) \circ e \\
&= \mathbb{G}(\rho \circ \mu \circ \mathbb{M}i \circ \varepsilon) \circ e \\
&= \mathbb{G}\rho \circ \nu \circ \eta \\
&= \mathbb{G}\rho \\
&= \mathbb{G}\big(\rho \circ \mu \circ \mathbb{M}\varepsilon)\big) \\
&= \mathbb{G}\big(\sigma \circ \mathbb{M}\rho \circ \mathbb{M}(\varepsilon \circ \mathrm{id})\big) \\
&= \mathbb{G}\big(\sigma \circ \mathbb{M}\big((\rho \circ \varepsilon) \circ i \circ \mathbb{F}e\big)\big) \\
&= \mathbb{G}\big(\sigma \circ \mathbb{M}\big(i \circ \mathbb{FG}(\rho \circ \varepsilon) \circ \mathbb{F}e\big)\big) \\
&= \mathbb{G}(\sigma \circ \mathbb{M}i) \circ \mathbb{GMF}(\mathbb{G}\rho \circ \eta) \\
&= \mathbb{G}(\sigma \circ \mathbb{M}i) \circ \mathbb{P}\tau\,.
\end{aligned}$$

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

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

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

To define the counit, let $\mathfrak{A} = \langle A, \pi\rangle$ be an $\mathbb{M}$-algebra and let $\rho = \rho_{\mathbb{G}A}$:

$\mathbb{MF}A \to \mathfrak{Q}_{\widehat{\mathbb{G}}A}$ be the coequaliser used in the definition of $\widehat{\mathbb{F}}(\widehat{\mathbb{G}}\mathfrak{A})$. Since

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

and $\pi \circ \mathbb{M}i$ is an $\mathbb{M}$-morphism, we can use universality of ρ to find a unique $\mathbb{M}$-morphism $v : \widehat{\mathbb{F}}\widehat{\mathbb{G}}\mathfrak{A} \to \mathfrak{A}$ 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 $v \circ \widehat{\mathbb{F}}\tau = \mathrm{id}$ and $\widehat{\mathbb{G}}v \circ \tau = \mathrm{id}$. For the second equation, we have

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

For the first one, we have

$$\begin{aligned}
\upsilon \circ \widehat{\mathbb{F}}\tau \circ \rho &= \upsilon \circ \rho \circ \mathbb{MF}\tau \\
&= \sigma \circ \mathbb{M}i \circ \mathbb{MF}\tau \\
&= \sigma \circ \mathbb{M}(i \circ \mathbb{FG}\rho \circ \mathbb{F}\eta) \\
&= \sigma \circ \mathbb{M}(\rho \circ i \circ \mathbb{F}\eta) \\
&= \rho \circ \mu \circ \mathbb{M}(i \circ \mathbb{F}\eta) \\
&= \rho \circ \mu \circ \mathbb{M}(i \circ \mathbb{FG}\varepsilon \circ \mathbb{F}e) \\
&= \rho \circ \mu \circ \mathbb{M}(\varepsilon \circ i \circ \mathbb{F}e) \\
&= \rho \circ \mu \circ \mathbb{M}(\varepsilon \circ \mathrm{id}) \\
&= \rho \circ \mathrm{id} \\
&= \rho\,.
\end{aligned}$$

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

Distributive Laws

Our final 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 $\mathbb{P}$. An equivalent way of looking at such an expansion is by finding a lift of $\mathbb{P}$ from the base category to $\mathsf{Alg}(\mathbb{M})$. It turns out that, in order for this to work, the two monads $\mathbb{M}$ and $\mathbb{P}$ need to be compatible: there needs to be what is called a *distributive law* between them.

Definition 6.4. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$ be monads.

(a) A natural transformation $\delta : \mathbb{MP} \Rightarrow \mathbb{PM}$ is a *distributive law* if

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

$$\begin{tikzcd}
\mathbb{M}\mathbb{M}\mathbb{P}A \arrow[r, "\mathbb{M}\delta"] \arrow[d, "\mu"'] & \mathbb{M}\mathbb{P}\mathbb{M}A \arrow[r, "\delta"] & \mathbb{P}\mathbb{M}\mathbb{M}A \arrow[d, "\mathbb{P}\mu"] \\
\mathbb{M}\mathbb{P}A \arrow[rr, "\delta"'] & & \mathbb{P}\mathbb{M}A
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathbb{P}A \arrow[d, "\varepsilon"'] \arrow[dr, "\mathbb{P}\varepsilon"] & \\
\mathbb{M}\mathbb{P}A \arrow[r, "\delta"'] & \mathbb{P}\mathbb{M}A
\end{tikzcd}$$

$$\begin{tikzcd}
\mathbb{M}\mathbb{P}\mathbb{P}A \arrow[r, "\delta"] \arrow[d, "\mathbb{M}\nu"'] & \mathbb{P}\mathbb{M}\mathbb{P}A \arrow[r, "\mathbb{P}\delta"] & \mathbb{P}\mathbb{P}\mathbb{M}A \arrow[d, "\nu"] \\
\mathbb{M}\mathbb{P}A \arrow[rr, "\delta"'] & & \mathbb{P}\mathbb{M}A
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathbb{M}A \arrow[d, "\mathbb{M}\eta"'] \arrow[dr, "\eta"] & \\
\mathbb{M}\mathbb{P}A \arrow[r, "\delta"'] & \mathbb{P}\mathbb{M}A
\end{tikzcd}$$

(b) A natural transformation $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ is *half a distributive law* if

$$-1\delta \circ \mu = \mathbb{P}\mu \circ \delta \circ \mathbb{M}\delta$$
$$\delta \circ \varepsilon = \mathbb{P}\varepsilon\,. \qquad \lrcorner$$

Example. Let $\mathbb{M}A := A^*$ be the monad for finite words and $\mathbb{A}A := A^*/{\approx}$ 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 $\mathbb{A}A$ as sums $a + \cdots + b$ and those of $\mathbb{M}A$ as products. A distributive law $\delta : \mathbb{M}\mathbb{A} \Rightarrow \mathbb{A}\mathbb{M}$ maps a product of sums

$$\sum_{i<n_0} a_{0i} \cdot \sum_{i<n_1} a_{1i} \cdot \cdots \cdot \sum_{i<n_{k-1}} a_{ki}$$

to the sum

$$\sum_{\sigma} a_{0,\sigma(0)} \cdots a_{k-1,\sigma(k-1)}$$

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

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

We start with a simple version of the construction we are interested in.

Theorem 6.5. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be half a distributive law. Then there exists a functor $\widehat{\mathbb{P}} : \mathsf{Alg}(\mathbb{M}) \to \mathsf{Alg}(\mathbb{M})$ mapping an $\mathbb{M}$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the $\mathbb{M}$-algebra $\mathbb{P}\mathfrak{A} := \langle A, \mathbb{P}\pi \circ \delta \rangle$ and mapping an $\mathbb{M}$-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ to the $\mathbb{M}$-morphism $\mathbb{P}\varphi : \mathbb{P}\mathfrak{A} \to \mathbb{P}\mathfrak{B}$.*

Proof. We start by checking that $\mathbb{P}\mathfrak{A}$ is an $\mathbb{M}$-algebra. Set $\hat{\pi} := \mathbb{P}\pi \circ \delta$. Then

$$\begin{aligned}
\hat{\pi} \circ \varepsilon &= \mathbb{P}\pi \circ \delta \circ \varepsilon \\
&= \mathbb{P}\pi \circ \mathbb{P}\varepsilon \\
&= \mathrm{id}\,, \\
\hat{\pi} \circ \mathbb{M}\hat{\pi} &= \mathbb{P}\pi \circ \delta \circ \mathbb{M}(\mathbb{P}\pi \circ \delta) \\
&= \mathbb{P}\pi \circ \mathbb{P}\mathbb{M}\pi \circ \delta \circ \mathbb{M}\delta \\
&= \mathbb{P}(\pi \circ \mu) \circ \delta \circ \mathbb{M}\delta \\
&= \mathbb{P}\pi \circ \delta \circ \mu \\
&= \hat{\pi} \circ \mu\,.
\end{aligned}$$

Finally, let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be an $\mathbb{M}$-morphism. To see that $\mathbb{P}\varphi$ is also an $\mathbb{M}$-morphism, note that

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

□

The downside of this simple version is that it does not imply that $\mathbb{P}$ forms a monad on $\mathsf{Alg}(\mathbb{M})$. If we want a monad, we need a full distributive law.

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 life a monad $\mathbb{P}$ from the base category to the category of $\mathbb{M}$-algebras.

Definition 6.6. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$ be monads on some category $\mathcal{C}$ and let $\mathbb{U} : \mathsf{Alg}(\mathbb{M}) \to \mathcal{C}$ be the forgetful functor mapping an $\mathbb{M}$-algebra to its universe.

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

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

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

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

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

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

Theorem 6.7 (Beck). *Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$ be monads on the category $\mathcal{C}$. There exist bijections between the following objects:*

(1) *distributive laws $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$;*

(2) *liftings $\hat{\mathbb{P}}$ of $\mathbb{P}$ to the category of $\mathbb{M}$-algebras;*

(3) *extensions $\widehat{\mathbb{M}}$ of $\mathbb{M}$ to the Kleisli category* $\mathsf{Free}(\mathbb{P})$*;*

(4) *functions κ such that*

  (m1) *$\langle \mathbb{P}\mathbb{M}, \kappa, \eta \circ \varepsilon \rangle$ is a monad,*

  (m2) *the functions $\mathbb{P}\varepsilon$ and η induce morphisms of monads $\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ and $\mathbb{M} \Rightarrow \mathbb{P}\mathbb{M}$,*

  (m3) *κ satisfies the* middle unit law: $\kappa \circ \mathbb{P}(\varepsilon \circ \eta) = \mathrm{id}$;

(5) *functions κ such that*

  (c1) *$\langle \mathbb{P}\mathbb{M}, \kappa, \eta \circ \varepsilon \rangle$ is a monad,* (c2) $\kappa \circ \mathbb{PM}\eta = \mathbb{M}\mu$,

(c3) $\kappa \circ \mathbb{P}\varepsilon = \nu$,

(c4) $\nu \circ \mathbb{P}\kappa = \kappa \circ \nu$,

(c5) $\mathbb{P}\mu \circ \kappa = \kappa \circ \mathbb{PMP}\mu$.

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

$$\begin{aligned} \widehat{\mathbb{M}}\langle \mathbb{P}C, \nu\rangle &:= \langle \mathbb{PM}C, \nu\rangle\,, && \text{for an algebra } \langle \mathbb{P}C, \nu\rangle \in \mathsf{Free}(\mathbb{P})\,, \\ \widehat{\mathbb{M}}\varphi &:= \nu \circ \mathbb{P}\delta \circ \mathbb{PM}(\varphi \circ \eta)\,, && \text{for a morphism } \varphi : \mathfrak{A} \to \mathfrak{B}\,. \end{aligned}$$

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

$$\begin{aligned} \widehat{\mathbb{M}}\varphi \circ \nu &= \nu \circ \mathbb{P}\delta \circ \mathbb{PM}(\varphi \circ \eta) \circ \nu \\ &= \nu \circ \nu \circ \mathbb{PP}\delta \circ \mathbb{PPM}(\varphi \circ \eta) \\ &= \nu \circ \mathbb{P}\nu \circ \mathbb{PP}\delta \circ \mathbb{PPM}(\varphi \circ \eta) \\ &= \nu \circ \mathbb{P}(\nu \circ \mathbb{P}\delta \circ \mathbb{PM}(\varphi \circ \eta)) = \nu \circ \mathbb{P}\widehat{\mathbb{M}}\varphi\,. \end{aligned}$$

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

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

Similarly, for a morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$, we have

$$\begin{aligned} \widehat{\mathbb{M}}\mathbb{F}_{\mathbb{P}}\varphi &= \nu \circ \mathbb{P}\delta \circ \mathbb{PM}(\mathbb{P}\varphi \circ \eta) \\ &= \nu \circ \mathbb{P}(\delta \circ \mathbb{MP}\varphi \circ \mathbb{M}\eta) \\ &= \nu \circ \mathbb{PPM}\varphi \circ \mathbb{P}(\delta \circ \mathbb{M}\eta) \\ &= \mathbb{PM}\varphi \circ \nu \circ \mathbb{P}\eta = \mathbb{PM}\varphi = \mathbb{F}_{\mathbb{P}}\mathbb{M}\varphi\,. \end{aligned}$$

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

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

(3) ⇒ (1) Given an extension $\widehat{\mathbb{M}}$ of $\mathbb{M}$ to $\mathsf{Free}(\mathbb{P})$, we set

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

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

$$\widehat{\mathbb{M}} \circ \mathbb{P} = \mathbb{P} \circ \mathbb{M}\,, \qquad \hat{\mu} = \mathbb{P}\mu\,, \qquad \hat{\varepsilon} = \mathbb{P}\varepsilon\,.$$

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

$$\nu \circ \mathbb{P}\widehat{\mathbb{M}}\varphi = \widehat{\mathbb{M}}\varphi \circ \nu\,, \qquad \text{for every morphism } \varphi : \mathbb{F}_{\mathbb{P}}A \to \mathbb{F}_{\mathbb{P}}B\,.$$

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

Then

$$
\begin{aligned}
\delta \circ \mathbb{M}\mathbb{P}f &= \widehat{\mathbb{M}}\nu \circ \eta \circ \mathbb{M}\mathbb{P}f \\
&= \widehat{\mathbb{M}}\nu \circ \mathbb{P}\mathbb{M}\mathbb{P}f \circ \eta \\
&= \widehat{\mathbb{M}}\nu \circ \widehat{\mathbb{M}}\mathbb{P}\mathbb{P}f \circ \eta \\
&= \widehat{\mathbb{M}}(\nu \circ \mathbb{P}\mathbb{P}f) \circ \eta \\
&= \widehat{\mathbb{M}}(\mathbb{P}f \circ \nu) \circ \eta \\
&= \widehat{\mathbb{M}}\mathbb{P}f \circ \widehat{\mathbb{M}}\nu \circ \eta = \mathbb{P}\mathbb{M}f \circ \delta \,.
\end{aligned}
$$

Furthermore, δ is a distributive law since

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

$$\begin{aligned}
\delta \circ \mathbb{M}\nu &= \widehat{\mathbb{M}}\nu \circ \eta \circ \mathbb{M}\nu \\
&= \widehat{\mathbb{M}}\nu \circ \mathbb{P}\mathbb{M}\nu \circ \eta \\
&= \widehat{\mathbb{M}}\nu \circ \widehat{\mathbb{M}}\mathbb{P}\nu \circ \eta \\
&= \widehat{\mathbb{M}}(\nu \circ \mathbb{P}\nu) \circ \eta \\
&= \widehat{\mathbb{M}}(\nu \circ \nu) \circ \eta \\
&= \widehat{\mathbb{M}}\nu \circ \widehat{\mathbb{M}}\nu \circ \eta \\
&= \widehat{\mathbb{M}}\nu \circ \nu \circ \mathbb{P}\eta \circ \widehat{\mathbb{M}}\nu \circ \eta \\
&= \nu \circ \mathbb{P}\widehat{\mathbb{M}}\nu \circ \mathbb{P}\eta \circ \widehat{\mathbb{M}}\nu \circ \eta \\
&= \nu \circ \mathbb{P}(\widehat{\mathbb{M}}\nu \circ \eta) \circ \widehat{\mathbb{M}}\nu \circ \eta = \nu \circ \mathbb{P}\delta \circ \delta\,,
\end{aligned}$$

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

$$\begin{aligned}
\delta \circ \mathbb{M}\eta &= \widehat{\mathbb{M}}\nu \circ \eta \circ \mathbb{M}\eta \\
&= \widehat{\mathbb{M}}\nu \circ \mathbb{P}\mathbb{M}\eta \circ \eta \\
&= \widehat{\mathbb{M}}\nu \circ \widehat{\mathbb{M}}\mathbb{P}\eta \circ \eta \\
&= \widehat{\mathbb{M}}(\nu \circ \mathbb{P}\eta) \circ \eta = \widehat{\mathbb{M}}\mathrm{id} \circ \eta = \eta\,.
\end{aligned}$$

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

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

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

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

Furthermore, for morphism a $\varphi : \mathfrak{A} \to \mathfrak{B}$ of $\mathbb{M}$-algebras, we have

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

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

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

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

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

$$\hat{\mu} : \mathbb{M}\mathbb{P}\mathbb{M}A \to \mathbb{P}\mathbb{M}A \quad \text{and} \quad \hat{\hat{\mu}} : \mathbb{M}\mathbb{P}\mathbb{P}\mathbb{M}A \to \mathbb{P}\mathbb{P}\mathbb{M}A\,.$$

We claim that

$$\kappa := \nu \circ \mathbb{P}\hat{\mu}$$

is the desired morphism.

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

$$\hat{\mu} \circ \mathbb{M}\hat{\mu} = \hat{\mu} \circ \mu \quad \text{and} \quad \hat{\mu} \circ \varepsilon = \mathrm{id}\,.$$

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

$$\mathbb{P}\hat{\mu} \circ \hat{\mu} = \hat{\hat{\mu}} \circ \mathbb{M}\mathbb{P}\hat{\mu}\,.$$

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

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

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

It remains to check the monad laws.

$$\begin{aligned}
\kappa \circ \kappa &= \nu \circ \mathbb{P}\hat{\mu} \circ \nu \circ \mathbb{P}\hat{\mu} \\
&= \nu \circ \nu \circ \mathbb{P}\mathbb{P}\hat{\mu} \circ \mathbb{P}\hat{\mu} \\
&= \nu \circ \mathbb{P}\nu \circ \mathbb{P}\hat{\hat{\mu}} \circ \mathbb{P}\mathbb{M}\mathbb{P}\hat{\mu} \\
&= \nu \circ \mathbb{P}\hat{\mu} \circ \mathbb{P}\mathbb{M}\nu \circ \mathbb{P}\mathbb{M}\mathbb{P}\hat{\mu} \\
&= \kappa \circ \mathbb{P}\mathbb{M}\kappa\,, \\
\kappa \circ \eta \circ \varepsilon &= \nu \circ \mathbb{P}\hat{\mu} \circ \mathbb{P}\varepsilon \circ \eta = \nu \circ \eta = \mathrm{id}\,, \\
\kappa \circ \mathbb{P}\mathbb{M}(\eta \circ \varepsilon) &= \nu \circ \mathbb{P}\hat{\mu} \circ \mathbb{P}\mathbb{M}(\eta \circ \varepsilon) \\
&= \nu \circ \mathbb{P}(\hat{\mu} \circ \mathbb{M}\eta) \circ \mathbb{P}\mathbb{M}\varepsilon \\
&= \nu \circ \mathbb{P}(\eta \circ \mu) \circ \mathbb{P}\mathbb{M}\varepsilon \\
&= \mathrm{id}\,.
\end{aligned}$$

For (M2), we have

$$\kappa \circ \mathbb{P}\varepsilon \circ \mathbb{P}\mathbb{P}\varepsilon = \nu \circ \mathbb{P}\hat{\mu} \circ \mathbb{P}\varepsilon \circ \mathbb{P}\mathbb{P}\varepsilon = \nu \circ \mathbb{P}\mathbb{P}\varepsilon = \mathbb{P}\varepsilon \circ \nu\,,$$

$$\kappa \circ \eta \circ \mathbb{M}\eta = \nu \circ \mathbb{P}\hat{\mu} \circ \eta \circ \mathbb{M}\eta = \nu \circ \eta \circ \hat{\mu} \circ \mathbb{M}\eta = \hat{\mu} \circ \mathbb{M}\eta = \eta \circ \mu\,.$$

For (M3), we have

$$\kappa \circ \mathbb{P}(\varepsilon \circ \eta) = \nu \circ \mathbb{P}\hat{\mu} \circ \mathbb{P}(\varepsilon \circ \eta) = \nu \circ \mathbb{P}\eta = \mathrm{id}\,.$$

(4) ⇒ (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

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

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

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

For (C2), we have

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

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

$$
\begin{aligned}
\mathbb{P}\mu \circ \kappa &= \kappa \circ \mathbb{PM}\eta \circ \kappa \\
&= \kappa \circ \kappa \circ \mathbb{PMPM}\eta \\
&= \kappa \circ \mathbb{PM}\kappa \circ \mathbb{PMPM}\eta \\
&= \kappa \circ \mathbb{PM}(\kappa \circ \mathbb{PM}\eta) \\
&= \kappa \circ \mathbb{PMP}\mu \,.
\end{aligned}
$$

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

$$\delta := \kappa \circ \eta \circ \mathbb{MP}\varepsilon \,.$$

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

$$
\begin{aligned}
\delta \circ \mu &= \kappa \circ \eta \circ \mathbb{MP}\varepsilon \circ \mu \\
&= \kappa \circ \eta \circ \mu \circ \mathbb{MMP}\varepsilon \\
&= \kappa \circ \mathbb{P}\mu \circ \eta \circ \mathbb{MMP}\varepsilon \\
&= \kappa \circ (\kappa \circ \mathbb{PM}\eta) \circ \eta \circ \mathbb{MMP}\varepsilon \\
&= \kappa \circ \mathbb{PM}\kappa \circ \mathbb{PM}\eta \circ \eta \circ \mathbb{MMP}\varepsilon \\
&= \kappa \circ \eta \circ \mathbb{M}\kappa \circ \mathbb{M}\eta \circ \mathbb{MMP}\varepsilon \\
&= \kappa \circ \eta \circ \mathbb{MP}(\mu \circ \varepsilon) \circ \mathbb{M}\delta \\
&= \kappa \circ \mathbb{PMP}\mu \circ \eta \circ \mathbb{MP}\varepsilon \circ \mathbb{M}\delta \\
&= \mathbb{P}\mu \circ \kappa \circ \eta \circ \mathbb{MP}\varepsilon \circ \mathbb{M}\delta \\
&= \mathbb{P}\mu \circ \delta \circ \mathbb{M}\delta
\end{aligned}
$$

$$\begin{aligned}
\delta \circ \mathbb{M}\nu &= \kappa \circ \eta \circ \mathbb{MP}\varepsilon \circ \mathbb{M}\nu \\
&= \kappa \circ \eta \circ \mathbb{MP}\varepsilon \circ \mathbb{M}(\kappa \circ \mathbb{P}\varepsilon) \\
&= \kappa \circ \mathbb{PMP}\varepsilon \circ \mathbb{PM}\kappa \circ \eta\mathbb{MP}\varepsilon \\
&= \kappa \circ \mathbb{PM}\kappa \circ \mathbb{PMPMP}\varepsilon \circ \eta \circ \mathbb{MP}\varepsilon\,, \\
&= \kappa \circ \kappa \circ \mathbb{PMPMP}\varepsilon \circ \eta \circ \mathbb{MP}\varepsilon\,, \\
&= \kappa \circ \mathbb{PMP}\varepsilon \circ \kappa \circ \eta \circ \mathbb{MP}\varepsilon\,, \\
&= \kappa \circ \nu \circ \mathbb{P}\eta \circ \mathbb{PMP}\varepsilon \circ \kappa \circ \eta \circ \mathbb{MP}\varepsilon\,, \\
&= \nu \circ \mathbb{P}(\kappa \circ \eta \circ \mathbb{MP}\varepsilon) \circ \kappa \circ \eta \circ \mathbb{MP}\varepsilon\,, \\
&= \nu \circ \mathbb{P}\delta \circ \delta\,,
\end{aligned}$$

$$\begin{aligned}
\delta \circ \varepsilon &= \kappa \circ \eta \circ \mathbb{MP}\varepsilon \circ \varepsilon \\
&= \kappa \circ \eta \circ \varepsilon \circ \mathbb{P}\varepsilon \\
&= \kappa \circ \zeta \circ \mathbb{P}\varepsilon \\
&= \mathbb{P}\varepsilon\,,
\end{aligned}$$

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

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

$$\widehat{\mathbb{M}}\varphi = \nu \circ \delta \circ \mathbb{PM}(\varphi \circ \eta)\,.$$

This functor is then mapped back to

$$\begin{aligned}
\delta' &= \widehat{\mathbb{M}}\nu \circ \eta \\
&= \nu \circ \mathbb{P}\delta \circ \mathbb{P}\mathbb{M}(\nu \circ \eta) \circ \eta \\
&= \nu \circ \mathbb{P}\delta \circ \eta \\
&= \nu \circ \mathbb{P}\delta \circ \delta \circ \mathbb{M}\eta \\
&= \delta \circ \mathbb{M}\nu \circ \mathbb{M}\eta \\
&= \delta \,.
\end{aligned}$$

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

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

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

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

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 $\hat{\mathbb{P}}$ with $\hat{\mu} = \mathbb{P}\mu \circ \delta$. The second step, maps this to the morphism $\kappa = \nu \circ \mathbb{P}\hat{\mu}$. The third step is the identity, and the last step maps κ to $\kappa \circ \eta \circ \mathbb{MP}\varepsilon$. Composing these steps we obtain

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

as desired.

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

$$\begin{aligned}
\nu \circ \mathbb{P}\hat{\mu} &= \nu \circ \mathbb{P}(\mathbb{P}\mu \circ \delta) \\
&= \nu \circ \mathbb{PP}\mu \circ \mathbb{P}(\kappa \circ \eta \circ \mathbb{MP}\varepsilon) \\
&= \mathbb{P}\mu \circ \nu \circ \mathbb{P}\kappa \circ \mathbb{P}\eta \circ \mathbb{PMP}\varepsilon \\
&= \mathbb{P}\mu \circ \kappa \circ \nu \circ \mathbb{P}\eta \circ \mathbb{PMP}\varepsilon \\
&= \mathbb{P}\mu \circ \kappa \circ \mathbb{PMP}\varepsilon \\
&= \kappa \circ \mathbb{PMP}\mu \circ \mathbb{PMP}\varepsilon \\
&= \kappa\,.
\end{aligned}$$

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

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

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

$\mathbb{M}$-algebras and we have

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

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

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

as desired. □

Corollary 6.8. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$.*

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

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

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

$$\mathbb{P}\pi \circ \delta : \mathbb{M}\mathbb{P}A \to \mathbb{P}A .$$

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

$$\nu \circ \mathbb{P}\mathbb{P}\pi \circ \mathbb{P}\delta : \mathbb{P}\mathbb{M}\mathbb{P}A \to \mathbb{P}A .$$

Remark. What do we do if we can to combine two monads $\mathbb{M}$ and $\mathbb{P}$ but there is no distributive law between them? In this case we can use the free monad construction and use the more complicated monad $(\mathbb{P}\mathbb{M})^*$. ⌟

We conclude this section with two simple observations of how to transfer a distributive law from one setting to another one. The first one concerns the transfer along a morphism of monads. A similar result holds for monomorphic ρ.

Lemma 6.9. *Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$, $\langle \mathbb{M}', \mu', \varepsilon' \rangle$, and $\langle \mathbb{P}, \nu, \eta \rangle$ be monads, $\rho : \mathbb{M} \Rightarrow \mathbb{M}'$ a morphism of monads, $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ a distributive law, and $\delta' = (\delta'_A)_A$ a family of functions such that*

$$\delta' \circ \rho = \mathbb{P}\rho \circ \delta \,.$$

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

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

$$\begin{aligned}
\mathbb{P}\mathbb{M}'f \circ \delta' \circ \rho &= \mathbb{P}\mathbb{M}'f \circ \mathbb{P}\rho \circ \delta \\
&= \mathbb{P}\rho \circ \mathbb{P}\mathbb{M}f \circ \delta \\
&= \mathbb{P}\rho \circ \delta \circ \mathbb{M}\mathbb{P}f = \delta' \circ \mathbb{M}'\mathbb{P}f \,.
\end{aligned}$$

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

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

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

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

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

Our second result shows that faithful functors reflect distributive laws.

Lemma 6.10. *Let $\mathbb{F} : \mathcal{C}^0 \to \mathcal{C}^1$ be a faithful functor, let $\langle \mathbb{M}^i, \mu^i, \varepsilon^i \rangle$ and $\langle \mathbb{P}^i, \nu^i, \eta^i \rangle$ be two monads on $\mathcal{C}^i$, for $i < 2$, and let $\delta = (\delta_A)_A$ be a family of morphisms $\delta_A : \mathbb{M}^0\mathbb{P}^0 A \to \mathbb{P}^0\mathbb{M}^0 A$. Suppose that*

- $\mathbb{M}^1 \circ \mathbb{F} = \mathbb{F} \circ \mathbb{M}^0$ *and* $\mathbb{P}^1 \circ \mathbb{F} = \mathbb{F} \circ \mathbb{P}^0$,
- $\mu^1 = \mathbb{F}\mu^0$, $\varepsilon^1 = \mathbb{F}\varepsilon^0$, $\nu^1 = \mathbb{F}\nu^0$, *and* $\eta^1 = \mathbb{F}\eta^0$.

If $\mathbb{F}\delta$ is a distributive law, then so is δ.

Proof. Suppose that $\mathbb{F}\delta$ is a distributive law. Then $\mathbb{F}\delta$ is a natural transformation since

$$\begin{aligned}\mathbb{F}(\mathbb{P}^0\mathbb{M}^0 f \circ \delta) &= \mathbb{F}\mathbb{P}^0\mathbb{M}^0 f \circ \mathbb{F}\delta \\ &= \mathbb{P}^1\mathbb{M}^1\mathbb{F} f \circ \mathbb{F}\delta \\ &= \mathbb{F}\delta \circ \mathbb{M}^1\mathbb{P}^1\mathbb{F} f \\ &= \mathbb{F}\delta \circ \mathbb{F}\mathbb{M}^0\mathbb{P}^0 f = \mathbb{F}(\delta \circ \mathbb{M}^0\mathbb{P}^0 f)\,.\end{aligned}$$

Since $\mathbb{F}$ is faithful, it follows that

$$\mathbb{P}^0\mathbb{M}^0 f \circ \delta = \delta \circ \mathbb{M}^0\mathbb{P}^0 f\,.$$

For the axioms, note that

$$\mathbb{F}(\delta \circ \varepsilon^0) = \mathbb{F}\delta \circ \varepsilon^1 = \mathbb{P}^1\varepsilon^1 = \mathbb{P}^0\mathbb{F}\varepsilon^0 = \mathbb{F}\mathbb{P}^0\varepsilon^0,$$

which implies by faithfulness of $\mathbb{F}$ that $\delta \circ \varepsilon^0 = \mathbb{P}^0\varepsilon^0$, and similarly for the other three axioms. □

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. They are a generalisation of the topological categories introduced in [1]. For a treatment of polynomial functors in a general category-theoretical setting, see [23]. Our exposition is based on [39, 18]. The theorem of Beck was originally proved in [4], with some of the conditions provided later on by other authors.

II. Algebra

1. Factorisations

To further develop our algebraic machinery we need the notions of a subalgebra 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. To generalise this to other categories we have to find analogues for the notion of a surjective/injective function.

Equalisers and Coequalisers

We start with an attempt that will not quite work out, but that will provide useful intuition. Our candidates for injective functions are equalisers and those for surjective functions are coequalisers.

Example. In the category Pos, a function $f : A \to B$ is a coequaliser if, and only if, it is a quotient map, i.e., surjective function such that the ordering of B is the image of the ordering of A. To see this, suppose that $f : A \to B$ is a quotient map. Then f is the coequaliser of the two projections $p, q : X \to A$, where

$$X := \{ \langle a, a' \rangle \in A \times A \mid f(a) = f(a') \} .$$

Similarly, $f : A \to B$ is an equaliser if, and only if, it is an embedding, i.e., an injective function satisfying

$$a \leq a' \quad \text{iff} \quad f(a) \leq f(a') , \quad \text{for all } a, a' \in A .$$

To see this, suppose that $f : A \to B$ is an embedding. Then f is the equaliser of the two inclusion maps $i, j : B \to Y$, where Y is the partial order obtained from $B + B$ by identifying, for every $a \in A$, the two copies of $f(a) \in B$. For the ordering $\sqsubseteq$ of Y we take the transitive closure of the relation $\leq_0 \cup \leq_1$ where $\leq_0$ and $\leq_1$ are the orderings of the two copies of B.

Let us quickly check that the relation $\sqsubseteq$ is indeed a partial ordering. First, note that

$$\sqsubseteq\ =\ (\leq_0 \circ \leq_1) \cup (\leq_1 \circ \leq_0)$$

since

$$a \leq_0 b \leq_1 c \leq_0 d \quad \text{implies} \quad b, c \in \operatorname{rng} f\,.$$

Hence, $a \leq_0 b \leq_0 c \leq_0 d$, which implies that $a \leq_0 d$.

To see that $\sqsubseteq$ is antisymmetric, suppose that $a \sqsubseteq b \sqsubseteq a$. We distinguish two cases.

- If $a \leq_0 c \leq_1 b$ and $b \leq_0 d \leq_1 a$, then $a, b, c, d \in \operatorname{rng} f$, which implies that $a \leq_0 b \leq_0 a$. Hence, $a = b$.
- If $a \leq_0 c \leq_1 b$ and $b \leq_1 d \leq_0 a$, then we have $c, d \in \operatorname{rng} f$, $d \leq_0 a \leq_0 c$, and $c \leq_1 b \leq_1 d$. Hence, $c = d$, which implies that $a = c$ and $b = c$. In particular, $a = b$. ⌟

Definition 1.1. A *kernel pair* of a morphism $f : A \to B$ is a pair $u, u' : X \to A$ of morphisms such that u, u' form the pullback of f along itself. Analogously, a *cokernel pair* of f consists of a pushout $v, v' : B \to Y$ of f along itself. ⌟

Examples. (a) In Set, the kernel pair of $f : A \to B$ consists of the two projections $X \to A$ where

$$X := \{\, \langle a, a' \rangle \in A \times A \mid f(a) = f(a') \,\}\,.$$

The cokernel pair consists of the two embeddings $v, v' : B \to Y$ where

$$Y = (B \setminus f[A]) + f[A] + (B \setminus f[A])\,.$$

(b) In Pos, the kernel pair of a monotone function $f : A \to B$ is the same as its kernel pair in Set (where $X \subseteq A \times A$ is equipped with the ordering of the product). To get the cokernel pair of f we start with its cokernel pair $v_0, v'_0 : B \to Y_0$ in Set. We equip Y_0 with the relation $\sqsubseteq$, the transitive closure of $v_0[\leq] \cup v'_0[\leq]$, where $\leq$ is the ordering of B. In general, $\sqsubseteq$ is not a partial order, but only a preorder. Let $\rho : Y_0 \to Y$ be the quotient map merging the elements in every equivalence class of $\sqsubseteq$. Then $Y \in \mathsf{Pos}$ and $v := \rho \circ v_0$ and $v' := \rho \circ v'_0$ form the cokernel pair of f.

(c) In Top, kernel pairs and cokernel pairs are constructed as in Set, by equipping the sets X and Y with a suitable topology. ⌟

Lemma 1.2. *Let $f : A \to B$ be a morphism.*

(a) *Every kernel pair of f consists of epimorphisms.*

(b) *The following statements are equivalent.*

   (1) *f is a monomorphism.*

   (2) *There is a function $u : X \to A$ such that u, u is a kernel pair of f.*

   (3) *$\mathrm{id}_A, \mathrm{id}_A$ is a kernel pair of f.*

Proof. (a) Let $u, u' : X \to A$ be a kernel pair of f. Since $f \circ \mathrm{id}_A = f \circ \mathrm{id}_A$ and u, u' is the pullback, we can find a unique morphism $\delta : A \to X$ satisfying

$$u \circ \delta = \mathrm{id}_A \quad \text{and} \quad u' \circ \delta = \mathrm{id}_A .$$

In particular, u and u' have a right-inverse, which implies that they are epimorphisms.

(b) (2) $\Rightarrow$ (1) Suppose that $u = u'$ and consider two morphisms $g, g' : C \to A$ with $f \circ g = f \circ g'$. Since u, u' is the pullback, there exists a unique morphism $\delta : C \to X$

$$u \circ \delta = g \quad \text{and} \quad u' \circ \delta = g' .$$

Hence,

$$g = u \circ \delta = u' \circ \delta = g' .$$

(1) $\Rightarrow$ (3) Clearly we have $f \circ \mathrm{id} = f \circ \mathrm{id}$. To see that id, id is the pullback, consider two morphisms $g, g' : Z \to A$ with $f \circ g = f \circ g'$. If f is a

monomorphism, it follows that $g = g'$. Hence, $g : Z \to A$ is the unique function satisfying

$$\mathrm{id}_A \circ g = g \quad \text{and} \quad \mathrm{id}_A \circ g = g' .$$

(3) ⇒ (2) Suppose that $\mathrm{id}_A, \mathrm{id}_A : A \to A$ is a kernel pair of f and let $u, u' : X \to A$ be another kernel pair. As limits are universal, we can find a (unique) isomorphism $\sigma : A \to X$ such that

$$\mathrm{id} = u \circ \sigma \quad \text{and} \quad \mathrm{id} = u' \circ \sigma .$$

Hence, $u = \sigma^{-1} = u'$. □

If the category in question has kernel pairs and coequalisers, we obtain a bijective correspondence between them. By duality, the same holds for cokernel pairs and equalisers.

Lemma 1.3. *Let $\mathcal{C}$ be a category with finite limits and colimits.*

(a) *Let $f : A \to B$ be a coequaliser and $g : B \to C$ an arbitrary morphism. Then f is the coequaliser of the kernel pair of $g \circ f$.*

(b) *Every kernel pair is the kernel pair of its coequaliser.*

Proof. (a) Let $f : A \to B$ be the coequaliser of $u, u' : X \to A$ and let $v, v' : Y \to A$ be the kernel pair of $g \circ f$. Since $g \circ f \circ u = g \circ f \circ u'$ and v, v' is the pullback, there exists a unique morphisms $\varphi : X \to Y$ satisfying

$$u = v \circ \varphi \quad \text{and} \quad u' = v' \circ \varphi .$$

To see that f is the coequaliser of v, v', consider a second morphism $h : A \to D$ with $h \circ v = h \circ v'$. Then

$$h \circ u = h \circ v \circ \varphi = h \circ v' \circ \varphi = h \circ u' .$$

As f is the coequaliser of u, u', it follows that there exists a unique morphism $\psi : B \to D$ satisfying $\psi \circ f = h$.

(b) Let $v, v' : Y \to A$ be the kernel pair of $g : A \to C$ and let $f : A \to B$ be the coequaliser of v, v'. Since $g \circ v = g \circ v'$, there exists a unique morphism

$\psi : B \to C$ satisfying $\psi \circ f = g$. To see that v, v' is the kernel pair of f, consider morphisms $u, u' : X \to A$ with $f \circ u = f \circ u'$. Then

$$g \circ u = \psi \circ f \circ u = \psi \circ f \circ u' = g \circ u' .$$

As v, v' is the kernel pair of g, it follows that there exists a unique morphism $\varphi : X \to Y$ satisfying $u = v \circ \varphi$ and $u' = v' \circ \varphi$. □

Remark. Note that (a) implies in particular that f is the coequaliser of its own kernel pair. It follows that the function mapping each coequaliser to its kernel pair forms a bijection between coequalisers and kernel pairs. ⌟

Every morphism can be factorised into equalisers and coequalisers, but we need factorisations with three parts.

Theorem 1.4. *Let $\mathcal{C}$ be a category with finite limits and colimits. Every morphism $f : A \to B$ has a factorisation $f = i \circ \varphi \circ q$ with the following properties.*

(a) *q is a coequaliser and i an equaliser.*

(b) *Given any other factorisation $f = i' \circ \varphi' \circ q'$ with a coequaliser q' and an equaliser i', there exist unique isomorphisms σ and τ such that*

$$i' \circ \tau = i , \quad \varphi' \circ \sigma = \tau \circ \varphi , \quad q' = \sigma \circ q .$$

$$\begin{tikzcd}
 & X \arrow[r, "\varphi"] \arrow[dd, "\sigma"] & Y \arrow[dr, "i"] \arrow[dd, "\tau"] & \\
A \arrow[ur, "q"] \arrow[dr, "q'"'] & & & B \\
 & X' \arrow[r, "\varphi'"'] & Y' \arrow[ur, "i'"'] &
\end{tikzcd}$$

(c) *For every factorisation $f = m \circ e$ of f into a monomorphism m and an epimorphism e, there exist unique morphisms σ and τ such that $e = \sigma \circ q$ and $m = \varphi \circ \tau$.*

Proof. Fix $f : A \to B$. Let $u, u' : X \to A$ the pull-back of f along itself, $v, v' : B \to Y$ the push-out of f along itself, let $q : A \to C$ be the coequaliser of u and u', and let $i : D \to B$ be the equaliser of v and v'.

(a) We claim there exists a (unique) morphism $\varphi : C \to D$ such that $f = i \circ \varphi \circ q$.

Since u and u' are the pull-back of f along itself, we have $f \circ u = f \circ u'$. Hence, f factorises through the coequaliser of u and u'. That is, there exists a function $g : C \to B$ with $f = g \circ q$.

$$\begin{tikzcd}
X \arrow[r, shift left, "u"] \arrow[r, shift right, "u'"'] & A \arrow[r, "f"] \arrow[d, "q"'] & B \arrow[r, shift left, "v"] \arrow[r, shift right, "v'"'] & Y \\
& C \arrow[ur, "g"] \arrow[r, "\varphi"'] & D \arrow[u, "i"'] &
\end{tikzcd}$$

Similarly, we have $v \circ f = v' \circ f$, which implies that

$$v \circ g \circ q = v' \circ g \circ q .$$

Since coequalisers are epimorphisms, it follows that $v \circ g = v' \circ g$. Consequently, g factorises through the equaliser of v and v' and there exists a morphism $\varphi : C \to D$ such that $g = i \circ \varphi$. It follows that

$$f = g \circ q = i \circ \varphi \circ q .$$

(b) Let $f = i' \circ \varphi' \circ q'$ be another factorisation where i' is an equaliser and q' a coequaliser. By Lemma 1.3 (a), q and q' are the coequalisers of the kernel pair of f. Dually, i and i' are the equalisers of the cokernel pair of f. As limits and colimits are unique, there exist unique isomorphisms σ and τ such that

$$i' \circ \tau = i \quad \text{and} \quad q' = \sigma \circ q .$$

For the third equation, note that

$$i' \circ \tau \circ \varphi \circ q = i \circ \varphi \circ q = f = i' \circ \varphi' \circ q' = i' \circ \varphi' \circ \sigma \circ q .$$

As i' is a monomorphism and q an epimorphism, it follows that $\tau \circ \varphi = \varphi' \circ \sigma$.

(c) Suppose that $f = m \circ e$ where $m : E \to B$ is a monomorphism and $e : A \to E$ an epimorphism. Then

$$m \circ e \circ u = f \circ u = f \circ u' = m \circ e \circ u' .$$

$$\begin{tikzcd}
X \arrow[r, shift left, "u"] \arrow[r, shift right, "u'"'] & A \arrow[rr, "f"] \arrow[dr, "e"] \arrow[dd, "q"'] & & B \arrow[r, shift left, "v"] \arrow[r, shift right, "v'"'] & Y \\
& & E \arrow[ur, "m"] \arrow[dr, "\tau"] & & \\
& C \arrow[ur, "\sigma"] \arrow[rr, "\varphi"'] & & D \arrow[uu, "i"'] &
\end{tikzcd}$$

Since m is a monomorphism, it follows that $e \circ u = e \circ u'$. As q is the coequaliser of u and u', we can therefore find a unique morphism $\sigma : C \to E$ such that $e = \sigma \circ q$. The existence of τ now follows by duality. □

Factorisation Systems

The problem with using factorisations into equalisers and coequalisers is that we usually want our notions of 'an injective function' and 'a surjective function' to be closed under composition. But the composition of two equalisers is not necessarily an equaliser (although in most of the categories we are interested in, this is indeed the case). Before choosing suitable generalisations, let us first take a look at the problem abstractly. We start by listing all the conditions we would like our two classes of morphisms to exhibit.

Definition 1.5. Let $\mathcal{C}$ be a category and E and M two sets of morphisms.

(a) An *EM-factorisation* of a morphism $\varphi : A \to B$ consists of two morphisms $\varepsilon \in E$ and $\mu \in M$ such that $\varphi = \mu \circ \varepsilon$.

(b) 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 $\varphi : A \to B$ 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

$$\mu \circ \delta = \psi\,, \qquad \delta \circ \varepsilon = \varphi\,.$$

$$\begin{tikzcd} A \arrow[r, "\varepsilon"] \arrow[d, "\varphi"'] & B \arrow[d, "\psi"] \arrow[dl, dashed, "\delta"'] \\ C \arrow[r, "\mu"'] & D \end{tikzcd}$$

(FS5) is often referred to as the *diagonal fill-in property.*

(c) 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 $\operatorname{im}\varphi$ and $\ker\varphi$, respectively. (We will see below that this factorisation is unique, up to isomorphism.) ⌟

Remark. Axiom (FS5) might look a bit arbitrary at first. It tells us that, if a morphism f factorises through an epimorphism $e \in E$ and through a monomorphism $m \in M$, then f is of the form $m \circ \delta \circ e$, for a unique map δ. This property can be seen as the category-theoretical analogue of the Factorisation Lemma from Universal Algebra. The following reformulation makes it more apparent: given morphisms $\varepsilon : A \to B$, $\varphi : A \to C$, $\psi : B \to C$ with $\varepsilon \in E$ satisfying $\varphi = \psi \circ \varepsilon$, it follows that

$$\ker\varphi = (\ker\psi) \circ \varepsilon \quad\text{and}\quad \operatorname{im}\varphi = \operatorname{im}\psi\,.$$

$$\begin{tikzcd} A \arrow[r, "\varepsilon"] \arrow[dr, "\varphi"'] & B \arrow[d, "\psi"] \\ & C \end{tikzcd} \qquad \begin{tikzcd} A \arrow[rr, "\varepsilon"] \arrow[dr, "\ker\varphi"'] & & B \arrow[dl, "\ker\psi"] \\ & X \arrow[dr, "\operatorname{im}\varphi = \operatorname{im}\psi"'] & \\ & & C \end{tikzcd}$$

⌟

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 $\varepsilon : A \to B$ satisfying

$$\varepsilon(a) \leq \varepsilon(b) \quad \Leftrightarrow \quad a' \leq b'\,, \quad \text{for some } a' \in \varepsilon^{-1}(a) \text{ and } b' \in \varepsilon^{-1}(b)\,.$$

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 $\mu : A \to B$ 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

$$\begin{aligned} \varepsilon(a) \leq \varepsilon(a') \quad &\Rightarrow \quad \psi(\varepsilon(a)) \leq \psi(\varepsilon(a')) \\ &\Rightarrow \quad \mu(\varphi(a)) \leq \mu(\varphi(a')) \\ &\Rightarrow \quad \varphi(a) \leq \varphi(a')\,, \end{aligned}$$

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 $\mathfrak{G}$-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 $\mathcal{D}^{\Xi}$. The proofs are straightforward.

Lemma 1.6. *Let $\langle E, M\rangle$ be a factorisation system for $\mathcal{C}$.*

(a) *$\langle M, E\rangle$ is a factorisation system for $\mathcal{C}^{\mathrm{op}}$.*

(b) *$\langle E^{\Xi}, M^{\Xi}\rangle$ is a factorisation system for $\mathcal{C}^{\Xi}$.*

Let us collect a few basic consequences of the axioms.

Lemma 1.7. *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 $\mu = \mu' \circ \sigma$.*

$$\begin{tikzcd} & B \arrow[dr, "\mu"] \arrow[dd, "\sigma"] & \\ A \arrow[ur, "\varepsilon"] \arrow[dr, "\varepsilon'"'] & & C \\ & B' \arrow[ur, "\mu'"'] & \end{tikzcd}$$

(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 $\delta' \circ \delta = \mathrm{id}$. In the same way, one can show that $\delta \circ \delta' = \mathrm{id}$.

(b) By (FS1) and (FS2), every isomorphism belongs to both E and M. Conversely, let $\sigma \in E \cap M$. Then $\mathrm{id} \circ \sigma = \sigma \circ \mathrm{id}$ and, by (FS5), there exists a unique morphism δ such that $\sigma \circ \delta = \mathrm{id}$ and $\delta \circ \sigma = \mathrm{id}$. 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

$$\mathrm{id} \circ (q \circ f) = \mu \circ \varepsilon \circ f = (\mu \circ \mu') \circ \varepsilon'.$$

By (FS5), we obtain a unique morphism δ such that

$$\delta \circ q \circ f = \varepsilon' \quad \text{and} \quad \mu \circ \mu' \circ \delta = \mathrm{id}\,.$$

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. □

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 1.8. Let $\varepsilon : A \to B$ and $\mu : C \to D$ be morphisms. We write $\varepsilon \perp \mu$ if, for all morphisms $\varphi : A \to C$ and $\psi : B \to D$ with $\psi \circ \varepsilon = \mu \circ \varphi$, there exists a unique morphism $\delta : B \to C$ such that $\delta \circ \varepsilon = \varphi$ and $\mu \circ \delta = \psi$. ⌟

Lemma 1.9. *Let $\langle E, M \rangle$ be a factorisation system.*

(a) $E = {}^{\perp}M := \{\, \varepsilon \mid \varepsilon \perp \mu \textit{ for all } \mu \in M \,\}$

(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 $\mathrm{id} \circ \eta = \mu \circ \varepsilon$, there exists a unique morphism δ with

$$\delta \circ \eta = \varepsilon \quad \text{and} \quad \mu \circ \delta = \mathrm{id}\,.$$

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 $\zeta = \mathrm{id}$ and $\zeta = \delta \circ \mu$ both satisfy these equations, it follows that $\delta \circ \mu = \mathrm{id}$. Hence, μ is an isomorphism with inverse δ. This implies that $\mu \in E$ and, thus, $\eta = \mu \circ \varepsilon \in E$. ☐

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 $M = E^{\perp}$. ⌟

Exercise 1.2. Prove that $\mathrm{id} \perp \mu$ implies that μ is a monomorphism, and $\varepsilon \perp \mathrm{id}$ implies that ε is an epimorphism. ⌟

The following remark saves us some work when checking the diagonal fill-in property.

Lemma 1.10. *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. □

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 1.11. *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 1.10 (b) that $\varphi = \delta \circ \varepsilon$.

(d) By (b) it is sufficient to prove that every morphism $\mu : A \to B$ with a left inverse $\varepsilon : B \to A$ is an equaliser. We claim that μ is the equaliser of $\mu \circ \varepsilon$ and id_B. Clearly,

$$\mu \circ \varepsilon \circ \mu = \mu \circ \mathrm{id} = \mathrm{id} \circ \mu \,.$$

For universality, suppose that $\varphi : C \to B$ is a morphism with $\mu \circ \varepsilon \circ \varphi = \mathrm{id} \circ \varphi$. Then $\varphi = \mu \circ (\varepsilon \circ \varphi)$ and this factorisation is unique since $\varphi = \mu \circ \psi$ implies that $\psi = \varepsilon \circ \mu \circ \psi = \varepsilon \circ \varphi$. □

Next, we consider closure under limits.

Definition 1.12. 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 $\tau_i : D(i) \to E(i) \in F$, for all i, the canonical map $\lim D \to \lim E$ 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 $\operatorname{colim} D \to \operatorname{colim} E$ 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. ⌟

Note that our terminology is not entirely consistent: closure under pullbacks is not the same as closure under limits that are pullbacks.

Lemma 1.13. *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 $\varepsilon_i : D(i) \to F(i)$ is an E-morphism, for each index i, and let $\hat{\varepsilon} : \operatorname{colim} D \to \operatorname{colim} F$ be the corresponding morphism between the colimits. Consider a commuting square $\psi \circ \hat{\varepsilon} = \mu \circ \varphi$ with $\mu : A \to B$ in M. We have to find a diagonal morphism $\hat{\delta} : \operatorname{colim} F \to A$. Applying (FS5) to the square $(\psi \circ \lambda_i) \circ \varepsilon_i = \mu \circ (\varphi \circ \kappa_i)$, we obtain a diagonal morphism $\delta_i : F(i) \to A$. Let $\hat{\delta} : \operatorname{colim} F \to A$ be the unique morphism with $\hat{\delta} \circ \lambda_i = \delta_i$. Then

$$\hat{\delta} \circ \hat{\varepsilon} \circ \kappa_i = \hat{\delta} \circ \lambda_i \circ \varepsilon_i = \delta_i \circ \varepsilon_i = \varphi \circ \kappa_i \,,$$

and $\quad \mu \circ \hat{\delta} \circ \lambda_i = \mu \circ \delta_i = \psi \circ \lambda_i \,.$

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 $\varepsilon : A \to B$ be an E-morphism, let $g : C \to D$ be the pushout of ε along $f : A \to C$, and let $f' : B \to D$ 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 1.10 (b) that δ is unique and that $\delta \circ \varepsilon = \varphi$. □

Corollary 1.14. *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. □

We have already seen above that a factorisation system $\langle E, M \rangle$ on $\mathcal{D}$ induces one on $\mathcal{D}^{\Xi}$. Let us show next that this factorisation system in turn induces one for $\mathbb{M}$-algebras.

Lemma 1.15. *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*

$$\begin{aligned} E^{\mathbb{M}} &:= \{\, \varepsilon \in E^{\Xi} \mid \varepsilon \text{ an } \mathbb{M}\text{-morphism} \,\}, \\ M^{\mathbb{M}} &:= \{\, \mu \in M^{\Xi} \mid \mu \text{ an } \mathbb{M}\text{-morphism} \,\} \end{aligned}$$

form a factorisation system $\langle E^{\mathbb{M}}, M^{\mathbb{M}} \rangle$ on $\mathsf{Alg}(\mathbb{M})$.

Proof. (FS1) If ε is an epimorphism of $\mathcal{D}^{\Xi}$ and an $\mathbb{M}$-morphism, it is also an epimorphism of $\mathsf{Alg}(\mathbb{M})$. Furthermore, every isomorphism of $\mathsf{Alg}(\mathbb{M})$ is also an isomorphism of $\mathcal{D}^{\Xi}$.

(FS2) analogous.

(FS3) E^{Ξ} and M^{Ξ} are closed under composition.

(FS4) Let $\varphi : \mathfrak{A} \to \mathfrak{B}$ 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 $\sigma : \mathbb{M}C \to C$ 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 $\varepsilon : A \to C$ 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. □

The Standard Factorisation System

The following factorisation system is the one we will use in the rest of this book.

Definition 1.16. 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. ⌟

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. ⌟

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 1.17. *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 $\varepsilon : \lim C \to \lim D$ 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 $\mathbb{V}\varepsilon : \lim(\mathbb{V} \circ C) \to \lim(\mathbb{V} \circ D)$ 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

$f : i \to j$ 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, $b \in \operatorname{rng} \varepsilon$.

(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, $\varepsilon' \in E$. ☐

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 $\mu : C \to A$. 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 $j : \mathbb{J}X \to A$ be the morphism corresponding to the inclusion map $X \to \mathbb{V}A$ via the adjunction.

(a) The *subobject* of A generated by X is the domain of the image

$$\langle\!\langle X \rangle\!\rangle_A := \operatorname{dom}(\operatorname{im} j)\,.$$

(b) Similarly, for an $\mathbb{M}$-algebra $\mathfrak{A}$ and $X \subseteq A$, we define the *subalgebra* of $\mathfrak{A}$ generated by X by

$$\langle\!\langle X \rangle\!\rangle_{\mathfrak{A}} := \langle C, \pi_0 \rangle \quad \text{where} \quad C := \operatorname{dom}(\operatorname{im}(\pi \circ \mathbb{M}j))$$

and the product π_0 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 $\pi_0 : \mathbb{M}C \to C$ turning $\langle\!\langle X \rangle\!\rangle_{\mathfrak{A}} = \langle C, \pi_0 \rangle$ into an $\mathbb{M}$-algebra and the inclusion map $\langle\!\langle X \rangle\!\rangle_{\mathfrak{A}} \to \mathfrak{A}$ 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 $\pi \circ \mathbb{M}j : \mathbb{M}X \to A$ and let C be the domain of μ. Similarly, let $\mu' \circ \varepsilon'$ be the EM-factorisation of $\pi \circ \mathbb{M}\mu : \mathbb{M}C \to A$ and let C' be the domain of μ'.

$$\begin{tikzcd}
\mathbb{M}X \arrow[r, "\mathbb{M}i"'] \arrow[rrr, bend left, "\mathbb{M}j"] \arrow[dr, "\varepsilon"] & \mathbb{M}C \arrow[r, "\mathbb{M}i'"'] \arrow[rr, bend left, "\mathbb{M}\mu"] \arrow[dr, "\varepsilon'"] & \mathbb{M}C' \arrow[r, "\mathbb{M}\mu'"'] & \mathbb{M}A \arrow[d, "\pi"'] \\
X \arrow[u, "\mathrm{sing}"] \arrow[r, "i"] \arrow[rrr, bend right, "j"'] & C \arrow[u, "\mathrm{sing}"'] \arrow[r, "i'"] \arrow[rr, bend right, "\mu"'] & C' \arrow[r, "\mu'"] & A
\end{tikzcd}$$

Set $i := \varepsilon \circ \mathrm{sing}$ and $i' := \varepsilon' \circ \mathrm{sing}$. Then

$$\begin{aligned}
\mu \circ i &= \mu \circ \varepsilon \circ \mathrm{sing} &&= \pi \circ \mathbb{M}j \circ \mathrm{sing} &&= \pi \circ \mathrm{sing} \circ j &&= j\,, \\
\mu' \circ i' &= \mu' \circ \varepsilon' \circ \mathrm{sing} &&= \pi \circ \mathbb{M}\mu \circ \mathrm{sing} &&= \pi \circ \mathrm{sing} \circ \mu &&= \mu\,.
\end{aligned}$$

We claim that i' is an isomorphism. For the proof, note that the two EM-factorisations

$$\mu \circ (\varepsilon \circ \mathrm{flat}) = \mu' \circ (\varepsilon' \circ \mathbb{M}\varepsilon)$$

must be isomorphic. Hence, there exists a morphism $k : C' \to C$ such that

$$\mu \circ k = \mu' \quad \text{and} \quad \varepsilon \circ \mathrm{flat} = k \circ (\varepsilon' \circ \mathbb{M}\varepsilon)\,.$$

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

$$\pi_o \circ \mathrm{sing} = k \circ \varepsilon' \circ \mathrm{sing} = k \circ i' = \mathrm{id}\,.$$

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

$$\begin{aligned} \mu \circ \pi_o \circ \mathrm{flat} &= \pi \circ \mathbb{M}\mu \circ \mathrm{flat} \\ &= \pi \circ \mathrm{flat} \circ \mathbb{M}\mathbb{M}\mu \\ &= \pi \circ \mathbb{M}\pi \circ \mathbb{M}\mathbb{M}\mu \\ &= \pi \circ \mathbb{M}(\mu \circ \pi_o) = \mu \circ \pi_o \circ \mathbb{M}\pi_o\,. \end{aligned}$$

Since $\mu \in M$, it follows that $\pi_o \circ \mathrm{flat} = \pi_o \circ \mathbb{M}\pi_o$. Finally, to see that the inclusion $\mu : C \to A$ is a morphism of $\mathbb{M}$-algebras, note that

$$\mu \circ \pi_o = \mu \circ k \circ \varepsilon' = \mu' \circ \varepsilon' = \pi \circ \mathbb{M}\mu\,. \qquad \square$$

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 $D : \mathcal{I} \to \mathcal{D}$ and a bijective morphism $\sigma : \operatorname{colim} D \to A$.*

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 $j_X : \mathbb{J}X \to A$ 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 $\sigma : \operatorname{colim} D \to A$ 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 $\operatorname{colim}(\mathbb{V} \circ D) \subseteq \bigcup_X X = \mathbb{V}A$. For injectivity, fix two elements $a, b \in \mathbb{V} \operatorname{colim} D$ with $a \neq b$. As $\mathcal{I}$ is κ-filtered, we can find some index $X \in \mathcal{I}$ with $a, b \in \operatorname{rng} \mathbb{V}\lambda_X = \mathbb{V}\lambda_X[\mathbb{V}\langle\!\langle X \rangle\!\rangle_A]$. Consequently, there are $a', b' \in \langle\!\langle X \rangle\!\rangle_A$ with $a = \lambda_X(a')$ and $b = \lambda_X(b')$. Hence,

$$a \neq b \quad \text{implies} \quad \sigma(a) = a' \neq b' = \sigma(b)\,. \qquad \square$$

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 $\langle E, M \rangle$ and let κ a regular cardinal.

(a) The *canonical κ-subobject diagram* $D : \mathcal{I} \to \mathcal{D}$ 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 $i : C \to A$ 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 $i : C \to A$ to its domain C and a morphism $\gamma : i \to j$ 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 $D : \mathcal{I} \to \mathcal{D}$, the morphisms $i \in \mathcal{I}$ form a limiting cocone from D to A. ⌟

Example. The categories Set, Pos, Met, and $\mathfrak{B}$-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, $|C| < \kappa$. 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. ⌟

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 $D : \mathcal{I} \to \mathcal{C}$ be a diagram and $(\mu_i)_i$ a cocone from D to some object B. We say that a morphism $\varphi : A \to B$ *factorises through* $(\mu_i)_i$ if there exists an index $i \in I$ such that

$$\varphi = \mu_i \circ \psi\,, \quad \text{for some morphism } \psi : A \to D(i)\,.$$

We say that such a factorisation is *essentially unique* if, given two morphisms $\psi, \psi' : A \to D(i)$ with $\mu_i \circ \psi = \varphi = \mu_i \circ \psi'$, there exist a morphism $f : i \to k$ of $\mathcal{I}$ such that $Df \circ \psi = Df \circ \psi'$.

(b) A diagram $D : \mathcal{I} \to \mathcal{C}$ 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 $D : \mathcal{I} \to \mathcal{C}$ that is κ-filtered and κ-M-generating, every morphism $f : A \to \operatorname{colim} D$ factorises essentially uniquely through $(\lambda_i)_i$. ⌟

Remark. Note that the canonical κ-subobject diagram is κ-M-generating and κ-directed. ⌟

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 $D : \mathcal{I} \to \mathcal{D}$ such that $\operatorname{colim} D = A$ 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 $\mathrm{id} : A \to A$ factorises through $(\lambda_i)_i$. Hence, there is some $i \in \mathcal{I}$ and some function $\gamma : A \to D(i)$ such that

$$\mathrm{id} = \lambda_i \circ \gamma\,.$$

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 $j : X \to \mathbb{V}D(i)$ be the inclusion map and $j_* : \mathbb{J}X \to D(i)$ 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 $\left|\mathbb{V}\lambda_i[X]\right| \leq |X| < \kappa$, the claim follows.

($\Leftarrow$) Suppose that $A = \langle\!\langle X \rangle\!\rangle_A$, for some set of size $|X| < \kappa$. Let $j : X \to \mathbb{V}A$ be the corresponding inclusion morphism and $j_* : \mathbb{J}X \to A$ 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 $D : \mathcal{I} \to \mathsf{Set}$ that is κ-filtered and κ-M-generating. Let B be its colimit and $(\lambda_i)_i$ the corresponding limiting cocone. Given a morphism $f : A \to B$, we set $f_\circ := \mathbb{V}f \circ j : X \to \mathbb{V}B$. 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 $f_\circ(x) \in \operatorname{rng} \mathbb{V}\lambda_{i_x}$. (This is possible since surjectivity of $\bigoplus_i \lambda_i$ implies that $\mathbb{V}B = \bigcup_i \operatorname{rng} \mathbb{V}\lambda_i$.) Since $\mathcal{I}$ is κ-filtered, we can find some index $k \in \mathcal{I}$ and morphisms $g_x : i_x \to k$. Then

$$f_\circ(x) \in \operatorname{rng} \mathbb{V}\lambda_{i_x} = \operatorname{rng} \mathbb{V}(\lambda_k \circ Dg_x) \subseteq \operatorname{rng} \mathbb{V}\lambda_k\,, \quad \text{for all } x \in X\,.$$

Consequently, there exists a function $f' : X \to \mathbb{V}D(k)$ such that

$$\mathbb{V}\lambda_k(f'(x)) = f_\circ(x)\,, \quad \text{for all } x \in X\,.$$

Via the adjunction $\mathbb{J} \dashv \mathbb{V}$ this function induces a morphism $f_* : \mathbb{J}X \to D(k)$. It follows that

$$\mathbb{V}\lambda_k \circ f' = f_\circ = \mathbb{V}f \circ j \quad \text{is associated with} \quad \lambda_k \circ f_* = f \circ j_*\,.$$

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 $\delta : A \to D(k)$ 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 unqiue.

Suppose that $f = \lambda_k \circ g$, for some $g : A \to D(k)$. 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 $\delta = g$. □

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 $D : \kappa \to \mathsf{Top}$ 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 $\operatorname{colim} D = X + \kappa$ 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 $X \to X + \kappa$ is continuous, but it does not factorise through any $D(i)$ since the set P is closed in $D(i)$ but not in $\mathfrak{X}$. ⌟

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 $\mathbb{F} : \mathcal{C} \to \mathcal{D}$ *preserves* limits of a certain kind if, whenever $D : \mathcal{I} \to \mathcal{C}$ 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 $D : \mathcal{I} \to \mathcal{C}$ 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. ⌟

Example. The forgetful functors from Pos, Top, Met, and $\mathfrak{G}$-Set to Set preserve and reflect κ-filtered colimits. ⌟

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, $D : \mathcal{I} \to \mathcal{D}^\Xi$ a κ-filtered diagram, and X a set of size $|X| < \kappa$. There exists a unique bijective morphism*

$$\sigma : \operatorname{colim}(D^X) \to (\operatorname{colim} D)^X$$

satisfying

$$\sigma \circ \mu_i = \lambda_i^X\,, \quad \textit{for all } i \in \mathcal{I}\,,$$

where $(\lambda_i)_i$ is the limiting cocone of D and $(\mu_i)_i$ the one of D^X.

Proof. Set $A := \operatorname{colim} D$ and $B := \operatorname{colim} D^X$. 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 $\sigma : B \to A^X$ 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 $a_x \in A = \operatorname{colim} D$, 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

$$(a_x)_x = \lambda_k^X(b) = (\sigma \circ \mu_k)(b) \in \operatorname{rng} \sigma\,.$$

For injectivity, consider elements $b, b' \in B$ with $\sigma(b) = \sigma(b')$. Since $b, b' \in \operatorname{colim} D^X$, 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 $f_x : i \to k_x$ with

$$D f_x(p_x(a)) = D f_x(p_x(a')).$$

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 use denote this morphism by $f : i \to k$. Then $D^X f(a) = D^X f(a')$, which implies that

$$b = \mu_i(a) = \mu_k(D^X f(a)) = \mu_k(D^X f(a')) = \mu_i(a') = b'. \qquad \square$$

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 $(-)^X : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ preserves κ-filtered colimits.*

Proof. Let $D : \mathcal{I} \to \mathcal{D}^\Xi$ be a κ-filtered diagram with colimit $A := \operatorname{colim} D$ and limiting cocone $(\lambda_i)_i$. Set $B := \operatorname{colim} D^X$ 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 $\sigma : B \to A^X$ 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 $\tau : A^X \to B$ of σ, then

$$\mathbb{V}\sigma \circ \mathbb{V}\tau = \mathbb{V}(\sigma \circ \tau) = \mathrm{id}$$

implies by bijectivity of $\mathbb{V}\sigma$ that $\mathbb{V}\tau = (\mathbb{V}\sigma)^{-1}$. Hence,

$$\mathbb{V}(\tau \circ \sigma) = \mathbb{V}\tau \circ \mathbb{V}\sigma = \mathrm{id}\,.$$

As $\mathbb{V}$ is faithful, we obtain $\tau \circ \sigma = \mathrm{id}$. Hence, σ is an isomorphism with inverse τ.

It remains to construct τ. Let $E : \mathcal{K} \to \mathcal{D}$ be the canonical κ-subobject diagram of A^X. By assumption, we have $\operatorname{colim} E = A^X$. Let $(\kappa_k)_k$ be the corresponding limiting cocone and let $p_x : A^X \to A$ be the projection to the x-th component. We obtain morphisms $p_x \circ \kappa_k : E(k) \to A$, 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 $p_x \circ \kappa_k = \lambda_{i_{k,x}} \circ \eta^x_k$, for some $i_{k,x} \in \mathcal{I}$ and some morphism $\eta^x_k : E(k) \to D(i_{k,x})$.

$$\begin{tikzcd}
D^X(i(k)) \arrow[d, "\mu_{i(k)}"'] \arrow[dr, "\lambda^X_{i(k)}"] & E(k) \arrow[l, "\eta_k"'] \arrow[r, "\eta^x_k"] \arrow[d, "\kappa_k"] & D(i(k)) \arrow[d, "\lambda_{i(k)}"] \\
B \arrow[r, "\sigma"'] & A^X \arrow[r, "p_x"'] & A
\end{tikzcd}$$

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

$$p_x \circ \kappa_k = \lambda_{i(k)} \circ \eta^x_k\,, \qquad \text{for all } x \in X\,.$$

Set $\eta_k := \langle \eta_k^x \rangle_{x \in X} : E(k) \to D(i(k))^X$. Then

$$\begin{aligned}
\lambda_{i(k)}^X \circ \eta_k &= \lambda_{i(k)}^X \circ \langle \eta_k^x \rangle_x \\
&= \langle \lambda_{i(k)}^X \circ \eta_k^x \rangle_x \\
&= \langle p_x \circ \kappa_k \rangle_x \\
&= \langle p_x \rangle_x \circ \kappa_k = \mathrm{id} \circ \kappa_k = \kappa_k \,.
\end{aligned}$$

We obtain morphisms $\mu_{i(k)} \circ \eta_k : E(k) \to B$. To see that these form a cocone form E to B, consider indices $k \leq l$ in $\mathcal{K}$. As $\mathcal{I}$ is κ-filtered, we can fix morphisms $f : i(k) \to m$ and $g : i(l) \to m$ in $\mathcal{I}$ for some m. Then

$$\lambda_m \circ Dg \circ \eta_l^x \circ E(k,l) = \lambda_{i(l)} \circ \eta_l^x \circ E(k,l)$$

and
$$\lambda_m \circ Df \circ \eta_k^x = \lambda_{i(k)} \circ \eta_k^x$$

are two factorisations of $p_x \circ \kappa_k$. By essential uniqueness, it follows that there is some morphism $h : m \to n$ 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 $\tau : A^X \to B$ with

$$\tau \circ \kappa_k = \mu_{i(k)} \circ \eta_k \,, \qquad \text{for all } k \,.$$

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 $\sigma \circ \tau = \mathrm{id}$. Hence, τ is indeed the right inverse of σ. □

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 $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ of arity at most κ preserves κ-filtered colimits.*

Proof. Let $D : \mathcal{I} \to \mathcal{D}^{\Xi}$ be κ-filtered and suppose that $\mathbb{F}X = \sum_{k \in K} X^{\mathrm{dom}(k)}$. Then

$$\begin{aligned}
\mathrm{colim}_i(\mathbb{F}D(i)) &= \mathrm{colim}_i \sum_{k \in K} D(i)^{\mathrm{dom}(k)} \\
&= \sum_{k \in K} \mathrm{colim}_i\, D(i)^{\mathrm{dom}(k)} \\
&= \sum_{k \in K} (\mathrm{colim}\, D)^{\mathrm{dom}(k)} = \mathbb{F}(\mathrm{colim}\, D)\,,
\end{aligned}$$

where the third step follows by Proposition 2.9. □

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 $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ a polynomial functor of arity at most κ. For every κ-filtered diagram $D : \mathcal{I} \to \mathcal{D}^{\Xi}$, there exists a unique bijective morphism*

$$\sigma : \mathrm{colim}\,(\mathbb{F} \circ D) \to \mathbb{F}(\mathrm{colim}\, D)\,.$$

Proof. Let $D : \mathcal{I} \to \mathcal{D}^{\Xi}$ be κ-filtered and suppose that $\mathbb{F}X = \sum_{k \in K} X^{\mathrm{dom}(k)}$. Then

$$\mathrm{colim}_i(\mathbb{F}D(i)) = \mathrm{colim}_i \sum_{k \in K} D(i)^{\mathrm{dom}(k)} = \sum_{k \in K} \mathrm{colim}_i\, D(i)^{\mathrm{dom}(k)}$$

and $$\mathbb{F}(\mathrm{colim}\, D) = \sum_{k \in K} (\mathrm{colim}\, D)^{\mathrm{dom}(k)}.$$

We can therefore obtain the desired morphism from the former to the latter by Proposition 2.8. □

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 $\mathbb{F} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ 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 $\pi : \mathbb{M}A \to A$ by a restriction $\pi \restriction \mathbb{M}^\circ A : \mathbb{M}^\circ A \to A$ 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 $\mathbb{M}^\circ$-algebra $\mathfrak{A}|_\rho := \langle A, \pi \circ \rho \rangle$. If ρ is understood, we also speak of the *$\mathbb{M}^\circ$-reduct* of $\mathfrak{A}$.

(b) For an $\mathbb{M}$-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$, we define $\varphi|_\rho : \mathfrak{A}|_\rho \to \mathfrak{B}|_\rho$ by

$$\varphi|_\rho := \varphi\,.$$ ⌟

Lemma 3.2. *Let $\rho : \langle \mathbb{M}^\circ, \mu^\circ, \varepsilon^\circ \rangle \to \langle \mathbb{M}, \mu, \varepsilon \rangle$ be morphism of monads. Then $|_\rho : \mathsf{Alg}(\mathbb{M}) \to \mathsf{Alg}(\mathbb{M}^\circ)$ is a functor.*

Proof. Let $\mathfrak{A}$ be an $\mathbb{M}$-algebra. Then $\mathfrak{A}|_\rho$ is an $\mathbb{M}^\circ$-algebra since, by the axioms of a morphism of monads,

$$(\pi \circ \rho) \circ \varepsilon^\circ = \pi \circ \varepsilon = \mathrm{id}\,,$$

$$\begin{aligned}(\pi \circ \rho) \circ \mu^\circ &= \pi \circ \mu \circ \rho \circ \mathbb{M}^\circ \rho \\ &= \pi \circ \mathbb{M}\pi \circ \rho \circ \mathbb{M}^\circ \rho = (\pi \circ \rho) \circ \mathbb{M}^\circ(\pi \circ \rho)\,.\end{aligned}$$

Similarly, let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be an $\mathbb{M}$-morphism. Then $\varphi|_\rho$ is an $\mathbb{M}^\circ$-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 replacing a monad $\mathbb{M}$ by a submonad $\mathbb{M}^\circ$, we can also go the other way.

Proposition 3.3. *Let $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ be a morphism of monads and suppose that $\mathsf{Alg}(\mathbb{M})$ has coequalisers. Then the reduct functor $|_\rho : \mathsf{Alg}(\mathbb{M}) \Rightarrow \mathsf{Alg}(\mathbb{M}^\circ)$ has a left adjoint $-^\uparrow : \mathsf{Alg}(\mathbb{M}^\circ) \Rightarrow \mathsf{Alg}(\mathbb{M})$.*

Proof. The idea is to define $\mathfrak{A}^\uparrow$ by taking the free algebra $\mathbb{M}A$ and factorising it by all identities holding in $\mathfrak{A}$. Thus, let $q : \mathbb{M}A \to \mathfrak{Q}$ with $\mathfrak{Q} = \langle Q, \sigma \rangle$ be the coequaliser (in $\mathsf{Alg}(\mathbb{M})$) of

$$\mathbb{M}\pi,\ \mathrm{flat} \circ \mathbb{M}\rho : \mathbb{M}\mathbb{M}^\circ A \to \mathbb{M}A\,.$$

(To see that this is well-defined, note that $\mathbb{M}f \circ \mathrm{flat} = \mathrm{flat} \circ \mathbb{M}\mathbb{M}f$, for every function f. Therefore, $\mathbb{M}\pi$ and $\mathbb{M}\rho$ are $\mathbb{M}$-morphisms.) We set $\mathfrak{A}^\uparrow := \mathfrak{Q}$.

$$\begin{tikzcd}
\mathbb{M}\mathbb{M}^\circ A \arrow[r, shift left, "\mathbb{M}\pi"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\rho"'] \arrow[dr, "\mathbb{M}\rho"'] & \mathbb{M}A \arrow[r, "q"] & Q \\
& \mathbb{M}\mathbb{M}A \arrow[u, "\mathrm{flat}"'] \arrow[r, "\mathbb{M}q"'] & \mathbb{M}Q \arrow[u, "\sigma"']
\end{tikzcd}$$

To define the action of $-^\uparrow$ on morphisms, let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be an $\mathbb{M}^\circ$-morphism. Then φ induces morphisms between the coequaliser diagrams associated with $\mathfrak{A}^\uparrow$ and $\mathfrak{B}^\uparrow$. We denote by $\varphi^\uparrow$ the corresponding morphism between the quotients.

$$\begin{tikzcd}
\mathbb{M}\mathbb{M}^\circ A \arrow[r, shift left, "\mathbb{M}\pi"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\rho"'] \arrow[d, "\mathbb{M}\mathbb{M}^\circ \varphi"'] & \mathbb{M}A \arrow[r, "q_A"] \arrow[d, "\mathbb{M}\varphi"] & Q_A \arrow[d, "\varphi^\uparrow"] \\
\mathbb{M}\mathbb{M}^\circ B \arrow[r, shift left, "\mathbb{M}\pi"] \arrow[r, shift right, "\mathrm{flat} \circ \mathbb{M}\rho"'] & \mathbb{M}B \arrow[r, "q_B"'] & Q_B
\end{tikzcd}$$

Note that this immediately implies that $q : \mathbb{M} \Rightarrow (-)^{\uparrow}$ is a natural transformation since, by definition of $\varphi^{\uparrow}$, we have

$$\varphi^{\uparrow} \circ q = q \circ \mathbb{M}\varphi \,.$$

Furthermore, every component q_A of q is an epimorphism: given morphisms $f, g : \mathfrak{Q} \to \mathfrak{C}$ with $f \circ q = g \circ q$, we can apply the universality of q to the morphism $\psi := f \circ q$, to obtain a unique morphism $h : \mathfrak{Q} \to \mathfrak{C}$ such that $h \circ q = \psi$. This implies that $f = h = g$.

Having defined $-^{\uparrow}$, it remains to show that it is the left adjoint of $|_{\rho}$. We define the unit $\tau : \mathrm{Id} \Rightarrow (-|_{\rho}) \circ (-^{\uparrow})$ of the adjunction by $\tau := q \circ \mathrm{sing}$. Note that τ is an $\mathbb{M}^{\circ}$-morphism $\mathfrak{A} \to \mathfrak{Q}|_{\rho}$ since

$$\begin{aligned}
\tau \circ \pi &= q \circ \mathrm{sing} \circ \pi \\
&= q \circ \mathbb{M}\pi \circ \mathrm{sing} \\
&= q \circ \mathrm{flat} \circ \mathbb{M}\rho \circ \mathrm{sing} \\
&= q \circ \mathrm{flat} \circ \mathrm{sing} \circ \rho \\
&= q \circ \rho \\
&= q \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} \circ \rho \\
&= \sigma \circ \mathbb{M}q \circ \mathbb{M}\mathrm{sing} \circ \rho \\
&= \sigma \circ \mathbb{M}\tau \circ \rho \\
&= (\sigma \circ \rho) \circ \mathbb{M}^{\circ}\tau \,.
\end{aligned}$$

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

$$\begin{aligned}
\tau \circ \varphi &= q \circ \mathrm{sing} \circ \varphi \\
&= q \circ \mathbb{M}\varphi \circ \mathrm{sing} \\
&= \varphi^{\uparrow} \circ q \circ \mathrm{sing} = \varphi^{\uparrow} \circ \tau \,,
\end{aligned}$$

where the third step follows from the fact that $q : \mathbb{M} \Rightarrow (-)^{\uparrow}$ is a natural transformation.

To define the counit, let $\mathfrak{A} = \langle A, \pi_{+} \rangle$ be an $\mathbb{M}$-algebra with reduct $\mathfrak{A}|_{\rho} = \langle A, \pi \rangle$, and let $q = q_A : \mathbb{M}A \to \mathfrak{Q}_{\mathfrak{A}|_{\rho}}$ be the coequaliser used in the definition of $(\mathfrak{A}|_\rho)^\uparrow$. Since

$$\pi_+ \circ \mathbb{M}\pi = \pi_+ \circ \mathbb{M}(\pi_+ \circ \rho) = \pi_+ \circ \mathrm{flat} \circ \mathbb{M}\rho\,,$$

there exists a unique morphism $v : Q_{\mathfrak{A}|_\rho} \to A$ with

$$v \circ q = \pi_+\,.$$

This morphism is a morphism $v : (\mathfrak{A}|_\rho)^\uparrow \to \mathfrak{A}$ of $\mathbb{M}$-algebras since

$$\begin{aligned} v \circ \sigma \circ \mathbb{M}q &= v \circ q \circ \mathrm{flat} \\ &= \pi_+ \circ \mathrm{flat} \\ &= \pi \circ \mathbb{M}\pi_+ = \pi \circ \mathbb{M}v \circ \mathbb{M}q\,. \end{aligned}$$

As $\mathbb{M}q$ is an epimorphism, we obtain $v \circ \sigma = \pi_+ \circ \mathbb{M}v$, as desired.

To prove that $-^\uparrow$ and $|_\rho$ form an adjunction, it is now sufficient to show that

$$v \circ \tau^\uparrow = \mathrm{id} \quad \text{and} \quad v|_\rho \circ \tau = \mathrm{id}\,.$$

For the second equation, let $\mathfrak{A} = \langle A, \pi_+\rangle$ be an $\mathbb{M}$-algebra with reduct $\mathfrak{A}|_\rho = \langle A, \pi\rangle$ where $\pi = \pi_+ \circ \rho$. Note that we have shown above that $\tau \circ \pi = q \circ \rho$. Hence,

$$v|_\rho \circ \tau \circ \pi = v \circ \tau \circ \pi = v \circ q \circ \rho = \pi \circ \rho\,.$$

Since π is the product of a $\mathbb{M}^\circ$-algebra, it is surjective. Consequently, it follows that $v|_\rho \circ \tau = \mathrm{id}$.

For the first equation, let $\mathfrak{A}$ be a $\mathbb{M}^\circ$-algebra with $\mathfrak{A}^\uparrow = \mathfrak{Q} = \langle Q, \sigma\rangle$ as above. Suppose that $((\mathfrak{A}^\uparrow)|_\rho)^\uparrow = \mathfrak{Q}' = \langle Q', \sigma'\rangle$, and let $q : \mathbb{M}A \to Q$ and $q' : \mathbb{M}Q \to Q'$ be the quotient morphisms used to define $\mathfrak{Q}$ and $\mathfrak{Q}'$. By definition, $\tau^\uparrow$ is the unique morphism $\mathfrak{Q} \to \mathfrak{Q}'$ satisfying $\tau^\uparrow \circ q = q' \circ \mathbb{M}\tau$.

Hence,

$$\begin{aligned}
v \circ \tau^{\uparrow} \circ q &= v \circ q' \circ \mathbb{M}\tau \\
&= \sigma \circ \mathbb{M}\tau \\
&= \sigma \circ \mathbb{M}(q \circ \mathrm{sing}) \\
&= q \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} \\
&= q\,.
\end{aligned}$$

Since q is an epimorphism, we obtain $v \circ \tau^{\uparrow} = q$. $\square$

Open Question. *In which cases is the unit morphism $\tau : \mathfrak{A} \to \mathfrak{A}^{\uparrow}|_{\rho}$ injective?*

Instead of operations, we can also add or remove sorts to an algebra.

Definition 3.4. Let $\Delta \subseteq \Xi$ be sets of sorts.

(a) The *expansion* of $A = (A_\xi)_{\xi \in \Delta} \in \mathcal{D}^{\Delta}$ to $\mathcal{D}^{\Xi}$ is the set $A \uparrow \Xi \in \mathcal{D}^{\Xi}$ defined by

$$(A \uparrow \Xi)_\xi := \begin{cases} A_\xi & \text{if } \xi \in \Delta\,, \\ 0 & \text{otherwise}\,, \end{cases}$$

where 0 denotes the initial object of $\mathcal{D}$. Similarly, we define the expansion $f \uparrow \Xi : A \uparrow \Xi \to B \uparrow \Xi$ of a morphism $f = (f_\xi)_{\xi \in \Delta} : A \to B$ by

$$(f \uparrow \Xi)_\xi := \begin{cases} f_\xi & \text{if } \xi \in \Delta\,, \\ \mathrm{id}_0 & \text{otherwise}\,. \end{cases}$$

(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 $f : A \to B$ in $\mathcal{D}^{\Xi}$, we denote the induced function $A|_\Delta \to B|_\Delta$ by $f|_\Delta$.

(c) The Δ-reduct of a functor $\mathbb{M} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ is defined by

$$\mathbb{M}|_\Delta A := (\mathbb{M}(A \uparrow \Xi))|_\Delta \quad \text{and} \quad \mathbb{M}|_\Delta f := (\mathbb{M}(f \uparrow \Xi))|_\Delta .$$

(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 $\mathcal{C}|_\Delta := \{ \mathfrak{A}|_\Delta \mid \mathfrak{A} \in \mathcal{C} \}$. ⌟

Example. We can see an ω-semigroup as an $\mathbb{M}$-algebra on Set^{Ξ} where $\Xi := \{1, \infty\}$ and $\mathbb{M}$ is the monad with

$$\mathbb{M}\langle X_1, X_\infty \rangle = \langle X_1,\ X_1^* X_\infty + X_1^\omega \rangle .$$

The reduct $\mathbb{M}|_{\{1\}} X_1 = X_1^+$ is the monad for semigroups, while $\mathbb{M}|_{\{\infty\}} X_\infty = X_\infty$ is just the identity monad. Given an ω-semigroup $\mathfrak{S} = \langle S_1, S_\infty, \pi \rangle$, the corresponding reduct are the associated semigroup $\mathfrak{S}|_{\{1\}} = \langle S_1, \pi_1 \rangle$ and the set $\mathfrak{S}|_{\{\infty\}} = \langle S_\infty, \mathrm{id} \rangle$.

Conversely, we can expand a given semigroup $\mathfrak{S} = \langle S_1, \pi \rangle$ to an ω-semigroup as follows. We take the free ω-semigroup $\langle S_1^\omega, S_1^+, \pi \rangle$ generated by S_1 and quotient by all equations holding in $\mathfrak{S}$. For instance, suppose that $\mathfrak{S}$ has two elements a, b satisfying the relations

$$ab = b , \quad a^4 = a , \quad b^7 = b^3 .$$

In the quotient we then have to identify the sequences $(ab)^\omega = b^\omega$, for example. Up to these identifications, only 8 infinite sequences remain:

$$b^\omega \quad \text{and} \quad b^k a^\omega , \quad \text{for } 0 \leq k \leq 6 .$$

Hence, we can extend $\langle S_1, \pi \rangle$ to an ω-semigroup $\langle S_1, S_\infty, \pi \rangle$ where S_∞ has 8 elements. ⌟

Let us check that these notions are well-behaved.

Lemma 3.5. *Let $\Delta \subseteq \Xi$ be sets of sorts and $\mathbb{M}$ a monad on $\mathcal{D}^{\Xi}$.*

(a) *The reduct functor $|_\Delta : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Delta}$ has the left adjoint $- \uparrow \Xi : \mathcal{D}^{\Delta} \to \mathcal{D}^{\Xi}$.*

(b) $\langle \mathbb{M}|_\Delta, (\mathrm{flat}|_\Delta \upharpoonright \mathbb{M}|_\Delta \mathbb{M}|_\Delta A|_\Delta), \mathrm{sing}|_\Delta \rangle$ *is a monad on* $\mathcal{D}^\Delta$.

(c) *The reduct functor induces a functor* $|_\Delta : \mathsf{Alg}(\mathbb{M}) \Rightarrow \mathsf{Alg}(\mathbb{M}|_\Delta)$.

(d) *If* $\mathsf{Alg}(\mathbb{M})$ *has coequalisers, the functor* $|_\Delta$ *from* (c) *has a left adjoint* $-^\uparrow : \mathsf{Alg}(\mathbb{M}|_\Delta) \Rightarrow \mathsf{Alg}(\mathbb{M})$.

(e) *If* $\mathsf{Alg}(\mathbb{M})$ *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 $A \in \mathcal{D}^\Xi$,
- $\varphi|_\Delta : A|_\Delta \to B|_\Delta$, for $\varphi : A \to B$, 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 $f : A \uparrow \Xi \to B$ to $f|_\Delta : A \to B|_\Delta$. 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 $f : A \uparrow \Xi \to B$, $g : A \to A'$, and $h : B \to B'$, 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 $e := \mathrm{id} : \mathrm{Id} \Rightarrow (- \uparrow \Xi)|_\Delta$ which forms the unit. For the counit, we define morphisms $i : A|_\Delta \uparrow \Xi \to A$ by choosing $i_\xi := \mathrm{id}_{A_\xi}$, for $\xi \in \Delta$, while, for $\xi \in \Xi \smallsetminus \Delta$, $i_\xi : 0 \to A_\xi$ is the unique morphism from the initial object. Then i forms a natural transformation $i : -|_\Delta \uparrow \Xi \Rightarrow \mathrm{Id}$, which is the counit of the adjunction.

(b), (c), (d) follow immediately by Proposition I.6.3.

(e) Let $\varepsilon : A \to B$ be an E-morphism. By Lemma 1.11, the coequalisers $\rho_A : \mathbb{M}A \to \mathfrak{A}^{\uparrow}$ and $\rho_B : \mathbb{M}B \to \mathfrak{B}^{\uparrow}$ 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 1.7 that $f^{\uparrow} \in E$. □

If the monad $\mathbb{M}$ is sufficiently well-behaved one would expect that $\mathfrak{A}^{\uparrow}|_{\Delta} = \mathfrak{A}$.

Open Question. *Under which conditions on the monad $\mathbb{M}$ is the unit $v : \mathfrak{A} \to \mathfrak{A}^{\uparrow}|_{\Delta}$ of the adjunction in (d) an isomorphism?*

4. Bialgebras

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

Definition 4.1. Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$.

(a) A $\mathbb{M}$,$\mathbb{P}$-*bialgebra* $\mathfrak{A} = \langle A, \pi, \rho \rangle$ consists of an object A and two morphisms $\pi : \mathbb{M}A \to A$ and $\rho : \mathbb{P}A \to A$ such that $\langle A, \pi \rangle$ forms an $\mathbb{M}$-algebra and $\langle A, \rho \rangle$ a $\mathbb{P}$-algebra.

(b) A $\mathbb{M}$,$\mathbb{P}$-bialgebra $\mathfrak{A}$ is a δ-*bialgebra* if

$$\pi \circ \mathbb{M}\rho = \rho \circ \mathbb{P}\pi \circ \delta\,.$$

(c) A *morphism of* $\mathbb{M}$,$\mathbb{P}$-*bialgebras* is a morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ that is both an $\mathbb{M}$-morphism and a $\mathbb{P}$-morphism. A *morphism of* δ-*bialgebras* is the same thing as a morphism of $\mathbb{M}$,$\mathbb{P}$-bialgebras.

(d) We denote the category of all δ-bialgebras and their morphisms by $\mathsf{Bialg}(\delta)$. ⌟

Example. Let $\delta : \mathbb{M}\mathbb{A} \Rightarrow \mathbb{A}\mathbb{M}$ be the distributive law from the example on page 75. 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 one. If we write the former multiplicatively and the latter additively, the bialgebra axiom

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

reduces to the familiar distributive law

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

Hence, δ-bialgebras are the same as semirings. ⌟

It turns out that a δ-bialgebra is just another formalisation for a $\mathbb{P}\mathbb{M}$-algebra.

Theorem 4.2. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$.*

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

$$\sigma = \rho \circ \mathbb{P}\pi \,.$$

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

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

(c) *This correspondence between $\mathbb{P}\mathbb{M}$-algebras and δ-bialgebras is bijective.*

(d) *A morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ is a $\mathbb{P}\mathbb{M}$-morphism if, and only if, it is a δ-bialgebra morphism.*

Proof. Remember form Theorem I.6.7 that the monad $\mathbb{P}\mathbb{M}$ has the unit map $\eta \circ \varepsilon$ and the product

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

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

$$\begin{aligned}
\sigma \circ (\eta \circ \varepsilon) &= \rho \circ \mathbb{P}\pi \circ \eta \circ \varepsilon \\
&= \rho \circ \eta \circ \pi \circ \varepsilon \\
&= \mathrm{id} \circ \mathrm{id}\,, \\
\sigma \circ \mathbb{PM}\sigma &= \rho \circ \mathbb{P}\pi \circ \mathbb{PM}(\rho \circ \mathbb{P}\pi) \\
&= \rho \circ \mathbb{P}(\pi \circ \mathbb{M}\rho) \circ \mathbb{PMP}\pi \\
&= \rho \circ \mathbb{P}(\rho \circ \mathbb{P}\pi \circ \delta) \circ \mathbb{PMP}\pi \\
&= \rho \circ \mathbb{P}\rho \circ \mathbb{PP}\pi \circ \mathbb{PPM}\pi \circ \mathbb{P}\delta \\
&= \rho \circ \nu \circ \mathbb{PP}(\pi \circ \mu) \circ \mathbb{P}\delta \\
&= (\rho \circ \mathbb{P}\pi) \circ (\nu \circ \mathbb{PP}\mu \circ \mathbb{P}\delta) \\
&= \sigma \circ \kappa\,.
\end{aligned}$$

(b) Let $\mathfrak{A} = \langle A, \sigma \rangle$ be a $\mathbb{PM}$-algebra. Then $\langle A, \sigma \circ \eta \rangle$ is the η-reduct of $\mathfrak{A}$. By Lemma 3.2, it is therefore also an $\mathbb{M}$-algebra. Furthermore, $\langle A, \sigma \circ \mathbb{P}\varepsilon \rangle$ is a $\mathbb{P}$-algebra since

$$\begin{aligned}
\rho \circ \eta &= \sigma \circ \mathbb{P}\varepsilon \circ \eta \\
&= \sigma \circ \eta \circ \varepsilon \\
&= \mathrm{id}\,, \\
\rho \circ \mathbb{P}\rho &= \sigma \circ \mathbb{P}\varepsilon \circ \mathbb{P}(\sigma \circ \mathbb{P}\varepsilon) \\
&= \sigma \circ \mathbb{PM}(\sigma \circ \mathbb{P}\varepsilon) \circ \mathbb{P}\varepsilon \\
&= \sigma \circ \kappa \circ \mathbb{P}(\mathbb{MP}\varepsilon \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{PP}\mu \circ \mathbb{P}\delta \circ \mathbb{P}(\mathbb{MP}\varepsilon \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{P}(\mathbb{P}\mu \circ \mathbb{PM}\varepsilon \circ \delta \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{P}(\delta \circ \varepsilon) \\
&= \sigma \circ \nu \circ \mathbb{PP}\varepsilon \\
&= \sigma \circ \mathbb{P}\varepsilon \circ \nu \\
&= \rho \circ \nu\,.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
\rho \circ \mathbb{P}\pi \circ \delta &= (\sigma \circ \mathbb{P}\varepsilon) \circ \mathbb{P}(\sigma \circ \eta) \circ \delta \\
&= \sigma \circ \mathbb{P}\mathbb{M}\sigma \circ \mathbb{P}\varepsilon \circ \mathbb{P}\eta \circ \delta \\
&= \sigma \circ \kappa \circ \mathbb{P}\varepsilon \circ \mathbb{P}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{P}\mathbb{P}\mu \circ \mathbb{P}\delta \circ \mathbb{P}\varepsilon \circ \mathbb{P}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{P}\mathbb{P}\mu \circ \mathbb{P}\mathbb{P}\varepsilon \circ \mathbb{P}\eta \circ \delta \\
&= \sigma \circ \nu \circ \mathbb{P}\eta \circ \delta \\
&= \sigma \circ \delta \\
&= \sigma \circ \mathbb{P}(\mu \circ \mathbb{M}\varepsilon) \circ \delta \\
&= \sigma \circ \mathbb{P}\mu \circ \delta \circ \mathbb{M}\mathbb{P}\varepsilon \\
&= \sigma \circ \mathbb{P}\mu \circ (\nu \circ \eta) \circ \delta \circ \mathbb{M}\mathbb{P}\varepsilon \\
&= \sigma \circ \nu \circ \mathbb{P}\mathbb{P}\mu \circ \mathbb{P}\delta \circ \eta \circ \mathbb{M}\mathbb{P}\varepsilon \\
&= \sigma \circ \kappa \circ \eta \circ \mathbb{M}\mathbb{P}\varepsilon \\
&= \sigma \circ \mathbb{P}\mathbb{M}\sigma \circ \eta \circ \mathbb{M}\mathbb{P}\varepsilon \\
&= (\sigma \circ \eta) \circ \mathbb{M}(\sigma \circ \mathbb{P}\varepsilon) \\
&= \pi \circ \mathbb{M}\rho\,.
\end{aligned}$$

(c) We have to show that the mappings from (a) and (b) are inverse to each other. First, let $\langle A, \sigma\rangle$ be a $\mathbb{P}\mathbb{M}$-algebra. The corresponding δ-bialgebra has products $\pi = \sigma \circ \eta$ and $\rho = \sigma \circ \mathbb{P}\varepsilon$. Mapping it back, we obtain the product

$$\begin{aligned}
\sigma' &= \rho \circ \mathbb{P}\pi \\
&= (\sigma \circ \mathbb{P}\varepsilon) \circ \mathbb{P}(\sigma \circ \eta) \\
&= \sigma \circ \mathbb{P}(\mathbb{M}\sigma \circ \varepsilon \circ \eta) \\
&= \sigma \circ \kappa \circ \mathbb{P}(\varepsilon \circ \eta) \\
&= \sigma\,,
\end{aligned}$$

where the last step follows by Theorem I.6.7 (4).

Conversely, let $\langle A, \pi, \rho\rangle$ be a δ-bialgebra. The corresponding $\mathbb{PM}$-algebra has the product $\sigma = \rho \circ \mathbb{P}\pi$. Mapping it back, we obtain the products

$$\begin{aligned}
\pi' &= \sigma \circ \eta \\
&= \rho \circ \mathbb{P}\pi \circ \eta \\
&= \rho \circ \eta \circ \pi \\
&= \pi\,, \\
\rho' &= \sigma \circ \mathbb{P}\varepsilon \\
&= \rho \circ \mathbb{P}\pi \circ \mathbb{P}\varepsilon \\
&= \rho\,.
\end{aligned}$$

(d) Let $\langle A, \sigma\rangle$ and $\mathfrak{B} = \langle B, \sigma'\rangle$ be $\mathbb{PM}$-algebras, and let $\langle A, \pi, \rho\rangle$ and $\langle B, \pi', \rho'\rangle$ be the corresponding δ-algebras. If $\varphi : A \to B$ is a $\mathbb{PM}$-morphism, we have

$$\begin{aligned}
\varphi \circ \pi &= \varphi \circ \sigma \circ \eta \\
&= \sigma' \circ \mathbb{PM}\varphi \circ \eta \\
&= \sigma' \circ \eta \circ \mathbb{M}\varphi \\
&= \pi' \circ \mathbb{M}\varphi\,, \\
\varphi \circ \rho &= \varphi \circ \sigma \circ \mathbb{P}\varepsilon \\
&= \sigma' \circ \mathbb{PM}\varphi \circ \mathbb{P}\varepsilon \\
&= \sigma' \circ \mathbb{P}\varepsilon \circ \mathbb{P}\varphi \\
&= \rho' \circ \mathbb{P}\varphi\,.
\end{aligned}$$

Conversely, if $\varphi : A \to B$ is a δ-bialgebra morphism, we have

$$\begin{aligned}
\varphi \circ \sigma &= \varphi \circ \rho \circ \mathbb{P}\pi \\
&= \rho' \circ \mathbb{P}\varphi \circ \mathbb{P}\pi \\
&= \rho' \circ \mathbb{P}\pi \circ \mathbb{PM}\varphi \\
&= \sigma' \circ \mathbb{PM}\varphi\,. \qquad \square
\end{aligned}$$

As a consequence the functor $\mathsf{Alg}(\mathbb{M}) \to \mathsf{Alg}(\mathbb{PM})$ from Corollary I.6.8 can be used to turn every $\mathbb{M}$-algebra into a δ-bialgebra.

Theorem 4.3. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law. The forgetful functor* $\mathsf{Bialg}(\delta) \to \mathsf{Alg}(\mathbb{M})$ *has a left adjoint, which maps an $\mathbb{M}$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the δ-bialgebra $\mathbb{P}\mathfrak{A} := \langle \mathbb{P}A, \mathbb{P}\pi \circ \delta, \nu \rangle$. The unit of this adjunction is given by the unit morphism η of the monad $\mathbb{P}$.*

Proof. Let $\mathbb{W} : \mathsf{Bialg}(\delta) \to \mathsf{Alg}(\mathbb{M})$ be the forgetful functor and let $\widehat{\mathbb{P}} : \mathsf{Alg}(\mathbb{M}) \to \mathsf{Bialg}(\delta)$ be its supposed left adjoint.

First, note that $\widehat{\mathbb{P}}$ is the functor $\mathsf{Alg}(\mathbb{M}) \to \mathsf{Alg}(\mathbb{P}\mathbb{M}) \to \mathsf{Bialg}(\delta)$ that is induced by $\mathbb{P}$ according to Corollary I.6.8 (c) and Theorem 4.2 (b). To see that, note that it maps an $\mathbb{M}$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ first to the $\mathbb{P}\mathbb{M}$-algebra $\langle \mathbb{P}A, \sigma \rangle$ with

$$\sigma = \nu \circ \mathbb{P}\mathbb{P}\pi \circ \mathbb{P}\delta\,,$$

and then to the δ-bialgebra $\langle \mathbb{P}A, \hat{\pi}, \hat{\rho} \rangle$ where

$$\begin{aligned}
\hat{\pi} &= \sigma \circ \eta \\
&= \nu \circ \mathbb{P}\mathbb{P}\pi \circ \mathbb{P}\delta \circ \eta \\
&= \mathbb{P}\pi \circ \nu \circ \mathbb{P}\delta \circ \delta \circ \mathbb{M}\eta \\
&= \mathbb{P}\pi \circ \delta \circ \mathbb{M}\nu \circ \mathbb{M}\eta \\
&= \mathbb{P}\pi \circ \delta\,, \\
\hat{\rho} &= \sigma \circ \mathbb{P}\varepsilon \\
&= \nu \circ \mathbb{P}\mathbb{P}\pi \circ \mathbb{P}\delta \circ \mathbb{P}\varepsilon \\
&= \nu \circ \mathbb{P}\mathbb{P}\pi \circ \mathbb{P}\mathbb{P}\varepsilon \\
&= \nu\,.
\end{aligned}$$

To prove that $\widehat{\mathbb{P}} \dashv \mathbb{W}$, 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

$$\begin{aligned}
\mathbb{W}\widehat{\mathbb{P}}\mathfrak{A} &= \langle \mathbb{P}A, \mathbb{P}\pi \circ \delta \rangle = \mathbb{P}\mathfrak{A}\,, \\
\widehat{\mathbb{P}}\mathbb{W}\mathfrak{B} &= \langle \mathbb{P}B, \mathbb{P}\pi' \circ \delta, \nu \rangle\,.
\end{aligned}$$

We claim that the unit of the adjunction is given by $\eta : \mathrm{Id} \to \mathbb{W}\widehat{\mathbb{P}}$ and its counit by $\tau : \widehat{\mathbb{P}}\mathbb{W} \to \mathrm{Id}$ where $\tau_{\mathfrak{B}} := \rho'$ is the $\mathbb{P}$-product of $\mathfrak{B}$. Note that they satisfy the identities

$$\tau \circ \mathbb{P}\eta = \nu \circ \mathbb{P}\eta = \mathrm{id} \quad \text{and} \quad \mathbb{W}\tau \circ \eta = \rho' \circ \eta = \mathrm{id}$$

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

$$(\mathbb{P}\pi \circ \delta) \circ \mathbb{M}\eta = \mathbb{P}\pi \circ \eta = \eta \circ \pi\,,$$

while τ is a morphism of δ-bialgebras since

$$\begin{aligned}
\tau \circ (\mathbb{P}\pi' \circ \delta) &= \rho' \circ \mathbb{P}\pi' \circ \delta \\
&= \pi' \circ \mathbb{M}\rho' \\
&= \pi' \circ \mathbb{M}\tau\,, \\
\tau \circ \nu &= \rho' \circ \nu \\
&= \rho' \circ \mathbb{M}\rho' \\
&= \rho' \circ \mathbb{M}\tau\,.
\end{aligned}$$

For naturality, we consider two morphisms $\varphi : \langle A_0, \pi_0 \rangle \to \langle A_1, \pi_1 \rangle$ and $\psi : \langle B_0, \pi'_0, \rho'_0 \rangle \to \langle B_1, \pi'_1, \rho'_1 \rangle$. Then

$$\eta \circ \varphi = \mathbb{P}\varphi \circ \eta \quad \text{and} \quad \rho'_1 \circ \mathbb{P}\psi = \psi \circ \rho'_0$$

follow from naturality of η and the fact that ψ is a $\mathbb{P}$-morphism. □

Corollary 4.4. *Let $\mathfrak{A} = \langle A, \pi \rangle$ be an $\mathbb{M}$-algebra, $\mathfrak{B} = \langle B, \pi', \rho' \rangle$ a δ-bialgebra, and $\varphi : \mathfrak{A} \to \mathfrak{B}$ an $\mathbb{M}$-morphism. Then*

$$\hat{\varphi} := \pi' \circ \mathbb{P}\varphi$$

is the unique δ-bialgebra morphism $\hat{\varphi} : \mathbb{P}\mathfrak{A} \to \mathfrak{B}$ such that

$$\varphi = \hat{\varphi} \circ \eta\,. \qquad \begin{tikzcd} \mathbb{P}\mathfrak{A} \arrow[dr, "\hat{\varphi}"] & \\ \mathfrak{A} \arrow[u, "\eta"] \arrow[r, "\varphi"'] & \mathfrak{B} \end{tikzcd}$$

Let us make two remarks that are sometimes helpful. The first one can be used to prove that a given bialgebra is a δ-bialgebra.

Lemma 4.5. *Let $\delta : \mathbb{MP} \Rightarrow \mathbb{PM}$ be a distributive law, suppose that $\mathbb{M}$ and $\mathbb{P}$ preserve surjectivity, and let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be a surjective $\mathbb{M},\mathbb{P}$-morphism between two $\mathbb{M},\mathbb{P}$-bialgebras. If $\mathfrak{A}$ is a δ-bialgebra, so is $\mathfrak{B}$.*

Proof. We have

$$\begin{aligned}\pi \circ \mathbb{M}\rho \circ \mathbb{MP}\varphi &= \pi \circ \mathbb{M}(\varphi \circ \rho)\\ &= \varphi \circ \pi \circ \mathbb{M}\rho\\ &= \varphi \circ \rho \circ \mathbb{P}\pi \circ \delta\\ &= \rho \circ \mathbb{P}(\varphi \circ \pi) \circ \delta\\ &= \rho \circ \mathbb{P}(\pi \circ \mathbb{M}\varphi) \circ \delta\\ &= \rho \circ \mathbb{P}\pi \circ \delta \circ \mathbb{MP}\varphi\,.\end{aligned}$$

Since $\mathbb{MP}\varphi$ is surjective, it follows that

$$\pi \circ \mathbb{M}\rho = \rho \circ \mathbb{P}\pi \circ \delta\,. \qquad \square$$

The second one can be used to show that a $\mathbb{P}$-morphism is in fact an $\mathbb{M},\mathbb{P}$-morphism.

Lemma 4.6. *Let $\delta : \mathbb{MP} \Rightarrow \mathbb{PM}$ be a distributive law, $\mathfrak{A}$ and $\mathfrak{B}$ δ-bialgebras, and $C \subseteq A$ a set generating $\mathfrak{A}$ via the $\mathbb{P}$-product ρ. A $\mathbb{P}$-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ is an $\mathbb{M},\mathbb{P}$-morphism if, and only if,*

$$\varphi(\pi(t)) = \pi(\mathbb{M}\varphi(t))\,, \quad \textit{for all } t \in \mathbb{M}C\,.$$

Proof. $(\Rightarrow)$ is trivial. For $(\Leftarrow)$, fix a term $t \in \mathbb{M}A$ and let $i : C \to B$ be the inclusion map. As C is a set of generators, there exists a term $s \in \mathbb{MP}C$ with $t = \mathbb{M}(\rho \circ \mathbb{P}i)(s)$. Since

$$\varphi \circ \pi \circ \mathbb{M}i = \pi \circ \mathbb{M}(\varphi \circ i)\,,$$

we have

$$\begin{aligned}
\varphi(\pi(t)) &= (\varphi \circ \pi \circ \mathbb{M}(\rho \circ \mathbb{P}i))(s) \\
&= (\varphi \circ \rho \circ \mathbb{P}\pi \circ \delta \circ \mathbb{M}\mathbb{P}i)(s) \\
&= (\rho \circ \mathbb{P}\varphi \circ \mathbb{P}\pi \circ \delta \circ \mathbb{M}\mathbb{P}i)(s) \\
&= (\rho \circ \mathbb{P}\varphi \circ \mathbb{P}\pi \circ \mathbb{P}\mathbb{M}i \circ \delta)(s) \\
&= (\rho \circ \mathbb{P}\pi \circ \mathbb{P}\mathbb{M}\varphi \circ \mathbb{P}\mathbb{M}i \circ \delta)(s) \\
&= (\rho \circ \mathbb{P}\pi \circ \delta \circ \mathbb{M}\mathbb{P}(\varphi \circ i))(s) \\
&= (\pi \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{P}(\varphi \circ i))(s) \\
&= (\pi \circ \mathbb{M}(\varphi \circ \rho \circ \mathbb{P}i))(s) \\
&= \pi(\mathbb{M}\varphi(t)) . \qquad \square
\end{aligned}$$

For our second construction of δ-bialgebras, we consider an embedding of an $\mathbb{M}$-algebras into a $\mathbb{P}$-algebra.

Definition 4.7. Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$.

A function $j : \mathfrak{A} \to \mathfrak{B}$ from an $\mathbb{M}$-bialgebra $\mathfrak{A} = \langle A, \pi \rangle$ to a $\mathbb{P}$-algebra $\mathfrak{B} = \langle B, \rho \rangle$ is *δ-distributive* if there exists a function $\sigma : \mathbb{M}\langle\!\langle \operatorname{rng} j \rangle\!\rangle_{\mathfrak{B}} \to B$ such that

$$\sigma \circ \mathbb{M}(\rho \circ \mathbb{P}j) = \rho \circ \mathbb{P}(j \circ \pi) \circ \delta .$$

We call σ the *product induced by j*. ⌟

A typical example of a δ-distributive function is the unit $\eta : \mathrm{Id} \Rightarrow \mathbb{P}$ of the monad $\mathbb{P}$.

Lemma 4.8. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law. The function $\eta : \mathsf{Alg}(\mathbb{M}) \Rightarrow \mathsf{Alg}(\mathbb{P})$ induced by the unit $\mathrm{Id} \Rightarrow \mathbb{P}$ of $\mathbb{P}$ is δ-distributive.*

Proof. By Corollary I.6.8 (c) and Theorem 4.2 (b), an $\mathbb{M}$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ is mapped by $\mathbb{P}$ to the $\mathbb{M}$,$\mathbb{P}$-bialgebra $\mathbb{P}\mathfrak{A} = \langle \mathbb{P}A, \hat{\pi}, \hat{\rho} \rangle$ with products

$$\hat{\pi} := \mathbb{P}\pi \circ \delta \quad \text{and} \quad \hat{\rho} := \nu ,$$

where $\nu : \mathbb{P}\mathbb{P} \Rightarrow \mathbb{P}$ is the multiplication of $\mathbb{P}$.

To see that η is δ-distributive, we set $\sigma := \hat{\pi}$. Then

$$\begin{aligned}\sigma \circ \mathbb{M}(\hat{\rho} \circ \mathbb{P}\eta) &= \mathbb{P}\pi \circ \delta \circ \mathbb{M}(\nu \circ \mathbb{P}\eta) \\ &= \mathbb{P}\pi \circ \delta \\ &= (\nu \circ \mathbb{P}\eta) \circ \mathbb{P}\pi \circ \delta \\ &= \hat{\rho} \circ \mathbb{P}(\eta \circ \pi) \circ \delta\,.\end{aligned}$$

(Note that this statement also follows immediately from Theorem 4.3, Lemma 4.9, and Proposition 4.10 (2) below.) □

Let us also remark that δ-distributive functions are closed under restrictions.

Lemma 4.9. *Let $j : \mathfrak{A} \to \mathfrak{B}$ be δ-distributive with induced product σ and let $\varphi : \mathfrak{C} \to \mathfrak{A}$ be an $\mathbb{M}$-morphism. Then $j \circ \varphi : \mathfrak{C} \to \mathfrak{B}$ is also δ-distributive with the same induced product σ.*

Proof. We have

$$\begin{aligned}\sigma \circ \mathbb{M}(\rho \circ \mathbb{P}(j \circ \varphi)) &= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mathbb{M}\mathbb{P}\varphi \\ &= \rho \circ \mathbb{P}(j \circ \pi) \circ \mathbb{P}\mathbb{M}\varphi \circ \delta \\ &= \rho \circ \mathbb{P}(j \circ \varphi \circ \pi) \circ \delta\,,\end{aligned}$$

where the last step follows since φ is a $\mathbb{M}$-morphism. □

Using the notion of δ-distributivity we can characterise δ-bialgebras as follows.

Proposition 4.10. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law and $\mathfrak{A} = \langle A, \pi, \rho \rangle$ an $\mathbb{M}$,$\mathbb{P}$-bialgebra. The following statements are equivalent.*

(1) *$\mathfrak{A}$ is a δ-bialgebra.*

(2) *ρ is an $\mathbb{M}$-morphism $\mathbb{P}\mathfrak{A} \to \mathfrak{A}$.*

(3) *$\mathfrak{A}$ can be expanded to a $\mathbb{P}\mathbb{M}$-algebra with product $\hat{\pi} := \rho \circ \mathbb{P}\pi$.*

(4) *The identity map* $\mathrm{id} : \mathfrak{A} \to \mathfrak{A}$ *is* δ*-distributive.*

(5) *There exists an* $\mathbb{M}$*-morphism* $j : \mathfrak{C} \to \mathfrak{A}$ *such that* j *is* δ*-distributive with induced product* π *and the function* $\rho \circ \mathbb{P}j : \mathbb{P}C \to A$ *belongs to some set* E *of epimorphisms that is preserved by both* $\mathbb{M}$ *and* $\mathbb{P}$*.*

Proof. Suppose that the monads are $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$.

(1) $\Rightarrow$ (3) follows by Theorem 4.2 (a).

(3) $\Rightarrow$ (1) follows by Theorem 4.2 (b) since

$$\begin{aligned} \hat{\pi} \circ \eta &= \rho \circ \mathbb{P}\pi \circ \eta = \rho \circ \eta \circ \pi = \pi \,, \\ \hat{\pi} \circ \mathbb{P}\varepsilon &= \rho \circ \mathbb{P}\pi \circ \mathbb{P}\varepsilon = \rho \,. \end{aligned}$$

(1) $\Leftrightarrow$ (2) Recall from Corollary I.6.8 that $\mathbb{P}\mathfrak{A}$ is an $\mathbb{M}$-algebra with product $\hat{\pi} := \mathbb{P}\pi \circ \delta$. Consequently,

$$\begin{aligned} & \rho \text{ is an } \mathbb{M}\text{-morphism } \mathbb{P}\mathfrak{A} \to \mathfrak{A} \\ \text{iff}\quad & \pi \circ \mathbb{M}\rho = \rho \circ \hat{\pi} \\ \text{iff}\quad & \pi \circ \mathbb{M}\rho = \rho \circ \mathbb{P}\pi \circ \delta \\ \text{iff}\quad & \mathfrak{A} \text{ is a } \delta\text{-bialgebra.} \end{aligned}$$

(1) $\Rightarrow$ (5) We take for j the identity morphism $\mathrm{id} : \mathfrak{A} \to \mathfrak{A}$. Then $\rho \circ \mathbb{P}\mathrm{id} = \rho$ is surjective since $\rho \circ \eta = \mathrm{id}$. Furthermore, setting $\sigma := \pi$, we have

$$\begin{aligned} \sigma \circ \mathbb{M}(\rho \circ \mathbb{P}j) &= \pi \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{P}j \\ &= \rho \circ \mathbb{P}\pi \circ \delta \circ \mathbb{M}\mathbb{P}j \\ &= \rho \circ \mathbb{P}\pi \circ \mathbb{P}\mathbb{M}j \circ \delta \\ &= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \,. \end{aligned}$$

(5) $\Rightarrow$ (4) Let $j : \mathfrak{C} \to \mathfrak{A}$ be δ-distributive with induced product π. Then

$$\begin{aligned} \pi \circ \mathbb{M}(\rho \circ \mathbb{P}\mathrm{id}) \circ \mathbb{M}\mathbb{P}(\rho \circ \mathbb{P}j) &= \pi \circ \mathbb{M}(\rho \circ \mathbb{P}(\rho \circ \mathbb{P}j)) \\ &= \pi \circ \mathbb{M}(\rho \circ \nu \circ \mathbb{P}\mathbb{P}j) \\ &= \pi \circ \mathbb{M}(\rho \circ \mathbb{P}j \circ \nu) \end{aligned}$$

$$
\begin{aligned}
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mathbb{M}\nu \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \nu \circ \mathbb{P}\delta \circ \delta \\
&= \rho \circ \nu \circ \mathbb{P}\mathbb{P}(j \circ \pi) \circ \mathbb{P}\delta \circ \delta \\
&= \rho \circ \mathbb{P}(\rho \circ \mathbb{P}(j \circ \pi) \circ \delta) \circ \delta \\
&= \rho \circ \mathbb{P}(\pi \circ \mathbb{M}(\rho \circ \mathbb{P}j)) \circ \delta \\
&= \rho \circ \mathbb{P}(\mathrm{id} \circ \pi) \circ \delta \circ \mathbb{M}\mathbb{P}(\rho \circ \mathbb{P}j)\,.
\end{aligned}
$$

Since $\rho \circ \mathbb{P}j \in E$ and $\mathbb{M}$ and $\mathbb{P}$ preserve E, it follows that

$$\pi \circ \mathbb{M}(\rho \circ \mathbb{P}\mathrm{id}) = \rho \circ \mathbb{P}(\mathrm{id} \circ \pi) \circ \delta\,. \qquad \square$$

Corollary 4.11. *Let $\mathfrak{A} = \langle A, \pi, \rho \rangle$ be a δ-bialgebra. Then $\rho : \mathbb{P}\mathfrak{A} \to \mathfrak{A}$ is a morphism of $\mathbb{M}$,$\mathbb{P}$-bialgebras.*

Proof. We have already shown in Proposition 4.10 that ρ is an $\mathbb{M}$-morphism. The fact that it is also a $\mathbb{P}$-morphism follows from the associative law.

$$\rho \circ \mathbb{P}\rho = \rho \circ \nu\,. \qquad \square$$

Next let us show that, if we have a δ-distributive function $j : \mathfrak{A} \to \mathfrak{B}$ whose range $\mathrm{rng}\, j$ generates $\mathfrak{B}$ (via the $\mathbb{P}$-algebra product), we can lift the $\mathbb{M}$-algebra product from $\mathfrak{A}$ to $\mathfrak{B}$.

Proposition 4.12. *Let $\delta : \mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ be a distributive law between the monads $\langle \mathbb{M}, \mu, \varepsilon \rangle$ and $\langle \mathbb{P}, \nu, \eta \rangle$ and suppose that E is a class of epimorphisms preserved by both $\mathbb{M}$ and $\mathbb{P}$. Let $j : \mathfrak{A} \to \mathfrak{B}$ be a δ-distributive function from an $\mathbb{M}$,$\mathbb{P}$-bialgebra $\mathfrak{A} = \langle A, \pi, \rho \rangle$ to a $\mathbb{P}$-algebra $\mathfrak{B} = \langle B, \rho \rangle$ such that $\rho \circ \mathbb{P}j \in E$.*

Then $\mathfrak{B}$ has a unique expansion to a δ-bialgebra such that j is a morphism of δ-bialgebras.

Proof. Let $\sigma : \mathbb{M}B \to B$ be the product induced by j. We start by proving that $\langle B, \sigma \rangle$ is an $\mathbb{M}$-algebra. Note that

$$\begin{aligned}
\sigma \circ \varepsilon \circ (\rho \circ \mathbb{P}j) &= \sigma \circ \mathbb{M}(\rho \circ \eta) \circ \varepsilon \circ (\rho \circ \mathbb{P}j) \\
&= \sigma \circ \mathbb{M}(\rho \circ \eta) \circ \mathbb{M}(\rho \circ \mathbb{P}j) \circ \varepsilon \\
&= \sigma \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{P}(\rho \circ \mathbb{P}j) \circ \mathbb{M}\eta \circ \varepsilon \\
&= \sigma \circ \mathbb{M}\rho \circ \mathbb{M}(\nu \circ \mathbb{P}\mathbb{P}j) \circ \mathbb{M}\eta \circ \varepsilon \\
&= \sigma \circ \mathbb{M}\rho \circ \mathbb{M}(\mathbb{P}j \circ \nu) \circ \mathbb{M}\eta \circ \varepsilon \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mathbb{M}\nu \circ \mathbb{M}\eta \circ \varepsilon \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \varepsilon \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \mathbb{P}\varepsilon \\
&= \rho \circ \mathbb{P}j\,,
\end{aligned}$$

$$\begin{aligned}
\sigma \circ \mathbb{M}\sigma \circ \mathbb{M}\mathbb{M}(\rho \circ \mathbb{P}j) &= \sigma \circ \mathbb{M}(j \circ \rho \circ \mathbb{P}\pi \circ \delta) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{P}j) \circ \mathbb{M}\mathbb{P}\pi \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mathbb{M}\mathbb{P}\pi \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \mathbb{P}\mathbb{M}\pi \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \mathbb{P}\mu \circ \delta \circ \mathbb{M}\delta \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mu \\
&= \sigma \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{P}j \circ \mu \\
&= \sigma \circ \mu \circ \mathbb{M}\mathbb{M}(\rho \circ \mathbb{P}j)\,.
\end{aligned}$$

As $\rho \circ \mathbb{P}j \in E$ and $\mathbb{M}\mathbb{M}(\rho \circ \mathbb{P}j) \in E$ are epimorphisms, it follows that

$$\sigma \circ \varepsilon = \mathrm{id} \quad \text{and} \quad \sigma \circ \mathbb{M}\sigma = \sigma \circ \mu\,.$$

To see that j is an $\mathbb{M}$-morphism, note that

$$\begin{aligned}
\sigma \circ \mathbb{M}j &= \sigma \circ \mathbb{M}(\rho \circ \eta \circ j) \\
&= \sigma \circ \mathbb{M}(\rho \circ \mathbb{P}j \circ \eta) \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \delta \circ \mathbb{M}\eta \\
&= \rho \circ \mathbb{P}(j \circ \pi) \circ \eta
\end{aligned}$$

$$\begin{aligned} &= \rho \circ \eta \circ j \circ \pi \\ &= j \circ \pi \,. \end{aligned}$$

The fact that $\langle B, \sigma, \rho \rangle$ is a δ-bialgebra now follows by Proposition 4.10 (5). Hence, it remains to prove uniqueness. Suppose that $\mathfrak{B}' := \langle B, \sigma', \rho \rangle$ is another expansion of $\mathfrak{B}$ to a δ-bialgebra. Since $\mathfrak{B}$ and $\mathfrak{B}'$ are both δ-bialgebras, it follows that

$$\sigma \circ \mathbb{M}(\rho \circ \mathbb{P}j) = (\rho \circ \mathbb{P}\pi \circ \delta) \circ \mathbb{M}\mathbb{P}j = \sigma' \circ \mathbb{M}(\rho \circ \mathbb{P}j)\,.$$

Hence, the fact that $\mathbb{M}(\rho \circ \mathbb{P}j) \in E$ is an epimorphisms implies that $\sigma = \sigma'$. □

5. 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 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 5.1. Let $\mathcal{C}$ be a category with a factorisation system $\langle E, M \rangle$.

(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 $\theta : A \to B$ on A, we denote the *quotient* by $A/\theta := B$.

(c) We define a contravariant functor $\mathrm{Cong} : \mathcal{C} \to \mathsf{Pos}$ as follows. For an object $A \in \mathcal{C}$, $\mathrm{Cong}(A)$ is the partial order of congruences of A. For a morphism $\varphi : A \to B$ and a congruence $\theta : B \to C$, we define

$$\mathrm{Cong}(\varphi)(\theta) := \ker(\theta \circ \varphi) .$$

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 $\mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$ for the former and $\mathrm{Cong}(A)$ 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 5.2. *Let $\varphi : A \to B$, $q : B \to C$, and $q' : B \to C'$ be morphisms. Then*

$$\ker q \leq \ker q' \quad \textit{implies} \quad \ker(q \circ \varphi) \leq \ker(q' \circ \varphi) .$$

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, $\ker(q \circ \varphi) = \varepsilon \leq \varepsilon' = \ker(q' \circ \varphi)$. □

Computing infima and suprema of congruences is fortunately straightforward.

Lemma 5.3. *Suppose that $\mathcal{C}$ has products of size κ and let $\theta_i : A \to B_i$, 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 $p_k : \prod_i B_i \to B_k$ 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 $q : A \to C$ is some epimorphism with

$$q \leq \theta_i , \quad \text{for all } i < \kappa .$$

Then there exist morphisms $f_i : C \to B_i$ 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 $g : C \to \prod_i B_i$ with $p_i \circ g = f_i$, for all i. It follows that

$$p_i \circ g \circ q = f_i \circ q_i = \theta_i .$$

Hence, $g \circ q = \langle \theta_i \rangle_i$, which implies that $q \leq \langle \theta_i \rangle_i$, as desired. □

Exercise 5.1. Show that $\sup_{i<\kappa} \theta_i = \ker \bigoplus_i \theta_i$, if the coproduct is defined. ⌟

Lemma 5.4. *Let $\mathcal{C}$ be a category with products of size less than κ. Then* $\mathrm{Cong} : \mathcal{C} \to \mathsf{CPO}_\kappa$ *is a contravariant functor.*

Proof. Let $f : A \to B$ be a morphism and let H be a set of congruences of B of size $|H| < \kappa$. Then

$$\begin{aligned}
\inf \{ \mathrm{Cong}(f)(\eta) \mid \eta \in H \} &= \big\langle \mathrm{Cong}(f)(\eta) \big\rangle_{\eta \in H} \\
&= \ker \langle \eta \circ f \rangle_{\eta \in H} \\
&= \ker (\langle \eta \rangle_{\eta \in H} \circ f) \\
&= \ker ((\inf H) \circ f) \\
&= \mathrm{Cong}(f)(\inf H) .
\end{aligned}$$

□

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 $\mathrm{Cong}(A)$ 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 5.5. 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 $f : A \to B$ 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 5.6. *Let $A \in \mathsf{Set}$ be a set. $\mathrm{Cong}(A)$ is isomorphic to the lattice of all equivalence relations on A ordered by inclusion. This isomorphism maps a surjective function $\varepsilon : A \to C$ 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,

$$\begin{aligned} \langle a, b\rangle \in \ker \varepsilon \quad &\Rightarrow \quad \varepsilon(a) = \varepsilon(b) \\ &\Rightarrow \quad \varepsilon'(a) = \varepsilon'(b) \quad \Rightarrow \quad \langle a, b\rangle \in \ker \varepsilon'\,. \end{aligned}$$

Conversely, If $\ker \varepsilon \subseteq \ker \varepsilon'$, then

$$\varepsilon(a) = \varepsilon(b) \quad \text{implies} \quad \varepsilon'(a) = \varepsilon'(b)\,.$$

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 $p_\sim : A \to A/{\sim}$ is the projection mapping each $a \in A$ to its $\sim$-class $[a]_\sim$. □

Congruences in Pos

Next, let us take a look at Pos. As seen above there are two natural factorisation system 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 $\varphi : A \to B$ 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 5.7. Let $\langle A, \leq \rangle$ be a partial order.

(a) A *congruence ordering* on A is a preorder $\sqsubseteq\, \subseteq A \times A$ with $\leq\, \subseteq\, \sqsubseteq$.

(b) The *kernel* of a function $f : A \to B$ 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* $q : A \to A/{\sqsubseteq}$ maps $a \in A$ to $[a]_\sqsubseteq$. ⌟

As above, we obtain the following correspondence. The proof is similar to that for Set.

Proposition 5.8. *Let $A \in \mathsf{Pos}$ be a partial order.* $\mathrm{Cong}(A)$ *is isomorphic to the lattice of all congruence orderings on A ordered by inclusion. This isomorphism maps a surjective monotone function $\varepsilon : A \to C$ to its kernel* $\ker \varepsilon$.

In this terminology, the diagonal fill-in property takes the following form.

Lemma 5.9 (Factorisation Lemma). *Let $f : A \to B$ and $g : A \to C$ be functions in Pos and assume that f is surjective. Then*

$$\ker f \subseteq \ker g \quad \text{iff} \quad g = h \circ f\,, \quad \textit{for some } h : B \to C\,.$$

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 $g = g \circ e = g \circ r \circ f = h \circ f$. ☐

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 $\mathfrak{A}$ a Σ-algebra. A congruence of $\mathfrak{A}$ is an equivalence relation $\sim \subseteq A \times A$ such that, for every function $f \in \Sigma$,

$$a_i \sim b_i\,, \quad \text{for all } i\,, \qquad \text{implies} \qquad f(\bar{a}) \sim f(\bar{b})\,.$$

This is equivalent to $\sim$ inducing a subalgebra of $\mathfrak{A} \times \mathfrak{A}$. ⌟

The definition of congruences for $\mathbb{M}$-algebras is similar. We have already seen in Lemma 1.15 how to lift a factorisation system from $\mathcal{D}$ to $\mathsf{Alg}(\mathbb{M})$. This immediately gives us the notion of a congruence for $\mathbb{M}$-algebras.

Definition 5.10. An $\mathbb{M}$-*congruence* is a congruence in the category $\mathsf{Alg}(\mathbb{M})$. ⌟

We start with a criterion for an congruence $\varepsilon \in \mathcal{D}^{\Xi}$ to be an $\mathbb{M}$-congruence.

Lemma 5.11. *Let $\langle E, M\rangle$ be a factorisation system where E and M are closed under the monad $\mathbb{M}$. Let $\mathfrak{A}$ be an $\mathbb{M}$-algebra and $f : A \to B$ a morphism of $\mathcal{D}^{\Xi}$. 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 $\varepsilon : \mathfrak{A} \to \mathfrak{C}$ is an $\mathbb{M}$-algebra morphism. Then

$$\pi \circ \mathbb{M}\varepsilon = \varepsilon \circ \pi \qquad \text{implies} \qquad \mathbb{M}\varepsilon \leq \varepsilon \circ \pi\,.$$

Note that $\mathbb{M}f = \mathbb{M}\mu \circ \mathbb{M}\varepsilon$ is the EM-factorisation of $\mathbb{M}f$. Furthermore, by Lemma 1.11, $\pi \circ \mathrm{sing} = \mathrm{id}$ 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 $\varepsilon := \ker f : A \to C$. Since $\mathbb{M}\varepsilon = \ker \mathbb{M}f \leq \ker(f \circ \pi) = \varepsilon \circ \pi$, there exists a morphism $\rho : \mathbb{M}C \to A$ with $\rho \circ \mathbb{M}\varepsilon = \varepsilon \circ \pi$. Note that

$$\begin{aligned}\rho \circ \mathbb{M}\rho \circ \mathbb{M}\mathbb{M}\varepsilon &= \rho \circ \mathbb{M}(\varepsilon \circ \pi)\\ &= \varepsilon \circ \pi \circ \mathbb{M}\pi\\ &= \varepsilon \circ \pi \circ \mathrm{flat}\\ &= \rho \circ \mathbb{M}\varepsilon \circ \mathrm{flat} = \rho \circ \mathrm{flat} \circ \mathbb{M}\mathbb{M}\varepsilon\,.\end{aligned}$$

As $\mathbb{M}\mathbb{M}\varepsilon$ is an epimorphism, it therefore follows that $\rho \circ \mathbb{M}\rho = \rho \circ \mathrm{flat}$. Consequently, $\mathfrak{C} := \langle C, \rho\rangle$ is an $\mathbb{M}$-algebra and $\varepsilon = \ker f : \mathfrak{A} \to \mathfrak{C}$ a morphism of $\mathbb{M}$-algebras. □

Lemma 5.12. $\mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$ *is closed (as a subset of* $\mathrm{Cong}(A)$*) under all infima that exists in* $\mathrm{Cong}(A)$*.*

Proof. Let $\Theta \subseteq \mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$ be a set with infimum $\theta := \inf \Theta$ (in $\mathrm{Cong}(A)$). By Lemma 5.3, we have

$$\theta = \langle \varepsilon\rangle_{\varepsilon\in\Theta} : \mathfrak{A} \to \prod_{\varepsilon\in\Theta} \mathfrak{A}/\varepsilon\,.$$

It remains to show that θ is a morphism of $\mathbb{M}$-algebras. Let π_ε be the product of $\mathfrak{A}/\varepsilon$, π_Π the one of $\prod_\varepsilon \mathfrak{A}/\varepsilon$, and let $p_\eta : \prod_\varepsilon A/\varepsilon \to A/\eta$ and $q_\eta : \prod_\varepsilon \mathbb{M}A/\varepsilon \to \mathbb{M}A/\eta$ be the projections. For every $\eta \in \Theta$, we have

$$\begin{aligned}q_\eta \circ \langle \mathbb{M}p_\varepsilon\rangle_\varepsilon \circ \mathbb{M}\langle\varepsilon\rangle_\varepsilon &= \mathbb{M}p_\eta \circ \mathbb{M}\langle\varepsilon\rangle_\varepsilon\\ &= \mathbb{M}(p_\eta \circ \langle\varepsilon\rangle_\varepsilon) = \mathbb{M}\eta = q_\eta \circ \langle\mathbb{M}\varepsilon\rangle_\varepsilon\,.\end{aligned}$$

Consequently, $\langle \mathbb{M}p_\varepsilon\rangle_\varepsilon \circ \mathbb{M}\langle\varepsilon\rangle_\varepsilon = \langle\mathbb{M}\varepsilon\rangle_\varepsilon$, which implies that

$$\begin{aligned}\theta \circ \pi &= \langle\varepsilon\rangle_\varepsilon \circ \pi\\ &= \langle\varepsilon \circ \pi\rangle_\varepsilon\\ &= \langle\pi_\varepsilon \circ \mathbb{M}\varepsilon\rangle_\varepsilon\\ &= \prod_\varepsilon \pi_\varepsilon \circ \langle\mathbb{M}\varepsilon\rangle_\varepsilon\end{aligned}$$

$$\begin{aligned}
&= \prod_{\varepsilon} \pi_\varepsilon \circ \langle \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon \\
&= \langle \pi_\varepsilon \circ \mathbb{M}p_\varepsilon \rangle_\varepsilon \circ \mathbb{M}\langle \varepsilon \rangle_\varepsilon \\
&= \langle p_\varepsilon \circ \pi_\Pi \rangle_\varepsilon \circ \mathbb{M}\theta \\
&= \pi_\Pi \circ \mathbb{M}\theta\,.
\end{aligned}$$
□

6. 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 6.1. Let $\mathcal{C}$ be a category. An object $P \in \mathcal{C}$ is *projective* with respect to a morphism $\varepsilon : A \to B$ if, for every $\varphi : P \to B$, there exists a morphism $\hat{\varphi} : P \to A$ such that $\varphi = \varepsilon \circ \hat{\varphi}$.

$$\begin{tikzcd}
P \arrow[d, dashed, "\hat{\varphi}"'] \arrow[dr, "\varphi"] & \\
A \arrow[r, "\varepsilon"'] & B
\end{tikzcd}$$
⌟

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 6.2. *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 $\varepsilon : A \to B$ if, and only if, $X = \varnothing$ or $\mathbb{V}\varepsilon$ is surjective.*

Proof. ($\Leftarrow$) If $X = \varnothing$, then X is initial in Set. Since left adjoints preserve colimits, it follows that $\mathbb{J}X$ is initial in $\mathcal{C}$. Consequently, if $f : \mathbb{J}X \to B$ is the unique morphism to B, then $\varepsilon \circ g = f$, where $g : \mathbb{J}X \to A$ is the unique morphism to A.

Now suppose that $\mathbb{V}\varepsilon$ is surjective. Then it has a right inverse $r : \mathbb{V}B \to \mathbb{V}A$. To show projectivity of $\mathbb{J}X$, consider a morphism $f : \mathbb{J}X \to B$. Let $f_0 : X \to \mathbb{V}B$ be its image under the adjunction and let $g : \mathbb{J}X \to A$ be the image of $g_0 := r \circ f_0$. As the image of $\varepsilon \circ g$ under the adjunction is

$$\mathbb{V}\varepsilon \circ g_0 = \mathbb{V}\varepsilon \circ r \circ f_0 = f_0 \,,$$

it follows that $\varepsilon \circ g = f$.

($\Rightarrow$) Suppose that $X \neq \varnothing$ and that $\mathbb{V}\varepsilon$ is not surjective. Then there exists a function $f_0 : X \to \mathbb{V}B$ with $\operatorname{rng} f_0 \cap \operatorname{rng} \mathbb{V}\varepsilon = \varnothing$. Let $f : \mathbb{J}X \to B$ be the image of f_0 under the adjunction. We claim that there is no morphism $g : \mathbb{J}X \to A$ with $\varepsilon \circ g = f$. For a contradiction, suppose otherwise. Let $g_0 : X \to \mathbb{V}A$ be the image of g under the adjunction. Then $\mathbb{V}\varepsilon \circ g_0 = f_0$, 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 6.1. 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 $A \to B$ is jointly epimorphic. (For the dual concept, see Definition III.1.9.)

(a) Show that, in the categories Set, Pos, or 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 $\varepsilon : A \to B$, then ε is an epimorphism. ⌟

More generally, we can lift projectiveness along every adjunction, but only in one direction.

Lemma 6.3. *Let $\mathbb{J} \dashv \mathbb{V}$ be an adjunction where $\mathbb{V}$ preserves epimorphisms and let $\varepsilon : A \to B$, 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 $\varepsilon : A \to B$ is an epimorphism and let $\psi : \mathbb{J}X \to B$ Then ψ corresponds to a morphism $\psi' : X \to \mathbb{V}B$. As X is projective with respect to $\mathbb{V}\varepsilon$, we can find a morphism $\rho' : X \to \mathbb{V}A$ with $\mathbb{V}\varepsilon \circ \rho = \psi'$. This morphism corresponds to a morphism $\rho : \mathbb{J}X \to A$. We claim that $\varepsilon \circ \rho = \psi$.

Note that $\psi \in \mathcal{D}(\mathbb{J}X, B)$ corresponds to $\psi' \in \mathsf{Set}(X, \mathbb{V}B)$, and $\rho \in \mathcal{D}(\mathbb{J}X, A)$ corresponds to $\rho' \in \mathsf{Set}(X, \mathbb{V}A)$. 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 $\psi = \varepsilon \circ \rho$. □

Corollary 6.4. *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 Set every set $X \in \mathsf{Set}$ is projective with respect to all surjective functions.

(b) follows by applying the lemma to (a) and the adjunction between $\mathcal{D}^{\Xi}$ and $\mathsf{Alg}(\mathbb{M})$ that is canonically associated with the monad $\mathbb{M}$. □

Definition 6.5. 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 1.7 carry over to the set $E_{\mathcal{P}}$.

Lemma 6.6. *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 (FSI), 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 $p : A \to B$ and $q : B \to C$. By Lemma 1.7, 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 $\varphi : P \to C$. Since $q \circ p \in E_{\mathcal{P}}$, we can find a morphism $\psi : P \to A$ 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 $\varepsilon_i : A_i \to B_i$, 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 $\varphi : P \to \prod_i B_i$. Let $q_i : \prod_k B_k \to B_i$ be the projection. Then $q_i \circ \varphi : P \to B_i$ and $\varepsilon_i \in E_{\mathcal{P}}$ implies that there is some morphism $\psi_i : P \to A_i$ with $\varepsilon_i \circ \psi_i = q_i \circ \varphi$. It follows that

$$\prod_i \varepsilon_i \circ \langle \psi_i \rangle_i = \langle q_i \circ \varphi \rangle_i = \varphi\,. \qquad \square$$

Example. Let $\mathcal{D}$ be a discrete category and $\mathcal{P} := \{\, \mathbb{J}X \mid X \in \mathsf{Set} \,\}$. We have shown in Corollary 6.4 that $E_{\mathcal{P}}$ contains all surjective E-morphisms. ⌟

Varieties

Let us start simply by considering classes closed under subobjects and quotients.

Definition 6.7. 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 $P \to A$, for some $P \in \mathcal{P}$.

(b) We denote the closure of $\mathcal{K}$ under $\mathcal{P}$-generated M-subobjects by

$$\mathsf{S}(\mathcal{K}) := \{\, A \mid e : A \to K\,,\ e \in M\,,\ K \in \mathcal{K}\,,\ A \text{ is } \mathcal{P}\text{-generated}\,\}\,.$$

(c) We denote the closure of $\mathcal{K}$ under all E_*-quotients by

$$\mathsf{H}(\mathcal{K}) := \{\, A \mid q : K \to A\,,\ q \in E_*\,,\ K \in \mathcal{K}\,\}\,.$$ ⌟

Example. For $\mathcal{C} = \mathsf{Alg}(\mathbb{M})$ 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. ⌟

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 6.8. *If $E_* = E$, then*

$$\mathsf{S} \circ \mathsf{H} \subseteq \mathsf{H} \circ \mathsf{S}\,.$$

Proof. Suppose that $A \in \mathsf{SH}(\mathcal{K})$. Then there exist some E_*-morphism $q : C \to B$, an M-morphism $e : A \to B$, and an E-morphism $\rho : P \to A$, for some $C \in \mathcal{K}$ and $P \in \mathcal{P}$. Since P is projective with respect to q, we can find a morphism $f : P \to C$ 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 $g : P \to D$ such that

$$q' \circ g = \rho \quad \text{and} \quad e' \circ g = f\,.$$

By Lemma 1.13, 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 $e' \circ i : K \to C$ belongs to M and $h : P \to K$ to E. Furthermore, by Lemma 1.7, $\rho, h \in E$ implies that $q' \circ i : K \to A$ also belongs to E. Consequently, $K \in \mathsf{S}(\mathcal{K})$ and $A \in \mathsf{HS}(\mathcal{K})$. ☐

Proposition 6.9. *Let $\mathcal{C}$ be a category, $\mathcal{P}, \mathcal{K}$ two classes of objects, and $E_* = E$. Then $\mathsf{HS}(\mathcal{K})$ is the closure of class $\mathcal{K}$ under E-quotients and $\mathcal{P}$-generated M-subobjects.*

Proof. $\mathsf{HS}(\mathcal{K})$ is closed under both operations since $\mathsf{HHS}(\mathcal{K}) = \mathsf{HS}(\mathcal{K})$ and $\mathsf{SHS}(\mathcal{K}) \subseteq \mathsf{HSS}(\mathcal{K}) = \mathsf{HS}(\mathcal{K})$. □

Next let us add products to our closure operations. We are both interested in finite products and in arbitrary ones.

Definition 6.10. 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

$$\mathsf{P}_\kappa(\mathcal{K}) := \{\, \prod_{i \in I} A_i \mid A_i \in \mathcal{K}\,,\ |I| < \kappa \,\}\,.$$

(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. ⌟

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. ⌟

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 6.11. *If E_* is closed under products of size less than κ, then*

$$\mathsf{P}_\kappa \mathsf{H} \subseteq \mathsf{HP}_\kappa \quad \textit{and} \quad \mathsf{SP}_\kappa \mathsf{S} \subseteq \mathsf{SP}_\kappa\,.$$

Proof. For the first inclusion, let $A := \prod_{i \in I} B_i$ be a product with $B_i \in \mathsf{H}(\mathcal{K})$ and fix E_*-morphisms $\varepsilon_i : B_i \to C_i$ with $C_i \in \mathcal{K}$. By Lemma 6.6, the product $\prod_i \varepsilon_i : \prod_i B_i \to \prod_i C_i$ belongs to E_*. Consequently, A is an E_*-quotient of $\prod_i C_i \in \mathsf{P}_\kappa(\mathcal{K})$.

For the second inclusion, let A be a $\mathcal{P}$-generated M-subobject of a product $B := \prod_{i \in I} C_i$ with $C_i \in \mathsf{S}(\mathcal{K})$. Fix an M-morphism $\mu : A \to B$ and M-morphisms $e_i : C_i \to D_i$ with $D_i \in \mathcal{K}$. By Lemma 1.13, the product $\prod_i e_i : \prod_i C_i \to \prod_i D_i$ belongs to M. Hence, so does $\prod_i e_i \circ \mu$. Consequently, A is a $\mathcal{P}$-generated M-subobject of $\prod_i D_i \in \mathsf{P}_\kappa(\mathcal{K})$. □

Theorem 6.12. *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, $\mathcal{V} = \mathsf{HSP}_\infty(\mathcal{V})$.*

(b) *A class $\mathcal{V}$ is a pseudo-variety if, and only if, $\mathcal{V} = \mathsf{HSP}_{\aleph_0}(\mathcal{V})$.*

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,

$$\begin{aligned}
\mathsf{HSP}_\kappa \circ \mathsf{HSP}_\kappa &\subseteq \mathsf{HSHP}_\kappa\mathsf{SP}_\kappa \\
&\subseteq \mathsf{HHSP}_\kappa\mathsf{SP}_\kappa \\
&= \mathsf{HSP}_\kappa\mathsf{SP}_\kappa \subseteq \mathsf{HSP}_\kappa\mathsf{P}_\kappa = \mathsf{HSP}_\kappa\,.
\end{aligned}$$

□

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. Many results can be found in [15, 16]. The section on bialgebras is based on [4], and the abstract variety theorem in Section 6 on [32].

Part B.

Language Theory

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* $\langle \wp(\Sigma^*), \cup, \cdot, \varnothing, \{\langle\rangle\} \rangle$,
- matrices $\langle S^{n \times n}, +, \cdot, 0, I \rangle$ over some semiring S.

A language is then a function $\mathbb{M}\Sigma \to \Omega$ for some alphabet Σ, or, more generally, a function $A \to \Omega$ 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 $\omega : \Omega^n \to \Omega$ and $\kappa_0, \dots, \kappa_{n-1} : X \to \Omega$, we set

$$\omega[\bar{\kappa}] := \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle, \quad \text{for } \kappa_0, \dots, \kappa_{n-1} : X \to \Omega .$$

An Ω*-operation* of arity n is a function of the form

$$\hat{\omega} : \mathcal{D}(X, \Omega)^n \to \mathcal{D}(X, \Omega) : \bar{\kappa} \mapsto \omega[\bar{\kappa}], \quad \text{for some } \omega : \Omega^n \to \Omega .$$

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

$$\mathsf{Pos}(X, \{0, 1\})^n \to \mathsf{Pos}(X, \{0, 1\})$$

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 $A \to \Omega$ 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 $\mu : A \to B$ and every morphism $\varphi : A \to S$, there exists a morphism $\hat{\varphi} : B \to S$ with $\hat{\varphi} \circ \mu = \varphi$. ⌟

Examples. (a) In Set and $\mathfrak{S}$-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 $f : A \to S$ with $A \subseteq B$ to a function $\hat{f} : B \to S$ 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 $\mathrm{id} : S \to S$ has no extension to $C \to S$, 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. ⌟

Finiteness 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 $\mathbb{V}A$ is finite.

(b) A set $A \in \mathcal{D}$ is *strongly finite-dimensional* if there exists an M-morphism $e : A \to \Omega^d$, for some $d < \omega$.

(c) A set $A \in \mathcal{D}$ is *weakly finite-dimensional* if there exists a surjective E-morphism $\varepsilon : \tilde{A} \to A$ and an M-morphism $\mu : \tilde{A} \to \Omega^d$ 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 $\mathfrak{A}$ 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 $\mathcal{D} = \mathsf{Pos}$. ⌟

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 $\mathcal{P} := \{ \mathbb{J}X \mid X \in \mathsf{Set} \}$ and $E_* := \{ \varepsilon \in E \mid \varepsilon \text{ surjective} \}$. Then every object A is $\mathcal{P}$-generated since the morphisms of the counit $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ belong to E. Consequently, it follows by Theorem II.6.12 that the closure of Ω under the above operations is given by $\mathsf{HSP}_{\aleph_0}(\{\Omega\})$, which is precisely the class of weakly finite-dimensional sets. □

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 $\mathrm{Cong}(A)$ satisfies the *descending chain condition*, that is, if every strictly descending sequence $\varepsilon_0 > \varepsilon_1 > \dots$ of congruences of A is finite. ⌟

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 $I \subseteq M$ 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. ⌟

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 $A \in \mathsf{Pos}$, there exists a surjective map from A to at least one of the above orders.

(c) Prove that an order $A \in \mathsf{Pos}$ is noetherian if, and only if, it is finite. ⌟

The closure properties of noetherian sets are quite weak.

Lemma 1.6. *If A is noetherian, so is every quotient of A.*

Proof. Let $q : A \to B$ 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. □

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 $X \in \mathsf{Set}$, 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 $\rho_{ij} : A/\varepsilon_i \to A/\varepsilon_j$ 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,

$$\mathbb{V}\rho_{ij} : \mathbb{V}(A/\varepsilon_i) \to \mathbb{V}(A/\varepsilon_j) \text{ is bijective}, \quad \text{for all } i, j \geq k\,.$$

As $\mathbb{V}(A/\varepsilon_i)$ is finite, it follows by assumption on $\mathcal{D}$ that there exists an infinite set I such that

$$\rho_{ij} : A/\varepsilon_i \cong A/\varepsilon_j \text{ is an isomorphism}, \quad \text{for all } i, j \in I\,.$$

Hence, $\varepsilon_j \leq \varepsilon_i$. A contradiction. □

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 $e : A \to \Omega^\kappa$ 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 $p_w : \Omega^\kappa \to \Omega^w$ be the projection to the components in w. Note that

$$u \subseteq w \quad \text{implies} \quad p_u \geq p_v\,.$$

Since A is noetherian, we can therefore find some finite set $w \subseteq \kappa$ such that $\ker(p_w \circ e) = \ker e = \mathrm{id}$. Consequently, $p_w \circ e \in M$ and $p_w \circ e : A \to \Omega^w$ is the desired embedding. □

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 $f, g : A \to B$,

$$k \circ f = k \circ g\,, \quad \text{for all } k : B \to C\,, \quad \text{implies} \quad f = g\,.$$

Given a factorisation system $\langle E, M\rangle$, we call C an *EM-cogenerator* if, for every $A \in \mathcal{D}$,

$$\Big(\inf_{k:A\to C} \ker k\Big) \leq \mathrm{id}_A\,,$$

where the infimum in computed in $\mathrm{Cong}(A)$. ⌟

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 $A \to \Omega^\kappa$, 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

$$\varphi := \langle k\rangle_{k:A\to\Omega} : A \to \Omega^{\mathcal{D}(A,\Omega)}$$

is a monomorphism.

(b) Let φ be the morphism from (a). By Lemma II.5.3, we have $\ker\varphi = \inf_k \ker k$. As Ω is an EM-cogenerator, it follows that $\ker\varphi = \mathrm{id}$. Hence, $\varphi \in M$. □

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 $A \to \Omega^\kappa$, for some κ. If A is noetherian, we can therefore use Lemma 1.8 to show that it is finite-dimensional. □

The existence of EM-cogenerators depends on the category in question. We start with Pos.

Lemma 1.12. *If $\Omega \in \mathsf{Pos}$ 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 $\mu_c : A \to \Omega$ by

$$\mu_c(a) := \begin{cases} v & \text{if } a \geq c\,, \\ u & \text{otherwise}\,. \end{cases}$$

Note that μ_c is monotone since

$$\begin{aligned} \mu_c(a) \nleq \mu_c(b) \quad &\Rightarrow \quad \mu_c(a) = u \quad \text{and} \quad \mu_c(b) = v \\ &\Rightarrow \quad a \geq c \quad \text{and} \quad b \ngeq c \\ &\Rightarrow \quad a \nleq b\,. \end{aligned}$$

To see that $\mu := \langle \mu_c \rangle_{c \in A}$ is an embedding $A \to \Omega^A$, 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

$$\begin{aligned} a \nleq b \quad &\Rightarrow \quad \mu_a(a) = v \quad \text{and} \quad \mu_a(b) = u \\ &\Rightarrow \quad \mu_a(a) \nleq \mu_a(b)\,. \end{aligned}$$ □

Corollary 1.13. *Suppose that $\mathcal{D} = \mathsf{Pos}$ 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 $\mathfrak{S} = \{0, 1\} \in \mathsf{Top}$ be Sierpiński space (where the closed sets are $\varnothing$, $\{0\}$, and $\{0, 1\}$). For a sober space $\Omega \in \mathsf{Top}$ the following statements are equivalent.*

(1) *Every sober space $\mathfrak{X} \in \mathsf{Top}$ can be embedded (via an injective continuous map) into Ω^κ, for some κ.*

(2) *There exists an injective continuous map $\mathfrak{S} \to \Omega$.*

Proof. (1) $\Rightarrow$ (2) Since $\mathfrak{S}$ is sober, we can use (1) to find an embedding $e : \mathfrak{S} \to \Omega^\kappa$, for some κ. Set $x := e(0)$ and $y := e(1)$, and let $p_i : \Omega^\kappa \to \Omega$ 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 $f : \mathfrak{S} \to \Omega$ 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 $\varnothing$, $\{0\}$, 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 $\mathfrak{S} \to \Omega$ induces an embedding $\mathfrak{S}^\kappa \to \Omega^\kappa$. Hence, it is sufficient to prove that every sober space $\mathfrak{X}$ can be embedded into $\mathfrak{S}^\kappa$, for some κ. Fix $\mathfrak{X}$. For a closed set $C \subseteq X$, let $\mu_C : X \to S$ 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 $\mu_C^{-1}[\varnothing] = \varnothing$, $\mu_C^{-1}[\{0\}] = C$, and $\mu_C^{-1}[S] = X$ are all closed. Setting $e := \langle \mu_C \rangle_C$ we obtain a continuous map $\mathfrak{X} \to \mathfrak{S}^T$, 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 $e(x) \neq e(y)$. □

Lemma 1.15. *Let $\Omega = [0,1] \in \mathsf{Met}$ with the metric $d(x,y) := |x - y|$. Every metric space $\mathfrak{X}$ can be embedded into Ω^X.*

Proof. Let $\mathfrak{X} \in \mathsf{Met}$. For every $z \in X$, we define a function $f_z : \mathfrak{X} \to \Omega$ via $f_z(x) := d(z,x)$. We claim that f_z is non-expansive and that the function $e := \langle f_z \rangle_{z \in X} : \mathfrak{X} \to \Omega^X$ is the desired embedding. For the first claim, note that

$$d(f_z(x), f_z(y)) = \big|d(z,x) - d(z,y)\big| \leq d(x,y)\,,$$

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

$$\begin{aligned} d(x,y) &\leq \big|d(x,x) - d(x,y)\big| \\ &= \big|f_x(x) - f_x(y)\big| \\ &\leq d\big(f_x(x), f_x(y)\big) \\ &\leq \sup_z d\big(f_z(x), f_z(y)\big) = d(e(x), e(y)) \leq d(x,y)\,. \end{aligned}$$

□

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 $q : \mathbb{J}X \to A$, for some finite set X. ⌟

Lemma 1.17. *Every finite set A is quasi-finite. If the unit $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{VJ}$ of the adjunction is an isomorphism, the converse also holds.*

Proof. By Lemma I.2.3, the counit morphism $\iota_A : \mathbb{JV}A \to A$ is surjective.

For the second statement, suppose that $\varepsilon_X : X \to \mathbb{VJ}X$ is an isomorphism and let $f : \mathbb{J}X \to A$ be a surjective function with finite X. Then $\mathbb{V}f \circ \varepsilon_X : X \to \mathbb{V}A$ 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 $\mathfrak{G}$-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. ⌟

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 $q : A \to B$ and $f : \mathbb{J}X \to A$ are surjective, so is $q \circ f : \mathbb{J}X \to B$.

For subsets, suppose that $f : \mathbb{J}X \to A$ is surjective and let $i : C \to A$ be a monomorphism. Then $\mathbb{V}i : \mathbb{V}C \to \mathbb{V}A$ is injective and so is its pullback $j : Z \to \mathbb{VJ}X$ along the function $\hat{f} : X \to \mathbb{V}A$ corresponding to f via the adjunction. Let $q : Z \to \mathbb{V}C$ be the other morphism of the pullback square and let $\hat{q} : \mathbb{J}Z \to C$ 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 $\iota : \mathbb{JV}C \to C$ is surjective. Hence, so is $\hat{q} = \iota \circ \mathbb{J}q$.

For the empty product, note that $\mathbb{V}$ preserves products. Hence, $\mathbb{V}1_{\mathcal{D}} = 1_{\mathsf{Set}}$ and the unique map $u : 1_{\mathsf{Set}} \to \mathbb{V}1_{\mathcal{D}}$ corresponds to a morphism $\hat{u} : \mathbb{J}1_{\mathsf{Set}} \to 1_{\mathcal{D}}$, which is clearly surjective.

For binary products, suppose that $f : \mathbb{J}X \to A$ and $g : \mathbb{J}Y \to B$ are surjective. Let $p : X \times Y \to X$ and $q : X \times Y \to Y$ be the two projections. Note that p and q are surjetive 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 $\langle f \circ p, g \circ q \rangle : \mathbb{J}(X \times Y) \to A \times B$. $\square$

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 $\mu : \tilde{A} \to \Omega^d$ and a surjective E-morphism $\varepsilon : \tilde{A} \to A$ with finite d. Then it follows by the preceding lemma that Ω^d, $\tilde{A}$, and A are also quasi-finite. □

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) *$\langle \mathbb{M}, \mathrm{flat}, \mathrm{sing} \rangle$ is a polynomial monad on $\mathcal{D}^{\Xi}$ 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 $\kappa : \mathbb{M}_\xi \Sigma \to \Omega$, 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 $f : A \to B$ *recognises* a function $\kappa : A_\xi \to \Omega$ if

$$\kappa = \mu \circ f\,, \quad \text{for some function } \mu : B_\xi \to \Omega\,.$$

(e) Let $f : A \to B$ and $\kappa : B_\xi \to \Omega$ be functions. The *inverse image* of κ is

$$f^{-1}[\kappa] := \kappa \circ f : A_\xi \to \Omega\,.$$

There are two common special cases. Functions of the form $f = \mathbb{M}g$, for some morphism $f : \Sigma \to \Gamma$ of Alph, are called *relabellings*. In this case, we call we call $f^{-1}[\kappa]$ an *inverse relabelling* of κ. Similarly, if $f : \mathbb{M}\Sigma \to \mathbb{M}\Gamma$ 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 $\mathcal{D} = \mathsf{Pos}$, a language $\kappa : \mathbb{M}_\xi\Sigma \to \{0, 1\}$ of course corresponds to a subset $K := \kappa^{-1}(1) \subseteq \mathbb{M}_\xi\Sigma$. Such a language is recognised by $f : \mathbb{M}\Sigma \to A$ if there is an upwards closed set $P \subseteq A$ such that $K = f^{-1}[P]$.

The inverse image of K under a function $g : \mathbb{M}\Sigma \to C$ 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 $\mathbb{J}\mathbb{V}C$. 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

$$\mathcal{K}[C] := \Big\{\, \kappa : \mathbb{M}_\xi C \to \Omega \;\Big|\; \xi \in \Xi,\ \kappa \circ \mathbb{M}\iota \in \mathcal{K}[\mathbb{J}\mathbb{V}C] \,\Big\}\,,$$

where $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ is the counit of the adjunction $\mathbb{J} \dashv \mathbb{V}$. ⌟

Example. Let Ω be a semiring and Σ an alphabet. A *weighted automaton* $\mathcal{A} = \langle \mathfrak{Q}, \Sigma, \delta, i, t \rangle$ (for words) over Ω consists of a finitely generated Ω-module $\mathfrak{Q}$, an *initial state* $i \in Q$, a *final form* $t : Q \to \Omega$, and a *transition function* $\delta : \Sigma^* \to \mathrm{End}(\mathfrak{Q})$, which is a monoid homomorphism into the monoid $\mathrm{End}(\mathfrak{Q})$ of all linear maps $\mathfrak{Q} \to \mathfrak{Q}$. Such an automaton recognises the language $\kappa : \Sigma^* \to \Omega$ defined by

$$\kappa(w) := t(\delta(w)(i))\,.$$

Note that this language is recognised by the morphism $\delta : \Sigma^* \to \mathrm{End}(\mathfrak{Q})$ since

$$\kappa = \mu \circ \delta \quad \text{where} \quad \mu(f) := t(f(i))\,. \qquad \lrcorner$$

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 $\mathbb{L}^\circ$ that maps each set A to a *language algebra* $\mathbb{L}^\circ A$, whose elements are just the Ω-languages $\kappa : A \to \Omega$.

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 $A \to \Omega$ and whose operations are all Ω-operations.

(b) For a function $f : A \to B$ in $\mathcal{D}$, we denote by $\mathbb{L}^\circ f : \mathbb{L}^\circ B \to \mathbb{L}^\circ A$ 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

$$\mathbb{L}^\circ A := \langle \mathcal{P}(A), \cup, \cap, {}^- \rangle\,.$$

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 $f : A \to B$.

To see that we only need the above operations, note that every function $[2]^n \to [2]$ can be expressed using disjunction, conjunction, and negation. Similarly, every monotone function $[2]^n \to [2]$ 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 $\varphi : \mathbb{M}\Sigma \to \mathfrak{A}$ if, and only if, $\kappa = \varphi_\xi^{-1}[\mu]$, for some $\mu \in \mathbb{L}^\circ A_\xi$.

(c) Note that, for $f : A \to B$, the function $\mathbb{L}^\circ f$ is indeed a homomorphism $\mathbb{L}^\circ B \to \mathbb{L}^\circ A$ since, for each Ω-operation $\omega : \Omega^n \to \Omega$,

$$\begin{aligned}\omega[\mathbb{L}^\circ f(\kappa_0), \dots, \mathbb{L}^\circ f(\kappa_{n-1})] &= \omega \circ \langle \kappa_0 \circ f, \dots, \kappa_{n-1} \circ f \rangle \\ &= \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle \circ f \\ &= \mathbb{L}^\circ f(\omega[\kappa_0, \dots, \kappa_{n-1}]) .\end{aligned}$$

⌟

Lemma 2.4. *If $\mu : A \to B$ is an M-morphism, then $\mathbb{L}^\circ \mu : \mathbb{L}^\circ B \to \mathbb{L}^\circ A$ is surjective.*

Proof. Let $\kappa : A \to \Omega$ be long to $\mathbb{L}^\circ A$. Since Ω is injective, there exists some function $\lambda : B \to \Omega$ with $\lambda \circ \mu = \kappa$. Thus, $\kappa = \mathbb{L}^\circ \mu(\lambda)$. □

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 $\mu : A \to \Omega^d$ and let $\kappa : A \to \Omega$ be an element of $\mathbb{L}^\circ A$. Since Ω is injective, there exists a function $\omega : \Omega^d \to \Omega$ with $\kappa = \omega \circ \mu$. Set $\lambda_i := p_i \circ \mu$ where $p_i : \Omega^d \to \Omega$ 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 $\mu_0, \dots, \mu_{n-1}$ of $\mathbb{L}^\circ A$. By assumption, we can find an Ω-operation ω with $\mathrm{id} = \omega[\bar\mu']$ for some subset $\bar\mu'$ of the generators. Suppose that $\bar\mu' = \mu_0, \dots, \mu_{d-1}$ and set

$$e := \langle \mu_0, \dots, \mu_{d-1} \rangle : A \to \Omega^d .$$

Since $\omega \circ e = \omega[\bar\mu'] = \mathrm{id}$, the function e has a left-inverse. This implies that $e \in M$. □

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 $\varepsilon : A \to B$ and $\eta : A \to C$ 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 $\rho : B \to C$ 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 $e : C \to \Omega^d$, for some $d < \omega$. Let $p_i : \Omega^d \to \Omega$ be the projection to the i-th component. Note that $p_i \circ e : C \to \Omega$ 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 $\delta : B \to C$ with $\delta \circ \varepsilon = \eta$ and $e \circ \delta = \langle \kappa_i \rangle_i$. In particular, $\varepsilon \leq \eta$. ☐

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 $p : A \times B \to A$ and $q : A \times B \to B$ 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 $e : A \to \Omega^m$ and $f : B \to \Omega^n$. Then $e \times f \in M$ and, since Ω is injective, we can find a function $\omega : \Omega^{m+n} \to \Omega$ with $\kappa = \omega \circ (e \times f)$. Consequently,

$$\begin{aligned}\kappa &= \omega \circ (e \times f) \\ &= \omega\big[(u_0 \circ e \circ p), \dots, (u_{m-1} \circ e \circ p), \\ &\qquad (v_0 \circ f \circ q), \dots, (v_{n-1} \circ f \circ q)\big],\end{aligned}$$

where $u_i : \Omega^m \to \Omega$ and $v_i : \Omega^n \to \Omega$ 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

$$\kappa \in \big\langle\!\big\langle p^{-1}[\mathbb{L}^\circ A] \cup q^{-1}[\mathbb{L}^\circ B] \big\rangle\!\big\rangle . \qquad \square$$

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 $\kappa : A_\xi \to \Omega$ 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 $\kappa : A_\xi \to \Omega$ a language.*

(a) *An E-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ 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 $\varepsilon : \mathfrak{A} \to \mathfrak{A}/\varepsilon$ 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 $\mathfrak{A}/\varepsilon$ 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 $\mu : B_\xi \to \Omega$ 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 $\mu := i \circ \rho$. □

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 $\kappa \in A_\xi \to \Omega$ be a language and $\varphi : \mathfrak{A} \to \mathfrak{M}$ an E-morphism of $\mathbb{M}$-algebras recognising κ. We call φ the *minimal morphism* of κ if, for every surjective morphism $\psi : \mathfrak{A} \to \mathfrak{B}$ recognising κ, there exists a unique morphism $\rho : \mathfrak{B} \to \mathfrak{M}$ such that $\varphi = \rho \circ \psi$. In this case we call $\mathfrak{M}$ the *minimal algebra* of κ. ⌟

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 37. 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 $\varphi : \mathbb{C}\{a\} \to \mathfrak{B}$ 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

$$w_n := \eta a^n \eta\,, \qquad \text{for } n < \omega$$

where η denotes a dense order labelled with a, and let $h_n : \mathbb{C}\{a\} \to \mathbb{C}\{a\}$ 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 $\sigma_n : \mathbb{C}\mathbb{C}\{a\} \to \mathbb{C}\{a\}$ 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 $h_n : \mathbb{C}\{a\} \to \mathfrak{A}_n$ a morphism.

For a contradiction, suppose that K has a minimal morphism $\varphi : \mathbb{C}\{a\} \to \mathfrak{M}$. Then there are morphisms $\rho_n : \mathfrak{A}_n \to \mathfrak{M}$ 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 $\kappa \in A_\xi \to \Omega$ has a minimal algebra if, and only if, there exists a greatest $\mathbb{M}$-congruence $\varepsilon \leq \ker \kappa$.*

Proof. ($\Leftarrow$) Let $\eta : \mathfrak{A} \to \mathfrak{M}$ 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, $\varepsilon : \mathfrak{A} \to \mathfrak{A}/\varepsilon$ recognises κ. By minimality of η, we can therefore find a morphism $\rho : \mathfrak{A}/\varepsilon \to \mathfrak{M}$ with $\eta = \rho \circ \varepsilon$. This implies that $\varepsilon \leq \ker(\rho \circ \varepsilon) = \ker \eta$.

($\Rightarrow$) We claim that $\varepsilon : \mathfrak{A} \to \mathfrak{A}/\varepsilon$ is the minimal morphism of κ. Hence, suppose that $\psi : \mathfrak{A} \to \mathfrak{C}$ 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 $\rho : \mathfrak{C} \to \mathfrak{A}/\varepsilon$ with $\rho \circ \psi = \varepsilon$. Uniqueness of ρ follows from the fact that ψ is surjective. □

We will prove the existence of minimal algebras in the case $\mathcal{D} = \mathsf{Pos}$ where we can work with congruence orderings. (The exact same proofs work for Set and $\mathfrak{S}$-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 $\mathrm{TC}(\theta)$. ⌟

Lemma 3.5. *Let $\mathcal{D} = \mathsf{Pos}$ and let $\mathfrak{A}$ be an $\mathbb{M}$-algebra.*

(a) $\mathrm{Cong}(A)$ *forms a complete lattice where*

$$\inf \Theta = \bigcap \Theta \quad \text{and} \quad \sup \Theta = \mathrm{TC}(\bigcup \Theta)\,, \quad \text{for } \Theta \subseteq \mathrm{Cong}(A)\,.$$

(b) $\mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$ *forms a complete lattice where*

$$\inf \Theta = \bigcap \Theta\,, \quad \text{for } \Theta \subseteq \mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})\,.$$

Proof. Part (a) is obvious and (b) follows by Lemma II.5.12. □

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 $A \in \mathsf{Pos}^{\Xi}$, and let $\sigma \subseteq \mathrm{TC}(\bigcup \Theta)$ be finite. Then*

$$s\ \sigma^{\mathbb{M}}\ t \quad \text{implies} \quad s = r_0\ \theta_0^{\mathbb{M}} \cdots \theta_{n-1}^{\mathbb{M}}\ r_n = t\,,$$
$$\text{for some } \theta_0, \dots, \theta_{n-1} \in \Theta \text{ and } r_0, \dots, r_n \in \mathbb{M}A\,.$$

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 $p, q : A \times A \to A$ 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 $g_k : \sigma \to A$, 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

$$\mathbb{M}g_k(u) = r_0\ \theta_0^{\mathbb{M}} \cdots \theta_{n-1}^{\mathbb{M}}\ r_n = \mathbb{M}g_{k+1}(u)\,.$$

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 $h_l : \sigma \to A \times A$, 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

$$\begin{aligned} \mathbb{M}g_k(u) = \mathbb{M}p(v_0)\ &\theta_0\ \mathbb{M}q(v_0) = \mathbb{M}p(v_1) \\ &\theta_1\ \mathbb{M}q(v_1) = \cdots \\ &\theta_{m-2}\ \mathbb{M}q(v_{m-2}) = \mathbb{M}p(v_{m-1}) \\ &\theta_{m-1}\ \mathbb{M}q(v_{m-1}) = \mathbb{M}g_{k+1}(u)\,, \end{aligned}$$

as desired. □

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 $\sqsubseteq_0, \sqsubseteq_1 \in \mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$. If $\sqsubseteq_0$ has finitary index, then $\mathrm{TC}(\sqsubseteq_0 \cup \sqsubseteq_1) \in \mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$.*

Proof. Let $\sqsubseteq := \mathrm{TC}(\sqsubseteq_0 \cup \sqsubseteq_1)$. Clearly, $\sqsubseteq$ is reflexive, transitive, and it contains $\leq$. Hence, we only need to prove the congruence property, that is, we have to show that

$$s \in \mathbb{M}\sqsubseteq \quad \text{implies} \quad \pi(\mathbb{M}p(s)) \sqsubseteq \pi(\mathbb{M}q(s)),$$

where $p, q : A \times A \to A$ are the two projections.

First, consider the case where $s \in \mathbb{M}\sigma$, for some finite $\sigma \subseteq \sqsubseteq$. Then we can use Lemma 3.6 to find a sequence

$$\mathbb{M}p(s) = r_0 \sqsubseteq_{i_0}^{\mathbb{M}} \cdots \sqsubseteq_{i_{n-1}}^{\mathbb{M}} r_n = \mathbb{M}q(s),$$

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

$$\pi(\mathbb{M}p(s)) = \pi(r_0) \sqsubseteq_{i_0} \cdots \sqsubseteq_{i_{n-1}} \pi(r_n) = \pi(\mathbb{M}q(s)).$$

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 $\gamma_0 : A/\sqsubseteq_0 \to A$ (not necessarily monotone) and let $\gamma : A \to A$ be the composition of the quotient map $A \to A/\sqsubseteq_0$ 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 $\sigma := \sqsubseteq \cap (\operatorname{rng}\gamma \times \operatorname{rng}\gamma)$ 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

$$\pi(\mathbb{M}p(s)) \sqsubseteq_0 \pi(\mathbb{M}(\gamma \circ p)(s)) \sqsubseteq \pi(\mathbb{M}(\gamma \circ q)(s)) \sqsubseteq_0 \pi(\mathbb{M}q(s)) \square$$

To obtain a stronger statement, we need to make additional assumptions on the monad $\mathbb{M}$.

Definition 3.8. A functor $\mathbb{M} : \mathsf{Pos}^{\Xi} \to \mathsf{Pos}^{\Xi}$ is *finitary* if it preserves directed colimits, that is, if

$$\mathbb{M}(\operatorname{colim} D) \cong \operatorname{colim}(\mathbb{M} \circ D)\,,$$

for every directed diagram $D : I \to \mathsf{Pos}^{\Xi}$. ⌟

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^{Ξ} 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^{Ξ} and let Θ be a set of $\mathbb{M}$-congruences on an $\mathbb{M}$-algebra $\mathfrak{A}$. Then $\mathrm{TC}(\bigcup \Theta)$ is also an $\mathbb{M}$-congruence on $\mathfrak{A}$.*

Proof. Set $\sqsubseteq\, := \mathrm{TC}(\bigcup \Theta)$. 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 $p, q : A \times A \to A$ 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

$$\mathbb{M}p(s) = r_0 \; \theta_0^{\mathbb{M}} \cdots \theta_{n-1}^{\mathbb{M}} \; r_n = \mathbb{M}q(s)\,,$$

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

$$\pi(\mathbb{M}p(s)) = \pi(r_0)\; \theta_0 \cdots \theta_{n-1}\; \pi(r_n) = \pi(\mathbb{M}q(s))\,.$$

This implies that $\pi(\mathbb{M}p(s)) \sqsubseteq \pi(\mathbb{M}q(s))$. □

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^{Ξ} where Ξ is finite. Every language $\kappa \in \mathbb{M}_\xi A \to \Omega$ that is recognised by a surjective morphism $\varphi : \mathbb{M}A \to \mathfrak{B}$ into a finitary $\mathbb{M}$-algebra has a minimal morphism.*

Proof. By Lemma 3.3, it is sufficient to prove that the set

$$\Theta := \{\, \theta \in \mathrm{Cong}_{\mathbb{M}}(\mathfrak{A}) \mid \theta \subseteq \ker \kappa \,\}$$

has a greatest element. First, note that Θ contains the congruence $\sigma := \ker \varphi$, which has finitary index. Furthermore, the set

$$\Theta_0 := \{\, \theta \in \Theta \mid \sigma \subseteq \theta \,\}\,.$$

is finite since every congruence in Θ_0 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 $\rho \in \Theta_0$ (the transitive closure of the union of all congruences in Θ_0).

We claim that ρ is also the greatest element of Θ. Hence, let $\theta \in \Theta$. Then Proposition 3.7 implies that $\hat\theta := \mathrm{TC}(\theta \cup \sigma) \in \Theta_0$. Consequently, $\theta \subseteq \hat\theta \subseteq \rho$. □

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 $\mathfrak{A}$, every set Ω, and every $\xi \in \Xi$, every Ω-language $\kappa : A_\xi \to \Omega$ has a minimal morphism.*

(2) *$\mathbb{M}$ is finitary.*

(3) *For every $\mathbb{M}$-algebra $\mathfrak{A}$, the set $\mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$ of all $\mathbb{M}$-congruences is closed, as a subset of $\mathrm{Cong}(A)$, under arbitrary joins.*

Proof. (2) ⇒ (3) follows by Proposition 3.9.

(3) ⇒ (1) Let θ be the supremum (in $\mathrm{Cong}(A)$) 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) ⇒ (2) Fix $A \in \mathsf{Set}$ and let A' be a disjoint copy of A. For every finite $C \subseteq A$ with copy $C' \subseteq A'$, we consider the function $\sigma_C : A + A' \to A + A'$ 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 $q : \mathbb{M}(A + A') \to \mathbb{M}(A + A')/{\sqsubseteq}$ 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 $\eta : \mathbb{M}(A + A') \to \mathfrak{M}$. 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 $g : M \to \mathbb{M}(A + A')/\sqsubseteq$ 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 $\rho : \operatorname{rng} \mathbb{M}\sigma_C \to \mathfrak{M}$ 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 $\iota : A \to A'$ be the functions mapping $a \in A$ to $a' \in A'$. For $a \in A$, it follows that

$$\mathbb{M}\sigma_{\{a\}}(\operatorname{sing}(a)) = \operatorname{sing}(a) = \mathbb{M}\sigma_{\{a\}}(\mathbb{M}\iota(\operatorname{sing}(a)))\,.$$

Consequently, $\operatorname{sing}(a) \equiv \mathbb{M}\iota'(\operatorname{sing}(a))$. Since $\sqsubseteq$ is a congruence ordering and $\mathbb{M}A$ is generated by $\operatorname{rng} \operatorname{sing}$, it follows that

$$s \equiv \mathbb{M}\iota(s)\,, \quad \text{for all } s \in \mathbb{M}A\,.$$

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

$$\mathbb{M}\sigma_{C_s}(s) = \mathbb{M}\sigma_{C_s}(\mathbb{M}\iota(s))\,.$$

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 $\tau : A + A' \to A$ 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 $\operatorname{rng}(\tau \circ \sigma_C \circ \iota) = C$ is finite. $\square$

Example. The word monad $\mathbb{M}X = X^+$ has arity $\aleph_0$ and is therefore finitary. Consequently, every Ω-language $\kappa : \Sigma^+ \to \Omega$ has a minimal algebra. ⌟

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 $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ 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

$$\mathrm{Syn}(K) := \Sigma^*/{\approx_K}\,,$$

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 $\Sigma^* \to \Sigma^*/{\approx_K}$ 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 $\square_\zeta$ 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 $\sigma_a : \mathbb{M}(A + \Box_\zeta) \to \mathfrak{A}$ is the unique morphism that extends the function $s_a : A + \Box_\zeta \to A$ given by

$$s_a(\Box_\zeta) := a \quad \text{and} \quad s_a(c) := c\,, \quad \text{for } c \in A\,.$$

In the case where $\mathfrak{A} = \mathbb{M}\Sigma$ is a free $\mathbb{M}$-algebra, we will also consider elements $p \in \mathbb{M}(\Sigma + \Box_\zeta)$ as contexts, by identifying them with their image under $\mathbb{M}(\mathrm{sing} + \mathrm{id})$.

(b) An *internal operation* of $\mathfrak{A}$ is a function $A_\zeta \to A_\xi$ of the form

$$a \mapsto \pi(p[a])\,, \quad \text{for some context } p \in \mathbb{M}_\xi(A + \Box_\zeta)\,.$$

In the following, we will often identify a context with the associated internal operation.

(c) The *composition* of two contexts $p \in \mathbb{M}_\xi(A+\Box_\zeta)$ and $q \in \mathbb{M}_\zeta(A+\Box_\eta)$ is the context obtained from p by replacing every occurrence of $\Box_\zeta$ by a copy of q. Formally,

$$pq := \hat{p}[q] \in \mathbb{M}_\xi(A + \Box_\eta)\,,$$

where $\hat{p} := \mathbb{M}(\mathrm{sing} + \mathrm{id})(p)$ and the expression $\hat{p}[q]$ is evaluated in the $\mathbb{M}$-algebra $\mathbb{M}(A + \Box_\eta)$.

(d) A *derivative* of $\kappa : A_\xi \to \Omega$ is a function of the form

$$p^{-1}[\kappa] := \kappa \circ p : A_\zeta \to \Omega\,, \quad \text{where } p \in \mathbb{M}_\xi(A + \Box_\zeta) \text{ is a context.} \quad \lrcorner$$

We can now lift the language functor $\mathbb{L}^\circ$ from $\mathcal{D}$ to $\mathsf{Alg}(\mathbb{M})$ 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 $(\mathbb{L}^\circ A_\xi)_{\xi\in\Xi} \in \mathsf{Set}^\Xi$ 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 $\mathbb{L}^\circ p : \mathbb{L}^\circ A_\xi \to \mathbb{L}^\circ A_\zeta$, for contexts $p \in \mathbb{M}_\xi(A + \Box_\zeta)$, $\xi, \zeta \in \Xi$.

(b) For a morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ 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 $\varphi : \mathbb{L}\mathfrak{A} \to \mathbb{L}\mathfrak{B}$ 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 $\varnothing$ and A; and (iv) derivatives $a^{-1}K$ and Ka^{-1}, for $a \in A$. ⌟

Lemma 4.3. *Let $\varphi : \mathfrak{A} \to \mathfrak{B}$ be a morphism of $\mathbb{M}$-algebras and $p \in \mathbb{M}_\xi(A + \square_\zeta)$ a context.*

(a) *The image $q := \mathbb{M}(\varphi + \mathrm{id})(p) \in \mathbb{M}_\xi(B + \square_\zeta)$ 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 $\mathbb{M}(\varphi + \mathrm{id})$ on contexts is also an E-morphism.*

Proof. (a) Note that

$$\begin{aligned}
(\hat{q} \circ \varphi)(s) &= \hat{q}(\varphi(s)) \\
&= \pi\big(\mathbb{M}(\varphi + \mathrm{id})(p)[\varphi(s)]\big) \\
&= \pi\big(\mathbb{M}\varphi(p[s])\big) \\
&= \varphi\big(\mathrm{flat}(p[s])\big) \\
&= (\varphi \circ \hat{p})(s) \,.
\end{aligned}$$

(b) follows by Lemma II.1.13 (a) and our assumption on $\mathbb{M}$. □

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 $\omega : \Omega^n \to \Omega$, every context p, and all $\kappa_0, \dots, \kappa_{n-1} \in \mathbb{L}A$.

Proof. By definition,

$$\begin{aligned} p^{-1}[\omega[\bar{\kappa}]] &= \omega \circ \langle \kappa_0, \dots, \kappa_{n-1} \rangle \circ p \\ &= \omega \circ \langle \kappa_0 \circ p, \dots, \kappa_{n-1} \circ p \rangle \\ &= \omega[p^{-1}[\kappa_0], \dots, p^{-1}[\kappa_{n-1}]]. \end{aligned}$$

□

Definition 4.5. Let $\mathfrak{A}$ be an $\mathbb{M}$-algebra and $\kappa : A_\xi \to \Omega$.

(a) The *syntactic morphism* of κ is

$$\mathrm{syn}_\kappa := \ker \rho,$$

where $\rho = (\rho_\zeta)_{\zeta \in \Xi}$ is the morphism with components

$$\rho_\zeta := \langle \kappa \circ p \rangle_{p \in \mathbb{M}_\zeta(A + \square)} : A_\zeta \to \prod_p \Omega.$$

(b) We call $\ker \mathrm{syn}_\kappa$ the *syntactic congruence* of κ.

(c) The *syntactic residue* is the unique function $\mathrm{res}_\kappa : \mathrm{Syn}_\xi(\kappa) \to \Omega$ satisfying $\kappa = \mathrm{res}_\kappa \circ \mathrm{syn}_{\kappa,\xi}$.

(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 $\mathrm{Syn}(\kappa)$. ⌟

Remark. (a) Using the general definition of the kernel of a morphism, we have $\ker \mathrm{syn}_\kappa = \mathrm{syn}_\kappa$. 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 $\mathrm{Syn}(\kappa)$ exists, the syntactic morphism $\mathrm{syn}_\kappa : \mathfrak{A} \to \mathrm{Syn}(\kappa)$ is a morphism of $\mathbb{M}$-algebras.

(c) Uniqueness of res_κ follows from the fact that syn_κ is an epimorphism. ⌟

In the category Pos the definitions simplify as follows.

Lemma 4.6. *Suppose that $\mathcal{D} = \mathsf{Pos}$ and let $\preceq_\kappa$ be the syntactic congruence of $\kappa : A_\xi \to \Omega$. Then*

$$\begin{aligned} a \preceq_\kappa b \quad &\text{iff} \quad \kappa(p[a]) \leq \kappa(p[b])\,, \quad \text{for all } p \in \mathbb{M}_\xi(A + \Box_\zeta)\,, \\ &\text{iff} \quad \mu(a) \leq \mu(b)\,, \quad \text{for all } \mu \in \langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}}\,. \end{aligned}$$

Proof. The first equivalence follows immediately by unravelling the definitions. Since $p^{-1}[\kappa] \in \langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}}$, it therefore remains to prove that

$$\kappa(p[a]) \leq \kappa(p[b])\,, \quad \text{for all } p \in \mathbb{M}_\xi(A + \Box_\zeta)\,,$$

implies that

$$\mu(a) \leq \mu(b)\,, \quad \text{for all } \mu \in \langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}}\,.$$

Hence, fix $\mu \in \langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}}$. 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 $\omega : \Omega \to \Omega$ and some context p. Since ω is monotone,

$$\kappa(p[a]) \leq \kappa(p[b]) \quad \text{implies} \quad \omega[\kappa](p[a]) \leq \omega[\kappa](p[b])\,. \qquad \Box$$

Remark. In the unweighted case for $\mathcal{D} = \mathsf{Pos}$, we obtain for $K \subseteq A$

$$a \preceq_K b \quad \text{iff} \quad p[a] \in K \Rightarrow p[b] \in K\,, \quad \text{for all contexts } p\,.$$

For $\mathcal{D} = \mathsf{Set}$, we obtain the familiar definition

$$a \approx_K b \quad \text{iff} \quad p[a] \in K \Leftrightarrow p[b] \in K\,, \quad \text{for all contexts } p\,.$$ ⌟

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 $\preceq_K$ has the following 5 classes.

$$\begin{tikzcd}[column sep=small]
 & \langle\rangle \arrow[dl, dash] \arrow[dr, dash] & \\
a^+ \arrow[dr, dash] & & b^+ \arrow[dl, dash] \\
 & a^+b^+ \arrow[d, dash] & \\
 & (a+b)^*ba(b+a)^* &
\end{tikzcd}$$

(b) Let us consider the $\mathbb{N}$-language $\kappa : \Sigma^* \to \mathbb{N}$ 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 $\leq_\kappa$ is a linear order with

$$w \leq_\kappa w' \quad \text{iff} \quad |w|_a \leq |w'|_a \,. \qquad \lrcorner$$

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 $\varphi : \Lambda X \to \mathfrak{A}$ 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 $\varphi = i \circ \mathrm{syn}_k$, for some embedding $i : \mathrm{Syn}(K) \to \mathfrak{A}$. $\lrcorner$

Let us collect a few basic properties of the syntactic congruence. We start by showing that our definitions make sense.

Lemma 4.7. *Let $\kappa : A_\xi \to \Omega$.*

(a) $\mathrm{syn}_\kappa : A \to \mathrm{Syn}(\kappa)$ *recognises* κ.

(b) $\mathrm{res}_\kappa : \mathrm{Syn}_\xi(\kappa) \to \Omega$ *exists.*

(c) $\ker \mathrm{syn}_\kappa \leq \ker \kappa$.

(d) $\ker \mathrm{syn}_\kappa = \inf_p \ker(\kappa \circ p)$

Proof. (d) Let $\rho_\xi := \langle \kappa \circ p \rangle_p : A_\xi \to \prod_p \Omega$ be the morphisms used to define syn_κ. Then

$$\ker \mathrm{syn}_k = \ker \rho = \ker \langle \kappa \circ p \rangle_p = \inf_p \ker(\kappa \circ p) .$$

This proves (d).

(c) We have $\kappa \circ p_\circ = \kappa$, for the trivial context $p_\circ := \mathrm{sing}(\square)$. Hence, (d) implies (c).

(a), (b) By (c) we can find some function μ with $\ker \kappa = \mu \circ \ker \mathrm{syn}_\kappa$. This implies that

$$\kappa = \mathrm{im}\, \kappa \circ \ker \kappa = \mathrm{im}\, \kappa \circ \mu \circ \mathrm{syn}_\kappa .$$

Hence, syn_κ recognises κ and $\mathrm{res}_k := \mathrm{im}\, \kappa \circ \mu$ exists. ☐

The following lemma contains the key property of the syntactic congruence.

Lemma 4.8. *Let* $\kappa : A_\xi \to \Omega$ *be a language such that* syn_κ *is an* $\mathbb{M}$*-congruence, and let* $\varphi : \mathfrak{A} \to \mathfrak{B}$ *be an* E*-morphism of* $\mathbb{M}$*-algebras. The following statements are equivalent.*

(1) φ *recognises* κ.

(2) $\ker \varphi \leq \ker \mathrm{syn}_\kappa$

(3) $\ker \varphi \leq \ker \kappa$

Proof. (1) ⇔ (3) has already been proved in Lemma 3.1 (a).

(1) ⇒ (2) If $\kappa = \mu \circ \varphi$, for some $\mu : B_\xi \to \Omega$, then

$$\varphi = \ker \varphi \leq \ker(\mu \circ \varphi) = \ker \kappa = \ker \mathrm{syn}_\kappa .$$

(2) ⇒ (1) Since $\varphi = \ker\varphi \leq \ker \mathrm{syn}_\kappa = \mathrm{syn}_\kappa$, there exists a function $\rho : B \to \mathrm{Syn}(\kappa)$ such that $\mathrm{syn}_\kappa = \rho \circ \varphi$. Hence,

$$\kappa = \mathrm{res}_\kappa \circ \mathrm{syn}_\kappa = \mathrm{res}_\kappa \circ \rho \circ \varphi\,.$$

which implies that φ recognises κ. □

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 $\kappa : A_\xi \to \Omega$ be a language that has a syntactic algebra. For every E-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ recognising κ, there exists a unique morphism $\rho : \mathfrak{B} \to \mathrm{Syn}(\kappa)$ such that $\mathrm{syn}_\kappa = \rho \circ \varphi$.*

Proof. Suppose that φ recognises κ. By Lemma 4.8 we have $\ker\varphi \leq \ker \mathrm{syn}_\kappa$. Therefore, there exists some function $\rho : B \to \mathrm{Syn}(\kappa)$ with $\mathrm{syn}_\kappa = \rho\circ\varphi$. It is unique as φ is an epimorphism. Furthermore, ρ is a morphism of $\mathbb{M}$-algebras by Lemma I.5.6. □

Corollary 4.10. *Let $\mathfrak{A}$ be an $\mathbb{M}$-algebra and $\kappa : A_\xi \to \Omega$ a language with a syntactic algebra. Then $\mathrm{Syn}(\kappa)$ is equal to the minimal algebra of κ and $\mathrm{syn}_\kappa : \mathfrak{A} \to \mathrm{Syn}(\kappa)$ 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 $\kappa : A_\xi \to \Omega$ 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, $\mathrm{Syn}(\kappa)$, syn_κ, and res_κ. ⌟

We are particularly interested in languages whose syntactic algebra is finite-dimensional.

Lemma 4.12. *If $\mathrm{Syn}(\kappa)$ is sort-wise noetherian, it is sort-wise strongly finite-dimensional.*

Proof. Let $\rho_\xi : \mathbb{M}_\xi A \to \Omega^X$ be the morphism used to define $\mathrm{Syn}(\kappa)$. Then $\operatorname{im}\rho_\xi : \mathrm{Syn}_\xi(\kappa) \to \Omega^X$ is an M-morphism. Since $\mathrm{Syn}_\xi(\kappa)$ is noetherian, it follows by Lemma 1.8 that it is strongly finite-dimensional. □

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, $\kappa : A_\xi \to \Omega$ a function, and $p \in \mathbb{M}_\xi(A + \Box_\zeta)$ a context.*

(a) $\ker \mathrm{syn}_{\kappa,\zeta} \leq \ker\left(\mathrm{syn}_{\kappa,\xi} \circ p\right)$.

(b) $\ker \mathrm{syn}_\kappa \leq \ker \mathrm{syn}_{p^{-1}[\kappa]}$.

Proof. Let $\rho_\eta := \langle \kappa \circ q\rangle_q : A_\eta \to \prod_q \Omega$ be the morphisms used to define syn_κ, and let $u_r : \prod_q \Omega \to \Omega$ be the projection to the r-th component.

(a) Note that $\rho = \mu \circ \mathrm{syn}_\kappa$, for some $\mu \in M$. Consequently, we have

$$\ker \mathrm{syn}_\kappa = \ker \rho \quad \text{and} \quad \ker\left(\mathrm{syn}_\kappa \circ p\right) = \ker\left(\rho \circ p\right).$$

Setting $g := \langle u_{qp}\rangle_q$, 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

$$\ker \mathrm{syn}_{\kappa,\zeta} = \ker \rho_\zeta \leq \ker\left(\rho_\xi \circ p\right) = \ker\left(\mathrm{syn}_{\kappa,\xi} \circ p\right).$$

(b) Note that the morphism $\mathrm{syn}_{p^{-1}[\kappa]}$ is defined via

$$\sigma_\eta := \langle \kappa \circ p \circ q\rangle_q : A_\eta \to \prod_q \Omega .$$

Setting $g := \langle u_{pq}\rangle_q$, it follows that $\sigma_\eta = g \circ \rho_\eta$. Consequently, we have

$$\ker \mathrm{syn}_\kappa = \ker \rho \leq \ker\left(g \circ \rho\right) = \ker \sigma = \ker \mathrm{syn}_{p^{-1}[\kappa]} . \qquad \Box$$

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 $\kappa : A_\xi \to \Omega$ a language with a sort-wise noetherian syntactic algebra. Then*

$$\mathbb{L}\mathrm{Syn}(\kappa) = \langle\!\langle \mathrm{res}_\kappa \rangle\!\rangle_{\mathbb{L}}\,.$$

Proof. Let $\mu \in \mathbb{L}_\zeta \mathrm{Syn}(\kappa)$ and set $\lambda := \mu \circ \mathrm{syn}_\kappa$. As $\mathrm{Syn}_\zeta(\kappa)$ is noetherian, there exists a finite set $H \subseteq \mathbb{M}_\zeta(A + \Box_\xi)$ of contexts such that

$$\ker\,(q_H \circ \mathrm{im}\,\rho)|_{\mathrm{Syn}_\zeta(\kappa)} = \ker(\mathrm{im}\,\rho)|_{\mathrm{Syn}_\zeta(\kappa)}\,,$$

where $\rho_\xi := \langle \kappa \circ p \rangle_p : A_\xi \to \prod_p \Omega$ are the morphisms used to define syn_κ and $q_H : \prod_p \Omega \to \prod_{p \in H} \Omega$ is the projection to the components in H. It follows that

$$\ker\,(q_H \circ \rho)|_{A_\zeta} = \ker \rho|_{A_\zeta} = \mathrm{syn}_\kappa|_{A_\zeta} \leq (\mu \circ \mathrm{syn}_\kappa)|_{A_\zeta} = \ker \lambda\,.$$

As Ω is injective, there therefore exists a function $\omega : \Omega^H \to \Omega$ 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,

$$\begin{aligned}
\mu \circ \mathrm{syn}_\kappa = \lambda &= \omega \circ \big\langle \kappa \circ p \big\rangle_{p \in H} \\
&= \omega \circ \big\langle \mathrm{res}_\kappa \circ \mathrm{syn}_\kappa \circ p \big\rangle_{p \in H} \\
&= \omega \circ \big\langle \mathrm{res}_\kappa \circ \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})(p) \circ \mathrm{syn}_\kappa \big\rangle_{p \in H} \\
&= \omega \circ \big\langle \mathrm{res}_\kappa \circ p \big\rangle_{p \in \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})[H]} \circ \mathrm{syn}_\kappa\,,
\end{aligned}$$

where the fourth step follows by Lemma 4.3. As syn_κ is an epimorphism, it follows that

$$\mu = \omega\big[\langle p^{-1}[\mathrm{res}_\kappa] \rangle_{p \in \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})[H]}\big] \in \langle\!\langle \mathrm{res}_\kappa \rangle\!\rangle_{\mathbb{L}}\,. \qquad \square$$

We obtain the following preliminary characterisation of which languages are recognised by $\mathrm{Syn}(\kappa)$. 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 $\kappa : A_\xi \to \Phi$ and $\lambda : A_\zeta \to \Omega$ be languages such that κ has a syntactic algebra. The following statements are equivalent.*

(1) $\ker \mathrm{syn}_\kappa \leq \ker \mathrm{syn}_\lambda$

(2) $\mathrm{syn}_\kappa : \mathfrak{A} \to \mathrm{Syn}(\kappa)$ *recognises* λ.

(3) *Every morphism of $\mathbb{M}$-algebras recognising κ also recognises λ.*

If $\mathrm{Syn}(\kappa)$ is sort-wise noetherian, the following statement is also equivalent to those ones.

(4) $\lambda = \omega\big[p_0^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]\big]$, *for some operation $\omega : \Phi^n \to \Omega$ and finitely many contexts $p_0, \dots, p_{n-1}$.*

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 $\varphi : \mathfrak{A} \to \mathfrak{B}$ recognises κ. By Lemma 4.8, it follows that $\ker \varphi \leq \ker \mathrm{syn}_\kappa \leq \ker \mathrm{syn}_\lambda$, which, by the same lemma, implies that φ recognises λ.

(2) $\Rightarrow$ (5) Suppose that $\lambda = \mu \circ \mathrm{syn}_\kappa$ for some $\mu : \mathrm{Syn}_\zeta(\kappa) \to \Omega$. By Proposition 4.14, we have $\mu \in \langle\!\langle \mathrm{res}_\kappa \rangle\!\rangle_{\mathbb{L}}$. Hence, there is some operation ω of $\mathbb{L}\mathrm{Syn}(\kappa)$ with $\mu = \omega \circ \mathrm{res}_\kappa$. It follows that

$$\lambda = \mu \circ \mathrm{syn}_\kappa = \omega \circ \mathrm{res}_\kappa \circ \mathrm{syn}_\kappa = \omega \circ \kappa\,.$$

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

$$\lambda = \omega\big[p_0^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]\big],$$

for some Ω-operation ω and contexts $p_0, \dots, p_{n-1}$.

(4) ⇒ (2) We have

$$\begin{aligned}\lambda &= \omega[p_0^{-1}[\kappa], \dots, p_{n-1}[\kappa]] \\ &= \omega \circ \langle \kappa \circ p_0, \dots, \kappa \circ p_{n-1} \rangle \\ &= \omega \circ \langle \mathrm{res}_\kappa \circ \mathrm{syn}_\kappa \circ p_0, \dots, \mathrm{res}_\kappa \circ \mathrm{syn}_\kappa \circ p_{n-1} \rangle \\ &= \omega \circ \langle \mathrm{res}_\kappa \circ \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})(p_0) \circ \mathrm{syn}_\kappa, \dots, \\ &\qquad\qquad \mathrm{res}_\kappa \circ \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})(p_{n-1}) \circ \mathrm{syn}_\kappa \rangle \\ &= \omega \circ \langle \mathrm{res}_\kappa \circ \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})(p_0), \dots, \\ &\qquad\qquad \mathrm{res}_\kappa \circ \mathbb{M}(\mathrm{syn}_\kappa + \mathrm{id})(p_{n-1}) \rangle \\ &\qquad \circ \mathrm{syn}_\kappa\,,\end{aligned}$$

where the fourth step follows by Lemma 4.3. Hence, $\lambda = \mu \circ \mathrm{syn}_\kappa$, for some μ, and syn_κ recognises λ. □

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 $\omega : \{0,1\}^n \to \{0,1\}$ 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. ⌟

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 $\kappa : A_\xi \to \Phi$ be a Φ-language with a sort-wise noetherian syntactic algebra. An Ω-language $\lambda : B_\zeta \to \Omega$ is recognised by* $\mathrm{Syn}(\kappa)$ *if, and only if, it is of the form*

$$\lambda = \omega\big[\varphi^{-1}[p_0^{-1}[\kappa]], \dots, \varphi^{-1}[p_{n-1}^{-1}[\kappa]]\big]\,,$$

for some a morphism $\varphi : \mathfrak{B} \to \mathfrak{A}$, contexts $p_0, \dots, p_{n-1} \in \mathbb{M}(A + \square)$, and an operation $\omega : \Phi^n \to \Omega$.

Proof. Note that we can write

$$\omega\big[\varphi^{-1}[p_0^{-1}[\kappa]], \dots, \varphi^{-1}[p_{n-1}^{-1}[\kappa]]\big] = \omega \circ \langle p_i^{-1}[\kappa]\rangle_{i<n} \circ \varphi\,.$$

($\Leftarrow$) By Proposition 4.15, the morphism $\mathrm{syn}_\kappa : \mathfrak{A} \to \mathrm{Syn}(\kappa)$ recognises the language

$$\lambda' := \omega \circ \langle p_i^{-1}[\kappa]\rangle_{i<n}\,.$$

Consequently, $\mathrm{syn}_\kappa \circ \varphi : \mathbb{M}\Gamma \to \mathrm{Syn}(\kappa)$ recognises $\lambda' \circ \varphi = \lambda$.

($\Rightarrow$) Fix a morphism $\psi : \mathfrak{B} \to \mathrm{Syn}(\kappa)$ and a language $\mu : \mathrm{Syn}_\zeta(\kappa) \to \Omega$ such that $\lambda = \psi^{-1}[\mu]$. By Corollary II.6.4, there exists a morphism $\varphi : \mathfrak{B} \to \mathfrak{A}$ such that $\mathrm{syn}_\kappa \circ \varphi = \psi$. The language $\lambda' := \mathrm{syn}_\kappa^{-1}[\mu]$ 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 $\omega : \Phi^n \to \Phi$ such that

$$\lambda' = \omega[p_0^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]]\,.$$

It follows that

$$\lambda = \mu \circ \psi = \mu \circ \mathrm{syn}_\kappa \circ \varphi = \omega \circ \langle p_i^{-1}[\kappa]\rangle_{i<n} \circ \varphi\,. \qquad \square$$

Remark. In the unweighted case, the condition reads

$$L = \varphi^{-1}\Big[\bigcup_{i<m} \bigcap_{k<n_i} p_{ik}^{-1}[K]\Big]\,.$$

⌟

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 $\varepsilon : A \to C$ 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 $\varepsilon \leq \varepsilon \circ p$. □

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 $\varepsilon : A \to C$ is a morphism of $\mathbb{M}$-algebras if, and only if,*

$$\varepsilon \leq \varepsilon \circ p\,, \quad \textit{for all contexts } p \in \mathbb{M}(A + \square)\,.$$

Proof. $(\Rightarrow)$ was already proved in Lemma 4.17.

$(\Leftarrow)$ It is sufficient to prove that $\sqsubseteq\; := \ker \varepsilon$ is an $\mathbb{M}$-congruence. Hence, consider a term $u \in \mathbb{M}{\sqsubseteq}$. We have to show that

$$\pi(\mathbb{M}q(u)) \sqsubseteq \pi(\mathbb{M}q'(u))\,,$$

where $q, q' : A \times A \to A$ are the two projections. Since $\mathbb{M}$ is finitary, we have $u \in \mathbb{M}\sigma$ for some finite subset $\sigma \subseteq {\sqsubseteq}$. Let $\langle a_0, b_0\rangle, \dots, \langle a_{m-1}, b_{m-1}\rangle$ be an enumeration of σ. For $i < m$, set

$$\theta_i := \leq \cup \{\langle a_i, b_i\rangle\}\,.$$

Then $\sigma \subseteq \mathrm{TC}(\theta_0 \cup \dots \cup \theta_{m-1})$ and we can use Lemma 3.6 to find $r_0, \dots, r_n \in \mathbb{M}A$ and indices $k_0, \dots, k_{n-1} < m$ such that

$$\mathbb{M}q(u) = r_0\ \theta^{\mathbb{M}}_{k_0}\ r_1\ \theta^{\mathbb{M}}_{k_1} \cdots \theta^{\mathbb{M}}_{k_{m-1}}\ r_m = \mathbb{M}q'(u)\,.$$

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,

$$a_{k_i} \sqsubseteq b_{k_i} \quad \text{implies} \quad \pi(r_i) = p_i(a_{k_i}) \sqsubseteq p_i(b_{k_i}) = \pi(r_{i+1})\,.$$

Consequently, $\pi(\mathbb{M}q(u)) \sqsubseteq \dots \sqsubseteq \pi(\mathbb{M}q'(u))$. □

Theorem 4.19. *Let $\mathbb{M}$ be a finitary monad on Pos^{Ξ} and $\mathfrak{A}$ an $\mathbb{M}$-algebra. Every language $\kappa : A_{\xi} \to \Omega$ has a syntactic algebra.*

Proof. By Lemma 4.13 (a), we have

$$\mathrm{syn}_{\kappa} \leq \mathrm{syn}_{\kappa} \circ p \,, \qquad \text{for all contexts } p \,.$$

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 $\mathbb{M}^{\circ}$ 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

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

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 $\mathbb{M}^\circ$-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 $\mathbb{M}^\circ$-reduct are isomorphic.*

(b) *Let $\varphi : \mathfrak{A}^\circ \to \mathfrak{B}^\circ$ be a morphism of $\mathbb{M}^\circ$-algebras and assume that $\mathfrak{A}^\circ$ and $\mathfrak{B}^\circ$ are the $\mathbb{M}^\circ$-reducts of two $\mathbb{M}$-algebras $\mathfrak{A}, \mathfrak{B} \in \mathcal{C}$. Then φ is also a morphism $\mathfrak{A} \to \mathfrak{B}$ 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 $\mathbb{M}^\circ$-reduct $\mathfrak{A}^\circ$ 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 $\mathbb{M}^\circ$-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 $d : A \to \Delta$ 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 $p, q : A \times A \to A$ 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 $i := \langle \mathrm{id}, \varphi\rangle : A \to G$ 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 $p : A \times B \to A$ and $q : A \times B \to B$ 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,

$$\begin{aligned}\pi(\mathbb{M}\varphi(s)) &= \pi(\mathbb{M}q(t)) \\ &= q(\pi(t)) \\ &= q(\pi(t^\circ)) \\ &= \pi(\mathbb{M}^\circ q(t^\circ)) \\ &= \pi\big(\mathbb{M}^\circ(\varphi \circ p)(t^\circ)\big) \\ &= \varphi(p(\pi(t^\circ))) \\ &= \varphi(p(\pi(t))) \\ &= \varphi(\pi(\mathbb{M}p(t))) = \varphi(\pi(s))\,.\end{aligned}$$

(c) Clearly, every $\mathbb{M}$-algebra morphism is also an $\mathbb{M}^\circ$-algebra morphism. The converse follows from (b). □

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 $\mathsf{Alg}(\mathbb{M})$) 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 $\iota_X : \langle\!\langle X \rangle\!\rangle_{\mathfrak{A}} \to A$ be the inclusion map. Since $\mathcal{D}$ has canonical subobject diagrams, it follows that A is the limit of D in $\mathcal{D}^{\Xi}$ and that $(\iota_X)_X$ is the corresponding limiting cone.

It remains to show that it is also limiting in $\mathsf{Alg}(\mathbb{M})$. Hence, consider a cocone $(\mu_X)_X$ from D to some $\mathbb{M}$-algebra $\mathfrak{B}$. We obtain a unique morphism $\varphi : A \to B$ satisfying

$$\varphi \circ \iota_X = \mu_X\,, \quad \text{for all } X\,.$$

It remains to show that $\varphi : \mathfrak{A} \to \mathfrak{B}$ is a $\mathbb{M}$-morphism. By assumption, there exists a morphism of monads $\rho : \mathbb{M}^\circ \Rightarrow \mathbb{M}$ such that $\mathbb{M}^\circ$ is finitary and ρ is dense over $\mathcal{C}$. By Lemma 4.22, it is sufficient to prove that φ is a $\mathbb{M}^\circ$-morphism $\mathfrak{A}^\circ \to \mathfrak{B}^\circ$ between the respective ρ-reducts. Let $s \in \mathbb{M}^\circ A$. As $\mathbb{M}^\circ$ 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 $\varphi \restriction \langle\!\langle X \rangle\!\rangle_{\mathfrak{A}} = \varphi \circ \iota_X = \mu_X$. ☐

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 $\kappa : A_\xi \to \Omega$ is recognised by some E-morphism $\varphi : \mathfrak{A} \to \mathfrak{C}$ with $\mathfrak{C} \in \mathcal{C}$, then $\ker \mathrm{syn}_\kappa$ 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 $\mathrm{rng}\,\varphi$. By Theorem 4.19, syn_μ is a morphism of $\mathbb{M}^\circ$-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

$$\mathrm{syn}_\kappa = \mathrm{syn}_\mu \circ \varphi\,.$$

Let

$$\rho := \langle \kappa \circ p \rangle_p : A \to \prod_p \Omega\,,$$
$$\sigma := \langle \mu \circ p \rangle_p : C \to \prod_p \Omega$$

be the morphisms such that $\mathrm{syn}_\kappa := \ker \rho$ and $\mathrm{syn}_\mu := \ker \sigma$. 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

$$\mathrm{syn}_\kappa = \ker \rho = \ker(\sigma \circ \varphi) = \ker \sigma \circ \varphi = \mathrm{syn}_\mu \circ \varphi\,. \qquad \square$$

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 $\pi : \mathbb{O}\Omega \to \Omega$, 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 $\mathrm{Lin}(X)$ 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 $\mathrm{DL}(X)$ 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 $\mathcal{D}^{\Xi}$ by $\mathbb{O}^{\dagger} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$. 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 $\delta : \mathbb{M}\mathbb{O}^{\dagger} \to \mathbb{O}^{\dagger}\mathbb{M}$. 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 $\kappa : \mathbb{M}_{\xi}\Sigma \to \Omega$ to a unique morphism of $\mathbb{O}$-algebras $\hat{\kappa} : \mathbb{O}\mathbb{M}_{\xi}\Sigma \to \Omega$ with $\hat{\kappa} \circ \varepsilon = \kappa$. Furthermore, if $\varphi : \mathbb{O}^{\dagger}\mathbb{M}\Sigma \to \mathfrak{A}$ 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 $\mu : A_{\xi} \to \Omega$ 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 $\kappa : \mathbb{O}\mathbb{M}_{\xi}\Sigma \to \Omega$ 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 $\mu : A_{\xi} \to \Omega$ is an $\mathbb{O}$-algebra morphism.

Let us take a look at how the syntactic $\mathbb{O}^{\dagger}\mathbb{M}$-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 $\kappa : \mathbb{M}_{\xi}\Sigma \to \Omega$ be an Ω-language and $\hat{\kappa} : \mathbb{O}\mathbb{M}_{\xi}\Sigma \to \Omega$ its lifting to $\mathbb{O}^{\dagger}\mathbb{M}\Sigma$. There exists an embedding $e : \mathrm{Syn}(\kappa) \to \mathrm{Syn}(\hat{\kappa})$ of $\mathbb{M}$-algebras such that $\mathrm{Syn}(\hat{\kappa})$ is generated (as an $\mathbb{O}^{\dagger}$-algebra) by $\operatorname{rng} e$ and e satisfies*

$$\hat{\mu} \circ e = \mu \quad \textit{and} \quad e \circ \mathrm{syn}_{\kappa} = \mathrm{syn}_{\hat{\kappa}} \circ \varepsilon\,,$$

where $\mu : \mathrm{Syn}_{\xi}(\kappa) \to \Omega$ and $\hat{\mu} : \mathrm{Syn}_{\xi}(\hat{\kappa}) \to \Omega$ are the functions such that $\kappa = \mu \circ \mathrm{syn}_{\kappa}$ and $\hat{\kappa} = \hat{\mu} \circ \mathrm{syn}_{\hat{\kappa}}$.

Proof. Let $\pi : \mathbb{O}\Omega \to \Omega$ be the $\mathbb{O}$-algebra product of Ω. Since

$$\pi \circ \mathbb{O}\mu \circ \mathbb{O}^{\dagger}\mathrm{syn}_{\kappa} = \pi \circ \mathbb{O}\kappa = \pi \circ \mathbb{O}(\hat{\kappa} \circ \varepsilon) = \hat{\kappa} \circ \pi \circ \mathbb{O}\varepsilon = \hat{\kappa}\,,$$

the morphism $\mathbb{O}^\dagger \mathrm{syn}_\kappa$ recognises $\hat{\kappa}$. Consequently, there exists an $\mathbb{O}^\dagger\mathbb{M}$-algebra morphism $\varphi : \mathbb{O}^\dagger \mathrm{Syn}(\kappa) \to \mathrm{Syn}(\hat{\kappa})$ with

$$\varphi \circ \mathbb{O}^\dagger \mathrm{syn}_\kappa = \mathrm{syn}_{\hat{\kappa}}\,.$$

Setting $e := \varphi \circ \varepsilon$ it follows that

$$e \circ \mathrm{syn}_\kappa = \varphi \circ \varepsilon \circ \mathrm{syn}_\kappa = \varphi \circ \mathbb{O}\mathrm{syn}_\kappa \circ \varepsilon = \mathrm{syn}_{\hat{\kappa}} \circ \varepsilon\,.$$

As syn_κ is an epimorphism, this implies that

$$\mathrm{rng}\, e = \mathrm{rng}\,(\varphi \circ \varepsilon) = \mathrm{rng}\,(\varphi \circ \varepsilon \circ \mathrm{syn}_\kappa) = \mathrm{rng}\,(\mathrm{syn}_{\hat{\kappa}} \circ \varepsilon)\,.$$

Hence, e factorises through $\mathfrak{S}$ and we obtain a morphism $e_0 : \mathrm{Syn}(\kappa) \to \mathfrak{S}$ with

$$i \circ e_0 = e \quad \text{and} \quad e_0 \circ \mathrm{syn}_\kappa = \eta\,.$$

Furthermore, $\mathrm{Syn}(\hat{\kappa})$ is generated by $\mathrm{rng}\, e = \mathrm{rng}\,(\mathrm{syn}_{\hat{\kappa}} \circ \varepsilon)$ since the morphism $\mathrm{syn}_{\hat{\kappa}}$ is an epimorphism and its domain $\mathbb{O}^\dagger\mathbb{M}\Sigma$ is generated by $\mathrm{rng}\,\varepsilon$.

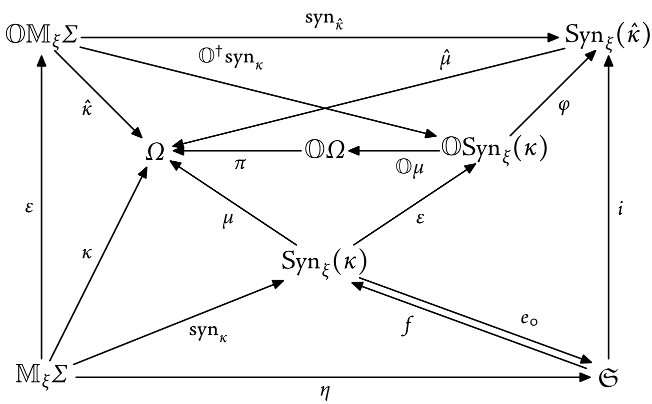

To prove that e is an embedding of $\mathbb{M}$-algebras, we show that e_0 is an isomorphism by finding an inverse. Note that $\mathbb{O}^\dagger\mathbb{M}\Sigma$ is a $\mathbb{M}$-algebra with product $\mathbb{O}^\dagger\mathrm{flat} \circ \delta$. Since

$$(\mathbb{O}^\dagger\mathrm{flat} \circ \delta) \circ \mathbb{M}\varepsilon = \mathbb{O}^\dagger\mathrm{flat} \circ \varepsilon = \varepsilon \circ \mathrm{flat}\,,$$

it therefore follows that $\varepsilon : \mathbb{M}\Sigma \to \mathbb{O}\mathbb{M}\Sigma$ is a morphism of $\mathbb{M}$-algebras. Hence, so is $\eta := \mathrm{syn}_{\hat\kappa} \circ \varepsilon$. Let $\mathfrak{S} \subseteq \mathrm{Syn}(\hat\kappa)$ be the $\mathbb{M}$-subalgebra of $\mathrm{Syn}(\hat\kappa)$ induces by $\mathrm{rng}\,\eta$ and let $i : \mathfrak{S} \to \mathrm{Syn}(\hat\kappa)$ be the corresponding inclusion map. Since

$$\hat\mu \circ \mathrm{syn}_{\hat\kappa} \circ \varepsilon = \hat\kappa \circ \varepsilon = \kappa\,,$$

the morphism $\eta : \mathbb{M}\Sigma \to \mathfrak{S}$ recognises κ. Therefore, there exists an $\mathbb{M}$-algebra morphism $f : \mathfrak{S} \to \mathrm{Syn}(\kappa)$ with

$$f \circ \eta = \mathrm{syn}_\kappa\,.$$

To see that f is the desired inverse of e_0, note that

$$f \circ e_0 \circ \mathrm{syn}_\kappa = f \circ \eta = \mathrm{syn}_\kappa \quad \text{and} \quad e_0 \circ f \circ \eta = e_0 \circ \mathrm{syn}_\kappa = \eta\,.$$

By surjectivity of syn_κ and η, we obtain $f \circ e_0 = \mathrm{id}$ and $e_0 \circ f = \mathrm{id}$. Consequently, $e_0 : \mathrm{Syn}(\kappa) \cong \mathfrak{S}$ is an isomorphism and $e = i \circ e_0 : \mathrm{Syn}(\kappa) \to \mathrm{Syn}(\hat\kappa)$ an embedding (of $\mathbb{M}$-algebras).

Finally, note that

$$\hat\mu \circ e \circ \mathrm{syn}_\kappa = \hat\mu \circ \mathrm{syn}_{\hat\kappa} \circ \varepsilon = \hat\kappa \circ \varepsilon = \kappa = \mu \circ \mathrm{syn}_\kappa\,.$$

By surjectivity of syn_κ, this implies that $\hat\mu \circ e = \mu$. □

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 $\Sigma \in \mathsf{Alph}$.*
- *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 $\varphi : \mathfrak{T} \to \mathfrak{A}$ recognises a language κ with $\mathfrak{A} \in \mathcal{R}$, then so does $\ker \varphi : \mathfrak{T} \to \mathfrak{R}$, and $\mathfrak{R}$ is $\mathcal{T}$-generated. Hence, $\mathfrak{R} \in \mathcal{R}$. ⌟

In Section II.6, we defined a variety to be a class $\mathcal{V} \subseteq \mathsf{Alg}(\mathbb{M})$ closed under E_*-quotients and $\mathcal{P}$-generated M-subobjects of arbitrary products, while a pseudo-variety is a class $\mathcal{V} \subseteq \mathsf{Alg}(\mathbb{M})$ 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 $\varepsilon : \mathfrak{C}|_\Delta \to \mathfrak{A}|_\Delta$. ⌟

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 $\mathsf{Alg}(\mathbb{M})$ induced by all algebras of the form $\mathbb{M}\Sigma$ with $\Sigma \in \mathsf{Alph}$ to the category of language algebras such that

(I) each $\mathcal{K}[\mathfrak{T}]$ forms a subalgebra of $\mathbb{L}\mathfrak{T}$,

(II) for $\varphi : \mathfrak{S} \to \mathfrak{T}$, the morphism $\mathcal{K}[\varphi] : \mathcal{K}[\mathfrak{T}] \to \mathcal{K}[\mathfrak{S}]$ is the corresponding restriction of $\mathbb{L}\varphi : \mathbb{L}\mathfrak{T} \to \mathbb{L}\mathfrak{S}$.

(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 $\mathcal{R}$. ⌟

Remark. (a) Note that every variety of languages $\mathcal{K}$ is uniquely determined by the domains of the algebras $\mathcal{K}[\mathfrak{X}]$ the algebraic structure being induced by that of $\mathbb{L}\mathfrak{X}$. 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 $\mathcal{K}[\mathfrak{S}] \to \mathcal{K}[\mathfrak{X}]$ 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. ⌟

We start with a characterisation of varieties in terms of closure operations like in Theorem II.6.12.

Definition 5.3. For a class $\mathcal{C}$ of $\mathbb{M}$-algebras we set

$$\begin{aligned}
\mathsf{C}(\mathcal{C}) &:= \{\, \mathfrak{A} \mid \mathfrak{A} \text{ a sort-accumulation point of some } \mathfrak{B} \in \mathcal{C} \,\}\,, \\
\mathsf{H}_0(\mathcal{C}) &:= \{\, \mathfrak{A} \mid \mathfrak{A} \text{ image of an } E\text{-morhpisms of some } \mathfrak{B} \in \mathcal{C} \,\}\,, \\
\mathsf{H}(\mathcal{C}) &:= \mathsf{H}_0(\mathcal{C}) \cap \mathcal{R}\,, \\
\mathsf{S}(\mathcal{C}) &:= \{\, \mathfrak{A} \mid \mathfrak{A} \text{ a } \mathcal{T}\text{-generated } M\text{-subalgebra of some } \mathfrak{B} \in \mathcal{C} \,\}\,, \\
\mathsf{P}(\mathcal{C}) &:= \{\, \mathfrak{A} \mid \mathfrak{A} \text{ a finite product of algebras in } \mathcal{C} \,\}\,.
\end{aligned}$$

⌟

Proposition 5.4. *Suppose that* $\mathsf{Alg}(\mathbb{M})$ *has coequalisers and let* $\mathcal{V} \subseteq \mathcal{R}$ *be a class of algebras.*

(a) $(\mathsf{HSP})^2(\mathcal{V}) = \mathsf{HSP}(\mathcal{V})$.

(b) *The following conditions are equivalent.*

  (1) $\mathcal{V}$ *is an* $\mathcal{R}$*-variety.*

  (2) $\mathcal{V} = \mathsf{HSP}(\mathcal{V})$ *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*

$$\mathfrak{A} \in \mathcal{V} \quad \text{iff} \quad \mathfrak{A}|_\Delta \in \mathcal{V}|_\Delta\,, \quad \textit{for all finite } \Delta \subseteq \Xi\,.$$

Proof. (a) We have shown in Theorem II.6.12 that (in the notation of the current section) $(\mathsf{H}_\circ\mathsf{SP})^2 = \mathsf{H}_\circ\mathsf{SP}$. Since $\mathcal{R}$ is closed under SP, it follows that

$$(\mathsf{HSP})^2 = (\mathsf{H}_\circ\mathsf{SP} \cap \mathcal{R})^2 = (\mathsf{H}_\circ\mathsf{SP})^2 \cap \mathcal{R} = \mathsf{H}_\circ\mathsf{SP} \cap \mathcal{R} = \mathsf{HSP}\,.$$

(b) (1) ⇒ (2) immediately follows from the closure properties of an $\mathcal{R}$-variety and (2) ⇒ (1) follows by (a).

(3) ⇒ (2) It is sufficient to show that $\mathsf{CHSP}(\mathcal{V}) = \mathcal{V}$. Hence, let $\mathfrak{A} \in \mathsf{CHSP}(\mathcal{V})$. Then $\mathfrak{A}|_\Delta \in \mathsf{HSP}(\mathcal{V}|_\Delta) = \mathcal{V}|_\Delta$, for all finite $\Delta \subseteq \Xi$. Consequently, $\mathfrak{A} \in \mathcal{V}$.

(2) ⇒ (3) If $\mathfrak{A}$ is an algebra with $\mathfrak{A}|_\Delta \in \mathcal{V}|_\Delta$, for all finite $\Delta \subseteq \Xi$. Then $\mathsf{C}(\mathcal{V}) = \mathcal{V}$ implies that $\mathfrak{A} \in \mathcal{V}$. For the other claim, note that

$$\mathcal{V} = \mathsf{HSP}(\mathcal{V}) \quad \text{implies} \quad \mathcal{V}|_\Delta = \mathsf{HSP}(\mathcal{V}|_\Delta)\,,$$

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. □

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 $\kappa : T_\xi \to \Omega$ a language where $\mathfrak{T} \in \mathcal{T}$ 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,* $\mathrm{Syn}(\kappa) \in \mathcal{V}$.

Proof. ($\Leftarrow$) is trivial since $\mathrm{syn}_\kappa : \mathfrak{T} \to \mathrm{Syn}(\kappa)$ recognises κ. For ($\Rightarrow$), consider a morphism $\varphi : \mathfrak{T} \to \mathfrak{A}$ recognising κ with $\mathfrak{A} \in \mathcal{V}$. Set $\mathfrak{B} := \mathfrak{T}/\ker\varphi$. Since $\mathrm{im}\,\varphi : \mathfrak{B} \to \mathfrak{A}$ is an M-morphism and $\ker\varphi : \mathfrak{T} \to \mathfrak{B}$ an E-morphism, it follows that $\mathfrak{B} \in \mathcal{V}$. Furthermore, $\ker\varphi$ recognises κ. As $\mathrm{Syn}(\kappa)$ is minimal, we can find a morphism $\rho : \mathfrak{A} \to \mathrm{Syn}(\kappa)$ with $\mathrm{syn}_\kappa = \rho \circ \ker\varphi$. Since syn_κ and $\ker\varphi$ are E-morphisms, it follows by Lemma II.1.7 (c) that $\rho \in E$. Hence, $\mathrm{Syn}(\kappa) \in \mathsf{H}(\mathcal{V}) = \mathcal{V}$. ☐

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.

(1) We start by proving that $\mathcal{K}[\mathfrak{T}]$ forms a subalgebra of $\mathbb{L}\mathfrak{T}$, that is, that

$$\langle\!\langle \mathcal{K}[\mathfrak{T}] \rangle\!\rangle_{\mathbb{L}} = \mathcal{K}[\mathfrak{T}]\,.$$

Since all operations of $\mathbb{L}\mathfrak{T}$ have finite arity, it is sufficient to show that

$$\langle\!\langle C \rangle\!\rangle_{\mathbb{L}} \subseteq \mathcal{K}[\mathfrak{T}]\,, \qquad \text{for all finite } C \subseteq \mathcal{K}[\mathfrak{T}]\,.$$

Hence, consider languages $\kappa_0, \dots, \kappa_{n-1} \in \mathcal{K}[\mathfrak{T}]$ for $\mathfrak{T} \in \mathcal{T}$. Then there are morphisms $\varphi_i : \mathfrak{T} \to \mathfrak{A}^i$ 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 $\lambda \in \mathcal{K}[\mathfrak{T}]$.

Let $\mathfrak{T} \to^e \mathfrak{B} \to^m \prod_{i<n} \mathfrak{A}^i$ 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 $\lambda \in \mathcal{K}[\mathfrak{T}]$ it is therefore sufficient to show that the morphism $e : \mathfrak{T} \to \mathfrak{B}$ recognises λ. Note that, being a morphism language algebras, $\mathbb{L}e$ commutes with all operations. Consequently, we

have

$$\begin{aligned}\lambda \in \langle\!\langle \kappa_0, \dots, \kappa_{n-1} \rangle\!\rangle_{\mathbb{L}} &= \big\langle\!\big\langle \mathbb{L}\varphi_0(\mu_0), \dots, \mathbb{L}\varphi_{n-1}(\mu_{n-1}) \big\rangle\!\big\rangle_{\mathbb{L}} \\ &= \big\langle\!\big\langle \mathbb{L}(\varphi_0 \times \dots \times \varphi_{n-1})(\bar{\mu}) \big\rangle\!\big\rangle_{\mathbb{L}} \\ &= \big\langle\!\big\langle \mathbb{L}(m \circ e)(\bar{\mu}) \big\rangle\!\big\rangle_{\mathbb{L}} \\ &= \big\langle\!\big\langle \mathbb{L}e\mathbb{L}m(\bar{\mu}) \big\rangle\!\big\rangle_{\mathbb{L}} \\ &= \mathbb{L}e\big\langle\!\big\langle \mathbb{L}m(\bar{\mu}) \big\rangle\!\big\rangle_{\mathbb{L}} ,\end{aligned}$$

which implies that $\lambda = \mathbb{L}e(\nu)$, for some $\nu \in \mathbb{L}B$.

(II) It remains to prove that, for every morphism $\psi : \mathfrak{S} \to \mathfrak{T}$ between algebras in $\mathcal{T}$, the morphism $\mathbb{L}\psi : \mathbb{L}\mathfrak{T} \to \mathbb{L}\mathfrak{S}$ restricts to a morphism $\mathcal{K}[\mathfrak{T}] \to \mathcal{K}[\mathfrak{S}]$. Hence, fix $\kappa \in \mathcal{K}_\xi[\mathfrak{T}]$. By assumption, there is some algebra $\mathfrak{A} \in \mathcal{V}$ recognising κ. Suppose that $\kappa = \mathbb{L}\varphi(\mu)$ where $\varphi : \mathfrak{T} \to \mathfrak{A}$ 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 $\varphi \circ \psi : \mathfrak{S} \to \mathfrak{A}$ recognises $\mathbb{L}\psi(\kappa)$. This implies that $\mathbb{L}\psi(\kappa) \in \mathcal{K}_\xi[\mathfrak{S}]$. ☐

It remains to prove the converse direction of the correspondence. We start with two lemmas.

Lemma 5.7. *Let $q : \mathfrak{A} \to \mathfrak{B}$ be an E-morphism and let $\mathfrak{T} \in \mathcal{T}$. Every language $\kappa : T_\xi \to \Omega$ recognised by $\mathfrak{B}$ is also recognised by $\mathfrak{A}$.*

Proof. Suppose that $\kappa = \psi^{-1}[\mu]$ where $\psi : \mathfrak{T} \to \mathfrak{B}$ and $\mu : B_\xi \to \Omega$. Since $\mathfrak{T}$ is projective with respect to E-morphisms, we can find a morphism $\varphi : \mathfrak{T} \to \mathfrak{A}$ 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 $\mathbb{L}q(\mu) \in \mathbb{L}_\xi A$. ☐

Lemma 5.8. *Let $\mathfrak{T} \in \mathcal{T}$ and suppose that $\kappa : T_\xi \to \Omega$ is recognised by an algebra $\mathfrak{C}$ such that there exists an M-morphism $\sigma : \mathfrak{C} \to \prod_{i<n} \mathfrak{A}^i$ 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 $\lambda_j : T_\xi \to \Omega$ is recognised by some factor $\mathfrak{A}^{k_j}$.

Proof. Suppose that $\kappa = \mathbb{L}\varphi(\mu)$ for $\varphi : \mathfrak{T} \to \mathfrak{C}$ and $\mu : C_\xi \to \Omega$. By Lemma 2.4, the function $\mathbb{L}\sigma$ is surjective. Fix some $\nu \in \mathbb{L}_\xi \prod_i A^i$ with $\mathbb{L}\sigma(\nu) = \mu$. Since $\mathfrak{A}^i \in \mathcal{R}$, the sets A^i_ξ 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

$$\nu = \omega\big[\mathbb{L}p_0(\lambda_0), \dots, \mathbb{L}p_{n-1}(\lambda_{n-1})\big],$$

where $p_i : \prod_j A^j \to A^i$ are the projections. Consequently,

$$\begin{aligned}
\kappa &= \mathbb{L}\varphi(\mu) \\
&= \mathbb{L}\varphi\mathbb{L}\sigma(\nu) \\
&= \mathbb{L}(\sigma \circ \varphi)\big(\omega\big[\mathbb{L}p_0(\lambda_0), \dots, \mathbb{L}p_{n-1}(\lambda_{n-1})\big]\big) \\
&= \omega\big[\mathbb{L}(\sigma \circ \varphi \circ p_0)(\lambda_0), \dots, \mathbb{L}(\sigma \circ \varphi \circ p_{n-1})(\lambda_{n-1})\big]\big),
\end{aligned}$$

where each language $\mathbb{L}(\sigma \circ \varphi \circ p_j)(\lambda_j)$ is recognised by some factor $\mathfrak{A}^{k_j}$. □

Theorem 5.9. *Let $\mathcal{K}$ be a variety of languages such that every $\kappa \in \mathcal{K}$ has a minimal algebra* $\mathrm{Syn}(\kappa) \in \mathcal{R}$*. 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 $\mathcal{S} :=$* $\{\, \mathrm{Syn}(\kappa) \mid \kappa \in \mathcal{K} \,\}$.

Proof. ($\Rightarrow$) Every language $\kappa \in \mathcal{K}$ is recognised by $\mathrm{Syn}(\kappa)$, which belongs to $\mathcal{V}$.

($\Leftarrow$) Set $\mathcal{V}_0 := \mathcal{S}$ and $\mathcal{V}_{n+1} := \mathsf{CHSP}(\mathcal{V}_n)$, 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 $n = 0$, suppose that $\lambda \in \mathbb{L}_\zeta S$ is recognised by $\mathrm{Syn}(\kappa)$ 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 $\varphi : \mathfrak{S} \to \mathrm{Syn}(\kappa)$ and a lanuage $\mu \in \mathbb{L}_\zeta\mathrm{Syn}(\kappa)$. As $\mathrm{Syn}(\kappa)$ is sort-wise noetherian, it follows by Proposition 4.14 that $\mu \in \langle\!\langle \mathrm{res}_\kappa \rangle\!\rangle_{\mathbb{L}}$. Thus,

$$\lambda = \mathbb{L}\varphi(\mu) \in \mathbb{L}\varphi[(\langle\!\langle \mathrm{res}_\kappa \rangle\!\rangle_{\mathbb{L}})_\zeta] \subseteq \mathbb{L}\varphi[\mathcal{K}_\zeta[\mathfrak{T}]] \subseteq \mathcal{K}_\zeta[\mathfrak{S}].$$

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 $\varphi : \mathfrak{T} \to \mathfrak{B}$ be a morphism recognising $\kappa = \mathbb{L}\varphi(\mu)$ with $\mu \in \mathbb{L}_\xi B$. Suppose that $\mathfrak{T} = \mathbb{M}\Sigma$ for $\Sigma \in \mathsf{Alph}$ 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 $\mathsf{HSP}(\mathcal{V}_n)$, we can find an algebra $\mathfrak{A} \in \mathsf{HSP}(\mathcal{S})$ and an E-morphism $\varepsilon : \mathfrak{A}|_\Delta \to \mathfrak{B}|_\Delta$. Since $\mathfrak{T}|_\Delta$ is projective with respect to E-morphisms, there exists a unique morphism $\psi_0 : \mathfrak{T}|_\Delta \to \mathfrak{A}|_\Delta$ with $\varepsilon \circ \psi_0 = \varphi|_\Delta$. We turn ψ_0 into a morphism $\psi : \mathfrak{T} \to \mathfrak{A}$ as follows. For $c \in \Sigma$, we set

$$\psi(\operatorname{sing}(c)) := \psi_0(\operatorname{sing}(c)) .$$

As $\mathfrak{T}$ is freely generated by $\operatorname{rng}\operatorname{sing}$, this uniquely determines a morphism $\psi : \mathfrak{T} \to \mathfrak{A}$ of $\mathbb{M}$-algebras. For $s \in \mathbb{M}_\xi\Sigma$, it follows that

$$\psi(s) = (\pi \circ \mathbb{M}(\psi_0 \circ \operatorname{sing}))(s) = (\psi_0 \circ \pi \circ \mathbb{M}\operatorname{sing})(s) = \psi_0(s) ,$$

where the second step follows from the fact that ψ_0 is a morphism of $\mathbb{M}|_\Delta$-algebras and $\xi \in \Delta$. Consequently,

$$(\mu \circ \varepsilon \circ \psi)(s) = (\mu \circ \varepsilon \circ \psi_0)(s) = (\mu \circ \varphi)(s) = \kappa(s) .$$

Thus, κ is recognised by $\mathfrak{A} \in \mathsf{HSP}(\mathcal{V}_n)$. Using Lemmas 5.7 and 5.8 and the closure properties of $\mathcal{K}$, it follows that $\kappa \in \mathcal{K}$. □

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, $\operatorname{Syn}(\kappa) \in \mathcal{V}$, for every language κ recognised by $\mathfrak{A}$.*

Proof. (⇒) If κ is recognised by $\mathfrak{A} \in \mathcal{V}$, it follows by Lemma 5.5 that $\operatorname{Syn}(\kappa) \in \mathcal{V}$.

($\Leftarrow$) Suppose that $\mathrm{Syn}(\kappa) \in \mathcal{V}$, for every language κ recognised by $\mathfrak{A}$. As $\mathfrak{A} \in \mathcal{A}$ is $\mathcal{T}$-generated, there exists an E-morphism $\varepsilon : \mathfrak{T} \to \mathfrak{A}$ 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 $\mathfrak{B}|_\Delta \to \mathfrak{A}|_\Delta$ 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 $A|_\Delta$ generates $\mathfrak{A}$. Furthermore, by Lemma 2.5 we can find finite sets H_ξ generating $\mathbb{L}_\xi A$, for $\xi \in \Xi$. For every $h \in H_\xi$ with $\xi \in \Delta$, we consider the language $\kappa_h := \mathbb{L}\varepsilon(h)$ and the morphism

$$\rho := \langle \mathrm{syn}_{\kappa_h} \rangle_h : \mathfrak{T} \to \prod_{\xi \in \Delta} \prod_{h \in H_\xi} \mathrm{Syn}(\kappa_h)\,.$$

Let $\mathfrak{T} \to^\eta \mathfrak{B} \to^\mu \prod_h \mathrm{Syn}(\kappa_h)$ 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 $p_h : \prod_{i \in H_\xi} \mathrm{Syn}(\kappa_i) \to \mathrm{Syn}(\kappa_h)$ be the projection. Then we have

$$\begin{aligned}\mathbb{L}\varepsilon(h) = \kappa_h = \mathbb{L}\mathrm{syn}_{\kappa_h}(\mathrm{res}_{\kappa_h}) \\ = \mathbb{L}(p_h \circ \mu \circ \eta)(\mathrm{res}_{\kappa_h}) = \mathbb{L}\eta\big(\mathbb{L}(p_h \circ \mu)(\mathrm{res}_{\kappa_h})\big)\,,\end{aligned}$$

where $\mathbb{L}(p_h \circ \mu)(\mathrm{res}_{\kappa_h}) \in \mathbb{L}_\xi B$. Consequently,

$$\mathbb{L}\varepsilon[\mathbb{L}_\xi A] = \mathbb{L}\varepsilon[\langle\!\langle H_\xi \rangle\!\rangle_{\mathbb{L}}] = \langle\!\langle \mathbb{L}\varepsilon[H_\xi] \rangle\!\rangle_{\mathbb{L}} \subseteq \mathbb{L}\eta[\mathbb{L}_\xi B]\,.$$

Since A_ξ is strongly finite-dimensional, we can therefore use Lemma 2.7 to find a function $q_\xi : B_\xi \to A_\xi$ with $\varepsilon_\xi = q_\xi \circ \eta_\xi$. Combining these into a single function $q : B|_\Delta \to A|_\Delta$ we obtain $\varepsilon|_\Delta = q \circ \eta|_\Delta$. By Lemma I.5.6, this implies that q is in fact a morphism $\mathfrak{B}|_\Delta \to \mathfrak{A}|_\Delta$ of $\mathbb{M}|_\Delta$-algebras. Furthermore, it follows by Lemma II.1.7 that $q \in E$. □

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* $\mathrm{Syn}(\kappa) \in \mathcal{V}$.

(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* $\{\, \mathrm{Syn}(\kappa) \mid \kappa \in \mathcal{K} \,\}$.

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 $\mathrm{Syn}(\kappa) \in \mathcal{V}$. By (2), this implies that $\kappa \in \mathcal{K}$. Conversely, if $\mathfrak{A}$ only recognises languages in $\mathcal{K}$, (2) implies that $\mathrm{Syn}(\kappa) \in \mathcal{V}$ for all languages κ recognised by $\mathfrak{A}$. By Lemma 5.10 it follows that $\mathfrak{A} \in \mathcal{V}$.

(3) $\Rightarrow$ (4) Let $\mathcal{V}_0$ be the $\mathcal{R}$-variety generated by $\{\, \mathrm{Syn}(\kappa) \mid \kappa \in \mathcal{K} \,\}$. 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 $\mathfrak{A} \in \mathcal{V}_0$. Hence, $\mathcal{V} \subseteq \mathcal{V}_0$.

For the converse inclusion it is sufficient to prove that $\mathrm{Syn}(\kappa) \in \mathcal{V}$, for all $\kappa \in \mathcal{K}$. Hence, fix a language $\kappa \in \mathcal{K}[\mathfrak{T}]$. It follows by Proposition 4.16 that every language $\lambda \in \mathbb{L}\mathfrak{S}$ recognised by $\mathrm{Syn}(\kappa)$ is of the form

$$\lambda = \mathbb{L}\varphi(\mu)\,, \qquad \text{for some } \mu \in \langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}} \text{ and } \varphi : \mathfrak{S} \to \mathfrak{T}\,.$$

Since $\mathcal{K}[\mathfrak{T}]$ is a language algebra, we have $\langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}} \subseteq \mathcal{K}[\mathfrak{T}]$. Consequently,

$$\lambda \in \mathbb{L}\varphi[\langle\!\langle \kappa \rangle\!\rangle_{\mathbb{L}}] \subseteq \mathbb{L}\varphi[\mathcal{K}[\mathfrak{T}]] \subseteq \mathcal{K}[\mathfrak{S}]\,.$$

We have shown that all languages recognised by $\mathrm{Syn}(\kappa)$ belong to $\mathcal{K}$. By assumption, this implies that $\mathrm{Syn}(\kappa) \in \mathcal{V}$. $\square$

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 $\mathcal{D} = \mathsf{Pos}$.

The Monad $\widehat{\mathbb{M}}_{\mathcal{A}}$

In this section we will make the following additional assumptions.

Conventions.

- *$\mathcal{T} := \{\, \mathbb{M}\mathbb{J}X \mid X$ finite $\}$*
- *$\mathcal{R}$ is a class of $\mathbb{M}$-algebras that is closed under $\mathcal{T}$-generated 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 $\iota : \mathbb{M}X \to \widehat{\mathbb{M}}X$. 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' $\beta : X \to A$. Let us denote the resulting value by $\operatorname{val}(t;\beta)$. For ordinary terms $t \in \mathbb{M}X$, this value should of course correspond to the value of t in $\mathfrak{A}$. Thus,

$$\operatorname{val}(\iota(t);\beta) = \pi(\mathbb{M}\beta(t)),$$

where $\pi(\mathbb{M}\beta(t))$ is the canonical extension of $\beta : X \to A$ to $\mathbb{M}X \to A$. Furthermore, $\operatorname{val}(t;\beta)$ should be compatible with morphisms of $\mathbb{M}$-algebras. That is,

$$\operatorname{val}(t;\varphi \circ \beta) = \varphi(\operatorname{val}(t;\beta)), \quad \text{for every morphism } \varphi : \mathfrak{A} \to \mathfrak{B}.$$

This leads to the following construction. We work in the category of all morphisms $\mathbb{M}X \to \mathfrak{A}$. In this category we consider the diagram of all $\beta : \mathbb{M}X \to \mathfrak{A}$ where $\mathfrak{A}$ belongs to a given class $\mathcal{R}$ of 'recognisers'. Then we take for $\iota : \mathbb{M}X \to \widehat{\mathbb{M}}X$ the limit. The morphisms $\widehat{\mathbb{M}}X \to \mathfrak{A}$ of the corresponding limiting cone can then be taken as our evaluation maps. The formal construction is as follows.

Definition 6.1. Let $\mathcal{R} \subseteq \mathsf{Alg}(\mathbb{M})$ be a subcategory of $\mathbb{M}$-algebras and $X \in \mathcal{D}$. We denote the comma category $(\mathbb{M}X \downarrow \mathsf{Alg}(\mathbb{M}))$ by $\mathcal{C}$, the subcategory $(\mathbb{M}X \downarrow \mathcal{R})$ by $\mathcal{C}_\circ$, and the inclusion diagram by $D : \mathcal{C}_\circ \to \mathcal{C}$.

(a) We denote by $\iota_\mathcal{R} : \mathbb{M}X \to \widehat{\mathbb{M}}_\mathcal{R}X$ the limit $\iota_\mathcal{R} := \lim D$ of D, and the limiting cone by $(\mathrm{val}_\mathcal{R}(-;\beta))_{\beta \in \mathcal{C}_\circ}$. If $\mathcal{R}$ is the category of all finitary $\mathbb{M}$-algebras, we drop the subscript and simply write $\widehat{\mathbb{M}}$, ι, and $\mathrm{val}(-;\beta)$.

(b) We turn $\widehat{\mathbb{M}}_\mathcal{R}$ into a functor as follows. Given $f : X \to Y$, the family $(\mathrm{val}(-;\beta \circ \mathbb{M}f))_\beta$ (where β ranges over all morphisms $\beta : \mathbb{M}Y \to \mathfrak{A} \in \mathcal{R}$) forms a cone from $\widehat{\mathbb{M}}X$ to D. As the cone $(\mathrm{val}(-;\beta))_\beta$ is limiting, there exists a unique function $f' : \widehat{\mathbb{M}}X \to \widehat{\mathbb{M}}Y$ such that

$$\mathrm{val}(-;\beta \circ \mathbb{M}f) = \mathrm{val}(-;\beta) \circ f', \quad \text{for all } \beta : \mathbb{M}Y \to \mathfrak{A} \in \mathcal{R}.$$

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 $\mathbb{U} : \mathcal{R} \to \mathcal{D}^\Xi$ 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 $\iota_\mathcal{R} : \mathbb{M}X \to \widehat{\mathbb{M}}_\mathcal{R}X$ exists.*

Proof. By assumption, $\mathcal{D}$ has arbitrary limits. It follows by Proposition I.5.7 that so has $\mathsf{Alg}(\mathbb{M})$. Now, let $D : \mathcal{C}_\circ \to \mathcal{C}$ be the diagram defining $\iota_\mathcal{R} : \mathbb{M}X \to \widehat{\mathbb{M}}_\mathcal{R}X$ and let $U : \mathcal{C} \to \mathsf{Alg}(\mathbb{M})$ be the forgetful functor mapping

$\beta : \mathbb{M}X \to \mathfrak{A}$ to the codomain $\mathfrak{A}$. Since $\mathcal{R}$ forms a set (up to isomorphism), so does the index category $\mathcal{C}_o$. As $\mathsf{Alg}(\mathbb{M})$ is complete, $U \circ D$ therefore has a limit $\mathfrak{T}$. 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 $\varphi : \mathbb{M}X \to \mathfrak{T}$ such that $\lambda_\beta \circ \varphi = \beta$, for all β. It is now straightforward to check that $\varphi : \mathbb{M}X \to \mathfrak{T}$ is the limit of D and $(\lambda_\beta)_\beta$ is the corresponding limiting cone. □

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, $\beta : \mathbb{M}X \to \mathfrak{A}$, $\varphi : \mathfrak{A} \to \mathfrak{B}$, and $f : Y \to X$ morphisms, and $s, t \in \widehat{\mathbb{M}}_\mathcal{R} X$.*

(a) $\mathrm{val}_\mathcal{R}(-;\beta) \circ \iota_\mathcal{R} = \beta$
(b) $\varphi \circ \mathrm{val}_\mathcal{R}(-;\beta) = \mathrm{val}_\mathcal{R}(-;\varphi \circ \beta)$
(c) $\mathrm{val}_\mathcal{R}(-;\beta) \circ \widehat{\mathbb{M}}_\mathcal{R} f = \mathrm{val}_\mathcal{R}(-;\beta \circ \mathbb{M}f)$
(d) $\big\langle \mathrm{val}_\mathcal{R}(-;\beta)\big\rangle_\beta = \mathrm{id}$.
(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 $\mathrm{val}_\mathcal{R}(\hat{s};\beta) = \beta(s)$.*

Proof. (a) By the definition of a cone, $\mathrm{val}_\mathcal{R}(-;\beta)$ is a morphism from $\iota_\mathcal{R} : \mathbb{M}X \to \widehat{\mathbb{M}}_\mathcal{R} X$ to $\beta : \mathbb{M}X \to \mathfrak{A}$. This is equivalent to (a).

(b) In the comma category, $\varphi : \mathfrak{A} \to \mathfrak{B}$ corresponds to a morphism from $\beta : \mathbb{M}X \to \mathfrak{A}$ to $\varphi \circ \beta : \mathbb{M}X \to \mathfrak{B}$. Hence, (b) holds again by definition of a cone.

(c) holds be 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 $\beta : \mathbb{M}X \to \mathfrak{A}$ satisfying

$$a_\gamma = \varphi(a_\beta)\,, \quad \text{for all } \varphi : \mathfrak{A} \to \mathfrak{B} \text{ with } \gamma = \varphi \circ \beta\,.$$

Then the function $\mathrm{val}_\mathcal{R}(-;\beta)$ is simply the projection to the component a_β. Consequently,

$$\big\langle \mathrm{val}_\mathcal{R}(-;\beta)\big\rangle_\beta = \mathrm{id}\,.$$

(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 $\operatorname{rng} \mathrm{val}_{\mathcal{R}}(-;\beta_0) \supseteq \operatorname{rng} \beta_0$ 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 $\beta_0(s) = \mathrm{val}_{\mathcal{R}}(\hat{s};\beta_0)$. By (b), it follows that

$$\beta(s) = i(\beta_0(s)) = i(\mathrm{val}_{\mathcal{R}}(\hat{s};\beta_0)) = \mathrm{val}_{\mathcal{R}}(\hat{s}; i \circ \beta_0) = \mathrm{val}_{\mathcal{R}}(\hat{s};\beta)\,. \square$$

Corollary 6.4. *Let X be a set and $f, g : C \to \widehat{\mathbb{M}}X$ functions.*

$$f = g \quad \text{iff} \quad \mathrm{val}(-;\beta) \circ f = \mathrm{val}(-;\beta) \circ g\,, \quad \text{for all } \beta : \mathbb{M}X \to A\,.$$

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

$$\begin{aligned} f = g \quad &\text{iff} \quad \langle \mathrm{val}_{\mathcal{R}}(-;\beta)\rangle_\beta \circ f = \langle \mathrm{val}_{\mathcal{R}}(-;\beta)\rangle_\beta \circ g \\ &\text{iff} \quad \langle \mathrm{val}_{\mathcal{R}}(-;\beta) \circ f\rangle_\beta = \langle \mathrm{val}_{\mathcal{R}}(-;\beta) \circ g\rangle_\beta\,. \end{aligned} \qquad \square$$

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 $\varepsilon := \iota_{\mathcal{R}} \circ \mathrm{sing}$ 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*

$$\mathrm{val}(-;\beta) \circ \mu = \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big)\,, \quad \text{for all } \beta\,.$$

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 $f : X \to Y$. For every $\beta : \mathbb{M}Y \to \mathfrak{A} \in \mathcal{R}$, Lemma 6.3 (c) implies that

$$\begin{aligned} \mathrm{val}(-;\beta) \circ \widehat{\mathbb{M}}f \circ \iota &= \mathrm{val}(-;\beta \circ \mathbb{M}f) \circ \iota \\ &= \beta \circ \mathbb{M}f = \mathrm{val}(-;\beta) \circ \iota \circ \mathbb{M}f\,. \end{aligned}$$

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 $\beta : \mathbb{M}X \to \mathfrak{A}$ with $\mathfrak{A} \in \mathcal{R}$, we have

$$\begin{aligned}
\beta &= \beta \circ \pi \circ \mathrm{sing} \\
&= \pi \circ \mathbb{M}\beta \circ \mathrm{sing} \\
&= \pi \circ \mathbb{M}\mathrm{val}(-;\beta) \circ \mathbb{M}\iota \circ \mathrm{sing} \\
&= \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big) \circ \iota \circ \mathbb{M}\iota \circ \mathrm{sing}\,.
\end{aligned}$$

Furthermore, for two such morphisms $\alpha : \mathbb{M}X \to \mathfrak{A}$ and $\beta : \mathbb{M}X \to \mathfrak{B}$ and a morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ with $\beta = \varphi \circ \alpha$, we have

$$\begin{aligned}
\varphi \circ \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\alpha)\big) &= \mathrm{val}\big(-; \varphi \circ \pi \circ \mathbb{M}\mathrm{val}(-;\alpha)\big) \\
&= \mathrm{val}\big(-; \pi \circ \mathbb{M}\varphi \circ \mathbb{M}\mathrm{val}(-;\alpha)\big) \\
&= \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\varphi \circ \alpha)\big) \\
&= \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big)\,.
\end{aligned}$$

Consequently, the morphisms $\big(\mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big)\big)_\beta$ form a cone from

$$\iota \circ \mathbb{M}\iota \circ \mathrm{sing} : \mathbb{M}X \to \widehat{\mathbb{M}}\widehat{\mathbb{M}}X$$

to the diagram $(\mathbb{M}X \downarrow \mathcal{R})$. As $\iota : \mathbb{M}X \to \widehat{\mathbb{M}}X$ is the limit of this cone, there exists a unique map $\mu : \widehat{\mathbb{M}}\widehat{\mathbb{M}}X \to \widehat{\mathbb{M}}X$ such that

$$\begin{aligned}
&\mu \circ \iota \circ \mathbb{M}\iota \circ \mathrm{sing} = \iota \\
\text{and}\quad &\mathrm{val}(-;\beta) \circ \mu = \mathrm{val}\big(-; \pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big)\,, \qquad \text{for all } \beta\,.
\end{aligned}$$

Note that the first of these equations follows from the second one since, for every β,

$$\begin{aligned}
&\mathrm{val}(-;\beta) \circ \mu \circ \iota \circ \mathbb{M}\iota \circ \mathrm{sing} \\
&= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big) \circ \iota \circ \mathbb{M}\iota \circ \mathrm{sing} \\
&= \pi \circ \mathbb{M}\mathrm{val}(-;\beta) \circ \mathbb{M}\iota \circ \mathrm{sing} \\
&= \pi \circ \mathbb{M}\beta \circ \mathrm{sing} \\
&= \beta \circ \pi \circ \mathrm{sing} \\
&= \beta \\
&= \mathrm{val}(-;\beta) \circ \iota\,,
\end{aligned}$$

which, by Corollary 6.4, implies that $\mu \circ \iota \circ \mathbb{M}\iota \circ \mathrm{sing} = \iota$.

Let us start by showing that these morphisms μ form a natural transformation. Hence, fix a function $f : X \to Y$. For every $\beta : \mathbb{M}Y \to \mathfrak{A}$, we have

$$\begin{aligned}
\mathrm{val}(-;\beta) \circ \mu \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}f &= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big) \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}f \\
&= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta) \circ \mathbb{M}\widehat{\mathbb{M}}f\big) \\
&= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta \circ \mathbb{M}f)\big) \\
&= \mathrm{val}(-;\beta \circ \mathbb{M}f) \circ \mu \\
&= \mathrm{val}(-;\beta) \circ \widehat{\mathbb{M}}f \circ \mu\,.
\end{aligned}$$

By Corollary 6.4, this implies that $\mu \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}f = \widehat{\mathbb{M}}f \circ \mu$.

The fact that $\varepsilon := \iota \circ \mathrm{sing}$ 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 $\beta : \mathbb{M}X \to \mathfrak{A}$, we have

$$\begin{aligned}
\mathrm{val}(-;\beta) \circ \mu \circ \varepsilon &= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big) \circ \iota \circ \mathrm{sing} \\
&= \pi \circ \mathbb{M}\mathrm{val}(-;\beta) \circ \mathrm{sing} \\
&= \mathrm{val}(-;\beta) \circ \pi \circ \mathrm{sing} \\
&= \mathrm{val}(-;\beta)\,,
\end{aligned}$$

$$\begin{aligned}
\mathrm{val}(-;\beta)\circ\mu\circ\widehat{\mathbb{M}}\varepsilon &= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\big)\circ\widehat{\mathbb{M}}\varepsilon \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\circ\mathbb{M}\varepsilon\big) \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\big(\mathrm{val}(-;\beta)\circ\iota\circ\mathrm{sing}\big)\big) \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\big(\beta\circ\mathrm{sing}\big)\big) \\
&= \mathrm{val}\big(-;\beta\circ\pi\circ\mathbb{M}\mathrm{sing}\big) \\
&= \mathrm{val}(-;\beta)\,, \\
\text{and}\quad \mathrm{val}(-;\beta)\circ\mu\circ\widehat{\mathbb{M}}\mu &= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\big)\circ\widehat{\mathbb{M}}\mu \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\circ\mathbb{M}\mu\big) \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\big(\mathrm{val}(-;\beta)\circ\mu\big)\big) \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\big)\big) \\
&= \mathrm{val}\big(-;\pi\circ\mathbb{M}\mathrm{val}(-;\beta)\big)\circ\mu \\
&= \mathrm{val}(-;\beta)\circ\mu\circ\mu\,.
\end{aligned}$$

By Corollary 6.4, this implies that

$$\mu\circ\varepsilon = \mathrm{id}\,,\quad \mu\circ\widehat{\mathbb{M}}\varepsilon = \mathrm{id}\,,\quad\text{and}\quad \mu\circ\widehat{\mathbb{M}}\mu = \mu\circ\mu\,. \qquad\square$$

The next lemma states that, without loss of generality, we may assume that the morphisms $\beta : \mathbb{M}X \to \mathfrak{A}$ 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) $\mathcal{C}_\circ = (\mathbb{M}X \downarrow \mathcal{R})$ *is cofiltered.*

(b) *In the definition of $\widehat{\mathbb{M}}_{\mathcal{R}}X$, we can restrict the category $\mathcal{C}_\circ$ to the surjective morphisms without changing the result.*

Proof. (a) There are two axioms to check. First, let $\alpha : \mathbb{M}X \to \mathfrak{A}$ and $\beta : \mathbb{M}X \to \mathfrak{B}$ be two objects of $\mathcal{C}_\circ$. We have to find some $\gamma : \mathbb{M}X \to \mathfrak{C}$ and morphisms $\varphi : \gamma \to \alpha$ and $\psi : \gamma \to \beta$. Set $\gamma := \langle\alpha,\beta\rangle : \mathbb{M}X \to \mathfrak{A}\times\mathfrak{B}$ 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, $\mathfrak{C} \in \mathcal{R}$, $\operatorname{im} \gamma \in \mathcal{C}_0$, and we have morphisms $p : \operatorname{im} \gamma \to \alpha$ and $q : \operatorname{im} \gamma \to \beta$, where $p : C \to A$ and $q : C \to B$ are the two projections.

For the second axiom, consider two morphisms $\varphi, \psi : \alpha \to \beta$ with $\alpha : \mathbb{M}X \to \mathfrak{A}$ and $\beta : \mathbb{M}X \to \mathfrak{B}$ in $\mathcal{C}_0$. 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, $\operatorname{rng} \alpha \subseteq C$. Let $\alpha = i \circ \alpha_0$ be the EM-factorisation of α and let $\mathfrak{D}$ be the codomain of α_0. Then $\mathfrak{D}$ is a $\mathcal{T}$-generated M-subalgebra of $\mathfrak{A}$, which implies that $\mathfrak{D} \in \mathcal{R}$. Furthermore, $i : \alpha_0 \to \alpha$ satisfies $\varphi \circ i = \psi \circ i$.

(b) Let $\mathcal{C}_{00}$ be the full subcategory of $\mathcal{C}_0 = (\mathbb{M}X \downarrow \mathcal{R})$ consisting of all morphisms $\beta : \mathbb{M}X \to \mathfrak{A}$ that are surjective. It is sufficient to show that the inclusion $\mathcal{C}_{00} \to \mathcal{C}_0$ if final, which means, we have to establish the following two properties.

(I) Every $\beta \in \mathcal{C}_0$ factorises through some $\beta_0 \in \mathcal{C}_{00}$.

(II) For all $\alpha, \alpha' \in \mathcal{C}_{00}$, $\beta \in \mathcal{C}_0$, and all morphisms $\varphi : \alpha \to \beta$ and $\varphi' : \alpha' \to \beta$, there is some $\gamma \in \mathcal{C}_{00}$ with morphisms $\psi : \gamma \to \alpha$ and $\psi' : \gamma \to \alpha'$ such that $\varphi \circ \psi = \varphi' \circ \psi'$.

(I) Given $\beta : \mathbb{M}X \to \mathfrak{A}$, let $\mathfrak{A}_0$ be the subalgebra of $\mathfrak{A}$ induced by $\operatorname{rng} \beta$, let $i : \mathfrak{A}_0 \to \mathfrak{A}$ be the inclusion function, and $\beta_0 : \mathbb{M}X \to \mathfrak{A}_0$ be the corestriction of β. Then $\beta = i \circ \beta_0$. Since $\mathfrak{A}_0$ is finitely generated by $\beta_0[X]$ and $\mathcal{R}$ is closed under finitely generated subalgebras, we have $\mathfrak{A}_0 \in \mathcal{R}$ and $\beta_0 \in \mathcal{C}_0$.

(II) Consider $\alpha : \mathbb{M}X \to \mathfrak{A}$, $\alpha' : \mathbb{M}X \to \mathfrak{A}'$ in $\mathcal{C}_{00}$, $\beta : \mathbb{M}X \to \mathfrak{B}$ in $\mathcal{C}_0$, and $\varphi : \alpha \to \beta$ and $\varphi' : \alpha' \to \beta$. Let $\mathfrak{C}$ be the subalgebra of $\mathfrak{A} \times \mathfrak{A}'$ induced by the range of $\gamma := \langle \alpha, \alpha' \rangle : \mathbb{M}X \to \mathfrak{A} \times \mathfrak{A}'$. As above, it follows that $\mathfrak{C} \in \mathcal{R}$ and $\gamma \in \mathcal{C}_{\circ\circ}$. The two projections $p : \mathfrak{C} \to \mathfrak{A}$ and $p' : \mathfrak{C} \to \mathfrak{A}'$ are morphisms of $\mathcal{C}_{\circ\circ}$ 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 $\mathcal{R}_\rho$ the class of their ρ-reducts. There exists a unique morphism $\varphi : \mathbb{N} \Rightarrow \widehat{\mathbb{M}}_{\mathcal{R}_\rho}$ of monads such that $\iota_{\mathcal{R}_\rho} = \varphi \circ \rho$.*

Proof. Fix a set X. For every morphism $\beta : \mathbb{M}X \to \mathfrak{A}|_\rho$ with $\mathfrak{A} \in \mathcal{R}$, we define $\mu_\beta : \mathbb{N}X \to \mathfrak{A}$ by

$$\mu_\beta := \pi \circ \mathbb{N}(\beta \circ \mathrm{sing})\,.$$

To see that μ_β is a morphism of $\mathbb{N}$-algebras, note that

$$\begin{aligned} \pi \circ \mathbb{N}\mu_\beta &= \pi \circ \mathbb{N}(\pi \circ \mathbb{N}(\beta \circ \mathrm{sing})) \\ &= \pi \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{N}(\beta \circ \mathrm{sing}) \\ &= \pi \circ \mathrm{flat} \circ \mathbb{N}\mathbb{N}(\beta \circ \mathrm{sing}) \\ &= \pi \circ \mathbb{N}(\beta \circ \mathrm{sing}) \circ \mathrm{flat} = \mu_\beta \circ \mathrm{flat}\,. \end{aligned}$$

We claim that $(\mu_\beta)_\beta$ is a cone from $\rho : \mathbb{M}X \to \mathbb{N}X$ to $(\mathbb{M}X \downarrow \mathcal{R}_\rho)$. First, we have

$$\begin{aligned} \mu_\beta \circ \rho &= \pi \circ \mathbb{N}(\beta \circ \mathrm{sing}) \circ \rho \\ &= (\pi \circ \rho) \circ \mathbb{M}(\beta \circ \mathrm{sing}) \\ &= \beta \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} \\ &= \beta \end{aligned}$$

where the third step follows from the fact that $\pi \circ \rho$ is the product of the $\mathbb{M}$-algebra $\mathfrak{A}|_\rho$, while $\mathrm{flat} : \mathbb{M}\mathbb{M}X \to \mathbb{M}X$ is the product of $\mathbb{M}X$. Consequently, μ_β is a morphism $\rho \to \beta$ of the comma category.

Furthermore, for a morphism $\psi : \beta \to \gamma$ of the comma category, we have

$$\begin{aligned}\psi \circ \mu_\beta &= \psi \circ \pi \circ \mathbb{N}(\beta \circ \mathrm{sing})\\ &= \pi \circ \mathbb{N}(\psi \circ \beta \circ \mathrm{sing})\\ &= \pi \circ \mathbb{N}(\gamma \circ \mathrm{sing})\\ &= \mu_\gamma\,,\end{aligned}$$

where the second step follows since the morphism $\psi : \beta \to \gamma$ is induced by a morphism $\psi : \mathrm{codom}(\beta) \to \mathrm{codom}(\gamma)$ of $\mathbb{N}$-algebras.

It follows that $(\mu_\beta)_\beta$ is a cone. As $(\mathrm{val}(-;\beta))_\beta$ is the limiting cone, we obtain a unique morphism $\varphi : \mathbb{N}X \to \widehat{\mathbb{M}}_{\mathcal{R}_\rho}X$ 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 $f : X \to Y$. Note that, by definition, φ is the unique function such that

$$\mathrm{val}(-;\beta) \circ \varphi = \mu_\beta\,.$$

Consequently, for each $\beta : \mathbb{M}X \to \mathfrak{A}|_\varphi$ with $\mathfrak{A} \in \mathcal{R}$, we have

$$\begin{aligned}\mathrm{val}(-;\beta) \circ \varphi \circ \mathbb{N}f &= \mu_\beta \circ \mathbb{N}f\\ &= \pi \circ \mathbb{N}(\beta \circ \mathrm{sing}) \circ \mathbb{N}f\\ &= \pi \circ \mathbb{N}(\beta \circ \mathrm{sing} \circ f)\\ &= \pi \circ \mathbb{N}(\beta \circ \mathbb{M}f \circ \mathrm{sing})\\ &= \mu_{\beta\circ\mathbb{M}f}\\ &= \mathrm{val}(-;\beta \circ \mathbb{M}f) \circ \varphi = \mathrm{val}(-;\beta) \circ \widehat{\mathbb{M}}_{\mathcal{R}_\rho}f \circ \varphi\,.\end{aligned}$$

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 $\varepsilon = \iota_{\mathcal{R}_\rho} \circ \mathrm{sing}$ be the unit map of $\widehat{\mathbb{M}}_{\mathcal{R}_\rho}$. Then

$$\varepsilon = \iota_{\mathcal{R}_\rho} \circ \mathrm{sing} = \varphi \circ \rho \circ \mathrm{sing} = \varphi \circ \mathrm{sing}\,,$$

where the last step follows from the fact that ρ is a morphism of monads. For the second equation, consider a morphism $\beta : \mathbb{M}X \to \mathfrak{A}|_\rho$ 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

$$\mathrm{val}(-;\beta) \circ \mu = \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big)\,.$$

Hence,

$$\begin{aligned}
\mathrm{val}(-;\beta) \circ \mu \circ \varphi \circ \mathbb{N}\varphi &= \mathrm{val}\big(-;\pi \circ \mathbb{M}\mathrm{val}(-;\beta)\big) \circ \varphi \circ \mathbb{N}\varphi \\
&= \mu_{\pi \circ \mathbb{M}\mathrm{val}(-;\beta)} \circ \mathbb{N}\varphi \\
&= \pi \circ \mathbb{N}(\pi \circ \mathbb{M}\mathrm{val}(-;\beta) \circ \mathrm{sing}) \circ \mathbb{N}\varphi \\
&= \pi \circ \mathbb{N}(\pi \circ \mathrm{sing} \circ \mathrm{val}(-;\beta) \circ \varphi) \\
&= \pi \circ \mathbb{N}\mu_\beta \\
&= \pi \circ \mathbb{N}(\pi \circ \mathbb{N}(\beta \circ \mathrm{sing})) \\
&= \pi \circ \mathbb{N}\pi \circ \mathbb{N}\mathbb{N}(\beta \circ \mathrm{sing}) \\
&= \pi \circ \mathrm{flat} \circ \mathbb{N}\mathbb{N}(\beta \circ \mathrm{sing}) \\
&= \pi \circ \mathbb{N}(\beta \circ \mathrm{sing}) \circ \mathrm{flat} \\
&= \mu_\beta \circ \mathrm{flat} \\
&= \mathrm{val}(-;\beta) \circ \varphi \circ \mathrm{flat}\,.
\end{aligned}$$

By Corollary 6.4, it therefore follows that $\mu \circ \varphi \circ \mathbb{N}\varphi = \varphi \circ \mathrm{flat}$. □

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 $\mathcal{D} = \mathsf{Pos}$ 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 $A \in \mathsf{Pos}$ equipped with a topology that is compact and has the following separation property: for every pair of elements $a, b \in A$ with $a \nleq b$, 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 $f : A \to B$ 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). ⌟

Remark. Every Priestley space is a Stone space, i.e., compact, Hausdorff, and totally disconnected. ⌟

Theorem 6.9 (Priestley). *The category* PSp *is equivalent to* $\mathsf{Dist}^{\mathrm{op}}$.

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 $\mu_i : A \to B_i$ 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 $D : I \to \mathsf{Pos}^{\Xi}$ 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 $\mathbb{M} : \mathsf{Pos}^{\Xi} \to \mathsf{Pos}^{\Xi}$ for which we have defined a lifting to $\mathsf{PSp}^{\Xi} \to \mathsf{PSp}^{\Xi}$, we write $\mathsf{PAlg}(\mathbb{M})$ for the category of $\mathbb{M}$-algebras in PSp^{Ξ}. ⌟

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 $(\mathrm{val}(-;\beta))_\beta$, respectively. Then it follows by Lemma 6.3 (e) that the embedding $\iota_{\mathcal{R}} : \mathbb{M}X \to \widehat{\mathbb{M}}_{\mathcal{R}}X$ 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 $(x^{n!})_{n<\omega}$. ⌟

Lemma 6.11. *The forgetful functor* $\mathbb{U} : \mathsf{PSp}^{\Xi} \to \mathsf{Pos}^{\Xi}$ *reflects limits. More precisely, the limit* $\lim D$ *of a diagram* $D : I \to \mathsf{PSp}^{\Xi}$ *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 $a_l = Df(a_k)$, for all I-morphisms $f : k \to l$. Hence, $B_\xi = \bigcap_f H_f$ where

$$H_f := \left\{ (a_i)_i \in \textstyle\prod_i D_\xi(i) \;\middle|\; Df(a_k) = a_l \right\}, \quad \text{for } f : k \to l\,.$$

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 from

$$\left(\mu_k^{-1}[(Df)^{-1}[C]] \cap \mu_l^{-1}[C]\right) \cup \left(\mu_k^{-1}[(Df)^{-1}[C']] \cap \mu_l^{-1}[C']\right),$$

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 $a \nleq b$ in B. By the definition of the ordering of a limit in Pos^{Ξ}, there exists an index $i \in I$ with $\mu_i(a) \nleq \mu_i(b)$. 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 $f : A \to B$ (in Pos^{Ξ}) such that $\lambda_i = \mu_i \circ f$, for all i. Similarly, there exists a unique morphism $g : B \to A$ 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 applying the same universality argument two more times to obtain $f \circ g = \mathrm{id}$ and $g \circ f = \mathrm{id}$. Therefore, B and A with the cone topology are isomorphic as topological space. □

The following version of compactness and its corollary below contain our key topology-based argument.

Proposition 6.12. *Let $E : I \to \mathsf{PSp}$ be a cofiltered diagram. If all spaces $E(i)$, $i \in I$, are non-empty, so is the limit* $\lim E$.

Proof. Let $\mu_i : \lim E \to E(i)$, $i \in I$, be the morphisms of the limiting cone. Note that, by the Theorem of Tychonoff, the product $\prod_{i \in I} E_\xi(i)$ compact and Hausdorff.

For each morphism $f : k \to l$ of I, consider the set

$$H_f := \left\{ (a_i)_i \in \textstyle\prod_i E(i) \;\middle|\; Ef(a_k) = a_l \right\}.$$

In the proof of Lemma 6.11, we have shown that the sets 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

$$H_{f_0} \cap \dots \cap H_{f_n} = \varnothing .$$

Suppose that $f_i : k_i \to l_i$. Since I is cofiltered, there exists some $m \in I$ and morphisms $g_i : m \to k_i$ 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 $(b_i)_i \in H_{f_0} \cap \cdots \cap H_{f_n} \neq \varnothing$. A contradiction. $\square$

The following consequence is what we will need below.

Lemma 6.13. *Let $D : I \to \mathsf{PSp}^{\Xi}$ be a cofiltered diagram and $(\mu_i)_i$ a cone from $A \in \mathsf{PSp}^{\Xi}$ to D where each $\mu_i : A \to D(i)$ is surjective. The induced morphism $\varphi : A \to \lim D$ is surjective.*

Proof. Fix an element $c \in \lim D$ and let $(\lambda_i)_i$ be the limiting cone. To show that $\varphi^{-1}(c) \neq \varnothing$, we define a diagram $E : I \to \mathsf{PSp}$ 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, $E(i) \in \mathsf{PSp}$. For a morphism $f : i \to j$ of I, we let $Ef : E(i) \to E(j)$ be the inclusion map. This is well-defined, since

$$\begin{aligned} a \in E(i) \quad &\Rightarrow \quad \mu_i(a) = \lambda_i(c) \\ &\Rightarrow \quad \mu_j(a) = Df(\mu_i(a)) = Df(\lambda_i(c)) = \lambda_j(c) \\ &\Rightarrow \quad a \in E(j) . \end{aligned}$$

By Proposition 6.12, the limit $B := \lim E \neq \varnothing$ is non-empty. Fix $b \in B$. Note that the inclusion maps $E(i) \to A_\xi$ constitute a natural transformation $\sigma : E \Rightarrow A_\xi^I$, where A_ξ^I denotes the constant diagram $I \to \mathsf{PSp}$ with value A_ξ and identity maps everywhere. Let $\psi : B \to A_\xi$ be the induced map between the correpsoning 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* $\mathcal{V} \subseteq \mathcal{R} \subseteq \mathsf{Alg}(\mathbb{M})$.

(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* $\beta : \mathbb{M}X \to \mathfrak{A}$ *where* $\mathfrak{A} \in \mathcal{V}$ *and* $X \in \mathcal{D}^{\Xi}$.

$$\begin{tikzcd}
& \widehat{\mathbb{M}}_{\mathcal{R}}X \arrow[rr, "\rho"] \arrow[dr, "\mathrm{val}_{\mathcal{R}}(-;\beta)"'] & & \widehat{\mathbb{M}}_{\mathcal{V}}X \arrow[dl, "\mathrm{val}_{\mathcal{V}}(-;\beta)"] \\
\mathbb{M}_{\mathcal{R}}X \arrow[ur, "\iota_{\mathcal{R}}"] \arrow[rr, "\beta"'] & & \mathfrak{A} &
\end{tikzcd}$$

(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* $\rho_X : \widehat{\mathbb{M}}_{\mathcal{R}}X \to \widehat{\mathbb{M}}_{\mathcal{V}}X$ *is surjective.*

Proof. (a) Fix a set $X \in \mathcal{D}^{\Xi}$. The family $(\mathrm{val}_{\mathcal{R}}(-;\beta))_{\beta \in (\mathbb{M}X \downarrow \mathcal{V})}$ forms a cone from $\widehat{\mathbb{M}}_{\mathcal{R}}X$ to the diagram defining $\widehat{\mathbb{M}}_{\mathcal{V}}X$. As $(\mathrm{val}_{\mathcal{V}}(-;\beta))_{\beta \in (\mathbb{M}X \downarrow \mathcal{V})}$ is limiting, there exists a unique map $\rho_X : \widehat{\mathbb{M}}_{\mathcal{R}}X \to \widehat{\mathbb{M}}_{\mathcal{V}}X$ such that

$$\mathrm{val}_{\mathcal{V}}(-;\beta) \circ \rho_X = \mathrm{val}_{\mathcal{R}}(-;\beta)\,, \qquad \text{for all } \beta : \mathbb{M}X \to \mathfrak{A}\,.$$

As the equation $\mathrm{val}_{\mathcal{V}}(-;\beta) \circ \iota_{\mathcal{V}} = \beta$ 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

$f : X \to Y$. Then

$$\begin{aligned}
\mathrm{val}_{\mathcal{R}}(-;\beta) \circ \widehat{\mathbb{M}}_{\mathcal{R}} f \circ \rho &= \mathrm{val}_{\mathcal{R}}(-;\beta \circ \mathbb{M}f) \circ \rho \\
&= \mathrm{val}_{\mathcal{V}}(-;\beta \circ \mathbb{M}f) \\
&= \mathrm{val}_{\mathcal{V}}(-;\beta) \circ \widehat{\mathbb{M}}_{\mathcal{V}} f \\
&= \mathrm{val}_{\mathcal{R}}(-;\beta) \circ \rho \circ \widehat{\mathbb{M}}_{\mathcal{V}} f\,.
\end{aligned}$$

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 $\beta : \mathbb{M}X \to \mathfrak{A}$ with $\mathfrak{A} \in \mathcal{V}$, we have

$$\begin{aligned}
\mathrm{val}_{\mathcal{V}}(-;\beta) \circ \rho \circ \mu_{\mathcal{R}} &= \mathrm{val}_{\mathcal{R}}(-;\beta) \circ \mu_{\mathcal{R}} \\
&= \mathrm{val}_{\mathcal{R}}\big(-;\pi \circ \mathbb{M}\mathrm{val}_{\mathcal{R}}(-;\beta)\big) \\
&= \mathrm{val}_{\mathcal{V}}\big(-;\pi \circ \mathbb{M}\mathrm{val}_{\mathcal{R}}(-;\beta)\big) \circ \rho \\
&= \mathrm{val}_{\mathcal{V}}\big(-;\pi \circ \mathbb{M}\mathrm{val}_{\mathcal{V}}(-;\beta) \circ \mathbb{M}\rho\big) \circ \rho \\
&= \mathrm{val}_{\mathcal{V}}\big(-;\pi \circ \mathbb{M}\mathrm{val}_{\mathcal{V}}(-;\beta)\big) \circ \widehat{\mathbb{M}}\rho \circ \rho \\
&= \mathrm{val}_{\mathcal{V}}(-;\beta) \circ \mu_{\mathcal{V}} \circ \widehat{\mathbb{M}}\rho \circ \rho
\end{aligned}$$

$$\text{and}\quad \begin{aligned}
\mathrm{val}_{\mathcal{V}}(-;\beta) \circ \rho \circ \varepsilon_{\mathcal{R}} &= \mathrm{val}_{\mathcal{V}}(-;\beta) \circ \rho \circ \iota_{\mathcal{R}} \circ \mathrm{sing} \\
&= \mathrm{val}_{\mathcal{R}}(-;\beta) \circ \iota_{\mathcal{R}} \circ \mathrm{sing} \\
&= \beta \circ \mathrm{sing} \\
&= \mathrm{val}_{\mathcal{V}}(-;\beta) \circ \iota_{\mathcal{V}} \circ \mathrm{sing} = \mathrm{val}_{\mathcal{V}}(-;\beta) \circ \varepsilon_{\mathcal{V}}\,.
\end{aligned}$$

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 $\beta : \mathbb{M}X \to \mathfrak{A}$ 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 $\mathbb{M}_\xi X \neq \varnothing$. By Lemma 6.3 (e), it follows that these are exactly the same sorts ξ with $\widehat{\mathbb{M}}_{\mathcal{V},\xi}X \neq \varnothing$, $\widehat{\mathbb{M}}_{\mathcal{R},\xi}X \neq \varnothing$, and with $A_\xi \neq \varnothing$, 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 $\mathrm{val}_{\mathcal{V}}(-;\beta)$ and $\mathrm{val}_{\mathcal{R}}(-;\beta)$ are continuous. Furthermore, since we restricted the diagram to surjective maps β, $\mathrm{val}_{\mathcal{R}}(-;\beta) \circ \iota = \beta$ implies that the value maps $\mathrm{val}_{\mathcal{R}}(-;\beta)$ 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 $\rho : \widehat{\mathbb{M}}_{\mathcal{R}}X \to \widehat{\mathbb{M}}_{\mathcal{V}}X$ is surjective. □

Finitely 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 $D : I \to \mathcal{C}$ with limit C and limiting cone $(\lambda_i)_{i\in I}$, and for every morphism $f : C \to A$, there exists an index $k \in I$ and an essentially unique morphism $g : D(k) \to A$ such that $f = g \circ \lambda_k$. Essentially uniqueness here means that, if $g' : D(k) \to A$ is another morphism with $f = g' \circ \lambda_k$, then there exists an I-morphisms $h : l \to k$ with $g \circ Dh = g' \circ Dh$. ⌟

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 $(\mathsf{Dist}^{\Xi})^{\mathrm{op}}$. 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 $D : I \to \mathsf{Dist}^{\Xi}$ be a directed diagram with colimit $\mathfrak{N}$ and limiting cocone $(\lambda_i)_{i \in I}$, and let $\varphi : \mathfrak{L} \to \mathfrak{N}$ be a lattice homomorphism. Since L is finite there exists some index $i \in I$ such that $\operatorname{rng} \varphi \subseteq \operatorname{rng} \lambda_i$. For every $a \in L$, fix an element $f(a) \in \lambda_i^{-1}(\varphi(a))$. This defines a function $f : L \to D(i)$ 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 $a, b \in L$,

$$\begin{aligned} & D(i,l)(f(a) \sqcap f(b)) = D(i,l)(f(a \sqcap b)) \\ \text{and}\quad & D(i,l)(f(a) \sqcup f(b)) = D(i,l)(f(a \sqcup b)) . \end{aligned}$$

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 $\mu' : \mathfrak{L} \to D(l')$ with $\lambda_{l'} \circ \mu' = \varphi$. Fixing some 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. □

It remains to transfer this result from PSp^{Ξ} to $\mathsf{PAlg}(\widehat{\mathbb{M}})$.

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 $I : \mathcal{R} \to \mathsf{PSp}^{\Xi}$ maping an $\mathbb{M}$-algebra $\mathfrak{A} \in \mathcal{R}$ to its universe A (equipped with the discrete topology). By definition, this means that

$$\widehat{\mathbb{M}}_{\mathcal{R}} = \mathrm{Ran}_I \, I$$

is the *right Kan extension* of I along itself. Furthermore, we can compute such a Kan extension as

$$(\mathrm{Ran}_I \, I)(X) = \int_{\mathfrak{A} \in \mathcal{R}} (I\mathfrak{A})^{\mathsf{PSp}^{\Xi}(X, I\mathfrak{A})} = \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathsf{PSp}^{\Xi}(X, A)} ,$$

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

$$\mathsf{PSp}^{\Xi}(B, A^X) \cong \mathsf{Set}^{\Xi}(X, \mathsf{PSp}^{\Xi}(B, A)),$$

for sets $X \in \mathsf{Set}^{\Xi}$ and spaces $A, B \in \mathsf{PSp}^{\Xi}$. For a fixed space A, it follows that $A^{(-)} : (\mathsf{Set}^{\Xi})^{\mathrm{op}} \to \mathsf{PSp}^{\Xi}$ is the right adjoint of the hom-functor $\mathsf{PSp}^{\Xi}(-, A) : \mathsf{PSp}^{\Xi} \to (\mathsf{Set}^{\Xi})^{\mathrm{op}}$. This in particular implies that $A^{(-)}$ preserves all limits. Furthermore we have seen in Lemma 6.16, that a space $A \in \mathsf{PSp}^{\Xi}$ is finitely copresentable if, and only if, the hom-functor $\mathsf{PSp}^{\Xi}(-, A)$ 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 PSp^{Ξ}, it follows that the composition $A^{\mathsf{PSp}^{\Xi}(-,A)}$ preserves cofiltered limits, for every $\mathfrak{A} \in \mathcal{R}$.

Given a cofiltered diagram $D : J \to \mathsf{PSp}^{\Xi}$, we therefore have

$$\begin{aligned}
\widehat{\mathbb{M}}_{\mathcal{R}}(\lim D) &= (\mathrm{Ran}_I\, I)(\lim D) \\
&= \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathsf{PSp}^{\Xi}(\lim_{j \in J} D(j), A)} \\
&= \int_{\mathfrak{A} \in \mathcal{R}} \lim_{j \in J} A^{\mathsf{PSp}^{\Xi}(D(j), A)} \\
&= \lim_{j \in J} \int_{\mathfrak{A} \in \mathcal{R}} A^{\mathsf{PSp}^{\Xi}(D(j), A)} = \lim_{j \in J} \widehat{\mathbb{M}}_{\mathcal{R}} D(j),
\end{aligned}$$

where the fourth step follows by the fact that an end is a limit and limits commute. □

Lemma 6.19. *Let $\mathcal{C}$ be a category, $\mathbb{M} : \mathcal{C} \to \mathcal{C}$ a monad preserving cofiltered limits, and $\mathfrak{A}$ an $\mathbb{M}$-algebra with finitely copresentable domain A. Then $\mathfrak{A}$ is finitely copresentable in* $\mathsf{Alg}(\mathbb{M})$.

Proof. Fix a cofiltered diagram $D : I \to \mathsf{Alg}(\mathbb{M})$ with limit $\mathfrak{C}$ and a limiting cone $(\lambda_i)_{i \in I}$, and let $\varphi : \mathfrak{C} \to \mathfrak{A}$ 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 $f : D(k) \to A$ such that $\varphi = f \circ \lambda_k$. Since φ and λ_k are morphisms of $\mathbb{M}$-algebras, we have

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

$$\begin{tikzcd}
\mathbb{M}C \arrow[dr, "\pi"] \arrow[ddd, "\mathbb{M}\varphi"'] \arrow[drrr, bend left, "\mathbb{M}\lambda_k"] \arrow[drrrr, bend left, "\mathbb{M}\lambda_m"] & & & & \\
& C \arrow[r, "\lambda_k"] \arrow[d, "\varphi"'] & D(k) \arrow[dl, "f"] & \mathbb{M}D(k) \arrow[l, "\pi"'] \arrow[ddlll, bend left, "\mathbb{M}f"] & \mathbb{M}D(m) \arrow[l, "\mathbb{M}Dh"'] \arrow[dlll, bend left, "f'"] \\
& A & & & \\
\mathbb{M}A \arrow[uur, "\pi"] & & & &
\end{tikzcd}$$

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 $f' : \mathbb{M}D(m) \to A$ 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 $g : m' \to m$ and $h : m' \to k$. Replacing m by m' and f' by $f' \circ Dg$, we may therefore assume w.l.o.g. that there exists an I-morphism $h : m \to k$. By essential uniqueness of f', we can find a morphism $h' : m' \to m$

$$f \circ \pi \circ \mathbb{M}Dh \circ \mathbb{M}Dh' = f' \circ \mathbb{M}Dh' = \pi \circ \mathbb{M}f \circ \mathbb{M}Dh \circ \mathbb{M}Dh'\,.$$

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 $\qquad \varphi = f \circ \lambda_k = f \circ Dh \circ \lambda_m = \psi \circ \lambda_m\,.$

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. □

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 $X \in \mathsf{PSp}^{\Delta}$, the $\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta}$-algebra $\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta} X$ is finitely copresentable in $\mathsf{PAlg}(\widehat{\mathbb{M}}_{\mathcal{R}}|_{\Delta})$.*

Proof. By Proposition 6.17, the set X (with the discrete topology) is finitely copresentable in PSp^{Δ}. As we have shown in Propositon 6.18 that $\widehat{\mathbb{M}}|_{\Delta}$ preserves cofiltered limits, the claim therefore follows by Lemma 6.19. □

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 $\mathcal{D} = \mathsf{Pos}$.

Convention. *In this section, we denote by $\mathcal{R}$ the class of all finitary $\mathbb{M}$-algebras.*

Definition 7.1. Let $X \in \mathsf{Set}^{\Xi}$ 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

$$\mathrm{val}_{\mathcal{R}}(s; \beta) \leq \mathrm{val}_{\mathcal{R}}(t; \beta)\,, \qquad \text{for all } \beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A}\,.$$

We write $\mathfrak{A} \vDash s \leq t$ to denote this fact.

(c) The $\mathbb{M}$-*theory* $\mathrm{Th}(\mathcal{V})$ 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 $\mathcal{V}$.

$$\mathrm{Mod}_{\mathcal{V}}(\Phi) := \left\{ \mathfrak{A} \in \mathcal{V} \;\middle|\; \mathfrak{A} \vDash s \leq t \text{ for all } s \leq t \in \Phi \right\}.$$ ⌟

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*

$$s \leq t \in \mathrm{Th}(\mathcal{V}) \quad \text{iff} \quad \rho_{\mathcal{V}}(s) \leq \rho_{\mathcal{V}}(t),$$

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

$$\begin{aligned}
&\mathfrak{A} \vDash s \leq t, && \text{for all } \mathfrak{A} \in \mathcal{V} \\
\text{iff}\quad &\mathrm{val}_{\mathcal{R}}(s;\beta) \leq \mathrm{val}_{\mathcal{R}}(t;\beta), && \text{for all } \beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A} \in \mathcal{V} \\
\text{iff}\quad &\mathrm{val}_{\mathcal{V}}(\rho_{\mathcal{V}}(s);\beta) \leq \mathrm{val}_{\mathcal{V}}(\rho_{\mathcal{V}}(t);\beta), && \text{for all } \beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A} \in \mathcal{V} \\
\text{iff}\quad &\rho_{\mathcal{V}}(s) \leq \rho_{\mathcal{V}}(t). &&
\end{aligned}$$

□

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 $\mathrm{Mod}_{\mathcal{V}}(\Phi)$ 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 $\mathfrak{B}^i \in \mathrm{Mod}_{\mathcal{V}}(\Phi)$. Let $p_k : \prod_i B^i \to B^k$ be the projection. For $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A}$ it follows that

$$p_k(\mathrm{val}(s;\beta)) = \mathrm{val}(s; p_k \circ \beta) \leq \mathrm{val}(t; p_k \circ \beta) = p_k(\mathrm{val}(t;\beta)),$$

where the second step follows from the fact that $\mathfrak{B}^k \vDash s \leq t$. As the ordering of the product is defined component-wise, this implies that $\mathrm{val}(s;\beta) \leq \mathrm{val}(t;\beta)$. Consequently, $\mathfrak{A} \in \mathrm{Mod}_{\mathcal{V}}(\Phi)$.

Next, consider a quotient $q : \mathfrak{B} \to \mathfrak{A}$ with $\mathfrak{B} \in \mathrm{Mod}_{\mathcal{V}}(\Phi)$. Fix $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A}$. Since q is surjective, we can use Corollary II.6.4 to find some $\gamma : \mathbb{M}\mathbb{J}X \to \mathfrak{A}$ with $\beta = q \circ \gamma$. Then

$$\begin{aligned}
\mathrm{val}(s;\beta) = \mathrm{val}(s; q \circ \gamma) &= q(\mathrm{val}(s;\gamma)) \\
&\leq q(\mathrm{val}(t;\gamma)) = \mathrm{val}(t; q \circ \gamma) = \mathrm{val}(t;\beta),
\end{aligned}$$

where the third step follows by monotonicity of q and the fact that $\mathfrak{B} \models s \leq t$. Consequently, $\mathfrak{A} \in \mathrm{Mod}_{\mathcal{V}}(\Phi)$.

Finally, suppose that $\mathfrak{A}$ is a sort-accumulation point of $\mathrm{Mod}_{\mathcal{V}}(\Phi)$. Fix $s \leq t \in \Phi$ over X and $\beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A}$. We have to show that

$$\mathrm{val}_{\mathcal{V}}(s;\beta) \leq \mathrm{val}_{\mathcal{V}}(t;\beta)\,.$$

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 $\mathfrak{B} \in \mathrm{Mod}_{\mathcal{V}}(\Phi)$ and a surjective morphism $\mu : \mathfrak{B}|_{\Delta} \to \mathfrak{A}|_{\Delta}$. By Corollary II.6.4, we can find a morphism $\gamma : \mathbb{M}|_{\Delta}\mathbb{J}X \to \mathfrak{B}|_{\Delta}$ 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})

$$\begin{aligned}
\mathrm{val}_{\mathcal{V}}(s;\beta|_{\Delta}) &= \mathrm{val}_{\mathcal{V}}(s;\mu \circ \gamma) \\
&= \mu(\mathrm{val}_{\mathcal{V}}(s;\gamma)) \\
&\leq \mu(\mathrm{val}_{\mathcal{V}}(t;\beta)) \\
&= \mathrm{val}_{\mathcal{V}}(t;\mu \circ \gamma) = \mathrm{val}_{\mathcal{V}}(t;\beta|_{\Delta})\,.
\end{aligned}$$

Since $\mathrm{val}_{\mathcal{V}}(-;\beta|_{\Delta}) = \mathrm{val}_{\mathcal{V}}(-;\beta) \restriction \widehat{\mathbb{M}}|_{\Delta}\mathbb{J}X$, it follows that $\mathfrak{A} \models s \leq t$. ☐

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 $\beta : \mathbb{M}\mathbb{J}X \to \mathfrak{A}$, for some finite set X. The claim follows since $\mathrm{val}(-;\beta) \circ \iota = \beta$ implies that that $\mathrm{val}(-;\beta) : \widehat{\mathbb{M}}_{\mathcal{V}}\mathbb{J}X \to \mathfrak{A}$ is also surjective.

($\supseteq$) Let $\mathfrak{A}$ be finitary and $\varphi : \widehat{\mathbb{M}}_{\mathcal{V}}\mathbb{J}X \to \mathfrak{A}$ 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 $\Delta \subseteq \Xi$ there is some algebra $\mathfrak{B} \in \mathcal{V}$ and a surjective morphism $\mathfrak{B}|_{\Delta} \to \mathfrak{A}|_{\Delta}$. Hence, fix $\Delta \subseteq \Xi$. 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 $\mathsf{PAlg}(\widehat{\mathbb{M}}_{\mathcal{V}}|_{\Delta})$. Therefore, there exists an algebra $\mathfrak{B} \in \mathcal{V}$ and morphisms $\beta : \mathbb{M}\mathbb{J}X \to \mathfrak{B}$ and $\mu : \mathfrak{B}|_{\Delta} \to \mathfrak{A}|_{\Delta}$ such that $\varphi|_{\Delta} = \mu \circ \mathrm{val}(-;\beta)|_{\Delta}$. Since $\varphi|_{\Delta}$ is surjective, so is μ. Consequently, $\mathfrak{A}|_{\Delta}$ is a quotient of $\mathfrak{B}|_{\Delta}$ and $\mathfrak{B} \in \mathcal{V}$. □

Corollary 7.5. *Let $\mathcal{V}$ and $\mathcal{W}$ be $\mathcal{R}$-varieties.*

(a) $\mathcal{V} \subseteq \mathcal{W} \quad$ iff $\quad \mathrm{Th}(\mathcal{V}) \supseteq \mathrm{Th}(\mathcal{W})$.

(b) $\mathrm{Mod}(\mathrm{Th}(\mathcal{V})) = \mathcal{V}$.

Proof. (a) ($\Rightarrow$) follows immediately by definition. For ($\Leftarrow$), let $\rho_{\mathcal{V},X} : \widehat{\mathbb{M}}\mathbb{J}X \to \widehat{\mathbb{M}}_{\mathcal{V}}\mathbb{J}X$ and $\rho_{\mathcal{W},X} : \widehat{\mathbb{M}}\mathbb{J}X \to \widehat{\mathbb{M}}_{\mathcal{W}}\mathbb{J}X$ be the morphisms from Theorem 6.14. It follows by Lemma 7.2 that

$$\mathrm{Th}(\mathcal{W}) \subseteq \mathrm{Th}(\mathcal{V}) \quad \text{implies} \quad \ker \rho_{\mathcal{W},X} \subseteq \ker \rho_{\mathcal{V},X}\,.$$

Hence, we can use the Factorisation Lemma to find a morphism $q_X : \widehat{\mathbb{M}}_{\mathcal{W}}\mathbb{J}X \to \widehat{\mathbb{M}}_{\mathcal{V}}\mathbb{J}X$ 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 $\mathcal{W} := \mathrm{Mod}(\mathrm{Th}(\mathcal{V}))$ is an $\mathcal{R}$-variety. We have to show that $\mathcal{V} = \mathcal{W}$.

($\subseteq$) Let $\mathfrak{A} \in \mathcal{V}$. Then we have $\mathfrak{A} \vDash s \leq t$, for every $s \leq t$ in $\mathrm{Th}(\mathcal{V})$. This implies that $\mathfrak{A} \in \mathrm{Mod}(\mathrm{Th}(\mathcal{V})) = \mathcal{W}$.

($\supseteq$) By (a) it is sufficient to prove that $\mathrm{Th}(\mathcal{W}) \supseteq \mathrm{Th}(\mathcal{V})$. Hence, let $s \leq t$ be in $\mathrm{Th}(\mathcal{V})$. Then $\mathfrak{A} \vDash s \leq t$, for all $\mathfrak{A} \in \mathrm{Mod}(\mathrm{Th}(\mathcal{V})) = \mathcal{W}$, which implies that $s \leq t$ belongs to $\mathrm{Th}(\mathcal{W})$. □

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 $\mathcal{V} = \mathrm{Mod}_{\mathcal{R}}(\Phi)$, 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 $\mathcal{V} = \mathrm{Mod}_{\mathcal{R}}(\mathrm{Th}(\mathcal{V}))$. $\square$

Notes

Monadic frameworks for formal language theory were initially put forward by Bojańczyk [9]. Later on they have been generalised and fleshed out in [38, 42, 32, 10, 7].

The material on minimal algebras in Section 3 is taken from lecture notes by Bojańczyk [10]. Theorem 3.11 is unpublished work of Bojańczyk and Plotkin.

The sections on pro-finitary terms and Reiterman's Theorem are based on [20, 42]. An introduction to Priestley spaces can be found, for instance, in Chapter 11 of [21]. For a thorough introduction to profinite groups, see [37].

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 $\langle L, \mathcal{M}, \mathrm{Mod}\rangle$ consisting of a set $L \in \mathsf{Set}^{\Xi}$ of *formulae*, a class $\mathcal{M} \in \mathcal{D}^{\Xi}$ of *models*, and a *model function*

$$\mathrm{Mod} : L \to \mathcal{D}^{\Xi}(\mathcal{M}, \tilde{\Omega}) .$$

We call $\mathrm{Mod}(\varphi) : \mathcal{M}_\xi \to \Omega$ 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 $\langle L, \mathcal{M}, \mathrm{Mod}\rangle$, we define the *satisfaction function*

$$\langle\!\langle \cdot ; \cdot \rangle\!\rangle : \mathbb{V}\mathcal{M} \times L \to \mathbb{V}\tilde{\Omega} \quad \text{by} \quad \langle\!\langle M ; \varphi \rangle\!\rangle := \mathrm{Mod}(\varphi)(M) ,$$

and the *theory function*

$$\mathrm{Th}_L := \langle \mathrm{Mod}(\varphi) \rangle_{\varphi \in L} : \mathcal{M} \to \tilde{\Omega}^L .$$

We call $\mathrm{Th}_L(M)$ the *L-theory* of M.

(c) The congruence

$$\mathrm{eq}(L) := \ker \mathrm{Th}_L : \mathcal{M} \to \mathcal{M}/(\ker \mathrm{Th}_L) \subseteq \tilde{\Omega}^L .$$

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 $\langle\!\langle \cdot ; \cdot \rangle\!\rangle : \mathcal{M} \times L \to \{0, 1\}$ by the relation $\vDash \subseteq \mathcal{M} \times L$ which is defined by

$$M \vDash \varphi \quad :\text{iff} \quad \langle\!\langle M; \varphi \rangle\!\rangle = 1 .$$

Then we have

$$\mathrm{Th}_L(M) = \{ \varphi \in L \mid M \vDash \varphi \} ,$$
$$\text{and} \quad \mathrm{Mod}(\varphi) = \{ M \in \mathcal{M} \mid M \vDash \varphi \} .$$

Note that $\mathrm{Mod}(\varphi)$ is a set, while $\mathrm{Th}_L(M) \in \mathcal{D}$. For $\mathcal{D} = \mathsf{Pos}$, the ordering on $\mathrm{Th}_L(M)$ is given by

$$\begin{aligned} \varphi \leq \psi \quad &\text{iff} \quad N \vDash \varphi \;\Rightarrow\; N \vDash \psi , \quad \text{for all } N \in \mathcal{M} , \\ &\text{iff} \quad \varphi \text{ implies } \psi . \end{aligned}$$

(b) In Set, L-equivalence $\mathrm{eq}(L)$ reduces to the well-known relation $\equiv_L$ given by

$$M \equiv_L N \quad :\text{iff} \quad \langle\!\langle M; \varphi \rangle\!\rangle = \langle\!\langle N; \varphi \rangle\!\rangle , \quad \text{for all } \varphi \in L .$$

In the unweighted case, this reads

$$M \equiv_L N \quad :\text{iff} \quad M \vDash \varphi \Leftrightarrow N \vDash \varphi , \quad \text{for all } \varphi \in L .$$

In Pos, we obtain a preorder $\sqsubseteq_L$ instead, which is given by

$$M \sqsubseteq_L N \quad :\text{iff} \quad \langle\!\langle M; \varphi \rangle\!\rangle \leq \langle\!\langle N; \varphi \rangle\!\rangle , \quad \text{for all } \varphi \in L ,$$
$$\text{or} \quad M \sqsubseteq_L N \quad :\text{iff} \quad M \vDash \varphi \Rightarrow N \vDash \varphi , \quad \text{for all } \varphi \in L .$$ ⌟

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

$$w \vDash \mathcal{A} \quad :\text{iff} \quad \mathcal{A} \text{ accepts the input } w\,.$$

(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

$$\big\langle \mathrm{MSO}[\Sigma, X_1, X_2], \mathsf{Alg}[\Sigma, X_1, X_2], \mathrm{Mod}\big\rangle$$

where $\mathrm{MSO}[\Sigma, X_1, X_2]$ 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 $\mathsf{Alg}[\Sigma, X_1, X_2]$ is the set of all triples $\langle \mathfrak{A}, \beta_1, \beta_2\rangle$ where $\mathfrak{A}$ is a Σ-structure and $\beta_1 : X \to A$ and $\beta_2 : X \to \wp(A)$ 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 $M \vDash \varphi$ 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 used weak definability below.

Definition 1.2. Let $\mathcal{L} = \langle L, \mathcal{M}, \mathrm{Mod}\rangle$ and $\mathcal{L}' = \langle L', \mathcal{M}', \mathrm{Mod}\rangle$ be Ω-valued logics.

(a) The *extension of $\mathcal{L}$ by Ω-operations* is the logic

$$\mathcal{L}[\Omega] := \big\langle L[\Omega], \mathcal{M}, \mathrm{Mod}\big\rangle\,,$$

where $L[\Omega]$ is the set of all pairs $\langle \omega, \bar\varphi\rangle$ such that $\omega : \Omega^n \to \Omega$ is an Ω-operation of some arity $n < \omega$ and $\bar\varphi \in L^n$ an n-tuple of formulae. We use the syntax $\omega[\bar\varphi]$ for such a tuple. The satisfaction function is defined by

$$\mathrm{Mod}(\omega[\bar\varphi]) := \omega\big[\mathrm{Mod}(\varphi_0), \dots, \mathrm{Mod}(\varphi_{n-1})\big]\,.$$

(b) A class $\mathcal{C} : \mathcal{M}_\xi \to \Omega$ is *strongly L-definable* if $\mathcal{C} = \mathrm{Mod}(\varphi)$, for some $\varphi \in L_\xi$. It is *(weakly) L-definable* if it is strongly $L[\Omega]$-definable.

(c) A *strong morphism* $\langle \lambda, \mu \rangle : \mathcal{L} \to \mathcal{L}'$ consists of functions $\lambda : L \to L'$ and $\mu : \mathcal{M}' \to \mathcal{M}$ such that

$$\langle\!\langle M'; \lambda(\varphi) \rangle\!\rangle = \langle\!\langle \mu(M'); \varphi \rangle\!\rangle\,, \quad \text{for all } \varphi \in L \text{ and } M' \in \mathcal{M}'.$$

(d) A *(weak) morphism* $\langle \lambda, \mu \rangle : \mathcal{L} \to \mathcal{L}'$ is a strong morphism $\mathcal{L} \to \mathcal{L}'[\Omega]$. We denote the category of all logics and their (weak) morphisms by wLog. ⌟

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 $\mathcal{D} = \mathsf{Pos}$ and $\Omega = \{0, 1\}$, this is the case if the logic in question is closed under finite meets and joins. Similarly, for $\mathcal{D} = \mathsf{Set}$ 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 $\mathrm{MSO}[\Gamma, \varnothing, \varnothing] \to \mathrm{MSO}[\Sigma, \varnothing, \varnothing]$ since we can construct, for every formula $\varphi \in \mathrm{MSO}[\Gamma, \varnothing, \varnothing]$, some formula $\varphi^\tau \in \mathrm{MSO}[\Sigma, \varnothing, \varnothing]$ with

$$\tau(\mathfrak{A}) \vDash \varphi \quad \text{iff} \quad \mathfrak{A} \vDash \varphi^\tau, \quad \text{for all } \Sigma\text{-structures } \mathfrak{A}\,.$$

(Just replace in φ every atomic formula containing a relation symbol $R \in \Gamma$ by the formula from τ defining R.) ⌟

Let us isolate a few simple conditions for when a class of models is definable.

Lemma 1.3. *Let $\langle L, \mathcal{M}, \mathrm{Mod} \rangle$ and $\langle L', \mathcal{M}', \mathrm{Mod} \rangle$ be Ω-valued logics.*

(a) *A class $\mathcal{C} : \mathcal{M}_\xi \to \Omega$ is L-definable if, and only if,*

$$\mathrm{eq}_\xi(\Delta) \leq \mathcal{C}\,, \quad \textit{for some finite } \Delta \subseteq L_\xi\,.$$

(b) *For sort-wise finite sets $\Delta \subseteq L$ and $\Delta' \subseteq L'$, and a function $f : \mathcal{M} \to \mathcal{M}'$ the following two statements are equivalent:*

(1) $\mathrm{eq}_\xi(\Delta) \leq \ker\left(\mathrm{eq}_\xi(\Delta') \circ f\right)$

(2) *If $\mathcal{C} : \mathcal{M}'_\xi \to \Omega$ is Δ'-definable, then $f^{-1}[\mathcal{C}] : \mathcal{M}_\xi \to \Omega$ is Δ-definable.*

Proof. (a) ($\Rightarrow$) Suppose that

$$\mathcal{C} = \omega[\mathrm{Mod}(\varphi_0), \dots, \mathrm{Mod}(\varphi_{n-1})] .$$

For $\Delta := \{\varphi_0, \dots, \varphi_{n-1}\}$, it follows that

$$\begin{aligned}\mathcal{C} &= \omega \circ \langle \mathrm{Mod}(\varphi_0), \dots, \mathrm{Mod}(\varphi_{n-1}) \rangle \\ &= \omega \circ \mathrm{Th}_\Delta \\ &= (\omega \circ \mathrm{im}\,\mathrm{Th}_\Delta) \circ \ker \mathrm{Th}_\Delta = (\omega \circ \mathrm{im}\,\mathrm{Th}_\Delta) \circ \mathrm{eq}_\xi(\Delta) .\end{aligned}$$

Consequently, $\mathrm{eq}_\xi(\Delta) \leq \mathcal{C}$.

($\Leftarrow$) Set $\Theta := \mathcal{M}/\mathrm{eq}(\Delta)$ and fix a morphism $\omega_0 : \Theta_\xi \to \Omega$ such that $\mathcal{C} = \omega_0 \circ \mathrm{eq}_\xi(\Delta)$. Since $\mathrm{im}\,\mathrm{Th}_{\Delta,\xi} : \Theta_\xi \to \Omega^\Delta$ is an M-morphism and Ω is injective, we can extend ω_0 to a function $\omega : \Omega^\Delta \to \Omega$ with $\omega \circ \mathrm{im}\,\mathrm{Th}_{\Delta,\xi} = \omega_0$. Consequently,

$$\begin{aligned}\mathcal{C} &= \omega_0 \circ \mathrm{eq}_\xi(\Delta) \\ &= \omega \circ \mathrm{im}\,\mathrm{Th}_{\Delta,\xi} \circ \ker \mathrm{Th}_{\Delta,\xi} \\ &= \omega \circ \mathrm{Th}_{\Delta,\xi} \\ &= \omega[\mathrm{Mod}(\varphi_0), \dots, \mathrm{Mod}(\varphi_{n-1})] ,\end{aligned}$$

where $\varphi_0, \dots, \varphi_{n-1}$ is an enumeration of Δ_ξ.

(b) (1) $\Rightarrow$ (2) Suppose that $\mathcal{C} : \mathcal{M}'_\xi \to \Omega$ is Δ'-definable. By (1) and (a), it then follows that

$$\mathrm{eq}_\xi(\Delta) \leq \ker\left(\mathrm{eq}_\xi(\Delta') \circ f\right) \leq \ker\left(\mathcal{C} \circ f\right) .$$

Hence, (a) (applied to the logic Δ instead of L) implies that $f^{-1}[\mathcal{C}]$ is Δ-definable.

(2) ⇒ (1) Suppose that $\mathrm{eq}_\xi(\Delta) \leq \ker\big(\mathrm{eq}_\xi(\Delta') \circ f\big)$. Note that every class $\mathrm{Mod}(\varphi)$ with $\varphi \in \Delta'$ is trivially Δ'-definable. By assumption, it therefore follows that $f^{-1}[\mathrm{Mod}(\varphi)]$ is Δ-definable. Hence, (a) implies that

$$\mathrm{eq}_\xi(\Delta) \leq \ker\big(\mathrm{Mod}(\varphi) \circ f\big)\,, \qquad \text{for all } \varphi \in \Delta'_\xi .$$

Consequently,

$$\begin{aligned}
\mathrm{eq}_\xi(\Delta) &\leq \prod_{\varphi \in \Delta'_\xi} \ker\big(\mathrm{Mod}(\varphi) \circ f\big) \\
&= \ker\big(\langle \mathrm{Mod}(\varphi)\rangle_{\varphi \in \Delta'_\xi} \circ f\big) \\
&= \ker\big(\mathrm{Th}_{\Delta',\xi} \circ f\big) \\
&= \ker\big(\mathrm{eq}_\xi(\Delta') \circ f\big)\,. \qquad \square
\end{aligned}$$

Remark. In the category Pos, the above conditions simplify to the following ones.

(a) $M \sqsubseteq_\Delta N \quad \text{implies} \quad \mathcal{C}(M) \leq \mathcal{C}(N)\,.$

(b) $M \sqsubseteq_\Delta N \quad \text{implies} \quad f(M) \sqsubseteq_{\Delta'} f(N)\,.$

For Set, we just have to replace $\sqsubseteq_\Delta$ by $\equiv_\Delta$, 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 $\mathcal{L} = \langle L, \mathcal{M}, \mathrm{Mod}\rangle$ and $\mathcal{L}' = \langle L', \mathcal{M}', \mathrm{Mod}\rangle$ be Ω-valued logics and $\mu : \mathcal{M}' \to \mathcal{M}$ a function. The following statements are equivalent.*

(1) *There exists a function $\lambda : L \to L'$ such that $\langle \lambda, \mu\rangle : \mathcal{L} \to \mathcal{L}'$ is a morphism of logics.*

(2) *If $\mathcal{C} : \mathcal{M}_\xi \to \Omega$ is L-definable, then $\mu^{-1}[\mathcal{C}] : \mathcal{M}'_\xi \to \Omega$ is L'-definable.*

(3) *For every sort-wise finite $\Delta \subseteq L$, there exists a sort-wise finite $\Delta' \subseteq L'$ such that*

$$\mathrm{eq}(\Delta') \leq \ker\big(\mathrm{eq}(\Delta) \circ \mu\big)\,.$$

Proof. (3) ⇒ (2) Suppose that $\mathcal{C} : \mathcal{M}_\xi \to \Omega$ 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

$$\mathrm{eq}(\Delta') \leq \ker\left(\mathrm{eq}(\Delta) \circ \mu\right).$$

By Lemma 1.3 (b) this implies that $\mathcal{C} \circ \mu$ is Δ'-definable. In particular, it is L'-definable.

(2) ⇒ (1) We define $\lambda : L \to L'$ as follows. For each $\varphi \in L$, the class $\mathrm{Mod}(\varphi)$ is obviously L-definable. By assumption it follows that the preimage $\mathrm{Mod}(\varphi) \circ \mu$ is L'-definable, i.e, there are formulae $\psi'_0, \dots, \psi'_n \in L'$ and an Ω-operation $\omega : \Omega^n \to \Omega$ such that

$$\mathrm{Mod}(\varphi) \circ \mu = \omega[\mathrm{Mod}(\psi_0), \dots, \mathrm{Mod}(\psi_{n-1})].$$

We can therefore set $\lambda(\varphi) := \langle \omega, \psi_0, \dots, \psi_{n-1} \rangle$.

(1) ⇒ (3) Given $\Delta \subseteq L$, we set

$$\Delta' := \bigcup \left\{ \{\psi_0, \dots, \psi_{n-1}\} \;\middle|\; \varphi \in \Delta\,,\ \lambda(\varphi) = \langle \omega, \psi_0, \dots, \psi_{n-1} \rangle \right\}.$$

Then

$$\mathrm{Mod}(\varphi) \circ \mu = \omega[\mathrm{Mod}(\psi_0), \dots, \mathrm{Mod}(\psi_{n-1})]\,, \quad \text{for every } \varphi \in \Delta\,.$$

Consequently, there exists a function $\chi : \Omega^{\Delta'} \to \Omega^{\Delta}$ such that

$$\langle \mathrm{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ \mu = \chi \circ \langle \mathrm{Mod}(\psi) \rangle_{\psi \in \Delta'}\,.$$

Hence,

$$\mathrm{Th}_{\Delta'} = \langle \mathrm{Mod}(\psi) \rangle_{\psi \in \Delta'} \leq \langle \mathrm{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ \mu = \mathrm{Th}_{\Delta} \circ \mu\,,$$

which implies that

$$\mathrm{eq}(\Delta') = \ker \mathrm{Th}_{\Delta'} \leq \ker\left(\mathrm{Th}_{\Delta} \circ \mu\right) = \mathrm{eq}(\Delta) \circ \mu\,. \qquad \square$$

Here, we are mainly interested in logics whose class of models is of the form $\mathcal{M} = \mathbb{M}\Sigma$ with $\Sigma \in \mathsf{Alph}$, 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 $L : \mathsf{Alph} \to \mathsf{wLog}$ such that

- for every alphabet Σ, the image $L[\Sigma]$ is a logic over Σ,
- for every function $f : \Sigma \to \Gamma$, the image $L[f]$ is a morphism $\langle \lambda, \mu \rangle : L[\Sigma] \to L[\Gamma]$ 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 $\kappa : \mathbb{M}A \to \Omega$ *L-definable*, if the language $\kappa \circ \mathbb{M}\iota : \mathbb{M}\mathbb{J}\mathbb{V}A \to \Omega$ is $L[\mathbb{J}\mathbb{V}A]$-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 $\mathrm{MSO}[\hat{\Sigma}, \varnothing, \varnothing]$ 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

$$\mathrm{eq}(\Delta[\Sigma]) \text{ is an } \mathbb{M}\text{-congruence on } \mathbb{M}\Sigma\,, \quad \text{for all alphabets } \Sigma\,. \qquad \lrcorner$$

Example. For words $u_0, \dots, u_{n-1}, v_0, \dots, v_{n-1} \in \Sigma^*$ we have

$$u_i \equiv_{\mathrm{MSO}_m} v_i\,, \quad \text{for all } i\,, \quad \text{implies} \quad u_0 \cdots u_{n-1} \equiv_{\mathrm{MSO}_m} v_0 \cdots v_{n-1}\,,$$

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 $\mathbb{M}A = A^*$. $\lrcorner$

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 $\Theta_\Delta : \mathsf{Alph} \to \mathsf{Alg}(\mathbb{M})$, and*
- *an epimorphic natural transformation $\theta_\Delta : (\mathbb{M} \restriction \mathsf{Alph}) \Rightarrow \Theta_\Delta$*

such that

$$(\theta_\Delta)_\Sigma = \mathrm{eq}(\Delta[\Sigma])\,, \quad \textit{for every } \Sigma\,.$$

Proof. $(\Leftarrow)$ Given $\Phi \subseteq L$, choose $\Phi \subseteq \Delta \subseteq L$ such that $\mathrm{eq}(\Delta[\Sigma]) = \theta_\Delta$. Since θ_Δ is a morphism of $\mathbb{M}$-algebras, it follows that $\mathrm{eq}(\Delta[\Sigma])$ is an $\mathbb{M}$-congruence.

($\Rightarrow$) Given $\Phi \subseteq L$, choose $\Phi \subseteq \Delta \subseteq L$ such that $\mathrm{eq}(\Delta[\Sigma])$ is an $\mathbb{M}$-congruence, for all Σ. Set

$$\Theta_\Delta \Sigma := \mathbb{M}\Sigma/\mathrm{eq}(\Delta[\Sigma]) \quad \text{and} \quad (\theta_\Delta)_\Sigma := \mathrm{eq}(\Delta[\Sigma]) : \mathbb{M}\Sigma \to \Theta_\Delta \Sigma \,.$$

Given a function $f : \Sigma \to \Gamma$, we define the morphism $\Theta_\Delta f : \Theta_\Delta \Sigma \to \Theta_\Delta \Gamma$ 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

$$\mathrm{eq}(\Delta[\Sigma]) \leq \ker\left(\mathrm{eq}(\Delta[\Gamma]) \circ \mathbb{M}f\right).$$

Consequently, there exists some function $\psi : \Theta_\Delta \Sigma \to \Theta_\Delta \Gamma$ with

$$\psi \circ \theta_\Delta = \psi \circ \mathrm{eq}(\Delta[\Sigma]) = \mathrm{eq}(\Delta[\Gamma]) = \theta_\Delta \circ \mathbb{M}f \,.$$

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 $f : \Sigma \to \Gamma$ and $g : \Gamma \to \Upsilon$. 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 $\Theta_\Delta(g \circ f) = \Theta_\Delta g \circ \Theta_\Delta f$. □

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

$$\ker \theta_\Delta = \mathrm{eq}(\Delta) \leq \ker \mathrm{Mod}(\varphi)\,, \quad \text{for all } \varphi \in \Delta\,.$$

For every $\varphi \in \Delta$, we can therefore find a function $\mu_\varphi : \Theta_\Delta \Sigma \to \Omega$ such that

$$\mu_\varphi \circ \theta_\Delta = \mathrm{Mod}(\varphi)\,.$$

It is sufficient to prove that $e := \langle \mu_\varphi \rangle_{\varphi \in \Delta} : \Theta_\Delta \Sigma \to \Omega^\Delta$ belongs to M. Note that

$$e \circ \theta_\Delta = \langle \mathrm{Mod}(\varphi) \rangle_{\varphi \in \Delta} = \mathrm{Th}_\Delta$$

implies that

$$\ker(e \circ \theta_\Delta) = \ker \mathrm{Th}_\Delta = \mathrm{eq}(\Delta) = \theta_\Delta\,.$$

Consequently, $e \in M$. □

It follows immediately from the definition that the theory algebra $\Theta_\Delta \Sigma$ recognises every Δ-definable language.

Lemma 2.4. *The morphism $\theta_\Delta : \mathbb{M}\Sigma \to \Theta_\Delta \Sigma$ recognises every Δ-definable language $\kappa : \mathbb{M}\Sigma \to \Omega$.*

Proof. Suppose that $\kappa = \omega[\mathrm{Mod}(\varphi_0), \dots, \mathrm{Mod}(\varphi_{n-1})]$ for some formulae $\varphi_0, \dots, \varphi_{n-1} \in \Delta$ and an Ω-operation $\omega : \Omega^n \to \Omega$. Since

$$\theta_\Delta = \mathrm{eq}(\Delta[\Sigma]) \leq \mathrm{eq}(\{\varphi_i\})\,, \quad \text{for all } i < n\,,$$

we have

$$\begin{aligned}
\theta_\Delta &\leq \inf_{i<n} \mathrm{eq}(\{\varphi_i\}) \\
&= \inf_{i<n} \ker \mathrm{Mod}(\varphi_i) \\
&= \ker \langle \mathrm{Mod}(\varphi_i) \rangle_{i<n} \\
&\leq \ker(\omega \circ \langle \mathrm{Mod}(\varphi_i) \rangle_{i<n}) = \ker \kappa\,.
\end{aligned}$$

Consequently, we can use Lemma III.3.1 (a) to show that θ_Δ recognises κ. □

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 $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ be a language with a syntactic algebra. Then κ is Δ-definable if, and only if,*

$$\mathrm{eq}(\Delta) \leq \ker \mathrm{syn}_\kappa \,.$$

Proof. ($\Rightarrow$) As θ_Δ recognises κ, we can use Theorem III.4.9 to find a function $\rho : \Theta_\Delta\Sigma \to \mathrm{Syn}(\kappa)$ with $\rho \circ \theta_\Delta = \mathrm{syn}_\kappa$. It follows that

$$\mathrm{eq}(\Delta) = \ker \theta_\Delta \leq \ker (\rho \circ \theta_\Delta) = \ker \mathrm{syn}_\kappa \,.$$

($\Leftarrow$) By Lemma III.4.7, we have

$$\mathrm{eq}(\Delta) \leq \ker \mathrm{syn}_\kappa \leq \ker \kappa \leq \kappa \,.$$

Therefore Lemma 1.3 implies that κ is Δ-definable. □

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 $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ be Δ-definable. If the minimal algebra* $\mathrm{Syn}(\kappa)$ *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 $\rho : \Theta_\Delta\Sigma \to \mathrm{Syn}(\kappa)$ with $\rho\circ\theta_\Delta = \mathrm{syn}_\kappa$. 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. □

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 $f : \Sigma \to \Gamma$ 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 $\kappa : \mathbb{M}_\xi\Gamma \to \Omega$. □

Lemma 2.8. *Let L be an $\mathbb{M}$-compositional family of logics, and let $\Delta \subseteq L$ be a subfamily such that* $\mathrm{eq}(\Delta)$ *is an $\mathbb{M}$-congruence. Then*

$$\mathrm{eq}(\Delta) \leq \mathrm{eq}(\Delta) \circ p\,, \quad \textit{for all contexts } p \in \mathbb{M}(\Sigma + \Box)\,.$$

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

$$\mathrm{eq}(\Delta) = \theta_\Delta \leq \ker(\theta_\Delta \circ p) = \ker(\mathrm{eq}(\Delta) \circ p)\,. \qquad \Box$$

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

$$\tilde{\theta}_L := \langle \mathrm{syn}_{\mathrm{Mod}(\varphi)} \rangle_{\varphi \in L[\Sigma]} : \mathbb{M}\Sigma \to \prod_{\varphi \in L[\Sigma]} \mathrm{Syn}(\mathrm{Mod}(\varphi))\,.$$ ⌟

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) $\mathrm{eq}(\Delta) = \ker \tilde{\theta}_\Delta(t)\,.$

(3) $\mathrm{eq}(\Delta) \leq \ker(\mathrm{eq}(\Delta) \circ p)\,, \quad$ *for every context p.*

Proof. (1) ⇔ (3) We have shown in Lemma 1.3 (b) that

$$\mathrm{eq}(\Delta[\Sigma]) \leq \ker(\Delta[\Sigma] \circ p)\,,$$

if, and only if, $p^{-1}[\kappa] : \mathbb{M}_\zeta\Sigma \to \Omega$ is Δ-definable, for every Δ-definable $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$.

(2) $\Leftrightarrow$ (3) First, note that $\mathrm{eq}(\Delta) \circ p = \mathrm{eq}(\Delta)$ for the empty context $p = \mathrm{sing}(\Box)$. Consequently, (3) is equivalent to

$$\mathrm{eq}(\Delta) = \inf \{\, \ker(\mathrm{eq}(\Delta) \circ p) \mid p \text{ a context} \,\}\,.$$

The equivalence of (2) and (3) therefore follows from the fact that

$$\begin{aligned}
\ker \tilde{\theta}_\Delta &= \inf \{\, \ker \mathrm{syn}_{\mathrm{Mod}(\varphi)} \mid \varphi \in \Delta \,\} \\
&= \inf \{\, \ker(\mathrm{Mod}(\varphi) \circ p) \mid \varphi \in \Delta\,,\ p \text{ context} \,\} \\
&= \inf \{\, \ker(\langle \mathrm{Mod}(\varphi) \rangle_{\varphi \in \Delta} \circ p) \mid p \text{ context} \,\} \\
&= \inf \{\, \ker(\mathrm{Th}_\Delta \circ p) \mid p \text{ context} \,\} \\
&= \inf_p \ker(\mathrm{eq}(\Delta) \circ p)\,,
\end{aligned}$$

where the second step follows by Lemma III.4.7. $\Box$

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 $\mathrm{eq}(\Delta) = \ker \tilde{\theta}_\Delta$. In particular, $\mathrm{eq}(\Delta)$ is an $\mathbb{M}$-congruence. $\Box$

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 $\mu : A_\xi \to \Omega$, the composition $\mu \circ \pi \restriction \mathbb{M}_\xi \llbracket C$ 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. ⌟

Remark. Suppose that $\mathcal{D} = \mathsf{Pos}$ and $\Omega = \{0, 1\}$. A subset C of $\mathfrak{A}$ is L-definably embedded if, for every subset $P \subseteq A_\xi$, the preimage

$$\pi^{-1}[\Uparrow P] \cap \mathbb{M}C \text{ is } L\text{-definable.}$$

Furthermore, being weakly finite-dimensional is the same as being finite. ⌟

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 $C \subseteq \mathbb{V}A$ is L-definably embedded in $\mathfrak{A}$ if, and only if, there exists a sort-wise finite set $\Delta \subseteq L[C]$ such that*

$$\mathrm{eq}(\Delta) \restriction \mathbb{M}\mathbb{J}C \leq \ker(\pi \restriction \mathbb{M}\mathbb{J}C)\,.$$

Proof. ($\Leftarrow$) Let Δ be a sort-wise finite set such that

$$\mathrm{eq}(\Delta) \restriction \mathbb{M}\mathbb{J}C \leq \ker(\pi \restriction \mathbb{M}\mathbb{J}C)\,.$$

For $\mu : A_\xi \to \Omega$, it follows that

$$\mathrm{eq}(\Delta) \restriction \mathbb{M}\mathbb{J}C \leq \ker(\pi \restriction \mathbb{M}\mathbb{J}C) \leq \ker(\mu \circ \pi \restriction \mathbb{M}\mathbb{J}C)\,.$$

By Lemma 1.3 (a) this means that $\mu \circ \pi \restriction \mathbb{M}_\xi C$ is L-definable.

($\Rightarrow$) Since A_ξ is weakly finite-dimensional, there exist morphisms $q_\xi : \tilde{A}_\xi \to A$ and $e_\xi : \tilde{A}_\xi \to \Omega^{d_\xi}$, 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 $j : C \to \mathbb{V}A$ be the inclusion, and set

$$\pi_0 := \pi \circ \mathbb{M}(q \circ \iota \circ \mathbb{J}j) : \mathbb{M}\mathbb{J}C \to A$$

be the restriction of the product. Since $\mathbb{MJ}C$ is projective with respect to E-morphisms and $q \in E$, we can find morphisms $\tilde{\pi}_\xi : \mathbb{M}_\xi\mathbb{J}C \to \tilde{A}_\xi$ with

$$q_\xi \circ \tilde{\pi}_\xi = \pi_0 , \quad \text{for all } \xi .$$

Let $p_i : \Omega^{d_\xi} \to \Omega$ 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

$$\mathrm{eq}_\xi(\Delta) \leq \mathrm{eq}_\xi(\Delta_{\xi,i}) \leq \ker\left(p_i \circ e_\xi \circ \tilde{\pi}_\xi\right),$$

where the second inequality follows by Lemma 1.3 (a). Hence,

$$\begin{aligned} \mathrm{eq}_\xi(\Delta) &\leq \inf_i \ker\left(p_i \circ e_\xi \circ \tilde{\pi}_\xi\right) \\ &= \ker\left(e_\xi \circ \tilde{\pi}_\xi\right) \\ &= \ker \tilde{\pi}_\xi \\ &\leq \ker\left(q_\xi \circ \tilde{\pi}_\xi\right) = \ker(\pi_0)_\xi , \end{aligned}$$

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 $C \subseteq \mathbb{V}A$ 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 $i : D \to C$ and $j : \mathbb{J}C \to A$ be the respective inclusion maps. For $\mu : A_\xi \to \Omega$ it follows that

$$\mu \circ \pi \restriction \mathbb{M}_\xi D = \mu \circ \pi \circ \mathbb{M}j \circ \mathbb{MJ}i = (\mu \circ \pi \circ \mathbb{M}j) \circ \mathbb{MJ}i .$$

By assumption, $\mu \circ \pi \circ \mathbb{M}j$ is L-definable. Consequently, Lemma 2.7 implies that so is $\mu \circ \pi \circ \mathbb{M}j \circ \mathbb{M}\mathbb{J}i$.

(b) By (a) it is sufficient to consider the case where $D = \mathbb{V}\langle\!\langle C \rangle\!\rangle_{\mathfrak{A}}$. Let $i : C \to D$ and $j : \mathbb{J}D \to A$ 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 $f : D \to \mathbb{V}\mathbb{M}\mathbb{J}C$ with

$$\mathbb{V}(\pi \circ \mathbb{M}(j \circ \mathbb{J}i)) \circ f = \mathrm{id}_D\,.$$

Via the adjunction, we obtain a function $f^* : \mathbb{J}D \to \mathbb{M}\mathbb{J}C$. Let $\varphi : \mathbb{M}\mathbb{J}D \to \mathbb{M}\mathbb{J}C$ be the (unique) extension of f^* to a morphism of $\mathbb{M}$-algebras. Then

$$\begin{aligned} \pi \circ \mathbb{M}j \circ \mathbb{M}\mathbb{J}i \circ \varphi \circ \mathrm{sing} &= \pi \circ \mathbb{M}j \circ \mathbb{M}\mathbb{J}i \circ f \\ &= \mathrm{id} \\ &= \pi \circ \mathbb{M}j \circ \mathrm{sing}\,. \end{aligned}$$

Since morphisms of $\mathbb{M}$-algebras are uniquely determined by their restriction to $\mathrm{rng}\,\mathrm{sing}$, it follows that

$$\pi \circ \mathbb{M}j \circ \mathbb{M}\mathbb{J}i \circ \varphi = \pi \circ \mathbb{M}j\,.$$

To show that D is L-definably embedded, consider a function $\mu : A_\xi \to \Omega$. We have to show that $\mu \circ \pi \circ \mathbb{M}j : \mathbb{M}\mathbb{J}D \to \Omega$ is L-definable. As C is L-definably embedded, we know that $\mu \circ \pi \circ \mathbb{M}(j \circ \mathbb{J}i) : \mathbb{M}_\xi\mathbb{J}C \to \Omega$ 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. □

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 $\mu : A_\xi \to \Omega$ such that the composition $\kappa := \mu \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C$ is not L-definable. Thus, the restriction $\pi \upharpoonright \mathbb{M}\mathbb{J}C : \mathbb{M}\mathbb{J}C \to \mathfrak{A}$ of the product is a morphism recognising the non-L-definable language κ.

($\Rightarrow$) Let $\varphi : \mathbb{M}\Sigma \to \mathfrak{A}$ be a morphism and $\mu : A_\xi \to \Omega$ a function. We have to show that $\kappa := \mu \circ \varphi$ is L-definable. By assumption, the set $C := \operatorname{rng}(\varphi \circ \operatorname{sing})$ is L-definably embedded in $\mathfrak{A}$. By definition, this implies that the function $\mu \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C : \mathbb{M}\mathbb{J}C \to \Omega$ is L-definable. As Σ is of the form $\mathbb{J}X$, we can use Corollary II.6.4 to find a function $\hat{\varphi} : \mathbb{M}\Sigma \to \mathbb{M}\mathbb{J}C$ satisfying $\varphi = \pi \circ \hat{\varphi}$. Hence,

$$\kappa = \mu \circ \varphi = \mu \circ \pi \circ \hat{\varphi}\,.$$

Since $\mu \circ \pi \upharpoonright \mathbb{M}\mathbb{J}C$ 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. □

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 $\pi_A : \mathbb{M}A \to A$ and $\pi_B : \mathbb{M}B \to B$ be the respective products. Given a finite set $C \subseteq \mathbb{V}A$ and a function $\mu : A_\xi \to \Omega$, consider the language $\kappa := \mu \circ \pi_A \upharpoonright \mathbb{M}\mathbb{J}C$. Since $\pi_A \upharpoonright \mathbb{M}\mathbb{J}C = \pi_B \upharpoonright \mathbb{M}\mathbb{J}C$, we have $\kappa = \mu \circ \pi_B \upharpoonright \mathbb{M}\mathbb{J}C$ 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 $\mu : A_\xi \to \Omega$, 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 $\omega : \Omega^\circ \to \Omega : \langle\rangle \mapsto \mu(1_\xi)$. 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 $p_k : \prod_i B^i \to B^k$ 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

$$\mathrm{eq}(\Delta_k) \restriction \mathbb{MJ}D^k \leq \ker \pi \restriction \mathbb{MJ}D^k .$$

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

$$\mathrm{eq}(\Delta'_k) \leq \ker\left(\mathrm{eq}(\Delta_k) \circ \mathbb{M}p_k\right).$$

Setting $\Delta := \Delta'_0 \cup \dots \cup \Delta'_{n-1}$, it follows that

$$\begin{aligned} \mathrm{eq}(\Delta) \restriction \mathbb{MJ}C &\leq \ker\left(\mathrm{eq}(\Delta_k) \circ \mathbb{M}p_k\right) \restriction \mathbb{MJ}C \\ &\leq \ker\left(\pi \circ \mathbb{M}p_k\right) \restriction \mathbb{MJ}C \\ &= \ker\left(p_k \circ \pi\right) \restriction \mathbb{MJ}C . \end{aligned}$$

which implies that

$$\begin{aligned} \mathrm{eq}(\Delta) \restriction \mathbb{MJ}C &\leq \inf_k \ker\left(p_k \circ \pi\right) \restriction \mathbb{MJ}C \\ &= \ker\left(\langle p_k\rangle_k \circ \pi\right) \restriction \mathbb{MJ}C = \ker \pi \restriction \mathbb{MJ}C \end{aligned}$$

as desired. □

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 $q : \mathfrak{B}|_{\{\xi\}} \to \mathfrak{A}|_{\{\xi\}}$. In particular, there exists an E-morphism $B_\xi \to A_\xi$. 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 $\mu : A_\xi \to \Omega$. 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 $q : \mathfrak{B}|_\Delta \to \mathfrak{A}|_\Delta$ be an E-morphism. Then $\mathbb{V}q$ is surjective and has a right-inverse. Via the adjunction we obtain a function $f : \mathbb{J}\mathbb{V}A|_\Delta \to B|_\Delta$ such that $q \circ f = \iota$, where $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ is the counit of the adjunction. Setting $D := f[\mathbb{J}C]$, it follows that

$$\begin{aligned}
\mu \circ \pi \restriction \mathbb{M}\mathbb{J}C &= \mu \circ \pi \circ \mathbb{M}\iota \restriction \mathbb{M}\mathbb{J}C \\
&= \mu \circ \pi \circ \mathbb{M}q \circ \mathbb{M}f \restriction \mathbb{M}\mathbb{J}C \\
&= \mu \circ q \circ \pi \circ \mathbb{M}f \restriction \mathbb{M}\mathbb{J}C\,.
\end{aligned}$$

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, $\mu \circ \pi \restriction \mathbb{M}\mathbb{J}C$ is also L-definable. □

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 $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ is an L-definable language with a syntactic algebra, then* $\mathrm{Syn}(\kappa)$ *is locally L-definable.*

Proof. Let $C \subseteq \mathbb{V}\mathrm{Syn}(\kappa)$ be finite and fix $\mu : \mathrm{Syn}_\xi(\kappa) \to \Omega$. As $\lambda := \mu \circ \pi \restriction \mathbb{MJ}C : \mathbb{MJ}C \to \Omega$ is recognised by the restriction $\pi \restriction \mathbb{MJ}C : \mathbb{MJ}C \to \mathrm{Syn}(\kappa)$, it follows by Proposition III.4.16 that λ is of the form

$$\lambda = \varphi^{-1}[\omega[p_0^{-1}[\kappa], \dots, p_{n-1}^{-1}[\kappa]]],$$

for some morphism $\varphi : \mathbb{M}C \to \mathbb{M}\Sigma$, contexts $p_0, \dots, p_{n-1}$, and an Ω-operation $\omega : \Omega^n \to \Omega$. 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 $\mathrm{Syn}(\kappa)$. □

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 $\kappa : \mathbb{M}\Sigma \to \Omega$, the syntactic algebra* $\mathrm{Syn}(\kappa)$ *is locally L-definable.*

Proof. (1) ⇒ (2) follows by Lemma 4.1. For (2) ⇒ (1), fix an L-definable language $\kappa : \mathbb{M}_\xi\Gamma \to \Omega$. Then $\kappa = \mu \circ \mathrm{syn}_\kappa$, for some $\mu : \mathrm{Syn}_\xi(\kappa) \to \Omega$. For closure under inverse morphisms, consider $\varphi : \mathbb{M}\Sigma \to \mathbb{M}\Gamma$. Then

$$\kappa \circ \varphi = \mu \circ \mathrm{syn}_\kappa \circ \varphi,$$

is recognised by $\mathrm{syn}_\kappa \circ \varphi : \mathbb{M}\Sigma \to \mathrm{Syn}(\kappa)$. 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 $\mathrm{syn}_\kappa : \mathbb{M}\Gamma \to \mathrm{Syn}(\kappa)$. Hence, L-definability follows again by Theorem 3.4. □

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* $\mathrm{eq}(\Delta)$ *is an $\mathbb{M}$-congruence, the set* $\mathrm{rng}\,\mathbb{V}(\theta_\Delta \circ \mathrm{sing})$ *is L-definably embedded in $\Theta_\Delta\Sigma$.*

Proof. Let $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{V}\mathbb{J}$ and $\iota : \mathbb{J}\mathbb{V} \Rightarrow \mathrm{Id}$ be the unit and the counit of the adjunction. Suppose that $\Sigma = \mathbb{J}X$ and let

$$f := \mathbb{V}(\theta_\Delta \circ \mathrm{sing}) \circ \varepsilon : X \to \mathbb{V}\Theta_\Delta\Sigma$$

be the morphism corresponding to $\theta_\Delta \circ \mathrm{sing} : \Sigma \to \Theta_\Delta\Sigma$ via the adjunction. Let $C := \mathrm{rng}\, f$ and choose a right-inverse $g : C \to X$ of $f : X \to C$. Finally, let $i : C \to \mathbb{V}\Theta_\Delta\Sigma$ be the inclusion and $\pi_\circ := \pi \circ \mathbb{M}(\iota \circ \mathbb{J}i) : \mathbb{M}\mathbb{J}C \to \Theta_\Delta\Sigma$ the restriction of the product to $\mathbb{M}\mathbb{J}C$. Then

$$\begin{aligned}
\theta_\Delta \circ \mathbb{M}\mathbb{J}g &= \theta_\Delta \circ (\mathrm{flat} \circ \mathbb{M}\mathrm{sing}) \circ \mathbb{M}\mathbb{J}g \\
&= \pi \circ \mathbb{M}\theta_\Delta \circ \mathbb{M}\mathrm{sing} \circ \mathbb{M}(\iota \circ \mathbb{J}\varepsilon \circ \mathbb{J}g) \\
&= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}\mathbb{V}(\theta_\Delta \circ \mathrm{sing}) \circ \mathbb{J}(\varepsilon \circ g)) \\
&= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}(f \circ g)) \\
&= \pi \circ \mathbb{M}(\iota \circ \mathbb{J}i) \\
&= \pi_\circ\,.
\end{aligned}$$

To show that C is L-definably embedded in $\Theta_\Delta\Sigma$, let $\mu : \Theta_\Delta\Sigma \to \Omega$. By Lemma 1.3,

$$\mathrm{eq}(\Delta) = \ker \theta_\Delta \leq \ker(\mu \circ \theta_\Delta)$$

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

$$\mu \circ \pi_\circ = \mu \circ \theta_\Delta \circ \mathbb{M}\mathbb{J}g = \kappa \circ \mathbb{M}\mathbb{J}g\,.$$

is also L-definable. □

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 $\kappa : \mathbb{M}_\xi \Sigma \to \Omega$, the syntactic algebra $\mathrm{Syn}(\kappa)$ is L-definable.*

(5) *For every L-definable language $\kappa : \mathbb{M}_\xi \Sigma \to \Omega$, the syntactic algebra $\mathrm{Syn}(\kappa)$ 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 $\mathrm{Syn}(\kappa)$ is finitely generated and weakly finite-dimensional. For the former, note that $\mathrm{Syn}(\kappa)$ is generated by the finite set $\mathrm{rng}\,(\mathrm{syn}_\kappa \circ \mathrm{sing})$. 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 $\rho : \Theta_\Delta \Sigma \to \mathrm{Syn}(\kappa)$. Since $\Theta_\Delta \Sigma$ is strongly finite-dimensional, it follows that $\mathrm{Syn}(\kappa)$ is weakly finite-dimensional.

(2) $\Rightarrow$ (3) First, note that $\Theta_\Delta \Sigma$ is generated by the finite set

$$C := \mathrm{rng}\,\mathbb{V}(\theta_\Delta \circ \mathrm{sing})\,.$$

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 $\varphi : \mathbb{M}\Sigma \to \mathbb{M}\Gamma$ and an L-definable language $\kappa : \mathbb{M}_\xi \Gamma \to \Omega$. We will construct two sort-wise finite sets $\Delta, \Delta' \subseteq L$ such that κ is $\Delta[\Gamma]$-definable and

$$\mathrm{eq}(\Delta'[\Sigma]) \leq \ker\,(\mathrm{eq}(\Delta[\Gamma]) \circ \varphi)\,.$$

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 $\mathrm{eq}(\Delta)$ is an $\mathbb{M}$-congruence. Set

$$f := \theta_\Delta \circ \varphi \circ \mathrm{sing} : \Sigma \to \Theta_\Delta \Gamma \quad \text{and} \quad C := \mathrm{rng}\, \mathbb{V} f\,.$$

Then we have

$$\begin{aligned} \pi \circ \mathbb{M} f &= \pi \circ \mathbb{M}(\theta_\Delta \circ \varphi \circ \mathrm{sing}) \\ &= \theta_\Delta \circ \varphi \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} = \theta_\Delta \circ \varphi\,. \end{aligned}$$

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

$$\mathrm{eq}(\Psi) \restriction \mathbb{M}\mathbb{J}C \leq \ker\,(\pi \restriction \mathbb{M}\mathbb{J}C)\,.$$

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

$$\mathrm{eq}(\Delta'[\Sigma]) \leq \ker\,(\mathrm{eq}(\Psi) \circ \mathbb{M} f)\,.$$

Consequently,

$$\begin{aligned} \mathrm{eq}_\xi(\Delta'[\Sigma]) &\leq \ker\,(\mathrm{eq}(\Psi) \circ \mathbb{M} f) \\ &\leq \ker\,(\pi \circ \mathbb{M} f) \\ &= \ker\,(\theta_\Delta \circ \varphi) = \ker\,(\mathrm{eq}(\Delta[\Gamma]) \circ \varphi)\,. \end{aligned}$$ □

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 $\kappa : \mathbb{M}_\xi \Sigma \to \Omega$ 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) ⇒ (2) is trivial and (2) ⇒ (1) has already been proved in Theorem 3.4. For (1) ⇒ (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 $\theta_\Delta : \mathbb{M}\Sigma \to \Theta_\Delta \Sigma$ recognises κ. □

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. (⇒) follows by Corollary 4.5.

(⇐) Let $\Phi \subseteq L$ be finite. For every formula $\varphi \in \Phi$, we fix a morphism $\eta_\varphi : \mathbb{M}\Sigma \to \mathfrak{A}_\varphi$ to a sort-wise strongly finite-dimensional L-definable algebra $\mathfrak{A}_\varphi$ that recognises the language $\mathrm{Mod}(\varphi)$. 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 $e_\xi : A_\xi \to \Omega^{d_\xi}$ 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 $p_i : \Omega^{d_\xi} \to \Omega$ is the projection to the i-th component. Since L is closed under inverse relabellings, there exist formulae $\psi_{\xi,i}$ such that

$$s \vDash \psi_{\xi,i} \quad \text{iff} \quad \mathbb{M}(\eta \circ \mathrm{sing})(s) \vDash \psi'_{\xi,i}\,.$$

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 $\mu_\varphi : A_\xi \to \Omega$, for $\varphi \in \Phi$, such that

$$\mathrm{Mod}(\varphi) = \mu_\varphi \circ \eta\,.$$

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 $\rho = \big\langle \langle \mu_\varphi \rangle_\varphi, \langle \mu_{\psi_{\xi,i}} \rangle_{\xi,i} \big\rangle \in M$. Consequently,

$$\begin{aligned} \mathrm{eq}(\Delta) &= \ker \mathrm{Th}_\Delta \\ &= \ker \langle \mathrm{Mod}(\varphi) \rangle_{\varphi\in\Delta} \\ &= \ker \langle \mu_\varphi \circ \eta \rangle_{\varphi\in\Delta} \\ &= \ker \big(\langle \mu_\varphi \rangle_{\varphi\in\Delta} \circ \eta \big) = \ker \eta\,, \end{aligned}$$

which implies that $\mathrm{eq}(\Delta)$ is an $\mathbb{M}$-congruence. □

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 $\mathrm{eq}(\Delta)$ 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 $X \subseteq \Xi$ 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 $e_\xi : A_\xi \to \Omega^{d_\xi}$, for some $d_\xi < \omega$. Let $p_i : \Omega^{d_\xi} \to \Omega$ 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 $p_i \circ e_\xi \circ \pi \restriction \mathbb{MJ}C$ with $\xi \in X$ and $i < d_\xi$ is Δ-definable. Then

$$\ker \theta_\Delta \restriction \mathbb{MJ}C = \mathrm{eq}(\Delta) \restriction \mathbb{MJ}C \leq \ker (p_i \circ e_\xi \circ \pi) \restriction \mathbb{MJ}C$$

implies that

$$\begin{aligned} \ker \theta_\Delta \restriction \mathbb{MJ}C &\leq \inf_{\xi,i} \ker (p_i \circ e_\xi \circ \pi) \restriction \mathbb{MJ}C \\ &= \inf_{\xi} \ker (e_\xi \circ \pi) \restriction \mathbb{MJ}C \\ &= \inf_{\xi} \ker \pi \restriction \mathbb{MJ}C \\ &= \ker \pi \restriction \mathbb{MJ}C\,, \end{aligned}$$

where the second step follows since $\langle p_i \rangle_i = \mathrm{id}$ and the third one since $e_\xi \in M$. Consequently, $\ker \theta_\Delta|_X \leq \ker \pi|_X$ and there exist a morphism $q : \Theta_\Delta \mathbb{J}C|_X \to \mathfrak{A}|_X$ such that $\pi = q \circ \theta_\Delta$. In particular, $\mathfrak{A}|_X$ is a quotient of $\Theta_\Delta \mathbb{J}C|_X$ and $\Theta_\Delta \mathbb{J}C \in \mathcal{V}$. □

Corollary 4.8. *Suppose that $\mathcal{D} = \mathsf{Set}$ or $\mathcal{D} = \mathsf{Pos}$, 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 $\mathrm{Th}(\mathcal{S})$.*

Proof. We distinguish three cases. First assume that $\mathcal{D} = \mathsf{Pos}$ 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 $\mathcal{W} := \mathrm{Mod}(\mathrm{Th}(\mathcal{S}))$ is also a $\mathcal{R}$-variety containing $\mathcal{S}$. Consequently, $\mathcal{V} \subseteq \mathcal{W}$. Furthermore, $\mathcal{S} \subseteq \mathcal{V}$ implies $\mathrm{Th}(\mathcal{S}) \supseteq \mathrm{Th}(\mathcal{V})$. Hence, it follows by Corollary III.7.5 that

$$\mathcal{W} = \mathrm{Mod}(\mathrm{Th}(\mathcal{S})) \subseteq \mathrm{Mod}(\mathrm{Th}(\mathcal{V})) = \mathcal{V}\,.$$

If $\mathcal{D} = \mathsf{Set}$, we can use the same proof. It therefore remains to consider the case where $\mathcal{D} = \mathsf{Pos}$ and Ω is an antichain. Since $\mathbb{M}$ is polynomial, it follows that every set of the form $\mathbb{M}\Sigma$, for $\Sigma \in \mathsf{Alph}$, is also an antichain. Furthermore, every algebra $\mathfrak{A}$ recognising some language $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ must be an antichain as, otherwise, there can be no functions $A \to \Omega$. Consequently, the problem reduces to the case where $\mathcal{D} = \mathsf{Set}$. □

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 $\kappa : \mathbb{M}_\xi\Sigma \to \Omega$ 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) $\mathrm{Syn}(\kappa)$ *is L-definable.*

(4) $\mathrm{Syn}(\kappa)$ *is a quotient of $\Theta_\Delta\Gamma$, for some Δ and Γ.*

(5) $\mathrm{syn}_\kappa = \rho \circ \theta_\Delta$, *for some Δ and a surjective morphism $\rho : \Theta_\Delta\Sigma \to \mathrm{Syn}(\kappa)$.*

(6) *κ is recognised by $\Theta_\Delta\Gamma$, for some Δ and Γ.*

(7) *$\theta_\Delta : \mathbb{M}\Sigma \to \Theta_\Delta\Sigma$ recognises κ, for some Δ.*

(8) $\mathrm{eq}(\Delta) \leq \ker \kappa$, *for some* Δ.

(9) $\mathrm{eq}(\Delta) \leq \ker \mathrm{syn}_\kappa$, *for some* Δ.

If $\mathcal{D} = \mathsf{Set}$ *or* $\mathcal{D} = \mathsf{Pos}$, *and if* $\Omega \in \mathcal{D}$ *is finite, the above statements are equivalent to the following one.*

(10) $\mathrm{Syn}(\kappa)$ *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) ⇔ (2) has been proved in Corollary 4.5; (1) ⇔ (8) in Lemma 1.3; (3) ⇔ (10) follows by Corollary 4.8.

(1) ⇒ (7) was proved in Lemma 2.4.

(7) ⇒ (8) Suppose that $\kappa = \mu \circ \theta_\Delta$ for some function $\mu : \Theta_\Delta \Sigma \to \Omega$. Then

$$\mathrm{eq}(\Delta) = \ker \theta_\Delta \leq \ker (\mu \circ \theta_\Delta) = \ker \kappa\,.$$

(1) ⇒ (9) follows by Corollary 2.5.

(9) ⇔ (5) Note that $\mathrm{eq}(\Delta) = \ker \theta_\Delta$. For a finite set $\Delta \subseteq L$, it therefore follows that

$$\begin{aligned} \mathrm{eq}(\Delta) \leq \ker \mathrm{syn}_\kappa \quad &\text{iff} \quad \ker \theta_\Delta \leq \ker \mathrm{syn}_\kappa \\ &\text{iff} \quad \mathrm{syn}_\kappa = \rho \circ \theta_\Delta\,, \quad \text{for some morphism } \rho\,, \end{aligned}$$

where the fact that ρ is a morphism of $\mathbb{M}$-algebras follows by Lemma I.5.6.

(5) ⇒ (4) is trivial.

(4) ⇒ (6) Since $\mathrm{syn}_\kappa : \mathbb{M}\Sigma \to \mathrm{Syn}(\kappa)$ recognises κ, the claim follows by Lemma III.5.7.

(6) ⇒ (2) follows by Theorem 4.4 (3).

(1) ⇒ (3) follows by Theorem 4.4 (4).

(3) ⇒ (2) holds as $\mathrm{syn}_\kappa : \mathbb{M}\Sigma \to \mathrm{Syn}(\kappa)$ recognises κ. □

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 $\Phi \subseteq \Delta \subseteq L$ with the following property. For every $\varphi \in \Delta$, there exists a finite subset $\Delta_0 \subseteq \Delta$ and a formula $\psi \in L[\Omega]$ satisfying

$$\mathrm{Mod}(\varphi) \circ \mathrm{flat} = \mathrm{Mod}(\psi) \circ \mathbb{M}\mathrm{Th}_{\Delta_0} .$$ ⌟

Remark. Note that $\mathbb{M}$-compositionality implies that

$$\mathbb{M}\mathrm{Th}_{\Delta}(s) \quad \text{determines} \quad \mathrm{Th}_{\Delta}(\mathrm{flat}(s)) .$$

Strong $\mathbb{M}$-compositionality strenghens this implication by requiring that the theory of $\mathrm{flat}(s)$ can be computed by an L-formmula. ⌟

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 $\mathrm{eq}(\Delta)$ is an $\mathbb{M}$-congruence, i.e., that

$$\ker \mathbb{M}\mathrm{eq}(\Delta) \leq \ker\,(\mathrm{eq}(\Delta) \circ \mathrm{flat}) .$$

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

$$\mathrm{Mod}(\varphi) \circ \mathrm{flat} = \mathrm{Mod}(\psi_\varphi) \circ \mathbb{M}\mathrm{Th}_{\Delta_\varphi} .$$

It follows that

$$\mathrm{eq}(\Delta) \circ \mathrm{flat} = \big\langle \mathrm{Mod}(\varphi) \big\rangle_{\varphi \in \Delta} \circ \mathrm{flat} = \big\langle \mathrm{Mod}(\psi_\varphi) \circ \mathbb{M}\mathrm{Th}_{\Delta_\varphi} \big\rangle_{\varphi \in \Delta} .$$

Since $\Delta_\varphi \subseteq \Delta$, there are functions p_φ with $\mathrm{Th}_{\Delta_\varphi} = p_\varphi \circ \mathrm{eq}(\Delta)$. Consequently,

$$\begin{aligned} \mathrm{eq}(\Delta) \circ \mathrm{flat} &= \big\langle \mathrm{Mod}(\psi_\varphi) \circ \mathbb{M}(p_\varphi \circ \mathrm{eq}(\Delta)) \big\rangle_{\varphi \in \Delta} \\ &= \big\langle \mathrm{Mod}(\psi_\varphi) \circ \mathbb{M}p_\varphi \big\rangle_{\varphi \in \Delta} \circ \mathbb{M}\mathrm{eq}(\Delta) , \end{aligned}$$

which implies that $\ker \mathbb{M}\mathrm{eq}(\Delta) \leq \ker(\mathrm{eq}(\Delta) \circ \mathrm{flat})$.

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 $C := \mathrm{rng}(\theta_\Delta \circ \mathrm{sing})$. Let $i : C \to \Theta_\Delta\Sigma$ be the inclusion map and let $q : \Sigma \to C$ be the function such that $i \circ q = \theta_\Delta \circ \mathrm{sing}$. We have to show that, for every function $\mu : \Theta_\Delta\Sigma \to \Omega$, the composition $\mu \circ \pi \circ \mathbb{M}i : \mathbb{M}C \to \Omega$ is L-definable.

Hence, fix $\mu : \Theta_\Delta\Sigma \to \Omega$. Then $\mu \circ \theta_\Delta : \mathbb{M}\Sigma \to \Omega$ is Δ-definable. Let $\varphi \in \Delta[\Omega]$ be a formula with $\mathrm{Mod}(\varphi) = \mu \circ \theta_\Delta$. By strong $\mathbb{M}$-compositionality, there is some formula ψ and a set $\Delta_0 \subseteq \Delta$ such that

$$\mathrm{Mod}(\varphi) \circ \mathrm{flat} = \mathrm{Mod}(\psi) \circ \mathbb{M}\mathrm{Th}_{\Delta_0} .$$

Since $\Delta_0 \subseteq \Delta$, we can find a function p with $\mathrm{Th}_{\Delta_0} = p \circ \theta_\Delta$. It follows that

$$\begin{aligned}(\mu \circ \pi \circ \mathbb{M}i) \circ \mathbb{M}q &= \mu \circ \pi \circ \mathbb{M}(\theta_\Delta \circ \mathrm{sing}) \\ &= \mu \circ \theta_\Delta \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} \\ &= \mathrm{Mod}(\varphi) \circ \mathrm{flat} \circ \mathbb{M}\mathrm{sing} \\ &= \mathrm{Mod}(\psi) \circ \mathbb{M}\mathrm{Th}_{\Delta_0} \circ \mathbb{M}\mathrm{sing} \\ &= \mathrm{Mod}(\psi) \circ \mathbb{M}(p \circ \theta_\Delta) \circ \mathbb{M}\mathrm{sing} \\ &= \mathrm{Mod}(\psi) \circ \mathbb{M}(p \circ i \circ q) .\end{aligned}$$

Since q (and hence $\mathbb{M}q$) is surjective, it follows that

$$\mu \circ \pi \circ \mathbb{M}i = \mathrm{Mod}(\psi) \circ \mathbb{M}(p \circ i) .$$

Thus $\mu \circ \pi \circ \mathbb{M}i$ is an inverse relabelling of an L-definable language and, therefore, also L-definable. □

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., [3]), the notion of an *abstract elementary class* introduced by Shelah (see, e.g., [2]), and the theory of *institutions* developed by Goguen and Burstall (see, e.g., [22]). 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.

Part C.

Applications

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* $\mathcal{D} = \mathsf{Pos}$.

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 $\varnothing$ 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 $\mathrm{dom}_+(g)$ and by $\mathrm{dom}(g)$ 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 $\mathbb{T}_\xi A \subseteq \mathbb{T}^\times_\xi A$ 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 $\mathbb{F}_\xi A \subseteq \mathbb{F}^\times_\xi A$ 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.

(d) Finally, we denote by $\mathbb{G}^{\mathrm{fin}}A$, $\mathbb{R}^{\mathrm{fin}}A$, $\mathbb{F}^{\mathrm{fin}}A$, $\mathbb{F}^{\times\mathrm{fin}}A$, $\mathbb{T}^{\mathrm{fin}}A$, $\mathbb{T}^{\times\mathrm{fin}}A$, the set of all *finite* graphs, forests, or trees in $\mathbb{G}A$, $\mathbb{R}A$,.... ⌟

Remark. (a) To denote finite trees we will frequently use term notation. For instance, $a(x, b(y))$ is the tree

$$\begin{tikzcd} & a \arrow[dl, "x_0"'] \arrow[dr, "x_1"] & \\ x & & b \arrow[d, "x_0"] \\ & & y \end{tikzcd}$$

(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 0 for the empty forest. For instance, $a(x + c, 0, b(y)) + b(x)$ denotes the forest

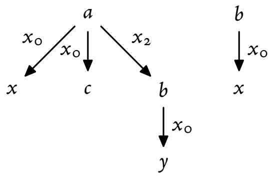

(b) For unranked forests, the edge labellings are usually only necessary when we consider forests in $s \in \mathbb{FF}A$ where the component forests $s(v)$ might contain several different variables. For forests $s \in \mathbb{F}\Sigma$ labelled by some alphabet Σ on the other hand, we could do without 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 implicitely 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 $\operatorname{gun}(g)$ 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 $\operatorname{gun}(g)$ are all the paths of length 0 (see Figure 1 for an example).

(b) For $a \in A_\xi$, we denote by $\operatorname{sing}(a)$ 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, $\operatorname{sing}(a) = a(x_0, \dots, x_{n-1})$ for $\xi = \{x_0, \dots, x_{n-1}\}$.

(c) The *flattening* $\operatorname{flat}(g)$ of a graph $g \in \mathbb{GG}A$ is defined as follows (see Figure 2 for an example). Let h be the disjoint union of all graphs $g(v)$, for $v \in \operatorname{dom}_+(g)$, and let R be the binary relation on $\operatorname{dom}_+(h)$ 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 $w \in \mathrm{dom}(g)$,
- v is a root of some component $g(w')$ where w' is the x-successor of v in g.

Then $\mathrm{flat}(g)$ is the graph obtained from h by

- deleting every vertex corresponding to a hole in $g(v)$ with $v \in \mathrm{dom}(g)$ 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 $\mathrm{flat} : \mathbb{G}\mathbb{G}A \to \mathbb{G}A$ and $\mathrm{sing} : A \to \mathbb{G}A$ 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: $\mathrm{flat}^{\mathbb{G}}$, $\mathrm{flat}^{\mathbb{R}}$, $\mathrm{flat}^{\mathbb{T}}$, etc.

(e) Finally, for the functors $\mathbb{T}^{\times}$ and $\mathbb{F}^{\times}$ we set $\mathrm{flat}^{\times} := \mathrm{gun} \circ \mathrm{flat}$ (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 follows from Lemma I.6.2 (a) by considering the corresponding inclusion morphisms. Clearly, we have $\mathrm{flat} \circ \mathrm{sing} = \mathrm{id}$ and $\mathrm{flat} \circ \mathbb{G}\mathrm{sing} = \mathrm{id}$. For associativity, let $G \in \mathbb{G}\mathbb{G}\mathbb{G}A$. By Proposition I.3.7 and the Remark on

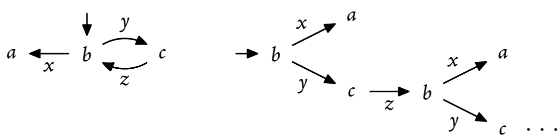

Figure 1.: A graph and its unravelling.

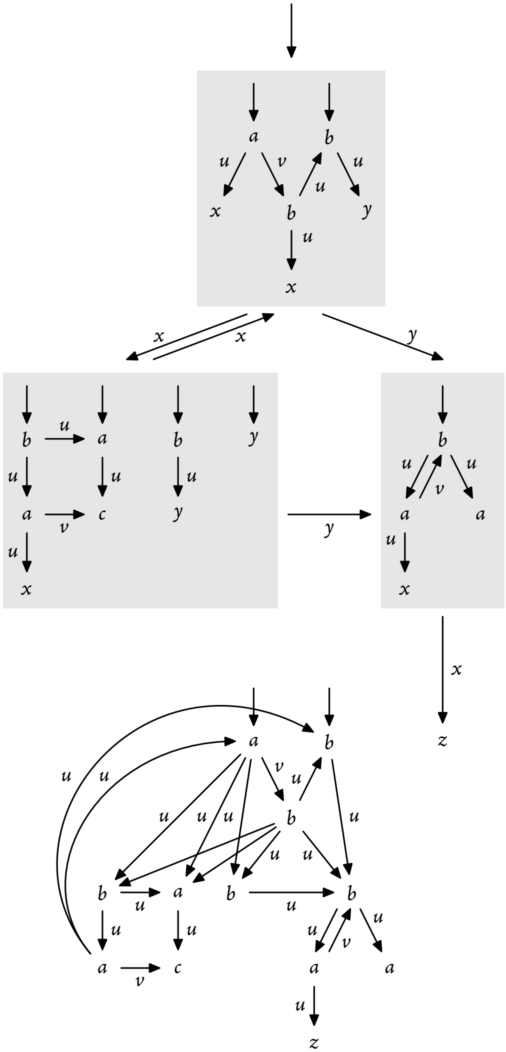

Figure 2.: A graph $g \in \mathbb{G}\mathbb{G}\{a, b, c\}$ and its flattening. The sorts of a, b, c are $\{u, v\}$, $\{u\}$, and $\varnothing$, respectively. The roots are marked by an incoming arrow.

page 416, we obtain labelling-preserving functions

$$\begin{aligned}
\mu &: \mathrm{dom}(\mathrm{flat}(G)) \to \sum_{v \in \mathrm{dom}(G)} \mathrm{dom}(G(v))\,, \\
\mu' &: \mathrm{dom}(\mathrm{flat}(\mathrm{flat}(G))) \to \sum_{w \in \mathrm{dom}(\mathrm{flat}(G))} \mathrm{dom}(\mathrm{flat}(G)(w))\,, \\
\mu_v &: \mathrm{dom}(\mathrm{flat}(G(v))) \to \sum_{u \in \mathrm{dom}(G(v))} \mathrm{dom}(G(v)(u))\,, \\
\mu'' &: \mathrm{dom}(\mathrm{flat}(\mathbb{G}\mathrm{flat}(G))) \to \sum_{v \in \mathrm{dom}(\mathbb{G}\mathrm{flat}(G))} \mathrm{dom}(\mathrm{flat}(G(v)))\,,
\end{aligned}$$

for $v \in \mathrm{dom}(G)$. As the flattening function for graphs does not erase or duplicated vertices (this is different or $\mathbb{T}^\times$ 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

$$\begin{aligned}
\mathrm{dom}(\mathrm{flat}(\mathrm{flat}(G)) &\to \sum_{w \in \mathrm{dom}(\mathrm{flat}(G))} \mathrm{dom}(\mathrm{flat}(G)(w)) \\
&\to \sum_{\langle v,u\rangle \in \sum_{v \in \mathrm{dom}(G)} \mathrm{dom}(G(v))} \mathrm{dom}(G(v)(u)) \\
&\to \sum_{v \in \mathrm{dom}(G)} \sum_{u \in \mathrm{dom}(G(u)} \mathrm{dom}(G(v)(u)) \\
&\to \sum_{v \in \mathrm{dom}(\mathbb{G}\mathrm{flat}(G))} \sum_{u \in \mathrm{dom}(G(u)} \mathrm{dom}(G(v)(u)) \\
&\to \sum_{v \in \mathrm{dom}(\mathbb{G}\mathrm{flat}(G))} \mathrm{dom}(\mathrm{flat}(G(v))) \\
&\to \mathrm{dom}(\mathrm{flat}(\mathbb{G}\mathrm{flat}(G)))
\end{aligned}$$

that preserves the labelling and the edge relation. This implies that the graphs $\mathrm{flat}(\mathrm{flat}(G))$ and $\mathrm{flat}(\mathbb{G}\mathrm{flat}(G))$ are in fact equal. □

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. $\langle \mathbb{T}^\times, \mathrm{flat}^\times, \mathrm{sing}\rangle$ *and* $\langle \mathbb{F}^\times, \mathrm{flat}^\times, \mathrm{sing}\rangle$ *form monads and the unravelling maps* $\mathrm{gun} : \mathbb{R} \Rightarrow \mathbb{T}^\times$ *and* $\mathrm{gun} : \mathbb{G} \Rightarrow \mathbb{F}^\times$ *are morphisms of monads.*

Proof. By Lemma I.6.2, it is sufficient to check that

$$\mathrm{sing} = \mathrm{gun} \circ \mathrm{sing} \quad \text{and} \quad \mathrm{flat}^\times \circ \mathrm{gun} \circ \mathbb{R}\mathrm{gun} = \mathrm{gun} \circ \mathrm{flat}$$

(and similarly for $\mathbb{F}^\times$ and $\mathbb{G}$.) The first equation immediately follows form the fact that $\mathrm{gun}(\mathrm{sing}(a)) = \mathrm{sing}(a)$. For the second one, note that the vertices of $\mathrm{gun}(\mathrm{flat}(g))$ correspond to the finite paths of $\mathrm{flat}(g)$, while those of $\mathrm{gun}(\mathrm{flat}(\mathrm{gun}(\mathbb{R}\mathrm{gun}(g))))$ correspond to those of $\mathrm{flat}(\mathrm{gun}(\mathbb{R}\mathrm{gun}(g)))$. Furthermore, every path α in a graph of the form $\mathrm{flat}(h)$ 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 $\mathrm{gun}(h)$ is the same as a path in h. Consequently, each path of $\mathrm{flat}(\mathrm{gun}(\mathbb{R}\mathrm{gun}(g)))$ 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

$$\mathrm{dom}_+(\mathrm{gun}(\mathrm{flat}(g))) \quad \text{and} \quad \mathrm{dom}_+(\mathrm{gun}(\mathrm{flat}(\mathrm{gun}(\mathbb{R}\mathrm{gun})))) .$$

As this bijection preserves the labelling it follows that

$$\mathrm{gun}(\mathrm{flat}(g)) = \mathrm{gun}(\mathrm{flat}(\mathrm{gun}(\mathbb{R}\mathrm{gun}))) . \qquad \square$$

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

$$\mathfrak{G} = \left\langle V, \leq_{\mathrm{pf}}, \leq_{\mathrm{so}}, (\mathrm{suc}_x)_{x\in X}, (P_c)_{c\in\Sigma+\xi}, R\right\rangle$$

where

- $V := \mathrm{dom}_+(g)$ is the set of vertices,
- $\leq_{\mathrm{pf}}$ the reflexive transitive closure of the edge relation,

- $\leq_{\mathrm{so}}$ 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) $\mathrm{MSO}[\leq_{\mathrm{so}}]$, monadic second-order logic with *successor ordering*, is MSO over structures of the above form where every relation can be used.

(b) $\mathrm{MSO}[\leq_{\mathrm{pf}}]$, monadic second-order logic with *forest ordering*, or *prefix ordering*, is MSO over the above structures, but without the relation $\leq_{\mathrm{so}}$.

(c) $\mathrm{MSO}[\mathrm{suc}]$, monadic second-order logic with *successor*, is MSO over the above structures, but without the relations $\leq_{\mathrm{pf}}$ and $\leq_{\mathrm{so}}$.

(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 $\mathrm{MSO}_m[\leq_{\mathrm{so}}]$, $\mathrm{MSO}_m[\leq_{\mathrm{pf}}]$,... 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 $\mathfrak{A}$ and $\mathfrak{B}$ 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

$$\mathfrak{A}, \bar{P}\bar{a} \equiv_L \mathfrak{B}, \bar{Q}\bar{b} \quad :\text{iff} \quad [\mathfrak{A} \vDash \varphi(\bar{P}, \bar{a}) \Leftrightarrow \mathfrak{B} \vDash \varphi(\bar{Q}, \bar{a})]\,, \text{ for all } \varphi \in L\,.$$ ⌟

We can compute this relation by induction on m as follows.

Proposition 1.8. *Let $\mathfrak{A}$ and $\mathfrak{B}$ 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*

$$\mathfrak{A}, \bar{a} \equiv_{\mathrm{FO}_{m+1}} \mathfrak{B}, \bar{b}$$

if, and only if,

- *for every $a' \in A$, there is some $b' \in B$ with $\mathfrak{A}, \bar{a}a' \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{b}b'$,*
- *for every $b' \in B$, there is some $a' \in A$ with $\mathfrak{A}, \bar{a}a' \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{b}b'$.*

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 $\mathfrak{A}$ and $\mathfrak{B}$ 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*

$$\mathfrak{A}, \bar{P}\bar{a} \equiv_{\mathrm{FO}_{m+1}} \mathfrak{B}, \bar{Q}\bar{b}$$

if, and only if,

- *for every $a' \in A$, there is some $b' \in B$ with $\mathfrak{A}, \bar{P}\bar{a}a' \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{Q}\bar{b}b'$,*
- *for every $b' \in B$, there is some $a' \in A$ with $\mathfrak{A}, \bar{P}\bar{a}a' \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{Q}\bar{b}b'$.*
- *for every $P' \subseteq A$, there is some $Q' \subseteq B$ with $\mathfrak{A}, \bar{P}P'\bar{a} \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{Q}Q'\bar{b}$,*
- *for every $Q' \subseteq B$, there is some $P' \subseteq A$ with $\mathfrak{A}, \bar{P}P'\bar{a} \equiv_{\mathrm{FO}_m} \mathfrak{B}, \bar{Q}Q'\bar{b}$.*

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

$$\delta^{\mathfrak{A}} := \{\, a \in A \mid \mathfrak{A} \vDash \delta(a) \,\}$$

and relations

$$\varphi^{\mathfrak{A}}_R := \{\, \bar{a} \in A^n \mid \mathfrak{A} \vDash \varphi_R(\bar{a}) \,\}, \quad \text{for } R \in \Gamma$$

(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*

$$\tau(\mathfrak{A}) \vDash \varphi \quad \text{iff} \quad \mathfrak{A} \vDash \varphi^\tau , \quad \textit{for all structures } \mathfrak{A} .$$

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*

$$\sum_{i \in \mathfrak{I}} \mathfrak{A}^i \vDash \varphi \quad \text{iff} \quad \langle \mathfrak{I}, [\![\chi_0]\!], \dots, [\![\chi_{n-1}]\!] \rangle \vDash \psi ,$$

where $[\![\chi]\!] := \{ i \in I \mid \mathfrak{A}^i \vDash \chi \}$.

2. Finite Forests

Before turning to the general case, let us take a brief look a languages of finite forests.

Definition 2.1. Let $\mathfrak{A}$ be an $\mathbb{F}^{\text{fin}}$-algebra.

(a) The *horizontal semigroup* associated with $\mathfrak{A}$ is the structure $\langle A_\varnothing, + \rangle$ where the *horizontal product* $+$ is defined by

$$a + b := \pi(s) , \quad \text{where } s \text{ is the forest consisting of the two trees } \operatorname{sing}(a) \text{ and } \operatorname{sing}(b) \text{ (in this order).}$$

(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* $+$ 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 2.2.

(a) *Every finitary* $\mathbb{F}^{\mathrm{fin}}$*-algebra is* $\mathrm{MSO}[\leq_{\mathrm{so}}]$*-definable.*

(b) *Every finitary* $\mathbb{F}^{\mathrm{fin}}$*-algebra with a commutative horizontal semigroup is* $\mathrm{CMSO}[\leq_{\mathrm{pf}}]$*-definable.*

(c) *Every finitary* $\mathbb{F}^{\mathrm{fin}}$*-algebra whose horizontal semigroup is commutative and aperiodic is* $\mathrm{MSO}[\leq_{\mathrm{pf}}]$*-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

$$s \vDash \varphi_c \quad \text{iff} \quad \pi(s) = c\,, \quad \text{for } s \in \mathbb{F}^{\mathrm{fin}}_\xi C\,.$$

The formula φ_c guesses sets $(Z_a)_a$ such that

$$Z_a := \{\, v \in \mathrm{dom}(s) \mid \eta(s|_v) = a \,\}\,, \quad \text{for } a \in \bigcup_{\zeta \subseteq \xi} A_\zeta\,,$$

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

$$\big\langle \{u_0, \dots, u_{m-1}\}, \leq_{\mathrm{so}}, \bar Z \big\rangle \vDash \vartheta_a \quad \text{iff} \quad b_0 + \dots + b_{m-1} = a\,.$$

Consequently, we can set

$$\varphi_c := \exists \bar Z \Big[Z_c(\mathrm{root}) \wedge \forall x \Big[\bigwedge_a Z_a x \leftrightarrow \bigvee_{b,d : b(d) = a} (P_b x \wedge \vartheta_d^{(x)}) \Big]\Big]\,,$$

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

$$\min\{n_a, k\} \quad \text{and} \quad n_a \bmod p\,, \quad \text{for every } a \in \bigcup_{\zeta \subseteq \xi} A_\xi\,,$$

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 □

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. ⌟

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 $\mathbb{C} : \mathcal{D} \to \mathcal{D}$ 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

$$\mathbb{C}X := \sum_I X^{\operatorname{dom}(I)}\,,$$

where I ranges over all countable linear orders.

(b) The flattening operation $\operatorname{flat} : \mathbb{C}\mathbb{C} \Rightarrow \mathbb{C}$ 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

$$\operatorname{dom}(\operatorname{flat}(S)) := \sum_{i \in \operatorname{dom}(S)} \operatorname{dom}(S(i)) \\ = \{ \langle i, k \rangle \mid i \in \operatorname{dom}(S),\ k \in \operatorname{dom}(S(i)) \}$$

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 $\operatorname{sing} : \mathrm{Id} \Rightarrow \mathbb{C}$ maps an element $a \in A$ to the linear order $\operatorname{sing}(a)$ 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 $\mathsf{Set}^{\{\varnothing\}}$. Then an alphabet $\Sigma \in \mathsf{Set}$ corresponds to a $\{\varnothing\}$-sorted set Σ' with $\Sigma'_\varnothing = \Sigma$. Furthermore, every vertex of a forest $s \in \mathbb{F}\Sigma'$ has sort $\varnothing$ and must therefore be a leaf. Consequently, s is just the linear order formed by its countably many roots. Thus, we obtain an embedding $\mathbb{C} \Rightarrow \mathbb{F}|_{\{\varnothing\}}$. 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 $\mathbb{F}|_{\{\varnothing\}}$-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 $\langle \mathbb{C}, \operatorname{flat}, \operatorname{sing} \rangle$ 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 $\langle \operatorname{dom}(s), \leq, (P_a)_{a \in A} \rangle$

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 $\mathbb{F}|_{\{\varnothing\}}$, this version of MSO corresponds to the variant $\mathrm{MSO}[\leq_{\mathrm{so}}]$ defined above. ⌟

Compositionality is straightforward.

Theorem 3.3. MSO *is strongly* $\mathbb{C}$*-compositional.*

Proof. First note that, for $s \in \mathbb{C}\mathbb{C}\Sigma$, we can write

$$\mathrm{flat}(s) \cong \tau\Big(\sum_{i \in \mathrm{dom}(s)} s(i)\Big)$$

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

$$\sum_{i \in \mathrm{dom}(s)} s(i) \vDash \varphi \quad \text{iff} \quad \mathbb{C}\theta_{\Delta_m}(s) \vDash \psi\,,$$

where Δ_m is the set of MSO-formulae of quantifier-rank at most m. □

Corollary 3.4. MSO *is varietal and* $\mathbb{C}$*-compositional.*

Proof. By Proposition IV.4.11. □

Next we turn to syntactic algebras. The proof utilises the submonad $\mathbb{C}^{\mathrm{reg}}$ 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 $a_0 \sqcup\!\sqcup \cdots \sqcup\!\sqcup\, a_{n-1}$.

(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 $s \times \omega^{\mathrm{op}}$, for $s \in \mathbb{C}A$,
- an n-ary operation $s_0 ш \cdots ш s_{n-1}$, for $s_0, \dots, s_{n-1} \in \mathbb{C}A$ and $n < \omega$.

$a \in A$ denotes the singleton order $\mathrm{sing}(a)$; $s + t$ 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 $s \times \omega^{\mathrm{op}}$ 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 Läuchli-Leonard operations. We denote the set of all such orders by $\mathbb{C}^{\mathrm{reg}} A$. ⌟

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

$$Q := (0 + 1)^* 1 .$$

Note that Q equipped with the lexicographic ordering is a dense linear order. We define a labelling $s : Q \to A$ by

$$s(0^n 1) := a_n \quad \text{and} \quad s(w10^n 1) := a_n , \qquad \text{for } n < \omega \text{ and } w \in \{0,1\}^* .$$

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 $w_0 + w_1 + \cdots \in K$ and $\cdots + w_1 + w_0 \in K$.*
- *If $s \in \mathbb{C}K$ such that $\mathbb{C}\pi(s)$ is the shuffle of a finite set, then $\mathrm{flat}(s) \in K$.*

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 $\mathrm{dom}(s)$.

$$x \sim y \quad :\text{iff} \quad x = y \text{ or the factor corresponding to the interval } (x, y] \text{ or } (y, x] \text{ is decomposable.}$$

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 $H = \varnothing$, 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 induction on 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. □

Theorem 3.7. *The inclusion $\mathbb{C}^{\mathrm{reg}} \Rightarrow \mathbb{C}$ is dense over the class of all finite $\mathbb{C}$-algebras.*

Proof. Let $s \in \mathbb{C}C$. Then $s \in K$, where K is the set from Lemma 3.6. By the inductive definition of K, we can construct a Läuchli-Leonard term s° with $\pi(s^\circ) = \pi(s)$. □

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 Läuchli-Leonard operations from trees to sets of trees.

Definition 3.9. (a) For an n-ary Läuchli-Leonard operation σ and languages $K_0, \dots, K_{n-1} \subseteq \mathbb{C}\Sigma$ we set

$$\sigma(\bar{K}) := \{ \sigma(\bar{s}) \mid s_i \in K_i \}.$$

(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 $\{\mathrm{sing}(a)\}$ with $a \in \Sigma$,
- boolean operations $K \cup L$, $K \cap L$, $\mathbb{C}\Sigma \setminus K$,
- the Läuchli-Leonard operations $K + L$, $K \times \omega$, $K \times \omega^{\mathrm{op}}$, $K_0 ш \cdots ш K_{n-1}$,
- 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 $\{\mathrm{sing}(a)\}$ 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 $P \cup Q = \mathrm{dom}(s)$ 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 $P \subseteq \mathrm{dom}(s)$. We use the fact that every subset $P \subseteq \mathrm{dom}(s)$ induces an equivalence relation on $\mathrm{dom}(s)$ 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 $P \subseteq \mathrm{dom}(s)$ 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 $K \times \omega^{\mathrm{op}}$. Finally, we have $s \in K_0 ш \cdots ш K_{n-1}$ 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. □

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*.

$$\begin{aligned} a \leq_{\mathsf{L}} b \quad &:\text{iff} \quad a = s + b\,, && \text{for some } s \in S\,,\\ a \leq_{\mathsf{R}} b \quad &:\text{iff} \quad a = b + s\,, && \text{for some } s \in S\,,\\ a \leq_{\mathsf{J}} b \quad &:\text{iff} \quad a = s + b + t\,, && \text{for some } s, t \in S\,. \end{aligned}$$

Let $\equiv_{\mathsf{L}}$, $\equiv_{\mathsf{R}}$, and $\equiv_{\mathsf{J}}$ be the corresponding equivalence relations and set

$$\equiv_{\mathsf{H}} \;:=\; \equiv_{\mathsf{L}} \cap \equiv_{\mathsf{R}}\,.$$

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) $(\equiv_{\mathsf{L}} \circ \equiv_{\mathsf{R}}) \;=\; \equiv_{\mathsf{J}} \;=\; (\equiv_{\mathsf{R}} \circ \equiv_{\mathsf{L}})$

(b) *$a \equiv_{\mathsf{J}} ab \equiv_{\mathsf{J}} b$ implies $a \equiv_{\mathsf{R}} ab$ and $ab \equiv_{\mathsf{L}} b$.*

(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 $b := s(0)$ be the first letter of s and let s' be the suffix such that $s = b + s'$. Since $b \equiv_{\mathrm{J}} e$, 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 $b \equiv_{\mathrm{R}} b'$, for some other element $b' \equiv_{\mathrm{J}} e$. 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 $t = u_0 + u_1 + \cdots$ and an element $f \in S$ such that

$$\pi(u_i + \cdots + u_j) = f\,, \quad \text{for all } 0 < i \leq j < \omega\,.$$

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

$$\pi(t) = b + f + f + \cdots = b + f \times \omega\,.$$

As $f \equiv_{\mathrm{J}} e$, 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 $g \equiv_{\mathrm{H}} e \equiv_{\mathrm{H}} h$. 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

$$g \times \omega = (g \times m) \times \omega = e \times \omega \quad \text{and} \quad h \times \omega = (h \times m) \times \omega = e \times \omega\,.$$

Consequently,

$$\pi(e+t) = g + h \times \omega = g + g \times \omega = g \times \omega = e \times \omega\,. \qquad \square$$

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 $\leq_{\mathrm{J}}$.*

(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

$$K_+ := K \cup K \times \omega \cup K \times \omega^{\mathrm{op}} \cup \big((K \cup J) + (K \cup J)\big)\,.$$

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 $u_0 + u_1 + \cdots$ with $u_i \in L$. Similarly, an *ω^{op}-suffix* of s is a suffix of the form $\cdots + v_1 + v_0$ 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

$$s = u + s_0 + \cdots + s_{n-1} + v\,,$$

where $s_0, \ldots, s_{n-1} \in K_+$, 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

$$K\,, \quad K \times \omega\,, \quad K \times \omega^{\mathrm{op}}\,, \quad \text{or} \quad \big((K \cup J) + (K \cup J)\big)\,.$$

(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 $x \in \operatorname{dom}(w)$, we can find an ω-prefix $u_0 + u_1 + \cdots$ containing x. If $x \in \operatorname{dom}(u_i)$, then u_{i+1} is a chain in L that lies entirely to the right of x. Consequently, we can find a cofinal sequence $v_0, v_1, \ldots$ 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 $w = v_0 + v_1 + \cdots \in L \times \omega$.

(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 $s_0, s_1, \ldots \in K_+$ and chains $t_1, t_2, \ldots$ 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 $(s_0 + s_1) + (s_2 + s_3) + \cdots$ with $s_i + s_{i+1} \in L$. Hence, s would have an ω-prefix $u + s_0 + s_1 + \cdots$ that is longer than u. A contradiction.

It remains to pove 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 $t_{i+1} = \cdots + w_2 + w_1 + w_0$ 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 $K \times \omega^{\mathrm{op}} \subseteq K_+$. 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_0 be the union of all prefixes p of s that belong to K, and let s_1 be the suffix of s such that $s = s_0 + s_1$. Then $s_0 \in K_+$. Set $a_0 := \pi(s_0)$ and $a_1 := \pi(s_1)$. If $a_0 \in J$ then $a_0 + a_1 \in J$ implies, by Lemma 3.12 (b), that the R-classes of a_0 and $a_0 + a_1 = \pi(s)$ coincide. By (1), a_0 is definable. Hence, so is its R-class.

It remains to consider the case where $a_0 \in F_0$. First, suppose that s_1 has a least element. Let c be its label. By maximality of s_0, we have $a_0 + c \in J$. As above, this implies that the R-classes of $a_0 + c$ and $\pi(s)$ coincide. Since $a_0 + c$ 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 $s_1 = \cdots + w_2 + w_1 + w_0$ where $\pi(w_i) = \pi(w_j)$, for $i, j \neq 0$. Let c be the value of this product. It follows that s_1 has a prefix with product $c \times \omega^{\mathrm{op}}$. As above, this implies that the R-class of $\pi(s)$ is equal to that of $a_0 + c \times \omega^{\mathrm{op}}$. It remains to define c. If $c \in F_0$, it is MSO-definable by inductive hypothesis. Otherwise, Lemma 3.13 implies that $c \times \omega^{\mathrm{op}} = e \times \omega^{\mathrm{op}}$, for any idempotent $e \in J$.

(b) Given $s \in \pi^{-1}[F]$, let $s = u + s_0 + \cdots + s_{n-1} + v$ 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 $\pi(s_0) + \cdots + \pi(s_{n-1})$ in MSO since we can obtain $\pi(s_i)$ via (1) and MSO allows us to evaluate finite products in a finite semigroup.

(c) Note that

$$\pi^{-1}[F] = \mathbb{C}S \setminus \pi^{-1}[I] \quad \text{where} \quad I := \pi^{-1}[S \setminus F]\,.$$

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 $N_a \subseteq \mathbb{C}S$. We call a chain $w \in \mathbb{C}S$ *N-critical* if it satisfies one of the following conditions.

- $w = \mathrm{sing}(a)$, for some $a \notin F$.
- $w = u + v$ with $u \in N_a$, $v \in N_b$ with $a + b \notin F$.
- $w = u_0 + u_1 + \cdots$ with $u_0, u_1, \dots \in N_a$ and $a \times \omega \notin F$.
- $w = \cdots + u_1 + u_0$ with $u_0, u_1, \dots \in N_a$ and $a \times \omega^{\mathrm{op}} \notin F$.
- $w \in N_{a_0} ш \cdots ш N_{a_{n-1}}$ with $a_0 ш \cdots ш a_{n-1} \notin F$.

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

$$s \vDash \varphi_a \quad \text{iff} \quad \pi(s) = a, \quad \text{for } s \in \pi^{-1}[F].$$

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 $\leq_{\mathrm{J}}$. 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

$$M := \{ s \in \mathbb{C}S \mid \pi(s) \in F \text{ or } s \text{ has an } N\text{-critical factor} \}$$

coincides with $\mathbb{C}S$. 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. □

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 $J := [a]_{\mathrm{J}}$ 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 $\pi^{-1}(a)$. □

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. ⌟

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 $\mathbb{T}_{\varnothing}\Sigma$.*

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

$$\mathrm{dom}(t) = \{0,1\}^* \quad \text{and} \quad t(v) := \begin{cases} a & \text{if } v \in P\,, \\ b & \text{otherwise}\,, \end{cases}$$

is antiregular.

Let s and s' be two subtrees of t rooted at u and u', respectively. Let $\tilde{u}$ be the mirror image of u. Then $u\tilde{u} \in L$ and $u'\tilde{u} \notin L$, which implies that $s(\tilde{u}) = a$ and $s'(\tilde{u}) = b$. In particular, s and s' are not isomorphic. □

Theorem 4.3. *Suppose that $\mathcal{D} = \mathsf{Pos}$ 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 $\xi \in \varXi$, let $\varDelta_\xi$ be the set of all finite trees in $\mathbb{T}_\xi\Sigma$. As Σ contains only binary elements, every leaf of a tree $t \in \varDelta_\xi$ must be labelled by a variable. Hence, t has at most $|\xi|$ leaves and, therefore, at most $|\xi| - 1$ internal vertices. This implies that $\varDelta_\xi$ is a finite set.

The domains of the desired algebra $\mathfrak{A}$ are

$$A_\xi := \varDelta_\xi \uplus \{\bot, *\}\,, \quad \text{for } \xi \in \varXi\,,$$

which we order such that $\bot$ 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 $\eta : \mathbb{T}^\times\Sigma \to A$ that will recognise L. For

$t \in \mathbb{T}^{\times}_{\xi}\Sigma$, we set

$$\eta(t) := \begin{cases} t & \text{if } t \in \Delta_{\xi}\,, \\ \bot & \text{if } t \text{ has a subtree without variables that is not densely} \\ & \quad \text{antiregular}\,, \\ ? & \text{every subtree of } t \text{ has a variable and some variable} \\ & \quad \text{occurs at least twice}\,, \\ * & t \text{ has a subtree without variables and every such subtree} \\ & \quad \text{is densely antiregular}\,. \end{cases}$$

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 $\mathrm{flat}(t|_v)$ is densely antiregular. For $t \in \mathbb{T}^{\times}_{\xi}A$, we then set

$$\pi(t) := \begin{cases} \bot & \text{if } t \text{ contains the label } \bot, \\ \bot & \text{if } t \text{ has a subtree without variables that is not} \\ & \quad \text{good}, \\ * & \text{if } t \text{ has a subtree without variables and every such} \\ & \quad \text{subtree is good}, \\ * & \text{if } t \text{ contains the label } * \text{ and every subtree of } t \text{ has} \\ & \quad \text{a variable}, \\ ? & \text{if } t \in \mathbb{T}^{\times}(\Delta + \{?\}),\ t \text{ contains the label } ?, \text{ and} \\ & \quad \text{every subtree of } t \text{ has a variable}, \\ ? & \text{if } t \in \mathbb{T}^{\times}\Delta, \text{ every subtree of } t \text{ has a variable, and} \\ & \quad \text{some variable appears more than once in } t, \\ \mathrm{flat}(t) & \text{if } t \in \mathbb{T}^{\times}\Delta \text{ is finite and every variable appears} \\ & \quad \text{exactly once}. \end{cases}$$

It is straightforward (but tedious) to check that these cases cover all possibilities, that the function π is monotone, and that is 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 $\eta : \mathbb{T}^\times \Sigma \to \mathfrak{A}$ an $\mathbb{T}^\times$-morphism. □

Definition 4.4. We call the algebra $\mathfrak{A}$ from the above proof the *Bojańczyk-Klin algebra.* ⌟

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 finite 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 $f : X^n \to X$ and $g_0, \dots, g_{n-1} : X^m \to X$, 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 $[3]^\xi \to [3]$ together with a special error value $\perp$. We turn

$A := (A_\xi)_\xi$ into a $\mathbb{T}^\times$-algebra by defining the following multiplication $\pi : \mathbb{T}^\times A \to A$. 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 $\bar{\beta} \in (B_\zeta)^\xi$ where the β_x are disjoint. If $\alpha = \langle x_{i_0}, \dots, x_{i_{k-1}} \rangle$, we set

$$\alpha(\bar{\beta}) := \beta_{x_{i_0}} \dots \beta_{x_{i_{k-1}}} ,$$

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 $\sigma, \tau : [n] \to [n]$ 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 $p : \mathfrak{B} \times \mathfrak{B} \to \mathfrak{B}$ be the projection to the first component. By looking at the first components, we see that

$$\begin{aligned}
&\pi\big(\mathbb{T}p(s_{\sigma(0)})\big)\dots\pi\big(\mathbb{T}p(s_{\sigma(n-1)})\big) \\
&= \gamma\big(\pi(\mathbb{T}p(s_0)),\dots,\pi(\mathbb{T}p(s_{n-1}))\big) \\
&= \alpha \\
&= \langle x_0,\dots,x_{2m-1}\rangle\,.
\end{aligned}$$

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. ⌟

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 $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_0, \Omega\rangle$ 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*,
- $q_0 \in Q$ the *initial state*,
- $\Omega : Q \to \omega$ 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 $s \in \mathbb{F}^{\times}_{\xi}\Sigma$ is a function $\rho : \mathrm{dom}_{+}(s) \to Q$ such that, for every vertex $v \in \mathrm{dom}_{+}(s)$,

$$\langle \rho(v), s(v), \bar{q} \rangle \in \Delta\,,$$

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,

- $\sigma(w) = q_0$, 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:

$$\liminf_{n\to\infty} \Omega(\rho(\beta_n)) \text{ is even,} \quad \text{where } \beta_n \text{ denotes the } n\text{-th vertex on } \beta.$$

(d) The language *recognised* by $\mathcal{A}$ is the set

$$L(\mathcal{A}) := \{\, s \in \mathbb{F}^{\times}_{\xi}\Sigma \mid \text{ there is an accepting run of } \mathcal{A} \text{ on } s \,\}\,.$$ ⌟

Example. Let $\Sigma = \{a, b\}$ where a and b both have arity 1, and let $K \subseteq \mathbb{F}^{\times}_{\varnothing}\Sigma$ the set of all forests with at least one occurrence of the letter a. The following automaton $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_0, \Omega \rangle$ recognises K. We use two states $Q := \{0, 1\}$ with initial state $q_0 := 1$. The input sort is $\xi := \varnothing$ and the priority function is the identity $\Omega(q) := q$. Finally, the transition relation is defined by

$$\Delta := \{\, \langle 1, a, q \rangle \mid q \in Q \,\} \cup \{\, \langle q, c, q \rangle \mid q \in Q,\ c \in \{a, b\} \,\}\,.$$

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 0. 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 $\leq_{\mathrm{pf}}$). 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 $s(v) = a$. ⌟

The behaviour of an automaton $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_0, \Omega\rangle$ on an input forest $s \in \mathbb{F}^{\times}_{\xi}\Sigma$ 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, $v \in \mathrm{dom}_+(s)$ a vertex, and $p \subseteq \mathrm{dom}_+(s)$ 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 $u \in \mathrm{dom}_+(s)$. 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 $f : p \to Q$ 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 $\langle r, q_0\rangle$ where r is the root place of s and q_0 the initial state.

Each round of the game proceeds in two phases. In the position $\langle p, q\rangle$, Automaton first chooses a function $f : p \to Q$ with $\sigma(f) = q$ and Pathfinder selects a vertex $u \in p$. Then Automaton picks a transition $\langle f(u), s(u), \bar{q}'\rangle \in \Delta$ 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 $\bar{q}'$.

The game continues until either a position is reached where the corresponding player cannot make a move, in which case this player looses, 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 $K \subseteq \mathbb{F}^{\times}_{\xi}\Sigma$ is* MSO*-definable if, and only if, it can be recognised by some automaton $\mathcal{A}$.*

Definition 5.4. A forest $s \in \mathbb{F}^{\times}_{\xi}\Sigma$ is *regular* if, up to isomorphism, s has only finitely many different subtrees. ⌟

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 $\mathrm{flat}(s)$ 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 $\mathcal{A} = \langle Q, \Sigma, \xi, \sigma, \Delta, q_0, \Omega\rangle$ be a forest automaton.

(a) A *partial run* of $\mathcal{A}$ on some forest $s \in \mathbb{F}^{\times}_{\zeta}\Sigma$ (where ζ might be different from ξ) is a function $\rho : \mathrm{dom}_{+}(s) \to Q$ such that

- ρ satisfies the transition relation

  $$\langle \rho(v), s(v), \bar{q}\rangle \in \Delta$$

  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 $s \in \mathbb{F}^{\times}_{\zeta}\Sigma$ is the tuple $\tau = \langle p, \bar{U}\rangle$ where p is the starting state of ρ and, for each $z \in \zeta$, 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. ⌟

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 $\mathrm{flat}(s)$. 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* MSO$[\leq_{\mathrm{so}}]$*-definable $\mathbb{M}$-algebras.*

Proof. Let $\mathbb{M}^{\mathrm{reg}}A \subseteq \mathbb{M}^\times A$ 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 $\mathbb{M}^{\mathrm{reg}} \Rightarrow \mathbb{M}^\times$ 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, $B \subseteq A_0 \times \cdots \times A_{n-1}$, and $t \in \mathbb{M}^\times B$ a forest with $\pi(t) = \bar{a}$. We have to find a regular forest $t^\circ \in \mathbb{M}^{\mathrm{reg}}B$ 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 $\mathcal{A}_i$, $\Omega_i : Q_i \to \omega$ its priority function, and $K_i := \operatorname{rng} \Omega_i$ the set of priorities used by $\mathcal{A}$. 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 $\sigma_i : \mathbb{V}B \to \mathbb{M}^\times \mathbb{V}C_i$, which we can extend to a morphism $\hat{\sigma}_i : \mathbb{M}^\times \mathbb{V}B \to \mathbb{M}^\times \mathbb{V}C_i$.

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

$$\begin{aligned}
\pi(\hat{\sigma}_i(s)) &= \pi(\mathrm{flat}(\mathbb{M}^\times \sigma_i(s))) \\
&= \pi(\mathbb{M}^\times \pi(\mathbb{M}^\times \sigma_i(s))) \\
&= \pi(\mathbb{M}^\times p_i(s)) \\
&= p_i(\pi(s))\,,
\end{aligned}$$

where $p_i : A_0 \times \cdots \times A_{n-1} \to A_i$ 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 $\bar{U} \in \prod_{i<n} \mathcal{P}(K_i \times Q_i)$, 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 $\langle \{ \langle 0, q_0^i \rangle \} \rangle_{i<n}$, where q_0^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 $(\bar{V}^x)_{x \in \xi}$ 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 $\bar{V}^x$.

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 $\bar{U}^0, \bar{V}^0, \bar{U}^1, \bar{V}^2, \dots$ and let $W_{ix}^l(q)$ be the sets used in the l-th turn by Automaton to determine the next position $\bar{V}^l = (\bar{V}_x^l)_x$. We call a sequence $k_0, q_0, k_1, q_1, k_2, q_0, \dots$ an *i-trace* of this play if $\langle k_0, q_0 \rangle \in U_i^0$ and, for all $l < \omega$,

$$\langle k_{l+1}, q_{l+1} \rangle \in W_{ix}^l(q_l) , \quad \text{for some } x \text{ with } \bar{U}^{l+1} = V_x^l .$$

We say that the play satisfies the parity condition if, for all $i < n$,

$$\liminf_{l < \omega} k_l \text{ is even,} \quad \text{for all } i\text{-traces } k_0, q_0, k_1, q_1, k_2, q_0, \dots .$$

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 $\mathcal{A}_i$. □

Remark. The Bojańczyk-Klin algebra shows that $\mathbb{M}^{\mathrm{reg}}$ is *not* densely embedded in $\mathbb{M}^\times$ over the class of all *finitary* $\mathbb{M}^\times$-algebras. ⌟

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 $\mathrm{MSO}_{(m)}$ 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 $\mathrm{MSO}_{(m)}$ is finite (up to logical equivalence). Let $\equiv_{(m)}$ be the equivalence relation which holds for two forests if they satisfy the same $\mathrm{MSO}_{(m)}$-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

$$S(v) \equiv_{(m)} T(v)\,, \quad \text{for all } v\,, \quad \text{implies} \quad \mathrm{flat}(S) \equiv_{(m)} \mathrm{flat}(T)\,.$$

For the proof, fix a formula $\varphi_{\mathcal{A},\tau} \in \mathrm{MSO}_{(m)}$ with $\mathrm{flat}(S) \vDash \varphi_{\mathcal{A},\tau}$. We have to show that $\mathrm{flat}(T)$ also satisfies $\varphi_{\mathcal{A},\tau}$, i.e., that there is a partial run of $\mathcal{A}$ on $\mathrm{flat}(T)$ 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 $\mathrm{flat}(T)$ 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 $u \in \mathrm{dom}_+(T)$. 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 $\langle p, \upsilon\rangle$ where p is a place and υ a profile, Automaton tries to show that there exists a partial run ρ on the subforest rooted at p

with profile υ. 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 υ. 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 $\bar{\lambda}$ for all the subforests attached to the copy of $T(u)$ in $\mathrm{flat}(T)$ such that the 'composition' of the profile of ρ and $\bar{\lambda}$ is equal to κ_u. This is done as follows.

Let $\mu = \langle p, \bar{U} \rangle$ 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.
- $\upsilon = \langle p, \bar{V} \rangle$ is the composition of μ and the profiles λ. Formally,

$$V_x = \big\{ \langle l, q' \rangle \bigm| \langle k, q, \lambda \rangle \in W_z\,,\ \lambda = \langle q, \bar{L} \rangle\,,\ \langle k', q' \rangle \in L_x\,,\ l = \min\{k, k'\} \big\}\,.$$

Given $\bar{W}$, 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 $\langle v, \upsilon \rangle$, then Automaton wins if, and only if, υ is of the form $\langle q, \bar{U} \rangle$ with $U_x = \{q\}$ and $U_z = \varnothing$, for $z \neq x$. Otherwise, Pathfinder wins. If the leaf is not labelled by a variable, then Automaton wins if he can choose $\mu = \langle p, \bar{U} \rangle$ such that $U_x = \varnothing$, for all x.

In the case where the game is infinite, Automaton wins if the sequence of pairs $\langle k_0, q_0, \lambda_0 \rangle, \langle k_1, q_1, \lambda_0 \rangle, \dots$ chosen by Pathfinder satisfies the parity condition

$$\liminf_{i<\omega} k_i \quad \text{is even}\,.$$

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 $\mathrm{flat}(T)$ with profile τ.

(Every partial run of $\mathcal{A}$ on $\mathrm{flat}(T)$ 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. □

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. ⌟

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 $\Delta_m := \mathrm{MSO}_{(m)}$ 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

$$\pi^{-1}(\sigma) \cap \mathbb{M}C\,.$$

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 $t \in \mathbb{M}C$, it follows that $\pi(t) = \sigma$ if, and only if, for every forest $s \in \mathbb{M}\Sigma$ 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. □

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.

The following observation can sometimes be used to simplify proofs of MSO-definability: we only need to check elements of arity at most one.

Proposition 5.12. *Let* L *be either* MSO *or* FO. *A finitary* $\mathbb{T}$*-algebra* $\mathfrak{A}$ *is* MSO*-definable if, and only if, it has a finite set* $C \subseteq A$ *of generators such that*

$$\pi^{-1}(a) \cap \mathbb{T}C \text{ is regular}, \quad \text{for every } a \in A \text{ of arity at most } 1\,.$$

Remark. Note that this statement fails for $\mathbb{T}^\times$-algebras. ⌟

Before giving the proof, we need to collect a few results about factorisations. We denote by $\mathcal{F}(t)$ the set of all factorisations T of t such that the trees $T(v)$ are singletons for all vertices v of T with more than one successor. The *height* of a factorisation T is the height of the tree T.

We call a tree $t \in \mathbb{T}A$ *reduced* if it has no non-trivial factor of arity at most one, that is, for every factorisation T of t and every vertex $v \in \mathrm{dom}(T)$ of arity at most one, we have $T(v) = \mathrm{sing}(a)$, for some $a \in A$. The important fact about reduced trees is that they are small.

Lemma 5.13. *Let* $\mathfrak{A}$ *be a* $\mathbb{T}$*-algebra and* $\xi \in \Xi$. *Every reduced tree* $t \in \mathbb{T}_\xi A$ *has height at most* $2|\xi|$.

Proof. We prove the claim by induction on $m := |\xi|$. For $\xi = \varnothing$, note that every reduced tree of sort ξ is of the form $\mathrm{sing}(a)$, for some $a \in A_\varnothing$. Hence, the height is 0. For the inductive step, suppose that $\xi \neq \varnothing$ and consider a reduced tree $t \in \mathbb{T}_\xi A$. We distinguish two cases.

First, suppose that the root has an arity greater than 1. As t is reduced, every subtree attached to the root must have fewer variables than t. By inductive hypothesis, their height is at most $2(m-1)$. Hence, the height of t is at most $2(m-1)+1$.

It remains to consider the case where the root has arity 1. As t is reduced, the successor must then have arity greater than 1. Hence, the attached subtree satisfies the above case, which means that its height is bounded by $2(m-1)+1$. Consequently, the height of t is at most $2(m-1)+2 = 2m$. □

Next we will show that the set $\mathcal{F}(t)$ of factorisations of t contains reduced trees. For the proof we will employ the following ordering on $\mathcal{F}(t)$. For $S, T \in \mathcal{F}(t)$, we set

$$\begin{aligned} S \sqsubseteq T \quad &:\text{iff} \quad \text{there is some } U \in \mathbb{TTT}A \text{ such that } S = \mathrm{flat}(U) \text{ and} \\ &\qquad\quad U(v) \text{ is a factorisation of } T(v), \text{ for } v \in \mathrm{dom}(U)\,. \end{aligned}$$

Lemma 5.14. *The set $\mathcal{F}(t)$ is inductively ordered by $\sqsubseteq$, i.e., every chain as an upper bound.*

Proof. Let $(T_i)_{i \in I}$ be an increasing sequence in $\mathcal{F}(t)$. We have to find an upper bound. Note that every factorisation T of t induces an equivalence relation $\approx_T$ on $\mathrm{dom}(t)$ by

$$\begin{aligned} u \approx_T v \quad &:\text{iff} \quad u \text{ and } v \text{ are vertices belonging to the same} \\ &\qquad\quad \text{factor } T(w)\,. \end{aligned}$$

It follows that

$$T_i \sqsubseteq T_j \quad \text{implies} \quad \approx_{T_i} \subseteq \approx_{T_j}\,, \qquad \text{for } i < j \text{ in } I\,.$$

Thus the sequence $(\approx_{T_i})_{i \in I}$ is increasing and the limit

$$\approx \;:=\; \bigcup_{i \in I} \approx_i$$

is an equivalence relation on $\mathrm{dom}(t)$ that corresponds to some factorisation T of t. We will show that $T \in \mathcal{F}(t)$. Thus T is the desired upper bound for $(T_i)_{i \in I}$.

To prove the claim, note that every $\approx$-class E is the union of an increasing sequence $(E_i)_{i \in I}$ of $\approx_{T_i}$-classes. Since each T_i belongs to $\mathcal{F}(t)$, every relation E_i is of one of the following two types.

(I) The class is a singleton.

(II) The class corresponds to a factor of arity at most one.

If there are arbitrarily large i such that E_i is of type (I), the sequence is constant and the limit E is also of type (I). Otherwise, the limit E is a union of classes of type (II) and, hence, is also of type (II). As this holds for all classes of $\approx$, it follows that $T \in \mathcal{F}(t)$. □

Lemma 5.15. *Let $\mathfrak{A}$ be a $\mathbb{T}$-algebra and $C \subseteq A$ a set such that $A_\varnothing \cup A_{\{z\}} \subseteq C$, for some variable z. Every tree $t \in \mathbb{T}_\xi C$ has a factorisation $T \in \mathcal{F}(t)$ such that*

(I) *T is reduced,*

(II) *the height of T is at most $2m$, and*

(III) *$\mathbb{T}\pi(T) \in \mathbb{T}C$.*

Proof. By Lemma 5.14, we can use Zorn's Lemma to find a maximal element $T \in \mathcal{F}(t)$. We claim that T is the desired factorisation.

(I) For a contradiction, suppose otherwise. Then there exists a factorisation U of T and a vertex $u \in \mathrm{dom}(U)$ of arity at most one such that $U(u)$ is not a singleton. Let T' be the tree obtained from T by replacing the factor $U(u)$ by its product. Then, $T \sqsubset T'$ and T is not maximal.

(II) follows from (I) by Lemma 5.13.

(III) Note that every factor $T(v)$ is either a singleton with label in C or a tree of arity at most one. Renaming the edge labels of T, we may assume that every factor $T(v)$ of the latter kind belongs to $\mathbb{T}_{\{z\}}C$. Since $A_\varnothing \cup A_{\{z\}} \subseteq C$, it therefore follows that $\pi(T(v)) \in C$. Consequently, $\mathbb{T}\pi(T) \in \mathbb{T}C$. □

Proof of Proposition 5.12. For the nontrivial direction, suppose that $\mathfrak{A}$ is an algebra as in the proposition and let $C \subseteq A$ be the corresponding set of generators. To prove that $\mathfrak{A}$ is L-definable, we fix an element $a \in A_\xi$, for an arbitray $\xi \in \Xi$. We have to show that $\pi^{-1}(a) \cap \mathbb{T}C$ is regular. Set $C' := C \cup A_\varnothing \cup A_{\{z\}}$, for some variable z, and fix $t \in \mathbb{T}C$. By Lemma 5.15, t has a factorisation $T \in \mathcal{F}(t)$ such that T is reduced, its height is at most $2m$, and $\mathbb{T}\pi(T) \in \mathbb{T}C'$. It follows that $\mathbb{T}\pi(T) \in H(a)$ where

$$H(a) := \{\, s \in \mathbb{T}C' \mid s \text{ has height at most } 2m \text{ and } \pi(s) = a \,\}\,.$$

Consequently, we have

$$\begin{aligned} \pi(t) = a \quad &\text{iff} \quad \pi(\mathrm{flat}(T)) = a \\ &\text{iff} \quad \pi(\mathbb{T}\pi(T)) = a \quad \text{iff} \quad \mathbb{T}\pi(T) \in H(a)\,. \end{aligned}$$

For every finite tree s, we will construct an L-formula ϑ_s such that

$$t \vDash \vartheta_s \quad \text{iff} \quad t \text{ has a factorisation } T \in \mathcal{F}(t) \text{ such that } \mathbb{T}\pi(T) = s\,.$$

Then it follows that

$$\pi(t) = a \quad \text{iff} \quad \mathbb{T}\pi(T) \in H(a) \quad \text{iff} \quad t \vDash \bigvee_{s \in H(a)} \vartheta_s\,,$$

as desired. Hence, it remains to construct the formulae ϑ_s.

First, note that we can encode a factorisation T of t by a set Z that contains the root of each factor $T(v)$. As T has bounded height, the set Z has bounded size. Hence, we can quantify over an enumeration $\bar{z}$ of Z. We can therefore set

$$\begin{aligned} \vartheta_s := \exists \bar{z}\Big[&\text{`}\bar{z} \text{ encodes a factorisation } T \text{ in } \mathcal{F}(t)\text{'} \\ &\wedge \bigwedge_{v \in \mathrm{dom}(s)} \text{`the factor } T(v) \text{ evaluates to } s(v)\text{'}\Big]\,. \end{aligned}$$

The first part of this formula is clearly expressible in L. For the second part, note that s is finite and each factor $T(v)$ is either a singleton or a term of arity at most one. In the first case it is trivial to compute the product. In the second case, we can use the formulae defining the sets $\pi^{-1}(a) \cap \mathbb{T}C$, for $a \in A_\varnothing \cup A_{\{z\}}$. □

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 $\mathrm{flat}(S) \equiv_m \mathrm{flat}(T)$.

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 $\mathrm{flat}(S)$ and a copy s of $S(v)$ in $\mathrm{flat}(S)$, we denote by $\bar a^s$ the tuple

$$a_i^s := \begin{cases} a_i & \text{if } a_i \in \mathrm{dom}_+(s)\,, \\ v & \text{if } v \text{ is a hole of } s \text{ and } v \leq_{\mathrm{pf}} a_i \text{ in } \mathrm{flat}(S)\,, \\ v & \text{if } v \text{ is the root of } s \text{ and } v \nleq_{\mathrm{pf}} a_i \text{ in } \mathrm{flat}(S)\,. \end{cases}$$

We use the same notation for parameters in $\mathrm{flat}(T)$. 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 $\mathrm{flat}(S)$ and $\bar b$ in $\mathrm{flat}(T)$,

$$\langle s, \bar a^s\rangle \equiv_m \langle t, \bar b^t\rangle\,, \quad \text{for all } v,\ \text{copies } s \text{ of } S(v),\ \text{and copies } t \text{ of } T(v),$$

implies that

$$\langle \mathrm{flat}(S), \bar a\rangle \equiv_m \langle \mathrm{flat}(T), \bar b\rangle\,.$$

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 $v_0, \dots, v_n$ be the path from the root v_0 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 $d_n \in \mathrm{dom}_+(T_n(v_n))$ 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^s 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

$$\langle \mathrm{flat}(S_n), \bar{a}, c \rangle \equiv_m \langle \mathrm{flat}(T_n), \bar{b}, d \rangle .$$

Since $\mathrm{flat}(S_n) = \mathrm{flat}(S)$ and $\mathrm{flat}(T_n) = \mathrm{flat}(T)$, the claim follows.

In the same way we can show that, for every choice of d in $\mathrm{flat}(T)$, we find a matching vertex c in $\mathrm{flat}(S)$. □

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 $\varphi : \mathbb{F}\Sigma \to \mathbb{F}\Gamma$ be a morphism of $\mathbb{F}$-algebras and let $\varphi_0 := \varphi \circ \mathrm{sing} : \Sigma \to \mathbb{F}\Gamma$ 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.

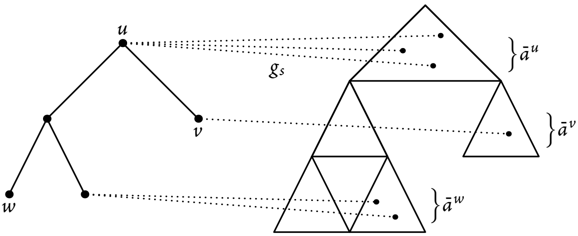

Figure 3.: Definition of g_s, I_u, and $\bar{a}^u$

Note that a forest of the form $\varphi(s) = \mathrm{flat}(\mathbb{F}\varphi_0(s))$ is obtained from s by replacing each vertex u by a forest $\varphi_0(s(u))$. For $s \in \mathbb{F}\Sigma$, we denote by $g_s : \mathrm{dom}_+(\varphi(s)) \to \mathrm{dom}_+(s)$ the function mapping a vertex u of $\varphi(s)$ to the vertex $v := g_s(u)$ such that the copy of the forest $\varphi_0(s(v))$ replacing v in $\varphi(s)$ contains u. (Note that this copy of $\varphi_0(s(v))$ 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 $\varphi_0(s(u))$.

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

$$\big\langle s, g_s(\bar{a})\big\rangle \equiv_m \big\langle t, g_t(\bar{b})\big\rangle$$
$$\text{and} \quad \big\langle \varphi_0(s(u)), \bar{a}^u\big\rangle \cong \big\langle \varphi_0(t(v)), \bar{b}^v\big\rangle\,, \quad \text{for all } u, v \text{ with } I_u = I_v \neq \varnothing\,,$$

implies

$$\langle \varphi(s), \bar{a}\rangle \equiv_m \langle \varphi(t), \bar{b}\rangle\,.$$

For $m = 0$, this is immediate. Hence, suppose that $m > 0$. We have to check the back-and-forth properties. Thus, let $c \in \mathrm{dom}_+(\varphi(s))$ and set

$u := g_s(c)$. Then there is some $v \in \mathrm{dom}_+(t)$ such that

$$\langle s, g_s(\bar{a}), g_s(c)\rangle \equiv_{m-1} \langle t, g_t(\bar{b}), v\rangle .$$

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 $I_u = I_v \neq \varnothing$, then there exists an isomorphism

$$\sigma : \left\langle \varphi_0(s(u)), \bar{a}^u \right\rangle \to \left\langle \varphi_0(t(v)), \bar{b}^v \right\rangle$$

and we can set $d := \sigma(c)$.

Otherwise, $I_u = I_v = \varnothing$ and $s(u) = t(v)$ implies that $\varphi_0(s(u)) \cong \varphi_0(t(v))$. Hence, can choose some element d of $\varphi_0(t(v))$ such that

$$\left\langle \varphi_0(s(u)), c \right\rangle \cong \left\langle \varphi_0(t(v)), d \right\rangle .$$

In both cases, it now follows that

$$\left\langle s, g_s(\bar{a}), g_s(c) \right\rangle \equiv_m \left\langle t, g_t(\bar{b}), g_t(d) \right\rangle$$

$$\text{and} \quad \left\langle \varphi_0(s(u)), \bar{a}^u c^u \right\rangle \cong \left\langle \varphi_0(t(v)), \bar{b}^v d^v \right\rangle , \quad \text{if } I_u = I_v \neq \varnothing ,$$

which, by inductive hypothesis, implies that

$$\langle \varphi(s), \bar{a}c\rangle \equiv_{m-1} \langle \varphi(t), \bar{b}d\rangle .$$

The other direction follows by symmetry. ☐

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 $\mathfrak{T} = \langle T, \mathrm{suc}_0, \mathrm{suc}_1, \leq_{\mathrm{pf}} \rangle$ 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.

$$\begin{aligned}
\mathrm{suc}(x, y) &:= \mathrm{suc}_0(x, y) \vee \mathrm{suc}_1(x, y) \\
\mathrm{zigzag}(x, y; u, v) &:= [\mathrm{suc}_0(x, y) \wedge \mathrm{suc}_1(u, v)] \\
&\quad\ \vee [\mathrm{suc}_1(x, y) \wedge \mathrm{suc}_0(u, v)] \\
\mathrm{probe}(x, y) &:= x \leq_{\mathrm{pf}} y \wedge \neg\exists z[\mathrm{suc}(y, z)] \\
&\quad\ \wedge \forall u \forall v \forall w[x \leq_{\mathrm{pf}} u \wedge \mathrm{suc}(u, v) \wedge \mathrm{suc}(v, w) \\
&\qquad\qquad \wedge w \leq_{\mathrm{pf}} y \to \mathrm{zigzag}(u, v; v, w)] \,.
\end{aligned}$$

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

$$\begin{aligned}
\mathrm{even}(x) := \exists y[\mathrm{probe}(x, y) \wedge{} & \\
& \exists u \exists v[x = y \vee [\mathrm{suc}(x, u) \wedge u \leq_{\mathrm{pf}} v \wedge \mathrm{suc}(v, y) \\
&\qquad \wedge \mathrm{zigzag}(x, u; v, y)]]] \,.
\end{aligned}$$

Consequently, the desired formula is

$$\begin{aligned}
&\forall x \forall y[\mathrm{suc}(x, y) \to [\mathrm{even}(x) \leftrightarrow \neg\mathrm{even}(y)]] \\
&\quad \wedge \exists x \forall y[x \leq_{\mathrm{pf}} y \wedge \mathrm{even}(x)] \,.
\end{aligned}$$

□

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 $\varphi := \mathbb{F}^{\times}\Sigma \to \mathbb{F}^{\times}\Gamma$ be the morphism mapping $a(x_0)$ to $b(x_0, x_0)$ 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 from $a^n(c)$ 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 [19]. Our presentation is based on [10]. For an introduction to Green's relations see, e.g., [34, 5]. The Bojańczyk-Klin algebra is from [11] and Theorem 4.5 from [44]. Introductions to automata on infinite trees can be found in [41, 28]. Our model of a forest automaton is a simplified version of the MSO-automata introduced in [43]. Proposition 6.2 and Lemma 6.3 are by Potthoff [35, 36] and its corollary by Bojańczyk and Michalewski [13].

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* $\mathcal{D} = \mathsf{Set}$.

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 MSO$[\leq_{\mathrm{pf}}]$ 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 $\vDash_t$ 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 $s|_{v_i} \vDash_\mathrm{t} \varphi_i$.

(c) Let $\mathcal{K}$ be a family of ∞-word languages. *Counting $\mathcal{K}$-temporal logic* $\mathrm{cTL}[\mathcal{K}]$ (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 ∞-word language $K \in \mathcal{K}[\Phi]$, and every positive integer n, $\mathsf{E}_n K$ 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 $\vDash_\mathrm{f}$ for forest formulae and one $\vDash_\mathrm{t}$ for tree formulae. In both cases boolean combinations are defined in the usual way. For a tree t, we define

$$\begin{aligned}
t \vDash_\mathrm{t} P_a \quad &:\text{iff} \quad \text{the root of } t \text{ has label } a \,, \\
t \vDash_\mathrm{t} \varphi \quad &:\text{iff} \quad t' \vDash_\mathrm{f} \varphi \,, \quad \text{for forest formulae } \varphi \,,
\end{aligned}$$

where t' denotes the forest obtained by removing the root from t. For a forest s we define

$$s \vDash_{\mathrm{f}} \mathsf{E}_n K \quad :\text{iff} \quad \text{there exist at least } n \text{ paths each of which has a } \Phi\text{-labelling in } K\,.$$

We do not require the paths above to be disjoint. They can even be prefixes of each other.

(d) *$\mathcal{K}$-temporal logic* $\mathrm{TL}[\mathcal{K}]$ is the fragment of $c\mathrm{TL}[\mathcal{K}]$ 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 $\mathcal{K}^\infty$, $\mathcal{K}^+$, and $\mathcal{K}^\omega$ 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 $c\mathrm{TL}[\mathcal{S}^x] := c\mathrm{TL}[\mathcal{K}^x]$ and $\mathrm{TL}[\mathcal{S}^x] := \mathrm{TL}[\mathcal{K}^x]$.

(f) The *modal rank* of a formula $\varphi \in c\mathrm{TL}[\mathcal{K}]$ 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 $\mathrm{TL}[\mathcal{K}]$ and $c\mathrm{TL}[\mathcal{K}]$. In may 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 $\mathbb{F}^{\mathrm{fb}}A$ the set of *finitely branching forests* in $\mathbb{F}^\times A$ 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* $\mathrm{TL}[\mathcal{K}]$ *or* $c\mathrm{TL}[\mathcal{K}]$ *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 MPL = cCTL*, FO = WMPL = cwCTL*, and cPDL = 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 TL[$\mathcal{K}$] and 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 $K_0, \dots, K_{n-1} \subseteq \mathbb{F}_\varnothing\Sigma$ be forest languages that form a partition of $\mathbb{F}_\varnothing\Sigma$.

(a) For $t \in \mathbb{F}_\varnothing\Sigma$, we denote by $t[\bar{K}]$ the forest where each vertex gets an additional label encoding to which language K_i the attached subtree belongs.

Formally, we define $t[\bar{K}] \in \mathbb{F}_\varnothing(\Sigma \times [n])$ by

$$t[\bar{K}](v) := \langle t(v),\ i\rangle \quad \text{where} \quad i \text{ is the unique index with } t|_v \in K_i\,.$$

(b) For $L \subseteq \mathbb{F}_\varnothing(\Sigma \times [n])$, we set

$$L[\bar{K}] := \{\, t \in \mathbb{F}_\varnothing\Sigma \mid t[\bar{K}] \in L \,\}\,.$$

(c) For a family $\mathcal{K}$ of languages, we denote by $\mathsf{Casc}(\mathcal{K})$ the smallest family of languages that contains $\mathcal{K}$, is closed under finite boolean operations, and satisfies

$$L, \bar{K} \in \mathsf{Casc}(\mathcal{K}) \quad \Rightarrow \quad L[\bar{K}] \in \mathsf{Casc}(\mathcal{K})\,. \qquad \lrcorner$$

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*

$$\{\, s \in \mathbb{F}_\varnothing\Sigma \mid s \text{ has at least } n \text{ paths with labelling in } K \,\}$$
$$\text{and} \quad \{\, s \in \mathbb{F}_\varnothing\Sigma \mid \text{the roof of } s \text{ has label } a \,\}\,,$$

for $K \in \mathcal{K}[\Sigma]$, $a \in \Sigma$, and $n < \omega$.

(a) *A language $L \subseteq \mathbb{F}_\varnothing\Sigma$ is $\mathrm{cTL}[\mathcal{K}]$-definable if, and only if, $L \in \mathsf{Casc}(\mathcal{M})$.*

(b) *If we modify the definition of $\mathcal{M}$ to only allow $n = 1$, we obtain an analogous characterisation of $\mathrm{TL}[\mathcal{K}]$-definable languages.*

Proof. We only prove the claim for $\mathrm{cTL}[\mathcal{K}]$. The second statement follows in exactly the same way.

($\Leftarrow$) Clearly, every language in $\mathcal{M}$ is $\mathrm{cTL}[\mathcal{K}]$-definable. Hence, the claim follows by induction if we can show that, if L and $\bar{K}$ are $\mathrm{cTL}[\mathcal{K}]$-definable languages, then so is $L[\bar{K}]$. Suppose that L is defined by the formula φ and K_i by ψ_i. Then we can define $L[\bar{K}]$ by the formula obtained from φ by replacing every atom of the form $P_{\langle a,i\rangle}$ by the formula $P_a \wedge \psi_i$.

($\Rightarrow$) Given a formula $\varphi \in \mathrm{cTL}[\mathcal{K}]$ we prove by induction on φ that it defines a language in $\mathsf{Casc}(\mathcal{M})$. For $\varphi = P_a$, the defined language is

$$\{\, s \in \mathbb{F}_\varnothing\Sigma \mid \text{the roof of } s \text{ has label } a \,\}\,,$$

which belongs to $\mathcal{M} \subseteq \mathsf{Casc}(\mathcal{M})$ by assumption. As $\mathsf{Casc}(\mathcal{M})$ is closed under finite boolean operations, it remains to consider formulae of the form $\varphi = \mathsf{E}_n K$ 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 $\mathsf{Casc}(\mathcal{M})$. By assumption, so does the language

$$M := \{\, s \in \mathbb{F}_\varnothing \Sigma \mid s \text{ has at least } n \text{ paths with labelling in } K \,\}\,.$$

Consequently, $M[\bar{L}] \in \mathsf{Casc}(\mathcal{M})$, which is the language defined by $\mathsf{E}_n K$. □

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 $\sigma : \mathrm{dom}(s) \to \mathrm{dom}(t)$ 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 $\mathrm{dom}(s) \to \mathrm{dom}(t)$.

(a) A language $K \subseteq \mathbb{T}\Sigma$ is *sibling-commutative* if it is closed under permutations. ⌟

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 $K \subseteq \mathbb{F}^{\mathrm{fb}}_\varnothing \Sigma$ is sibling-commutative if, and only if, its syntactic algebra $\mathrm{Syn}(K)$ satisfies the equations*

$$c + d = d + c \quad \textit{and} \quad a(x + y) = a(y + x)\,.$$

for all $a \in A_{\{x,y\}}$ and $c, d \in A_\varnothing$.

Proof. ($\Leftarrow$) If s is a permutation of t, we have $\mathrm{syn}_K(s) = \mathrm{syn}_K(t)$. Since syn_K recognises K it follows that $s \in K \Leftrightarrow t \in K$ and that K is sibling-commutative.

($\Rightarrow$) Given an element $a \in A_{\{x,y\}}$, we have to show that

$$a(x + y) \sim_K a(y + x)$$

(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, $a(x + y)$ and $a(y + x)$ are permutations of each other. As K is sibling-commutative we therefore have

$$p[a(x + y)] \in K \Leftrightarrow p[a(y + x)] \in K\,.$$

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.) ⌟

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 $Z \subseteq \mathrm{dom}(s) \times \mathrm{dom}(t)$ 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. ⌟

Theorem 2.4. *Let Σ be an alphabet where every symbol has arity 1, and let $K \subseteq \mathbb{F}_\varnothing \Sigma$ be a forest language with a syntactic algebra. Then K is bisimulation-invariant if, and only if,* Syn(K) *satisfies the following equations:*

$$c + c = c\,, \qquad\qquad a(x + x) = a(x)\,,$$
$$c + d = d + c\,, \qquad a(x_0 + x_1 + x_2 + x_3) = a(x_0 + x_2 + x_1 + x_3)\,,$$

for all $a \in \mathrm{Syn}_{\{x\}}(K)$ and $c, d \in \mathrm{Syn}_\varnothing(K)$.

Proof. We denote the syntactic congruence of K by $\sim_K$.

($\Rightarrow$) Given elements $c, d \in \mathrm{Syn}_\varnothing(K)$, we fix forests $s \in \mathrm{syn}_K^{-1}(c)$ and $t \in \mathrm{syn}_K^{-1}(d)$. 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 $a \in \mathrm{Syn}_{\{x\}}(K)$ and $s \in \mathrm{syn}_K^{-1}(a)$. 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 $a(x_0) = \mathrm{syn}_K(s) = \mathrm{syn}_K(s') = a(x_0 + x_0)$.

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 Syn(K) satisfies the four equations above and let s and s' be bisimilar forests. We claim that $\mathrm{syn}_K(s) = \mathrm{syn}_K(s')$, which implies that $s \in K \Leftrightarrow s' \in K$.

Fix a bisimulation relation $Z \subseteq \mathrm{dom}(s) \times \mathrm{dom}(s')$. W.l.o.g. we may assume that Z only relates vertices on the same level of the respective forests

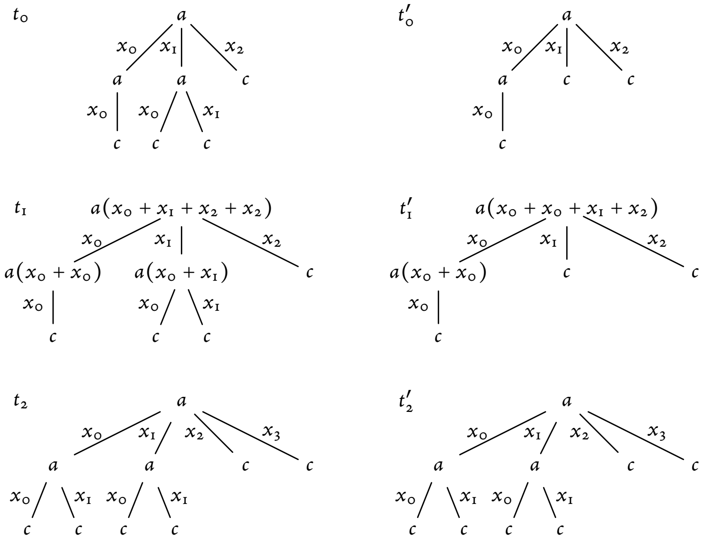

Figure 2.: Transforming bisimilar forests

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 $\mathrm{dom}(s) \cup \mathrm{dom}(s')$ 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 $\mathrm{Syn}(K)$. We define two new forests $t_0, t'_0 \in \mathbb{F}_\varnothing \mathrm{Syn}(K)$ with the same domains as, respectively, s and s' and the following labelling. If $v \in \mathrm{dom}(s)$ has the successors $u_0, \dots, u_{n-1}$, we set

$$t_0(v) := \mathrm{syn}_K(s(v))(x_0 + \dots + x_{n-1})$$

and we make u_i an x_i-successor of v in t_0. We obtain t'_0 from s' in the same way. By associativity it follows that $\pi(t_0) = \mathrm{syn}_K(s)$ and $\pi(t'_0) = \mathrm{syn}_K(s')$.

Next we make the shapes of the forests t_0 and t'_0 the same. Let t_1 and t'_1 be the forests with the same domains as t_0 and t'_0 and the following labelling. For every vertex v of t_0 with successors $u_0, \dots, u_{n-1}$ and labelling

$$t_0(v) = a(x_0 + \dots + x_{n-1}),$$

we set

$$t_1(v) := a(x_0 + \dots + x_0 + \dots + x_{n-1} + \dots + x_{n-1}),$$

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'_0. By the top right equation in the statement of the theorem, the value of the product is not affected by this modification. Hence, $\pi(t_1) = \pi(t_0)$ and $\pi(t'_1) = \pi(t'_0)$.

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 $u_0, \dots, u_{n-1}$ and label

$$t_1(v) = a(x_0 + \dots + x_0 + \dots + x_{n-1} + \dots + x_{n-1}),$$

we set

$$t_2(v) := a(x_0 + \dots + x_k + \dots + x_l + \dots + x_m),$$

where we number the variables from left-to-right, e.g., $a(x_0+x_0+x_1+x_2+x_2)$ becomes $a(x_0 + x_1 + x_2 + x_3 + x_4)$, 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 $\sigma : \mathrm{dom}(t) \to \mathrm{dom}(r')$ such that, for a vertex v of r with successors $u_0, \dots, u_{n-1}$,

$$r'(v) = r(v)(x_{\sigma_v(0)} + \dots + x_{\sigma_v(n-1)}),$$

where the function $\sigma_v : [n] \to [n]$ 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 $u_0, \dots, u_{n-1}$ and labelling

$$r(v) = a(x_0 + \dots + x_{n-1}),$$

we set

$$\hat{r}(v) := a(x_{\sigma_v(0)} + \dots + x_{\sigma_v(n-1)}),$$

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_0 there is some component of t'_0 with the same product. Therefore, we can write

$$\pi(t_0) = a_0 + \dots + a_{m-1} \quad \text{and} \quad \pi(t'_0) = b_0 + \dots + b_{n-1}$$

where the sets $\{a_0, \dots, a_{m-1}\}$ and $\{b_0, \dots, b_{m-1}\}$ coincide. Using the equations $c + c = c$ and $c + d = d + c$ we can therefore transform $\pi(t_0)$ into $\pi(t'_0)$. Consequently,

$$\mathrm{syn}_K(s) = \pi(t_0) = \pi(t'_0) = \mathrm{syn}_K(s').$$

As syn_K recognises K it follows that $s \in K \Leftrightarrow s' \in K$, as desired. □

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 $\mathrm{EF} = \mathrm{TL}[\mathcal{K}]$ and $\mathrm{cEF} = \mathrm{cTL}[\mathcal{K}]$ 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 $\mathrm{EF} := \mathrm{cEF}_1$ and $\mathrm{EF}^m := \mathrm{cEF}_1^m$. ⌟

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 the 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 $\mathbb{F}^{\mathrm{fb}}_{\varnothing} A$. Consequently, the equations below are all well-defined.

Theorem 3.2. *A language $K \subseteq \mathbb{F}^{\mathrm{fb}}_{\varnothing}\Sigma$ is definable in the logic cEF_n if, and only* *if, its syntactic algebra* $\mathrm{Syn}(K)$ *satisfies the following equations:*

$$
\begin{aligned}
c + d &= d + c & (a(x) + b(x))^\omega &= (ab(x))^\omega \\
(ab)^\pi &= b(ab)^\pi & (a(x) + c)^\omega &= (a(x + c))^\omega \\
a^\omega + a^\omega &= a^\omega & (a(x + c + c))^\omega &= (a(x + c))^\omega \\
(abb')^\omega &= (ab'b)^\omega & \big[a(b(x, y))^{\omega y}\big]^{\omega x} &= \big[ab(x, x)\big]^{\omega x} \\
(aab)^\omega &= (ab)^\omega & \big[a(x + bc + c)\big]^\omega &= \big[a(x + bc)\big]^\omega
\end{aligned}
$$

$$
\begin{aligned}
&a_\xi(c, \dots, c) + (n - |\xi|) \times c = a_\xi(c, \dots, c) + (n - |\xi| + 1) \times c\,, \\
&\big[a(x + (a(n \times x))^\pi(c))\big]^\omega = n \times (a(n \times x))^\pi(c)
\end{aligned}
$$

for all $a, b, b' \in \mathrm{Syn}_{\{x\}}(K)$, $c, d \in \mathrm{Syn}_\varnothing(K)$, $a_\xi \in \mathrm{Syn}_\xi(K)$ *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* $\mathbb{F}^{\mathrm{fb}}$*-algebra* $\mathfrak{A}$ *that is generated by* $A_\varnothing \cup A_{\{x\}}$*, 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 $N := m_0^{2m_1} + m_0$ where $m_0 := |A_\varnothing|$ 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

$(1)_n$ $\quad a_\xi(c, \dots, c) + (n - |\xi|) \times c = a_\xi(c, \dots, c) + (n - |\xi| + 1) \times c$

$(2)_n$ $\quad \big[a(x + (a(n \times x))^\pi(c))\big]^\omega = n \times (a(n \times x))^\pi(c)$

We have to show that $\mathfrak{A}$ also satisfies $(1)_N$ and $(2)_N$.

For $(2)_N$, note that $n \geq N \geq |A_\varnothing|$ implies that $N \times c = \pi \times c = n \times c$, for all $c \in A_\varnothing$. 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 $c \in A_\varnothing$. If $|\xi| \leq N - m_0$, then $N - |\xi| \geq m_0 = |A_\varnothing|$ 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 $|\xi| > N - m_0 = m_0^{2m_1}$. As $\mathfrak{A}$ is generated by $A_\varnothing \cup A_{\{x\}}$, there exists some forest $s \in \mathbb{F}_\xi(A_\varnothing \cup A_{\{x\}})$ 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

$$s = r(t_0 + \dots + t_{m_0^2 - 1})$$

where each subtree t_i contains some variable. Then we can find indices $i_0 < \dots < i_{m_0 - 1}$ such that $\pi(t_{i_0}(\bar{c})) = \dots = \pi(t_{i_{m_0-1}}(\bar{c}))$ (where $\bar{c}$ denotes as many copies of c as appear in the respective term). Hence, $(1)_n$ again implies that

$$\begin{aligned}
&a(\bar{c}) + (N - |\xi|) \times c \\
&= \pi(s(\bar{c})) + (N - |\xi|) \times c \\
&= \pi\big(r\big(t_0(\bar{c}) + \dots + t_{m_0^2-1}(\bar{c})\big)\big) + (N - |\xi|) \times c \\
&= \pi\big(r\big(t_0(\bar{c}) + \dots + t_{m_0^2-1}(\bar{c}) + n \times t_{i_0}(\bar{c})\big)\big) + (N - |\xi|) \times c \\
&= a(\bar{c}) + (N - |\xi|) \times c + c\,.
\end{aligned}$$

Note that a tree of height $h := m_1$ where every vertex has at most $d := m_0^2$ successors has at most $d^h = m_0^{2m_1}$ leaves. Hence, if s is not highly branching in the sense above, the fact that it contains $|\xi| > m_0^{2m_1}$ variables implies that there must be a chain $v_0 \prec \dots \prec v_{m_1}$ of vertices such that, for every $i < m_1$, there is some leaf u labelled by a variable with $v_{i-1} \prec u$ and $v_i \npreceq u$. (For $i = 0$, we omit the first condition.) Hence, we can decompose s as

$$s(\bar{c}) = r_0(\bar{c}, r_1(\bar{c}, \dots r_{m_1}(\bar{c})))\,.$$

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

$$\pi(r_{0i}) = \pi(r_{0j})\,.$$

Consequently, we can use pumping to obtain a term

$$\pi(s(\bar{c})) = \pi\big(r_{0i}(r_{ij})^n r_{jm_1}\big)$$

which contains at least n occurrences of c, and the claim follows again by $(1)_n$.

□

According to this lemma, we can check for cEF-definability of a language K, by computing its syntactic algebra $\mathrm{Syn}(K)$, 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 $K \subseteq \mathbb{F}^{\mathrm{fb}}_{\varnothing}\Sigma$ is definable in the logic* EF *if, and only if, the syntactic algebra $\mathrm{Syn}(K)$ satisfies the following equations:*

$$\begin{aligned} c + d &= d + c & (a(x) + b(x))^\omega &= (ab(x))^\omega \\ (ab)^\pi &= b(ab)^\pi & (a(x) + c)^\omega &= (a(x + c))^\omega \\ (abb')^\omega &= (ab'b)^\omega & (a(x + c + c))^\omega &= (a(x + c))^\omega \\ (aab)^\omega &= (ab)^\omega & \big[a(b(x, y))^{\omega y}\big]^{\omega x} &= [ab(x, x)]^{\omega x} \end{aligned}$$

$$ac = ac + c \qquad c = c + c \qquad [a(x + a^\pi c)]^\omega = a^\pi c\,,$$

for all $a, b, b' \in \mathrm{Syn}_{\{x\}}(K)$ and $c, d \in \mathrm{Syn}_{\varnothing}(K)$.

Corollary 3.7. *It is decidable whether a given regular language $K \subseteq \mathbb{F}^{\mathrm{fb}}_{\varnothing}\Sigma$ 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 $s, t \in \mathbb{F}^{\mathrm{fb}}\Sigma$, we define

$$s \approx^0_n t \quad :\text{iff} \quad \text{the roots of } s \text{ and } t \text{ have the same label,}$$

and, inductively, we set $s \approx^{m+1}_n t$ if,

- the roots of s and t have the same label,

- for every n-tuple $\bar{x}$ in $\mathrm{dom}(s)$ not containing the root, there is some n-tuple $\bar{y}$ in $\mathrm{dom}(t)$ not containing the root such that

$$s|_{x_i} \approx^m_n t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j\,, \qquad \text{for all } i, j < n\,,$$

- for every n-tuple $\bar{y}$ in $\mathrm{dom}(t)$ not containing the root, there is some n-tuple $\bar{x}$ in $\mathrm{dom}(s)$ not containing the root such that

$$s|_{x_i} \approx^m_n t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j\,, \qquad \text{for all } i, j < n\,.$$

To simplify notation, we will frequently write $x \approx^m_n y$ for vertices x and y instead of the more cumbersome $s|_x \approx^m_n t|_y$.

(b) For two forests $s, t \in \mathbb{F}^{\mathrm{fb}}\Sigma$ with possibly several components, we set $s \sim^{m+1}_n t$ if

- for every n-tuple $\bar{x}$ in s there is some n-tuple $\bar{y}$ in t such that

$$s|_{x_i} \approx^m_n t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j\,, \qquad \text{for all } i, j < n\,,$$

- for every n-tuple $\bar{y}$ in t there is some n-tuple $\bar{x}$ in s such that

$$s|_{x_i} \approx^m_n t|_{y_i} \quad \text{and} \quad x_i = x_j \Leftrightarrow y_i = y_j\,, \qquad \text{for all } i, j < n\,. \qquad \lrcorner$$

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* $\mathrm{tp}^m_n(s)$ of a tree $s \in \mathbb{F}\Sigma$ by

$$\mathrm{tp}^0_n(s) := a \quad \text{and} \quad \mathrm{tp}^{m+1}_n(s) := \langle a, \theta_s \rangle\,,$$

where a is the label at the root of s and

$$\theta_s := \big\{ \langle l, \sigma \rangle \bigm| l \leq n\,,\ x_0, \dots, x_{l-1} \in \mathrm{dom}(s) \text{ distinct, not equal to the root, } \sigma = \mathrm{tp}^m_n(s|_{x_0}) = \dots = \mathrm{tp}^m_n(s|_{x_{l-1}}) \big\}\,.$$

(b) For an arbitrary forest $s \in \mathbb{F}\Sigma$, we set

$$\mathrm{Tp}^{m+1}_n(s) := \theta_s\,,$$

where

$$\theta_s := \Big\{ \langle l, \sigma \rangle \Bigm| l \leq n\,,\ x_0, \dots, x_{l-1} \in \mathrm{dom}(s) \text{ distinct}\,,$$
$$\sigma = \mathrm{tp}^m_n(s|_{x_0}) = \dots = \mathrm{tp}^m_n(s|_{x_{l-1}}) \Big\}\,.$$ ⌟

A standard proof establishes the following equivalences.

Lemma 3.10. *Let $n, m < \omega$.*

(a) *For trees $s, t \in \mathbb{F}_\varnothing\Sigma$, the following statements are equivalent.*
   (1) $s \approx^m_n t$
   (2) $\mathrm{tp}^m_n(s) = \mathrm{tp}^m_n(t)$
   (3) $s \vDash \varphi \Leftrightarrow t \vDash \varphi\,,$ *for all* $\varphi \in \mathrm{cEF}^m_n\,.$

(b) *For arbitrary forests $s, t \in \mathbb{F}_\varnothing\Sigma$, the following statements are equivalent.*
   (1) $s \sim^m_n t$
   (2) $\mathrm{Tp}^m_n(s) = \mathrm{Tp}^m_n(t)$
   (3) $s \vDash \varphi \Leftrightarrow t \vDash \varphi\,,$ *for all* $\varphi \in \mathrm{cEF}^m_n\,.$

Proof. (a) (2) ⇒ (1) follows by a straightforward induction on m and (1) ⇒ (3) by induction on φ. For (3) ⇒ (2) it is sufficient to show that, for every type τ, there exists a formula $\chi_\tau \in \mathrm{EF}^m_n$ such that

$$s \vDash \chi_\tau \quad \text{iff} \quad \mathrm{tp}^m_n(s) = \tau\,, \quad \text{for every tree } s\,.$$

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

$$\chi_\tau := P_a \wedge \bigwedge_{\langle l,\sigma\rangle \in \theta} \mathsf{E}_l\,\chi_\sigma \wedge \bigwedge_{\langle l,\sigma\rangle \notin \theta} \neg\mathsf{E}_l\,\chi_\sigma\,.$$

(b) is proved in the same way. □

Corollary 3.11. *A language $L \subseteq \mathbb{F}\Sigma$ is cEF^m_n-definable if, and only if, it is regular and satisfies*

$$s \sim^m_n t \quad \textit{implies} \quad s \in L \Leftrightarrow t \in L\,, \quad \textit{for all regular forests } s, t \in \mathbb{F}_\varnothing\Sigma\,.$$

Proof. ($\Rightarrow$) follows by the implication (1) $\Rightarrow$ (3) of Lemma 3.10.

($\Leftarrow$) Set

$$\varphi := \bigvee \{ \chi_\tau \mid \tau = \mathrm{Tp}^m_n(s) \text{ for some regular forest } s \in L \} ,$$

where χ_τ are the formulae from the proof of Lemma 3.10. For a regular forest $t \in \mathbb{F}_\varnothing\Sigma$, it follows that

$$\begin{aligned} t \vDash \varphi \quad &\text{iff} \quad \mathrm{Tp}^m_n(t) = \mathrm{Tp}^m_n(s) , \quad \text{for some regular forest } s \in L , \\ &\text{iff} \quad t \sim^m_n s , \quad \text{for some regular forest } s \in L , \\ &\text{iff} \quad t \in L . \end{aligned}$$

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^m_n-definable. □

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 $A_\varnothing \cup A_{\{x\}}$, 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) $\big[a(b(x,y))^{\omega y}\big]^{\omega x} = [ab(x,x)]^{\omega x}$

(G9) $(abb')^\omega = (ab'b)^\omega$

(G10) $(aab)^\omega = (ab)^\omega$

(G11) $[a(x + bc + c)]^\omega = [a(x + bc)]^\omega$

$(\mathrm{G12})_n$ $[a(x + (a(n \times x))^\pi(c))]^\omega = n \times (a(n \times x))^\pi(c)$

where $a, b, b' \in A_{\{x\}}$, $c, d \in A_\varnothing$, $a_\xi \in A_\xi$, 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 $a, b \in A_\varnothing$, we define

$$a \leq_{\mathrm{L}} b \quad :\text{iff} \quad a = c(b) \quad \text{or} \quad a = b + d\,, \quad \text{for some } c \in A_{\{x\}}\,,\ d \in A_\varnothing\,.$$ ⌟

Lemma 3.14. *Let $\mathfrak{A}$ be an algebra for cEF_k.*

(a) *The relation $\leq_{\mathrm{L}}$ is antisymmetric.*

(b) *For $a \in A_{\{x\}}$, $c \in A_\varnothing$, we have*

$$c = c + c \quad \textit{implies} \quad ac = ac + c\,,$$
$$c = a(c, c) \quad \textit{implies} \quad c = c + c\,.$$

Proof. (a) For a contradiction, suppose that there are elements $a \neq b$ with $a \leq_{\mathrm{L}} b \leq_{\mathrm{L}} a$. 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 $(\mathrm{G1})_k$ or $(\mathrm{G2})$.

$$(1,i) \quad a = cb = cda = (cd)^\pi(a) = d(cd)^\pi(a) = da = b\,.$$
$$(1,ii) \quad a = cb = c(a + d) = (c(x + d))^\pi(a) = (c(x + d))^\pi(a) + d = a + d = b\,.$$
$$(2,i) \quad b = da = d(b + c) = (d(x + c))^\pi(b) = (d(x + c))^\pi(b) + c = b + c = a\,.$$
$$(2,ii) \quad a = b + c = a + d + c = a + k \times (d + c) = a + k \times (d + c) + d = a + d = b\,.$$

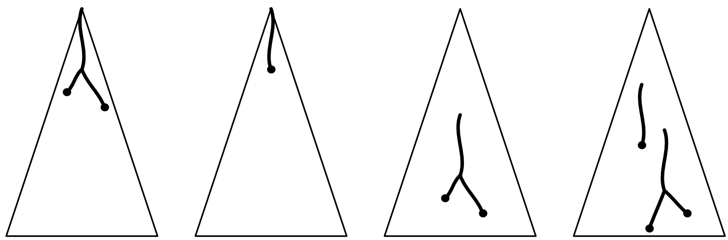

Figure 3.: A forest s with a convex set U (in bold) that has three close U-ends (on the left) and five far ones (on the right). The height is $h(s, U) = 2$.

(b) By $(\mathrm{GI})_k$ we have

$$\begin{aligned} c = c + c \quad &\text{implies} \quad ac = a(c+c) = a(k \times c) \\ & \qquad\qquad\quad = a(k \times c) + c = ac + c\,, \\ c = a(c,c) \quad &\text{implies} \quad c = a(c,c) = (a(x,c))^\pi(c) \\ & \qquad\qquad\quad = (a(x,c))^\pi(c) + c = c + c\,. \quad \square \end{aligned}$$

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 $u \leq_{\mathrm{pf}} v \leq_{\mathrm{pf}} w$ and $u, w \in U$ implies $v \in U$ (where $\leq_{\mathrm{pf}}$ denotes the forest order). We call the maximal elements (w.r.t. $\leq_{\mathrm{pf}}$) of U the *U-ends*. An U-end u is *close* if $u' \in U$, for all $u' \leq_{\mathrm{pf}} u$. 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, $U \subseteq \mathrm{dom}(s)$ a convex set that is closed under $\approx_n^m$, and set*

$$V := \{\, v \in \mathrm{dom}(t) \mid u \approx_n^m v \text{ for some } u \in U \,\}\,.$$

(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 $v <_{\mathrm{pf}} v' <_{\mathrm{pf}} v''$ of t with $v, v'' \in V$ and $v' \notin V$. Fix vertices $u <_{\mathrm{pf}} u' <_{\mathrm{pf}} u''$ 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 $u' >_{\mathrm{pf}} u_j$ with $u' \in U$. Fix a vertex $v' >_{\mathrm{pf}} v_j$ 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 $v \nleq_{\mathrm{pf}} v'$ 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 $u \leq_{\mathrm{pf}} u'$. Fix some $v' \geq_{\mathrm{pf}} v$ with $u' \approx_n^{m+1} v'$. Then $u' \in U$ implies $v' \in V$. By choice of v, there is some $v'' >_{\mathrm{pf}} v'$ with $v'' \in V$. Choose $u'' >_{\mathrm{pf}} u'$ 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.

$$\begin{aligned}
\langle s, U\rangle \asymp_0 \langle t, V\rangle \quad &:\text{iff} \quad N(s,U) = N(t,V)\,,\\
\langle s, U\rangle \asymp_{i+1} \langle t, V\rangle \quad &:\text{iff} \quad N(s,U) = N(t,V) \text{ and}\\
&\qquad\quad \#_\tau(s,U) = \#_\tau(t,V)\,,\\
&\qquad\qquad\qquad \text{for every } \asymp_i\text{-class } \tau,
\end{aligned}$$

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

$$h(s,U) := \begin{cases} 0 & \text{if } F(s;U) = \varnothing \\ 1 + \max\{\, h(s',U) \mid s' \in F(s,U) \,\} & \text{otherwise.} \end{cases}$$

By induction on l, we will prove the following claim:

$$(*)\quad s \sim_n^{m+l+2} t \text{ and } h(s,U) \leq l \text{ implies}$$
$$h(s,U) = h(t,V) \quad \text{and} \quad \langle s,U\rangle \asymp_l \langle t,V\rangle\,.$$

As $h(s,U) \leq |\hat{U}| < n$, it then follows that $\langle s,U\rangle \asymp_n \langle t,V\rangle$. In particular, $N(s,U) = N(t,V)$, as desired.

It thus remains to prove $(*)$. First, consider the case where $l = 0$. If $h(t,V) > 0$, there is some V-end v that is not close. Fix some vertex $v' <_{\mathrm{pf}} v$ with $v' \notin V$. Since $s \sim_n^{m+2} t$, we can find vertices $u' <_{\mathrm{pf}} u$ 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 $w \geq_{\mathrm{pf}} u$. But $u' <_{\mathrm{pf}} u \leq_{\mathrm{pf}} w$ implies that w is not close. Hence, $h(s,U) > 0$. A contradiction.

For the second part, suppose that $\langle s,U\rangle \not\asymp_0 \langle t,V\rangle$, 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 $u_0,\dots,u_m \in \mathrm{dom}(s)$ 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 $u' >_{\mathrm{pf}} u_j$. 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 $\langle s, U\rangle \not\asymp_{l+1} \langle t, V\rangle$. 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 $\asymp_l$-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 $\asymp_l$-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 $s_i \asymp_l t_i$, for $i \leq m$. Thus, s has at least $m + 1$ different subforest in the class τ. A contradiction. □

Corollary 3.16. *Let $s \sim_n^{m+n+2} t$ be forests such that, for every $c \in A_\varnothing$, the sets*

$$U_c := \{\, x \in \mathrm{dom}(s) \mid \pi(s|_x) = c \,\}$$
$$V_c := \{\, y \in \mathrm{dom}(t) \mid \pi(t|_y) = c \,\}$$

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 $c \in A_\varnothing$ 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 $(\mathrm{G1})_n$ 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 $(\mathrm{G1})_n$ 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. □

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 $\mathfrak{A}$ be an algebra for* cEF$_n$*. Then*

$$s \approx_n^{(n+3)(|A_\varnothing|+1)} t \quad \textit{implies} \quad \pi(s) = \pi(t) \,,$$

for all regular trees $s, t \in \mathbb{F}_\varnothing(A_\varnothing \cup A_{\{x\}})$.

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

$$s \approx_n^{f(m)} t \quad \text{implies} \quad \pi(t) = b \,,$$

where $f(m) := (n+3)(m+1)$. Set

$$\begin{aligned} S &:= \{\, x \in \mathrm{dom}(s) \mid \pi(s|_x) = b \,\} \,, \\ T &:= \{\, y \in \mathrm{dom}(t) \mid x \approx^{f(m-1)} y \text{ for some } x \in S \,\} \,. \end{aligned}$$

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

$$t|_y = a(t_0 + \cdots + t_{k-1} + t'_0 + \cdots + t'_{q-1})$$

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 $b <_{\mathrm{L}} d_i$. Hence,

$$\pi(t|_y) = a(d_0 + \cdots + d_{k-1} + q \times b)\,.$$

We have to show that this value is equal to b. Suppose that

$$s|_x = a(s_0 + \cdots + s_{l-1} + s'_0 + \cdots + s'_{p-1})\,,$$

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

$$\pi(s|_x) = a(c_0 + \cdots + c_{l-1} + p \times b)\,.$$

Since $x \in S$, we already know that this value is equal to b. Hence, it remains to show that

$$a(c_0 + \cdots + c_{l-1} + p \times b) = a(d_0 + \cdots + d_{k-1} + q \times b)\,.$$

For $c \in A_\varnothing$, let U_c be the set of all vertices $u >_{\mathrm{pf}} x$ such that $\pi(s|_u) = c$ and let V_c be the set of vertices $v >_{\mathrm{pf}} y$ with $\pi(t|_v) = c$. As $\leq_{\mathrm{L}}$ 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

$$c_0 + \cdots + c_{l-1} = d_0 + \cdots + d_{k-1}\,.$$

If $p = q$, we are done. Hence, we may assume that $p \neq q$. To conclude the proof, we set

$$U := \{ u \in S \mid x <_{\mathrm{pf}} u \} \quad \text{and} \quad V := \{ v \in T \mid y <_{\mathrm{pf}} v \}.$$

If $p > 0$, then $x \approx_n^{f(m)-1} y$ and $U \neq \varnothing$ implies $V \neq \varnothing$. 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 $d_i >_{\mathrm{L}} b$ 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 $v \in \mathrm{C}$, we can write r as a system of equations

$$x_v = a_v(x_{u_0} + \dots + x_{u_{l-1}} + c_0 + \dots + c_{q-1}), \quad \text{for } v \in \mathrm{C},$$

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 $A_\varnothing \cup A_{\{x\}}$ 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

$$\begin{aligned} &\big[a_0 \cdots a_{j-1}\big(x + d_0 + \dots + d_{k-1} + b\big)\big]^\omega \\ \text{or}\quad &\big[a_0 \cdots a_{j-1}\big(x + d_0 + \dots + d_{k-1}\big)\big]^\omega \end{aligned}$$

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

$$\big[a_0 \cdots a_{j-1}\big(x + c_0 + \dots + c_{l-1}\big)\big]^\omega$$

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 $b <_{\mathrm{L}} d_i$, it then follows that

$$\begin{aligned}\pi(t|_y) &= \big[a_0 \cdots a_{j-1}\big(x + d_0 + \cdots + d_{k-1} + b\big)\big]^{\omega} \\ &= \big[a_0 \cdots a_{j-1}(x + b)\big]^{\omega}.\end{aligned}$$

For every $i < j$, we fix some $z_i \in S$ with label a_i such that $x <_{\mathrm{pf}} z_i$ and some successor of z_i also belongs to S. Then

$$\pi(s|_{z_i}) = a_i(c^i_0 + \cdots + c^i_{l_i-1} + b + \cdots + b)\,,$$

for some $c^i_0, \dots, c^i_{l_i-1} >_{\mathrm{L}} b$. Since

$$\begin{aligned}b = \pi(s|_{z_i}) &= a_i(c^i_0 + \cdots + c^i_{l_i-1} + b + \cdots + b) \\ &\qquad\qquad \leq_{\mathrm{L}} c^i_0 + \cdots + c^i_{l_i+1} + b + \cdots + b \leq_{\mathrm{L}} b\end{aligned}$$

it follows by asymmetry of $\leq_{\mathrm{L}}$ that

$$c^i_0 + \cdots + c^i_{l_i+1} + b + \cdots + b = b$$

and $\quad a_i(b) = a_i(c^i_0 + \cdots + c^i_{l_i+1} + b + \cdots + b) = b\,.$

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 $(\mathrm{G}12)_n$, 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 $x \leq_{\mathrm{pf}} u_0 <_{\mathrm{pf}} \cdots <_{\mathrm{pf}} u_{n-1}$ each of which has some successor $v_i \in S$ with $v_i \nleq_{\mathrm{pf}} u_{i+1}$. 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 $(\mathrm{G1})_n$ that

$$\begin{aligned} b + b = \pi(s|_{u_0}) + b &= a''(b, \dots, b) + b \\ &= a''(b, \dots, b) = \pi(s|_{u_0}) = b\,. \end{aligned}$$

□

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*

$$s \sim_n^m t \quad \textit{implies} \quad \pi(s) = \pi(t)\,,$$

for all regular forests $s, t \in \mathbb{F}(A_\varnothing \cup A_{\{x\}})$.

Proof. ($\Leftarrow$) In each of the equations $(\mathrm{G1})_n$–$(\mathrm{G12})_n$, the two terms on both sides are $\sim_n^m$-equivalent.

($\Rightarrow$) By Proposition 3.17, there is some number m such that

$$s \approx_n^m t \quad \text{implies} \quad \pi(s) = \pi(t)\,,$$

for all regular trees $s, t \in \mathbb{F}(A_\varnothing \cup A_{\{x\}})$.

Suppose that $s, t \in \mathbb{F}(A_\varnothing \cup A_{\{x\}})$ 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 $\pi(s) = \pi(t)$. □

We complete the proof of Theorem 3.2 as follows.

Theorem 3.19. *A regular language $K \subseteq \mathbb{F}_\varnothing \Sigma$ is* cEF$_n$*-definable if, and only if, its syntactic algebra* Syn(K) *is an algebra for* cEF$_n$.

Proof. ($\Leftarrow$) Suppose that $\mathrm{Syn}(K)$ is an algebra for cEF_n. By Theorem 3.18, every language recognised by $\mathrm{Syn}(K)$ is invariant under $\sim^m_n$, 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 $\sim^m_n$-invariant, for some m. Thus $\sim^m_n$ is contained in the syntactic congruence of K, which means that $\mathrm{syn}_K : \mathbb{F}\Sigma \to \mathrm{Syn}(K)$ maps $\sim^m_n$-equivalent forests to the same value. Given forests $s, t \in \mathbb{F}(\mathrm{Syn}_0(K) \cup \mathrm{Syn}_1(K))$ with $s \sim^m_n t$, we can choose forests $s', t' \in \mathbb{F}\Sigma$ with $s' \sim^m_n t'$ and $s = \mathbb{F}\mathrm{syn}_K(s')$ and $t = \mathbb{F}\mathrm{syn}_K(t')$. Then

$$s \sim^m_n t \quad \text{implies} \quad \pi(s) = \mathrm{syn}_K(s') = \mathrm{syn}_K(t') = \pi(t)\,.$$

By Theorem 3.18, it follows that $\mathrm{Syn}(K)$ 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 $v \in \mathrm{dom}(t)$ of sort ξ, and a variable $x \in \xi$, we denote by $t{\downarrow}^x_v$ the subforest of t attached to the x-successors of v, i.e.,

$$t|_v = (t(v))\big((t{\downarrow}^x_v)_{x \in \xi}\big)\,.$$

(d) For a forest $t \in \mathbb{F}_\xi A$ and a tuple $\bar a \in (A_\varnothing)^\xi$, we denote by $t \Leftarrow \bar a \in \mathbb{F}_\varnothing A$ the forest obtained from t by replacing each variable x by the corresponding value a_x, for $x \in \xi$. ⌟

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 $t \in \mathbb{F}_\varnothing\Sigma$ and a morphism $\varphi : \mathbb{F}\Sigma \to \mathfrak{A}$, we define the *$\varphi$-annotation* $t \lhd \varphi \in \mathbb{F}_\varnothing(\Sigma \times A^\Delta_\varnothing)$ of t as the forest that adds to each vertex v the φ-images of the subtrees attached to successors of v. Formally,

$$(t \lhd \varphi)(v) := \left\langle t(v), (\varphi(t{\downarrow}^x_v))_{x\in\xi} \right\rangle, \quad \text{for } v \in \operatorname{dom}(t) \text{ of arity } \xi\,.$$

(b) The *cascade composition* of two morphisms $\alpha : \mathbb{F}\Sigma \to \mathfrak{A}$ and $\beta : \mathbb{F}(\Sigma \times A^\Delta_\varnothing) \to \mathfrak{B}$ is the function $\alpha \lhd \beta : \mathbb{F}_\varnothing\Sigma \to B_\varnothing$ defined by

$$(\beta \lhd \alpha)(t) := \beta(t \lhd \alpha)\,, \quad \text{for } t \in \mathbb{F}_\varnothing\Sigma\,.$$ ⌟

Thus, to simulate the cascading operation on languages, we can first transform a given input forest $t \in \mathbb{F}_\varnothing\Sigma$ to the forest $t \lhd \alpha$ and then compute its image under β. In other words, we can use the cascade composition $\beta \lhd \alpha$ to 'recognise' a language. But note that these notions are only defined for forests of sort $\varnothing$. In particular, $\beta \lhd \alpha$ is no morphism. To handle higher arities, we need to set up a bit of algebraic machinery.

The problem with generalising the labelling operation $t \lhd \alpha$ 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 $t \lhd \alpha$ every label from $A_\varnothing$ by a function $A^\Delta \to A_\varnothing$ 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

$$C := A \times B^{A^\Delta_\varnothing}$$

and the following product. Let $p : C \to A$ and $q : C \to B^{A^\Delta_\varnothing}$ 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 $\lambda : A^{\xi}_{\varnothing} \to \mathbb{F}_{\varnothing}(A^{\Delta}_{\varnothing})$ and $\rho : A^{\xi}_{\varnothing} \to \mathbb{F}B$ by

$$\begin{aligned}
&\rho(\bar{a})(v) := q(t(v))\big(\lambda(\bar{a})(v)\big)\,,\\
&\lambda(\bar{a})(v) := (\pi(s_x))_{x \in \xi}\,, \quad \text{where} \quad s_x := (\mathbb{F}p(t) \Leftarrow \bar{a}) \downarrow^x_v \ .
\end{aligned}$$

Then we set

$$\pi(t) := \big\langle \pi(\mathbb{F}p(t)),\ \pi \circ \rho \big\rangle\,. \qquad \lrcorner$$

Remark. Note that λ is the unique function such that

$$\big((\mathbb{F}p(t) \Leftarrow \bar{a}) \lhd \pi\big)(v) = \big\langle p(t(v)),\ \lambda(\bar{a})(v) \big\rangle\,. \qquad \lrcorner$$

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*

$$\alpha : \mathbb{F}\Sigma \to \mathfrak{A} \quad \textit{and} \quad \beta : \mathbb{F}(\Sigma \times A^{\Delta}_{\varnothing}) \to \mathfrak{B}$$

and morphisms $\varphi : \mathbb{F}\Sigma \to \mathfrak{A} \circ \mathfrak{B}$. This bijection respects the equation

$$p \circ \varphi = \alpha$$

where $p : \mathfrak{A} \circ \mathfrak{B} \to \mathfrak{A}$ is the projection to the first component.

Proof. Given α and β, we define φ by

$$\varphi(\operatorname{sing}(c)) := \langle \alpha(\operatorname{sing}(c)),\ \bar{a} \mapsto \beta(\operatorname{sing}(\langle c, \bar{a} \rangle))\,, \quad \text{for } c \in \Sigma\,.$$

Since $\mathbb{F}\Sigma$ is freely generated by the elements of the form $\operatorname{sing}(c)$ this defines a unique morphism. Conversely, given φ we set

$$\begin{aligned}
\alpha(\operatorname{sing}(c)) &:= p(\varphi(\operatorname{sing}(c)))\\
\beta(\operatorname{sing}(\langle c, \bar{a} \rangle)) &:= q\big(\varphi(\operatorname{sing}(c))\big)(\bar{a})\,.
\end{aligned}$$

We claim that these two constructions are inverse to each other. To keep the notation simple, we identify $\operatorname{sing}(c)$ 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

$$\begin{aligned} \varphi'(c) &= \big\langle \alpha(c),\ \bar{a} \mapsto \beta(c,\bar{a}) \big\rangle \\ &= \big\langle p(\varphi(c)),\ \bar{a} \mapsto q(\varphi(c))(\bar{a}) \big\rangle \\ &= \big\langle p(\varphi(c)),\ q(\varphi(c)) \big\rangle \\ &= \varphi(c)\,. \end{aligned}$$

□

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

$$\pi(\operatorname{sing}(\langle c, f \rangle)) = \langle c, \pi \circ \rho \rangle$$

where

$$\begin{aligned} \mathbb{F}p(t) \Leftarrow \bar{a} &= c(\bar{a})\,, \\ \lambda(\bar{a})(\langle\rangle) &= \bar{a}\,, \\ \rho(\bar{a}) &= \operatorname{sing}(f(\bar{a}))\,. \end{aligned}$$

Hence, $\pi(\rho(\bar{a})) = \pi(\operatorname{sing}(f(\bar{a}))) = f(\bar{a})$ and $\pi \circ \rho = f$, as desired.

It remains to prove that $\pi(\operatorname{flat}(T)) = \pi(\mathbb{F}\pi(T))$, 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 $\pi(\mathrm{flat}(T))$, and λ_* and ρ_* those for the outer product of $\pi(\mathbb{F}\pi(T))$. Let $R : A^{\xi}_{\varnothing} \to \mathbb{F}_{\xi}\mathbb{F}B$ be the function defined by

$$R(\bar{a}) \simeq_{\mathrm{sh}} T \quad \text{and} \quad R(\bar{a})(v) := \rho_v(\lambda_*(\bar{a})(v)), \quad \text{for } \bar{a} \in A^{\xi}_{\varnothing}.$$

We claim that

$$\hat{\rho} = \mathrm{flat} \circ R \quad \text{and} \quad \mathbb{F}\pi \circ R = \rho_*.$$

It then follows that

$$\begin{aligned} q(\pi(\mathrm{flat}(T))) = \pi \circ \hat{\rho} &= \pi \circ \mathrm{flat} \circ R \\ &= \pi \circ \mathbb{F}\pi \circ R = \pi \circ \rho_* = q(\pi(\mathbb{F}\pi(T))), \end{aligned}$$

as desired.

Hence, it remains to prove the claim. First, note that, for a vertex $v \in \mathrm{dom}(T)$ and a tuple $\bar{a} \in A^{\xi}_{\varnothing}$, we have

$$\begin{aligned} \rho_*(\bar{a})(v) &= q\big(\pi(T(v))\big)\big(\lambda_*(\bar{a})(v)\big) \\ &= \pi\big(\rho_v(\lambda_*(\bar{a})(v))\big) = \pi(R(\bar{a})(v)), \end{aligned}$$

which proves the second equation. For the other one, consider two vertices $v \in \mathrm{dom}(T)$ and $x \in \mathrm{dom}(T(v))$, and let x' be the corresponding vertex of $\mathrm{flat}(T)$. Then

$$\begin{aligned} \mathrm{flat}(R(\bar{a}))(x') &= R(\bar{a})(v)(x) \\ &= \rho_v(\lambda_*(\bar{a})(v))(x) \\ &= q\big(T(v)(x)\big)\big(\lambda_v(\lambda_*(\bar{a})(v))\big) \\ &= q\big(\mathrm{flat}(T)(x')\big)\big(\lambda_v(\lambda_*(\bar{a})(v))\big) \end{aligned}$$

and

$$\hat{\rho}(\bar{a})(x') = q(\mathrm{flat}(T)(x'))\big(\hat{\lambda}(\bar{a})(x')\big).$$

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

$$\begin{aligned}
s_v^z &:= (\mathbb{F}p(T(v)) \Leftarrow \lambda_*(\bar{a})(v))\downarrow_x^z &&= \mathbb{F}p(r_v^z) \Leftarrow \lambda_*(\bar{a})(v)\,,\\
S_*^z &:= (\mathbb{F}p(\mathbb{F}\pi(T)) \Leftarrow \bar{a})\downarrow_v^z &&= \mathbb{F}p(\mathbb{F}\pi(R_*^z)) \Leftarrow \bar{a}\,,\\
\hat{s}^z &:= (\mathbb{F}p(\pi(\mathrm{flat}(T))) \Leftarrow \bar{a})\downarrow_{x'}^z &&= \mathbb{F}p(\hat{r}^z) \Leftarrow \bar{a}\,,
\end{aligned}$$

and

$$r_v^z := T(v)\downarrow_x^z\,, \quad R_*^z := T\downarrow_v^z\,, \quad \hat{r}^z := \mathrm{flat}(T)\downarrow_{x'}^z\,,$$

Note that

$$\hat{r}^z = r_v^z\big((\mathrm{flat}(R_*^z))_z\big)\,.$$

Furthermore, it follows that

$$\begin{aligned}
\pi(S_*^z) &= \pi\big(\mathbb{F}p(\mathbb{F}\pi(R_*^z)) \Leftarrow \bar{a}\big)\\
&= \pi\big(\mathbb{F}\pi\big(\mathbb{F}\mathbb{F}p(R_*^z) \Leftarrow \mathrm{sing}(\bar{a})\big)\big)\\
&= \pi\big(\mathrm{flat}\big(\mathbb{F}\mathbb{F}p(R_*^z) \Leftarrow \mathrm{sing}(\bar{a})\big)\big)\\
&= \pi\big(\mathrm{flat}(\mathbb{F}\mathbb{F}p(R_*^z)) \Leftarrow \bar{a}\big)\,,\\
&= \pi\big(\mathbb{F}p(\mathrm{flat}(R_*^z)) \Leftarrow \bar{a}\big)\,.
\end{aligned}$$

Consequently,

$$\begin{aligned}
\pi(s^z_v) &= \pi\big(\mathbb{F}p(r^z_v) \Leftarrow \lambda_*(\bar a)(v)\big) \\
&= \pi\Big(\mathbb{F}p(r^z_v) \Leftarrow \big(\pi(\mathbb{F}p(\mathrm{flat}(R^y_*)) \Leftarrow \bar a)\big)_y\Big) \\
&= \pi\big((\mathbb{F}p(r^z_v))\big((\mathbb{F}p(\mathrm{flat}(R^y_*)) \Leftarrow \bar a)_y\big)\big) \\
&= \pi\big((\mathbb{F}p(r^z_v))\big((\mathbb{F}p(\mathrm{flat}(R^y_*)))_y\big) \Leftarrow \bar a\big) \\
&= \pi\big(\mathbb{F}p\big(r^z_v\big((\mathrm{flat}(R^y_*))_y\big) \Leftarrow \bar a\big)\big) \\
&= \pi\big(\mathbb{F}p(\hat r^z) \Leftarrow \bar a\big) \\
&= \pi(\hat s^z)\,,
\end{aligned}$$

as desired. □

Unfortunately, in general wreath products do not preserve MSO-definable as, 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

$$D := \{\, b(\bar c) \mid \langle a, b\rangle \in C,\ \bar c \in A^{\Delta}_{\varnothing} \,\}\,.$$

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

$$\begin{aligned}
\pi(t) &= a \quad &&\text{iff} \quad t \vDash \varphi_a \quad &&\text{for } t \in \mathbb{F}p[C]\,, \\
\pi(t) &= b \quad &&\text{iff} \quad t \vDash \psi_b \quad &&\text{for } t \in \mathbb{F}D\,.
\end{aligned}$$

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 $\bigvee_d P_{\langle c,d\rangle} x$ (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

$$t \vDash \varphi'_a \quad \text{iff} \quad p(\pi(t)) = a\,.$$

For the second component, we proceed as follows. Given an input forest $t \in \mathbb{F}_\xi C$, our formula will guess two families $(R_{\bar a})_{\bar a \in A^\xi_\varnothing}$ and $(L_{\bar a})_{\bar a \in A^\xi_\varnothing}$ 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 $L_{\bar a}$. □

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 $\mathfrak{A} \times \mathfrak{B} \to \mathfrak{A} \circ \mathfrak{B}$.*

(b) *The projection $p : \mathfrak{A} \circ \mathfrak{B} \to \mathfrak{A}$ 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 $\mathfrak{A} \circ \mathfrak{B}$. □

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

$$D_\xi = (A_\xi \times B_\xi^{(A_\varnothing)^\xi}) \times C_\xi^{(A_\varnothing \times B_\varnothing)^\xi}$$

and $\quad E_\xi = A_\xi \times (B_\xi \times C_\xi^{(B_\varnothing)^\xi})^{(A_\varnothing)^\xi}$.

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

$$\begin{aligned}
p &: D \to A \times B^{A_\varnothing^\Delta} & \hat{p} &: E \to A \\
q &: D \to C^{(A_\varnothing \times B_\varnothing)^\Delta} & \hat{q} &: E \to (B \times C^{B_\varnothing^\Delta})^{A_\varnothing^\Delta} \\
p' &: A \times B^{A_\varnothing^\Delta} \to A & \hat{p}' &: B \times C^{B_\varnothing^\Delta} \to B \\
q' &: A \times B^{A_\varnothing^\Delta} \to B^{A_\varnothing^\Delta} & \hat{q}' &: B \times C^{B_\varnothing^\Delta} \to C^{B_\varnothing^\Delta}
\end{aligned}$$

we can recast the definition of φ into the form of three equations:

$$\begin{aligned}
p'(p(d)) &= \hat{p}(\varphi(d))\,, \\
q'(p(d))(\bar{a}) &= \hat{p}'\big[\hat{q}(\varphi(d))(\bar{a})\big]\,, \\
q(d)(\bar{a}, \bar{b}) &= \hat{q}'\big[\hat{q}(\varphi(d))(\bar{a})\big](\bar{b})\,,
\end{aligned}$$

for $d \in D$, $\bar{a} \in (A_\varnothing)^\xi$, and $\bar{b} \in (B_\varnothing)^\xi$.

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 ρ, λ, $\hat{\rho}$, $\hat{\lambda}$, 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

$$\begin{aligned} \mathbb{F}\hat{p}(\mathbb{F}\varphi(t)) \Lleftarrow \bar{a} &= \mathbb{F}p'(\mathbb{F}p(t)) \Lleftarrow \bar{a} , \\ \mathbb{F}p'(\mathbb{F}p(t)) \Lleftarrow \bar{a} &= \mathbb{F}\hat{p}(\mathbb{F}\varphi(t)) \Lleftarrow \bar{a} , \end{aligned}$$

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

$$\begin{aligned} (*_1) \quad \rho(\bar{a})(v) &= q'\big(p(t(v))\big)\big(\lambda(\bar{a})(v)\big) \\ &= q'\big(p(t(v))\big)\big(\hat{\lambda}(\bar{a})(v)\big) \\ &= \hat{p}'\big[\hat{q}\big(\varphi(t(v))\big)\big(\hat{\lambda}(\bar{a})(v)\big)\big] = \hat{p}'\big(\hat{\rho}(\bar{a})(v)\big) , \\ (*_2)\ \rho'(\bar{a}\bar{b})(v) &= q(t(v))\big(\lambda'(\bar{a}\bar{b})(v)\big) \\ &= \hat{q}'\big[\hat{q}\big(\varphi(t(v))\big)\big(p'(\lambda'(\bar{a}\bar{b})(v))\big)\big]\big(q'(\lambda'(\bar{a}\bar{b})(v))\big) \\ &= \hat{q}'\big[\hat{q}\big(\varphi(t(v))\big)\big(\hat{\lambda}(\bar{a})(v)\big)\big]\big(q'(\lambda'(\bar{a}\bar{b})(v))\big) \\ &= \hat{q}'[\hat{\rho}(\bar{a})(v)]\big(q'(\lambda'(\bar{a}\bar{b})(v))\big) . \end{aligned}$$

We further claim that

$(*_3)$ $\pi(\rho'(\bar{a}\bar{b})) = \hat{q}'\big(\pi(\hat{\rho}(\bar{a}))\big)(\bar{b})$.

It then follows by $(*_1)$ and $(*_3)$ that

$$\begin{aligned}
\hat{p}'\big[\hat{q}\big(\varphi(\pi(t))\big)(\bar{a})\big] &= q'\big(p(\pi(t))\big)(\bar{a}) \\
&= q'\big(\pi(\mathbb{F}p(t))\big)(\bar{a}) \\
&= \pi(\rho(\bar{a})) \\
&= \pi\big(\mathbb{F}\hat{p}'(\hat{\rho}(\bar{a}))\big) \\
&= \hat{p}'\big(\pi(\hat{\rho}(\bar{a}))\big) \\
&= \hat{p}'\big[\hat{q}\big(\pi(\mathbb{F}\varphi(t))\big)(\bar{a})\big], \\
\hat{q}'\big[\hat{q}\big(\varphi(\pi(t))\big)(\bar{a})\big](\bar{b}) &= q(\pi(t))(\bar{a},\bar{b}) \\
&= \pi(\rho'(\bar{a}\bar{b})) \\
&= \hat{q}'\big(\pi(\hat{\rho}(\bar{a}))\big)(\bar{b}) \\
&= \hat{q}'\big[\hat{q}\big(\pi(\mathbb{F}\varphi(t))\big)(\bar{a})\big](\bar{b}),
\end{aligned}$$

which concludes the proof.

Thus, it remains to prove the above claim. Suppose that

$$\lambda'(\bar{a}\bar{b})(v) = \big(\pi(s'_x)\big)_x \quad \text{and} \quad \hat{\lambda}(\bar{a})(v) = \big(\pi(\hat{s}_x)\big)_x,$$

and let $\hat{\rho}_*, \hat{\lambda}_*$ and $\tilde{\rho}^x_v, \tilde{\lambda}^x_v$ be the functions used to define the products $\pi(\hat{\rho}(\bar{a}))$ and $\pi((\mathbb{F}p(t) \Leftarrow \bar{a}\bar{b}) \downarrow^x_v)$, respectively. By definition, we have

$$\begin{aligned}
\pi(\hat{\rho}_*(\bar{b})) &= \hat{q}'\big(\pi(\hat{\rho}(\bar{a}))\big)(\bar{b}), \\
\pi(\tilde{\rho}^x_v(\langle\rangle)) &= q'\big(\pi\big((\mathbb{F}p(t) \Leftarrow \bar{a}\bar{b}) \downarrow^x_v\big)\big)(\langle\rangle).
\end{aligned}$$

Furthermore, $p' \circ p = \hat{p} \circ \varphi$ implies that

$$\begin{aligned}
\tilde{\lambda}^x_v(\langle\rangle)(u) &= \Big(\pi\Big(\big(\mathbb{F}p'\big((\mathbb{F}p(t) \Leftarrow \bar{a}\bar{b}) \downarrow^x_v\big) \Leftarrow \langle\rangle\big) \downarrow^y_u\Big)\Big)_y \\
&= \Big(\pi\Big(\big(\mathbb{F}p'\mathbb{F}p(t) \Leftarrow \bar{a}\big) \downarrow^x_v \downarrow^y_u\Big)\Big)_y \\
&= \Big(\pi\Big(\big(\mathbb{F}\hat{p}\mathbb{F}\varphi(t) \Leftarrow \bar{a}\big) \downarrow^y_{vxu}\Big)\Big)_y \\
&= \hat{\lambda}(\bar{a})\downarrow^x_v(u).
\end{aligned}$$

Below we will show that

$$\tilde{\rho}^x_v(\langle\rangle)(u) = \hat{p}'\big[(\hat{\rho}(\bar{a}) \Leftarrow \langle\bar{b}(\langle\rangle), \bar{c}\rangle) \downarrow^x_v (u)\big] .$$

Then it follows that

$$\begin{aligned} q'(\lambda'(\bar{a}\bar{b})(v)) &= \big(q'\big(\pi\big((\mathbb{F}p(t) \Leftarrow \bar{a}\bar{b}) \downarrow^x_v\big)\big)(\langle\rangle)\big)_x \\ &= \big(\pi(\tilde{\rho}^x_v(\langle\rangle))\big)_x \\ &= \big(\pi\big(\mathbb{F}\hat{p}'(\hat{\rho}(\bar{a})) \downarrow^x_v\big)\big)_x \\ &= \hat{\lambda}_*(\bar{b})(v) , \end{aligned}$$

which implies that

$$\begin{aligned} \rho'(\bar{a}\bar{b})(v) &= \hat{q}'[\hat{\rho}(\bar{a})(v)]\big(q'(\lambda'(\bar{a}\bar{b})(v))\big) \\ &= \hat{q}'[\hat{\rho}(\bar{a})(v)]\big(\hat{\lambda}_*(\bar{b})(v)\big) \\ &= \hat{\rho}_*(\bar{b})(v) . \end{aligned}$$

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

$$\begin{aligned} &\hat{q}\big(\varphi\big((t \Leftarrow \langle\bar{a}\bar{b}, \bar{c}'\rangle) \downarrow^x_v (u)\big)\big)\big(\hat{\lambda}(\bar{a}) \downarrow^x_v (u)\big) \\ &\quad = \big(\hat{\rho}(\bar{a}) \Leftarrow \langle\bar{b}(\langle\rangle), \bar{c}\rangle\big) \downarrow^x_v (u) , \end{aligned}$$

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

$$\begin{aligned} \tilde{\rho}^x_v(\langle\rangle)(u) &= q\big((\mathbb{F}p(t) \Leftarrow \bar{a}\bar{b}) \downarrow^x_v (u)\big)(\tilde{\lambda}^x_v(\langle\rangle)(u)) \\ &= q\big(p\big((t \Leftarrow \langle\bar{a}\bar{b}\rangle\bar{c}) \downarrow^x_v (u)\big)\big)(\hat{\lambda}^x_v(\bar{a}) \downarrow^x_v (u)) \\ &= \hat{p}'\big[\hat{q}\big(\varphi\big((t \Leftarrow \langle\bar{a}\bar{b}\rangle\bar{c}) \downarrow^x_v (u)\big)\big)(\hat{\lambda}^x_v(\bar{a}) \downarrow^x_v (u))\big] \\ &= \hat{p}'\big[\big(\hat{\rho}(\bar{a}) \Leftarrow \langle\bar{b}(\langle\rangle), \bar{c}\rangle\big) \downarrow^x_v (u)\big] , \end{aligned}$$

as desired. First, consider a vertex u such that $t(vxu) = z$ is a variable. Then

$$\begin{aligned}
&\hat{q}\big(\varphi\big((t \Leftarrow \langle \bar{a}\bar{b}, \bar{c}'\rangle) \downarrow^x_v (u)\big)\big)\big(\hat{\lambda}(\bar{a}) \downarrow^x_v (u)\big) \\
&= \hat{q}\big(\varphi(\langle a_z b_z, c'_z\rangle)\big)(\langle\rangle) \\
&= \langle b_z(\langle\rangle), c_z\rangle \\
&= \big(\hat{\rho}(\bar{a}) \Leftarrow \langle \bar{b}(\langle\rangle), \bar{c}\rangle\big) \downarrow^x_v (u)\,.
\end{aligned}$$

If vxu is not lablled by a variable, we have

$$\begin{aligned}
&\hat{q}\big(\varphi\big((t \Leftarrow \langle \bar{a}\bar{b}, \bar{c}'\rangle) \downarrow^x_v (u)\big)\big)\big(\hat{\lambda}(\bar{a}) \downarrow^x_v (u)\big) \\
&= \hat{q}\big(\varphi(t(vxu))\big)\big(\hat{\lambda}(\bar{a})(vxu)\big) \\
&= \hat{\rho}(\bar{a})(vxu) \\
&= \big(\hat{\rho}(\bar{a}) \Leftarrow \langle \bar{b}(\langle\rangle), \bar{c}\rangle\big) \downarrow^x_v (u)\,.
\end{aligned}$$ ☐

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*

$L \in \mathsf{Casc}(\mathcal{K})$ iff *L 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 $\mathsf{Casc}(\mathcal{K})$. 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 $L, \bar{K} \in \mathsf{Casc}(\mathcal{K})$. By inductive hypothesis, we can find iterated wreath products $\mathfrak{A}$ and $\bar{\mathfrak{B}}$ such that L is recognised by a morphism to $\mathfrak{A}$ and K_i by one to $\mathfrak{B}_i$. It follows that $L[\bar{K}]$ is recognised by a morphism

$$\varphi : \mathbb{F}\Sigma \to (\mathfrak{B}_0 \times \cdots \times \mathfrak{B}_{n-1}) \circ \mathfrak{A}\,.$$

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

$$\mathbb{F}\Sigma \to \mathfrak{B}_0 \circ \cdots \circ \mathfrak{B}_{n-1} \circ \mathfrak{A}$$

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 $\varphi : \mathbb{F}\Sigma \to \mathfrak{A} \circ \mathfrak{B}$, $P \subseteq A_\varnothing \times B_\varnothing$, $\mathfrak{B} \in \mathcal{C}$, and $\mathfrak{A}$ is an iterated wreath product. Let p and q be the two projections and set

$$K_a := \{\, s \mid p(\varphi(s)) = a \,\}\,, \qquad \text{for } a \in A\,.$$

By inductive hypothesis, we know that $K_a \in \mathsf{Casc}(\mathcal{K})$. By Lemma 4.4, there exists a morphism $\zeta : \mathbb{F}(\Sigma \times A_\varnothing) \to \mathfrak{B}$ such that

$$q(\varphi(s)) = b \quad \text{iff} \quad \zeta(s \lhd (p \circ \varphi)) = b\,, \quad \text{for } s \in \mathbb{F}_\varnothing\Sigma \text{ and } b \in B_\varnothing\,.$$

Each language $M_b := \zeta^{-1}(b)$ is recognised by $\mathfrak{B}$ and, hence, belongs to $\mathcal{K}$. Furthermore, for $s \in \mathbb{F}_\varnothing\Sigma$, we have

$$\begin{aligned}
s \in L \quad &\text{iff} \quad \varphi(s) \in P \\
&\text{iff} \quad \text{there is some } \langle a, b\rangle \in P \text{ such that} \\
&\qquad\qquad p(\varphi(s)) = a \quad \text{and} \quad q(\varphi(s)) = b \\
&\text{iff} \quad \text{there is some } \langle a, b\rangle \in P \text{ such that} \\
&\qquad\qquad s \in K_a \quad \text{and} \quad \zeta(s \lhd (p \circ \varphi)) = b \\
&\text{iff} \quad \text{there is some } \langle a, b\rangle \in P \text{ such that} \\
&\qquad\qquad s \in K_a \quad \text{and} \quad s \lhd (p \circ \varphi) \in M_b \\
&\text{iff} \quad s \in \bigcup_{\langle a,b\rangle \in P} \big[K_a \cap M_b[\bar{K}]\big]\,.
\end{aligned}$$

As $\mathsf{Casc}(\mathcal{K})$ is closed under boolean operations, it follows that the latter language belongs to $\mathsf{Casc}(\mathcal{K})$. $\square$

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 $\mathrm{TL}[\mathcal{K}]$.

Definition 5.1. Let $\mathfrak{A}$ be an $\mathbb{F}$-algebra.

(a) For a path w (starting at some root) of a forest $s \in \mathbb{F}_{\varnothing} A_{\{x\}}$, we denote by $\mathrm{pr}_w \in \mathbb{F}A_{\{x\}}$ the following path-like forest. If w is infinite, pr_w is simply the restriction of s to the domain

$$\mathrm{dom}(\mathrm{pr}_w) := \{ v \in \mathrm{dom}(s) \mid v \text{ lies on } w \}\,.$$

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 $\mathrm{pr}_w \in \mathbb{F}A_{\{x\}}$ and the corresponding word in $A^{\infty}_{\{x\}}$.

(b) $\mathfrak{A}$ is *distributive* if, for all forests $s, t \in \mathbb{F}A$,

$$\{ \pi(\mathrm{pr}_w) \mid w \text{ a path of } s \} = \{ \pi(\mathrm{pr}_w) \mid w \text{ a path of } t \}$$

implies $\pi(s) = \pi(t)$.

(b) $\mathfrak{A}$ is *+-distributive* if, for all forests $s, t \in \mathbb{F}A$,

$$\{ \pi(\mathrm{pr}_w) \mid w \text{ a finite path of } s \} = \{ \pi(\mathrm{pr}_w) \mid w \text{ a finite path of } t \}$$

implies $\pi(s) = \pi(t)$. ⌟

Remark. Note that the set $\{ \pi(\mathrm{pr}_w) \mid w \text{ a path of } s \}$ may contain both elements of arity 0 and of arity 1. ⌟

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 evaluates products in ω-semigroups, there exists an automaton performing this check. □

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

$$\{ \pi(\mathrm{pr}_w) \mid w \text{ a path of } s \} = \{ \pi(\mathrm{pr}_w) \mid w \text{ a path of } t \}.$$

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* $\mathsf{E}K$ *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* $\langle A_{\{x\}}, A_\varnothing, \cdot \rangle$ *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* $\mathsf{E}K$ *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* $\langle A_{\{x\}}, A_\varnothing, \cdot \rangle$ *belongs to* $\mathcal{S}$.

Proof. (a) ($\Rightarrow$) Let $\varphi : \mathbb{F}\Sigma \to \mathfrak{A}$ recognise L where $\mathfrak{A}$ is as above. For every $a \in A_\varnothing \cup A_{\{x\}}$, we set

$$H_a := \big\{ \{ \pi(\mathrm{pr}_w) \mid w \text{ path in } t \} \bigm| t \in \pi^{-1}(a) \big\},$$
$$K_a := \big\{ w \in \Sigma^\infty \bigm| \pi(\mathbb{F}\varphi(w)) = a \big\}$$

(where we identify words $w \in \Sigma^\infty$ with forest that have a single branch). It follows that

$$\begin{aligned} \pi(t) = a \quad &\text{iff} \quad \{ \pi(\mathrm{pr}_w) \mid w \text{ path in } t \} \in H_a \\ &\text{iff} \quad t \vDash \bigvee_{I \in H_a} \Big[\bigwedge_{c \in I} \mathsf{E}K_c \wedge \bigwedge_{c \notin I} \neg \mathsf{E}K_c \Big]. \end{aligned}$$

($\Leftarrow$) Suppose that L is definable by a formula of this form. Then L is regular and recognised by a morphism $\mathrm{syn}_L : \mathbb{F}\Sigma \to \mathrm{Syn}(L)$ to its syntactic algebra. It is sufficient to prove that $\mathrm{Syn}(L)$ 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

$$\{\, \pi(\mathrm{pr}_w) \mid w \text{ a path of } s \,\} = \{\, \pi(\mathrm{pr}_w) \mid w \text{ a path of } t \,\}.$$

As $\mathrm{syn}_L : \mathbb{F}\Sigma \to \mathrm{Syn}(L)$ is surjective, it has a right inverse $\iota : \mathrm{Syn}(L) \to \mathbb{F}\Sigma$. We set

$$s' := \mathrm{flat}(\mathbb{F}\iota(s)) \quad \text{and} \quad t' := \mathrm{flat}(\mathbb{F}\iota(t)).$$

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

$$p[s'] \in L \quad \Leftrightarrow \quad p[t'] \in L.$$

Thus $s' \sim_L t'$, which implies that $\mathrm{syn}_L(s') = \mathrm{syn}_L(t')$. Therefore, we have

$$\begin{aligned}
\pi(s) &= \pi\big(\mathrm{syn}_L(\iota(s))\big) \\
&= \mathrm{syn}_L\big(\mathbb{F}\mathrm{flat}(\iota(s))\big) \\
&= \mathrm{syn}_L(s') \\
&= \mathrm{syn}_L(t') \\
&= \mathrm{syn}_L\big(\mathbb{F}\mathrm{flat}(\iota(t))\big) \\
&= \pi\big(\mathrm{syn}_L(\iota(t))\big) \\
&= \pi(t).
\end{aligned}$$

To show that the vertical ω-semigroup belongs to $\mathcal{S}$ consider two ∞-words u, v of elements of $\langle A_{\{x\}}, A_\varnothing, \cdot \rangle$ that should be equal according to the equations for $\mathcal{S}$. Analogously to above, we can show that the words

$$u' := \mathrm{flat}(\mathbb{F}\iota(u)) \quad \text{and} \quad v' := \mathrm{flat}(\mathbb{F}\iota(v))$$

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 $\langle s', t' \rangle \in \ker \mathrm{syn}_L$. $\square$

With this theorem we are able to give characterisations for various logics of the form $\mathrm{TL}[\mathcal{S}^\infty]$ and $\mathrm{TL}[\mathcal{S}^+]$. To do so we will make use of the following types of $\mathbb{F}$-algebras.

We denote by $\mathfrak{U}_{(1)}$ the $\mathbb{F}$-algebra with two elements 0_ξ and 1_ξ, for each sort ξ, where the product is simply the maximum function (using the order $0_\xi < 1_\xi$):

$$\pi(s) := \max \{\, s(v) \mid v \in \mathrm{dom}(s) \,\}\,.$$

$\mathfrak{U}^*_{(2)}$ is the $\mathbb{F}$-algebra with domains

$$(U^*_{(2)})_\xi := \{\top_\xi\} \cup \wp(\xi)\,, \qquad \text{for } \xi \in \Xi\,.$$

To define the product of a forest $s \in \mathbb{F}U^*_{(2)}$, let us call a vertex $v \in \mathrm{dom}(s)$ *reachable* if, for each $u \prec v$ and every variable x with $ux \preceq v$, 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 $\mathfrak{U}^\omega_{(2)}$. 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	$\mathfrak{U}_{(1)}$
wCTL	$\mathfrak{U}^*_{(2)}$
CTL	$\mathfrak{U}^*_{(2)}$, $\mathfrak{U}^\omega_{(2)}$
wCTL*	+-distributive aperiodic algebras
CTL*	distributive aperiodic algebras
PDL	+-distributive algebras

Proof. Note that the only forest languages recognised by a morphism $\varphi : \mathbb{F}\Sigma \to \mathfrak{U}_{(1)}$ are the empty language $\emptyset$, the full language $\mathbb{F}\Sigma$, the language

$$\{\, s \in \mathbb{F}_\emptyset\Sigma \mid \text{there is some } v \in \mathrm{dom}(s) \text{ with } s(v) \in \varphi^{-1}(1_\emptyset) \,\},$$

and its complement. These languages are exactly the languages of the form

$$\{\, s \in \mathbb{F}_\emptyset\Sigma \mid \text{there is some path with labelling in } \Sigma^* C \,\}, \quad \text{for } C \subseteq \Sigma,$$

which, by Propositions 1.3 and 1.5, coincide with the EF-definable ones.

Similarly, a morphism $\varphi : \mathbb{F}\Sigma \to \mathfrak{U}^*_{(2)}$ 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 $\varphi : \mathbb{F}\Sigma \to \mathfrak{U}^{\omega}_{(2)}$, 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 $c\mathrm{TL}[\mathcal{K}]$, 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 $L \subseteq \mathbb{F}_{\varnothing}\Sigma$ is an *$\mathcal{S}$-path language* if it is a finite boolean combination of languages of the form

$$B_{k,K} := \left\{ t \in \mathbb{F}_{\varnothing}\Sigma \;\middle|\; t \text{ has at least } k \text{ paths whose label belongs to } K \right\},$$

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 $a \in S$. 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

$$a \sqsubseteq_{\mathrm{R}} a' \quad \text{iff} \quad a = a'b\,, \quad \text{for some } b \in (A_{\{x\}})^1.$$

(b) Let $\widetilde{\Psi} \subseteq A^*_{\{x\}}$ be the set of all finite words $\langle a_0, \dots, a_{n-1} \rangle$ satisfying

$$a_0 \sqsupset_{\mathrm{R}} a_0 a_1 \sqsupset_{\mathrm{R}} \cdots \sqsupset_{\mathrm{R}} a_0 a_1 \cdots a_{n-1} ,$$

and set $\Psi_\xi := \widetilde{\Psi} \times A_\xi$, for $\xi \in \Xi$.. The *character* $\chi(w)$ of a word $w \in A^\infty_{\{x\}}$ 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

$$a_0 \cdots a_{i-1} a_i \sqsubset_{\mathrm{R}} a_0 \cdots a_{i-1} .$$

If no such factor exists, we take for w_i the rest of the word.

The *relative character* $\chi(w/\sigma)$ of $w \in A^\infty_{\{x\}}$ 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

$$\chi(w/\sigma) = \langle b', a_{k+1}, \dots, a_{n-1}, c \rangle . \qquad \lrcorner$$

Definition 6.3. Let $\mathfrak{A}$ be an $\mathbb{F}$-algebra and $s \in \mathbb{F}_\varnothing A_{\{x\}}$ 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

$$\widetilde{\mathrm{pr}}_{w,\sigma} := \chi(\mathrm{pr}_w/\sigma) \in \mathbb{F}_{\{x\}} A_{\{x\}} .$$

(c) For a given constant $N < \omega$, we denote by

$$\Pi_N(s) : X \to [N+1]$$

the function mapping a character $\sigma \in \Psi$ to the number of paths w of s such that $\chi(\mathrm{pr}_w) = \sigma$, 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 $H \subseteq \mathrm{dom}(s)$, 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 $\langle A_\varnothing, +\rangle$ is commutative and aperiodic, and there exists a constant $N < \omega$ such that

$$\Pi_N(s) = \Pi_N(t) \quad \text{implies} \quad \pi(s) = \pi(t)\,, \qquad \text{for } s, t \in \mathbb{F}_\varnothing A_{\{x\}}\,.$$

(b) $\mathfrak{A}$ is a *$+$-path algebra* if it is an ω-path algebra satisfying the stronger condition that

$$\Pi_N^+(s) = \Pi_N^+(t) \quad \text{implies} \quad \pi(s) = \pi(t)\,, \qquad \text{for } s, t \in \mathbb{F}_\varnothing A_{\{x\}}\,. \qquad ⌟$$

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 $b, c, d \in A_\varnothing$.*

- $a(b) + a(c+d) = a(b+c) + a(d)\,,$
- $aa = a$ *implies* $a(a(b+c)) = a(a(b)+c)\,.$

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 $a \sqsubseteq_{\mathrm{R}} aa$. Hence, $\chi(aac) = \langle a, c\rangle = \chi(ac)$. □

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. □

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 $\langle \mathrm{Syn}_{\{x\}}(L), \mathrm{Syn}_{\varnothing}(L), \cdot \rangle$ of $\mathrm{Syn}(L)$ belongs to $\mathcal{S}$.*

(b) $\mathrm{Syn}(L)$ *is an ω-path algebra.*

(c) *If L is an $\mathcal{S}^+$-path language,* $\mathrm{Syn}(L)$ *is even a $+$-path algebra.*

Proof. We use the notation $\sim_K := \ker \mathrm{syn}_K$ 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 $\mathbb{F}_\varnothing\Sigma$ or in $\mathbb{F}_{\{x\}}\Sigma$. Then it follows that $\langle \mathrm{Syn}_{\{x\}}(L), \mathrm{Syn}_\varnothing(L), \cdot \rangle$ 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 $\mathbb{F}_\varnothing\Sigma$.

(b) We start with the horizontal monoid. Consider two elements $a, b \in \mathrm{Syn}_\varnothing(L)$ and choose trees $s \in \mathrm{syn}_L^{-1}(a)$ and $t \in \mathrm{syn}_L^{-1}(b)$. 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 $s, t \in \mathbb{F}_\varnothing \mathrm{Syn}_{\{x\}}(L)$ satisfy $\Pi_N(s) = \Pi_N(t)$ where $N := \max_i k_i$. As syn_L is surjective, there exists a function $\varphi : \mathrm{Syn}(L) \to \mathbb{F}_\varnothing \Sigma$ such that $\mathrm{syn}_L \circ \varphi = \mathrm{id}$. We set

$$S := \mathbb{F}\varphi(s) \quad \text{and} \quad T := \mathbb{F}\varphi(t)\,.$$

Then $\Pi_N(\mathrm{flat}(S)) = \Pi_N(\mathrm{flat}(T))$. Consequently, $\mathrm{flat}(S)$ and $\mathrm{flat}(T)$ have the same number of paths in K_i, for $i < n$, (up to the number N) and we have

$$p[\mathrm{flat}(S)] \in L \Leftrightarrow p[\mathrm{flat}(T)] \in L\,, \quad \text{for all contexts } p\,.$$

Thus $\mathrm{flat}(S) \sim_L \mathrm{flat}(T)$ and it follows that

$$\begin{aligned} \pi(s) &= \pi(\mathbb{F}\mathrm{syn}_L(S)) \\ &= \mathrm{syn}_L(\mathrm{flat}(S)) \\ &= \mathrm{syn}_L(\mathrm{flat}(T)) \\ &= \pi(\mathbb{F}\mathrm{syn}_L(T)) = \pi(t)\,. \end{aligned}$$

(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 $\Pi_N^+(\mathrm{flat}(S)) = \Pi_N^+(\mathrm{flat}(T))$. But since the languages K_i contains only finite words, this is sufficient to imply that

$$p[\mathrm{flat}(S)] \in L \Leftrightarrow p[\mathrm{flat}(T)] \in L\,, \quad \text{for all contexts } p\,. \qquad \square$$

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 $u+a=u$. □

Lemma 6.9. *Let $\mathfrak{A}$ be a finitary $\mathbb{F}$-algebra where the horizontal semigroup $\langle A_\varnothing,+\rangle$ is commutative and aperiodic and let $a \in (A_{\{x\}})^1$ be $\sqsubseteq_{\mathsf{R}}$-minimal. Then*

$$ac = ac'\,, \quad \textit{for all } c,c' \in A_\varnothing\,.$$

Proof. By Lemma 6.8, there exists an element $u \in A_\varnothing$ such that $u+c=u$, for all $c \in A_\varnothing$. The element $\hat{a} := a(x+u)$ satisfies $\hat{a} \sqsubseteq_{\mathsf{R}} a$ and

$$\hat{a}c = a(c+u) = au\,, \quad \text{for all } c \in A_\varnothing\,.$$

By $\sqsubseteq_{\mathsf{R}}$-minimality of a, there are elements b and d such that $a(x)=\hat{a}(b(x)+d)$. Consequently,

$$ac = \hat{a}(bc+d) = \hat{a}u = \hat{a}(bc'+d) = ac'\,, \quad \text{for all } c,c' \in A_\varnothing\,. \qquad \square$$

Lemma 6.10. *Let $\mathfrak{A}$ be a finitary ω-path algebra and $a \in A_{\{x\}}$ an element.*

$$aa = a \quad \textit{implies} \quad a(x) = a(x + \pi \times a(0)) .$$

Proof. Lemma 6.5 (b) implies that $a(x) = aa(x) = a(x + a(0))$. Iterating this equation, we obtain $a(x) = a(x + \pi \times a(0))$. □

Lemma 6.11. *Let $\mathfrak{A}$ be a finitary ω-path algebra and $a, b, b' \in (A_{\{x\}})^1$. Then*

$$a = abb' \quad \textit{implies} \quad a(x) = a(x + b(0)) .$$

Proof. Below we will prove that the element $\hat{b} := (bb')^\pi$ satisfies the following equations.

(1) $\pi \times \hat{b}(0) = \pi \times \hat{b}(0) + b(0)$

(2) $a(x) = a(x + \pi \times \hat{b}(0))$

Then it follows that

$$\begin{aligned} a(x) &= a(x + \pi \times \hat{b}(0)) \\ &= a\big(x + \pi \times \hat{b}(0) + b(0)\big) \\ &= a\big((x + b(0)) + \pi \times \hat{b}(0)\big) \\ &= a(x + b(0)) , \end{aligned}$$

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

$$\begin{aligned} \pi \times \hat{b}(0) &= \pi \times \hat{b}(0) + \pi \times \hat{b}(0) \\ &= \pi \times \big(\hat{b}(0) + \hat{b}(0)\big) \\ &= \pi \times \big(bb'\hat{b}(0) + bb'\hat{b}(0)\big) \\ &= \pi \times \big(b(b'\hat{b}(0) + b'\hat{b}(0)) + b(0)\big) \\ &= \pi \times b(b'\hat{b}(0) + b'\hat{b}(0)) + \pi \times b(0) , \end{aligned}$$

where the fifth step follows by Lemma 6.5 (b). Consequently,

$$\begin{aligned}\pi \times \hat{b}(0) + b(0) &= \pi \times b\big(b'\hat{b}(0) + b'\hat{b}(0)\big) + \pi \times b(0) + b(0)\\ &= \pi \times b\big(b'\hat{b}(0) + b'\hat{b}(0)\big) + \pi \times b(0)\\ &= \pi \times \hat{b}(0)\,.\end{aligned}$$

(2) Set $d := \hat{b}(x) + \pi \times \hat{b}(0)$. Then

$$\begin{aligned}d &= \hat{b}(x) + \pi \times \hat{b}(0)\\ &= \hat{b}(x) + \pi \times \hat{b}(0) + \pi \times \hat{b}(0)\\ &= d(x) + \pi \times \hat{b}(0)\end{aligned}$$

and it follows that

$$d^\pi = dd^\pi = dd^\pi(x) + \pi \times \hat{b}(0) = d^\pi(x) + \pi \times \hat{b}(0)\,.$$

Consequently, we have

$$\begin{aligned}a = abb' = a\hat{b} = a\hat{b}\hat{b} &\\ &= a\hat{b}(\hat{b}(x) + \pi \times \hat{b}(0))\\ &= a(\hat{b}(x) + \pi \times \hat{b}(0))\\ &= ad\\ &= ad^\pi d^\pi\\ &= a(d^\pi d^\pi(x) + \pi \times \hat{b}(0))\\ &= ad^\pi d^\pi(x + \pi \times \hat{b}(0))\\ &= a(x + \pi \times \hat{b}(0))\,,\end{aligned}$$

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 $a = ad^\pi d^\pi$. □

Lemma 6.12. *Let $\mathfrak{A}$ be a finitary ω-path algebra. Then*

$$a = a(b(x) + c) \quad \textit{implies} \quad a = ab = a(x + c)\,.$$

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

$$d(x) := \hat{b}(x) + \pi \times \hat{b}(0) + \pi \times c\,.$$

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,

$$\begin{aligned}
a &= a(b(x) + c) \\
&= a(\hat{b}(\hat{b}(x) + c) + c) \\
&= a(\hat{b}(\hat{b}(x) + c + \pi \times (\hat{b}(0) + c)) + c) \\
&= a(\hat{b}(x) + c + \pi \times \hat{b}(0) + \pi \times c) \\
&= a(\hat{b}(x) + \pi \times \hat{b}(0) + \pi \times c) \\
&= ad \\
&= ad^{\pi} \\
&= ad^{\pi}(x + c) = a(x + c)\,,
\end{aligned}$$

where the third step follows by Lemma 6.10. □

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 $\sigma \sqsubseteq_{\mathrm{R}} \sigma a$. 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

$$\Pi_N\big(\sigma(s + \textstyle\sum_{i<m} p_i(0))\big) = \Pi_N\big(\sigma(t + \textstyle\sum_{i<n} q_i(0))\big)\,,$$

and the products $\pi(p_i)$ and $\pi(q_i)$ are σ-insignificant.

(c) A forest $s \in \mathbb{F}_\varnothing A$ 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. ⌟

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

$$\Pi_N\big(\sigma(s + \textstyle\sum_{i<m} p_i(0))\big) = \Pi_N\big(\sigma(t + \textstyle\sum_{i<n} q_i(0))\big)\,,$$

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

$$\begin{aligned}
\sigma a(x) &= \sigma(a(x) + c_0(0)) \\
&= \sigma(a(x) + c_0(0) + c_1(0)) \\
&= \dots \\
&= \sigma(a(x) + c_0(0) + \dots + c_{m-1}(0)) \\
&= \sigma(b(x) + d_0(0) + \dots + d_{n-1}(0)) \\
&= \dots \\
&= \sigma(b(x))\,.
\end{aligned}$$

(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. □

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 $H_0 \subseteq H \subseteq \mathrm{dom}(s)$ where H_0 is finite, we set

$$s[H/H_0] := s' + \sum_{v \in H_0} \widetilde{\mathrm{pr}}_{v,\sigma}(\mathrm{sf}_v)\,,$$

where s' is the forest obtained from s by removing all subtrees attached to some vertex in H. ⌟

Lemma 6.16. *Let $s \in \mathbb{F}_{\{x\}}A_{\{x\}}$ be a forest and $\sigma \in \Psi$ a character. Then*

$$\Pi_N(\sigma s/H_0) = \Pi_N(\sigma s/H)$$

implies that

$$\Pi_N\big(\sigma s[H/H_0]\big) = \Pi_N\big(\sigma(s + \textstyle\sum_{v \in H_0} \widetilde{\mathrm{pr}}_{v,\sigma}(0))\big)\,.$$

Proof. Set $t := \sigma(s + \sum_v \widetilde{\mathrm{pr}}_{v,\sigma}(0))$ and $t' := \sigma(s[H/H_0])$. 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 $s[H/H_0]$, it follows that there are vertices $v \preceq u$ of s with $v \in H \setminus H_0$ and $\tau = \chi(\sigma s|_u)$. If $\Pi_N(H_0) < N$, this would mean that $\Pi_N(H) > \Pi_N(H_0)$, which contradicts our assumption on H_0. Hence, $\Pi_N(H_0) = N$ and we have

$$N \geq \Pi_N(t)(\tau) > \Pi_N(t')(\tau) \geq \Pi_N(H_0)(\tau) = N\,.$$

A contradiction. □

Proposition 6.17. *Let $\mathfrak{A}$ be a finitary ω-path algebra and $\sigma \in \Psi$. For every forest $s \in \mathbb{F}_\varnothing A_{\{x\}}$, there exists a σ-separated forest $t \in \mathbb{F}_\varnothing A_{\{x\}}$ such that $s \sim_\sigma t$.*

Proof. Let us call a vertex $v \in \operatorname{dom}(s)$ *insignificant* if the product $\pi(\mathrm{pr}_v)$ is σ-insignificant. The set I of all insignificant vertices forms an initial subset of $\operatorname{dom}(s)$. The *frontier* H of s consists of the minimal (in the forest ordering) vertices of $\operatorname{dom}(s) \setminus I$. We choose a finite subset $H_0 \subseteq H$ such that

$$\Pi_N(\sigma s/H_0) = \Pi_N(\sigma s/H)\,.$$

Set

$$t' := s[H/H_0] \quad \text{and} \quad s' := s + \textstyle\sum_{v \in H_0} \widetilde{\mathrm{pr}}_{v,\sigma}(0)\,.$$

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

$$t' = s[H/H_0] = r' + \sum_{v \in H_0} \widetilde{\mathrm{pr}}_{v,\sigma}(\mathrm{sf}_v)\,,$$

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

$$t := r + \sum_{v \in H_0} \widetilde{\mathrm{pr}}_{v,\sigma}(\mathrm{sf}_v)$$

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 $\tau \in A^*_{\{x\}} A_\varnothing$, exactly $\mu(\tau)$ infinite branches of r' with character τ. Furthermore, fix a finite initial subset $I \subseteq \operatorname{dom}(r')$ containing, for every finite character $\tau \in A^*_{\{x\}}$, at least $\mu(\tau)$ vertices v such that $\chi(\sigma \mathrm{pr}_v) = \tau$. Let r'' be the forest obtained from r' by deleting every vertex that does not belong to $I \cup \bigcup B$. Then

$$\Pi_N(\sigma r'') = \mu = \Pi_N(\sigma r')\,.$$

Finally, set

$$r := \sum_{\beta \in B} r''|_\beta + \sum_{\nu \in I} \mathrm{pr}_\nu(0)\,.$$

By construction, we have

$$\Pi_N(\sigma r)(\tau) \geq \mu(\tau) = \Pi_N(\sigma r'')(\tau)\,.$$

Set $\sigma := \Pi_N(\sigma r)$. Then

$$\begin{aligned}
&\Pi_N\big(\sigma\big(r'' + \textstyle\sum_\tau(\sigma(\eta) - \mu(\eta)) \times \pi(\eta)(0)\big)\big)(\tau) \\
&= \mu(\tau) + \sigma(\tau) - \mu(\tau) \\
&= \sigma(\tau) \\
&= \Pi_N(\sigma r)\,.
\end{aligned}$$

Hence, $r \sim_\sigma r'' \sim_\sigma r'$, as desired. □

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 $c, c' \in A_\varnothing$, 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 $\approx_\sigma$-class of c by $[c]_\sigma$.

(b) We denote by $\Psi^{(\sigma)}_{\{x\}}$ the subset of $\Psi_{\{x\}}$ consisting of all characters τ whose product $\pi(\tau)$ is σ-significant.

(c) The *σ-profile* $\mathrm{pf}(s)$ of a σ-separated forest $s = \sum_{i<m} q_i + \sum_{i<n} p_i(r_i)$ consists of two functions

$$\mathrm{pf}_0(s) : \Psi_\varnothing \to [N+1] \quad \text{and} \quad \mathrm{pf}_1(s) : \Psi^{(\sigma)}_{\{x\}} \to [N+1] \times A_\varnothing/{\approx_\sigma}\,,$$

where the first one maps a character $\tau \in \Psi_\varnothing$ to the number

$$|\{ i < m \mid \chi(q_i) = \tau \}|$$

and the second one maps a character $\tau \in \Psi^{(\sigma)}_{\{x\}}$ to the pair $\langle k, [c]_\sigma \rangle$ where

$$k := |\{ i < n \mid \chi(p_i) = \tau \}| \quad \text{and} \quad c := \sum \big\{ \pi(r_i) \bigm| \chi(p_i) = \tau \big\} . \qquad \lrcorner$$

Lemma 6.19. *If s and s' are σ-separated forests, then*

$$\mathrm{pf}(s) = \mathrm{pf}(s') \quad \textit{implies} \quad \pi(\sigma s) = \pi(\sigma s') .$$

Proof. Suppose that

$$\mathrm{pf}_0(s)(\tau) = l_\tau \quad \text{and} \quad \mathrm{pf}_1(s)(\tau) = \langle k_\tau, [c_\tau]_\tau \rangle .$$

Then

$$\begin{aligned}
\pi(s) &= b + \sum_{\tau \in \Psi_0} l_\tau \times \tau + \sum_{\tau \in \Psi^{(\sigma)}_1} \sum_{i < k_\tau} \tau d^i_\tau \\
&= b + \sum_{\tau \in \Psi_0} l_\tau \times \tau + \sum_{\tau \in \Psi^{(\sigma)}_1} \big(\tau \hat{d}_\tau + (k_\tau - 1) \times \tau(0) \big) , \\
\pi(s') &= b' + \sum_{\tau \in \Psi_0} l_\tau \times \tau + \sum_{\tau \in \Psi^{(\sigma)}_1} \sum_{i < k_\tau} \tau e^i_\tau \\
&= b' + \sum_{\tau \in \Psi_0} l_\tau \times \tau + \sum_{\tau \in \Psi^{(\sigma)}_1} \big(\tau \hat{e}_\tau + (k_\tau - 1) \times \tau(0) \big) ,
\end{aligned}$$

where b and b' are sums of σ-insignificant elements, $\hat{d}_\tau := \sum_i d^i_\tau$, $\hat{e}_\tau := \sum_i e^i_\tau$, 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 $\approx_\sigma$, it therefore follows that

$$\hat{d}_\tau \approx_\sigma \hat{e}_\tau \quad \text{implies} \quad \sigma(\tau \hat{d}_\tau + a) = \sigma(\tau \hat{e}_\tau + a) , \quad \text{for all } a \in A_\varnothing .$$

In particular for

$$a := \sum_{\tau \in \Psi_0} l_\tau \times \tau + \sum_{\tau \in \Psi_1} (k_\tau - 1) \times \tau(0)\,,$$

we have

$$\begin{aligned}
\pi(\sigma(s)) &= \sigma\Big(a + b + \sum_{\tau \in \Psi_1} \tau \hat{d}_\tau\Big) \\
&= \sigma\Big(a + \sum_{\tau \in \Psi_1} \tau \hat{d}_\tau\Big) \\
&= \sigma\Big(a + \sum_{\tau \in \Psi_1} \tau \hat{e}_\tau\Big) \\
&= \sigma\Big(a + b' + \sum_{\tau \in \Psi_1} \tau \hat{e}_\tau\Big) = \pi(\sigma(s'))\,,
\end{aligned}$$

where the second and forth step follow by Lemma 6.11. □

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 propositon.

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 $\langle A_{\{x\}}, A_\varnothing, \cdot\rangle$ belongs to $\mathcal{S}$. Let $\varphi : \mathbb{T}\Sigma \to \mathfrak{A}$ be a morphism of $\mathbb{F}$-algebras, $\sigma \in \Psi_{\{x\}}$ a character, and $c \in A_\varnothing$ 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 $\sqsubseteq_{\mathsf{R}}$-class of σ. If σ is $\sqsubseteq_{\mathsf{R}}$-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-varitety $\mathcal{S}$.

For the inductive step, suppose that σ is not $\sqsubseteq_{\mathrm{R}}$-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 $s \in \mathbb{F}A$,

$$\begin{aligned} s \in M_{\tau,k} \quad &\text{iff} \quad \mathrm{pf}_0(s)(\tau) = k\,, \\ s \in M_{\tau,\langle k,[c]_\sigma\rangle} \quad &\text{iff} \quad \mathrm{pf}_1(s)(\tau) = \langle k, [c]_\sigma\rangle\,. \end{aligned}$$

and, for arbitrary forests $s, s' \in \mathbb{F}_\varnothing A_{\{x\}}$, $s \sim_\sigma s'$ 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

$$\begin{aligned} L_{\sigma,c} = \bigcup \Big\{ (\mathbb{F}\varphi_0)^{-1}[M_\mu] \;\Big|\; &\mu \text{ a profile such that} \\ &\mathrm{pf}(s) = \mu \;\Rightarrow\; \pi(\sigma s) = c \Big\}\,, \end{aligned}$$

where $\varphi_0 := \varphi \circ \mathrm{sing}$ is the restriction of φ to the alphabet Σ and

$$M_\mu := \bigcap_{\tau \in \Psi_0} M_{\tau,\mu_0(\tau)} \cap \bigcap_{\tau \in \Psi^{(\sigma)}_{\{x\}}} M_{\tau,\mu_1(\tau)}\,.$$

To prove this equality, we start by noting that, for a σ-separated forest $s \in \mathbb{F}A$,

$$\begin{aligned} s \in M_\mu \quad &\text{iff} \quad s \in M_{\tau,\mu_0(\tau)} \text{ and } s \in M_{\tau,\mu_1(\tau)}\,, \quad \text{for all suitable } \tau\,, \\ &\text{iff} \quad \mathrm{pf}_0(s)(\tau) = \mu_0(\tau) \text{ and } \mathrm{pf}_1(s)(\tau) = \mu_1(\tau)\,, \\ &\qquad\qquad\qquad\qquad\qquad\qquad \text{for all suitable } \tau\,, \\ &\text{iff} \quad \mathrm{pf}(s) = \mu\,. \end{aligned}$$

Now consider a forest $t \in \mathbb{F}\Sigma$ and set $s := \mathbb{F}\varphi_0(t)$. 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

$$\begin{aligned} t \in L_{\sigma,c} \quad &\text{iff} \quad s \in M_\mu\,, \quad \text{for some } \mu \text{ such that} \\ &\qquad\qquad \mathrm{pf}(r) = \mu \;\Rightarrow\; \pi(\sigma r) = c \\ &\text{iff} \quad s' \in M_\mu\,, \quad \text{for some } \mu \text{ such that} \\ &\qquad\qquad \mathrm{pf}(r) = \mu \;\Rightarrow\; \pi(\sigma r) = c \\ &\text{iff} \quad \mathrm{pf}(r) = \mathrm{pf}(s') \;\Rightarrow\; \pi(\sigma r) = c \\ &\text{iff} \quad \pi(\sigma s') = c \\ &\text{iff} \quad \sigma(\varphi(t)) = c\,, \end{aligned}$$

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 $\langle A_{\{x\}}, A_\varnothing, \cdot\rangle \in \mathcal{S}$ 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,

$$M_{\tau,k} = B_{k,K_\tau} \smallsetminus B_{k+1,K_\tau}$$

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 $s, s' \in \mathbb{F}_\varnothing A_{\{x\}}$. Before defining $M_{\tau,\langle k,[c]_\sigma\rangle}$, we show that the languages

$$N_{[c]_\sigma} := \{\, s \in \mathbb{F}_\varnothing A \mid \pi(s) \approx_\sigma c \,\}$$

are $\mathcal{S}$-path languages. Let H be the set of all σ-significant elements in $A_{\{x\}}$. Note that

$$\begin{aligned} s \in N_{[c]_\sigma} \quad & \text{iff} \quad \pi(s) \approx_\sigma c \\ & \text{iff} \quad \sigma a(\pi(s)) = \sigma a(c)\,, \quad \text{for all } a \in H\,, \\ & \text{iff} \quad s \in L_{\sigma a, \sigma a(c)}\,, \quad \text{for all } a \in H\,, \\ & \text{iff} \quad s \in \bigcap_{a \in H} L_{\sigma a, \sigma a(c)}\,. \end{aligned}$$

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 $B_{l,K_\tau K}$ where

$$K_\tau := \chi^{-1}(\tau)\,, \quad \text{for } \tau \in \Psi^{(\sigma)}_{\{x\}}\,.$$

For a path p with $\chi(p) = \tau \in \Psi^{(\sigma)}_{\{x\}}$ it then follows

$$ps \in N'_{\tau,[c]_\sigma} \quad \text{iff} \quad \pi(s) \approx_\sigma c\,,$$

Hence, we can set

$$M_{\tau,\langle k,[c]_\sigma\rangle} := (B_{k,K_\tau} \smallsetminus B_{k+1,K_\tau}) \cap N'_{\tau,[c]_\sigma}\,.$$

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 $s, s' \in \mathbb{F}_\varnothing A_{\{x\}}$.

Given a σ-separated forest s, the first term of the above intersection expresses that $\mathrm{pf}(s)(\tau) = \langle k, [d]_\sigma\rangle$, 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, $s = s_0 + \sum_{i<k} p_i(r_i)$ with $\chi(p_i) = \tau$. Set

$$r := \sum_i r_i \quad \text{and} \quad s' := s_0 + p_0(r) + \sum_{0<i<k} p_i(0)\,.$$

Then $\Pi_N(s) = \Pi_N(s')$, which implies that

$$\begin{aligned} s \in N'_{\tau,[c]_\sigma} \quad & \text{iff} \quad s' \in N'_{\tau,[c]_\sigma} \\ & \text{iff} \quad p_0(r) \in N'_{\tau,[c]_\sigma} \\ & \text{iff} \quad \pi(r) \approx_\sigma c \\ & \text{iff} \quad \mathrm{pf}_1(s)(\tau) = \langle k, [c]_\sigma \rangle\,. \end{aligned}$$

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 $\mathrm{pf}_0(s)(\tau)$ and $M_{\tau,k}$ to only count finite paths. ☐

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 $\langle A_{\{x\}}, A_\varnothing, \cdot \rangle$ 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 $\langle A_{\{x\}}, A_\varnothing, \cdot \rangle$ 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 $\mathfrak{N}^n_{(1)}$ be the $\mathbb{F}$-algebra with domains

$$(N^n_{(1)})_\xi := [n+1]\,, \quad \text{for } \xi \in \Xi\,.$$

The product of a forest s is simply the sum (up to n) of all labels $s(v)$, $v \in \mathrm{dom}(s)$.

For $n < \omega$, let $\mathfrak{N}^{*,n}_{(2)}$ be the $\mathbb{F}$-algebra with domains

$$(N^{*,n}_{(2)})_\xi := [n+1] \times \wp(\xi)\,, \quad \text{for } \xi \in \Xi\,.$$

To define the product of a forest s, let us call a vertex $v \in \operatorname{dom}(s)$ *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 $\mathfrak{N}^{\omega,n}_{(2)}$ with the domains

$$(N^{\omega,n}_{(2)})_\xi := [n+1] \times [n+1] \times \wp(\xi)\,,$$

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 compontents 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	$\mathfrak{N}^{n}_{(1)}$, $n < \omega$
cwCTL	$\mathfrak{N}^{*,n}_{(2)}$, $n < \omega$
cCTL	$\mathfrak{N}^{*,n}_{(2)}$, $\mathfrak{N}^{\omega,n}_{(2)}$, $n < \omega$
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 [14], where the corresponding material was developed for finite forests. The equivalence between MPL = cCTL* has been proven by Moller and Rabinovitch [33], while the equivalences FO = WMPL = cwCTL* can be found in [26].

Part D.

Advanced Topics

VII. Power Sets

POWER-SET OPERATIONS PLAY AN IMPORTANT role in language theory. In this chapter we develop the general theory and we present some applications.

1. Power-Set Functors

Although there is no abstract notion of a (covariant) power-set functor, we will present a fairly general definition below. Before doing so, let us present the most important examples. For the category Set, there are two canonical choices.

Definition 1.1. (a) The *covariant power-set functor* $\mathrm{Pw} : \mathsf{Set}^\Xi \to \mathsf{Set}^\Xi$ is defined by

$$\mathrm{Pw}_\xi(A) := \wp(A_\xi)\,, \quad \text{for } \xi \in \Xi\,.$$

The image of a function $f : A \to B$ is

$$\mathrm{Pw}(f)(S) := f[S]\,, \quad \text{for } S \subseteq A\,.$$

For an infinite cardinal κ, we denote by Pw_κ the subfunctor consisting of all sets of size less than κ.

(b) Let κ be a regular cardinal. We denote the class of all cardinals by Cn. The functor $\mathrm{Lin}_\kappa : \mathsf{Set}^\Xi \to \mathsf{Set}^\Xi$ maps a set A to the set $\mathrm{Lin}_\kappa(A)$ whose component of sort ξ consists of all functions $\mu : A_\xi \to \mathrm{Cn}$ satisfying

$$\sum_{a \in A_\xi} \mu(a) < \kappa\,.$$

We regard such functions as multi-sets, where the cardinal $\mu(a)$ indicates how many times the element a belongs to the set. Frequently, we will write multi-sets as linear combinations $\sum_{i \in I} a_i$ or $\sum_{i \in I} k_i \cdot a_i$ where $|I| < \kappa$, $k_i < \kappa$, and $a_i \in A_\xi$. (The elements a_i are not required to be distinct.)

The image of a function $f : A \to B$ is the function

$$\mathrm{Lin}_\kappa(f)(\mu) := \mu' \quad \text{where} \quad \mu'(b) = \sum_{a \in f^{-1}(b)} \mu(a)\,,$$

or, in sum notation,

$$\mathrm{Lin}_\kappa(f)\Big(\sum_{i \in I} k_i \cdot a_i\Big) := \sum_{i \in I} k_i \cdot f(a_i)\,.$$

For $\kappa = \aleph_0$, we omit the subscript and simply write Lin. ⌟

It is straightforward to check that these two functors form monads.

Lemma 1.2. *The functor* $\mathrm{Pw} : \mathsf{Set}^\Xi \to \mathsf{Set}^\Xi$ *forms a monad with multiplication*

$$\mathrm{union} : \mathrm{Pw} \circ \mathrm{Pw} \Rightarrow \mathrm{Pw} : H \mapsto \bigcup H\,,$$

and unit morphism

$$\mathrm{pt} : \mathrm{Id} \Rightarrow \mathrm{Pw} : a \mapsto \{a\}\,.$$

Proof. We have

$$\begin{aligned}
\mathrm{union}(\mathrm{pt}(I)) &= \bigcup \{I\} = I\,, \\
\mathrm{union}(\mathrm{Pw}(\mathrm{pt})(I)) &= \bigcup \{\, \{a\} \mid a \in I \,\} = I\,, \\
\mathrm{union}(\mathrm{Pw}(\mathrm{union})(H)) &= \bigcup \Big\{ \bigcup I \;\Big|\; I \in H \Big\} \\
&= \Big\{ a \;\Big|\; a \in \bigcup I,\ I \in H \Big\} \\
&= \Big\{ a \;\Big|\; a \in s,\ s \in I,\ I \in H \Big\} \\
&= \bigcup \Big\{ s \;\Big|\; s \in I,\ I \in H \Big\} \\
&= \bigcup \bigcup H \\
&= \mathrm{union}(\mathrm{union}(H))\,.
\end{aligned}$$

□

Exercise 1.1. Prove that, for a regular cardinal κ, Pw_κ forms a submonad of Pw. ⌟

Lemma 1.3. *If κ is a regular cardinal, the functor $\mathrm{Lin}_\kappa : \mathsf{Set}^\Xi \to \mathsf{Set}^\Xi$ forms a monad with multiplication*

$$\mathrm{sum} : \mathrm{Lin}_\kappa \circ \mathrm{Lin}_\kappa \Rightarrow \mathrm{Lin}_\kappa : \sum_{i \in I} \sum_{j \in J_i} k_{ij} \cdot a_{ij} \mapsto \sum_{\langle i,j \rangle \in \sum_{i \in I} J_i} k_{ij} \cdot a_{ij} \,,$$

(or in compact notation $\mathrm{sum}(f)(a) = \sum_\mu \mu(a) \cdot f(\mu)$*) and unit morphism*

$$\mathrm{pt} : \mathrm{Id} \Rightarrow \mathrm{Lin}_\kappa : a \mapsto a \,.$$

Proof. First, note that $\mathrm{sum}(\mu) \in \mathrm{Lin}_\kappa(A)$ since

$$\begin{aligned}
\sum_{a \in A} \sum(\mu)(a) &= \sum_{a \in A} \sum_\nu \nu(a) \cdot \mu(\nu) \\
&= \sum_\nu \Big(\sum_{a \in A} \nu(a) \Big) \cdot \mu(\nu) \\
&= \sum_\nu \lambda_\nu \cdot \mu(\nu) \,, \quad \text{where } \lambda_\nu := \sum_{a \in A} \nu(a) < \kappa \,.
\end{aligned}$$

This is a sum of fewer than κ cardinals that are all less than κ. As κ is regular, it follows that the sum is also less than κ.

It remains to check the monad axioms.

$$\mathrm{sum}(\mathrm{pt}(\mu))(a) = \sum_\nu \nu(a) \cdot \mathrm{pt}(\mu)(\nu) = \mu(a) \,,$$

$$\begin{aligned}
\mathrm{sum}(\mathrm{Lin}_\kappa(\mathrm{pt})(\mu))(a) &= \sum_\nu \nu(a) \cdot \mathrm{Lin}_\kappa(\mathrm{pt})(\mu)(\nu) \\
&= \sum_\nu \nu(a) \cdot \sum_{b \in \mathrm{pt}^{-1}(\nu)} \mu(b) \\
&= \sum_b \mathrm{pt}(b)(a) \cdot \mu(b) \\
&= \mu(a) \,,
\end{aligned}$$

$$\begin{aligned}
\mathrm{sum}(\mathrm{Lin}_\kappa(\mathrm{sum})(\mu))(a) &= \sum_\nu \nu(a) \cdot \mathrm{Lin}_\kappa(\mathrm{sum})(\mu)(\nu) \\
&= \sum_\nu \nu(a) \cdot \sum_{\lambda \in \mathrm{sum}^{-1}(\nu)} \mu(\lambda) \\
&= \sum_\lambda \mathrm{sum}(\lambda)(a) \cdot \mu(\lambda) \\
&= \sum_\lambda \sum_\nu \nu(a) \cdot \lambda(\nu) \cdot \mu(\lambda) \\
&= \sum_\nu \nu(a) \cdot \sum_\lambda \lambda(\nu) \cdot \mu(\lambda) \\
&= \sum_\nu \nu(a) \cdot \mathrm{sum}(\mu)(\nu) \\
&= \mathrm{sum}(\mathrm{sum}(\mu))(a)\,.
\end{aligned}$$ □

Remark. A $\mathrm{Lin}_{\aleph_0}$-algebra is nothing more than an abelian semigroup. ⌟

Example. Another example is the monad FD producing finite probability distributions, which is defined by

$$\begin{aligned}\mathrm{FD}(A) := \Big\{ \delta : A \to [0,1] \;\Big|\; &\textstyle\sum_{a \in A} \delta(a) = 1 \text{ and there are only} \\ &\text{finitely many } a \in A \text{ with } \delta(a) \neq 0 \Big\}\,.\end{aligned}$$

The unit is given by

$$\mathrm{pt}(a)(b) := \begin{cases} 1 & \text{if } a = b\,, \\ 0 & \text{otherwise}\,, \end{cases}$$

and the multiplication by

$$\mathrm{sum}(\Delta)(a) := \sum_{\delta \in \mathrm{FD}(A)} \Delta(\delta) \cdot \delta(a)\,.$$ ⌟

For the category Pos, the two most common power-set functors take the following form.

Definition 1.4. Let $A \in \mathsf{Pos}^{\Xi}$.

(a) For $X \subseteq A$, we write

$$\Uparrow X := \{ a \in A \mid a \geq x \text{ for some } x \in X \},$$
$$\text{and } \Downarrow X := \{ a \in A \mid a \leq x \text{ for some } x \in X \}.$$

For single elements $x \in A$, we omit the braces and simply write $\Uparrow x$ and $\Downarrow x$.

(b) The *(upward) power set* $\mathrm{Up}A$ of A is the ordered set with domains

$$\mathrm{Up}_\xi A := \{ I \subseteq A_\xi \mid I \text{ is upwards closed} \}, \quad \text{for } \xi \in \Xi,$$

and ordering

$$I \leq J \quad :\text{iff} \quad I \supseteq J, \quad \text{for } I, J \in \mathrm{Up}_\xi A.$$

For a function $f : A \to B$, we define $\mathrm{Up} f : \mathrm{Up}A \to \mathrm{Up}B$ by

$$\mathrm{Up} f(I) := \Uparrow f[I], \quad \text{for } I \in \mathrm{Up}A.$$

Given an infinite cardinal κ, we denote by $\mathrm{Up}_\kappa(A) \subseteq \mathrm{Up}(A)$ the subset of all sets of the form $\Uparrow I$ with $|I| < \kappa$.

(c) The *(downward) power set* $\mathrm{Dn}A$ of A is the ordered set with domains

$$\mathrm{Dn}_\xi A := \{ I \subseteq A_\xi \mid I \text{ is downwards closed} \}, \quad \text{for } \xi \in \Xi,$$

and ordering

$$I \leq J \quad :\text{iff} \quad I \subseteq J, \quad \text{for } I, J \in \mathrm{Dn}_\xi A.$$

For a function $f : A \to B$, we define $\mathrm{Dn} f : \mathrm{Dn}A \to \mathrm{Dn}B$ by

$$\mathrm{Dn} f(I) := \Downarrow f[I], \quad \text{for } I \in \mathrm{Dn}A.$$

Given an infinite cardinal κ, we denote by $\mathrm{Dn}_\kappa(A) \subseteq \mathrm{Dn}(A)$ the subset of all sets of the form $\Downarrow I$ with $|I| < \kappa$. ⌟

In the following we will state and prove most results only for the functor Up. The case of Dn can be handled in exactly the same way. Again it is straightforward to check that Up forms a monad on Pos^{Ξ}.

Proposition 1.5. *The functor* Up : $\mathsf{Pos}^{\Xi} \to \mathsf{Pos}^{\Xi}$ *forms a monad where the multiplication*

$$\mathrm{union} : \mathrm{UpUp}(A) \to \mathrm{Up}(A) : X \mapsto \bigcup X$$

is given by taking the union and the singleton function

$$\mathrm{pt} : A \to \mathrm{Up}(A) : a \mapsto \Uparrow\{a\}$$

is given by the principal filter operation.

Proof. Note that, for $I \in \mathrm{Up}(A)$, we have $\Uparrow I = \{\, K \mid K \subseteq I \,\}$. Therefore,

$$\begin{aligned}
\mathrm{union}(\mathrm{pt}(I)) &= \bigcup \Uparrow\{I\} = I\,, \\
\mathrm{union}(\mathrm{Up}(\mathrm{pt})(I)) &= \bigcup \Uparrow\{\, \{a\} \mid a \in I \,\} = I\,, \\
\mathrm{union}(\mathrm{Up}(\mathrm{union})(H)) &= \bigcup \Uparrow\Big\{ \bigcup I \;\Big|\; I \in H \Big\} \\
&= \Big\{ a \;\Big|\; a \in \bigcup I,\ I \in H \Big\} \\
&= \Big\{ a \;\Big|\; a \in s,\ s \in I,\ I \in H \Big\} \\
&= \bigcup \Big\{ s \;\Big|\; s \in I,\ I \in H \Big\} \\
&= \bigcup \bigcup H \\
&= \mathrm{union}(\mathrm{union}(H))\,.
\end{aligned}$$

□

Remark. Note that Dn and Up are isomorphic as functors, but not as monads. The corresponding natural isomorphism $c : \mathrm{Dn} \Rightarrow \mathrm{Up}$ maps a set $I \in \mathrm{Dn}(A)$ to its complement $A \setminus I$. But c is not a morphism of monads since the complement of $\Downarrow a$ is not of the form $\Uparrow a$. ⌟

The algebras for the monad Up are easy to characterise.

Lemma 1.6. *A pair* $\langle A, \pi \rangle$ *forms an* Up*-algebra if, and only if,* A *is a complete partial order and* $\pi = \inf$.

Proof. ($\Leftarrow$) Clearly,

$$\begin{aligned}(\inf \circ \mathrm{pt})(a) &= \inf \{ b \mid b \geq a \} = a\,,\\ (\inf \circ \mathrm{Up}(\inf))(H) &= \inf \{ \inf(I) \mid I \in H \}\\ &= \inf \bigcup_{I \in H} I\\ &= (\inf \circ \mathrm{union})(H)\,.\end{aligned}$$

($\Rightarrow$) For $I \in \mathrm{Up}(A)$, let $a \in I$ and let c be a lower bound of I. Then

$$\begin{aligned} I \leq {\Uparrow}a &\quad \text{implies} \quad \pi(I) \leq \pi({\Uparrow}a) = \pi(\mathrm{pt}(a)) = a\,,\\ {\Uparrow}c \leq I &\quad \text{implies} \quad \pi(I) \geq \pi({\Uparrow}c) = \pi(\mathrm{pt}(c)) = c\,.\end{aligned}$$

Consequently, $\pi(I)$ is a lower bound of I that is greater or equal to every other lower bound. Thus, $\pi(I) = \inf I$. □

Example. For compact Hausdorff spaces, we have the *Vietoris monad* Vt defined by

$$\mathrm{Vt}(\mathfrak{X}) := \{ C \subseteq X \mid C \text{ closed} \}\,,$$

where the topology on $\mathrm{Vt}(\mathfrak{X})$ is generated by the open subbasis consisting of the sets

$$\begin{aligned} C^+ &:= \{ U \in \mathrm{Vt}(\mathfrak{X}) \mid U \cap C = \varnothing \}\,,\\ C^- &:= \{ U \in \mathrm{Vt}(\mathfrak{X}) \mid U \nsubseteq C \}\,,\end{aligned}$$

for $C \in \mathrm{Vt}(\mathfrak{X})$. Multiplication and unit are given by

$$\mathrm{union}(A) := \bigcup A \quad \text{and} \quad \mathrm{pt}(x) := \{x\}\,. \qquad \lrcorner$$

To be able to treat the above functors in a uniform way, we introduce the following general notion of a power-set functor. It is relatively straightforward to define a generalised notion of a contravariant power-set functor: we can fix a set $\varPhi$ of weights and use the hom-functor $\mathcal{C}(-, \varPhi)$. To get a covariant version of such a functor, we need additional assumptions. The simplest one is to assume that $\varPhi$ is equipped with the structure of a semiring.

Definition 1.7. Let $\mathcal{C}$ be a concrete category with forgetful functor $\mathbb{V} : \mathcal{C} \to \mathsf{Set}$ and let $\Phi \in \mathcal{C}$ be an object that is equipped with the structure of a commutative semiring. We assume in Φ that all infinite sums used below are defined and that the product of Φ distributes over all such sums. A monad $\langle \mathbb{P}, \mathrm{union}, \mathrm{pt} \rangle$ is a *generalised power-set monad* over Φ if

- we can identify every element $s \in \mathbb{P}A$ with a function $s : A \to \Phi$ (which we will do so tacitly from now on), that is, we have $\mathbb{V}\mathbb{P}A \subseteq \mathcal{C}(A, \Phi)$,
- the action of $\mathbb{P}$ on morphisms is given by

  $$\mathbb{P}f(s)(b) = \sum_{a \in A} s(a) \cdot (\mathrm{pt} \circ f)(a)(b)\,,$$

  for $f : A \to B$, $s \in \mathbb{P}A$, and $b \in B$,
- the product of $\mathbb{P}$ satisfies

  $$\mathrm{union}(s)(a) := \sum_{t \in \mathbb{P}A} s(t) \cdot t(a)\,, \quad \text{for } s \in \mathbb{P}\mathbb{P}A \text{ and } a \in A\,.$$

⌟

Example. There is a generalised power-set monad consisting only of the singletons. We set

$$\mathbb{P}A := \mathbb{J}\{\, \mathrm{pt}(a) \mid a \in A \,\}\,,$$

where

$$\mathrm{pt}(a)(b) := \begin{cases} 1 & \text{if } a = b\,, \\ 0 & \text{otherwise}\,. \end{cases}$$

The multiplication is defined by

$$\mathrm{union}(\mathrm{pt}(\mathrm{pt}(a))) := \mathrm{pt}(a)\,.$$

⌟

Examples. All of the monads introduced above can be understood as generalised power-set monads.

(a) For $\mathrm{Pw} : \mathsf{Set} \to \mathsf{Set}$, we can set use the semiring $\Phi := \{0, 1\}$ with maximum and minimum for, respectively, addition and multiplication. Then we obtain an isomorphism

$$\mathrm{Pw}(A) \cong \mathsf{Set}(A, \Phi)$$

that maps a subset $S \subseteq A$ to its characteristic function. Under this isomorphism the monad operations become

$$\begin{aligned}
\mathrm{union}(s)(a) &= \sum_{t \in \mathsf{Set}(A,\Phi)} s(t) \cdot t(a)\,, \\
\mathrm{pt}(a)(b) &= \begin{cases} 1 & \text{if } a = b\,, \\ 0 & \text{otherwise}\,. \end{cases}
\end{aligned}$$

For the requirement on $\mathbb{P}f$ in the definition above, note that

$$\begin{aligned}
\mathbb{P}f(s)(b) = 1 \quad &\text{iff} \quad b \in \{\, f(a) \mid s(a) = 1 \,\} \\
&\text{iff} \quad \text{there is some } a \in A \text{ with } s(a) = 1 \text{ and} \\
&\phantom{\text{iff}} \quad \mathrm{pt}(f(a))(b) = 1 \\
&\text{iff} \quad \sum_{a \in A} s(a) \cdot (\mathrm{pt} \circ f)(a)(b) = 1\,.
\end{aligned}$$

(b) For $\mathrm{Lin}_\kappa : \mathsf{Set} \to \mathsf{Set}$, we can set use the semiring $\Phi := \mathrm{Cn}_\kappa$ consisting of all cardinals less than or equal to κ with the usual cardinal arithmetic for addition and multiplication. (Infinite sums and products with a value greater than κ evaluate to κ.) Then we obtain

$$\mathrm{Lin}_\kappa(A) = \{\, s \in \mathsf{Set}(A, \Phi) \mid \textstyle\sum_{a \in A} s(a) < \kappa \,\}\,.$$

In this case, $\mathbb{P}$ is a proper subfunctor of $\mathsf{Set}(-, \Phi)$. According to Lemma 1.3, the monad operations are

$$\begin{aligned}
\mathrm{union}(s)(a) &= \sum_{t \in \mathbb{P}A} s(t) \cdot t(a)\,, \\
\mathrm{pt}(a)(b) &= \begin{cases} 1 & \text{if } a = b\,, \\ 0 & \text{otherwise}\,. \end{cases}
\end{aligned}$$

For the requirement on $\mathbb{P}f$ in the definition above, note that

$$\begin{aligned}\mathbb{P}f(s)(b) &= \sum_{a \in f^{-1}(b)} s(a) \\ &= \sum_{a \in A} s(a) \cdot \mathrm{pt}(f(a))(b) \\ &= \sum_{a \in A} s(a) \cdot (\mathrm{pt} \circ f)(a)(b) .\end{aligned}$$

(c) For Up : $\mathsf{Pos} \to \mathsf{Pos}$, we can set use the semiring $\Phi := \{0, 1\}$ with the natural ordering and with maximum and minimum for, respectively, addition and multiplication. Then we obtain an isomorphism

$$\mathrm{Up}(A) \cong \mathsf{Pos}(A, \Phi)^{\mathrm{op}}$$

that maps a subset $S \subseteq A$ to its characteristic function, where $\mathsf{Pos}(A, [2])^{\mathrm{op}}$ denotes the set $\mathsf{Pos}(A, [2])$ equipped with the *opposite* of the usual ordering. Under this isomorphism the monad operations become

$$\begin{aligned}\mathrm{union}(s)(a) &= \sum_{t \in \mathsf{Set}(A, \Phi)} s(t) \cdot t(a) , \\ \mathrm{pt}(a)(b) &= \begin{cases} 1 & \text{if } a \leq b , \\ 0 & \text{otherwise} . \end{cases}\end{aligned}$$

For the requirement on $\mathbb{P}f$ in the definition above, note that

$$\begin{aligned}\mathbb{P}f(s)(b) = 1 \quad &\text{iff} \quad b \geq f(a) \text{ for some } a \text{ with } s(a) = 1 \\ &\text{iff} \quad \text{there is some } a \in A \text{ with } s(a) = 1 \text{ and} \\ &\qquad\ \ \mathrm{pt}(f(a))(b) = 1 \\ &\text{iff} \quad \sum_{a \in A} s(a) \cdot (\mathrm{pt} \circ f)(a)(b) = 1 .\end{aligned}$$

(d) For Dn : $\mathsf{Pos} \to \mathsf{Pos}$, we can set use the semiring $\Phi := \{0, 1\}$ with the natural ordering and with minimum and maximum for, respectively, addition and multiplication. (Note that these are swapped, so that 0 is the neutral element for multiplication and 1 the one for addition.) Then we obtain an isomorphism

$$\mathsf{Dn}(A) \cong \mathsf{Pos}(A, \Phi)^{\mathrm{op}}$$

that maps a subset $s \subseteq A$ to the function

$$s(a) := \begin{cases} 0 & \text{if } a \in s\,, \\ 1 & \text{otherwise}\,. \end{cases}$$

Under this isomorphism the monad operations become

$$\mathrm{union}(s)(a) = \sum_{t \in \mathsf{Set}(A,\Phi)} s(t) \cdot t(a)\,,$$

$$\mathrm{pt}(a)(b) = \begin{cases} 0 & \text{if } a \geq b\,, \\ 1 & \text{otherwise}\,. \end{cases}$$

For the requirement on $\mathbb{P}f$ in the definition above, note that

$$\begin{aligned} \mathbb{P}f(s)(b) = 0 \quad & \text{iff} \quad b \leq f(a) \text{ for some } a \text{ with } s(a) = 0 \\ & \text{iff} \quad \text{there is some } a \in A \text{ with } s(a) = 0 \text{ and} \\ & \qquad\ \mathrm{pt}(f(a))(b) = 0 \\ & \text{iff} \quad \sum_{a \in A} s(a) \cdot (\mathrm{pt} \circ f)(a)(b) = 0\,. \end{aligned}$$

⌟

2. Linear Monads

The goal of this section is to determine whether or not there exists a distributive law between a given polynomial monad and one of the above power-set monads. We will isolate a property of a polynomial monad called *linearity* that characterises the existence of such a distributive law. Intuitively, linearity requires that the multiplication $\mathbb{M}\mathbb{M} \Rightarrow \mathbb{M}$ does not duplicate labels. Before we can give the formal definition, we need to take a look at the special form the multiplication morphism for a polynomial functor takes.

Remark. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a polynomial monad with $\mathbb{M}X = \sum_{i \in I} X^{\mathrm{dom}(i)}$. We have seen in Lemma I.3.4 that the composition $\mathbb{M} \circ \mathbb{M}$ is the polynomial functor given by

$$\mathbb{M}\mathbb{M}X = \sum_{i \in I} \sum_{g : \mathrm{dom}(i) \to I} X^{\sum_{v \in \mathrm{dom}(i)} \mathrm{dom}(g(v))}.$$

Thus $\mathbb{M}\mathbb{M}X = \sum_{j \in J} X^{E_j}$ where

$$J := \sum_{i \in I} I^{\mathrm{dom}(i)} \quad \text{and} \quad E_{\langle i, g \rangle} := \sum_{v \in \mathrm{dom}(i)} \mathrm{dom}(g(v)) .$$

Further, note that the identity functor Id is polynomial since

$$\mathrm{Id}(A) = \sum_{\xi \in \Xi} A^{1_\xi},$$

where 1_ξ is a set with a single element, which has sort ξ. Therefore, we can apply Proposition I.3.9 to the natural transformations $\mu : \mathbb{M}\mathbb{M} \Rightarrow \mathbb{M}$ and $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{M}$ and we obtain the corresponding morphisms

$$\begin{aligned} \langle f, (\varphi_j)_{j \in J} \rangle &: (\mathrm{dom}(i))_{i \in I} \to (E_j)_{j \in J} \\ \langle h, (\psi_\xi)_{\xi \in \Xi} \rangle &: (\mathrm{dom}(i))_{i \in I} \to (1_\xi)_{\xi \in \Xi} \end{aligned}$$

of $\Pi(\mathsf{Set})$. Thus,

$$\begin{aligned} &h : \Xi \to I , \qquad &&\psi_\xi : \mathrm{dom}(h(\xi)) \to 1_\xi , \quad &&\text{for } \xi \in \Xi , \\ &f : J \to I , &&\varphi_j : \mathrm{dom}(f(j)) \to E_j , &&\text{for } j \in J . \end{aligned}$$

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

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

⌟

Definition 2.1. Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a polynomial monad and let $\langle f, (\varphi_j)_{j \in J} \rangle$ and $\langle h, (\psi_\xi)_{\xi \in \Xi} \rangle$ be the functions corresponding to the natural transformations $\mu : \mathbb{M}\mathbb{M} \Rightarrow \mathbb{M}$ and $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{M}$ as above. We call $\langle \mathbb{M}, \mu, \varepsilon \rangle$ *linear* if, for all indices j and ξ, the maps φ_j and ψ_ξ are bijective. ⌟

Example. The monads $\mathbb{G}$, $\mathbb{R}$, $\mathbb{F}$, and $\mathbb{T}$ are linear since each vertex of $\mathrm{flat}(g)$ corresponds to exactly one vertex of exactly one component $g(v)$. The monads $\mathbb{F}^\times$ and $\mathbb{T}^\times$ on the other hand are not, since their multiplication duplicates labels: substituting $b(z)$ for x in $a(x,x)$ creates two copies of b. ⌟

To derive our distributive law for linear monads, we start with two technical lemmas. The first one works for arbitrary polynomial functors, the second one requires linearity.

Lemma 2.2. *Let* $\langle \mathbb{P}, \mathrm{union}, \mathrm{pt}\rangle$ *be a generalised power-set monad on* $\mathcal{D}$ *and* $\mathbb{M} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ *a polynomial functor of arity* κ. *Suppose that*

- Φ *is commutative, has products of size less than* κ, *and such products distribute over arbitrary sums, and*
- *for every* $A \in \mathcal{D}^\Xi$, *there exists a function* $\delta_A : \mathbb{MP}A \to \mathbb{PM}A$ *satisfying*

$$\delta_A(s)(t) = \begin{cases} \prod\limits_{v \in \mathrm{dom}(s)} \big(s(v)\big)(t(v)) & \text{if } s \simeq_{\mathrm{sh}} t\,, \\ 0 & \text{otherwise}\,, \end{cases}$$

$$\delta_A \circ \mathbb{M}\mathrm{pt} = \mathrm{pt}\,,$$

for $s \in \mathbb{MP}A$ *and* $t \in \mathbb{M}A$.

Then $\mathbb{M} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ *has an extension* $\widehat{\mathbb{M}}$ *to* $\mathsf{Free}(\mathbb{P})$ *such that*

$$\widehat{\mathbb{M}}f = \mathrm{union} \circ \mathbb{P}(\delta_B \circ \mathbb{M}(f \circ \mathrm{pt}))\,,$$

for every $\mathbb{P}$*-morphism* $f : \mathbb{P}A \to \mathbb{P}B$.

Proof. For simplicity we will work in this proof with the representation of $\mathsf{Free}(\mathbb{P})$ where we use the same objects as $\mathcal{D}^\Xi$ and morphisms of the form $A \to \mathbb{P}B$. Then the composition of two morphisms $f : A \to \mathbb{P}B$ and $g : B \to \mathbb{P}C$ is given by $\mathrm{union} \circ \mathbb{P}g \circ f : A \to \mathbb{P}C$. Using this convention, we define $\widehat{\mathbb{M}} : \mathsf{Free}(\mathbb{P}) \to \mathsf{Free}(\mathbb{P})$ by

$$\begin{aligned} \widehat{\mathbb{M}}A &:= \mathbb{M}A\,, && \text{for objects } A \in \mathcal{D}^\Xi, \\ \widehat{\mathbb{M}}f &:= \delta_B \circ \mathbb{M}f\,, && \text{for morphisms } f : A \to \mathbb{P}B\,. \end{aligned}$$

Let us start by showing that $\widehat{\mathbb{M}}$ extends $\mathbb{M}$. For objects A this is trivial since we have $\widehat{\mathbb{M}}A = \mathbb{M}A$ by definition. For morphisms, note that a function $f : A \to B$ of $\mathcal{D}^{\Xi}$ corresponds to the morphism $\mathrm{pt} \circ f : A \to \mathbb{P}B$ of $\mathsf{Free}(\mathbb{P})$, and that

$$\widehat{\mathbb{M}}(\mathrm{pt} \circ f) = \delta_B \circ \mathbb{M}(\mathrm{pt} \circ f) = \mathrm{pt} \circ \mathbb{M}f\,.$$

It remains to show that $\widehat{\mathbb{M}}$ is a functor. We start with two remarks. First, note that

$$\begin{aligned}
\widehat{\mathbb{M}}g(p)(t) &= \prod_{v \in \mathrm{dom}(p)} g(p(v))(t(v)) \\
&= \prod_{v \in \mathrm{dom}(p)} (\mathrm{union} \circ \mathrm{pt} \circ g)(p(v))(t(v)) \\
&= \prod_{v \in \mathrm{dom}(p)} \mathrm{union}\big((\mathrm{pt} \circ g)(p(v))\big)(t(v)) \\
&= \prod_{v \in \mathrm{dom}(p)} \sum_{q \in \mathbb{P}C} (\mathrm{pt} \circ g)(p(v))(q) \cdot q(t(v))\,.
\end{aligned}$$

Second, recall that the composition of $f : A \to \mathbb{P}B$ and $g : B \to \mathbb{P}C$ in $\mathsf{Free}(\mathbb{P})$ is given by $\mathrm{union} \circ \mathbb{P}g \circ f : A \to \mathbb{P}C$, and that, for $a \in A$ and $c \in C$, we have

$$\begin{aligned}
(\mathrm{union} \circ \mathbb{P}g \circ f)(a)(c) &= \sum_{t \in \mathbb{P}C} (\mathbb{P}g \circ f)(a)(t) \cdot t(c) \\
&= \sum_{t \in \mathbb{P}C} \Big[\sum_{b \in B} f(a)(b) \cdot (\mathrm{pt} \circ g)(b)(t)\Big] \cdot t(c) \\
&= \sum_{b \in B} f(a)(b) \cdot \Big[\sum_{t \in \mathbb{P}C} (\mathrm{pt} \circ g)(b)(t) \cdot t(c)\Big] \\
&= \sum_{b \in B} f(a)(b) \cdot \mathrm{union}((\mathrm{pt} \circ g)(b))(c) \\
&= \sum_{b \in B} f(a)(b) \cdot (\mathrm{union} \circ \mathrm{pt})(g(b))(c) \\
&= \sum_{b \in B} f(a)(b) \cdot g(b)(c)\,.
\end{aligned}$$

By distributivity and commutativitiy of the semiring Φ, it therefore follows that

$$
\begin{aligned}
&\widehat{\mathbb{M}}(\mathrm{union} \circ \mathbb{P}g \circ f)(s)(t) \\
&= (\delta_C \circ \mathrm{union} \circ \mathbb{P}g \circ f)(s)(t) \\
&= \prod_{v \in \mathrm{dom}(s)} (\mathrm{union} \circ \mathbb{P}g \circ f)(s(v))(t(v)) \\
&= \prod_{v \in \mathrm{dom}(s)} \sum_{q \in \mathbb{P}C} (\mathbb{P}g \circ f)(s(v))(q) \cdot q(t(v)) \\
&= \prod_{v \in \mathrm{dom}(s)} \sum_{q \in \mathbb{P}C} \sum_{b \in B} f(s(v))(b) \cdot (\mathrm{pt} \circ g)(b)(q) \cdot q(t(v)) \\
&= \sum_{\substack{p \in \mathbb{M}B \\ p \simeq_{\mathrm{sh}} s}} \prod_{v \in \mathrm{dom}(s)} \sum_{q \in \mathbb{P}C} f(s(v))(p(v)) \cdot (\mathrm{pt} \circ g)(p(v))(q) \cdot q(t(v)) \\
&= \sum_{\substack{p \in \mathbb{M}B \\ p \simeq_{\mathrm{sh}} s}} \prod_{v \in \mathrm{dom}(s)} f(s(v))(p(v)) \cdot \sum_{q \in \mathbb{P}C} (\mathrm{pt} \circ g)(p(v))(q) \cdot q(t(v)) \\
&= \sum_{\substack{p \in \mathbb{M}B \\ p \simeq_{\mathrm{sh}} s}} \Big[\prod_{v \in \mathrm{dom}(s)} f(s(v))(p(v)) \Big] \\
&\qquad\qquad \cdot \Big[\prod_{v \in \mathrm{dom}(p)} \sum_{q \in \mathbb{P}C} (\mathrm{pt} \circ g)(p(v))(q) \cdot q(t(v)) \Big] \\
&= \sum_{p \in \mathbb{M}B} \widehat{\mathbb{M}}f(s)(p) \cdot \widehat{\mathbb{M}}g(p)(t) \\
&= (\mathrm{union} \circ \mathbb{P}\widehat{\mathbb{M}}g \circ \widehat{\mathbb{M}}f)(s)(t)\,,
\end{aligned}
$$

where the last two steps follow from the above remarks. □

Lemma 2.3. *Let $\langle \mathbb{P}, \mathrm{union}, \mathrm{pt} \rangle$ be a generalised power-set monad on $\mathcal{D}$ and $\langle \mathbb{M}, \mu, \varepsilon \rangle$ a linear monad on $\mathcal{D}^{\Xi}$ that satisfies the assumptions of Lemma 2.2. Let $\widehat{\mathbb{M}}$ be the corresponding extension of $\mathbb{M}$ to $\mathsf{Free}(\mathbb{P})$, and let $\varphi : \mathbb{P}A \to \mathbb{P}B$ be a morphism of $\mathsf{Free}(\mathbb{P})$. Then*

$$\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\varepsilon = \mathbb{P}\varepsilon \circ \varphi \quad \text{and} \quad \widehat{\mathbb{M}}\varphi \circ \mathbb{P}\mu = \mathbb{P}\mu \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}\varphi\,.$$

Proof. Fix a morphism $\varphi : \mathbb{P}A \to \mathbb{P}B$ between free $\mathbb{P}$-algebras and set $\varphi_\circ := \varphi \circ \mathrm{pt} : A \to \mathbb{P}B$. By Lemma 2.2, we have

$$\begin{aligned}
\widehat{\mathbb{M}}\varphi \circ \mathrm{pt} &= \mathrm{union} \circ \mathbb{P}(\delta_B \circ \mathbb{M}(\varphi \circ \mathrm{pt})) \circ \mathrm{pt} \\
&= \mathrm{union} \circ \mathrm{pt} \circ \delta_B \circ \mathbb{M}(\varphi \circ \mathrm{pt}) \\
&= \delta_B \circ \mathbb{M}\varphi_\circ\,.
\end{aligned}$$

As $\mathbb{M}$ is linear, there is a unique vertex in $\mathrm{dom}(\varepsilon(a))$. We denote it by $*$.

(a) For $a \in A$ and $t \in \mathbb{M}B$, we have

$$\begin{aligned}
&(\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\varepsilon \circ \mathrm{pt})(a)(t) \\
&= (\widehat{\mathbb{M}}\varphi \circ \mathrm{pt} \circ \varepsilon)(a)(t) \\
&= (\delta \circ \mathbb{M}\varphi_\circ \circ \varepsilon)(a)(t) \\
&= (\delta \circ \varepsilon \circ \varphi_\circ)(a)(t) \\
&= \prod_{v \in \mathrm{dom}(\varepsilon(\varphi_\circ(a)))} \big(\varepsilon(\varphi_\circ(a))(v)\big)(t(v)) \\
&= \varepsilon(\varphi_\circ(a))(*)(t(*)) \\
&= (\varphi_\circ(a))(t(*)) \\
&= (\mathrm{union} \circ \mathbb{P}\mathrm{pt})(\varphi_\circ(a))(t(*)) \\
&= \sum_{s \in \mathbb{P}B} \mathbb{P}\mathrm{pt}(\varphi_\circ(a))(s) \cdot s(t(*)) \\
&= \sum_{s \in \mathbb{P}B} \sum_{b \in B} \varphi_\circ(a)(b) \cdot (\mathrm{pt} \circ \mathrm{pt})(b)(s) \cdot s(t(*)) \\
&= \sum_{b \in B} \varphi_\circ(a)(b) \cdot \sum_{s \in \mathbb{P}B} (\mathrm{pt} \circ \mathrm{pt})(b)(s) \cdot s(t(*)) \\
&= \sum_{b \in B} \varphi_\circ(a)(b) \cdot \mathrm{union}\big((\mathrm{pt} \circ \mathrm{pt})(b)\big)(t(*)) \\
&= \sum_{b \in B} \varphi_\circ(a)(b) \cdot \mathrm{pt}(b)(t(*)) \\
&= \sum_{b \in B} \varphi_\circ(a)(b) \cdot \mathrm{pt}\big(\varepsilon(b)(*)\big)(t(*)) \\
&= \sum_{b \in B} \varphi_\circ(a)(b) \cdot \prod_{v \in \mathrm{dom}(\varepsilon(b))} \mathrm{pt}\big(\varepsilon(b)(v)\big)(t(v))
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{b\in B} \varphi_\circ(a)(b) \cdot (\delta \circ \mathbb{M}\mathrm{pt} \circ \varepsilon)(b)(t) \\
&= \sum_{b\in B} \varphi_\circ(a)(b) \cdot (\mathrm{pt} \circ \varepsilon)(b)(t) \\
&= \mathbb{P}\varepsilon(\varphi_\circ(a))(t) \\
&= (\mathbb{P}\varepsilon \circ \varphi \circ \mathrm{pt})(a)(t)\,.
\end{aligned}
$$

Since $\mathbb{P}$-morphisms are determined by their restriction to $\operatorname{rng}\mathrm{pt}$, it follows that $\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\varepsilon = \mathbb{P}\varepsilon \circ \varphi$.

(b) For $s \in \mathbb{M}\mathbb{M}A$ and $t \in \mathbb{M}A$, we have

$$
\begin{aligned}
&(\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\mu \circ \mathrm{pt})(s)(t) \\
&= (\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\mu \circ \mathrm{pt})(s)(t) \\
&= (\widehat{\mathbb{M}}\varphi \circ \mathrm{pt} \circ \mu)(s)(t) \\
&= (\delta_B \circ \mathbb{M}(\varphi \circ \mathrm{pt}) \circ \mu)(s)(t) \\
&= (\delta_B \circ \mu \circ \mathbb{M}\mathbb{M}\varphi_\circ)(s)(t) \\
&= (\delta_B \circ \mathbb{M}\varphi_\circ \circ \mu)(s)(t) \\
&= (\mathbb{P}\mathrm{id} \circ \delta_B \circ \mathbb{M}\varphi_\circ \circ \mu)(s)(t) \\
&= \sum_{r\in\mathbb{M}A} (\delta \circ \mathbb{M}\varphi_\circ \circ \mu)(s)(r) \cdot \mathrm{pt}(r)(t) \\
&= \sum_{\substack{r\in\mathbb{M}A \\ r \simeq_{\mathrm{sh}} \mu(s)}} \Big[\prod_{w\in\mathrm{dom}(\mu(s))} \varphi_\circ\big(\mu(s)(w)\big)\big(r(w)\big)\Big] \cdot \mathrm{pt}(r)(t) \\
&= \sum_{\substack{r\in\mathbb{M}\mathbb{M}A \\ r \simeq_{\mathrm{sh}} s \\ r(v) \simeq_{\mathrm{sh}} s(v)}} \Big[\prod_{v\in\mathrm{dom}(s)} \prod_{u\in\mathrm{dom}(s(v))} \varphi_\circ\big(s(v)(u)\big)\big(r(v)(u)\big)\Big] \\
&\qquad \cdot (\mathrm{pt} \circ \mu)(r)(t) \\
&= \sum_{\substack{r\in\mathbb{M}\mathbb{M}A \\ r \simeq_{\mathrm{sh}} s \\ r(v) \simeq_{\mathrm{sh}} s(v)}} \Big[\prod_{v\in\mathrm{dom}(s)} \prod_{u\in\mathrm{dom}(s(v))} \big(\mathbb{M}\varphi_\circ(s(v))\big)(u)\big(r(v)(u)\big)\Big] \\
&\qquad \cdot (\mathrm{pt} \circ \mu)(r)(t)
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_{\substack{r \in \mathbb{M}\mathbb{M}A \\ r \simeq_{\mathrm{sh}} s \\ r(v) \simeq_{\mathrm{sh}} s(v)}} \Big[\prod_{v \in \mathrm{dom}(s)} \big((\delta_B \circ \mathbb{M}\varphi_\circ)(s(v))\big)(r(v)) \Big] \cdot (\mathrm{pt} \circ \mu)(r)(t) \\
&= \sum_{\substack{r \in \mathbb{M}\mathbb{M}A \\ r \simeq_{\mathrm{sh}} s \\ r(v) \simeq_{\mathrm{sh}} s(v)}} \Big[\prod_{v \in \mathrm{dom}(s)} \big((\mathbb{M}\delta_B \circ \mathbb{M}\mathbb{M}\varphi_\circ)(s)\big)(v)(r(v)) \Big] \\
&\qquad \cdot (\mathrm{pt} \circ \mu)(r)(t) \\
&= \sum_{r \in \mathbb{M}\mathbb{M}A} (\delta_{\mathbb{M}B} \circ \mathbb{M}\delta_B \circ \mathbb{M}\mathbb{M}\varphi_\circ)(s)(r) \cdot (\mathrm{pt} \circ \mu)(r)(t) \\
&= (\mathbb{P}\mu \circ \delta_{\mathbb{M}B} \circ \mathbb{M}\delta_B \circ \mathbb{M}\mathbb{M}\varphi_\circ)(s)(t) \\
&= (\mathbb{P}\mu \circ \delta_{\mathbb{M}B} \circ \mathbb{M}(\delta_B \circ \mathbb{M}(\varphi \circ \mathrm{pt})))(s)(t) \\
&= (\mathbb{P}\mu \circ \delta_{\mathbb{M}B} \circ \mathbb{M}(\widehat{\mathbb{M}}\varphi \circ \mathrm{pt}))(s)(t) \\
&= (\mathbb{P}\mu \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}\varphi \circ \mathrm{pt})(s)(t)\,,
\end{aligned}
$$

where in the 9-th step we use the fact that, by linearity of $\mathbb{M}$, there exists a label-preserving bijection between $\mathrm{dom}(\mu(s))$ and $\sum_v \mathrm{dom}(s(v))$. As above it follows that $\widehat{\mathbb{M}}\varphi \circ \mathbb{P}\mu = \mathbb{P}\mu \circ \widehat{\mathbb{M}}\widehat{\mathbb{M}}\varphi$. □

The central result of this section is the following existence statement of a distributive law.

Theorem 2.4. *Let $\langle \mathbb{P}, \mathrm{union}, \mathrm{pt} \rangle$ be a generalised power-set monad on $\mathcal{D}$ and $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a linear monad of arity κ. Suppose that*

- *Φ is commutative, has products of size less than κ, and such products distribute over arbitrary sums, and*
- *for every $A \in \mathcal{D}^{\Xi}$, there exists a function $\delta_A : \mathbb{M}\mathbb{P}A \to \mathbb{P}\mathbb{M}A$ satisfying*

$$
\delta_A(s)(t) = \begin{cases} \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)(t(v)) & \text{if } s \simeq_{\mathrm{sh}} t\,, \\ 0 & \text{otherwise}\,, \end{cases}
$$

$$
\delta_A \circ \mathbb{M}\mathrm{pt} = \mathrm{pt}\,,
$$

for $s \in \mathbb{M}\mathbb{P}A$ and $t \in \mathbb{M}A$.

Then the family of functions $\delta = (\delta_A)_A$ *forms a distributive law* $\mathbb{M} \circ \mathbb{P} \Rightarrow \mathbb{P} \circ \mathbb{M}$.

Proof. Let $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathsf{Free}(\mathbb{P})$ be the functor mapping a set A to the free algebra $\mathbb{P}A$ generated by it, and let $\widehat{\mathbb{M}}$ be the functor from Lemma 2.2. We claim that $\langle \widehat{\mathbb{M}}, \mathbb{F}\mu, \mathbb{F}\varepsilon \rangle$ is an extension of the monad $\langle \mathbb{M}, \mu, \varepsilon \rangle$ to $\mathsf{Free}(\mathbb{P})$. First, note that it follows by Lemma 2.3 that

$$\mathbb{F}\mu : \widehat{\mathbb{M}}\widehat{\mathbb{M}} \Rightarrow \widehat{\mathbb{M}} \quad \text{and} \quad \mathbb{F}\varepsilon : \mathrm{Id} \Rightarrow \widehat{\mathbb{M}}$$

are natural transformations. Hence, we only have to check the monad laws for $\langle \widehat{\mathbb{M}}, \mathbb{F}\mu, \mathbb{F}\varepsilon \rangle$.

$$\begin{aligned} \mathbb{F}\mu \circ \widehat{\mathbb{M}}\mathbb{F}\mu &= \mathbb{F}\mu \circ \mathbb{F}\mathbb{M}\mu \\ &= \mathbb{F}(\mu \circ \mathbb{M}\mu) = \mathbb{F}(\mu \circ \mu) = \mathbb{F}\mu \circ \mathbb{F}\mu\,, \\ \mathbb{F}\mu \circ \mathbb{F}\varepsilon &= \mathbb{F}(\mu \circ \varepsilon) = \mathrm{id}\,, \\ \mathbb{F}\mu \circ \widehat{\mathbb{M}}\mathbb{F}\varepsilon &= \mathbb{F}\mu \circ \mathbb{F}\mathbb{M}\varepsilon = \mathbb{F}(\mu \circ \mathbb{M}\varepsilon) = \mathrm{id}\,. \end{aligned}$$

Having found this extension $\widehat{\mathbb{M}}$, we can now use Theorem I.6.7 (and its proof) to obtain a distributive law $\mathbb{M}\mathbb{P} \Rightarrow \mathbb{P}\mathbb{M}$ consisting of the morphisms

$$\begin{aligned} \mathbb{V}\widehat{\mathbb{M}}\mathrm{union} \circ \mathrm{pt} &= \mathrm{union} \circ \mathbb{P}(\delta \circ \mathbb{M}(\mathrm{union} \circ \mathrm{pt})) \circ \mathrm{pt} \\ &= \mathrm{union} \circ \mathbb{P}\delta \circ \mathrm{pt} \\ &= \mathrm{union} \circ \mathrm{pt} \circ \delta \\ &= \delta\,, \end{aligned}$$

where $\mathbb{V} : \mathsf{Free}(\mathbb{P}) \to \mathcal{D}^{\Xi}$ is the forgetful functor. $\square$

Corollary 2.5. *Let* $\langle \mathbb{M}, \mu, \varepsilon \rangle$ *be a linear monad of arity* κ.

(a) *The family* $\mathrm{dist} = (\mathrm{dist}_A)_A$ *defined by*

$$\mathrm{dist}_A(t) := \{\, s \in \mathbb{M}A \mid s \in^{\mathbb{M}} t \,\}\,, \quad \textit{for } t \in \mathbb{M}\mathrm{Pw}(A)\,,$$

forms a distributive law $\mathbb{M} \circ \mathrm{Pw} \Rightarrow \mathrm{Pw} \circ \mathbb{M}$.

(b) *Let $\lambda \geq \kappa$ be a regular cardinal such that*

$$\xi < \lambda \quad \text{and} \quad \eta < \kappa \quad \text{implies} \quad \xi^\eta < \lambda\,.$$

Then the family $\mathrm{dist} = (\mathrm{dist}_A)_A$ *defined by*

$$\mathrm{dist}_A(s)(t) := \begin{cases} \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)\big(t(v)\big) & \text{if } s \simeq_{\mathrm{sh}} t\,, \\ 0 & \text{otherwise}\,, \end{cases}$$

for $s \in \mathbb{M}\mathrm{Lin}_\lambda(A)$ *and* $t \in \mathbb{M}A$, *forms a distributive law* $\mathbb{M} \circ \mathrm{Lin}_\lambda \Rightarrow \mathrm{Lin}_\lambda \circ \mathbb{M}$.

(c) *The family* $\mathrm{dist} = (\mathrm{dist}_A)_A$ *defined by*

$$\mathrm{dist}_A(t) := \{\, s \in \mathbb{M}A \mid s \in^{\mathbb{M}} t \,\}\,, \quad \text{for } t \in \mathbb{M}\mathrm{Up}(A)\,,$$

forms a distributive law $\mathbb{M} \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathbb{M}$.

(d) *The family* $\mathrm{dist} = (\mathrm{dist}_A)_A$ *defined by*

$$\mathrm{dist}_A(t) := \{\, s \in \mathbb{M}A \mid s \in^{\mathbb{M}} t \,\}\,, \quad \text{for } t \in \mathbb{M}\mathrm{Dn}(A)\,,$$

forms a distributive law $\mathbb{M} \circ \mathrm{Dn} \Rightarrow \mathrm{Dn} \circ \mathbb{M}$.

Proof. We have shown in the example on page 412 how to express the above power-set functors as a generalised power-set monad. It remains to check that these representations satisfy the assumptions of Theorem 2.4.

(a) For Pw we use the semiring $\Phi := \{0, 1\}$ with maximum and minimum for, respectively, addition and multiplication; and we encode a subset $S \subseteq A$ by its characteristic function $A \to \Phi$.

Clearly, Φ is commutative, it has arbitrary infinite sums and products, and these distribute. Furthermore, the category Set^{Ξ} obviously contains a function dist_A satisfying

$$\mathrm{dist}_A(s)(t) = \begin{cases} \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)\big(t(v)\big) & \text{if } s \simeq_{\mathrm{sh}} t\,, \\ 0 & \text{otherwise}\,. \end{cases}$$

Finally note that, for $s \simeq_{\mathrm{sh}} t$, we have

$$\begin{aligned}(\mathrm{dist}_A \circ \mathbb{M}\mathrm{pt})(s)(t) = 1 \quad &\text{iff} \quad \prod_{v \in \mathrm{dom}(s)} \mathrm{pt}(s(v))(t(v)) = 1 \\ &\text{iff} \quad s = t \\ &\text{iff} \quad \mathrm{pt}(s)(t) = 1\,.\end{aligned}$$

(b) For Lin_λ we use the semiring $\Phi := \mathrm{Cn}_\lambda$ with the usual cardinal arithmetic. This semiring is again commutative, it has arbitrary infinite sums and products, and these distribute. The category Set^Ξ obviously contains the functions dist_A but we have to check that $\mathrm{dist}_A(s) \in \mathrm{Lin}_\lambda(\mathbb{M}A)$, for all $s \in \mathbb{M}\mathrm{Lin}_\lambda(A)$. Hence, let $s \in \mathbb{M}\mathrm{Lin}_\lambda(A)$ and $t \in \mathbb{M}A$. As λ is regular, we have

$$\xi := \sup_{v \in \mathrm{dom}(s)} s(v)(t(v)) < \lambda \quad \text{and} \quad \eta := |\mathrm{dom}(s)| < \kappa\,,$$

which implies that

$$\mathrm{dist}_A(s)(t) = \prod_{v \in \mathrm{dom}(s)} s(v)(t(v)) \leq \xi^\eta < \lambda\,, \quad \text{for all } t \in \mathbb{M}A\,.$$

Furthermore, setting

$$P_s := \{\, a \in A \mid s(v)(a) > 0 \text{ for some } v \in \mathrm{dom}(t) \,\}\,,$$

we have

$$\mathrm{dist}_A(s)(t) > 0 \quad \text{implies} \quad t \simeq_{\mathrm{sh}} s \text{ and } t \in \mathbb{M}P_s\,.$$

Since $\mu := |P_s| < |\mathrm{dom}(s)| \cdot \lambda = \eta\lambda = \lambda$, there are at most

$$|P_s|^{|\mathrm{dom}(s)|} = \mu^\eta < \lambda$$

terms t with $\mathrm{dist}_A(s)(t) > 0$. By regularity of λ, it follows that

$$\sum_{t \in \mathbb{M}A} \mathrm{dist}_A(s)(t) < \lambda\,,$$

since this is a sum of less than λ cardinals of size less than λ.

It remains to check that $\mathrm{dist} \circ \mathbb{M}\mathrm{pt} = \mathrm{pt}$. Fixing $s, t \in \mathbb{M}A$ with $s \simeq_{\mathrm{sh}} t$, we have

$$\begin{aligned}(\mathrm{dist}_A \circ \mathbb{M}\mathrm{pt})(s)(t) &= \prod_{v \in \mathrm{dom}(s)} \mathrm{pt}(s(v))(t(v)) \\ &= \begin{cases} 1 & \text{if } s(v) = t(v) \text{ for all } v\,, \\ 0 & \text{otherwise}\,, \end{cases} \\ &= \mathrm{pt}(s)(t)\,.\end{aligned}$$

(c) For Up we use the semiring $\Phi := \{0, 1\}$ with maximum and minimum for, respectively, addition and multiplication; and we encode a subset $S \subseteq A$ by its characteristic function $A \to \Phi$. These characteristic functions are ordered by the opposite of the usual ordering.

The proof is analogous to that of (a), the only difference is that we have to show that dist_A is a morphism of Pos^{Ξ}, i.e., that it is monotone. Hence, consider elements $s, s' \in \mathbb{M}\mathrm{Up}(A)$ and $t, t' \in \mathbb{M}A$ with $s \simeq_{\mathrm{sh}} t$, $s \leq s'$, and $t \leq t'$. This means that

$$s(v)(a) \geq s'(v)(a)\,, \quad s(v)(a) \leq s'(v)(a')\,, \quad \text{and} \quad t(v) \leq t'(v)\,,$$

for $v \in \mathrm{dom}(s)$ and $a \leq a'$ in A. Consequently,

$$\begin{aligned}\mathrm{dist}_A(s)(t) &= \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)(t(v)) \\ &\geq \prod_{v \in \mathrm{dom}(s')} \big(s'(v)\big)(t(v)) = \mathrm{dist}_A(s')(t)\,, \\ \mathrm{dist}_A(s)(t) &= \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)(t(v)) \\ &\leq \prod_{v \in \mathrm{dom}(s)} \big(s(v)\big)(t'(v)) = \mathrm{dist}_A(s)(t')\,,\end{aligned}$$

which implies that $\mathrm{dist}_A(s) \leq \mathrm{dist}_A(s')$ and $\mathrm{dist}_A(s)(t) \leq \mathrm{dist}_A(s)(t')$.

(d) For Dn we use the semiring $\Phi := \{0, 1\}$ with minimum and maximum for, respectively, addition and multiplication; and we encode a subset $S \subseteq A$

by the function $A \to \Phi$ mapping the members of S to 0. These characteristic functions are ordered by the opposite of the usual ordering.

As in (c), we only have to show that dist_A is monotone. The argument is exactly the same as in (c). □

Remark. Let us note that cardinals λ as in (b) above do exist. If $\kappa = \aleph_0$, we can use, e.g., $\lambda := \aleph_0$. For arbitrary κ, we can take

$$\lambda := (2^{<\mu})^+, \quad \text{for any } \mu \geq \kappa\,,$$

since $\xi < \lambda$ and $\eta < \kappa \leq \mu$ implies that

$$\xi^\eta \leq (2^{<\mu})^\eta = 2^{<\mu} < \lambda\,.$$

($2^{<\mu}$ denotes the cardinal $\sup\{\, 2^\eta \mid \eta < \mu \,\}$.) For instance, for the monads $\mathbb{T}$ or $\mathbb{F}$ whose arity is $\kappa = \aleph_1$, we can use $\lambda := (2^{\aleph_0})^+$. ⌟

For the functor Up, we can strengthen Theorem 2.4 in two ways: (I) the distributive law dist is unique and (II) there is no distributive law for non-linear monads. We start with the former.

Theorem 2.6. *Let* $\mathbb{M}$ *be a polynomial monad on* Pos^Ξ *and* $\delta : \mathbb{M}\mathrm{Up} \Rightarrow \mathrm{Up}\mathbb{M}$ *a distributive law. Then* $\delta = \mathrm{dist}$.

Proof. ($\supseteq$) Since δ is monotone, we have

$$\begin{aligned}
\delta(t) &\leq \inf\{\, \delta(s) \mid s \geq t \,\} \\
&\leq \inf\big\{\, \delta(\mathbb{M}\mathrm{pt}(r)) \bigm| \mathbb{M}\mathrm{pt}(r) \geq t \,\big\} \\
&= \inf\big\{\, \mathrm{pt}(r) \bigm| \mathbb{M}\mathrm{pt}(r)(v) \geq t(v) \text{ for all } v \,\big\} \\
&= \inf\big\{\, \mathrm{pt}(r) \bigm| \mathrm{pt}(r(v)) \geq t(v) \text{ for all } v \,\big\} \\
&= \bigcup\big\{\, \mathrm{pt}(r) \bigm| \mathrm{pt}(r(v)) \subseteq t(v) \text{ for all } v \,\big\} \\
&= \bigcup\big\{\, \mathrm{pt}(r) \bigm| r(v) \in t(v) \text{ for all } v \,\big\} \\
&= \Uparrow\{\, r \mid r \in^{\mathbb{M}} t \,\} \\
&= \mathrm{dist}(t)\,.
\end{aligned}$$

($\subseteq$) Suppose that $s \in \delta(t)$ for $t \in \mathbb{M}\mathrm{Up}A$. To prove that $s \in \mathrm{dist}(t)$ it is sufficient to show that $s(v) \in t(v)$, for all v. Hence, fix $v \in \mathrm{dom}(t)$ and let $\theta : A \to [2]$ be the map with

$$\theta(a) := \begin{cases} 1 & \text{if } a \in t(v)\,, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathbb{M}\mathrm{Up}\theta(t)(v) = \mathrm{Up}\theta(t(v)) = \{1\}$. Since $[2]$ is well-ordered, we can find some $r \in \mathbb{M}[2]$ such that $\mathbb{M}\mathrm{Up}\theta(t) = \mathbb{M}\mathrm{pt}(r)$. It follows that

$$\mathrm{Up}\mathbb{M}\theta(\delta(t)) = \delta(\mathbb{M}\mathrm{Up}\theta(t)) = \delta(\mathbb{M}\mathrm{pt}(t)) = \mathrm{pt}(r)\,.$$

Consequently,

$$\theta(s(v)) = \mathbb{M}\theta(s)(v) \geq r(v) = 1 \quad \text{implies} \quad s(v) \in t(v)\,. \qquad \square$$

As a consequence, we obtain the following strengthening of Theorem 2.4.

Theorem 2.7. *Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a polynomial monad on Pos^{Ξ}. There exists a distributive law $\delta : \mathbb{M}\mathrm{Up} \Rightarrow \mathrm{Up}\mathbb{M}$ if, and only if, $\mathbb{M}$ is linear.*

Proof. ($\Leftarrow$) has already been proved in Theorem 2.4.

($\Rightarrow$) Suppose that $\mathbb{M}$ is not linear and let $\langle f, (\varphi_j)_{j \in J} \rangle$ and $\langle h, (\psi_\xi)_{\xi \in \Xi} \rangle$ be the functions corresponding to the natural transformations $\mu : \mathbb{M}\mathbb{M} \Rightarrow \mathbb{M}$ and $\varepsilon : \mathrm{Id} \Rightarrow \mathbb{M}$ as in the definition of linearity. By Theorem 2.6, it is sufficient to show that dist is not a distributive law. For a contradiction, suppose otherwise.

By assumption, there is some index j or ξ such that φ_j is not injective or ψ_ξ not bijective. First, assume that $\varphi_j : \mathrm{dom}(f(j)) \to E_j$ is not injective, for some index j. Then there are two positions $u, v \in \mathrm{dom}(f(j))$ with $\varphi_j(u) = \varphi_j(v)$. Set $w := \varphi_j(u)$, Let A be a set with at least two elements a and b of the same sort as these positions (and trivial ordering), and let $s \in \mathbb{M}\mathbb{M}\mathrm{Up}\mathbb{M}A$ be such that $\mathrm{dom}(s) = E_j$,

$$s(w) := \{\varepsilon(a), \varepsilon(b)\} \quad \text{and} \quad s(x) = \{\varepsilon(c_x)\}\,, \quad \text{for all } x \neq w\,.$$

By Theorem I.6.7, $\langle \mathrm{Up}\mathbb{M}A, \mathrm{Up}(\mu) \circ \mathrm{dist} \rangle$ is an $\mathbb{M}$-algebra with product $\pi := \mathrm{Up}(\mu) \circ \mathrm{dist}$. Note that

$$\begin{aligned}
&\{ \langle t(u), t(v) \rangle \mid t \in \mathrm{dist}(\mu(s)) \} \\
&= \{ \langle t(u), t(v) \rangle \mid t \in^{\mathbb{M}} \mu(s) \} \\
&= \{ \langle p, q \rangle \mid p \in \mu(s)(u),\ q \in \mu(s)(v) \} \\
&= \{ \langle p, q \rangle \mid p, q \in s(w) \} \\
&= \{ \langle \varepsilon(a), \varepsilon(a) \rangle,\ \langle \varepsilon(a), \varepsilon(b) \rangle,\ \langle \varepsilon(b), \varepsilon(a) \rangle,\ \langle \varepsilon(b), \varepsilon(b) \rangle \} .
\end{aligned}$$

Similarly,

$$\begin{aligned}
&\{ t(w) \mid t \in \mathrm{dist}(\mathbb{M}\pi(s)) \} \\
&= \{ t(w) \mid t \in^{\mathbb{M}} \mathbb{M}\pi(s) \} \\
&= \{ p \mid p \in \pi(s(w)) \} \\
&= \{ p \mid p \in \mathrm{Up}(\mu)(\mathrm{dist}(s(w))) \} \\
&= \{ \mu(\varepsilon(a)), \mu(\varepsilon(b)) \} \\
&= \{ a, b \} .
\end{aligned}$$

Since every $t \in \mathrm{dist}(\mu(s))$ is of the form $t = \mathbb{M}\varepsilon(t_0)$, for some $t_0 \in \mathbb{M}A$, it follows that

$$\begin{aligned}
&\{ \langle t(u), t(v) \rangle \mid t \in \mathrm{Up}(\mu)(\mathrm{dist}(\mu(s))) \} \\
&= \{ \langle \mu(t(u)), \mu(t(v)) \rangle \mid t \in \mathrm{dist}(\mu(s)) \} \\
&= \{ \langle a, a \rangle,\ \langle a, b \rangle,\ \langle b, a \rangle,\ \langle b, b \rangle \} .
\end{aligned}$$

But

$$\begin{aligned}
&\{ \langle t(u), t(v) \rangle \mid t \in \mathrm{Up}(\mu)(\mathrm{dist}(\mathbb{M}\pi(s))) \} \\
&= \{ \langle t(w), t(w) \rangle \mid t \in \mathrm{dist}(\mathbb{M}\pi(s)) \} \\
&= \{ \langle a, a \rangle,\ \langle b, b \rangle \} .
\end{aligned}$$

Thus $\pi(\mu(s)) \neq \pi(\mathbb{M}\pi(s))$. A contradiction.

It remains to consider the case where ψ_ξ is not bijective, for some ξ. Then, for every element a of sort ξ, the domain $D := \mathrm{dom}(\varepsilon(a))$ is either empty or of size at least 2. Let $A := \{a, b\}$ be a set with two elements of sort ξ and the trivial ordering. If D is empty, we set $s := \varepsilon(a)$ and $t := \varepsilon(b)$. Then

$$\mathrm{dom}(\varepsilon(s)) = \varnothing = \mathrm{dom}(\varepsilon(t)) \quad \text{implies} \quad \varepsilon(s) = \varepsilon(t)\,.$$

Hence, $s = \mu(\varepsilon(s)) = \mu(\varepsilon(t)) = t$. A contradiction.

Consequently, D must have at least two elements and $\varepsilon(a) : D \to \{a\}$ is the constant function with value a. Note that $A \in \mathrm{Up}A$ and

$$\begin{aligned}
\mathrm{Up}(\varepsilon)(A) &= \{\varepsilon(a), \varepsilon(b)\} \\
&= \{\, s \mid s : D \to \{a, b\} \text{ a constant function} \,\}\,, \\
\mathrm{dist}(\varepsilon(A)) &= \{\, s \mid s \in^{\mathbb{M}} \varepsilon(A) \,\} \\
&= \{\, s \mid s : D \to \{a, b\} \,\}\,.
\end{aligned}$$

As $|D| > 1$, there exist non-constant functions $D \to \{a, b\}$. This implies that $\mathrm{dist} \circ \varepsilon \neq \mathrm{Up}(\varepsilon)$, a violation of one of the axioms of a distributive law. □

3. Closure Under Projection

As a simple application of the distributive law for linear monads, let us prove that MSO-definable algebras are closed under the functor Dn, which one minor technical caveat: algebras of the form $\mathrm{Dn}(\mathfrak{A})$ need not be finitely generated. Hence, we have to take a finitely generated subalgebra instead.

Proposition 3.1. *If $\mathfrak{A}$ is an* MSO*-definable $\mathbb{T}$-algebra, then so is every finitely-generated subalgebra of* $\mathrm{Dn}(\mathfrak{A})$ *and* $\mathrm{Up}(\mathfrak{A})$.

Proof. Clearly, if A_ξ is finite, so is $\mathrm{Dn}(A_\xi)$. For the second condition, fix a finite set $C \subseteq \mathrm{Dn}(A)$ and an element $I \in \mathrm{Dn}(A)$. Then $D := \bigcup C$ is also finite and, as $\mathfrak{A}$ is MSO-definable, there exist MSO-formulae φ_a, for $a \in A$, such that

$$t \vDash \varphi_a \quad \text{iff} \quad \pi(t) \geq a\,, \quad \text{for every } t \in \mathbb{T}D\,.$$

Let ζ be the sort of I. For $s \in \mathbb{T}C$ it then follows that

$$\begin{aligned} \pi(s) \geq I \quad &\text{iff} \quad \text{for every } a \in I \text{ there exists } t \in^{\mathbb{T}} s \text{ with } \pi(t) \geq a \\ &\text{iff} \quad \text{for every } a \in I \text{ there exists } t \in^{\mathbb{T}} s \text{ with } t \vDash \varphi_a \,. \end{aligned}$$

Since the relation $\in^{\mathbb{T}}$ is MSO-definable, so is the above statement. □

Remark. With a suitable definition of MSO, this statement generalises to other linear monads, like $\mathbb{F}$, $\mathbb{G}$, $\mathbb{R}$, etc. ⌟

One reason why we are interest in closure under Dn is that this is what is needed to prove that the family of recognisable languages is closed under projections.

Proposition 3.2. *Let $\langle \mathbb{M}, \mu, \varepsilon \rangle$ be a linear monad on Pos^{Ξ} and $\mathcal{C}$ a class of $\mathbb{M}$-algebras such that*

$$\mathfrak{A} \in \mathcal{C} \quad \textit{implies} \quad \mathfrak{B} \in \mathcal{C}\,, \quad \textit{for every finitely-generated subalgebra } \mathfrak{B} \subseteq \mathrm{Dn}(\mathfrak{A})\,.$$

If $K \subseteq \mathbb{M}\Sigma$ is recognised by a morphism to an algebra in $\mathcal{C}$ then so is its image $\mathbb{M}f[K] \subseteq \mathbb{M}\Gamma$, for every $f : \Sigma \to \Gamma$.

Proof. We define the *pseudo-inverse* $g^- : B \to \mathrm{Dn}(A)$ of a map $g : A \to B$ by

$$g^-(b) := \{\, a \in A \mid g(a) \leq b \,\}\,.$$

It follows that

$$\mathrm{Dn}(g) \circ g^- = \mathrm{pt}\,.$$

Suppose that $K = \varphi^{-1}[P]$ for some morphism $\varphi : \mathbb{M}\Sigma \to \mathfrak{A}$ with $\mathfrak{A} \in \mathcal{C}$ and some upwards closed $P \subseteq A$. Note that

$$\begin{aligned} \mathrm{dist}(\mathbb{M}f^-(t)) &= \Uparrow\big\{\, s \;\big|\; s \in^{\mathbb{M}} \mathbb{M}f^-(t) \,\big\} \\ &= \Uparrow\big\{\, s \;\big|\; s \in \{\, r \mid f(r(v)) \geq t(v) \,\} \,\big\} \\ &= \Uparrow\{\, r \mid \mathbb{M}f(r) \geq t \,\} \\ &= (\mathbb{M}f)^-(t)\,. \end{aligned}$$

Hence,

$$\begin{aligned}
t \in \mathbb{M}f[K] \quad &\text{iff} \quad \text{there is some } s \in K \text{ with } t = \mathbb{M}f(s) \\
&\text{iff} \quad \text{there is some } s \in K \text{ with } t \geq \mathbb{M}f(s) \\
&\text{iff} \quad \text{there is some } s \in K \text{ with } s \in (\mathbb{M}f)^{-}(t) \\
&\text{iff} \quad (\mathbb{M}f)^{-}(t) \cap K \neq \varnothing \\
&\text{iff} \quad (\mathrm{Dn}(\varphi) \circ (\mathbb{M}f)^{-})(t) \cap P \neq \varnothing \\
&\text{iff} \quad (\mathrm{Dn}(\varphi) \circ \mathrm{dist} \circ \mathbb{M}f^{-})(t) \cap P \neq \varnothing \\
&\text{iff} \quad \psi(t) \in Q\,,
\end{aligned}$$

where $\psi : \mathbb{M}\Gamma \to \mathrm{Dn}(A)$ and $Q \subseteq \mathrm{Dn}(A)$ are given by

$$\psi := \mathrm{Dn}(\varphi) \circ \mathrm{dist} \circ \mathbb{M}f^{-} \quad \text{and} \quad Q := \{\, I \in \mathrm{Dn}(A) \mid I \cap P \neq \varnothing \,\}\,.$$

Note that Q is upwards closed. Let $\mathfrak{B} \subseteq \mathrm{Dn}(\mathfrak{A})$ be the subalgebra induced by $\operatorname{rng} \psi$. As $\mathbb{M}\Gamma$ is finitely generated, so is $\mathfrak{B}$. Hence, $\mathfrak{B} \in \mathcal{C}$.

To conclude the proof, it is therefore sufficient to show that $\psi : \mathbb{M}\Gamma \to \mathfrak{B}$ is a morphism of $\mathbb{M}$-algebras. As $\mathbb{M}f^{-}$ is such a morphism, we only need to consider $\psi_0 := \mathrm{Dn}(\varphi) \circ \mathrm{dist}$. By Corollary I.6.8 (b), the product of $\mathrm{Dn}(\mathfrak{A})$ is given by $\hat{\pi} := \mathrm{Dn}(\pi) \circ \mathrm{dist}$. Furthermore,

$$\begin{aligned}
\hat{\pi} \circ \mathbb{M}\psi_0 &= \mathrm{Dn}(\pi) \circ \mathrm{dist} \circ \mathbb{M}(\mathrm{Dn}(\varphi) \circ \mathrm{dist}) \\
&= \mathrm{Dn}(\pi) \circ \mathrm{Dn}(\mathbb{M}\varphi) \circ \mathrm{dist} \circ \mathbb{M}\mathrm{dist} \\
&= \mathrm{Dn}(\pi \circ \mathbb{M}\varphi) \circ \mathrm{dist} \circ \mathbb{M}\mathrm{dist} \\
&= \mathrm{Dn}(\varphi \circ \mu) \circ \mathrm{dist} \circ \mathbb{M}\mathrm{dist} \\
&= \mathrm{Dn}(\varphi) \circ \mathrm{Dn}(\mu) \circ \mathrm{dist} \circ \mathbb{M}\mathrm{dist} \\
&= \mathrm{Dn}(\varphi) \circ \mathrm{dist} \circ \mu \\
&= \psi_0 \circ \mu\,.
\end{aligned}$$

□

4. Non-Linear Trees

Unfortunately, the monad $\mathbb{T}^{\times} : \mathsf{Pos}^{\Xi} \to \mathsf{Pos}^{\Xi}$ is not linear. As we have seen in Theorem 2.7, this means that there does not exist a distributive law

$\mathbb{T}^\times \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathbb{T}^\times$. Nevertheless, for certain arguments involving infinite trees it would be very convenient to have such a law. In this section, we will prove a weaker property that can sometimes be used instead. We will show that we can lift Up to the class of free $\mathbb{T}^\times$-algebras.

We start with some technical remarks considering sorts. Below we will need to deal with trees with infinitely many different variables, that is, we have to work in the category Pos^{Ξ_+} instead of Pos^{Ξ} where $\Xi_+ := \wp(\omega)$. It is straightforward to extend the monads $\mathbb{R}$, $\mathbb{T}$, and $\mathbb{T}^\times$ to this more general setting. We will denote them by the same letters to keep notation simple.

The Action On The Variables

As noted above, the problem with finding a distributive law for $\mathbb{T}^\times$ is that this monad is not linear: its multiplication contains an unravelling operation gun which can duplicate arguments for variables appearing multiple times. To continue we need a variant of this operation that also modifies the variables of the given graph.

Definition 4.1. Let $g \in \mathbb{R}_\zeta A$ be a graph.

(a) For a surjective function $\sigma : \zeta \to \xi$, we denote by ${}^\sigma g \in \mathbb{R}_\xi A$ the graph obtained from g by replacing each variable x by $\sigma(x)$.

(b) We set

$$\mathrm{un}(g) := \langle \sigma, t \rangle ,$$

where t is the tree obtained from the unravelling $\mathrm{gun}(g)$ by renaming the variables so that each of them appears exactly once (note that this changes the sort) and σ is the function such that ${}^\sigma t = \mathrm{gun}(g)$. (To make this well-defined, we can fix a standard well-ordering on the domain, say, the length-lexicographic one, and we number the variables in increasing order with respect to this ordering, i.e., if $v_0 <_{\mathrm{llex}} v_1 <_{\mathrm{llex}} \cdots$ is an enumeration of all vertices labelled by a variable, we set $t(v_i) := x_i$, where $x_0, x_1, \ldots$ is some fixed sequence of variables.)

(c) We denote by $\mathbb{T}^\circ A$ the set of trees $t \in \mathbb{T}^\times A$ such that $\mathrm{un}(t) = \langle \mathrm{id}, t \rangle$. Let $\iota : \mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$ be the inclusion. (In actual calculations we will frequently omit ι to keep the notation simple.) ⌟

Remark. Note that the operation un can introduce infinitely many different variables. This is the reason why we have to work in $\mathsf{Pos}^{\Xi+}$. ⌟

Example. $\mathrm{un}(a(x,y,x)) = \langle \sigma, a(x_0,x_1,x_2)\rangle$ where σ maps x_0,x_1,x_2 to x,y,x. Then ${}^{\sigma}a(x_0,x_1,x_2) = a(x,y,z)$. ⌟

To make sense of the type of the above operations, we introduce the following monad where every element is annotated by some function renaming the variables.

Definition 4.2. (a) We define a functor $\mathbb{X} : \mathsf{Pos}^{\Xi+} \to \mathsf{Pos}^{\Xi+}$ as follows. For $A \in \mathsf{Pos}^{\Xi+}$, we set

$$\mathbb{X}_\xi A := \{\, \langle \sigma, a\rangle \mid a \in A_\zeta\,,\ \sigma : \zeta \to \xi \text{ surjective} \,\}\,.$$

We define the order on $\mathbb{X}_\xi A$ by

$$\langle \sigma, a\rangle \leq \langle \tau, b\rangle \quad :\text{iff} \quad \sigma = \tau \quad \text{and} \quad a \leq b\,.$$

For a morphism $f : A \to B$, we define $\mathbb{X}f : \mathbb{X}A \to \mathbb{X}B$ by

$$\mathbb{X}f(\langle \sigma, a\rangle) := \langle \sigma, f(a)\rangle\,.$$

(b) We define functions $\mathrm{comp} : \mathbb{X}\mathbb{X}A \to \mathbb{X}A$ and $\mathrm{in} : A \to \mathbb{X}A$ by

$$\mathrm{comp}(\langle \tau, \langle \sigma, a\rangle\rangle) := \langle \tau \circ \sigma, a\rangle \quad \text{and} \quad \mathrm{in}(a) := \langle \mathrm{id}, a\rangle\,.$$ ⌟

Lemma 4.3. *$\langle \mathbb{X}, \mathrm{comp}, \mathrm{in}\rangle$ and $\langle \mathbb{T}^\circ, \mathrm{flat}, \mathrm{sing}\rangle$ are monads.*

Exercise 4.1. Show that $\iota : \mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$ is the equaliser of $\mathrm{un} : \mathbb{T}^\times \Rightarrow \mathbb{X}\mathbb{T}^\circ$ and $\mathrm{in} : \mathbb{T}^\times \Rightarrow \mathbb{X}\mathbb{T}^\circ$ in the category of all endofunctors on $\mathcal{D}$. ⌟

The set $\mathbb{T}^\times A$ carries a canonical structure of an $\mathbb{X}$-algebra.

Definition 4.4. For $\langle \sigma, t\rangle \in \mathbb{X}\mathbb{T}^\times A$, we define the *reconstitution operation*

$$\mathrm{re}(\langle \sigma, t\rangle) := {}^{\sigma}t \in \mathbb{T}^\times A\,.$$

We denote its restriction to $\mathbb{X}\mathbb{T}^\circ$ by $\mathrm{re}_\circ := \mathrm{re} \circ \mathbb{X}\iota : \mathbb{X}\mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$. ⌟

The unravelling operation on trees can now be formalised using the following two natural transformations.

Lemma 4.5. *The inclusion morphism $\iota : \mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$ is a morphism of monads. The functions*

$$\mathrm{un} : \mathbb{T}^\times \Rightarrow \mathbb{X}\mathbb{T}^\circ, \quad \mathrm{re}_\circ : \mathbb{X}\mathbb{T}^\circ \Rightarrow \mathbb{T}^\times, \quad \textit{and} \quad \mathrm{re} : \mathbb{X}\mathbb{T}^\times \Rightarrow \mathbb{T}^\times$$

form natural transformations satisfying the following equations.

(a) $\mathrm{re}_\circ \circ \mathrm{un} = \mathrm{id}$

(b) $\mathrm{un} \circ \mathrm{re} = \mathrm{comp} \circ \mathbb{X}\mathrm{un}$

(c) $\mathrm{un} \circ \iota = \mathrm{in}$

(d) $\mathrm{re}_\circ \circ \mathrm{comp} = \mathrm{re} \circ \mathbb{X}\mathrm{re}_\circ$

(e) $\mathrm{flat}^\times \circ \mathrm{re}_\circ = \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota)$

(f) $\mathrm{re} \circ \mathrm{in} = \mathrm{id}$

(g) $\mathrm{un} \circ \mathrm{re}_\circ = \mathrm{id}$

Proof. The fact that ι is a morphism of monads is straightforward. To see that un is natural, it is sufficient to note that

$$\mathrm{un}(t) = \langle \sigma, s \rangle \quad \text{iff} \quad \mathrm{un}(\mathbb{T}^\times f(t)) = \langle \sigma, \mathbb{T}^\circ f(s) \rangle,$$

for every function $f : A \to B$. For re, we have

$$\begin{aligned} \mathbb{T}^\times f(\mathrm{re}(\langle \sigma, t \rangle)) &= \mathbb{T}^\times f({}^\sigma t) \\ &= {}^\sigma\big(\mathbb{T}^\times f(t)\big) \\ &= \mathrm{re}(\langle \sigma, \mathbb{T}^\times f(t) \rangle) = \mathrm{re}(\mathbb{X}\mathbb{T}^\times f(\langle \sigma, t \rangle)). \end{aligned}$$

Since $\mathrm{re}_\circ = \mathrm{re} \circ \mathbb{X}\iota$, this implies that $\mathrm{re}_\circ$ is natural as well.

(a) Note that $\mathrm{re}_\circ \circ \mathrm{un} = \mathrm{id}$ holds since

$$\mathrm{un}(t) = \langle \sigma, s \rangle \quad \text{implies} \quad {}^\sigma s = t, \quad \text{for trees } t \in \mathbb{T}^\times A.$$

(b) Suppose that $\mathrm{un}(t) = \langle \sigma, s \rangle$ and $\mathrm{un}({}^{\tau}t) = \langle \rho, r \rangle$. Then

$$ {}^{\tau\circ\sigma}s = {}^{\tau}t = {}^{\rho}r \,. $$

In particular, s and r only differ in the labelling of the variables. But $s, r \in \mathbb{T}^{\circ} A$ implies that the variables appear in the same order in both trees. Hence, $s = r$ and it follows that $\tau \circ \sigma = \rho$. Consequently,

$$ \begin{aligned} \mathrm{un}(\mathrm{re}(\langle \tau, t \rangle)) &= \langle \rho, r \rangle \\ &= \langle \tau \circ \sigma, s \rangle \\ &= \mathrm{comp}(\langle \tau, \langle \sigma, s \rangle \rangle) = \mathrm{comp}(\mathbb{X}\mathrm{un}(\langle \tau, t \rangle)) \,. \end{aligned} $$

(c)–(f) We have

$$ \begin{aligned} \mathrm{un}(\iota(t)) &= \langle \mathrm{id}, t \rangle = \mathrm{in}(t) \,, \\ \mathrm{re}_{\circ}\big(\mathrm{comp}\big(\langle \sigma, \langle \tau, t \rangle \rangle\big)\big) &= \mathrm{re}_{\circ}(\langle \sigma \circ \tau, t \rangle) \\ &= {}^{\sigma\circ\tau}\iota(t) \\ &= {}^{\sigma}({}^{\tau}\iota(t)) \\ &= {}^{\sigma}\mathrm{re}_{\circ}(\langle \tau, t \rangle) \\ &= \mathrm{re}\big(\big\langle \sigma, \mathrm{re}_{\circ}(\langle \tau, t \rangle)\big\rangle\big) \\ &= \mathrm{re}\big(\mathbb{X}\mathrm{re}_{\circ}\big(\langle \sigma, \langle \tau, t \rangle \rangle\big)\big) \,, \\ \mathrm{flat}^{\times}(\mathrm{re}_{\circ}(\langle \sigma, t \rangle)) &= \mathrm{flat}^{\times}({}^{\sigma}\iota(t)) \\ &= {}^{\sigma}(\mathrm{flat}^{\times} \circ \iota)(t) \\ &= \mathrm{re}\big(\big\langle \sigma, (\mathrm{flat}^{\times} \circ \iota)(t) \big\rangle\big) \\ &= (\mathrm{re} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \iota))(\langle \sigma, t \rangle) \,, \\ \mathrm{re}(\mathrm{in}(t)) &= \mathrm{re}(\langle \mathrm{id}, t \rangle) = {}^{\mathrm{id}}t = t \,. \end{aligned} $$

(g) By (c), we have

$$ \mathrm{un} \circ \mathrm{re}_{\circ} = \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}\iota = \mathrm{comp} \circ \mathbb{X}\mathrm{un} \circ \mathbb{X}\iota = \mathrm{comp} \circ \mathbb{X}\mathrm{in} = \mathrm{id} \,. \qquad \square $$

We can understand point (a) of this lemma as saying that $\mathbb{T}^\times$ is a retract of $\mathbb{X}\mathbb{T}^\circ$, but only as functors, not necessarily as monads. For the latter we first have to establish that $\mathbb{X}\mathbb{T}^\circ$ forms a monad and that the operations un and $\mathrm{re}_\circ$ are morphisms of monads.

Proposition 4.6.

(a) $\mathbb{X}\mathbb{T}^\circ$ *forms a monad with multiplication*

$$\mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) : \mathbb{X}\mathbb{T}^\circ\mathbb{X}\mathbb{T}^\circ \Rightarrow \mathbb{X}\mathbb{T}^\circ$$

and unit

$$\mathrm{in} \circ \mathrm{sing} : \mathrm{Id} \Rightarrow \mathbb{X}\mathbb{T}^\circ \,.$$

(b) $\mathrm{re}_\circ : \mathbb{X}\mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$ *and* $\mathrm{un} : \mathbb{T}^\times \Rightarrow \mathbb{X}\mathbb{T}^\circ$ *are isomorphisms of monads.*

(c) $\mathrm{in} : \mathbb{T}^\circ \Rightarrow \mathbb{X}\mathbb{T}^\circ$ *is an injective morphism of monads.*

Proof. (a), (b) By Lemma 4.5 (c), (e), and (a), we have

$$\begin{aligned}
\mathrm{re}_\circ \circ \mathrm{in} \circ \mathrm{sing} &= \mathrm{re}_\circ \circ \mathrm{un} \circ \iota \circ \mathrm{sing} = \iota \circ \mathrm{sing} = \mathrm{sing}^\times, \\
\mathrm{flat}^\times \circ \mathrm{re}_\circ \circ \mathbb{X}\mathbb{T}^\circ \mathrm{re}_\circ &= \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) \\
&= \mathrm{re}_\circ \circ \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) \,.
\end{aligned}$$

As $\mathrm{re}_\circ$ is a surjective natural transformation, most of the claim therefore follows by Lemma I.6.2. It only remains to check that un is also a morphism of monads. For this, note that by Lemma 4.5 (c), (a), and (e) we have

$$\begin{aligned}
\mathrm{in} \circ \mathrm{sing} &= \mathrm{un} \circ \iota \circ \mathrm{sing} = \mathrm{un} \circ \mathrm{sing}^\times \,, \\
\mathrm{un} \circ \mathrm{flat}^\times &= \mathrm{un} \circ \mathrm{flat}^\times \circ \mathrm{re}_\circ \circ \mathrm{un} \\
&= \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota) \circ \mathrm{un} \\
&= \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota) \circ \mathrm{un} \circ \mathbb{T}^\times(\mathrm{re}_\circ \circ \mathrm{un}) \\
&= \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) \circ \mathrm{un} \circ \mathbb{T}^\times \mathrm{un} \,.
\end{aligned}$$

(c) As un and ι are morphisms of monads, so is $\mathrm{un} \circ \iota = \mathrm{in}$. □

Corollary 4.7. $\mathbb{T}^\times \cong \mathbb{X}\mathbb{T}^\circ$ *(as monads)*

Proposition 4.8. *The functions*

$$\mathrm{dist}^{\mathbb{X}}(\langle\sigma, I\rangle) := \Uparrow\{\, \langle\sigma, a\rangle \mid a \in I \,\}$$

form a distributive law $\mathrm{dist}^{\mathbb{X}} : \mathbb{X} \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathbb{X}$.

Proof. Note that

$$\mathbb{X}_\zeta X \cong \sum_{\sigma : \xi \to \zeta \text{ surjective}} X_\xi$$

is polynomial. Furthermore $\mathbb{X}$ is linear since the sets

$$\mathrm{dom}(\mathrm{comp}(s)) \quad \text{and} \quad \sum_{v \in \mathrm{dom}(s)} \mathrm{dom}(s(v))\,, \quad \text{for } s \in \mathbb{X}\mathbb{X}A\,,$$

are both singletons. Consequently, the claim follows by Theorem 2.4. □

One could hope to construct a distributive law $\mathbb{T}^\circ\mathbb{X} \Rightarrow \mathbb{X}\mathbb{T}^\circ$ by applying the Theorem of Beck to the monad structure on $\mathbb{X}\mathbb{T}^\circ$. This does not work for the following reason.

Lemma 4.9. *The natural transformation* $\mathbb{X}\mathrm{sing} : \mathbb{X} \Rightarrow \mathbb{X}\mathbb{T}^\circ$ *is* not *a morphism of monads.*

Proof. The following of the two axioms fails:

$$\mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) \circ \mathbb{X}\mathrm{sing} \circ \mathbb{X}\mathbb{X}\mathrm{sing} \neq \mathbb{X}\mathrm{sing} \circ \mathrm{comp}.$$

To see this, fix $\langle\sigma, \langle\tau, a\rangle\rangle \in \mathbb{X}\mathbb{X}A$. Then

$$\begin{aligned}
&(\mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ) \circ \mathbb{X}\mathrm{sing} \circ \mathbb{X}\mathbb{X}\mathrm{sing})(\langle\sigma, \langle\tau, a\rangle\rangle) \\
&= \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathbb{T}^\circ \mathrm{re}_\circ \circ \mathrm{sing} \circ \mathbb{X}\mathrm{sing})(\langle\sigma, \langle\tau, a\rangle\rangle) \\
&= \mathbb{X}(\mathrm{flat}^\times \circ \iota \circ \mathrm{sing} \circ \mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing})(\langle\sigma, \langle\tau, a\rangle\rangle) \\
&= \mathbb{X}(\mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing})(\langle\sigma, \langle\tau, a\rangle\rangle) \\
&= \langle\sigma, {}^\tau\mathrm{sing}(a)\rangle\,,
\end{aligned}$$

whereas

$$\begin{aligned}
&(\mathbb{X}\mathrm{sing} \circ \mathrm{comp})(\langle \sigma, \langle \tau, a\rangle\rangle) \\
&= \mathbb{X}\mathrm{sing}(\langle \sigma \circ \tau, a\rangle) \\
&= \langle \sigma \circ \tau, \mathrm{sing}(a)\rangle\,.
\end{aligned}$$

For $\tau \neq \mathrm{id}$, these two values are different. $\square$

Exercise 4.2. Prove that the monad $\mathbb{X}\mathbb{T}^{\circ}$ satisfies all of the other conditions in the Theorem of Beck. ⌟

Graphs And Unravellings

The next step is to transfer the unravelling operation from $\mathbb{T}^{\times}A$ to arbitrary sets.

Definition 4.10. (a) An *unravelling structure* $\langle A, \mathrm{re}, \mathrm{un}\rangle$ consists of a set $A \in \mathsf{Pos}^{\Xi_+}$ equipped with two functions

$$\mathrm{re} : \mathbb{X}A \to A \quad \text{and} \quad \mathrm{un} : A \to \mathbb{X}A$$

such that $\langle A, \mathrm{re}\rangle$ forms an $\mathbb{X}$-algebra while un satisfies

$$\mathbb{X}\mathrm{un} \circ \mathrm{un} = \mathbb{X}\mathrm{in} \circ \mathrm{un} \quad \text{and} \quad \mathrm{re} \circ \mathrm{un} = \mathrm{id}\,.$$

We call $\mathrm{un}(a)$ the *unravelling* of a. To keep notation simple, we write

$${}^{\sigma}a := \mathrm{re}(\langle \sigma, a\rangle)\,.$$

(b) A *morphism of unravelling structures* is a function $\varphi : A \to B$ satisfying

$$\mathrm{un} \circ \varphi = \mathbb{X}\varphi \circ \mathrm{un} \quad \text{and} \quad \varphi \circ \mathrm{re} = \mathrm{re} \circ \mathbb{X}\varphi\,.$$

(c) The *free unravelling structure* generated by a set X is $\langle \mathbb{X}X, \mathrm{comp}, \mathbb{X}\mathrm{in}\rangle$. ⌟

Clearly, the operations re and un defined above for trees $t \in \mathbb{T}^\times A$ induce an unravelling structure on $\mathbb{T}^\times A$. But note that this is not the case for $\mathbb{R}A$ since we have $\mathrm{re}(\mathrm{un}(g)) \neq g$, for every $g \in \mathbb{R}A$ that is not a tree.

Example. For each $\mathbb{T}^\times$-algebra $\mathfrak{A} = \langle A, \pi \rangle$, we can equip the universe A with the *trivial* unravelling structure where

$$\mathrm{un} := \mathrm{in} \quad \text{and} \quad {}^\sigma a := \pi({}^\sigma \mathrm{sing}(a))\,. \qquad \lrcorner$$

Remark. (a) Note that the monad multiplication $\mathrm{flat}^\times$ is *not* a morphism of unravelling structures since $\mathrm{un} \circ \mathrm{flat}^\times \neq \mathbb{X}\mathrm{flat}^\times \circ \mathrm{un}$. In what follows we will therefore *not* work in the category of unravelling structures and their morphisms. Instead we will work in the weaker category of unravelling structures with arbitrary monotone maps as morphisms.

(b) Intuitively the axioms for $\mathrm{un} : A \to \mathbb{X}A$ say that $\langle A, \mathrm{un} \rangle$ forms a coalgebra for the 'comonad' $\langle \mathbb{X}, \mathbb{X}\mathrm{in}, \mathrm{re} \rangle$. But note that re is not a natural transformation. So the definition is 'local' to the set A. $\lrcorner$

As a technical tool we use the following generalisation of the unravelling relation for graphs where we do not only unravel the graph itself but also each label. The intuition is as follows. Suppose we are given a relation $\theta \subseteq A \times B$ and a graph $h \in \mathbb{R}B$. We construct an (unravelled) graph $g \in \mathbb{R}A$ as follows. Starting at the root v, we pick some element $c \;\theta\; h(v)$, and label $g(v)$ by the unravelling of c. Then we recursively choose labellings for the successors. Note that the shapes of g and h are different since we are unravelling g, so the labels in g might have a higher arity than the corresponding ones in h. Therefore, we simultaneously construct a graph homomorphism $\varphi : g \to h$ to keep track of which vertices of g correspond to which ones of h.

To simplify the definition, we will split the construction into two stages. In the first step we apply the unravelling operation to every label of h, resulting in a graph $\mathbb{R}\mathrm{un}(h) \in \mathbb{R}\mathbb{X}B$. What is then left for the second step is the following relation, which does the choosing of the label and the unravelling of the tree. What makes this operation complicated is the fact that the unravelling depends on the chosen label, while the label may depend on which copy (produced by previous unravelling steps) of a vertex we are at. So we cannot separate the second stage into two independent phases.

Definition 4.11. (a) Let $g \in \mathbb{R}_\xi A$ and $h \in \mathbb{R}_\zeta B$. A *graph homomorphism* is a function $\varphi : \mathrm{dom}_+(g) \to \mathrm{dom}_+(h)$ such that

- φ maps the root of g to the root of h;
- $\varphi(u)$ is a successor of $\varphi(v)$ if, and only if, u is a successor of v (not necessarily with the same edge labelling); and
- $\varphi(v)$ is labelled by a variable if, and only if, v is labelled by one.

(b) Suppose that $\varphi : g \to h$ is a surjective graph homomorphism and let $v \in \mathrm{dom}_+(g)$ be a vertex of sort ξ with successors $(u_x)_{x\in\xi}$ and suppose that $\varphi(v)$ has sort ζ. We denote by $\varphi_{/v} : \xi \to \zeta$ the function such that

$$\varphi(u_x) \text{ is the } \varphi_{/v}(x)\text{-successor of } \varphi(v)\,.$$

(c) Let $s \in \mathbb{R}_\xi A$, $t \in \mathbb{R}_\zeta B$, and $\theta \subseteq \mathbb{X}A \times B$. We write

$$\varphi, \sigma : s\ \theta^{\mathrm{sel}}\ t$$

if the following conditions are satisfied.

- $s \in \mathbb{T}^\circ A$
- $\varphi : s \to t$ is a surjective graph homomorphism.
- $\sigma : \xi \to \zeta$ is surjective.
- $\langle \varphi_{/v}, s(v)\rangle\ \theta\ t(\varphi(v))\,, \quad$ for every $v \in \mathrm{dom}(s)$.
- $\sigma(s(v)) = t(\varphi(v))\,, \quad$ if $s(v) = x$ is a variable. ⌟

We are mostly interested in the cases where θ is either the identity $=$ or set membership $\in$. The resulting relations are

$$\begin{aligned}
&\varphi, \sigma : s =^{\mathrm{sel}} t\,, && \text{for } s \in \mathbb{T}^\times A \text{ and } t \in \mathbb{T}^\times \mathbb{X}A\,,\\
&\varphi, \sigma : s \in^{\mathrm{sel}} t\,, && \text{for } s \in \mathbb{T}^\times A \text{ and } t \in \mathbb{T}^\times \mathrm{Up}(\mathbb{X}A)\,.
\end{aligned}$$

Combining them with the unravelling operation as explained above, we obtain the relations

$$\begin{aligned}
&\varphi, \sigma : s =^{\mathrm{un}} t && :\text{iff} \quad \varphi, \sigma : s =^{\mathrm{sel}} \mathbb{R}\mathrm{un}(t)\,,\\
&\varphi, \sigma : s \in^{\mathrm{un}} t\,, && :\text{iff} \quad \varphi, \sigma : s \in^{\mathrm{sel}} \mathbb{R}\mathrm{Up}(\mathrm{un})(t)\,.
\end{aligned}$$

Example. We have $\varphi, \sigma : g \in^{\mathrm{un}} h$ where g is the tree on the left, h the one on the right, $\varphi : g \to h$ is the obvious homomorphism, and $\sigma : \{x, y, z\} \to \{x\}$.

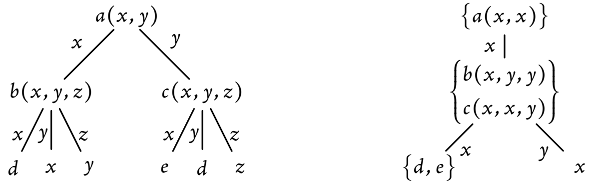

⌟

Remark. (a) For every graph g, there exists a canonical graph homomorphism $\varphi : \mathrm{gun}(g) \to g$.

(b) Note that

$$\varphi, \sigma : g =^{\mathrm{sel}} k \quad \text{and} \quad k \,\theta^{\mathbb{R}}\, h \quad \text{implies} \quad \varphi, \sigma : g \,\theta^{\mathrm{sel}}\, h\,,$$

but the converse is generally not true since the function φ does not need to be injective and we can choose different values $\langle \varphi_{/u}, c_u \rangle$, $\langle \varphi_{/v}, c_v \rangle \;\theta\; h(w)$ for $u, v \in \varphi^{-1}(w)$. For this reason, we cannot reduce the relation $\in^{\mathrm{sel}}$ to the much simpler $=^{\mathrm{sel}}$. ⌟

Let us derive an algebraic description of the relation $\varphi, \sigma : s =^{\mathrm{sel}} t$ that is much easier to work with. We introduce a function sun satisfying

$$\langle \sigma, s \rangle = \mathrm{sun}(t) \quad \text{iff} \quad \varphi, \sigma : s =^{\mathrm{sel}} t\,, \quad \text{for some } \varphi\,,$$

and a similar function dun associated with the relation $=^{\mathrm{un}}$.

Definition 4.12. (a) For a set A, we define the *strong unravelling operation* $\mathrm{sun} : \mathbb{T}^{\times}\mathbb{X}A \to \mathbb{X}\mathbb{T}^{\circ}A$ by

$$\mathrm{sun} := \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing})\,.$$

(b) For an unravelling structure A, we define the *deep unravelling operation* $\mathrm{dun} : \mathbb{T}^{\times}A \to \mathbb{X}\mathbb{T}^{\circ}A$ by

$$\mathrm{dun} := \mathrm{sun} \circ \mathbb{T}^{\times}\mathrm{un}\,.$$

⌟

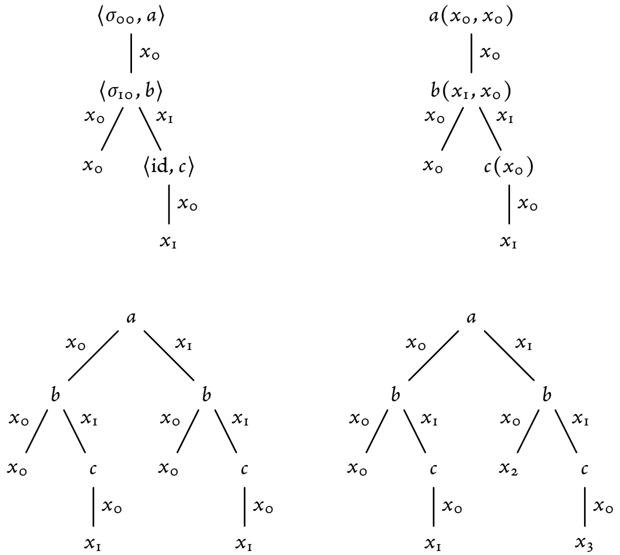

Figure 1.: An example of $\mathrm{sun}(t)$

Example. To understand the definition of sun, let us consider the following tree $t \in \mathbb{T}^\times \mathbb{X} A$ in Figure 1. We have depicted t itself, the intermediate terms $t' := \mathbb{T}^\times(\mathrm{re}_0 \circ \mathbb{X}\mathrm{sing})(t)$ and $t'' := \mathrm{flat}(t')$, and the end result $\mathrm{sun}(t)$. Here $a, b \in A_{\{x_0, x_1\}}$, $c \in A_{\{x_0\}}$, and σ_{ij} denotes the function mapping $x_0 \mapsto x_i$ and $x_1 \mapsto x_j$. ⌟

Let us check that the above definitions have the desired effect.

Lemma 4.13. *We have*

$$\begin{aligned} \langle \sigma, s \rangle = \mathrm{sun}(t) \quad &\text{iff} \quad \varphi, \sigma \colon s =^{\mathrm{sel}} t\,, \quad \textit{for some } \varphi\,, \\ \langle \sigma, s \rangle = \mathrm{dun}(t) \quad &\text{iff} \quad \varphi, \sigma \colon s =^{\mathrm{un}} t\,, \quad \textit{for some } \varphi\,. \end{aligned}$$

Proof. We only have to prove the first equivalence. Then the second one follows by definition of dun and $=^{\mathrm{un}}$. Given a tree t, we set

$$r := \mathbb{R}(\mathrm{re}_0 \circ \mathbb{X}\mathrm{sing})(t) \quad \text{and} \quad \langle \sigma, s \rangle := \mathrm{un}(\mathrm{flat}^\times(r))\,.$$

Let $\varphi : \mathrm{dom}_+(\mathrm{flat}^\times(r)) \to \mathrm{dom}_+(t)$ be the graph homomorphism induced by the canonical map

$$\mathrm{dom}(\mathrm{flat}^\times(r)) \to \sum_{v \in \mathrm{dom}(r)} \mathrm{dom}(r(v))\,,$$

and fix φ', σ', s' with $\varphi', \sigma' : s' =^{\mathrm{sel}} t$. It is sufficient to show that

$$\varphi = \varphi'\,, \quad \sigma = \sigma'\,, \quad \text{and} \quad s = s'\,.$$

We start by proving that $\varphi(v) = \varphi'(v)$ and $s(v) = s'(v)$. We proceed by induction on v. For the root $v = \langle\rangle$ of $\mathrm{flat}^\times(r)$, we have $\varphi(\langle\rangle) = \langle\rangle = \varphi'(\langle\rangle)$.

For the inductive step, suppose that we have already shown that $\varphi(v) = \varphi'(v)$. We will prove that $s(v) = s'(v)$ and that $\varphi(u) = \varphi'(u)$, for every successor u of v. By definition of $=^{\mathrm{sel}}$, we have

$$t(\varphi'(v)) = \langle \varphi'_{/v}, s'(v) \rangle\,, \quad \text{for } v \in \mathrm{dom}_+(s')\,.$$

This implies that

$$r(\varphi'(v)) = (\mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing})\big(\langle \varphi'_{/v}, s'(v)\rangle\big) = {}^{\varphi'_{/v}}\mathrm{sing}(s'(v))\,.$$

Consequently,

$$s(v) = \mathrm{flat}^\times(r)(v) = r(\varphi(v))(\langle\rangle) = r(\varphi'(v))(\langle\rangle) = s'(v)\,.$$

To complete the induction, it remains to show that $\varphi_{/v} = \varphi'_{/v}$. Let $(u_x)_x$ be the successors of v in s and let $(w_y)_y$ be the successors of $\varphi(v)$ in r. Then

$$r(\varphi(v)) = {}^{\varphi'_{/v}}\mathrm{sing}(s(v))$$

implies that the x-successor of v in s corresponds (via φ) to the $\varphi'_{/v}(x)$-successor of $\varphi(v)$ in r, that is,

$$\varphi(u_x) = w_{\varphi'_{/v}(x)}\,.$$

But, by definition of $\varphi_{/v}$, we also have $\varphi(u_x) = w_{\varphi_{/v}(x)}$. Hence,

$$\varphi_{/v}(x) = \varphi'_{/v}(x)\,.$$

This completes the induction. To finish the proof it remains to show that $\sigma = \sigma'$ and that $s(v) = s'(v)$, for all $v \in \mathrm{dom}_+(s) \smallsetminus \mathrm{dom}(s)$. For the latter, note that the vertices of s carrying a variable are the same as those of s' carrying one. Since the variable labelling is determined by the ordering of these vertices with respect to the length-lexicographic order, it follows that the two labellings coincide.

Hence, let v be such a vertex. Then

$$\begin{aligned}\sigma(s(v)) = \mathrm{flat}^\times(r)(v) = r(\varphi(v)) &= t(\varphi(v))\\ &= \sigma'(s'(v)) = \sigma'(s(v))\,.\end{aligned}$$

Thus, $\sigma(x) = \sigma'(x)$, for all x, which implies that $\sigma = \sigma'$ □

Let us collect a few basic properties of the operations we have just introduced.

Lemma 4.14.

(a) $\mathbb{X}(\mathrm{un} \circ \mathrm{flat}^\times \circ \iota) \circ \mathrm{dun} = \mathbb{X}(\mathrm{in} \circ \mathrm{flat}^\times \circ \iota) \circ \mathrm{dun}$

(b) $\mathrm{flat}^\times \circ \mathrm{re} \circ \mathrm{dun} = \mathrm{flat}^\times$

(c) $\mathrm{un} \circ \mathrm{flat}^\times = \mathbb{X}(\mathrm{flat}^\times \circ \iota) \circ \mathrm{dun}$

(d) $\mathrm{sun} \circ \mathbb{T}^\times \mathrm{in} = \mathrm{un}$

(e) $\mathrm{sun} \circ \mathrm{sing}^\times = \mathbb{X}\mathrm{sing}$

Proof. (a) Let $\langle \sigma, s \rangle = \mathrm{dun}(t)$. According to Lemma 4.13, this means that $\varphi, \sigma : s =^{\mathrm{un}} t$. By definition of $=^{\mathrm{un}}$, we therefore have

$$\mathrm{un}(t(\varphi(v))) = \langle \varphi_{/v}, s(v) \rangle\,, \quad \text{for all } v \in \mathrm{dom}(s)\,.$$

In particular, $s(v) \in \mathbb{T}^\circ A$ and, therefore, $s \in \mathbb{T}^\circ\mathbb{T}^\circ A$. This implies that $\mathrm{flat}(s) \in \mathbb{T}^\circ A$. Hence, $\mathrm{un}(\mathrm{flat}(s)) = \langle \mathrm{id}, \mathrm{flat}(s) \rangle$ and we have

$$\begin{aligned} \mathbb{X}(\mathrm{un} \circ \mathrm{flat})(\mathrm{dun}(t)) &= \langle \sigma, \mathrm{un}(\mathrm{flat}(s)) \rangle \\ &= \langle \sigma, \langle \mathrm{id}, \mathrm{flat}(s) \rangle \rangle \\ &= \langle \sigma, \mathrm{in}(\mathrm{flat}(s)) \rangle = \mathbb{X}(\mathrm{in} \circ \mathrm{flat})(\mathrm{dun}(t))\,. \end{aligned}$$

(b) By Lemma 4.5 it follows that

$$\begin{aligned} \mathrm{flat}^\times \circ \mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un} &= \mathrm{re}_\circ \circ \mathbb{X}(\mathrm{flat}^\times \circ \iota) \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un} \\ &= \mathrm{re}_\circ \circ \mathbb{X}(\mathrm{flat}^\times \circ \mathrm{sing}^\times) \circ \mathrm{un} \\ &= \mathrm{re}_\circ \circ \mathrm{un} \\ &= \mathrm{id}\,. \end{aligned}$$

Consequently,

$$\begin{aligned} \mathrm{flat}^\times \circ \mathrm{re} \circ \mathrm{dun} &= \mathrm{flat}^\times \circ \mathrm{re} \circ \mathrm{un} \circ \mathrm{flat}^\times \circ \mathbb{T}^\times(\mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \mathrm{flat}^\times \circ \mathrm{flat}^\times \circ \mathbb{T}^\times(\mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \mathrm{flat}^\times \circ \mathbb{T}^\times\mathrm{flat}^\times \circ \mathbb{T}^\times(\mathrm{re}_\circ \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \mathrm{flat}^\times \circ \mathbb{T}^\times\mathrm{id} \\ &= \mathrm{flat}^\times\,. \end{aligned}$$

(c) By (a) and Lemma 4.5, we have

$$\begin{aligned}
\mathbb{X}(\mathrm{flat}^{\times} \circ \iota) \circ \mathrm{dun} &= \mathrm{comp} \circ \mathbb{X}(\mathrm{in} \circ \mathrm{flat}^{\times} \circ \iota) \circ \mathrm{dun} \\
&= \mathrm{comp} \circ \mathbb{X}(\mathrm{in} \circ \mathrm{flat}^{\times} \circ \iota) \circ \mathrm{dun} \\
&= \mathrm{comp} \circ \mathbb{X}\mathrm{un} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \iota) \circ \mathrm{dun} \\
&= \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \iota) \circ \mathrm{dun} \\
&= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathrm{re}_{\circ} \circ \mathrm{dun} \\
&= \mathrm{un} \circ \mathrm{flat}^{\times} .
\end{aligned}$$

(d), (e) We have

$$\begin{aligned}
\mathrm{sun} \circ \mathbb{T}^{\times}\mathrm{in} &= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing}) \circ \mathbb{T}^{\times}\mathrm{in} \\
&= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathrm{in} \circ \mathrm{sing}) \\
&= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}\mathrm{sing} \\
&= \mathrm{un} \,, \\
\mathrm{sun} \circ \mathrm{sing}^{\times} &= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing}) \circ \mathrm{sing}^{\times} \\
&= \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathrm{sing}^{\times} \circ \mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \\
&= \mathrm{un} \circ \mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \\
&= \mathbb{X}\mathrm{sing} .
\end{aligned}$$

□

In Lemma 4.13, we have found an algebraic characterisation of the relations $=^{\mathrm{sel}}$ and $=^{\mathrm{un}}$ in terms of the operations sun and dun. Unfortunately, there does not seem to exist a similar purely algebraic characterisation of the relation $\in^{\mathrm{sel}}$. Instead, we define the corresponding operation directly in terms of $\in^{\mathrm{sel}}$.

Definition 4.15. The *selection operation* $\mathrm{sel} : \mathbb{T}^{\times} \circ \mathrm{Up} \circ \mathbb{X} \Rightarrow \mathrm{Up} \circ \mathbb{X} \circ \mathbb{T}^{\circ}$ is defined by

$$\mathrm{sel}(t) := \{\, \langle \sigma, s \rangle \mid \varphi, \sigma : s \in^{\mathrm{sel}} t \,\}, \quad \text{for } t \in \mathbb{T}^{\times}\mathrm{Up}(\mathbb{X}A) .$$ ⌟

The properties of this operation are as follows.

Lemma 4.16.

(a) $\mathrm{sel} : \mathbb{T}^\times \circ \mathrm{Up} \circ \mathbb{X} \Rightarrow \mathrm{Up} \circ \mathbb{X} \circ \mathbb{T}^\circ$ *is a natural transformation on* Pos^{Ξ_+}.

(b) $\mathrm{sel} \circ \mathbb{T}^\times \mathrm{pt} = \mathrm{pt} \circ \mathrm{sun}$

(c) $\mathrm{sel} \circ \mathrm{sing}^\times = \mathrm{Up}(\mathbb{X}\mathrm{sing})$

(d) $\mathrm{sel} \circ \mathbb{T}^\times(\mathrm{pt} \circ \mathrm{in}) = \mathrm{pt} \circ \mathrm{un}$

(e) $\mathrm{Up}(\mathrm{dun} \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) = \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})$

Proof. (a) Let $f : A \to B$. By definition of $\in^{\mathrm{sel}}$, we have

$$\begin{array}{lll}
 & \varphi, \sigma : s \in^{\mathrm{sel}} \mathbb{T}^\times \mathrm{Up}(\mathbb{X}f)(t) & \\
\text{iff} & \langle \varphi_{/v}, s(v) \rangle \in \mathrm{Up}(\mathbb{X}f)(t(\varphi(v))), & \text{for all } v, \\
\text{iff} & s(v) \geq f(r(v)) \quad \text{and} \quad \langle \varphi_{/v}, r(v) \rangle \in t(\varphi(v)), & \text{for all } v, \\
\text{iff} & s \geq \mathbb{T}^\circ f(r) \quad \text{and} \quad \varphi, \sigma : r \in^{\mathrm{sel}} t. &
\end{array}$$

This implies that $\mathrm{sel}(\mathbb{T}^\times \mathrm{Up}(\mathbb{X}f)(t)) = \mathrm{Up}(\mathbb{X}\mathbb{T}^\circ f)(\mathrm{sel}(t))$.

(b) To simplify notation, we will leave the universal quantification over vertices v implicit in the expressions below. Let $t \in \mathbb{T}^\times \mathbb{X} A$. Then

$$\begin{aligned}
\mathrm{sel}(\mathbb{T}^\times \mathrm{pt}(t)) &= \Uparrow\{ \langle \sigma, s \rangle \mid \varphi, \sigma : s \in^{\mathrm{sel}} \mathbb{T}^\times \mathrm{pt}(t) \} \\
&= \Uparrow\big\{ \langle \sigma, s \rangle \,\big|\, \langle \varphi_{/v}, s(v) \rangle \in \mathrm{pt}(t(\varphi(v))) \text{ or} \\
&\qquad\qquad [s(v) = x \text{ and } \mathbb{T}^\times \mathrm{pt}(t)(\varphi(v)) = \sigma(x)] \big\} \\
&= \Uparrow\big\{ \langle \sigma, s \rangle \,\big|\, \langle \varphi_{/v}, s(v) \rangle \geq t(\varphi(v)) \text{ or} \\
&\qquad\qquad [s(v) = x \text{ and } t(\varphi(v)) = \sigma(x)] \big\} \\
&= \Uparrow\big\{ \langle \sigma, s \rangle \,\big|\, \langle \varphi_{/v}, s(v) \rangle = t(\varphi(v)) \text{ or} \\
&\qquad\qquad [s(v) = x \text{ and } t(\varphi(v)) = \sigma(x)] \big\} \\
&= \Uparrow\{ \langle \sigma, s \rangle \mid \varphi, \sigma : s =^{\mathrm{sel}} t \} \\
&= \Uparrow\{\mathrm{sun}(t)\} \\
&= \mathrm{pt}(\mathrm{sun}(t)).
\end{aligned}$$

(c) Let $I \in \mathrm{Up}(\mathbb{X}A)$. Then

$$\begin{aligned}\mathrm{sel}(\mathrm{sing}^\times(I)) &= \Uparrow\{\,\langle\sigma, s\rangle \mid \varphi, \sigma : s \in^{\mathrm{sel}} \mathrm{sing}^\times(I)\,\}\\ &= \Uparrow\{\,\langle\sigma, s\rangle \mid s = \mathrm{sing}(a)\,,\ \langle\tau, a\rangle \in I\,,\ \sigma = \tau\,\}\\ &= \Uparrow\{\,\langle\sigma, \mathrm{sing}(a)\rangle \mid \langle\sigma, a\rangle \in I\,\}\\ &= \mathrm{Up}(\mathbb{X}\mathrm{sing})(I)\,.\end{aligned}$$

(d) By (b) and Lemma 4.14 (d), we have

$$\mathrm{sel} \circ \mathbb{T}^\times(\mathrm{pt} \circ \mathrm{in}) = \mathrm{pt} \circ \mathrm{sun} \circ \mathbb{T}^\times\mathrm{in} = \mathrm{pt} \circ \mathrm{un}\,.$$

(e) Let $\langle\sigma, s\rangle \in \mathrm{sel}(\mathbb{T}^\times\mathrm{Up}(\mathrm{un})(t))$. Then $\varphi, \sigma : s \in^{\mathrm{sel}} \mathbb{T}^\times\mathrm{Up}(\mathrm{un})(t)$, which implies that

$$\langle\varphi_{/v}, s(v)\rangle \in \mathrm{un}(t(\varphi(v)))\,, \quad \text{for all } v \in \mathrm{dom}(s)\,.$$

Consequently, we have $\mathrm{un}(s(v)) = \langle\mathrm{id}, s(v)\rangle$, that is, $\mathbb{T}^\circ\mathrm{un}(s) = \mathbb{T}^\circ\mathrm{in}(s)$. Hence,

$$\begin{aligned}(\mathbb{T}^\circ\mathrm{un} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle) &= \mathbb{T}^\circ\mathrm{un}({}^\sigma s)\\ &= {}^\sigma\mathbb{T}^\circ\mathrm{un}(s)\\ &= {}^\sigma\mathbb{T}^\circ\mathrm{in}(s)\\ &= \mathbb{T}^\circ\mathrm{in}({}^\sigma s) = (\mathbb{T}^\circ\mathrm{in} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle)\,.\end{aligned}$$

Furthermore, $s \in \mathbb{T}^\circ A$ implies that $\mathrm{un}(s) = \langle\mathrm{id}, s\rangle$. It therefore follows by Lemma 4.14 (d) that

$$\begin{aligned}(\mathrm{dun} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle) &= (\mathrm{sun} \circ \mathbb{T}^\times\mathrm{un} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle)\\ &= (\mathrm{sun} \circ \mathbb{T}^\times\mathrm{in} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle)\\ &= (\mathrm{un} \circ \mathrm{re}_\circ)(\langle\sigma, s\rangle)\\ &= (\mathrm{comp} \circ \mathbb{X}\mathrm{un})(\langle\sigma, s\rangle)\\ &= (\mathrm{comp} \circ \mathbb{X}\mathrm{in})(\langle\sigma, s\rangle)\\ &= \langle\sigma, s\rangle\,.\end{aligned}$$

Consequently,

$$\mathrm{Up}(\mathrm{dun} \circ \mathrm{re}_\circ) \restriction (\mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})) = \mathrm{Up}(\mathrm{id}) \restriction (\mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un}))\,.$$

$\square$

We need one more equation concerning the operation sel whose proof is more involved: Lemma 4.18 below contains a commutation relation between sel and $\mathrm{flat}^\times$ that is similar to one of the axioms of a distributive law. The proof makes use of the following technical lemma.

Lemma 4.17. *Let $r \in \mathbb{T}^\times \mathbb{T}^\times A$ and $t \in \mathbb{T}^\times \mathbb{T}^\times B$ be trees, set $s := \mathrm{flat}(r)$, let*

$$\begin{aligned}
\chi &: \mathrm{dom}_+(s) \to \mathrm{dom}_+(\mathrm{flat}^\times(t))\,,\\
\varphi &: \mathrm{dom}_+(r) \to \mathrm{dom}_+(t)\,,\\
\psi_v &: \mathrm{dom}_+(r(v)) \to \mathrm{dom}_+(t(\varphi(v)))
\end{aligned}$$

be surjective graph homomorphisms, and let

$$\begin{aligned}
\lambda &: \mathrm{dom}_+(\mathrm{flat}^\times(t)) \to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v)) + [\mathrm{dom}_+(t) \setminus \mathrm{dom}(t)]\,,\\
\mu &: \mathrm{dom}_+(\mathrm{flat}(r)) \to \sum_{v \in \mathrm{dom}(r)} \mathrm{dom}(r(v)) + [\mathrm{dom}_+(r) \setminus \mathrm{dom}(r)]
\end{aligned}$$

be the functions induced by the canonical maps

$$\begin{aligned}
\mathrm{dom}(\mathrm{flat}^\times(t)) &\to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v))\\
\mathrm{dom}(\mathrm{flat}(r)) &\to \sum_{v \in \mathrm{dom}(r)} \mathrm{dom}(r(v))\,.
\end{aligned}$$

Then

$$\lambda(\chi(w)) = \langle \varphi(v), \psi_v(u)\rangle\,, \quad \textit{for every } w \in \mathrm{dom}(s) \textit{ with } \mu(w) = \langle v, u\rangle\,,$$

implies that

$$\chi_{/w} = (\psi_v)_{/u}\,, \quad \textit{for } \mu(w) = \langle v, u\rangle\,.$$

Proof. Consider a vertex $w \in \mathrm{dom}(s)$ with $\mu(w) = \langle v, u \rangle$ and an x-successor $\tilde{u}$ of u. Suppose that $\lambda(\chi(w)) = \langle v', u' \rangle$. First, let us consider the case where $\tilde{u} \in \mathrm{dom}(r(v))$. Let $\tilde{w}$ be the successor of w with $\mu(\tilde{w}) = \langle v, \tilde{u} \rangle$. By assumption, we have $\lambda(\chi(\tilde{w})) = \langle \varphi(v), \psi_v(\tilde{u}) \rangle$ and $\psi_v(\tilde{u})$ is the y-successor of $\psi_v(u)$ in $t(\varphi(v))$, for some y. By definition, it follows that $\chi_{/w}(x) = y$ and $(\psi_v)_{/u}(x) = y$.

It remains to consider the case where $\tilde{u} \notin \mathrm{dom}(r(v))$. Then $r(v)(\tilde{u}) = z$, for some variable z. Let v' be the z-successor of v, let $\langle\rangle$ be the root of $r(v')$, and let $\tilde{w}$ be the successor of w with $\mu(\tilde{w}) = \langle v', \langle\rangle \rangle$. Then $\lambda(\chi(w)) = \langle \varphi(v'), \psi_{v'}(\langle\rangle) \rangle$. Let y be the variable such that $\lambda(\varphi(v'), \psi_{v'}(\langle\rangle))$ is the y-successor of $\lambda(\varphi(v), \psi_v(u))$. Then $\chi_{/w}(x) = y$ and $(\psi_v)_{/u}(x) = y$. $\square$

Lemma 4.18. $\mathrm{sel} \circ \mathrm{flat}^\times = \mathrm{Up}(\mathbb{X}\mathrm{flat}) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{sel}$

Proof. Note that the canonical function

$$\mathrm{dom}(\mathrm{flat}^\times(t)) \to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v))$$

induces a function

$$\lambda : \mathrm{dom}_+(\mathrm{flat}^\times(t)) \to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v)) + [\mathrm{dom}_+(t) \setminus \mathrm{dom}(t)] \,.$$

Similarly, for a tree r (which we will specify below), we obtain a function

$$\mu : \mathrm{dom}_+(\mathrm{flat}(r)) \to \sum_{v \in \mathrm{dom}(r)} \mathrm{dom}(r(v)) + [\mathrm{dom}_+(r) \setminus \mathrm{dom}(r)] \,.$$

To prove the lemma, we check the two inclusions separately.

($\supseteq$) Suppose that $\langle \sigma, s \rangle \in \mathrm{Up}(\mathbb{X}\mathrm{flat})(\mathrm{sel}(\mathbb{T}^\times \mathrm{sel}(t)))$. Then

$$s = \mathrm{flat}(r) \,, \quad \text{for some} \quad \varphi, \sigma : r \in^{\mathrm{sel}} \mathbb{T}^\times \mathrm{sel}(t) \,.$$

For every vertex v of r, it follows that

$$\langle \varphi_{/v}, r(v) \rangle \in \mathrm{sel}(t(\varphi(v)))$$

or $\quad r(v) = x$ and $\mathrm{sel}(t(\varphi(v))) = \sigma(x)$.

This implies that

$$\psi_v, \varphi_{/v} : r(v) \in^{\mathrm{sel}} t(\varphi(v)) \quad \text{or} \quad r(v) = x \text{ and } t(\varphi(v)) = \sigma(x)\,,$$

for some homomorphism ψ_v. Let χ be the unique graph homomorphism satisfying the equations

$$\lambda(\chi(w)) = \begin{cases} \langle \varphi(v), \psi_v(u) \rangle & \text{if } \mu(w) = \langle v, u \rangle\,, \\ \varphi(v) & \text{if } \mu(w) = v\,, \end{cases}$$

where λ and μ are the homomorphisms defined above. We claim that $\chi, \sigma : s \in^{\mathrm{sel}} \mathrm{flat}^{\times}(t)$, which implies that $\langle \sigma, s \rangle \in \mathrm{sel}(\mathrm{flat}^{\times}(t))$.

Hence, fix a vertex $w \in \mathrm{dom}_+(s) = \mathrm{dom}_+(\mathrm{flat}(r))$. First, consider the case where $w \in \mathrm{dom}(s)$. Suppose that $\mu(w) = \langle v, u \rangle$. Then $\psi_v, \varphi_{/v} : r(v) \in^{\mathrm{sel}} t(\varphi(v))$ implies that

$$\langle (\psi_v)_{/u}, r(v)(u) \rangle \in t(\varphi(v))(\psi_v(u))\,.$$

Consequently, we have

$$\langle (\psi_v)_{/u}, s(w) \rangle \in t(\varphi(v))(\psi_v(u)) = \mathrm{flat}^{\times}(t)(\chi(w))\,.$$

Furthermore, we have $(\psi_v)_{/u} = \chi_{/w}$ by Lemma 4.17.

It remains to consider the case where $s(w) = x$ is a variable. Then $\mu(w) = v$, for some $v \in \mathrm{dom}_+(r)$, and $r(v) = x$ implies that $t(\varphi(v)) = \sigma(x)$. Hence,

$$\mathrm{flat}^{\times}(t)(\chi(w)) = t(\lambda(\chi(w))) = t(\varphi(v)) = \sigma(x)\,.$$

($\subseteq$) Suppose that $\langle \sigma, s \rangle \in \mathrm{sel}(\mathrm{flat}^{\times}(t))$. Then

$$\chi, \sigma : s \in^{\mathrm{sel}} \mathrm{flat}^{\times}(t)\,, \quad \text{for some } \chi\,.$$

We define a tree r with $\mathrm{flat}(r) = s$ as follows. Intuitively, we factorise s by cutting every edge $w \to w'$ such that the corresponding vertices $\chi(w)$

and $\chi(w')$ in $\mathrm{flat}^{\times}(t)$ belong to different components $t(v)$ and $t(v')$, i.e., if $\lambda(\chi(w)) = \langle v, u\rangle$ and $\lambda(\chi(w')) = \langle v', u'\rangle$ with $v \neq v'$. The formal definition is as follows. Let us call a vertex $w \in \mathrm{dom}_+(s)$ *principal* if its image under χ corresponds to the root of some conponent $t(v)$, or to a leaf, that is, if

$$\lambda(\chi(w)) = \langle v, \langle\rangle\rangle \quad \text{or} \quad \lambda(\chi(w)) = v\,, \quad \text{for some } v\,,$$

(where $\langle\rangle$ denotes the root of $t(v)$). We define the domain of r by

$$\mathrm{dom}_+(r) := \{\, w \in \mathrm{dom}_+(s) \mid w \text{ is principal}\,\}$$

and the edge relation as follows. Given a principal vertex w, let $w_0, \dots, w_{n-1}$ be an enumeration of all minimal principal vertices w' with $w < w'$. We make w_i an i-successor of w. (The precise labels i are not important, only the fact that they are pairwise distinct.) Finally, the labelling of r is given by

$$r(w) := \begin{cases} r_w & \text{if } w \in \mathrm{dom}(s)\,, \\ s(w) & \text{if } w \notin \mathrm{dom}(s)\,, \end{cases}$$

where r_w is the tree with

$$\begin{aligned} \mathrm{dom}_+(r_w) &:= \Big\{\, u \in \mathrm{dom}_+(s) \;\Big|\; w \leq u \text{ and there is no principal } w' \\ &\qquad\qquad\qquad\qquad \text{with } w < w' < u \,\Big\}\,, \\ r_w(u) &:= \begin{cases} s(u) & \text{if } u \notin \mathrm{dom}_+(r) \text{ or } u = w\,, \\ i & \text{if } u = w_i \in \mathrm{dom}_+(r) \text{ is the } i\text{-successor of } w \\ & \quad \text{in } r\,. \end{cases} \end{aligned}$$

By definition, it follows that $\mathrm{flat}(r) = s$ and that

$$\begin{aligned} \mu(w) &= \langle v, w\rangle\,, && \text{if } w \in \mathrm{dom}(s)\,, \text{ where } v \text{ is the maximal principal} \\ & && \qquad \text{vertex with } v \leq w\,, \\ \text{and} \quad \mu(w) &= w\,, && \text{if } w \notin \mathrm{dom}(s)\,. \end{aligned}$$

Let φ and ψ_v be the functions defined by the equations

$$\begin{aligned} \langle \varphi(v), \psi_v(u) \rangle &= \lambda(\chi(w)), && \text{for } \mu(w) = \langle v, u \rangle, \\ \varphi(w) &= \lambda(\chi(w)), && \text{if } w \in \mathrm{dom}_+(s) \smallsetminus \mathrm{dom}(s), \\ \psi_v(u) &= u'' && \text{if } u \in \mathrm{dom}_+(r(v)) \smallsetminus \mathrm{dom}(r(v)), \end{aligned}$$

where the vertex u'' in the last equation is chosen as follows. Given u, let u' be the predecessor of u and let x be the label of the edge $u' \to u$. Then u'' is the $(\psi_v)_{/u'}(x)$-successor of $\psi_v(u')$.

We claim that, for all v,

$$\psi_v, \varphi_{/v} : r(v) \in^{\mathrm{sel}} t(\varphi(v)) \quad \text{or} \quad r(v) = x \text{ and } t(\varphi(v)) = \sigma(x).$$

Then it follows that

$$\begin{aligned} &\langle \varphi_{/v}, r(v) \rangle \in \mathrm{sel}(t(\varphi(v))) \\ \text{or} \quad &r(v) = x \text{ and } \mathrm{sel}(t(\varphi(v))) = \sigma(x). \end{aligned}$$

Thus,

$$\langle \sigma, r \rangle \in \mathrm{sel}(\mathbb{T}^\times \mathrm{sel}(t)) \quad \text{and} \quad \langle \sigma, s \rangle \in \mathrm{Up}(\mathbb{X}\mathrm{flat})(\mathrm{sel}(\mathbb{T}^\times \mathrm{sel}(t))),$$

as desired. Hence, it remains to prove the above claim.

If $r(v) = x$ is a variable, we have $s(v) = r(v) = x$ and, therefore,

$$t(\varphi(v)) = t(\lambda(\chi(v))) = \mathrm{flat}^\times(t)(\chi(v)) = \sigma(x),$$

as desired. Otherwise, $v \in \mathrm{dom}(r)$ and we have to show that

$$\psi_v, \varphi_{/v} : r(v) \in^{\mathrm{sel}} t(\varphi(v)).$$

Note that $\chi, \sigma : s \in^{\mathrm{sel}} \mathrm{flat}^\times(t)$ implies that

$$\langle \chi_{/w}, s(w) \rangle \in \mathrm{flat}^\times(t)(\chi(w)), \quad \text{for all } w.$$

We distinguish two cases. If $u \in \mathrm{dom}(r(v))$, let $w \in \mathrm{dom}(s)$ be the vertex with $\mu(w) = \langle v, u \rangle$. Then

$$\langle \chi_{/w}, r(v)(u) \rangle = \langle \chi_{/w}, s(w) \rangle \in \mathrm{flat}^{\times}(t)(\chi(w)) = t(\varphi(v))(\psi_v(u))\,.$$

By Lemma 4.17, we have $\chi_{/w} = (\psi_v)_{/u}$, which implies that

$$\langle (\psi_v)_{/u}, r(v)(u) \rangle \in t(\varphi(v))(\psi_v(u))\,.$$

If $u \in \mathrm{dom}_+(r(v)) \setminus \mathrm{dom}(r(v))$ with label $r(v)(u) = x$, let v' be the x-successor of v. By definition of $\varphi_{/v}$, it follows that $\varphi(v')$ is the $\varphi_{/v}(x)$-successor of $\varphi(v)$ in t. This implies that $t(v)(\psi_v(u)) = \varphi_{/v}(x)$. □

Exercise 4.3. Prove that $\mathrm{sun} \circ \iota : \mathbb{T}^{\circ}\mathbb{X} \Rightarrow \mathbb{X}\mathbb{T}^{\circ}$ forms a distributive law on Pos^{Ξ_+}. ⌟

Exercise 4.4. Prove that $\mathbb{X}\iota \circ \mathrm{sun} : \mathbb{T}^{\times}\mathbb{X} \Rightarrow \mathbb{X}\mathbb{T}^{\times}$ satisfies all axioms of a distributive law except for the equation $(\mathbb{X}\iota \circ \mathrm{sun}) \circ \mathbb{T}^{\times}\mathrm{in} = \mathrm{in}$. ⌟

A Partial Distributive Law

The idea to find our partial distributive law is to work in the category of unravelling structures, although this does not solve our problems entirely. First of all, there is no obvious way to lift the functor Up to unravelling structures. Given an unravelling structure A, we can define an 'unravelling map' $\mathrm{Up}(\mathrm{un}) : \mathrm{Up}(A) \to \mathrm{Up}(\mathbb{X}A)$, but we would need one of the form $\mathrm{Up}(A) \to \mathbb{X}\mathrm{Up}(A)$, and there is no natural transformation $\mathrm{Up} \circ \mathbb{X} \Rightarrow \mathbb{X} \circ \mathrm{Up}$. The functor $\mathbb{T}^{\times}$ on the other hand *can* be lifted to the category of unravelling structures, but only in a trivial way: given A we can forget its unravelling structure, construct $\mathbb{T}^{\times}A$, and equip it with the canonical unravelling structure defined above (which does not depend on that of A). In particular, with this definition the monad multiplication $\mathrm{flat}^{\times}$ would not be a morphism of the resulting unravelling structure. What would be more useful would be a lift that uses deep unravelling dun as the unravelling operation on $\mathbb{T}^{\times}A$. But there is no corresponding reconstitution operation re satisfying $\mathrm{re} \circ \mathrm{dun} = \mathrm{id}$.

What we will do instead is to use an ad-hoc argument showing how to define a lift of Up to sufficiently well-behaved $\mathbb{T}^\times$-algebras. We are mainly interested in free $\mathbb{T}^\times$-algebras, but a slightly more abstract definition helps to make the proof more modular. We extract the needed properties of the algebras in question in the following technical definition.

Definition 4.19. We say that a $\mathbb{T}^\times$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ *supports unravelling* if its universe A can be equipped with an unravelling structure that satisfies the following conditions.

$$\begin{aligned} \pi \circ \mathrm{re} \circ \mathbb{X}\mathrm{sing}^\times &= \mathrm{re}\,, \\ \mathrm{un} \circ \mathrm{re} &= \mathrm{comp} \circ \mathbb{X}\mathrm{un}\,, \\ \mathbb{X}(\mathrm{un} \circ \pi \circ \iota) \circ \mathrm{dun} &= \mathbb{X}(\mathrm{in} \circ \pi \circ \iota) \circ \mathrm{dun}\,. \end{aligned}$$ ⌟

The following observation might help explaining the second of the above axioms.

Lemma 4.20. *An unravelling structure* $\langle A, \mathrm{re}, \mathrm{un} \rangle$ *is a retract of a free one if, and only if, it satisfies*

$$\mathrm{un} \circ \mathrm{re} = \mathrm{comp} \circ \mathbb{X}\mathrm{un}\,.$$

Proof. $(\Rightarrow)$ Suppose that there are morphisms of unravelling structures $\mu : A \to \mathbb{X}X$ and $\rho : \mathbb{X}X \to A$ with $\rho \circ \mu = \mathrm{id}$. Then

$$\begin{aligned} \mathrm{un} \circ \mathrm{re} &= \mathbb{X}\mu \circ \mathbb{X}\mathrm{in} \circ \mathrm{comp} \circ \mathbb{X}\mu \\ &= \mathbb{X}\rho \circ \mathrm{comp} \circ \mathbb{X}\mathbb{X}\mathrm{in} \circ \mathbb{X}\mu \\ &= \mathbb{X}\rho \circ \mathrm{comp} \circ \mathbb{X}\mathbb{X}\mu \circ \mathbb{X}\mathrm{un} \\ &= \mathbb{X}\rho \circ \mathbb{X}\mu \circ \mathrm{comp} \circ \mathbb{X}\mathrm{un} = \mathrm{comp} \circ \mathbb{X}\mathrm{un}\,. \end{aligned}$$

$(\Leftarrow)$ Set

$$K := \{\, a \in A \mid \langle \sigma, a \rangle \in \mathrm{rng}\,\mathrm{un} \,\}\,,$$

let $i : K \to A$ be the inclusion map, and let $\mu : A \to \mathbb{X}K$ be the corestriction of un. Then $\mathrm{un} = \mathbb{X}i \circ \mu$. We claim that μ is the desired embedding of A

into $\mathbb{X}K$ and that $\rho := \mathrm{re} \circ \mathbb{X}i : \mathbb{X}K \to A$ is the corresponding quotient map.

First, note that

$$\rho \circ \mu = \mathrm{re} \circ \mathbb{X}i \circ \mu = \mathrm{re} \circ \mathrm{un} = \mathrm{id}\,.$$

Hence, it remains to show that ρ and μ are morphisms of unravelling structures.

For μ, note that

$$\mathrm{comp} \circ \mathbb{X}\mu = \mathrm{comp} \circ \mathbb{X}\mathrm{un} = \mathrm{un} \circ \mathrm{re} = \mu \circ \mathrm{re}$$

and $$\mathbb{X}\mu \circ \mathrm{un} = \mathbb{X}\mathrm{un} \circ \mathrm{un} = \mathbb{X}\mathrm{in} \circ \mathrm{un} = \mathbb{X}\mathrm{in} \circ \mu\,,$$

where the former holds by assumption and the latter by the axioms of an unravelling structure.

Concerning ρ, we have $\mathrm{un}(a) = \langle \mathrm{id}, a\rangle$, for every $a \in K$. Consequently, $\mathrm{un} \circ i = \mathrm{in} \circ i$, which implies that

$$\begin{aligned}\mathrm{un} \circ \rho &= \mathrm{un} \circ \mathrm{re} \circ \mathbb{X}i \\ &= \mathrm{comp} \circ \mathbb{X}\mathrm{un} \circ \mathbb{X}i \\ &= \mathrm{comp} \circ \mathbb{X}\mathrm{in} \circ \mathbb{X}i = \mathbb{X}i = \mathbb{X}\rho \circ \mathbb{X}\mathrm{in}\,.\end{aligned}$$

Furthermore,

$$\begin{aligned}\rho \circ \mathrm{comp} &= \mathrm{re} \circ \mathbb{X}i \circ \mathrm{comp} \\ &= \mathrm{re} \circ \mathrm{comp} \circ \mathbb{X}\mathbb{X}i \\ &= \mathrm{re} \circ \mathbb{X}\mathrm{re} \circ \mathbb{X}\mathbb{X}i = \mathrm{re} \circ \mathbb{X}\rho\,.\end{aligned}$$ □

The intended target for this definition are the free algebras. We start by noting that these satisfy the above conditions.

Proposition 4.21. *The free $\mathbb{T}^\times$-algebra $\langle \mathbb{T}^\times A, \mathrm{flat}^\times\rangle$ supports unravelling.*

Proof. Using the operations un and re from Definitions 4.1 and 4.4, it follows by Lemma 4.5 (e) and (b), that

$$\begin{aligned}
\mathrm{flat}^{\times} \circ \mathrm{re} \circ \mathbb{X}\mathrm{sing} &= \mathrm{flat}^{\times} \circ \mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \\
&= \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \iota) \circ \mathbb{X}\mathrm{sing} \\
&= \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \mathrm{sing}^{\times}) = \mathrm{re}\,, \\
\mathrm{un} \circ \mathrm{re} &= \mathrm{comp} \circ \mathbb{X}\mathrm{un}\,,
\end{aligned}$$

while the third condition follows by Lemma 4.14 (a). □

For the proof below, let us collect a few basic properties of algebras that support unravelling.

Lemma 4.22. *Let $\mathfrak{A}$ be a $\mathbb{T}^{\times}$-algebra that supports unravelling.*

(a) $\pi \circ \mathrm{re} = \mathrm{re} \circ \mathbb{X}\pi$

(b) $\pi \circ \mathrm{re} \circ \mathrm{dun} = \pi$

(c) $\mathrm{Up}(\mathrm{un} \circ \pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) = \mathrm{Up}(\mathbb{X}\pi) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un})$

Proof. Below we will make freely use of the equations from Lemma 4.5.

(a) We have

$$\begin{aligned}
\pi \circ \mathrm{re} &= \pi \circ \mathrm{re} \circ \mathbb{X}(\mathrm{flat}^{\times} \circ \mathrm{sing}^{\times}) \\
&= \pi \circ \mathrm{flat}^{\times} \circ \mathrm{re} \circ \mathbb{X}\mathrm{sing}^{\times} \\
&= \pi \circ \mathbb{T}^{\times}\pi \circ \mathrm{re} \circ \mathbb{X}\mathrm{sing}^{\times} \\
&= \pi \circ \mathrm{re} \circ \mathbb{X}\mathbb{T}^{\circ}\pi \circ \mathbb{X}\mathrm{sing}^{\times} \\
&= \pi \circ \mathrm{re} \circ \mathbb{X}\mathrm{sing}^{\times} \circ \mathbb{X}\pi \\
&= \mathrm{re} \circ \mathbb{X}\pi\,,
\end{aligned}$$

where the last step follows from the fact that $\mathfrak{A}$ supports unravelling.

(b) Since

$$\pi \circ \mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un} = \mathrm{re}_{\circ} \circ \mathbb{X}\pi \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un} = \mathrm{re}_{\circ} \circ \mathrm{un} = \mathrm{id}\,,$$

we have

$$\begin{aligned}\pi \circ \mathrm{re} \circ \mathrm{dun} &= \pi \circ \mathrm{re} \circ \mathrm{un} \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \pi \circ \mathrm{flat}^{\times} \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \pi \circ \mathbb{T}^{\times}\pi \circ \mathbb{T}^{\times}(\mathrm{re}_{\circ} \circ \mathbb{X}\mathrm{sing} \circ \mathrm{un}) \\ &= \pi \circ \mathbb{T}^{\times}\mathrm{id} \\ &= \pi\,.\end{aligned}$$

(c) By (a), Lemma 4.16 (e), the fact that $\mathfrak{A}$ supports unravelling, and (b), we have

$$\begin{aligned}&\mathrm{Up}(\mathrm{un} \circ \pi \circ \mathrm{re}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathrm{un} \circ \mathrm{re} \circ \mathbb{X}\pi) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathrm{comp} \circ \mathbb{X}\mathrm{un} \circ \mathbb{X}\pi) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathrm{comp} \circ \mathbb{X}(\mathrm{un} \circ \pi) \circ \mathrm{dun} \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathrm{comp} \circ \mathbb{X}(\mathrm{in} \circ \pi) \circ \mathrm{dun} \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathrm{comp} \circ \mathbb{X}(\mathrm{in} \circ \pi)) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \\ &= \mathrm{Up}(\mathbb{X}\pi) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un})\,.\end{aligned}$$

□

Finally we can state our partial distributive law for Up and $\mathbb{T}^{\times}$ for algebras that support unravelling.

Proposition 4.23. *If* $\mathfrak{A} = \langle A, \pi\rangle$ *is a* $\mathbb{T}^{\times}$*-algebra supporting unravelling, we can form a* $\mathbb{T}^{\times}$*-algebra* $\mathrm{Up}(\mathfrak{A}) := \langle \mathrm{Up}(A), \hat{\pi}\rangle$ *with product*

$$\hat{\pi} := \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un})\,.$$

Furthermore, the function $\mathrm{pt} : A \to \mathrm{Up}(A)$ *induces an embedding* $\mathfrak{A} \to \mathrm{Up}(\mathfrak{A})$.

Proof. We have to check three equations. To see that pt is an embedding, note that

$$\begin{aligned}
\hat{\pi} \circ \mathbb{T}^{\times}\mathrm{pt} &= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \circ \mathbb{T}^{\times}\mathrm{pt} \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{pt} \circ \mathbb{T}^{\times}\mathrm{un} \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{pt} \circ \mathrm{sun} \circ \mathbb{T}^{\times}\mathrm{un} \\
&= \mathrm{pt} \circ \pi \circ \mathrm{re}_{\circ} \circ \mathrm{sun} \circ \mathbb{T}^{\times}\mathrm{un} \\
&= \mathrm{pt} \circ \pi \circ \mathrm{re}_{\circ} \circ \mathrm{dun} \\
&= \mathrm{pt} \circ \pi\,,
\end{aligned}$$

where the third step follows by Lemma 4.16 (b), the fifth one by definition of dun, and the sixth one by Lemma 4.22 (b). For the unit law, we have

$$\begin{aligned}
\hat{\pi} \circ \mathrm{sing}^{\times} &= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathbb{T}^{\times}\mathrm{Up}(\mathrm{un}) \circ \mathrm{sing}^{\times} \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{sel} \circ \mathrm{sing}^{\times} \circ \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_{\circ}) \circ \mathrm{Up}(\mathbb{X}\mathrm{sing}) \circ \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\mathrm{re}_{\circ} \circ \mathrm{un}) \\
&= \mathrm{id}\,,
\end{aligned}$$

where the third step follows by Lemma 4.16 (c) and the fifth one by the definition of an unravelling structure. Finally, for the associative law,

$$\begin{aligned}
\hat{\pi} \circ \mathbb{T}^\times \hat{\pi} &= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \\
&\qquad\qquad \circ \mathbb{T}^\times \big(\mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})\big) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \\
&\qquad\qquad \circ \mathbb{T}^\times \big(\mathrm{Up}(\mathrm{un} \circ \pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})\big) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times (\mathrm{Up}(\mathbb{X}\pi) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{Up}(\mathbb{X}\mathbb{T}^\circ \pi) \circ \mathrm{sel} \circ \mathbb{T}^\times (\mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un})) \\
&= \mathrm{Up}(\pi \circ \mathbb{T}^\times \pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{sel} \circ \mathbb{T}^\times \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\pi \circ \mathrm{flat}^\times \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{sel} \circ \mathbb{T}^\times \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ \circ \mathbb{X}\mathrm{flat}^\times) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{sel} \circ \mathbb{T}^\times \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathrm{flat}^\times \circ \mathbb{T}^\times \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \\
&= \mathrm{Up}(\pi \circ \mathrm{re}_\circ) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \circ \mathrm{flat}^\times \\
&= \hat{\pi} \circ \mathrm{flat}^\times .
\end{aligned}$$

where the third, fourth, seventh, and eighth step follows by, respectively, Lemmas 4.22 (e), 4.16 (a), 4.22 (a), and 4.18. □

The following consequence can be considered the main result of this section.

Theorem 4.24. *In* Pos^{Ξ}, *the set* $\mathrm{Up}(\mathbb{T}^\times A)$ *forms a* $\mathbb{T}^\times$*-algebra with product*

$$\hat{\pi}(t) := \Uparrow\{\, \mathrm{flat}^\times({}^{\sigma}s) \mid \varphi, \sigma : s \in^{\mathrm{un}} t \,\} .$$

Proof. We know by Proposition 4.23 that $\mathrm{Up}(\mathbb{T}^\times(A \uparrow \Xi_+))$ forms a $\mathbb{T}^\times$-algebra in Pos^{Ξ_+}. Since

$$\mathrm{Up}(\mathbb{T}^\times A) = \mathrm{Up}(\mathbb{T}^\times(A \uparrow \Xi_+))|_{\Xi} ,$$

the claim follows by Lemma II.3.5 (c). □

In order to strengthen this theorem to obtain a $(\mathrm{Up} \circ \mathbb{T}^\times)$-algebra, we would need to prove that $\mathrm{Up} \circ \mathbb{T}^\times$ forms a monad. The next result shows that the canonical choice for the corresponding monad multiplication does not work. (Note that this is not a simple consequence of Theorem 2.7 since it might be the case that, instead of condition (M1) of Theorem I.6.7 (4), it is (M2) or (M3) that is violated.)

Proposition 4.25. *The function* $\kappa : (\mathrm{Up} \circ \mathbb{T}^\times)(\mathrm{Up} \circ \mathbb{T}^\times)A \to (\mathrm{Up} \circ \mathbb{T}^\times)A$ *with*

$$\kappa(T) := \Uparrow\{\, \mathrm{flat}^\times({}^\sigma s) \mid \varphi, \sigma : s \in^{\mathrm{un}} t\,,\ t \in T \,\}$$

does not *satisfy the associative law*

$$\kappa \circ \kappa = \kappa \circ \mathrm{Up}(\mathbb{T}^\times \kappa)\,.$$

Proof. We use term notation $a(c), b(c,d), \dots$ for trees. Note that, for two sets

$$X = \{\, a_i(x_0, x_0) \mid i < m \,\} \quad \text{and} \quad Y = \{\, \mathrm{sing}^\times(c_i) \mid i < n \,\}$$

(where $a_i \in A_{\{x,y\}}$ and $c_i \in A_\varnothing$) we have

$$\begin{aligned}
\kappa(\{X(Y)\}) &= \{\, \mathrm{flat}^\times({}^\sigma s) \mid \varphi, \sigma : s \in^{\mathrm{un}} X(Y) \,\} \\
&= \Big\{\, \mathrm{flat}^\times({}^\sigma s) \;\Big|\; s = u(v,w),\ u = \mathrm{sing}^\times(a_i),\ i < m\,, \\
&\qquad\qquad v = \mathrm{sing}^\times(c_k),\ w = \mathrm{sing}^\times(c_l),\ k, l < n \,\Big\} \\
&= \{\, a_i(c_k, c_l) \mid i < m\,,\ k, l < n \,\}\,.
\end{aligned}$$

Similarly, if the $a_i \in A_1$ are unary, we obtain

$$\kappa(\{X(Y)\}) = \{\, a_i(c_k) \mid i < m\,,\ k < n \,\}\,.$$

Setting

$$\begin{aligned}
I &:= \{a(x_0, x_0)\}\,, & C &:= \{c\}\,, \\
J &:= \{b(x_0, x_0)\}\,, & D &:= \{d\}\,, \\
K &:= \{\mathrm{sing}^\times(I), \mathrm{sing}^\times(J)\}\,, & E &:= \{\mathrm{sing}^\times(C), \mathrm{sing}^\times(D)\}\,,
\end{aligned}$$

we obtain

$$\kappa(\{K(E)\}) = \{I(C), I(D), J(C), J(D)\}\,,$$

$$\kappa(\{I(C)\}) = \{a(c,c)\}\,, \qquad \kappa(\{I(D)\}) = \{a(d,d)\}\,,$$
$$\kappa(\{J(C)\}) = \{b(c,c)\}\,, \qquad \kappa(\{J(D)\}) = \{b(d,d)\}\,,$$

$$(\kappa \circ \kappa)(\{K(E)\}) = \{a(c,c), a(d,d), b(c,c), b(d,d)\}\,,$$

$$\kappa(K) = I \cup J \quad =: X\,,$$
$$\kappa(E) = C \cup D =: Y\,,$$

$$\mathrm{Up}(\mathbb{T}^\times\kappa)(\{K(E)\}) = \{X(Y)\}\,,$$

$$(\kappa \circ \mathrm{Up}(\mathbb{T}^\times\kappa))(\{K(E)\}) = \{\, u(v,w) \mid u \in \{a,b\}\,,\ v,w \in \{c,d\}\,\}\,.$$

Hence,

$$(\kappa \circ \kappa)(\{K(D)\}) \neq (\kappa \circ \mathrm{Up}(\mathbb{T}^\times\kappa))(\{K(D)\})\,.$$

(For instance, the tree $a(c,d)$ does belong to the right-hand side, but not to the left-hand one.) $\square$

5. Substitutions

As a first simple application of the tools we have developed in Section 4, let us take a look at substitutions for tree languages.

Definition 5.1. Let Σ be an alphabet.

(a) A *substitution* is a function $\sigma : X \to \mathrm{Up}(\mathbb{T}^\times\Sigma)$. We call σ *regular* if every $\sigma(x) \subseteq \mathbb{T}^\times\Sigma$ is a regular tree language.

(b) Every substitution $\sigma : X \to \mathrm{Up}(\mathbb{T}^\times\Sigma)$ induces two different function $\mathbb{T}^\times X \to \mathrm{Up}(\mathbb{T}^\times\Sigma)$. The *inside-out* morphism σ_{io} is defined by

$$\sigma_{\mathrm{io}}(t) := \{\, \mathrm{flat}^\times(s) \mid s \in^{\mathbb{R}} \mathbb{R}\sigma(t)\,\}\,,$$

while the *outside-in* morphism σ_{oi} is defined by

$$\sigma_{\mathrm{oi}}(t) := \{\, \mathrm{flat}^\times({}^{\sigma}s) \mid \varphi, \sigma : s \in^{\mathrm{un}} \mathbb{R}\sigma(t)\,\}\,.$$ ⌟

Remark. (a) Intuitively, the difference between these two variants is that, with the inside-out version σ_{io}, we have to choose the same image $s(u) \in \sigma(t(v))$ for every vertex u of s corresponding to $v \in \mathrm{dom}_+(t)$, while the outside-in σ_{oi} version allows us to choose a different tree for each of them. The former has the advantage of simplicity, but the latter turns out to be more natural from an algebraic perspective: we will show below that it forms a morphism of $\mathbb{T}^\times$-algebras.

(b) In the notation of Section 4, we can rewrite the above definitions as

$$\begin{aligned}
\sigma_{\mathrm{io}} &= \mathrm{Up}(\mathrm{flat}^\times) \circ \mathrm{dist} \circ \mathbb{T}^\times \sigma \,, \\
\sigma_{\mathrm{oi}} &= \mathrm{Up}(\mathrm{flat}^\times \circ \mathrm{re}_{\mathrm{o}}) \circ \mathrm{sel} \circ \mathbb{T}^\times (\mathrm{Up}(\mathrm{un}) \circ \sigma) \,.
\end{aligned}$$

Hence, σ_{io} is based on the failed distributive law dist, while σ_{oi} is based on the more successful attempt using the relation $\in^{\mathrm{un}}$. ⌟

For the next lemma, let us recall from Theorem 4.24 that $\mathrm{Up}(\mathbb{T}^\times \Sigma)$ indeed forms a $\mathbb{T}^\times$-algebra.

Lemma 5.2. *$\sigma_{\mathrm{oi}} : \mathbb{T}^\times X \to \mathrm{Up}(\mathbb{T}^\times \Sigma)$ is a morphism of $\mathbb{T}^\times$-algebras.*

Proof. According to Theorem 4.24, the product of the algebra $\mathrm{Up}(\mathbb{T}^\times \Sigma)$ is given by

$$\hat{\pi} := \mathrm{Up}(\mathrm{flat}^\times \circ \mathrm{re}_{\mathrm{o}}) \circ \mathrm{sel} \circ \mathbb{T}^\times \mathrm{Up}(\mathrm{un}) \,.$$

Hence, $\sigma_{\mathrm{oi}} = \hat{\pi} \circ \mathbb{T}^\times \sigma$ and it follows that

$$\begin{aligned}
\sigma_{\mathrm{oi}} \circ \mathrm{flat}^\times &= \hat{\pi} \circ \mathbb{T}^\times \sigma \circ \mathrm{flat}^\times \\
&= \hat{\pi} \circ \mathrm{flat}^\times \circ \mathbb{T}^\times \mathbb{T}^\times \sigma \\
&= \hat{\pi} \circ \mathbb{T}^\times \hat{\pi} \circ \mathbb{T}^\times \mathbb{T}^\times \sigma = \hat{\pi} \circ \mathbb{T}^\times \sigma_{\mathrm{oi}} \,.
\end{aligned}$$ □

Remark. Note that the function $\sigma_{\mathrm{io}} : \mathbb{T}^\times X \to \mathrm{Up}(\mathbb{T}^\times \Sigma)$ is *not* a morphism of $\mathbb{T}^\times$-algebras. ⌟

For the simpler inside-out substitutions, we can solve inequalities of the form $\rho_{\mathrm{io}}[L] \subseteq R$ as follows.

Theorem 5.3. *Let $L \subseteq \mathbb{T}^\times X$ and $R \subseteq \mathbb{T}^\times \Sigma$ be regular tree languages, $\sigma, \tau : X \to \mathrm{Up}(\mathbb{T}^\times \Sigma)$ regular substitutions, and let S be the set of all substitutions ρ such that*

$$\sigma \subseteq \rho \subseteq \tau \quad \textit{and} \quad \rho_{\mathrm{io}}[L] \subseteq R\,.$$

Then

(a) *S has finitely many maximal elements.*

(b) *Every maximal element of S is regular.*

(c) *We can effectively compute the maximal elements of S.*

Proof. Since R is regular, it is recognised by some morphism $\eta : \mathbb{T}^\times \Sigma \to \mathfrak{A}$ into an MSO-definable $\mathbb{T}^\times$-algebra $\mathfrak{A} = \langle A, \pi \rangle$. We define the *saturation* $\hat{\rho} : X \to \mathrm{Up}(\mathbb{T}^\times \Sigma)$ of a given substitution $\rho : X \to \mathrm{Up}(\mathbb{T}^\times \Sigma)$ by

$$\hat{\rho}(x) := \big\{ s \in \mathbb{T}^\times \Sigma \bigm| \eta(s) \in \mathrm{Up}(\eta)(\rho(x)) \big\}\,.$$

Then we have $\mathrm{Up}(\eta) \circ \hat{\rho} = \mathrm{Up}(\eta) \circ \rho$. Since we can phrase the definition of ρ_{io} as

$$\rho_{\mathrm{io}} = \mathrm{Up}(\mathrm{flat}^\times) \circ \mathrm{dist} \circ \mathbb{T}^\times \rho\,,$$

it follows that

$$\begin{aligned}
\mathrm{Up}(\eta) \circ \rho_{\mathrm{io}} &= \mathrm{Up}(\eta \circ \mathrm{flat}^\times) \circ \mathrm{dist} \circ \mathbb{T}^\times \rho \\
&= \mathrm{Up}(\pi \circ \mathbb{T}^\times \eta) \circ \mathrm{dist} \circ \mathbb{T}^\times \rho \\
&= \mathrm{Up}(\pi) \circ \mathrm{dist} \circ \mathbb{T}^\times \mathrm{Up}(\eta) \circ \mathbb{T}^\times \rho \\
&= \mathrm{Up}(\pi) \circ \mathrm{dist} \circ \mathbb{T}^\times \mathrm{Up}(\eta) \circ \mathbb{T}^\times \hat{\rho} \\
&= \mathrm{Up}(\pi \circ \mathbb{T}^\times \eta) \circ \mathrm{dist} \circ \mathbb{T}^\times \hat{\rho} \\
&= \mathrm{Up}(\eta \circ \mathrm{flat}^\times) \circ \mathrm{dist} \circ \mathbb{T}^\times \hat{\rho} \\
&= \mathrm{Up}(\eta) \circ \hat{\rho}_{\mathrm{io}}\,.
\end{aligned}$$

As $\eta(s) = \eta(t)$ implies $s \in R \Leftrightarrow t \in R$, it therefore follows that

$$\rho_{\mathrm{io}}(t) \subseteq R \quad \text{implies} \quad \hat{\rho}_{\mathrm{io}}(t) \subseteq R\,.$$

Since $\rho \subseteq \hat{\rho}$ this implies that the maximal elements of S satisfy $\rho = \hat{\rho} \cap \tau$. In particular, a substitution of this form is regular. This proves (b).

For (a), note that the number of substitutions of the form $\hat{\rho}$ is bounded by the number of functions $X \to \mathrm{Up}(A)$. As X is finite and A is sort-wise finite, there are only finitely many of them.

It remains to establish (c). We can enumerate all functions $X \to \mathrm{Up}(A)$. This gives an enumeration of all substitutions of the form $\hat{\rho}$. For each of them, we can check whether $\sigma \leq \hat{\rho} \cap \tau$. If so, $\hat{\rho} \cap \tau$ is a maximal element of S. Otherwise, it is not. □

The more complicated case of outside-in substitutions is still open.

Remark. There is one technical detail worth mentioning: the way we have defined substitutions, every tree in $\sigma(x)$, for $x \in X_\xi$, contains *all* variables in ξ. But usually one uses a more general notion of a substitution where the trees in $\sigma(x)$ can omit some or all of these variables. We can formalise this generalisation in our setting as follows.

We consider a substitution as a function $\sigma : X \to \mathrm{Up}(\mathbb{T}^{?}\Sigma)$, where $\mathbb{T}^{?}$ is the functor with

$$\mathbb{T}^{?}_{\xi} X := \sum_{\zeta \subseteq \xi} \mathbb{T}^{\times}_{\zeta} X$$

(for details, see Section VIII.1 below). We can extend the monad operation to $\mathbb{T}^{?}$ in the obvious way. As above we define two induced operations $\sigma_{\mathrm{io}}, \sigma_{\mathrm{oi}} : \mathbb{T}^{?}X \to \mathrm{Up}(\mathbb{T}^{?}\Sigma)$. The definition of the outside-in version is the same as above

$$\sigma_{\mathrm{oi}}(t) := \{\, \mathrm{flat}^{?}({}^{\sigma}s) \mid \varphi, \sigma : s \in^{\mathrm{un}} \mathbb{R}^{?}\sigma(t) \,\}$$

(where $\mathbb{R}^{?}$ is the corresponding variant of $\mathbb{R}$).

But the inside-out version is more complicated. The problem is that some sets $\sigma(x)$ might be empty, but a tree t might still have a non-empty image $\sigma_{\mathrm{io}}(t)$ because, for every vertex v with $\sigma(t(v)) = \varnothing$, there might be some vertex u higher up in the tree where we have chosen an element $s \in t(u)$ which omits the variable corresponding to the subtree containing v. The easiest way to formalise this process is to make the problem disappear by adding dummy elements to all sets $\sigma(x)$. Hence, fix some element $\perp \notin \Sigma$ and let $\mu : \mathrm{Up}(\Sigma) \to \mathrm{Up}(\Sigma + \{\perp\})$ be the function with

$$\mu(I) := I \cup \{\perp\}.$$

Then we set

$$\sigma_{\mathrm{io}}(t) := \left\{ \operatorname{flat}^{?}(s) \;\middle|\; s \in^{\mathbb{R}^{?}} \mathbb{R}^{?}(\mu \circ \sigma)(t)\,,\ \operatorname{flat}^{?}(s) \in \mathbb{T}^{?}\Sigma \right\}.$$

The proof of Theorem 5.3 can now straightforwardly be adapted to these new definitions. ⌟

6. Regular Expressions

As a second, more involved application of our results let us define regular expressions for languages of infinite trees. We consider tree languages of the form $L \subseteq \mathbb{T}^{\times}_{\xi}\Sigma$, for some alphabet Σ and some fixed sort $\xi \in \Xi$. Alphabets will always be assumed to be finite and unordered. Note that, if Σ is unordered, so is $\mathbb{T}^{\times}_{\xi}\Sigma$ and $\mathrm{Up}(\mathbb{T}^{\times}_{\xi}\Sigma)$ is just the power set. Hence, we can regard every language $L \subseteq \mathbb{T}^{\times}_{\xi}\Sigma$ as an element of $\mathrm{Up}(\mathbb{T}^{\times}_{\xi}\Sigma)$.

We aim for a characterisation of which elements of this set are regular languages. Towards this goal we introduce a few operations on $\mathrm{Up}(\mathbb{T}^{\times}\Sigma)$.

Before presenting the definition we need to deal with the problem that $\mathrm{Up} \circ \mathbb{T}^{\times}$ does not form a monad and that $\mathrm{Up}(\mathbb{T}^{\times}\Sigma)$ not a $(\mathrm{Up} \circ \mathbb{T}^{\times})$-algebra. For this reason we will work with $\mathbb{T}^{\times}$, Up-bialgebras $\mathfrak{A} = \langle A, \pi, \rho\rangle$. By Theorem 4.24, $\mathrm{Up}(\mathbb{T}^{\times}\Sigma)$ forms a $\mathbb{T}^{\times}$, Up-bialgebra with products.

$$\begin{aligned} \hat{\pi} &: \mathbb{T}^{\times}\mathrm{Up}(\mathbb{T}^{\times}\Sigma) \to \mathrm{Up}(\mathbb{T}^{\times}\Sigma) \\ \mathrm{comp} &: \mathrm{Up}(\mathrm{Up}(\mathbb{T}^{\times}\Sigma)) \to \mathrm{Up}(\mathbb{T}^{\times}\Sigma)\,. \end{aligned}$$

We use the following operations for our version of regular expressions:

- variables $x \in X$,
- letters of the alphabet $a \in \Sigma$,

- substitution $\cdot_x$, iteration $-^{+x}$, and ω-power $-^{\omega x}$ with respect to a single variable x,
- relabelling ${}^{\sigma}-$ of the variables,
- union $+$ and the empty language $\varnothing$.

The formal definition is as follows.

Definition 6.1. Given a bialgebra $\mathfrak{A} = \langle A, \pi, \rho \rangle$ we define the following operations.

(a) Each $a \in A_\xi$, induces an operation $a : A^\xi \to A$ by

$$a(\bar{b}) := \pi(s)\,,$$

where $s \in \mathbb{T}^\times_\xi A$ is the tree obtained form $\operatorname{sing}(a)$ by replacing each leaf with label $x \in \xi$ by the tree $\operatorname{sing}(b_x)$.

(b) For sorts $\xi, \zeta \in \Xi$ and a variable $x \in \xi$, we define a binary *substitution operation*

$$\cdot_x \;: A_\xi \times A_\zeta \to A_{(\xi \setminus \{x\}) \cup \zeta} \quad \text{by} \quad a \cdot_x b := \pi(s)\,,$$

where s is the tree obtained from $\operatorname{sing}(a)$ by replacing the leaf labelled x by the tree $\operatorname{sing}(b)$.

(c) For $a \in A_\xi$ and a surjective map $\sigma : \xi \to \zeta$, we set

$${}^{\sigma}a := \pi(s)\,,$$

where s is the tree obtained from $\operatorname{sing}(a)$ by replacing each label $x \in \xi$ by $\sigma(x)$.

(d) We define $+ : A_\xi \times A_\xi \to A_\xi$ and $\varnothing \in A_\xi$ by

$$a + b := \rho(\Uparrow\{a, b\}) \quad \text{and} \quad \varnothing := \rho(\varnothing)\,.$$

(e) Let $\zeta \in \Xi$. We call a tree s *ζ-trivial* if, for all $v \in \operatorname{dom}_+(s)$ and $z \in \zeta$, we have

$$s(v) = z \quad \text{iff} \quad v \text{ is an } z\text{-successor.}$$

(That is, all z-successors are labelled by z and there are no other occurrences of z.) For a finite sequence of elements $a_i \in A_{\xi_i}$, $i < n$, and a variable $x \in \zeta := \xi_0 \cup \dots \cup \xi_{n-1}$, we define the *$\omega$-power* and the *iteration* by

$$(a_0 + \dots + a_{n-1})^{\omega x} := \\ \rho\big(\{\, \pi(s) \mid s \in \mathbb{T}^{\times}_{\zeta \setminus \{x\}}\{a_0, \dots, a_{n-1}\} \text{ is } (\zeta \setminus \{x\})\text{-trivial} \,\}\big)\,,$$

$$(a_0 + \dots + a_{n-1})^{+x} := \\ \rho\big(\{\, \pi(s) \mid s \in \mathbb{T}^{\times}_{\zeta}\{a_0, \dots, a_{n-1}\} \text{ has finite height and it is } (\zeta \setminus \{x\})\text{-trivial} \,\}\big)\,.$$

(f) For a sort $\xi \in \Xi$ and a set Σ, the set $\mathbb{E}_\xi\Sigma$ of *regular expression* over Σ consists of all finite terms R that can be built up from variables and the operations (a)–(e) (for the bialgebra $\mathrm{Up}(\mathbb{T}^{\times}\Sigma)$), where

- we restrict the operations from (a) to those where $a = {\Uparrow}\mathrm{sing}(c)$, for some $c \in \Sigma$, and
- the free variables are exactly those in ξ.

We write $[\![R]\!] \subseteq \mathrm{Up}(\mathbb{T}^{\times}\Sigma)$ for the value of $R \in \mathbb{E}\Sigma$ in $\mathrm{Up}(\mathbb{T}^{\times}\Sigma)$. ⌟

Remark. The iteration and the ω-power in (e) have a built-in sum operation in order to support choices between terms of different sorts, which is not possible using the normal sum operation from (d). ⌟

Examples. We consider the alphabet $\Sigma = \{a, b, c\}$ where a and b have sort $\{x, y\}$ and c has sort $\varnothing$.

(a) A regular expression for the language $\mathbb{T}^{\times}\Sigma$ is

$$E := \big((a(x,y) + b(x,y) + c)^{\omega x}\big)^{\omega y}.$$

(b) An expression for the language of all trees with an infinite branch labelled by a is given by

$$R := \big(a(x,z) + a(z,y)\big)^{\omega z} \cdot_x E \cdot_y E\,.$$

(c) Finally, the following expression describes all trees containing the letter a.

$$S := a(x,y) \cdot_x E \cdot_y E + \big(b(x,z) + b(z,y)\big)^{+z} \cdot_z a(x,y) \cdot_x E \cdot_y E\,. \lrcorner$$

We still have to show that regular expressions capture the class of regular languages. Recall from Definitions V.5.1 and V.5.6 the notion of a forest automaton, a partial run, and its profile. We order profiles as follows.

Definition 6.2. Let $\mathcal{A}$ be a forest automaton. We define an ordering on profiles by

$$\langle p, \bar{U}\rangle \leq \langle p', \bar{U}'\rangle \quad :\text{iff} \quad p = p' \text{ and } U_z \subseteq U'_z \text{ for all } z \in \zeta\,.$$

If $\sigma \leq \tau$, we say that the profile σ is *bounded* by τ. ⌟

Theorem 6.3. *Let Σ be an alphabet. A language $L \subseteq \mathbb{T}^\times_\zeta \Sigma$ is regular if, and only if, $L = [\![R]\!]$, for some regular expression $R \in \mathbb{E}_\zeta \Sigma$.*

Proof. $(\Leftarrow)$ The class of all regular tree languages is closed under all operations that can appear in a regular expression.

$(\Rightarrow)$ Let $\mathcal{A} = \langle Q, \Sigma, \zeta, \Delta, q_{\mathrm{I}}, \Omega\rangle$ be an automaton recognising L and fix an enumeration $q_0, \dots, q_{n-1}$ of Q such that $\Omega(q_0) \geq \dots \geq \Omega(q_{n-1})$. For every profile τ of $\mathcal{A}$ and every number $k \leq n$, we will construct a regular expressions R^k_τ defining the language

$$[\![R^k_\tau]\!] = \Big\{ t \in \mathbb{T}^\times\Sigma \;\Big|\; \begin{array}{l}\text{there is a partial run on } t \text{ whose profile is}\\ \text{bounded by } \tau \text{ and whose internal states are}\\ \text{among } q_0, \dots, q_{k-1}\end{array} \Big\}\,.$$

Then we obtain the desired expression for L by setting

$$R := \sum_{\tau \in H} R^n_\tau\,,$$

where H is the set of all profiles $\tau = \langle q_{\mathrm{I}}, \bar{U}\rangle$ such that, for all $z \in \zeta$,

$$\langle k, p\rangle \in U_z \quad \text{implies} \quad \langle p, z\rangle \in \Delta\,.$$

We define the expressions R^k_τ by induction on k. For $k = 0$, we only need to consider runs without internal states. Hence, we can set

$$R^0_\tau := \sum \{ a(\bar{x}) \mid a \in \Sigma, \text{ there is a partial run on } \operatorname{sing}(a) \text{ whose profile is bounded by } \tau \}.$$

For the inductive step, suppose that $\tau = \langle p, \bar{U} \rangle$, let ξ be the sort of τ, and let $D := \operatorname{rng} \Omega$ be the set of priorities used by $\mathcal{A}$. We start with an expression describing runs starting with the state q_k and with only finitely many occurrences of q_k on each branch. For a set $\eta \subseteq \xi$ of variables, we write $\bar{U}|_\eta$ for the subtuple $(U_x)_{x \in \eta}$. Let $V := D \times \{q_k\}$, let $y_0, y_1, \dots$ be new variables not in ξ, and set

$$T^k_{\bar{U}} := R^k_{q_k, \bar{U}} + \sum_{\zeta \cup \eta_0 \cup \dots \cup \eta_{n-1} = \xi} (S^{\zeta,n}_0 + \dots + S^{\zeta,n}_{m-1})^{+y_0} \cdot_{y_0} R^k_{q_k, \bar{U}|_{\eta_0}} \cdot_{y_1} \dots \cdot_{y_{n-1}} R^k_{q_k, \bar{U}|_{\eta_{n-1}}},$$

where

- the sum ranges over all sequences $\zeta, \eta_0, \dots, \eta_{n-1}$ of subsets of ξ whose union is equal to ξ and such that $\eta_i \neq \eta_j$, for $i \neq j$, and
- $S^{\zeta,n}_0, \dots, S^{\zeta,n}_{m-1}$ is an enumeration of all expressions of the form

  $$R^k_{q_k, \bar{U}|_\eta V \dots V} \quad \text{where} \quad \eta \subseteq \zeta \text{ and } v \subseteq \{y_0, \dots, y_{n-1}\}$$

  (with $|v|$ copies of V that correspond to the variables $y \in v$).

Then $T^k_{\bar{U}}$ describes all trees that have a run with profile bounded by $\langle q_k, \bar{U} \rangle$ and such that every branch contains only finitely many occurrences of the state q_k.

Similarly, we obtain an expression for all such trees with possibly infinitely many occurrences of q_k by setting

$$\hat{T}^k_{\bar{U}} := T^k_{\bar{U}} + \sum_{\zeta \cup \eta_0 \cup \dots \cup \eta_{n-1} = \xi} (S^{\zeta,n}_0 + \dots + S^{\zeta,n}_{m-1})^{\omega z} \cdot_{y_0} T^k_{\bar{U}|_{\eta_0}} \cdot_{y_1} \dots \cdot_{y_{n-1}} T^k_{\bar{U}|_{\eta_{n-1}}},$$

where the $S_i^{\zeta,n}$ are defined as above, except that there is an additional copy of V corresponding to the variable z.

If $\Omega(q_k)$ is odd, we can now set

$$R_\tau^{k+1} := R_\tau^k + \sum_{\zeta \cup \eta_0 \cup \dots \cup \eta_{n-1} = \xi} R^k_{p,\bar{U}|_\zeta V \dots V} \cdot_{y_0} T^k_{\bar{U}|_{\eta_0}} \cdot_{y_1} \cdots \cdot_{y_{n-1}} T^k_{\bar{U}|_{\eta_{n-1}}} .$$

where the variables $y_0, \dots, y_{n-1}$ are the ones corresponding to the n copies of the set V. If $\Omega(q_k)$ is even, we instead use

$$R_\tau^{k+1} := R_\tau^k + \sum_{\zeta \cup \eta_0 \cup \dots \cup \eta_{n-1} = \xi} R^k_{p,\bar{U}|_\zeta V \dots V} \cdot_{y_0} \hat{T}^k_{\bar{U}|_{\eta_0}} \cdot_{y_1} \cdots \cdot_{y_{n-1}} \hat{T}^k_{\bar{U}|_{\eta_{n-1}}} . \square$$

7. Distributive Lattices

Note that Up,Dn-bialgebras are the same as complete lattices. Such a lattice is distributive if joins distribute over meets and vice versa. In this section, we show that distributivity for lattices corresponds to a distributive law for the functors Up and Dn. To avoid confusion we use superscripts in the monad operations to indicate which of the two monads $\langle \mathrm{Up}, \mathrm{union}^{\mathrm{Up}}, \mathrm{pt}^{\mathrm{Up}} \rangle$ and $\langle \mathrm{Dn}, \mathrm{union}^{\mathrm{Dn}}, \mathrm{pt}^{\mathrm{Dn}} \rangle$ we are talking about.

We follow our usual strategy of construction a distributive law by lifting the functor Dn to the Kleisli category of Up. When doing so, it turns out to be convenient to make a detour through a category of certain relations.

Definition 7.1. (a) We denote by $\mathcal{R}$ the category whose objects are all sets $A \in \mathsf{Pos}^\Xi$ and whose morphisms $R : A \to B$ are relations $R \subseteq A \times B$ satisfying

$$\begin{aligned} \langle a, b \rangle \in R \text{ and } a' \geq a \quad &\text{implies} \quad \langle a', b \rangle \in R , \\ \langle a, b \rangle \in R \text{ and } b' \leq b \quad &\text{implies} \quad \langle a, b' \rangle \in R . \end{aligned}$$

The composition of two relations $R : A \to B$ and $S : B \to C$ is given by

$$S \circ R := \{ \langle a, c \rangle \mid \langle a, b \rangle \in R,\ \langle b, c \rangle \in S \} .$$

(b) We say that a relation $R \in \mathcal{R}(A, B)$ *represents* the function $f : A \to \mathrm{Dn}(B)$ defined by

$$f(a) := \{\, b \in B \mid \langle a, b \rangle \in R \,\}\,.$$

(c) The *graph* of a function $f : A \to \mathrm{Dn}(B)$ is the relation

$$\mathrm{G}(f) := \{\, \langle a, b \rangle \in A \times B \mid b \in f(a) \,\}\,.$$ ⌟

Lemma 7.2. *The operation* G *mapping a function to its graph induces an isomorphism* $\mathsf{Free}(\mathrm{Dn}) \cong \mathcal{R}$.

Proof. We can uniquely represent each morphism $\varphi : \mathrm{Dn}(A) \to \mathrm{Dn}(B)$ by its restriction $\varphi \circ \mathrm{pt} : A \to \mathrm{Dn}(B)$. Using this representation of $\mathsf{Free}(\mathrm{Dn})$, the composition of two morphisms $f : A \to \mathrm{Dn}(B)$ and $g : B \to \mathrm{Dn}(C)$ is given by

$$g \,\square\, f := \mathrm{union} \circ \mathrm{Dn}(g) \circ f\,.$$

We start by checking that G induces a functor $\mathsf{Free}(\mathrm{Dn}) \to \mathcal{R}$. Let $f : A \to \mathrm{Dn}(B)$. To see that $\mathrm{G}(f) \in \mathcal{R}(A, B)$, consider a pair $\langle a, b \rangle \in \mathrm{G}(f)$ and elements $a' \geq a$ and $b' \leq b$. Then

- $b \in f(a) \subseteq f(a')$ implies $b \in f(a')$,
- $b \in f(a)$ implies $b' \in f(a)$.

Hence, $\langle a', b \rangle \in \mathrm{G}(f)$ and $\langle a, b' \rangle \in \mathrm{G}(f)$.

Since $\mathrm{G}(\mathrm{pt}) = \{\, \langle a, b \rangle \mid a \leq b \,\}$, G maps the identity to the identity. Finally, we have

$$\begin{aligned} \mathrm{G}(g \,\square\, f) &= \{\, \langle a, c \rangle \mid c \in (g \,\square\, f)(a) \,\} \\ &= \{\, \langle a, c \rangle \mid c \in g(b) \text{ for some } b \in f(a) \,\} \\ &= \{\, \langle a, c \rangle \mid \langle b, c \rangle \in \mathrm{G}(g) \text{ and } \langle a, b \rangle \in \mathrm{G}(f) \,\} \\ &= \mathrm{G}(g) \circ \mathrm{G}(f)\,. \end{aligned}$$

It remains to prove that this functor is bijective. Clearly, $f \neq g$ implies $\mathrm{G}(f) \neq \mathrm{G}(g)$. For surjectivity, let $R \in \mathcal{R}(A, B)$. We claim that $R = \mathrm{G}(f)$, for the function $f : A \to \mathrm{Dn}(B)$ defined by

$$f(a) := \{\, b \in B \mid \langle a, b\rangle \in R \,\}.$$

We have to check three conditions.

(I) f is monotone.
(II) The codomain of f is $\mathrm{Dn}(B)$.
(III) $\mathrm{G}(f) = R$.

(I) Let $a \leq a'$. Then $f(a) \subseteq f(a')$ since

$$b \in f(a) \quad\Rightarrow\quad \langle a, b\rangle \in R \quad\Rightarrow\quad \langle a', b\rangle \in R \quad\Rightarrow\quad b \in f(a').$$

(II) To see that $f(a)$ is downwards closed, let $b' \leq b$. Then

$$b \in f(a) \quad\Rightarrow\quad \langle a, b\rangle \in R \quad\Rightarrow\quad \langle a, b'\rangle \in R \quad\Rightarrow\quad b' \in f(a).$$

(III) We have

$$\mathrm{G}(f) = \{\, \langle a, b\rangle \mid b \in f(a) \,\} = \{\, \langle a, b\rangle \mid \langle a, b\rangle \in R \,\} = R. \qquad \square$$

The advantage of encoding Dn-morphisms $\varphi : \mathrm{Dn}(A) \to \mathrm{Dn}(B)$ by relations $R \subseteq A \times B$ is that the functor Dn vanishes that way. Hence, to find an extension of Up, we can simply apply Up to R and then translate back.

Lemma 7.3. *There exists an extension* $\widehat{\mathrm{Up}} : \mathsf{Free}(\mathrm{Dn}) \to \mathsf{Free}(\mathrm{Dn})$ *of* $\mathrm{Up} : \mathsf{Pos}^{\Xi} \to \mathsf{Pos}^{\Xi}$ *to* $\mathsf{Free}(\mathrm{Dn})$ *such that*

$$\begin{aligned}&\widehat{\mathrm{Up}}(\varphi)(\mathrm{pt}^{\mathrm{Dn}}(I)) = \\ &\qquad \{\, J \in \mathrm{Up}(B) \mid J \cap \varphi(\mathrm{pt}^{\mathrm{Dn}}(a)) \neq \varnothing \textit{ for all } a \in I \,\},\end{aligned}$$

for every Dn*-morphism* $\varphi : \mathrm{Dn}(A) \to \mathrm{Dn}(B)$ *and all* $I \in \mathrm{Up}(A)$.

Proof. By Lemma 7.2, it is sufficient to define a functor $\widehat{\mathrm{Up}} : \mathcal{R} \to \mathcal{R}$. We map a set $A \in \mathcal{R}$ to $\widehat{\mathrm{Up}}(A) := \mathrm{Up}(A)$ and a relation $R : A \to B$ to

$$\widehat{\mathrm{Up}}(R) := \{\, \langle I, J\rangle \in \mathrm{Up}(A) \times \mathrm{Up}(B) \;\big|\; I \subseteq R/J \,\},$$

where

$$R/J := \{\, a \in A \mid (\{a\} \times J) \cap R \neq \varnothing \,\}\,.$$

Let us start by showing that $\widehat{\mathrm{Up}}$ extends Up. For objects A this is trivial since we have $\widehat{\mathrm{Up}}(A) = \mathrm{Up}(A)$ by definition. For morphisms, note that a function $f : A \to B$ of Pos^{Ξ} corresponds to the morphism $\mathrm{pt}^{\mathrm{Dn}} \circ f : A \to \mathrm{Dn}(B)$ of $\mathsf{Free}(\mathrm{Dn})$, and that

$$\begin{aligned}
&\widehat{\mathrm{Up}}(\mathrm{G}(\mathrm{pt}^{\mathrm{Dn}} \circ f)) \\
&= \big\{\, \langle I, J\rangle \bigm| I \subseteq \mathrm{G}(\mathrm{pt}^{\mathrm{Dn}} \circ f)/J \,\big\} \\
&= \big\{\, \langle I, J\rangle \bigm| \text{for every } a \in I \text{ there is } b \in J \text{ with } b \in \mathrm{pt}^{\mathrm{Dn}}(f(a)) \,\big\} \\
&= \big\{\, \langle I, J\rangle \bigm| \text{for every } a \in I \text{ there is } b \in J \text{ with } b \leq f(a) \,\big\} \\
&= \big\{\, \langle I, J\rangle \bigm| J \supseteq f[I] \,\big\} \\
&= \big\{\, \langle I, J\rangle \bigm| J \leq \mathrm{Up}(f)(I) \,\big\} \\
&= \mathrm{G}(\mathrm{pt}^{\mathrm{Dn}} \circ \mathrm{Up}(f))\,.
\end{aligned}$$

Furthermore, for a Dn-morphism $\varphi : \mathrm{Dn}(A) \to \mathrm{Dn}(B)$ and a set $I \in \mathrm{Up}(A)$, we have

$$\begin{aligned}
&\widehat{\mathrm{Up}}(\mathrm{G}(\varphi \circ \mathrm{pt}^{\mathrm{Dn}}))(\mathrm{pt}^{\mathrm{Dn}}(I)) \\
&= \{\, J \mid I \subseteq \mathrm{G}(\varphi \circ \mathrm{pt}^{\mathrm{Dn}})/J \,\} \\
&= \{\, J \mid \text{for all } a \in I \text{ there is } b \in J \text{ with } \langle a, b\rangle \in \mathrm{G}(\varphi \circ \mathrm{pt}^{\mathrm{Dn}}) \,\} \\
&= \{\, J \mid \text{for all } a \in I \text{ there is } b \in J \text{ with } b \in (\varphi \circ \mathrm{pt}^{\mathrm{Dn}})(a) \,\} \\
&= \{\, J \mid J \cap (\varphi \circ \mathrm{pt}^{\mathrm{Dn}})(a) \neq \varnothing \text{ for all } a \in I \,\}\,.
\end{aligned}$$

It remains to show that $\widehat{\mathrm{Up}}$ is a functor. Let $R \in \mathcal{R}(A, B)$. To check that $\widehat{\mathrm{Up}}(R) \in \mathcal{R}(\mathrm{Up}(A), \mathrm{Up}(B))$, fix $\langle I, J\rangle \in \widehat{\mathrm{Up}}(R)$ and $I' \geq I$, $J' \leq J$. Then

$I' \subseteq I$ and $J' \supseteq J$. Hence,

$$\begin{aligned}
\langle I, J\rangle \in \widehat{\mathrm{Up}}(R) \quad &\Rightarrow \quad I \subseteq R/J \\
&\Rightarrow \quad I' \subseteq R/J \quad \Rightarrow \quad \langle I', J\rangle \in \widehat{\mathrm{Up}}(R)\,, \\
\langle I, J\rangle \in \widehat{\mathrm{Up}}(R) \quad &\Rightarrow \quad I \subseteq R/J \\
&\Rightarrow \quad I \subseteq R/J' \quad \Rightarrow \quad \langle I, J'\rangle \in \widehat{\mathrm{Up}}(R)\,,
\end{aligned}$$

as desired.

Furthermore, $\widehat{\mathrm{Up}}$ maps identities to identities since,

$$\begin{aligned}
\widehat{\mathrm{Up}}(G(\mathrm{pt}^{\mathrm{Dn}})) &= \big\{ \langle I, J\rangle \bigm| (\{a\} \times J) \cap G(\mathrm{pt}^{\mathrm{Dn}}) \neq \varnothing \text{ for all } a \in I \big\} \\
&= \big\{ \langle I, J\rangle \bigm| {\Downarrow} a \cap J \neq \varnothing \text{ for all } a \in I \big\} \\
&= \big\{ \langle I, J\rangle \bigm| a \in J \text{ for all } a \in I \big\} \\
&= \big\{ \langle I, J\rangle \bigm| I \subseteq J \big\} \\
&= G(\mathrm{pt}^{\mathrm{Dn}})\,.
\end{aligned}$$

Finally, we have to show that $\widehat{\mathrm{Up}}$ preserves composition. Let $R \subseteq A \times B$ and $S \subseteq B \times C$. For $K \in \mathrm{Up}(C)$, we have

$$\begin{aligned}
&(S \circ R)/K \\
&= \{ a \in A \mid (\{a\} \times K) \cap (S \circ R) \neq \varnothing \} \\
&= \{ a \in A \mid \text{there are } b \in B,\ c \in K \text{ with } \langle a, b\rangle \in R,\ \langle b, c\rangle \in S \} \\
&= \{ a \in A \mid \text{there is } b \in B \text{ with } b \in S/K \text{ and } \langle a, b\rangle \in R \} \\
&= \{ a \in A \mid \text{there is } J \text{ with } J \subseteq S/K \text{ and } a \in R/J \} \\
&= \{ a \in A \mid \text{there is } J \text{ with } J \subseteq S/K \text{ and } a \in R/J \} \\
&= R/(S/K)\,.
\end{aligned}$$

Consequently,

$$
\begin{aligned}
&\widehat{\mathrm{Up}}(S \circ R) \\
&= \{\, \langle I, K\rangle \mid I \subseteq (S \circ R)/K \,\} \\
&= \{\, \langle I, K\rangle \mid I \subseteq R/(S/K) \,\} \\
&= \{\, \langle I, K\rangle \mid \text{there is } J \subseteq S/K \text{ with } I \subseteq R/J \,\} \\
&= \{\, \langle I, K\rangle \mid \text{there is } \langle J, K\rangle \in \widehat{\mathrm{Up}}(S) \text{ and } \langle I, J\rangle \in \widehat{\mathrm{Up}}(R) \,\} \\
&= \widehat{\mathrm{Up}}(S) \circ \widehat{\mathrm{Up}}(R)\,. \qquad \square
\end{aligned}
$$

Lemma 7.4. *Let* $\widehat{\mathrm{Up}}$ *be the extension of* Up *to* $\mathsf{Free}(\mathrm{Dn})$ *from Lemma 7.3, and let* $\varphi : \mathrm{Dn}(A) \to \mathrm{Dn}(B)$ *be a morphism of* $\mathsf{Free}(\mathrm{Dn})$. *Then*

$$
\begin{aligned}
\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) &= \mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) \circ \varphi\,, \\
\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) &= \mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) \circ \widehat{\mathrm{Up}}(\widehat{\mathrm{Up}}(\varphi))\,.
\end{aligned}
$$

Proof. Fix a morphism $\varphi : \mathrm{Up}(A) \to \mathrm{Up}(B)$ between free Dn-algebras and set $\varphi_\circ := \varphi \circ \mathrm{pt}^{\mathrm{Dn}} : A \to \mathrm{Dn}(B)$.

(a) For $a \in A$, we have

$$
\begin{aligned}
&\big(\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) \circ \mathrm{pt}^{\mathrm{Dn}}\big)(a) \\
&= \big(\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{pt}^{\mathrm{Dn}} \circ \mathrm{pt}^{\mathrm{Up}}\big)(a) \\
&= \{\, J \mid J \cap \varphi_\circ(b) \neq \varnothing \text{ for all } b \in \mathrm{pt}^{\mathrm{Up}}(a) \,\} \\
&= \{\, J \mid J \cap \varphi_\circ(b) \neq \varnothing \text{ for all } b \geq a \,\} \\
&= \{\, J \mid J \cap \varphi_\circ(a) \neq \varnothing \,\} \\
&= \{\, J \mid b \in J \text{ for some } b \in \varphi_\circ(a) \,\} \\
&= \{\, J \mid \mathrm{pt}^{\mathrm{Up}}(b) \subseteq J \text{ for some } b \in \varphi_\circ(a) \,\} \\
&= \{\, J \mid \mathrm{pt}^{\mathrm{Up}}(b) \geq J \text{ for some } b \in \varphi_\circ(a) \,\} \\
&= \Downarrow\{\, \mathrm{pt}^{\mathrm{Up}}(b) \mid b \in \varphi_\circ(a) \,\} \\
&= \big(\mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) \circ \varphi \circ \mathrm{pt}^{\mathrm{Dn}}\big)(a)\,.
\end{aligned}
$$

Since Dn-morphisms are determined by their restriction to $\operatorname{rng} \mathrm{pt}^{\mathrm{Dn}}$, it follows that $\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) = \mathrm{Dn}(\mathrm{pt}^{\mathrm{Up}}) \circ \varphi$.

(b) For $H \in \mathrm{Up}(\mathrm{Up}(A))$, we have

$$\begin{aligned}
&\big(\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) \circ \mathrm{pt}^{\mathrm{Dn}}\big)(H) \\
&= \big(\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{pt}^{\mathrm{Dn}} \circ \mathrm{union}^{\mathrm{Up}}\big)(H) \\
&= \big\{\, J \,\big|\, J \cap \varphi_{\circ}(a) \neq \varnothing \text{ for all } a \in \textstyle\bigcup H \,\big\} \\
&= \big\{\, J \,\big|\, \text{for all } K \in H \text{ we have } J \cap \varphi_{\circ}(a) \neq \varnothing \text{ for all } a \in K \,\big\} \\
&= \big\{\, J \,\big|\, \text{for all } K \in H \text{ there is } I \in \mathrm{Up}(J) \text{ such that} \\
&\qquad\qquad I \cap \varphi_{\circ}(a) \neq \varnothing \text{ for all } a \in K \,\big\} \\
&= \big\{\, \textstyle\bigcup L \,\big|\, \text{for all } K \in H \text{ there is } I \in L \text{ such that} \\
&\qquad\qquad\quad I \cap \varphi_{\circ}(a) \neq \varnothing \text{ for all } a \in K \,\big\} \\
&= \big\{\, \textstyle\bigcup L \,\big|\, L \cap \big\{\, I \,\big|\, I \cap \varphi_{\circ}(a) \neq \varnothing \text{ for all } a \in K \,\big\} \neq \varnothing \text{ for all } K \in H \,\big\} \\
&= \big\{\, \textstyle\bigcup L \,\big|\, L \cap \widehat{\mathrm{Up}}(\varphi)(\mathrm{pt}^{\mathrm{Dn}}(K)) \neq \varnothing \text{ for all } K \in H \,\big\} \\
&= \big(\mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) \circ \widehat{\mathrm{Up}}(\widehat{\mathrm{Up}}(\varphi)) \circ \mathrm{pt}^{\mathrm{Dn}}\big)(H)\,.
\end{aligned}$$

As above it follows that

$$\widehat{\mathrm{Up}}(\varphi) \circ \mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) = \mathrm{Dn}(\mathrm{union}^{\mathrm{Up}}) \circ \widehat{\mathrm{Up}}(\widehat{\mathrm{Up}}(\varphi))\,. \qquad \square$$

Definition 7.5. We define $\mathrm{skolem}^{\mathrm{Dn}} : \mathrm{Dn} \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathrm{Dn}$ by

$$\mathrm{skolem}^{\mathrm{Dn}}(H) := \{\, \Downarrow I \mid I \cap J \neq \varnothing, \text{ for all } J \in H \,\}\,.$$

We define $\mathrm{skolem}^{\mathrm{Up}} : \mathrm{Up} \circ \mathrm{Dn} \Rightarrow \mathrm{Dn} \circ \mathrm{Up}$ by

$$\mathrm{skolem}^{\mathrm{Up}}(H) := \{\, \Uparrow I \mid I \cap J \neq \varnothing, \text{ for all } J \in H \,\}\,. \qquad ⌟$$

Theorem 7.6. $\mathrm{skolem}^{\mathrm{Dn}} : \mathrm{Dn} \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathrm{Dn}$ *and* $\mathrm{skolem}^{\mathrm{Up}} : \mathrm{Up} \circ \mathrm{Dn} \Rightarrow \mathrm{Dn} \circ \mathrm{Up}$ *are distributive laws.*

Proof. By duality, it is sufficient to check that one of these operations is a distributive law. We take $\mathrm{skolem}^{\mathrm{Up}}$. Let $\mathbb{F} : \mathcal{D}^{\Xi} \to \mathsf{Free}(\mathrm{Dn})$ be the functor mapping a set A to the free algebra $\mathrm{Dn}(A)$ generated by it, and let $\widehat{\mathrm{Up}}$ be the functor from Lemma 7.3. We claim that $\langle \widehat{\mathrm{Up}}, \mathbb{F}\mathrm{union}^{\mathrm{Up}}, \mathbb{F}\mathrm{pt}^{\mathrm{Up}} \rangle$ is an extension of the monad $\langle \mathrm{Up}, \mathrm{union}^{\mathrm{Up}}, \mathrm{pt}^{\mathrm{Up}} \rangle$ to $\mathsf{Free}(\mathrm{Dn})$. First, note that it follows by Lemma 7.4 that

$$\mathbb{F}\mathrm{union}^{\mathrm{Up}} : \widehat{\mathrm{Up}} \circ \widehat{\mathrm{Up}} \Rightarrow \widehat{\mathrm{Up}} \quad \text{and} \quad \mathbb{F}\mathrm{pt}^{\mathrm{Up}} : \mathrm{Id} \Rightarrow \widehat{\mathrm{Up}}$$

are natural transformations. Hence, we only have to check the monad laws for $\langle \widehat{\mathrm{Up}}, \mathbb{F}\mathrm{union}^{\mathrm{Up}}, \mathbb{F}\mathrm{pt}^{\mathrm{Up}} \rangle$.

$$\begin{aligned}
\mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \widehat{\mathrm{Up}}(\mathbb{F}\mathrm{union}^{\mathrm{Up}}) &= \mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \mathbb{F}\mathrm{Up}(\mathrm{union}^{\mathrm{Up}}) \\
&= \mathbb{F}(\mathrm{union}^{\mathrm{Up}} \circ \mathrm{Up}(\mathrm{union}^{\mathrm{Up}})) \\
&= \mathbb{F}(\mathrm{union}^{\mathrm{Up}} \circ \mathrm{union}^{\mathrm{Up}}) \\
&= \mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \mathbb{F}\mathrm{union}^{\mathrm{Up}}\,, \\
\mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \mathbb{F}\mathrm{pt}^{\mathrm{Up}} &= \mathbb{F}(\mathrm{union}^{\mathrm{Up}} \circ \mathrm{pt}^{\mathrm{Up}}) \\
&= \mathrm{id}\,, \\
\mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \widehat{\mathrm{Up}}(\mathbb{F}\mathrm{pt}^{\mathrm{Up}}) &= \mathbb{F}\mathrm{union}^{\mathrm{Up}} \circ \mathbb{F}\mathrm{Up}(\mathrm{pt}^{\mathrm{Up}}) \\
&= \mathbb{F}(\mathrm{union}^{\mathrm{Up}} \circ \mathrm{Up}(\mathrm{pt}^{\mathrm{Up}})) \\
&= \mathrm{id}\,.
\end{aligned}$$

Having found this extension $\widehat{\mathrm{Up}}$, we can now use Theorem I.6.7 (and its proof) to obtain a distributive law $\delta : \mathrm{Up} \circ \mathrm{Dn} \Rightarrow \mathrm{Dn} \circ \mathrm{Up}$ given by

$$\begin{aligned}
\delta(H) &= \big(\mathbb{V}\widehat{\mathrm{Up}}(\mathrm{union}^{\mathrm{Dn}}) \circ \mathrm{pt}^{\mathrm{Dn}}\big)(H) \\
&= \big\{\, J \in \mathrm{Up}(A) \bigm| J \cap \mathrm{union}^{\mathrm{Dn}}(\mathrm{pt}^{\mathrm{Dn}}(I)) \neq \varnothing \text{ for all } I \in H \,\big\} \\
&= \{\, J \in \mathrm{Up}(A) \mid J \cap I \neq \varnothing \text{ for all } I \in H \,\} \\
&= \mathrm{skolem}^{\mathrm{Up}}(H)\,,
\end{aligned}$$

for $H \in \mathrm{Up}(\mathrm{Dn}(A))$, where $\mathbb{V} : \mathsf{Free}(\mathrm{Dn}) \to \mathcal{D}^{\Xi}$ is the forgetful functor.

$\square$

Remark. A skolem-bialgebra is the same thing as a complete distributive lattice, i.e., a complete lattice were arbitrary meets distributive over arbitrary joins and vice versa. ⌟

Notes

There is an extensive literature on distributive laws for variants of the power-set functor. Our presentation is based on [27, 25, 12, 8]. Further relevant references include [30, 31, 24, 45].

Proposition 3.2 is based on Lemma 6.13 in [10]. Theorem 5.3 is due to [17].

Regular expressions for infinite trees seem to be folklore, but I have not found them anywhere in the literature (except for a few remarks in [40]). Regular except for finite trees can be found, e.g., in Section 2.4 of [29].

VIII. Branch-Continuous Algebras

1. Sublinear Trees

THE AIM OF THIS CHAPTER IS to develop a structure theory for $\mathbb{T}$-algebras that are generated in a certain way by an ω-semigroup. To simplify the technical development below we define a variant of trees where we are allowed to omit variables. The following functor $\mathbb{Y}$ is a variant of the monad $\mathbb{X}$ which we have introduced in Section VII.4 above. The difference between $\mathbb{X}$ and $\mathbb{Y}$ is that $\mathbb{X}$ uses surjective functions $\sigma : \zeta \to \xi$ between sorts, while $\mathbb{Y}$ uses inclusion maps $\zeta \to \xi$.

Definition 1.1. (a) The functor $\mathbb{Y} : \mathcal{D}^{\Xi} \to \mathcal{D}^{\Xi}$ is defined by

$$\mathbb{Y}_\xi A := \sum_{\zeta \subseteq \xi} A_\zeta\,, \quad \text{for } \xi \in \Xi\,.$$

We write the elements of $\mathbb{Y}_\xi A$ as $\langle \xi, a \rangle$ with $a \in A_\zeta$ and $\zeta \subseteq \xi$.

(b) The *support* of $\langle \xi, a \rangle \in \mathbb{Y}_\xi A$ with $a \in A_\zeta$ is the set

$$\operatorname{supp}(\langle \xi, a \rangle) := \zeta\,.$$

(c) The *canonical inclusion* $i : \mathrm{Id} \Rightarrow \mathbb{Y}$ is the natural transformation defined by

$$i(a) := \langle \xi, a \rangle\,, \quad \text{for } a \in A_\xi\,.$$

The *merging function* $m : \mathbb{Y}\mathbb{Y} \Rightarrow \mathbb{Y}$ is the natural transformation defined by

$$m(\langle \xi, \langle \zeta, a \rangle \rangle) := \langle \xi, a \rangle\,. \qquad ⌟$$

Lemma 1.2. *$\langle \mathbb{Y}, m, i\rangle$ is a monad.*

Proof. For $a \in A_\xi$ and $\xi \subseteq \zeta \subseteq \eta \subseteq \nu$, we have

$$\begin{aligned}
m(i(\langle \zeta, a\rangle)) &= m(\langle \zeta, \langle \zeta, a\rangle\rangle) = \langle \zeta, a\rangle\,,\\
m(\mathbb{Y}i(\langle \zeta, a\rangle)) &= m(\langle \zeta, \langle \xi, a\rangle\rangle) = \langle \zeta, a\rangle\,,\\
m(\mathbb{Y}m(\langle \nu, \langle \eta, \langle \zeta, a\rangle\rangle\rangle)) &= m(\langle \nu, \langle \eta, a\rangle\rangle)\\
&= \langle \nu, a\rangle\\
&= m(\langle \nu, \langle \zeta, a\rangle\rangle)\\
&= m(m(\langle \nu, \langle \eta, \langle \zeta, a\rangle\rangle\rangle))\,.
\end{aligned}$$ □

For trees, we can lift the monad structure from $\mathbb{T}$ to $\mathbb{YT}$.

Definition 1.3. (a) For a functor $\mathbb{F} : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$, we define $\mathbb{F}^? : \mathcal{D}^\Xi \to \mathcal{D}^\Xi$ by

$$\mathbb{F}^? := \mathbb{Y} \circ \mathbb{F}\,.$$

(b) The elements of $\mathbb{T}^?$ are called *sublinear trees*.

(c) Let prune : $\mathbb{TY} \Rightarrow \mathbb{YT}$ be the natural transformation mapping each tree $t \in \mathbb{T}_\xi \mathbb{Y}A$ to $\mathrm{prune}(t) = \langle \xi, s\rangle \in \mathbb{T}^? A$, where s is the tree obtained from t as follows. For every vertex $v \in \mathrm{dom}(t)$ with $t(v) = \langle \xi, a\rangle$ where $a \in A_\zeta$, we change the label to a and we remove all subtrees attached at an x-successor with $x \in \xi \setminus \zeta$. Since this operation might remove variables from t, we have $s \in \mathbb{T}_\eta A$, for some $\eta \subseteq \xi$.

(d) We define $\mathrm{sing}^? : \mathrm{Id} \Rightarrow \mathbb{T}^?$ and $\mathrm{flat}^? : \mathbb{T}^?\mathbb{T}^? \Rightarrow \mathbb{T}^?$ by

$$\mathrm{sing}^? := i \circ \mathrm{sing} \quad \text{and} \quad \mathrm{flat}^? := \mathbb{Y}\mathrm{flat} \circ m \circ \mathbb{Y}\mathrm{prune}\,.$$

where $i : \mathrm{Id} \Rightarrow \mathbb{Y}$ and $m : \mathbb{YY} \Rightarrow \mathbb{Y}$ are the functions from above.

(e) For a tree $t \in \mathbb{T}^?_\xi A$ and an injective function $\sigma : \xi \to \zeta$, we denote by ${}^\sigma t \in \mathbb{T}^?_\zeta A$ the tree obtained from t by replacing every variable x by $\sigma(x)$.

If $\mathfrak{A}$ is a $\mathbb{T}^?$-algebra and $a \in A_\xi$, we define

$${}^\sigma a := \pi({}^\sigma \mathrm{sing}^?(a))\,.$$ ⌟

Lemma 1.4. $\mathrm{prune} : \mathbb{T}\mathbb{Y} \Rightarrow \mathbb{Y}\mathbb{T}$ *is a distributive law.*

Proof. We start by checking that prune is a natural transformation. Let $f : A \to B$ and $t \in \mathbb{T}\mathbb{Y}A$. For $v \in \mathrm{dom}(t)$, we have

$$t(v) = \langle \xi, a \rangle \quad \text{and} \quad \mathbb{T}\mathbb{Y}f(t)(v) = \langle \xi, f(a) \rangle, \quad \text{for some } a \text{ and } \xi.$$

Since $\mathrm{supp}(\langle \xi, a \rangle) = \mathrm{supp}(\langle \xi, f(a) \rangle)$ it follows that we delete the same subtrees when constructing the trees $\mathrm{prune}(t)$ and $\mathrm{prune}(\mathbb{T}\mathbb{Y}f(t))$. Consequently,

$$\mathbb{Y}\mathbb{T}f(\mathrm{prune}(t)) = \mathrm{prune}(\mathbb{T}\mathbb{Y}f(t)).$$

It remains to check the axioms of a distributive law.

(i) $\mathrm{prune} \circ \mathrm{flat} = \mathbb{Y}\mathrm{flat} \circ \mathrm{prune} \circ \mathbb{T}\mathrm{prune}$

(ii) $\mathrm{prune} \circ \mathbb{T}m = m \circ \mathbb{Y}\mathrm{prune} \circ \mathrm{prune}$

(iii) $\mathrm{prune} \circ \mathrm{sing} = \mathbb{Y}\mathrm{sing}$

(iv) $\mathrm{prune} \circ \mathbb{T}i = i$

(iii) Let $a \in A_\zeta$ and $\xi \supseteq \zeta$. Then

$$\begin{aligned}(\mathrm{prune} \circ \mathrm{sing})(\langle \xi, a \rangle) &= \mathrm{prune}(\mathrm{sing}(\langle \xi, a \rangle)) \\ &= \langle \xi, \mathrm{sing}(a) \rangle \\ &= \mathbb{Y}\mathrm{sing}(\langle \xi, a \rangle).\end{aligned}$$

(iv) Let $t \in \mathbb{T}_\xi A$. Then

$$(\mathrm{prune} \circ \mathbb{T}i)(t) = \langle \xi, t \rangle = i(t).$$

(i) Let $t \in \mathbb{T}\mathbb{T}\mathbb{Y}A$. Let $\mu : \mathrm{dom}(\mathrm{flat}(t)) \to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v))$ be the canonical bijection. A vertex $v \in \mathrm{dom}(\mathrm{flat}(t))$ remains in $\mathrm{prune}(\mathrm{flat}(t))$ if, and only if,

$$x \in \mathrm{supp}(\mathrm{flat}(t)(u)), \quad \text{for all } u <_{\mathrm{pf}} v \text{ with } \mathrm{suc}_x(u) \leq_{\mathrm{pf}} v.$$

Similarly, $v \in \mathrm{dom}(t)$ belongs to the tree $(\mathrm{prune} \circ \mathbb{T}\mathrm{prune})(t)$ if

$$x \in \mathrm{supp}(\mathrm{prune}(t(u)))\,, \quad \text{for all } u <_{\mathrm{pf}} v \text{ with } \mathrm{suc}_x(u) \leq_{\mathrm{pf}} v\,.$$

Since $x \in \mathrm{supp}(\mathrm{prune}(t(u)))$ if the leaf w_x with label x is not removed by $\mathrm{prune}(t(u))$, the latter is equivalent to

$$y \in \mathrm{supp}(t(u)(u'))\,, \quad \text{for all } u <_{\mathrm{pf}} v \text{ with } \mathrm{suc}_x(u) \leq_{\mathrm{pf}} v \text{ and all } u' <_{\mathrm{pf}} w_x \text{ with}$$

It follows that v belongs to the tree $(\mathbb{Y}\mathrm{flat} \circ \mathrm{prune} \circ \mathbb{T}\mathrm{prune})(t)$ if

$$x \in \mathrm{supp}(t(u)(u'))\,, \quad \text{for all } u, u' \text{ with } \mu^{-1}(u, u') <_{\mathrm{pf}} v \text{ and } \mathrm{suc}_x(\mu^{-1}(u, u'))$$

Consequently, we have

$$\mathrm{prune}(\mathrm{flat}(t)) = (\mathbb{Y}\mathrm{flat} \circ \mathrm{prune} \circ \mathbb{T}\mathrm{prune})(t)\,.$$

(11) Let $t \in \mathbb{T}\mathbb{Y}\mathbb{Y}A$. For $v \in \mathrm{dom}(t)$, we have

$$t(v) = \langle \xi, \langle \zeta, a \rangle \rangle \quad \text{and} \quad \mathbb{Y}m(t)(v) = \langle \xi, a \rangle\,, \quad \text{for some } a \text{ and } \xi, \zeta\,.$$

Hence,

- $\mathrm{prune}(\mathbb{Y}m(t))$ removes all x-successors of v with $x \in \xi \setminus \mathrm{supp}(a)$,
- $\mathrm{prune}(t)$ removes all x-successors of v with $x \in \xi \setminus \zeta$, and
- $\mathbb{Y}\mathrm{prune}(\mathrm{prune}(t))$ removes all x-successors of v with $x \in \zeta \setminus \mathrm{supp}(a)$.

It follows that we remove the same successors in both trees. Consequently,

$$\mathrm{prune}(\mathbb{T}m(t)) = (m \circ \mathbb{Y}\mathrm{prune} \circ \mathrm{prune})(t)\,. \qquad \square$$

Corollary 1.5. *Let $\rho : \mathbb{T} \Rightarrow \mathbb{T}^{?}$ be the morphism of monads induced by the canonical inclusion $i : \mathrm{Id} \Rightarrow \mathbb{Y}$.*

(a) *$\langle \mathbb{T}^{?}, \mathrm{flat}^{?}, \mathrm{sing}^{?} \rangle$ is a monad on $\mathcal{D}^{\Xi}$.*

(b) *The reduct functor $-|_\rho : \mathsf{Alg}(\mathbb{T}^{?}) \to \mathsf{Alg}(\mathbb{T})$ has a left-adjoint $-^{?} : \mathsf{Alg}(\mathbb{T}) \to \mathsf{Alg}(\mathbb{T}^{?})$ mapping a $\mathbb{T}$-algebra $\mathfrak{A} = \langle A, \pi \rangle$ to the $\mathbb{T}^{?}$-algebra*

$$\mathfrak{A}^{?} := \langle \mathbb{Y}A,\ \mathbb{Y}\pi \circ m \circ \mathbb{Y}\mathrm{prune} \rangle\,.$$

Proof. (a) By Corollary I.6.8 (a), the functor $\mathbb{YT} = \mathbb{T}^?$ forms a monad with multiplication

$$\begin{aligned} m \circ \mathbb{Y}(\mathbb{Y}\mathrm{flat} \circ \mathrm{prune}) &= m \circ \mathbb{YY}\mathrm{flat} \circ \mathbb{Y}\mathrm{prune} \\ &= \mathbb{Y}\mathrm{flat} \circ m \circ \mathbb{Y}\mathrm{prune} \\ &= \mathrm{flat}^? \end{aligned}$$

and unit map

$$i \circ \mathrm{sing} = \mathrm{sing}^? .$$

(b) follows by Corollary I.6.8 (c) since

$$m \circ \mathbb{YY}\pi \circ \mathbb{Y}\mathrm{prune} = \mathbb{Y}\pi \circ m \circ \mathbb{Y}\mathrm{prune} . \qquad \square$$

Remark. It follows from the above results that a $\mathbb{T}^?$-algebra $\mathfrak{A}$ is the same as a prune-bialgebra, i.e., a $\mathbb{T}$-algebra that is equipped with an additional $\mathbb{Y}$-product which commutes with the $\mathbb{T}$-product. Note that a $\mathbb{Y}$-product is nothing but a family of functions $A_\xi \to A_\zeta$, for for all pairs of sorts $\xi \subseteq \zeta$. Thus, a $\mathbb{T}^?$-algebra is a $\mathbb{T}$-algebra equipped with additional functions $A_\xi \to A_\zeta$ that are compatible with the $\mathbb{T}$-product. ⌟

2. Semigroup-Like Algebras

Below we will be interested in ways an ω-semigroup can sit inside a $\mathbb{T}^?$-algebra and in $\mathbb{T}^?$-algebras generated by some ω-semigroup they contain. We start with the simplest case, $\mathbb{T}^?$-algebras that are generated by elements of arity at most 1. These provide the basic building blocks we will use in the following sections to build more complicated algebras.

Definition 2.1. A $\mathbb{T}^?$-algebra $\mathfrak{A}$ is *semigroup-like* if it is generated by the subset $A_\varnothing \cup A_{\{z\}}$. ⌟

The reason why we are working with the monad $\mathbb{T}^?$ is that the above definition becomes trivial for $\mathbb{T}$-algebras.

Exercise 2.1. Prove that a $\mathbb{T}^\times$-algebra $\mathfrak{A}$ is generated by the subset $A_\varnothing \cup A_{\{z\}}$ if, and only if, $A_\xi = \varnothing$ for all sorts ξ of size $|\xi| > 1$. ⌟

The aim of the remainder of this section is to give a more concrete characterisation of when an algebra is semigroup-like.

Definition 2.2. (a) The *word functor* $\mathbb{W} : \mathsf{Pos}^2 \to \mathsf{Pos}^2$ is defined by

$$\mathbb{W}\langle A_1, A_\omega\rangle := \langle A_1^+,\ A_1^* A_\omega + A_1^\omega\rangle\,.$$

(b) The natural transformation $\mathrm{cat} : \mathbb{W}\mathbb{W} \Rightarrow \mathbb{W}$ maps a sequence $(w_i)_i$ to their concatenation $w_0 w_1 \dots$. The natural transformation $\langle - \rangle : \mathrm{Id} \Rightarrow \mathbb{W}$ maps an element a to the sequence $\langle a\rangle$.

(c) $\mathbb{W}$-algebras are called *(ordered) ω-semigroups.* We use the usual notation for products in ω-semigroups. That is, we denote the product of two elements $a \in S_1$ and $b \in S_1 + S_\omega$ by $a \cdot b$ or just ab. Similarly, we write $\prod_{i<\alpha} a_i$ for an infinite product. ⌟

Lemma 2.3. $\langle \mathbb{W}, \mathrm{cat}, \langle - \rangle\rangle$ *forms a monad on* Pos^2.

Remark. Note that there exists a fully faithful embedding $\mathsf{Pos}^2 \to \mathsf{Pos}^\Xi$ mapping each set $\langle A_1, A_\omega\rangle$ to

$$(A'_\xi)_{\xi\in\Xi} \quad \text{with} \quad A'_\xi := \begin{cases} A_1 & \text{if } \xi = \{z\}\,, \\ A_\omega & \text{if } \xi = \varnothing\,, \\ \varnothing & \text{otherwise}\,, \end{cases}$$

where z is some arbitrary, but fixed variable.

This embedding can be used to lift every monad on Pos^2 to one on Pos^Ξ. In the following we will tacitly identify every set, function, and monad of Pos^2 with its image in Pos^Ξ.

Using this identification, we obtain an isomorphism of monads

$$\mathbb{W} \cong \mathbb{T}|_{\{\varnothing,\{z\}\}}\,.$$

In particular, using this isomorphism we can consider $\mathbb{W}$ a submonad of $\mathbb{T}$. ⌟

When we restrict a $\mathbb{T}$-algebra to the elements of arity 0 and 1 we obtain an ω-semigroup.

Definition 2.4. The *ω-semigroup associated* with a $\mathbb{T}^{?}$-algebra $\mathfrak{A}$ is the $\mathbb{W}$-algebra

$$\mathrm{SG}(\mathfrak{A}) = \mathfrak{A}|_{\{\varnothing,\{z\}\}}\,. \qquad \lrcorner$$

Conversely, we can associate with every ω-semigroup $\mathfrak{S}$ a $\mathbb{T}^{?}$-algebra which consists of elements of the form a or $a(x)$, for $a \in S$ and an optional variable x.

Definition 2.5. (a) We define a functor $\mathrm{TA} : \mathsf{Pos}^2 \to \mathsf{Pos}^{\Xi}$ as follows. For sets $S \in \mathsf{Pos}^2$, we set

$$\mathrm{TA}_\xi(S) := S_\omega + (S_1 \times \xi)\,, \quad \text{for } \xi \in \Xi\,.$$

The ordering on $\mathrm{TA}_\xi(S)$ is the one induced by the orderings of S_ω, S_1, and ξ, where we consider the last one to be equipped with the trivial order. We will use the more suggestive notation $a(x)$ for the elements of the form $\langle a, x\rangle \in S_1 \times \xi$.

For functions $f : S \to T$, we set

$$\mathrm{TA}(f)(a) := \begin{cases} f(a) & \text{if } a \in S_\omega\,,\\ \langle f(b), x\rangle & \text{if } a = \langle b, x\rangle \in S_1 \times \xi\,.\end{cases}$$

(b) We define a natural transformation $\mathrm{prune} : \mathbb{T}^{?} \circ \mathrm{TA} \Rightarrow \mathrm{TA} \circ \mathbb{W}$ as follows. Given $t \in \mathbb{T}^{?}_\xi \mathrm{TA}(S)$, let $\beta = (v_i)_i$ be the branch of t defined inductively as follows. We start with the root $v_0 := \langle\rangle$. For the inductive step, suppose that we have already defined v_i. If $t(v_i) \in S_\omega$ or $t(v_i) \in \xi$, we stop. Otherwise, $t(v_i) = \langle a_i, x_i\rangle \in S_1 \times \zeta$ and we choose for v_{i+1} the x_i-successor of v_i. Let $w := (a_i)_i$ be the sequence of labels along β. This sequence is of one of the following forms:

- an infinite sequence $a_0, a_1, \dots \in S_1^\omega$,
- a finite sequence $a_0, \dots, a_n \in S_1^* S_\omega$,

- a finite sequence $a_0, \dots, a_n \in S_1^+ \xi \cong S_1^+ \times \xi$.

In particular, $w \in \mathrm{TA}(\mathbb{W}S) = S_1^\omega + S_1^* S_\omega + S_1^+ \times \xi$. We set $\mathrm{prune}(t) := w$.

(c) The $\mathbb{T}^?$*-algebra* $\mathrm{TA}(\mathfrak{S})$ *associated with* an ω-semigroup $\mathfrak{S}$ is

$$\mathrm{TA}(\mathfrak{S}) := \langle \mathrm{TA}(\mathfrak{S}), \mathrm{TA}(\pi) \circ \mathrm{prune} \rangle ,$$

where $\pi : \mathbb{W}S \to S$ is the product of $\mathfrak{S}$. ⌟

Remark. Note that $\mathrm{TA}(S) \cong \mathbb{Y}\tilde{S}$, where $\tilde{S}$ is the set with components

$$\tilde{S}_\xi := \begin{cases} S_\omega & \text{if } \xi = \varnothing , \\ S_1 & \text{if } |\xi| = 1 , \\ \varnothing & \text{if } |\xi| > 1 . \end{cases}$$

Under this isomorphism, the prune function for TA turns into the one for $\mathbb{Y}$. That is the reason we have chosen the same name for both. ⌟

Exercise 2.2. Prove that $\mathrm{prune} : \mathbb{T}^? \circ \mathrm{TA} \Rightarrow \mathrm{TA} \circ \mathbb{W}$ is a natural transformation. ⌟

Proposition 2.6. $\mathrm{TA} \dashv \mathrm{SG}$ *is an adjunction between* $\mathsf{Alg}(\mathbb{T}^?)$ *and* $\mathsf{Alg}(\mathbb{W})$.

Proof. Since $\mathrm{SG} \cong |_{\{\varnothing,\{z\}\}}$ and $\mathbb{W} \cong \mathbb{T}|_{\{\varnothing,\{z\}\}}$, it follows directly by Lemma II.3.5 (c) that $\mathrm{SG} : \mathsf{Alg}(\mathbb{T}^?) \to \mathsf{Alg}(\mathbb{W})$ is a functor.

Before continuing let us establish the following two equations, which are similar to those for a distributive law.

$$\mathrm{prune} \circ \mathrm{sing} = \mathrm{TA}(\langle - \rangle)$$
$$\mathrm{prune} \circ \mathrm{flat} = \mathrm{TA}(\mathrm{cat}) \circ \mathrm{prune} \circ \mathbb{T}^? \mathrm{prune} .$$

The proof is very similar to that of Lemma 1.4. For the first equation, let $a \in \mathrm{TA}(S)$. Then

$$\mathrm{prune}(\mathrm{sing}(a)) = \begin{cases} \langle a \rangle & \text{if } a \in S_\omega , \\ \langle b \rangle (x) & \text{if } a = b(x) \in S_1 \times \xi . \end{cases}$$

For the second one, consider a tree $t \in \mathbb{T}^? \mathbb{T}^? \mathrm{TA}(S)$ and let β be the path of $\mathrm{flat}(t)$ used to define $\mathrm{prune}(\mathrm{flat}(t))$. The image of β under the canonical map

$$\mu : \mathrm{dom}(\mathrm{flat}(t)) \to \sum_{v \in \mathrm{dom}(t)} \mathrm{dom}(t(v))$$

induces a path $\gamma = (v_i)_i$ in t and, for each i, a path δ_i in $t(v_i)$. Furthermore, δ_i is the path defining $\mathrm{prune}(t(v_i))$ and γ is the one defining $\mathrm{prune}(\mathbb{T}^? \mathrm{prune}(t))$. Since the sequence of semigroup elements labelling β is the concatenation of the sequences labelling δ_i, the claim follows.

It is straightforward to check that TA preserves identity morphisms and composition of morphisms. To show that it is a functor $\mathsf{Alg}(\mathbb{W}) \to \mathsf{Alg}(\mathbb{T}^?)$ it is therefore sufficient to prove that TA maps $\mathbb{W}$-algebras to $\mathbb{T}^?$-algebras and $\mathbb{W}$-morphisms to $\mathbb{T}^?$-morphisms. Hence, let $\mathfrak{S} = \langle S, \pi \rangle$ be an ω-semigroup. To see that $\mathrm{TA}(\mathfrak{S})$ is a $\mathbb{T}^?$-algebra, note that

$$\begin{aligned} &\mathrm{TA}(\pi) \circ \mathrm{prune} \circ \mathrm{sing} \\ &= \mathrm{TA}(\pi) \circ \mathrm{TA}(\langle - \rangle) \\ &= \mathrm{id}\,, \end{aligned}$$

$$\begin{aligned} &(\mathrm{TA}(\pi) \circ \mathrm{prune}) \circ \mathbb{T}^?(\mathrm{TA}(\pi) \circ \mathrm{prune}) \\ &= \mathrm{TA}(\pi) \circ \mathrm{TA}(\mathbb{W}\pi) \circ \mathrm{prune} \circ \mathbb{T}^? \mathrm{prune} \\ &= \mathrm{TA}(\pi) \circ \mathrm{TA}(\mathrm{cat}) \circ \mathrm{prune} \circ \mathbb{T}^? \mathrm{prune} \\ &= (\mathrm{TA}(\pi) \circ \mathrm{prune}) \circ \mathrm{flat}\,. \end{aligned}$$

Let $\varphi : \mathfrak{S} \to \mathfrak{T}$ be a morphism of ω-semigroups. To see that $\mathrm{TA}(\varphi)$ is a $\mathbb{T}^?$-morphism, note that

$$\begin{aligned} \mathrm{TA}(\varphi) \circ (\mathrm{TA}(\pi) \circ \mathrm{prune}) &= \mathrm{TA}(\pi \circ \mathbb{W}\varphi) \circ \mathrm{prune} \\ &= (\mathrm{TA}(\pi) \circ \mathrm{prune}) \circ \mathbb{T}^? \mathrm{TA}(\varphi)\,. \end{aligned}$$

It remains to prove that TA is the left adjoint of SG. Note that every $\mathbb{T}^?$-morphism $\varphi : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{A}$ induces a $\mathbb{W}$-morphism $\hat{\varphi} : \mathfrak{S} \to \mathrm{SG}(\mathfrak{A})$

defined by

$$\hat{\varphi}(a) := \begin{cases} \varphi_{\varnothing}(a) & \text{if } a \in S_\omega\,, \\ \varphi_{\{z\}}(a(z)) & \text{if } a \in S_1\,. \end{cases}$$

Conversely, every $\mathbb{W}$-morphism $\psi : \mathfrak{S} \to \mathrm{SG}(\mathfrak{A})$ induces a $\mathbb{T}^{?}$-morphism $\tilde{\psi} : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{A}$ by

$$\tilde{\psi}(a) := \begin{cases} \psi(a) & \text{if } a \in S_\omega\,, \\ \psi(b)(z) & \text{if } a = b(z) \in S_1 \times \xi\,. \end{cases}$$

Since the two maps $\varphi \mapsto \hat{\varphi}$ and $\psi \mapsto \tilde{\psi}$ are inverse to each other, it follows that

$$\mathsf{Pos}^{\Xi}(\mathrm{TA}(\mathfrak{S}), \mathfrak{A}) \cong \mathsf{Pos}^{2}(\mathfrak{S}, \mathrm{SG}(\mathfrak{A}))\,.$$

To see that this bijection is natural in $\mathfrak{S}$ and $\mathfrak{A}$, consider morphisms $f : \mathfrak{S} \to \mathfrak{T}$, $g : \mathfrak{A} \to \mathfrak{B}$, $\psi : \mathfrak{T} \to \mathrm{SG}(\mathfrak{A})$, and $\psi' : \mathfrak{S} \to \mathrm{SG}(\mathfrak{A})$. We have to show that the morphism $\psi \circ f$ is mapped to $\tilde{\psi} \circ \mathrm{TA}(f)$, and that $\mathrm{SG}(g) \circ \psi'$ is mapped to $g \circ \tilde{\psi}'$. Both proofs are straithforward. □

After these preparations, we can characterise semigroup-like algebras as follows.

Proposition 2.7. *A $\mathbb{T}^{?}$-algebra $\mathfrak{A}$ is semigroup-like if, and only if, there exists a surjective $\mathbb{T}^{?}$-morphism $\varphi : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{A}$, for some ω-semigroup $\mathfrak{S}$.*

Proof. ($\Rightarrow$) The counit $\eta : \mathrm{TA}(\mathrm{SG}(\mathfrak{A})) \to \mathfrak{A}$ of the adjunction from Proposition 2.6 is surjective since $\operatorname{rng} \eta = \langle\!\langle A_{\varnothing} \cup A_{\{z\}} \rangle\!\rangle_{\mathfrak{A}} = A$.

($\Leftarrow$) Suppose that $\varphi : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{A}$ is surjective and set $\mathfrak{C} := \mathrm{TA}(\mathfrak{S})$. Then

$$\begin{aligned} \langle\!\langle A_{\varnothing} \cup A_{\{z\}} \rangle\!\rangle_{\mathfrak{A}} &= \big\langle\!\big\langle \varphi[C_{\varnothing} \cup C_{\{z\}}] \big\rangle\!\big\rangle_{\mathfrak{A}} \\ &= \varphi[\langle\!\langle C_{\varnothing} \cup C_{\{z\}} \rangle\!\rangle_{\mathfrak{C}}] = \varphi[C] = A\,. \end{aligned}$$

Hence, $\mathfrak{A}$ is semigroup-like. □

3. Meet-Distributive Algebras

Having a distributive law available for the functors Up and Dn, we can now study their bialgebras. Let us give them a special name.

Definition 3.1. Let $\mathbb{M}$ be a linear monad on Pos^{Ξ}.

(a) An $\mathbb{M}$,Up-bialgebra $\mathfrak{A}$ is *meet-distributive* if it is a dist-bialgebra. An Up-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ from an $\mathbb{M}$,Up-bialgebra $\mathfrak{A}$ to an Up-algebra $\mathfrak{B}$ is *meet-distributive* if it is dist-distributive.

(b) An $\mathbb{M}$,Dn-bialgebra $\mathfrak{A}$ is *join-distributive* if it is a dist-bialgebra. A Dn-morphism $\varphi : \mathfrak{A} \to \mathfrak{B}$ from an $\mathbb{M}$,Dn-bialgebra $\mathfrak{A}$ to a Dn-algebra $\mathfrak{B}$ is *join-distributive* if it is dist-distributive.

(c) A Dn,Up-bialgebra $\mathfrak{A}$ is *lattice-distributive* if it is skolem-distributive.

(d) Let $A \in \mathsf{Pos}^{\Xi}$ be a set. A subset $C \subseteq A$ is a set of *join-generators* if every element of A is a join of elements of C. Similarly, C is a set of *meet-generators* if every $a \in A$ is a meet of such elements. ⌟

Remark. Let $\mathbb{M}$ be a linear monad. Unpacking the definition, we see that an $\mathbb{M}$-algebra $\mathfrak{A}$ is meet-distributive if, and only if, A is completely ordered and the product commutes with meets in the sense that, for a term $T \in \mathbb{M}\mathrm{Up}(A)$, we have

$$\pi(t) = \inf\{\, \pi(s) \mid s \in^{\mathbb{M}} T \,\}\,, \quad \text{where } t := \mathbb{M}\mathrm{inf}(T)\,, \qquad ⌟$$

Exercise 3.1. Prove that a function $j : \mathfrak{A} \to \mathfrak{B}$ is meet-distributive if, and only if, for all terms $t, t' \in \mathbb{M}\mathrm{Up}(A)$,

$$\mathbb{M}\mathrm{inf}(\mathbb{M}\mathrm{Up}(j)(t)) \leq \mathbb{M}\mathrm{inf}(\mathbb{M}\mathrm{Up}(j)(t'))$$

implies

$$\inf\{\, j(\pi(s)) \mid s \in^{\mathbb{M}} t \,\} \leq \inf\{\, j(\pi(s')) \mid s' \in^{\mathbb{M}} t' \,\}\,. \qquad ⌟$$

A special property of the monads Up and Dn is that every Up-algebra has a canonical Dn-algebra product and vice versa. This is because all Up-algebras are of the form $\langle A, \inf \rangle$, for some completely ordered set A. This implies that A also has arbitrary joins and $\langle A, \sup \rangle$ forms a Dn-algebra.

Lemma 3.2. *Let $\mathbb{M}$ be a linear monad and $\mathfrak{A}$ an $\mathbb{M}$-algebra. The canonical embedding* $\mathrm{pt} : \mathfrak{A} \to \mathrm{Up}(\mathfrak{A})$ *preserves joins.*

Proof. For $S \subseteq A$, we have

$$\begin{aligned} \mathrm{pt}(\sup S) &= \{ c \in A \mid c \geq \sup S \} \\ &= \bigcap \{ \Uparrow a \mid a \in S \} \\ &= \bigcap \{ \mathrm{pt}(a) \mid a \in S \} = \sup \mathrm{Up}(\mathrm{pt})(S) . \end{aligned}$$ □

Finally, let us characterise when an algebra is of the form $\mathrm{Up}(\mathfrak{C})$.

Definition 3.3. Let A be an ordered set. An element $a \in A_\xi$ is *completely prime* if, for every set $F \subseteq A_\xi$,

$$\inf F \leq a \quad \text{implies} \quad b \leq a \text{ for some } b \in F .$$ ⌟

Theorem 3.4. *Let $\mathbb{M}$ be a linear monad. An $\mathbb{M}$-algebra $\mathfrak{A}$ is of the form*

$$\mathfrak{A} \cong \mathrm{Up}(\mathfrak{B}) , \quad \text{for some } \mathbb{M}\text{-algebra } \mathfrak{B} ,$$

if, and only if, $\mathfrak{A}$ is meet-distributive and there exists a subalgebra $\mathfrak{C} \subseteq \mathfrak{A}$ such that C is a set of meet-generators of A and every element $c \in C$ is completely prime in A.

Proof. ($\Rightarrow$) Suppose that $\mathfrak{A} = \mathrm{Up}(\mathfrak{B})$. Then $\mathfrak{A}$ is meet-distributive, the function $\mathrm{pt} : \mathfrak{B} \to \mathfrak{A}$ is an embedding of $\mathbb{M}$-algebras, and its range $\mathfrak{C} := \mathrm{pt}(\mathfrak{B})$ is a subalgebra of $\mathfrak{A}$. Furthermore, every element of A can be written as a infimum of elements of C. Finally, consider an element $c \in C_\xi$ and a set $F \subseteq A_\xi$ with $\inf F \leq c$. Then $c = \Uparrow b$, for some $b \in B$, and we have

$$\begin{aligned} \inf F \leq c \quad &\Rightarrow \quad \bigcap F \supseteq \Uparrow b \\ &\Rightarrow \quad b \in I , \quad \text{for some } I \in F , \\ &\Rightarrow \quad I \supseteq \Uparrow b \\ &\Rightarrow \quad I \leq c . \end{aligned}$$

($\Leftarrow$) Let $\mathfrak{C} \subseteq \mathfrak{A}$ be a subalgebra as above. Since $\mathfrak{A}$ is meet-distributive, the inclusion $\mathfrak{C} \to \mathfrak{A}$ induces a morphism $\varphi : \mathrm{Up}(\mathfrak{C}) \to \mathfrak{A}$ of $\mathbb{M}$,Up-bialgebras via Corollary II.4.4. We claim that φ is an isomorphism. As φ is a Up-morphism and every element of $I \in \mathrm{Up}(C)$ can be written as the infimum of all singletons $\Uparrow c \supseteq I$, we have

$$\varphi(I) = \inf I \,, \quad \text{for } I \in \mathrm{Up}(C)\,,$$

where the infimum is taken in A. We have to prove the following properties.

(I) $I \supseteq J \quad \text{iff} \quad \varphi(I) \leq \varphi(J)$

(II) φ is injective.

(III) φ is surjective.

(I) Since φ is a morphism of Pos^{Ξ}, it is monotone. For the converse, let $I, J \in \mathrm{Up}(C)$. Then

$$\begin{aligned}
\varphi(I) \leq \varphi(J) \quad &\Rightarrow \quad \inf I \leq \inf J \\
&\Rightarrow \quad \inf I \leq d\,, \quad \text{for all } d \in J\,, \\
&\Rightarrow \quad \text{for every } d \in J \text{ there is some } c \in I \text{ with } c \leq d \\
&\Rightarrow \quad I \supseteq J\,.
\end{aligned}$$

where the third step follows since the elements of C are completely prime.

(II) By (I),

$$\varphi(I) = \varphi(J) \quad \text{implies} \quad I \subseteq J \text{ and } J \subseteq I\,.$$

(III) holds as C is a set of meet-generators of A. □

4. Trace Algebras

The definition of an MSO-definable algebra is quite abstract and not very helpful when we want to study MSO-definability. Our aim here is to give a more explicit, purely algebraic characterisation. The class of algebras we are considering in this section is the following one.

Definition 4.1. A $\mathbb{T}^?$-algebra $\mathfrak{A}$ is a *trace algebra* if it is meet-distributive and it has a semigroup-like subalgebra $\mathfrak{U}$ such that U forms a set of meet-generators of $\mathfrak{A}$. ⌟

If follows from the definition that every element of a trace algebra can be written in the form

$$a_0(x_0) \sqcap \cdots \sqcap a_{m-1}(x_{m-1}) \sqcap b_0 \sqcap \cdots \sqcap b_{n-1}\,,$$

where $a_0, \dots, a_{m-1} \in A_{\{z\}}$, $b_0, \dots, b_{n-1} \in A_\varnothing$, and the $x_0, \dots, x_{m-1}$ are variables.

We start with a more concrete description of such algebras.

Proposition 4.2. *A $\mathbb{T}^{?}$-algebra $\mathfrak{A}$ is a trace algebra if, and only if, there exists a surjective $\mathbb{T}^{?}$,*Up*-morphism $\varphi : \mathrm{Up}(\mathrm{TA}(\mathfrak{S})) \to \mathfrak{A}$, for some ω-semigroup $\mathfrak{S}$.*

Proof. ($\Rightarrow$) Let $\mathfrak{A}$ be a trace algebra and let $\mathfrak{U}$ be the corresponding subalgebra. We can use Proposition 2.7, to find a surjective $\mathbb{T}^{?}$-morphism $\varphi_0 : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{U}$, for some ω-semigroup $\mathfrak{S}$. Using the adjunction from Theorem II.4.3, we can extend this morphism to a $\mathbb{T}^{?}$,Up-morphism $\varphi : \mathrm{Up}(\mathrm{TA}(\mathfrak{S})) \to \mathfrak{A}$. Since U is a set of meet-generators of A and $\operatorname{rng} \varphi$ includes U, it follows that φ is surjective.

($\Leftarrow$) Given a $\mathbb{T}^{?}$,Up-morphism $\varphi : \mathrm{Up}(\mathrm{TA}(\mathfrak{S})) \to \mathfrak{A}$, let $\mathfrak{U}$ be the subalgebra of $\mathfrak{A}$ induced by $\varphi[\mathrm{TA}(S)]$. By Proposition 2.7, $\mathfrak{U}$ is semigroup-like.

To see that U is a set of meet-generators of A, fix an element $a \in A$. Since φ is surjective, there is some $b \in \mathrm{Up}(\mathrm{TA}(S))$ with $\varphi(b) = a$. Furthermore, since $\mathrm{TA}(S)$ is a set of meet-generators of $\mathrm{Up}(\mathrm{TA}(S))$, we can find a set $I \subseteq \mathrm{TA}(S)$ with $\inf I = b$. Consequently,

$$a = \varphi(b) = \varphi(\inf I) = \inf \varphi[I]$$

is a meet of elements of U.

It remains to prove that $\mathfrak{A}$ is meet-distributive. By Theorem II.4.3, the algebra $\mathrm{Up}(\mathrm{TA}(\mathfrak{S}))$ is meet-distributive. Hence,

$$\begin{aligned}
\pi \circ \mathbb{T}^{?}\mathrm{inf} \circ \mathbb{T}^{?}\mathrm{Up}(\varphi) &= \pi \circ \mathbb{T}^{?}(\varphi \circ \mathrm{inf}) \\
&= \varphi \circ \pi \circ \mathbb{T}^{?}\mathrm{inf} \\
&= \varphi \circ \mathrm{inf} \circ \mathrm{Up}(\pi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(\varphi \circ \pi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(\pi \circ \mathbb{T}^{?}\varphi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(\pi \circ \mathbb{T}^{?}\varphi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(\pi) \circ \mathrm{dist} \circ \mathbb{T}^{?}\mathrm{Up}(\varphi)\,.
\end{aligned}$$

As φ is surjective and $\mathbb{T}^{?}$ and Up preserve surjectivity, it follows that

$$\pi \circ \mathbb{T}^{?}\mathrm{inf} = \mathrm{inf} \circ \mathrm{Up}(\pi) \circ \mathrm{dist}\,,$$

as desired. □

The product in a trace algebra takes a particularly simple form.

Definition 4.3. Let $\mathfrak{A}$ be a $\mathbb{T}^{?}$-algebra, $t \in \mathbb{T}^{?}A$ a tree, and let $\beta = (v_i)_{i<\alpha}$ be a path (finite or infinite) in t starting at the root.

A *trace* of t along β is a tree $s \in \mathbb{T}^{?}A$ that forms a single path of the same length as β and whose labelling satisfies

$$\begin{aligned}
{}^{\sigma_i}(s(u_i)) &\geq t(v_i)\,, &&\text{if } t(v_i) \text{ is not a variable,} \\
s(u_i) &= t(v_i)\,, &&\text{if } t(v_i) \text{ is a variable.}
\end{aligned}$$

where $(u_i)_{i<\alpha}$ is an increasing enumeration of $\mathrm{dom}(s)$ and the functions $\sigma_i : \zeta_i \to \xi_i$ are defined as follows. ζ_i is the sort of u_i and ξ_i the sort of v_i.

If $\zeta_i = \{z_i\}$, σ_i is the function mapping z_i to the variable x_i such that v_{i+1} is the x_i-successor of v_i. If $\zeta_i = \varnothing$, σ_i is the unique function $\varnothing \to \xi_i$. ⌟

Example. Let a, b be elements of sort $\{z\}$ and c, d ones of sort $\varnothing$. The tree

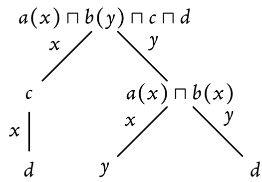

has the following traces.

$$c \sqcap d\,, \qquad a(c)\,, \qquad b\big(a(y) \sqcap b(y)\big)\,.$$

Furthermore, every sequence that is point-wise greater than one of these is also a trace. ⌟

Lemma 4.4. *Let $\mathfrak{A}$ be a semigroup-like $\mathbb{T}^{?}$-algebra. Every tree $t \in \mathbb{T}^{?}A$ has a trace $s \in \mathbb{T}^{?}A$ with $\pi(s) = \pi(t)$.*

Proof. Fix $t \in \mathbb{T}^{?}A$. By Proposition 2.7, there exists a surjective morphism $\varphi : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{A}$, for some ω-semigroup $\mathfrak{S}$. We fix a tree $t' \in \mathbb{T}^{?}\mathrm{TA}(S)$ with $\mathbb{T}^{?}\varphi(t') = t$. Let β be the branch of t' used to define $\mathrm{prune}(t')$ and let s' be the trace of t' along β where each vertex of s' has the same label as the corresponding vertex of t'. Then

$$\pi(s') = \mathrm{TA}(\pi)(\mathrm{prune}(t')) = \pi(t')\,.$$

Consequently, $s := \mathbb{T}^{?}\varphi(s')$ is a trace of t' with

$$\pi(s) = \varphi(\pi(s')) = \varphi(\pi(t')) = \pi(\mathbb{T}^{?}\varphi(t')) = \pi(t)\,. \qquad \square$$

Lemma 4.5. *Let $\mathfrak{A}$ be a trace algebra and let $\mathfrak{U} \subseteq \mathfrak{A}$ be the corresponding semigroup-like subalgebra.*

$$\pi(t) = \inf\{\, \pi(s) \mid s \in \mathbb{T}^{?}U \text{ a trace of } t \,\}\,, \quad \text{for } t \in \mathbb{T}^{?}A\,.$$

Proof. Given a tree $t \in \mathbb{T}^{?}A$, there exists some tree $T \in \mathbb{T}^{?}\mathrm{Up}(U)$ with $t = \mathbb{T}^{?}\inf(T)$. By meet-distributivity, it follows that

$$\pi(t) = \inf\{\, \pi(r) \mid r \in^{\mathbb{T}^{?}} T \,\}\,.$$

To conclude the proof, it is therefore sufficient to show that, for every $r \in^{\mathbb{T}^?} T$, there is some trace s of t with

$$s \in \mathbb{T}^? U \quad \text{and} \quad \pi(s) = \pi(r) .$$

Hence, fix $r \in^{\mathbb{T}^?} T$. By definition of T, we have $r \in \mathbb{T}^? U$ and $t \leq r$. Consequently, we can use Lemma 4.4 to find a trace s of r with $\pi(s) = \pi(r)$. The claim follows since $t \leq r$ implies that s is also a trace of t. □

As an immediate corollary, we obtain the following observation.

Proposition 4.6. *Every finitary trace algebra is* MSO*-definable.*

Proof. Let $\mathfrak{A}$ be a trace algebra, let $\mathfrak{U} \subseteq \mathfrak{A}$ be the corresponding semigroup-like subalgebra, and let $D \subseteq A$ be a finite set of generators of $\mathfrak{A}$. There exists an MSO-interpretation mapping every tree $t \in \mathbb{T}^? D$ to the set of all of its traces. Furthermore, the product of a trace can be computed by performing a product in the ω-semigroup $\mathrm{SG}(\mathfrak{A})$, which can be done in MSO. Consequently, we can compute the set

$$\{\, \pi(s) \mid s \in \mathbb{T}^? U \text{ a trace of } t \,\}$$

in MSO. According to Lemma 4.5, the infimum of this set is the desired value $\pi(t)$. □

We can characterise the class of MSO-definable $\mathbb{T}^?$-algebras in terms of certain trace algebras derived from automata.

Definition 4.7. Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ be a tree automaton.

(a) The *ω-semigroup $\mathfrak{S}_{\mathcal{A}} = \langle S_1, S_\omega \rangle$ associated with $\mathcal{A}$* has the domains

$$S_1 := Q \times D \times Q + \{\bot\} \quad \text{and} \quad S_\omega := Q + \{\bot\} ,$$

where $D := \operatorname{rng} \Omega$ is the set of priorities used by $\mathcal{A}$ and $\perp$ is used as an 'error element'. The multiplication is given by

$$\langle p, k, p' \rangle \cdot \langle q, l, q' \rangle := \begin{cases} \langle p, m, q' \rangle & \text{if } p' = q \text{ and } m := \min \{k, l\}\,, \\ \perp & \text{otherwise}\,, \end{cases}$$

$$\langle p, k, p' \rangle \cdot q := \begin{cases} p & \text{if } p' = q\,, \\ \perp & \text{otherwise}\,, \end{cases}$$

$$\prod_{i<\omega} \langle p_i, k_i, p'_i \rangle := \begin{cases} p_0 & \text{if } p'_i = p_{i+1} \text{ for all } i\,, \\ \perp & \text{otherwise}\,, \end{cases}$$

and every product involving $\perp$ evaluates to $\perp$.

(b) The *transition algebra* of $\mathcal{A}$ is the $\mathbb{T}^{?}$-algebra

$$\mathfrak{T}(\mathcal{A}) := \operatorname{Up}(\operatorname{TA}(\mathfrak{S}_{\mathcal{A}}))\,.$$

(a) Let $\tau = \langle p, \bar{u} \rangle$ be a profile with $u_x = \langle k_x, q_x \rangle$, for $x \in \xi$. An element $a \in T_\xi(\mathcal{A})$ *represents* τ if it is of the form

$$a = p \sqcap \prod_{x \in \xi} \langle p, k_x, q_x \rangle(x) \quad \text{or} \quad a = \prod_{x \in \xi} \langle p, k_x, q_x \rangle(x)\,.$$

(The latter only if $\xi \neq \varnothing$.) ⌟

The following observation immediately follows from Proposition 4.2.

Lemma 4.8. *The transition algebra $\mathfrak{T}(\mathcal{A})$ is a trace algebra.*

Exercise 4.1. We call a tree automaton $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \Omega \rangle$ *deterministic* if, for every state $p \in Q$ and every letter $a \in \Sigma$, there exists a unique transition $\langle r, b, \bar{q} \rangle \in \Delta$ with $r = p$ and $b = a$. Prove that every language $K \subseteq \mathbb{T}^{?}_{\xi}\Sigma$ recognised by a deterministic automaton is recognised by a finitary trace algebra. ⌟

We can use transition algebras to give a more concrete characterisation of when an algebra is MSO-definable. But the price we pay for this is that we need a slightly more general notion of recognition.

Definition 4.9. Let $\mathfrak{A}$ and $\mathfrak{B}$ be $\mathbb{T}^?$-algebras.

(a) A *relational morphism* $\mathfrak{R} : \mathfrak{A} \to \mathfrak{B}$ is a subalgebra $\mathfrak{R} \subseteq \mathfrak{A} \times \mathfrak{B}$ such that the projection $R \to A$ to the first component is surjective.

(b) A subset $K \subseteq A_\xi$ is *recognised* by a relational morphism $R : \mathfrak{A} \to \mathfrak{B}$ if there exists a set $P \subseteq B_\xi$ such that

$$K = \{\, a \in A_\xi \mid \langle a, b\rangle \in R \text{ for some } b \in P \,\}\,. \qquad \lrcorner$$

The relational morphisms we are interested in are of the following form.

Definition 4.10. Let $\mathfrak{A}$ be a $\mathbb{T}^?$-algebra, $\mathcal{A}$ an automaton, and let $C \subseteq A$ be a set of generators of $\mathfrak{A}$. We denote by $\mathfrak{R}(\mathfrak{A}, \mathcal{A}) : \mathbb{T}\mathfrak{A} \to \mathfrak{T}(\mathcal{A})$ the relational morphism with domains

$$R_\xi(\mathfrak{A}, \mathcal{A}) := \{\, \langle \mathbb{T}^? \pi(t), \tau\rangle \mid t \in \mathbb{T}^?_\xi \mathbb{T}^? C \text{ and } \tau \in \mathfrak{T}(\mathcal{A}) \text{ represents the profile of some partial run of } \mathcal{A} \text{ on the input tree } \mathrm{flat}(t) \,\}\,, \qquad \lrcorner$$

To see that $\mathfrak{R}(\mathfrak{A}, \mathcal{A})$ is really a relational morphism we have to check that the projection $R(\mathfrak{A}, \mathcal{A}) \to A$ is surjective.

Lemma 4.11. *Let $\mathfrak{A}$ be an* MSO*-definable $\mathbb{T}^?$-algebra and $\mathcal{A}$ an automaton that has at least one partial run on every tree.*

(a) *The projection $p : \mathfrak{R}(\mathfrak{A}, \mathcal{A}) \to \mathbb{T}^?\mathfrak{A}$ is surjective and every fibre $p^{-1}(a)$ is finite.*

(b) *$\mathfrak{R}(\mathfrak{A}, \mathcal{A}) : \mathbb{T}^?\mathfrak{A} \to \mathfrak{T}(\mathcal{A})$ is a relational morphism.*

Proof. (a) Consider an element $a \in A_\xi$. Fix some transition δ of $\mathcal{A}$ for the input letter a and let $\tau \in T(\mathcal{A})$ be (an element represeting) the profile corresponding to δ. Then $\langle a, \tau\rangle \in R(\mathfrak{A}, \mathcal{A})$ and $p(\langle a, \tau\rangle) = a$. Hence, $a \in \operatorname{rng} p$. For the second statement, note that every domain $R_\xi(\mathfrak{A}, \mathcal{A}) \subseteq A_\xi \times T_\xi(\mathcal{A})$ is finite. Hence, so is $p^{-1}(a) \subseteq R_\xi(\mathfrak{A}, \mathcal{A})$, for $a \in A_\xi$.

(b) We have already proved in (a) that the projection to $\mathfrak{A}$ is surjective. To see that $R(\mathfrak{A}, \mathcal{A})$ is closed under products, fix a finite set $C \subseteq A$ of generators and consider a tree $t \in \mathbb{T}^? R(\mathfrak{A}, \mathcal{A})$. Let $r \in \mathbb{T}^?\mathbb{T}^? A$ and $s \in \mathbb{T}^? T(\mathcal{A})$

be the projections of t to the two components. For every $v \in \mathrm{dom}(t)$, it follows that there is some tree $T(v) \in \mathbb{T}^? \mathbb{T}^? C$ with $\mathbb{T}^? \pi(T(v)) = r(v)$ and $s(v)$ is (an element represeting) the profile of some run of $\mathcal{A}$ and $\mathrm{flat}(T(v))$. Consequently, $\pi(s)$ is the profile of some run of $\mathcal{A}$ on $\mathrm{flat}(\mathrm{flat}(T))$ and

$$\mathbb{T}^? \pi(\mathrm{flat}(T)) = \mathrm{flat}(\mathbb{T}^? \mathbb{T}^? \pi(T)) = \mathrm{flat}(r) .$$

This implies that $\pi(t) = \langle \mathrm{flat}(r), \pi(s) \rangle \in R(\mathfrak{A}, \mathcal{A})$. □

We obtain the following characterisation of when a $\mathbb{T}^?$-algebra is MSO-definable.

Theorem 4.12. *Let $\mathfrak{A}$ be a finitary $\mathbb{T}^?$-algebra and $C \subseteq A$ a finite set of generators. $\mathfrak{A}$ is MSO-definable if, and only if, there exists a trace algebra $\mathfrak{D}$ and a relational morphism $\mathfrak{R} : \mathbb{T}^? \mathfrak{A} \to \mathfrak{D}$ such that*

- *every fibre $p^{-1}(a)$ of the projection $p : R(\mathfrak{A}, \mathcal{A}) \to A$ is finite, and*
- *the relational morphism $(\mathrm{id} \times \pi) \circ \mathbb{T}^? \mathfrak{R}(\mathfrak{A}, \mathcal{A}) : \mathbb{T}^? \mathfrak{A} \to \mathfrak{D}$ recognises every preimage*

$$\pi^{-1}(a) , \quad \text{for } a \in A_{\varnothing} \cup A_{\{z\}} .$$

Proof. ($\Rightarrow$) Let $\mathfrak{A}$ be MSO-definable. Then all preimages $\pi^{-1}(a) \cap \mathbb{T}^? C$, for $a \in A_{\varnothing} \cup A_{\{z\}}$, are regular. We can combine the corresponding automata into a single automaton $\mathcal{A}$ that, for every $a \in A_{\varnothing} \cup A_{\{z\}}$, is equipped with a distinguished state q_a such that, when starting in q_a, $\mathcal{A}$ recognises all trees in $\pi^{-1}(a) \cap \mathbb{T}^? C$. Set $\mathfrak{D} := \mathfrak{T}(\mathcal{A})$ and let $\mathfrak{R}(\mathfrak{A}, \mathcal{A})$ be the relational morphism from Definition 4.10. We have seen above that $\mathfrak{D}$ is a trace algebra, the projection $p : R(\mathfrak{A}, \mathcal{A}) \to A$ is surjective, and all fibres $p^{-1}(a)$ are finite.

To conclude the proof, consider an element $a \in A_\xi$ of sort $\xi = \varnothing$ or $\xi = \{z\}$. Let $P_a \subseteq D_\xi$ be the set of all (elements representing) profiles of the form $\langle q_a, \langle\rangle\rangle$ or $\langle q_a, \langle k, p\rangle\rangle$ (depending on whether $\xi = \varnothing$ or $\xi = \{z\}$) where the priority k is arbitrary and p is a state such that, when starting in p,

$\mathcal{A}$ accepts the singleton tree with label z. For $t \in \mathbb{T}^?_\xi A$, it follows that

$$\begin{aligned}
\pi(t) = a \quad & \text{iff} \quad t = \mathbb{T}^?\pi(T)\,, \text{ for some } T \in \mathbb{T}^?\mathbb{T}^?C \text{ with} \\
& \qquad\quad \pi(\mathrm{flat}(T)) = a\,, \\
& \text{iff} \quad \text{there exist } T \in \mathbb{T}^?\mathbb{T}^?C \text{ and an accepting run on } T \\
& \qquad\quad \text{starting in the state } q_a \text{ such that } t = \mathbb{T}^?\pi(T) \\
& \text{iff} \quad \langle t, \tau\rangle \in R(\mathfrak{A}, \mathcal{A}) \text{ for some } \tau \in P_a\,.
\end{aligned}$$

($\Leftarrow$) By Proposition V.5.12, it is sufficient to show that the preimages $\pi^{-1}(a) \cap \mathbb{T}^?C$ are regular for elements a of sort ξ with $|\xi| \leq 1$. Hence, let $a \in A_\varnothing \cup A_{\{z\}}$ and set

$$C' := \{\, \tau \mid \langle \mathrm{sing}(c), \tau\rangle \in R(\mathfrak{A}, \mathcal{A}) \text{ for some } c \in C \,\}\,.$$

Note that C' is a finite set since, by assumption, all fibres of p are finite. We start by proving that

$$\langle t, \tau\rangle \in R(\mathfrak{A}, \mathcal{A}) \quad \text{iff} \quad \langle t, s\rangle \in R'\,, \quad \text{for some } s \in \pi^{-1}(\tau) \cap \mathbb{T}^?C',$$

where

$$R' := \{\, \langle t, s\rangle \mid \langle \mathrm{sing}(t(v)), s(v)\rangle \in R(\mathfrak{A}, \mathcal{A}) \text{ for all } v \,\}\,.$$

($\subseteq$) Suppose that $\langle t, \tau\rangle \in R(\mathfrak{A}, \mathcal{A})$. Then there exists a tree $T \in \mathbb{T}^?\mathbb{T}^?C$ such that $t = \mathbb{T}^?\pi(T)$ and τ is (an element representing) the profile of some partial run ρ on $\mathrm{flat}(T)$. The restriction of ρ to $\mathrm{dom}(T(v))$ forms a partial run on $T(v)$. Let $s(v)$ be its profile. Then $s \in \mathbb{T}^?D$, $\pi(s) = \tau$, and

$$\langle \mathrm{sing}(t(v)), s(v)\rangle \in R(\mathfrak{A}, \mathcal{A})\,, \quad \text{for all } v \in \mathrm{dom}(T)\,.$$

Hence, $\langle t, s\rangle \in R'$. Furthermore, by definition of C' we have $s \in \mathbb{T}^?C'$.

($\supseteq$) Suppose that $\langle t, s\rangle \in R'$ and $s \in \pi^{-1}(\tau) \cap \mathbb{T}^?C'$. Let $v \in \mathrm{dom}(t)$. Since $\langle \mathrm{sing}(t(v)), s(v)\rangle \in R(\mathfrak{A}, \mathcal{A})$, there exists a tree $T(v) \in \mathbb{T}^?C$ with $\pi(T(v)) = t(v)$ such that $s(v)$ is the profile of some partial run on $T(v)$.

It follows that $t = \mathbb{T}^{?}\pi(T)$ and $\tau = \pi(s)$ is the profile of some partial run on $\mathrm{flat}(T)$. Hence, $\langle t, \tau \rangle \in R(\mathfrak{A}, \mathcal{A})$.

To conclude the proof note that, by assumption, there exists a (finite) set $P \subseteq D_\xi$ such that

$$\begin{aligned}
\pi^{-1}(a) &= \left\{ t \in \mathbb{T}^{?}_{\xi}A \;\middle|\; \langle t, \tau \rangle \in R(\mathfrak{A}, \mathcal{A}) \text{ for some } \tau \in P \right\} \\
&= \bigcup_{\tau \in P} \left\{ t \in \mathbb{T}^{?}_{\xi}A \;\middle|\; \langle t, \tau \rangle \in R(\mathfrak{A}, \mathcal{A}) \right\} \\
&= \bigcup_{\tau \in P} \left\{ t \in \mathbb{T}^{?}_{\xi}A \;\middle|\; \langle t, s \rangle \in R', \text{ for some } s \in \pi^{-1}(\tau) \cap \mathbb{T}^{?}C' \right\}.
\end{aligned}$$

By Proposition 4.6 and the fact that the class of MSO-definable algebras is closed under taking finitary subalgebras, it follows that each preimage $\pi^{-1}(\tau) \cap \mathbb{T}^{?}C'$ forms a regular language. Furthermore, regular tree languages are closed under projections and inverse projections. Hence, each term in the above union is regular and, therefore, so is the union itself. □

5. Game Algebras

We consider games between two players, called $\Diamond$ and $\Box$, moving a token along the edges of a directed graph. The vertices of this graph are partitioned into two classes, one for Player $\Diamond$ and one for Player $\Box$. The class of the current vertex determines which of the two player may choose the next move. Every edge of the graph and every leaf (i.e., a vertex without outgoing edges) is labelled by an element of some ω-semigroup $\mathfrak{S}$. The winner of a given play of the game is determined by the product of the labels along the path traversed during this play. As we want to compose games we also equip these graphs with distinguished leaves where we can attach other games.

Definition 5.1. Let $\mathfrak{S} = \langle S_1, S_\omega \rangle$ be an ω-semigroup.

(a) A *game graph* over $\mathfrak{S}$ of sort $\xi \in \Xi$ is a structure of the form

$$\mathfrak{G} = \left\langle V_\Diamond, V_\Box, (E_a)_{a \in S_1}, (P_c)_{c \in S_\omega}, v_0, \bar{u} \right\rangle$$

where

- $V := V_\Diamond + V_\Box$ is the set of *positions*,
- $V_\Diamond$ is the set of *positions for Player* $\Diamond$,
- $V_\Box$ is the set of *positions for Player* $\Box$,
- $E_a \subseteq V \times V$ is the set of edges with label $a \in S_1$,
- $P_c \subseteq V$ is the set of vertices labelled $c \in S_\omega$,
- $v_0 \in V$ is the *initial position*, and
- $\bar{u} \in V^\xi$ is a ξ-tuple of *holes* or *variables*.

We assume that each variable u_x is a leaf of $\mathfrak{G}$, that the predicates P_c contain only non-variable leaves, and that every non-variable leaf belongs to exactly one P_c. We set

$$\mathrm{GG}(\mathfrak{S}) := (\mathrm{GG}_\xi(\mathfrak{S}))_{\xi \in \Xi} \in \mathsf{Pos}^\Xi\,,$$

where $\mathrm{GG}_\xi(\mathfrak{S})$ is the set of all game graphs over $\mathfrak{S}$ of sort ξ. We equip $\mathrm{GG}(\mathfrak{S})$ with the trivial ordering where all elements are incomparable.

(b) A *play* of $\mathfrak{G}$ is a sequence of edges that forms a path starting in the initial position v_0, and that is either infinite, or it ends in some leaf. Let p be a play and let $(a_i)_{i<\alpha}$ be the sequence of labels of the edges in p. The *outcome* of p is the element of $\mathrm{TA}_\xi(\mathfrak{S})$ given by

$$\begin{cases} \prod_{i<\omega} a_i & \text{if } p \text{ is infinite}\,, \\ a_0 \cdots a_{\alpha-1} \cdot c & \text{if } p \text{ ends in a position belonging to } P_c\,, \\ \langle a_0 \cdots a_{\alpha-1}, x\rangle & \text{if } p \text{ ends in } u_x\,. \end{cases}$$

(c) A *strategy* for Player $\alpha \in \{\Diamond, \Box\}$ is a function that returns the next move for the player when given the sequence of previous moves as input. Formally, we consider a strategy as a function σ mapping each prefix of a play that ends in a position for Player α to one of the outgoing edges.

We say that a play p *conforms* to a strategy σ for Player α if, for every prefix p_0 of p that ends in a position for Player α, the edge $e := \sigma(p_0)$ returned by σ is the next edge in p, that is, if $p_0 e$ is also a prefix of p. ⌟

Example. A *parity game* is a game for the ω-semigroup $\mathfrak{D}_n$, for some $n < \omega$, where $\mathfrak{D}_n = \langle D^n_1, D^n_\omega \rangle$ has the domains

$$D^n_1 := [n] \quad \text{and} \quad D^n_\omega := [2],$$

and the product

$$k \cdot l := \min\{k, l\}, \qquad k \cdot b := b, \qquad \prod_{i<\omega} k_i := \big(\liminf_{i<\omega} k_i\big) \bmod 2,$$

for $k, l, k_i \in D^n_1$ and $b \in D^n_\omega$. ⌟

We can equip the set $\mathrm{GG}(\mathfrak{S})$ with the structure of an $\mathbb{R}$-algebra.

Definition 5.2. Let $g \in \mathbb{R}_\xi \mathrm{GG}(\mathfrak{S})$. The product $\pi(g) \in \mathrm{GG}_\xi(\mathfrak{S})$ is formed as follows. We take the disjoint union of all game graphs $g(v)$, for $v \in \mathrm{dom}(g)$, plus single vertices for each variable in g. For each component $g(v)$, we delete all variables and we redirect each edge ending in a variable x to the initial position of $g(u)$, where u is the x-successor of v. ⌟

Lemma 5.3. $\langle \mathrm{GG}(\mathfrak{S}), \pi \rangle$ *forms an $\mathbb{R}$-algebra.*

Frequently, we would like to abstract away from the precise shape of the game graph and only consider the available strategies for the two players. This leads to the following algebra.

Definition 5.4. Let $\mathfrak{S}$ be an ω-semigroup.

(a) We set

$$\mathrm{Game}(\mathfrak{S}) := (\mathrm{Dn} \circ \mathrm{Up} \circ \mathrm{TA})(\mathfrak{S}).$$

(b) The *outcome map* $\mathrm{out} : \mathrm{GG}(\mathfrak{S}) \to \mathrm{Game}(\mathfrak{S})$ maps a game graph $\mathfrak{G}$ to the element

$$\mathrm{out}(\mathfrak{G}) := \sup_\sigma \inf_\tau a_{\sigma\tau},$$

where σ ranges over all strategies for Player $\Diamond$, τ ranges of the strategies for Player $\Box$, and $a_{\sigma\tau} \in \mathrm{TA}(\mathfrak{S})$ is the outcome of the unique play conforming to the strategies σ and τ. ⌟

Example. The game graph

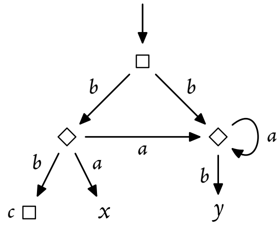

is mapped by the function out to the element

$$\Big[ba^{\omega} \sqcup \bigsqcup_{n<\omega} ba^{n}b(y)\Big] \sqcap \Big[bbc \sqcup ba(x) \sqcup ba^{\omega} \sqcup \bigsqcup_{n<\omega} ba^{n+1}b(y)\Big].$$

⌟

Lemma 5.5. *Let* Sgp_{ω} *be the category of* ω*-semigroups.* GG *and* Game *form functors* $\mathsf{Sgp}_{\omega} \to \mathsf{Alg}(\mathbb{R})$*, and* out : GG ⇒ Game *is a natural transformation.*

Exercise 5.1. (a) Prove that $\mathrm{GG}(\mathfrak{S})$ is generated by elements of the form

$$\Diamond(a_0(x_0), \dots, a_{n-1}(x_{n-1}))\,, \quad \Box(a_0(x_0), \dots, a_{n-1}(x_{n-1}))\,, \quad c\,,$$

for $a_0, \dots, a_{n-1} \in S_1$ and $c \in S_{\omega}$, where $\Diamond(a_0(x_0), \dots)$ denotes the game with an initial position belonging to Player $\Diamond$ and n outgoing edges labelled $a_0, \dots, a_{n-1}$ that lead to the variables $x_0, \dots, x_{n-1}$, respectively, and similarly for $\Box(a_0(x_0), \dots)$. c is a single position game with label c.

(b) Prove that $\mathrm{out} : \mathrm{GG}(\mathfrak{S}) \to \mathrm{Game}(\mathfrak{S})$ is the unique $\mathbb{R}$-morphism satisfying

$$\begin{aligned}
\mathrm{out}\big(\Diamond(a_0(x_0), \dots, a_{n-1}(x_{n-1}))\big) &:= a_0(x_0) \sqcup \dots \sqcup a_{n-1}(x_{n-1})\,,\\
\mathrm{out}\big(\Box(a_0(x_0), \dots, a_{n-1}(x_{n-1}))\big) &:= a_0(x_0) \sqcap \dots \sqcap a_{n-1}(x_{n-1})\,,\\
\mathrm{out}(c) &:= c\,,
\end{aligned}$$

for $a_0, \dots, a_{n-1} \in S_1$ and $c \in S_{\omega}$. ⌟

Recall the characterisation of automata in terms of Automaton-Pathfinder games from Section V.5. We can use the machinery of games developed above to turn this characterisation into a morphism $\mathbb{T}^{?}\Sigma \to \mathrm{Game}(\mathfrak{S})$.

Theorem 5.6. *For every tree automaton $\mathcal{A}$ there exists a finite ω-semigroup $\mathfrak{S}$ and a morphism $\varphi : \mathbb{T}^{?}\Sigma \to \mathrm{Game}(\mathfrak{S})$ recognising $L(\mathcal{A})$.*

Proof. Let $\mathfrak{S} = \mathfrak{S}_{\mathcal{A}}$ be the ω-semigroup from associated with $\mathcal{A}$ as in Definition 4.7. For a tree $t \in \mathbb{T}^{?}_{\xi}\Sigma$, we set

$$\varphi(t) := \sup \{ \tau \mid \tau \text{ the profile of some partial run } \rho \text{ on } t \} ,$$

where we encode a profile $\tau = \langle p, \bar{U} \rangle$ as a meet

$$p \sqcap \prod_{x \in \xi} \prod_{\langle k,q \rangle \in U_x} \langle p, k, q \rangle (x) .$$

Then it follows that

$$\varphi(t) = \mathrm{out}(\mathfrak{G}_t) ,$$

where $\mathfrak{G}_t$ is the Automaton-Pathfinder game on the input tree t. □

The game returned by the morphism φ from the preceding theorem is not quite the Automaton-Pathfinder game, but a related one. We can return the exact Automaton-Pathfinder game if we make the game algebra slightly more complicated. In the preceding theorem we have encoded the states of the automaton in the elements of the semigroup. But in the Automaton-Pathfinder game the states are part of the positions of the game. To model this in our algebra we will use games with several initial positions, one for each state of the automaton. This is equivalent to using tuples of games with one component per state. The formal definition is as follows.

Definition 5.7. Let $\mathfrak{A}$ be a $\mathbb{T}^{?}$-algebra and $\eta \in \Xi$ a set of sorts. The η-th *matrix power* $\mathfrak{A}^{[\eta]}$ of $\mathfrak{A}$ is the $\mathbb{T}^{?}$-algebra with domains

$$A^{[\eta]}_{\xi} := (A_{\eta \times \xi})^{\eta} , \quad \text{for } \xi \in \Xi .$$

The product is defined as follows. Given a tree $t \in \mathbb{T}^{?}A^{[\eta]}$, we construct a graph $g \in \mathbb{R}^{?}A$ with set of vertices

$$\mathrm{dom}(g) := \eta \times \mathrm{dom}(t)$$

and labelling

$$g(\langle y, v\rangle) := a_y \quad \text{where } \bar{a} := t(v)\,.$$

There is an $\langle y, x\rangle$-labelled edge from $\langle z, v\rangle$ to $\langle y, \mathrm{suc}_x(v)\rangle$. For $y \in \eta$, let $s_y := \mathrm{gun}(g)$ be the unravelling of g starting at the vertex $\langle y, \langle\rangle\rangle$. We set

$$\pi(g) := (\pi(s_y))_{y \in \eta}\,,$$

where the products on the right-hand side are computed in $\mathfrak{A}$. ⌟

Example. Let $\mathcal{A} = \langle Q, \Sigma, \varnothing, \Delta, q_0, \Omega\rangle$ be a tree automaton (we omit the product σ on Q since we are dealing with ranked trees, not unranked ones), and let $\mathfrak{D}_n$ be the ω-semigroup from the example on page 503. We define a $\mathbb{T}^?$-morphism $\mathrm{APG}_{\mathcal{A}} : \mathbb{T}^?\Sigma \to \mathrm{GG}(\mathfrak{D}_n)^{[Q]}$ by

$$\mathrm{APG}_{\mathcal{A}}(\mathrm{sing}(a)) := \bar{g}\,, \quad \text{for } a \in \Sigma_\xi\,,$$

where, for $p \in Q$, g_p is the game with a initial position belonging to Player $\Diamond$ (Automaton) with one outgoing edge for each transition $\langle p, a, \bar{q}\rangle \in \Delta$ matching the state p and the input letter a. The label of these edges is $\Omega(p)$. The edge belonging to $\langle p, a, \bar{q}\rangle$ leads to a position for Player $\Box$ (Pathfinder) with one outgoing edge, for each variable $x \in \xi$, and label $\Omega(q_x)$. These edges lead to the variable $\langle q_x, x\rangle \in Q \times \xi$. For every tree $t \in \mathbb{T}^?\Sigma$, it follows that $\mathrm{APG}_{\mathcal{A}}(t)$ is the Automaton-Pathfinder game of $\mathcal{A}$ on input t. ⌟

Branch-Continuity

In the same way we built trace algebras by taking meets of semigroup elements, we can construct more complicated algebras from semigroup elements by taking both meets and joins. It turns out that the resulting algebras are exactly the quotients of game algebras. The abstract definition is as follows.

Definition 5.8. Let $\mathfrak{A}$ be a $\mathbb{T}^?$-algebra.

(a) A subalgebra $\mathfrak{U} \subseteq \mathfrak{A}$ is a *skeleton* of $\mathfrak{A}$ if

- $\mathfrak{U}$ is semigroup-like,

- the embedding $U \to A$ is meet-distributive with induced product π,
- the closure of U under meets is a set of join-generators of $\mathfrak{A}$.

(b) $\mathfrak{A}$ is *branch-continuous* if it is join-distributive, lattice-distributive, and it has a skeleton. ⌟

Proposition 5.9. *Let $\mathfrak{A}$ be a $\mathbb{T}^?$-algebra. The following statements are equivalent.*

(1) *$\mathfrak{A}$ is branch-continuous.*

(2) *$\mathfrak{A}$ is completely ordered and there exists a surjective $\mathbb{T}^?$,Up,Dn-morphism $\varphi : \mathrm{Game}(\mathfrak{S}) \to \mathfrak{A}$, for some ω-semigroup $\mathfrak{S}$.*

(3) *$\mathfrak{A}$ is join-distributive, lattice-distributive, and there exists a $\mathbb{T}^?$,Up-morphism $\varphi : \mathfrak{D} \to \mathfrak{A}$ from some trace algebra $\mathfrak{D}$ such that $\operatorname{rng}\varphi$ is a set of join-generators of $\mathfrak{A}$.*

Proof. (1) ⇒ (3) Suppose that $\mathfrak{A}$ is branch-continuous. By definition, this implies that $\mathfrak{A}$ is join-distributive and lattice-distributive. Let $\mathfrak{U} \subseteq \mathfrak{A}$ be its skeleton and C the closure of U under meets. We have to check three conditions.

(I) The set C induces a subalgebra $\mathfrak{C}$ of $\mathfrak{A}$.

(II) The inclusion $i : \mathfrak{C} \to \mathfrak{A}$ is a $\mathbb{T}^?$,Up-morphism.

(III) $\mathfrak{C}$ is a trace algebra.

(I) Fix $t \in \mathbb{T}^? C$. As C is a set of join-generators, there exists some tree $s \in \mathbb{T}^? \mathrm{Up}(U)$ with $t = \mathbb{T}^? \inf(s)$. Since the inclusion $j : U \to A$ is meet-distributive with induced product π, it follows that

$$\pi(t) = \pi(\mathbb{T}^? \inf(s)) = \inf \mathrm{Up}(\pi)(\mathrm{dist}(s)) .$$

As $\mathrm{dist}(s) \in \mathrm{Up}(\mathbb{T}^? U)$ and U is closed under products, it follows that $\pi(t)$ is a meet of elements of U. Hence, $\pi(t) \in C$.

(II) As $\mathfrak{C}$ is a subalgebra of $\mathfrak{A}$, the inclusion is a $\mathbb{T}^?$-morphism. Hence, it remains to prove that it also preserves meets. Fix a set $X \subseteq C$. We have to show that the meet $\inf_A X$ of X in A is the same as its meet $\inf_C X$ in C. As $\inf_C X$ is a lower bound of X in A, we have

$$\inf_C X \leq \inf_A X .$$

For a contradiction, suppose that $\inf_C X < \inf_A X$. Then $\inf_A X \notin C$. By definition of C we can find, for every $x \in X$, some set $P_x \subseteq U$ with $\inf_A P_x = x$. Setting

$$P := \bigcup_{x \in X} P_x$$

it follows that

$$\inf_A X = \inf_A \{ \inf_A P_x \mid x \in X \} = \inf_A P \in C .$$

A contradiction.

(iii) By (ii) and the definition of C, $\mathfrak{U}$ is a semigroup-like subalgebra of $\mathfrak{C}$ and all elements of C can be written as meets of elements of U. Hence, it is sufficient to show that $\mathfrak{C}$ is meet-distributive. Since the inclusion $j : U \to A$ is meet-distributive with induced product π, so is its corestriction $k : U \to C$. Consequently, the meet-distributivity of $\mathfrak{C}$ follows by Proposition II.4.10.

(3) ⇒ (2) Suppose that $\mathfrak{A}$ is join-distributive, lattice-distributive, and that $\varphi : \mathfrak{D} \to \mathfrak{A}$ is a $\mathbb{T}^?$,Up-morphism as above. By Proposition 4.2, there exists a surjective $\mathbb{T}^?$,Up-morphism $\psi : \mathrm{Up}(\mathrm{TA}(\mathfrak{S})) \to \mathfrak{D}$, for some ω-semigroup $\mathfrak{S}$. Since $\mathfrak{A}$ is join-distributive, it is the reduct of a $\mathbb{T}^?$,Dn-bialgebra. We can therefore use the adjunction from Theorem II.4.3 to extend the $\mathbb{T}^?$-morphism $\varphi \circ \psi : \mathrm{Up}(\mathrm{TA}(\mathfrak{S})) \to \mathfrak{A}$ to an $\mathbb{T}^?$,Dn-morphism $\chi : \mathrm{Dn}(\mathrm{Up}(\mathrm{TA}(\mathfrak{S}))) \to \mathfrak{A}$ satisfying

$$\chi \circ \mathrm{pt} = \varphi \circ \psi .$$

To conclude the proof it remains to show that χ is surjective and that it preserves meets. For surjectivity, let $a \in A$. As C is a set of join-generators, there exists some set $X \subseteq C$ with $a = \sup X$. Since $X \subseteq \mathrm{rng}\, \chi$ and χ preserves joins, it follows that $a \in \mathrm{rng}\, \chi$.

To check that χ is an Up-morphism, we use Lemma II.4.6. According to this lemma, it is sufficient to prove that χ preserves meets of elements in $\mathrm{Up}(\mathrm{TA}(S))$. Hence, let $X \subseteq \mathrm{Up}(\mathrm{TA}(S))$. Since pt is injective and we have shown in (the dual of) Lemma 3.2 that it preserves meets, we have

$$\begin{aligned}\chi(\inf \mathrm{pt}[X]) = \chi(\mathrm{pt}(\inf X)) &= \varphi(\psi(\inf X)) \\ &= \inf \varphi(\psi[X]) = \inf \chi(\mathrm{pt}[X])\,.\end{aligned}$$

(2) ⇒ (1) It is sufficient to prove the following two claims.

(I) Each game algebra $\mathrm{Game}(\mathfrak{S})$ is branch-continuous.

(II) If $\varphi : \mathfrak{A} \to \mathfrak{B}$ is a surjective $\mathbb{T}^?$,Up,Dn-morphism and $\mathfrak{A}$ is branch-continuous, then so is $\mathfrak{B}$.

(I) It follows by Theorem II.4.3 (applied to the functors $\mathbb{P} :=$ Dn and $\mathbb{M} := \mathbb{T}^?$) that $\mathrm{Game}(\mathfrak{S})$ is a dist-bialgebra. In particular, it is a completely ordered $\mathbb{T}^?$-algebra. Similarly, lattice-distributivity of $\mathrm{Game}(\mathfrak{S})$ follows by Theorem II.4.3 (applied to the functors $\mathbb{P} :=$ Dn and $\mathbb{M} :=$ Up). Hence, it remains to prove that $\mathrm{Game}(\mathfrak{S})$ has a skeleton.

Let U be the image of the canonical embedding $\mathrm{TA}(\mathfrak{S}) \to \mathrm{Game}(\mathfrak{S})$ and let $C := \mathrm{Up}(U)$ be its closure under meets. Let $\mathfrak{U}$ and $\mathfrak{C} = \mathrm{Up}(\mathfrak{U})$ be the corresponding subalgebras. We have to check three properties.

- The fact that $\mathfrak{U}$ is semigroup-like follows by Proposition 2.7.
- The fact that the embedding $\mathrm{pt} : \mathfrak{U} \to \mathfrak{C}$ is meet-distributive follows by Lemma II.4.8 since $\mathfrak{C} \cong \mathrm{Up}(\mathfrak{U})$.
- Finally, the fact that C is a set of join-generators of $\mathrm{Game}(S)$ holds since $\mathrm{Game}(S) = \mathrm{Dn}(C)$.

(II) The fact that $\mathfrak{B}$ is join-distributive and lattice-distributive follows immediately by Lemma II.4.5 (applied to the two distributive laws dist : $\mathbb{T}^? \circ \mathrm{Dn} \Rightarrow \mathrm{Dn} \circ \mathbb{T}^?$ and skolem : $\mathrm{Dn} \circ \mathrm{Up} \Rightarrow \mathrm{Up} \circ \mathrm{Dn}$).

Hence, it remains to prove that $\mathfrak{B}$ has a skeleton. Let $\mathfrak{U}$ be the skeleton of $\mathfrak{A}$ and C its closure under meets. Set $V := \varphi[U]$ and $D := \varphi[C]$. We claim that D is a skeleton of $\mathfrak{B}$. As φ is a surjective Up,Dn-morphism, D is the closure of V under meets and B is the closure of D under joins.

To see that the subalgebra $\mathfrak{V}$ induced by V is semigroup-like, we use Proposition 2.7 to find a surjective $\mathbb{T}^?$-morphism $\rho : \mathrm{TA}(\mathfrak{S}) \to \mathfrak{U}$, for some ω-semigroup $\mathfrak{S}$. Then $\varphi \circ \rho$ is a surjective $\mathbb{T}^?$-morphism $\mathrm{TA}(\mathfrak{S}) \to \mathfrak{V}$ and the claim follows by Proposition 2.7.

Finally, we have to check that the embedding $j : \mathfrak{V} \to \mathfrak{B}$ is meet-distributive. As the embedding $i : \mathfrak{U} \to \mathfrak{A}$ is meet-distributive with induced product π, we have

$$\begin{aligned}
&\pi \circ \mathbb{T}^{?}\mathrm{inf} \circ \mathbb{T}^{?}\mathrm{Up}(j) \circ \mathbb{T}^{?}\mathrm{Up}(\varphi) \\
&= \pi \circ \mathbb{T}^{?}\mathrm{inf} \circ \mathbb{T}^{?}\mathrm{Up}(\varphi \circ i) \\
&= \pi \circ \mathbb{T}^{?}(\varphi \circ \mathrm{inf}) \circ \mathbb{T}^{?}\mathrm{Up}(i) \\
&= \varphi \circ \pi \circ \mathbb{T}^{?}\mathrm{inf} \circ \mathbb{T}^{?}\mathrm{Up}(i) \\
&= \varphi \circ \mathrm{inf} \circ \mathrm{Up}(i \circ \pi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(\varphi \circ i \circ \pi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(j \circ \varphi \circ \pi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(j \circ \pi \circ \mathbb{T}^{?}\varphi) \circ \mathrm{dist} \\
&= \mathrm{inf} \circ \mathrm{Up}(j \circ \pi) \circ \mathrm{dist} \circ \mathbb{T}^{?}\mathrm{Up}(\varphi) .
\end{aligned}$$

As φ is surjective and $\mathbb{T}^{?}$ and Up preserve surjectivity, it follows that

$$\pi \circ \mathbb{T}^{?}\mathrm{inf} \circ \mathbb{T}^{?}\mathrm{Up}(j) = \mathrm{inf} \circ \mathrm{Up}(j \circ \pi) \circ \mathrm{dist} ,$$

as desired. □

As for trace algebras, the product in a branch-continuous algebra can be reduced to a product in an ω-semigroup.

Lemma 5.10. *Let $\mathfrak{A}$ be a branch-continuous tree algebra, $\mathfrak{U} \subseteq \mathfrak{A}$ a skeleton of $\mathfrak{A}$, and C the closure of U under meets.*

(a) $\pi(t) = \sup \{ \pi(s) \mid s \in \mathbb{T}^{?}C ,\ s \leq^{\mathbb{T}^{?}} t \} ,\qquad$ *for $t \in \mathbb{T}^{?}A$,*

(b) $\pi(t) = \inf \{ \pi(u) \mid u \in \mathbb{T}^{?}U \text{ a trace of } t \} ,\qquad$ *for $t \in \mathbb{T}^{?}C$.*

Proof. (a) As C is a set of join-generators, we have

$$\pi(t) = \pi(\mathbb{T}^{?}\sup(s)) , \quad \text{where } s(v) := \{ c \in C \mid c \leq t(v) \} .$$

Consequently, the equation follows by join-distributivity.

(b) follows by Lemma 4.5 since the subalgebra $\mathfrak{C}$ of $\mathfrak{A}$ induced by C is a trace algebra. □

Since the conditions in the above lemma can be expressed in MSO, we obtain the following corollary.

Corollary 5.11. *Every finitary branch-continuous $\mathbb{T}^?$-algebra is* MSO*-definable.*

Note that the functor Game provides an adjunction between the category of ω-semigroups and the category of all branch-continuous algebras.

Proposition 5.12. *Let $\mathfrak{A}$ be a branch-continuous $\mathbb{T}^?$-algebra and let $\mathfrak{S}$ be an ω-semigroup. For every morphism $\varphi : \mathfrak{S} \to \mathrm{SG}(\mathfrak{A})$ of ω-semigroups, there exists a unique $\mathbb{T}^?$,Up,Dn-morphism $\hat{\varphi} : \mathrm{Game}(\mathfrak{S}) \to \mathfrak{A}$ extending φ.*

Proof. Let $\mathfrak{U} := \mathrm{TA}(\mathfrak{S})$. By Proposition 2.6, there exists a unique $\mathbb{T}^?$-morphism $\varphi_0 : \mathfrak{U} \to \mathfrak{A}$ extending φ. Using Theorem II.4.3 twice, we can first extend φ_0 to a unique $\mathbb{T}^?$,Up-morphism $\varphi_1 : \mathrm{Up}(\mathfrak{U}) \to \mathfrak{A}$ satisfying $\varphi_1 \circ \mathrm{pt}^{\mathrm{Up}} = \varphi_0$, and then to a unique $\mathbb{T}^?$,Dn-morphism $\hat{\varphi} : \mathrm{Dn}(\mathrm{Up}(\mathfrak{U})) \to \mathfrak{A}$ satisfying $\hat{\varphi} \circ \mathrm{pt}^{\mathrm{Dn}} = \varphi_1$. Finally, note that Lemma II.4.6 implies that $\hat{\varphi}$ preserves arbitrary meets. □

The reason why branch-continuous algebras are of interest is that they provide an alternative characterisation of the class of regular tree languages.

Theorem 5.13. *Let $K \subseteq \mathbb{T}^?\Sigma$. The following statements are equivalent.*

(1) *K is regular.*

(2) *K is recognised by a morphism into a finitary branch-continuous $\mathbb{T}^?$-algebra.*

(3) *K is recognised by an algebra of the form $\mathrm{Game}(\mathfrak{S})$, for some finite ω-semigroup $\mathfrak{S}$.*

Proof. (2) ⇒ (1) follows by Corollary 5.11; (1) ⇒ (3) by Theorem 5.6; and (3) ⇒ (2) by Proposition 5.9. □

By Theorem III.5.11, we obtain the following description of MSO-definable algebras.

Corollary 5.14. *The pseudo-variety of all* MSO*-definable* $\mathbb{T}^{?}$*-algebras is generated by the class of all algebras of the form* $\mathrm{Game}(\mathfrak{S})$*, for finite* $\mathfrak{S}$*.*

It follows that game algebras and, more generally, branch-continuous algebras play a similar role as the MSO-definable ones. The reason we usually work with the latter is that the former do not form a pseudo-variety: the class of branch-continuous algebras is not closed under finitely-generated subalgebras.

Notes

Most of the material in this chapter is new. Section 4 is taken from [6].

Recommended Literature

M. Bojańczyk, *Languages Recognises 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. Adámek, H. Herrlich, and G. E. Strecker, *Abstract and Concrete Categories*, John Wiley and Sons, 1990.

[2] J. T. Baldwin, *Categoricity*, American Mathematical Society, 2009.

[3] J. Barwise and S. Feferman, *Model-Theoretic Logics*, Springer, 1985.

[4] J. Beck, *Distributive Laws*, in Seminar on triples and categorical homology theory, B. Eckmann, ed., Lecture Notes in Mathematics 80, Springer, 1969, pp. 119–140.

[5] A. Blumensath, *Monadic Second-Order Model Theory*. unpublished lecture notes, `https://www.fi.muni.cz/~blumens/MSO.pdf`.

[6] ———, *Regular Tree Algebras*, Logical Methods in Computer Science, 16 (2020), pp. 16:1–16:25.

[7] ———, *Algebraic Language Theory for Eilenberg–Moore Algebras*, Logical Methods in Computer Science, 17 (2021), pp. 6:1–6:60.

[8] ———, *The Power-Set Operation for Tree Algebras*, Logical Methods in Computer Science, 19 (2023), pp. 9:1–9:47.

[9] M. Bojańczyk, *Recognisable languages over monads*. unpublished note, arXiv:1502.04898v1.

[10] ———, *Languages Recognises 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.

[11] M. Bojańczyk and B. Klin, *A Non-Regular Language of Infinite Trees That is Recognizable by a Finite Algebra*, Logical Methods in Computer Science, 15 (2019).

[12] M. Bojanczyk, B. Klin, and J. Salamanca, *Monadic monadic second order logic*. arXiv:2201.09969, unpublished.

[13] M. Bojańczyk and H. Michalewski, *Some connections between universal algebra and logics for trees*. arXiv:1703.04736, unpublished.

[14] M. Bojańczyk, H. Straubing, and I. Walukiewicz, *Wreath products of forest algebras, with applications to tree logics*, Logical Methods in Computer Science, 8 (2012).

[15] F. Borceux, *Handbook of Categorical Algebra*, vol. 1, Cambridge University Press, 1994.

[16] ———, *Handbook of Categorical Algebra*, vol. 2, Cambridge University Press, 1994.

[17] C. Camino, V. Diekert, B. Dundua, M. Marin, and G. Sénizergues, *Regular matching problems for infinite trees*, Logical Methods in Computer Science, 18 (2022), pp. 25:1–25:38.

[18] A. Carboni and P. T. Johnstone, *Connected Limits, Familial Representability and Artin Glueing*, Math. Struct. Comput. Sci., 5 (1995), pp. 441–459.

[19] 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.

[20] 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).

[21] B. A. Davey and H. A. Priestley, *Introduction to Lattices and Order*, Cambridge University Press, 2002.

[22] R. Diaconescu, *Institution-independent Model Theory*, Birkhäuser, 2008.

[23] N. Gambino and J. Kock, *Polynomial Functors and Polynomial Monads*, Math. Proc. Cambridge Phil. Soc., 154 (2013), pp. 153–192.

[24] R. Garner, *The Vietoris Monad and Weak Distributive Laws*, Appl. Categorical Struct., 28 (2020), pp. 339–354.

[25] A. Goy, D. Petrisan, and M. Aiguier, *Powerset-like monads weakly distribute over themselves in toposes and compact hausdorff spaces*, in 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12–16, 2021, Glasgow, Scotland (Virtual Conference), N. Bansal, E. Merelli, and J. Worrell, eds., vol. 198 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, pp. 132:1–132:14.

[26] T. Hafer and W. Thomas, *Computation Tree Logic CTL* and Path Quantifiers in the Monadic Theory of the Binary Tree*, in Automata, Languages and Programming, 14th International Colloquium, ICALP87, Karlsruhe, Germany, July 13–17, 1987, Proceedings, T. Ottmann, ed., vol. 267 of Lecture Notes in Computer Science, Springer, 1987, pp. 269–279.

[27] B. Jacobs, *Trace Semantics for Coalgebras*, in Proceedings of the Workshop on Coalgebraic Methods in Computer Science, CMCS 2004, Barcelona, Spain, March 27–29, 2004, J. A. and S. Milius, eds., vol. 106 of Electronic Notes in Theoretical Computer Science, 2004, pp. 167–184.

[28] C. Löding, *Automata on Infinite Trees*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 265–302.

[29] C. Löding and W. Thomas, *Automata on Finite Trees*, in Handbook of Automata Theory, J.-É. Pin, ed., European Mathematical Society, 2021, pp. 235–264.

[30] E. Manes and P. S. Mulry, *Monad compositions I: general constructions and recursive distributive laws*, Theory and Applications of Categories, 18 (2007), pp. 172–208.

[31] ———, *Monad compositions II: Kleisli strength*, Math. Struct. Comput. Sci., 18 (2008), pp. 613–643.

[32] 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.

[33] F. Moller and A. Rabinovitch, *Counting on CTL*: on the expressive power of monadic path logic*, Information and Computation, 184 (2003), pp. 147–159.

[34] J.-É. Pin, *Mathematical Foundations of Automata Theory*. unpublished lecture notes, `http://www.liafa.jussieu.fr/~jep/PDF/MPRI/MPRI.pdf`.

[35] A. Potthoff, *Logische Klassifizierung regulärer Baumsprachen*, Ph. D. Thesis, Universität Kiel, Kiel, 1994.

[36] ———, *First-order logic on finite trees*, in TAPSOFT'95: Theory and Practice of Software Development, 6th International Joint Conference CAAP/FASE, Aarhus, Denmark, May 22–26, 1995, Proceedings, P. D. Mosses, M. Nielsen, and M. I. Schwartzbach, eds., vol. 915 of Lecture Notes in Computer Science, Springer, 1995, pp. 125–139.

[37] L. Ribes and P. Zalesskii, *Profinite Groups*, Springer-Verlag, 2010.

[38] J. Salamanca, *Unveiling eilenberg-type correspondences: Birkhoff's theorem for (finite) algebras + duality*. unpublished note, arXiv:1702.02822v1.

[39] D. I. Spivak and N. Niu, *Polynomial Functors: A General Theory of Interaction*. lecture notes.

[40] W. Thomas, *Automata on Infinite Objects*, in Handbook of Theoretical Computer Science, J. van Leeuwen, ed., vol. B, Elsevier, Amsterdam, 1990, pp. 135–191.

[41] ———, *Languages, Automata, and Logic*, in Handbook of Formal Languages, G. Rozenberg and A. Salomaa, eds., vol. 3, Springer, New York, 1997, pp. 389–455.

[42] 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.

[43] I. Walukiewicz, *Monadic second-order logic on tree-like structures*, Theoretical Computer Science, 275 (2002), pp. 311–346.

[44] 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.

[45] M. Zwart and D. Marsden, *No-Go Theorems for Distributive Laws*, Log. Methods Comput. Sci., 18 (2022), pp. 13:1–13:61.

Symbol Index

Chapter I

$[n]$	$\{0, \dots, n-1\}$, 10
$\Uparrow X$	upwards closure, 10
$\Downarrow X$	downwards closure, 10
Set	category of sets, 10
Pos	category of posets, 10
Top	category of topological spaces, 10
Met	category of bounded metric spaces, 10
$\mathfrak{G}$-Set	category of $\mathfrak{G}$-sets, 11
$\mathbb{J}X$	discrete object, 12
$\mathbb{V}A$	underlying set, 12
A^X	power of A, 21
$\mathrm{dom}(s)$	domain of s, 24
$\Pi(\mathcal{C})$	free product completion, 28
$s\ \theta^{\mathbb{F}}\ t$	lift of θ, 41
$\simeq_{\mathrm{sh}}$	the same shape, 41

Chapter II

$\mathrm{im}\,\varphi$	image, 100
$\ker\varphi$	kernel, 100
$\varepsilon \perp \mu$	orthogonal morphisms, 103
$\langle\!\langle X \rangle\!\rangle_A$	generated subobject, 110
$\langle\!\langle X \rangle\!\rangle_{\mathfrak{A}}$	generated subalgebra, 110
$\mathfrak{A}_\rho$	ρ-reduct of $\mathfrak{A}$, 122
$A \uparrow \Xi$	expansions to sorts in Ξ, 126
A_Δ	reduct to sorts in Δ, 126
A/θ	quotient, 142
$\mathrm{Cong}(A)$	congruences of A, 143
$\mathrm{Cong}_{\mathbb{M}}(\mathfrak{A})$	$\mathbb{M}$-congruences of $\mathfrak{A}$, 143
$[a]_\sim$	equivalence class, 145
$A/\sqsubseteq$	quotient, 146
$[a]_\sqsubseteq$	$\sqsubseteq$-class, 146
$E_{\mathcal{P}}$	epimorphisms making every $P \in \mathcal{P}$ projective, 152
$\mathsf{S}(\mathcal{K})$	closure under subobjects, 154
$\mathsf{H}(\mathcal{K})$	closure under quotients, 154
$\mathsf{P}_\kappa(\mathcal{K})$	closure under products, 155

Chapter III

$f^{-1}[\kappa]$	inverse image of κ, 169
$\mathrm{TC}(\theta)$	transitive closure, 177
$\mathbb{L}\mathfrak{A}$	language algebra, 185
$\langle\!\langle C \rangle\!\rangle_{\mathbb{L}}$	subalgebra generated by C, 186
syn_κ	syntactic morphism, 187
$\mathrm{Syn}(\kappa)$	syntactic algebra, 187
$\mathcal{T}$	term algebras, 206

$\mathcal{R}$ recognisers, 206
$\mathfrak{A} \models s \leq t$ satisfaction, 237
$\mathrm{Th}(\mathcal{V})$ theory of $\mathcal{V}$, 237

Chapter IV

$\mathrm{Mod}(\varphi)$ models of φ, 243
$\langle\!\langle M; \varphi \rangle\!\rangle$ truth value of φ in the model M, 243
$\mathrm{Th}_L(M)$ L-theory, 243
$\mathrm{eq}(L)$ L-equivalence, 244
$\mathcal{L}[\Omega]$ Ω-extension, 245

Chapter V

Ξ sorts, 279
Ξ_+ extended sorts, 279
$\mathbb{G}$ nondeterministic rooted graphs, 279
$\mathrm{dom}_+(g)$ extended domain, 280
$\mathrm{dom}(g)$ domain, 280
$\mathbb{R}$ deterministic rooted graphs, 280
$\mathbb{F}^{\mathrm{fin}}A$ finite forests, 280
$\mathbb{F}^{\times\mathrm{fin}}A$ finite forests, 280
$\mathrm{gun}(g)$ unravelling, 281
sing singleton graph, 281
flat flattening for graphs and trees, 281
$\mathrm{MSO}[\leq_{\mathrm{so}}]$ MSO with successor ordering, 286
$\mathrm{MSO}[\leq_{\mathrm{pf}}]$ MSO with forest ordering, 286
$\mathrm{MSO}[\mathrm{suc}]$ MSO with successor relations, 286
CMSO counting MSO, 286
$\equiv_L$ logical equivalence, 287
$a + b$ horizontal product, 289
$a \cdot b$ vertical product, 289
$s_0 ш \cdots ш s_{n-1}$ shuffle, 293
$\mathbb{C}^{\mathrm{reg}}$ regular linear orders, 294
$a \leq_{\mathrm{L}} b$ L-order, 298
$a \leq_{\mathrm{R}} b$ R-order, 298
$a \leq_{\mathrm{J}} b$ J-order, 298
$a \equiv_{\mathrm{L}} b$ L-relation, 298
$a \equiv_{\mathrm{R}} b$ R-relation, 298
$a \equiv_{\mathrm{J}} b$ J-relation, 298
$\mathcal{F}(t)$ set of factorisations of t, 320

Chapter VI

ML modal logic, 331
CTL computation tree logic, 331
PDL propositional dynamic logic, 331
EF reachability fragment of CTL, 331
MPL monadic path logic, 331
MCL monadic chain logic, 331
cEF counting variant of EF, 331
cCTL^* counting variant of CTL^*, 331
cwCTL^* counting variant of wCTL^*, 331
cPDL counting variant of PDL, 331
$\mathrm{cTL}[\mathcal{K}]$ counting $\mathcal{K}$-temporal logic, 332
$\models_{\mathrm{f}}$ satisfiability for forests, 332

$\models_t$	satisfiability for trees, 332
$\mathrm{TL}[\mathcal{K}]$	$\mathcal{K}$-temporal logic, 333
$\mathbb{F}^{\mathrm{fb}}$	finitely branching forests, 333
$t[\bar{K}]$	labelling of t, 335
$L[\bar{K}]$	cascade operation, 335
$\mathsf{Casc}(\mathcal{K})$	cascade languages, 335
a^ω	ω-power, 342
a^π	idempotent power, 342
$\pi \times a$	idempotent multiple, 342
$a \approx_n^m t$	cEF$_n$-bisimulation, 346
$\mathrm{tp}_n^m(s)$	EF$_n$-type of s, 347
A^Δ	ξ-tuples, 362
B^A	set of functions, 362
$t{\downarrow}_v^x$	subforest of x-successors, 362
$t \Leftrightarrow \bar{a}$	substituting values for variables, 362
$t \lhd \varphi$	φ-labelling of t, 363
$\alpha \lhd \beta$	cascade composition, 363
$\mathfrak{A} \circ \mathfrak{B}$	wreath product, 363
pr_w	prefix corresponding to w, 376
$\mathfrak{S}^1$	adjoining a unit to a semigroup, 381
Ψ	set of characters, 382
$\chi(w)$	character of w, 382
$\chi(w/\sigma)$	relative character, 382
sf_v	subforest attached to v, 382
$\Pi_N(s)$	number of paths, 383
$\Pi_N^+(s)$	number of finite paths, 383
$\Pi_N(s/H)$	number of paths through H, 383
$\Pi_N^+(s/H)$	number of finite paths through H, 383
$s \sim_\sigma t$	equivalent w.r.t. a character, 390
$c \approx_\sigma d$	equivalent w.r.t. σ, 393
$[c]_\sigma$	$\approx_\sigma$-class, 393
$\Psi_{\{x\}}^{(\sigma)}$	σ-significant characters, 393
$\mathrm{pf}(s)$	profile of s, 393

Chapter VII

Pw	power-set functor, 405
Lin	linear combination functor, 405
${}^\sigma a$	variable substitution, 433
un	unravelling operation, 433
$\mathbb{T}^\circ$	linear trees with variables in standard ordering, 433
ι	inclusion $\mathbb{T}^\circ \Rightarrow \mathbb{T}^\times$, 433
$\mathbb{X}$	monad for variable substitutions, 434
comp	multiplication of $\mathbb{X}$, 434
in	unit of $\mathbb{X}$, 434
re	reconstitution operation, 434
re_0	reconstitution operation, 434
re	reconstitution operation, 439
un	unravelling operation, 439
$\varphi_{/v}$	action on the variables at vertex v, 441
$\varphi, \sigma : s\ \theta^{\mathrm{sel}}\ t$	unravelling relation, 441
$\mathrm{sun}(t)$	strong unravelling, 442
$\mathrm{dun}(t)$	deep unravelling, 442
$\mathrm{sel}(t)$	selection, 447
σ_{io}	inside-out morphism, 463
σ_{oi}	outside-in morphism, 463
$a \cdot_x b$	substitution, 468
${}^\sigma a$	variable renaming, 468
$a + b$	union, 468
$(a_0 + \cdots + a_{n-1})^{\omega x}$	ω-power, 469
$(a_0 + \cdots + a_{n-1})^{+x}$	iteration, 469

$[\![R]\!]$	value of a regular expression, 469
$\mathcal{R}$	category of relations, 472
skolem	distributive law for Dn and Up, 478

Chapter VIII

$\mathbb{Y}$	omit variables, 481
$\mathrm{supp}(a)$	support, 481
$\mathrm{sing}^?$	singleton map for $\mathbb{T}^?$, 482
$\mathrm{flat}^?$	flattening map for $\mathbb{T}^?$, 482
${}^{\sigma}a$	variable replacement, 482
cat	concatenation, 486
$\langle - \rangle$	singleton word, 486
$\mathrm{SG}(\mathfrak{A})$	ω-semigroup associated with $\mathfrak{A}$, 487
$\mathrm{TA}(\mathfrak{S})$	$\mathbb{T}^?$-algebra generated by $\mathfrak{S}$, 488
$\mathfrak{S}_{\mathcal{A}}$	ω-semigroup associated with the automaton $\mathcal{A}$, 497
$\mathfrak{T}(\mathcal{A})$	transition algebra, 498
$\mathfrak{R}(\mathfrak{A}, \mathcal{A})$	relational morphism for an automaton, 499
$\mathrm{GG}(\mathfrak{S})$	set of game graphs, 503
$\mathrm{Game}(\mathfrak{S})$	set of game abstractions, 504

Index

accepting run, 311
algebraic signature, 38
alphabet, 169
φ-annotation, 363
anti-chain logic, 291
antiregular tree, 304
aperiodic semigroup, 381
arctic semiring, 159
aritiy of a polynomial functor, 24
arity, 279
arity of a function symbol, 38
arity of a signature, 38
automaton, 310
Automaton-Pathfinder Game, 312
axiomatising a class of algebras, 237

back-and-forth argument, 287
$\mathbb{M}$,$\mathbb{P}$-bialgebra, 129
δ-bialgebra, 129
bisimilar forests, 337
bisimulation, 337
bisimulation-invariance, 337
Bojańczyk-Klin algebra, 307
bounded profile, 470
branch-continuous algebra, 508

canonical subobject diagram, 113
cascade composition, 363
category of relations, 472
category with congruences, 143
character of a word, 382
close end, 351
closure of morphisms pullback, 106
closure of morphisms pushouts, 106
closure of morphisms under (co-)limits, 106
cogenerator, 164
cokernel pair, 94
comma category, 10
complete binary tree, 328
completely prime element, 492
$\mathbb{M}$-compositional logic, 251
composition of contexts, 185
computation tree logic, 331
conforming to a strategy, 503
congruence, 142
$\mathbb{M}$-congruence, 148
congruence ordering, 146
context, 184
convex set, 351
counting monadic second-order logic, 286
counting $\mathcal{K}$-temporal logic, 332
counting temporal logics, 331
covariant power-set functor, 405

deep unravelling, 442
definable algebra, 257
definable class, 245
definable language, 250
definably embedded subset, 257
dense morphism of monads, 198
densely antiregular tree, 304
derivative of a language, 185
descending chain condition, 162
deterministic automaton, 498
deterministic rooted graphs, 280
diagonal fill-in property, 100
discrete category, 12
disjunctive category, 12
δ-distributive, 137
distributive $\mathbb{F}$-algebra, 376
distributive law, 74

x-edge, 280
EF-bisimulation, 346
U-end, 351
enough injectives, 160
L-equivalence, 244
essentially finitary, 314
essentially finitary monad, 198
essentially unique factorisation, 114
existence of syntactic algebras, 187
expansion, 126
Ω-extension of a logic, 245
extension of a monad, 77
extension system, 42

factor of a term, 41
factorisation, 41
EM-factorisation, 99
factorisation system, 99
factorising through a cocone, 114
family of languages, 169
family of logics, 250
far end, 351
finitary algebra, 161
finitary functor, 180
finitary signature, 38
finite-dimensional, 161
finitely branching forest, 333
first-order logic, 285
flattening, 36
flattening of a graph, 281
forest, 280
forest automaton, 310
forest formula, 332
formula, 243
free algebra, 48
free functor, 77
free unravelling structure, 439
function symbol, 38

$\mathfrak{G}$-set, 9
game graph, 502
generalised power-set monad, 412
generalised sum, 288
$\mathcal{P}$-generated object, 154
κ-M-generated object, 114
κ-M-generating diagram, 114
generator, 151
graph homomorphism, 441
graph of a function, 473
graph structures, 285
Green's relations, 298

H-class, 298
half a distributive law, 75
height of a factorisation, 320
hole, 184, 280

horizontal product, 289
horizontal semigroup of a forest algebra, 289

idempotent power, 342
image of a morphism, 100
induced product, 137
$\mathbb{M}$-inequality, 237
injective object, 160
inside-out, 463
σ-insignificant element, 390
internal operation, 185
interpretation, 288
inverse image, 169, 171
inverse morphism, 170
inverse relabelling, 170
iteration, 469
iteration of a language, 297

J-class, 298
J-homogeneous word, 298
join-distributive algebra, 491
join-distributive morphism, 491
join-generator, 491

kernel of a function in Pos, 146
kernel of a function in Set, 145
kernel of a morphism, 100
kernel pair, 94
Kleisli category, 77

L-class, 298
Φ-labelling, 332
Ω-language, 169
language algebra, 185
language recognised by an automaton, 311

lattice-distributive, 491
Läuchli-Leonard operation, 293
lift of a mond, 77
lift of a relation, 41
linear forest, 280
linear monad, 416
linear tree, 280
locally definable algebra, 257
logic, 243
logic over an alphabet, 250

matrix power, 506
meet-distributive algebra, 491
meet-distributive morphism, 491
meet-generator, 491
minimal algebra of a language, 175
minimal morphism of a language, 175
modal logic, 331
modal rank, 333
model, 243
models of a formula, 243
monadic chain logic, 331
monadic path logic, 331
monadic second-order logic, 285
monadic second-order logic for linear orders, 292
monadic second-order logic with forest ordering, 286
monadic second-order logic with successor, 286
monadic second-order logic with successor ordering, 286
morphism of $\mathbb{M}$,$\mathbb{P}$-bialgebras, 129
morphism of language algebras, 186
morphism of logics, 246
morphism of monads, 65
morphism of polynomial functors, 28

morphism of unravelling structures, 439

neutral element, 381
noetherian, 162
non-linear forest, 280
non-linear tree, 280
nondeterministic rooted graph, 279

ω-semigroup, 486
Ω-operation, 160
orbit finite set, 167
orbit-finite, 9
origin map, 32
outcome map, 504
outcome of a play, 503
outside-in, 463

parity condition, 311
parity game, 504
partial run of an automaton, 313
$+$-path algebra, 383
$\mathcal{S}$-path language, 381
ω-path algebra, 383
permutation of a forest, 336
place in a forest, 312
play of a game, 503
polynomial functor, 24
ω-power, 469
power operator, 21
power-set functor, 405
x-predecessor, 280
preserving a limit/colimit, 116
priority function, 310
σ-profile of a forest, 393
profile of a partial run, 313
projective object, 150
propositional dynamic logic, 331
pseudo-variety, 155

quasi-finite set, 167
quotient, 142

R-class, 298
ranked tree, 280
recognised by a relational morphism, 499
recognising a language, 169
recognising a language by an automaton, 311
reconstitution operation, 434
reduced path, 382
reduced tree, 320
Δ-reduct, 126
reduct of a monad, 65
reduct of an $\mathbb{M}$-algebra, 122
reflecting a limit/colimit, 117
regular expression for languages of countable chains, 297
regular expressions for trees, 468
regular forest, 313
regular linear order, 294
regular substitution, 463
relabelling, 170
relational morphism, 499
relative character, 382
representing a function by a relation, 473
representing a profile, 498
root, 279
rooted graph, 279
run of an automaton, 311

satisfaction function, 243
satisfying an $\mathbb{M}$-inequality, 237
selection operation, 447

ω-semigroup, 486
ω-semigroup associated with a $\mathbb{T}^?$-algebra, 487
semigroup-like algebra, 485
σ-separated forest, 390
shape map, 32
shape of a term, 41
shuffle, 293
sibling-commutative forest language, 336
Sierpiński space, 166
signature, 38
σ-significant element, 390
simple interpretation, 288
singleton operation, 36
skeleton of an algebra, 507
sort, 38
sort-accumulation point, 206
sort-wise finite family of logics, 250
sort-wise properties, 161
standard factorisation system, 109
starting state of a partial run, 313
state of an automaton, 310
strategy, 503
strong morphism of logics, 246
strong unravelling, 442
strongly $\mathbb{M}$-compositional logic, 273
strongly definable class, 245
strongly finite-dimensional, 161
Σ-structure, 45
subalgebra, 110
sublinear tree, 482
subobject, 110
substitution, 463
substitution operation, 468
x-successor, 280
support of an element, 481
supporting unravelling, 456
syntactic algebra, 187
syntactic congruence, 187
syntactic morphism, 187
syntactic partition, 332
syntactic residue, 187

$\mathcal{K}$-temporal logic, 333
term, 24
term algebra, 49
theory, 243
$\mathbb{M}$-theory of a class, 237
theory function, 243
trace algebra, 493
trace of a tree, 495
transition algebra of an automaton, 498
transition relation, 310
transitive closure, 177
tree, 280
tree formula, 332
ζ-trivial tree, 468
tropical semiring, 159
type of a function symbol, 38

universe, 45
unranked tree, 280
unravelling, 281, 439
unravelling structure, 439

Ω-valued logic, 243
variable, 279
varietal logic, 250
variety, 155
$\mathcal{R}$-variety, 206
variety of languages, 206
vertical product, 289
vertical semigroup of a forest algebra, 289
Vietoris monad, 411

weak coequaliser, 57
weak morphism of logics, 246
weakly definable class, 245
weakly finite-dimensional set, 161
well-defined factorisation system, 109
word functor, 486
wreath product, 363

The Roman and Fraktur alphabets							
A	a	$\mathfrak{A}$	$\mathfrak{a}$	N	n	$\mathfrak{N}$	$\mathfrak{n}$
B	b	$\mathfrak{B}$	$\mathfrak{b}$	O	o	$\mathfrak{O}$	$\mathfrak{o}$
C	c	$\mathfrak{C}$	$\mathfrak{c}$	P	p	$\mathfrak{P}$	$\mathfrak{p}$
D	d	$\mathfrak{D}$	$\mathfrak{d}$	Q	q	$\mathfrak{Q}$	$\mathfrak{q}$
E	e	$\mathfrak{E}$	$\mathfrak{e}$	R	r	$\mathfrak{R}$	$\mathfrak{r}$
F	f	$\mathfrak{F}$	$\mathfrak{f}$	S	s	$\mathfrak{S}$	ſ $\mathfrak{s}$
G	g	$\mathfrak{G}$	$\mathfrak{g}$	T	t	$\mathfrak{T}$	$\mathfrak{t}$
H	h	$\mathfrak{H}$	$\mathfrak{h}$	U	u	$\mathfrak{U}$	$\mathfrak{u}$
I	i	$\mathfrak{I}$	$\mathfrak{i}$	V	v	$\mathfrak{V}$	$\mathfrak{v}$
J	j	$\mathfrak{J}$	$\mathfrak{j}$	W	w	$\mathfrak{W}$	$\mathfrak{w}$
K	k	$\mathfrak{K}$	$\mathfrak{k}$	X	x	$\mathfrak{X}$	$\mathfrak{x}$
L	l	$\mathfrak{L}$	$\mathfrak{l}$	Y	y	$\mathfrak{Y}$	$\mathfrak{y}$
M	m	$\mathfrak{M}$	$\mathfrak{m}$	Z	z	$\mathfrak{Z}$	$\mathfrak{z}$

The Greek alphabet					
A	α	alpha	N	ν	nu
B	β	beta	Ξ	ξ	xi
Γ	γ	gamma	O	o	omicron
Δ	δ	delta	Π	π	pi
E	ε	epsilon	P	ρ	rho
Z	ζ	zeta	Σ	σ	sigma
H	η	eta	T	τ	tau
Θ	ϑ	theta	Υ	υ	upsilon
I	ι	iota	Φ	ϕ	phi
K	κ	kappa	X	χ	chi
Λ	λ	lambda	Ψ	ψ	psi
M	μ	mu	Ω	ω	omega