An Algebraic Proof of Rabin's Tree Theorem

Achim Blumensath*

13th February 2013

We continue the development of a theory of recognisability for infinite trees by introducing the equivalent of a Wilke algebra. As an application we give a new proof of the decidability of the monadic second-order theory of the infinite binary tree, a proof that does not use automata or games.

Keywords. monadic second-order logic; infinite trees; recognisability

1 Introduction

We continue to develop the theory of recognisability for languages of infinite trees that was introduced in [1]. The main result of [1] is a characterisation of the regular languages of infinite trees via homomorphisms into suitable algebras. Historically, one of the main advantages of such algebraic characterisations has been their suitability for deriving decision procedures for subclasses of regular languages. We hope that the theory developed in [1] will also be useful in this respect. One of the prerequisites for obtaining decision procedures is that the algebras one is dealing with are finitely representable. The main result of the present article states that the algebras introduced in [1], called *path-continuous ω-hyperclones,* are, in fact, finitely representable.

*Work partially supported by DFG grant BL 1127/2-1.

As an application of this result, we present a new proof of Rabin's Tree Theorem, which states that the monadic second-order theory of the infinite complete binary tree is decidable. The standard proof of this result is based on the translation of monadic second-order formulae into tree automata. The required automata-theoretic machinery in turn rests on two main results: (i) the positional determinacy of parity games and (ii) either a determinisation construction for Büchi automata, or an analogous result for tree automata (see [3, 4]).

For Büchi's Theorem – the corresponding result for the natural numbers with successor relation – there exists, besides the usual automata-theoretic proof, an alternative proof due to Shelah [5, 6]. It is purely combinatorial in nature and it is based on Feferman-Vaught like composition arguments for monadic second-order theories.

For a long time it has been an open problem to extend these results to the theory of the binary tree. What was missing in order to transfer Shelah's proof was a suitable variant of Ramsey's Theorem for trees. Such a theorem has recently been provided by Colcombet [2]. In this article we use Colcombet's result to prove that our algebras have finite representations. This fact is then used to give an alternative proof of Rabin's theorem without references to automata or games.

The outline of the article is as follows. After recalling the definition of an ω-hyperclone in Section 2, we develop the combinatorial machinery needed to obtain finite representations of ω-hyperclones in Sections 3 and 4. Section 5 contains a proof that ω-hyperclones can be represented by so-called *power-hyperclones* which are finitely representable. As an application, we use this result to give an alternative proof of Rabin's Theorem in Section 6.

2 Preliminaries

To fix notation, let $[n] := \{0, \dots, n-1\}$. The domain of a function $f : A \to B$ is $\mathrm{dom}(f) = A$ and its range is $\mathrm{rng}(f) \subseteq B$. We tacitly identify a tuple $\bar{a} = \langle a_0, \dots, a_{n-1} \rangle$ with the function $i \mapsto a_i$ and with the set $\{a_0, \dots, a_{n-1}\}$ of its components. We denote by $\omega^{<\omega}$ the set of all finite sequences of natural numbers. The empty sequence is $\langle\rangle$. We denote the prefix ordering on $\omega^{<\omega}$ by $\leq$:

$$x \leq y \quad \text{iff} \quad y = xu \text{ for some } u \in \omega^{<\omega}.$$

We denote structures with fraktur letters $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}, \dots$ and their universes by the corresponding Roman letters $A, B, C, \dots$. We will also deal with many-sorted structures. In the many-sorted case, we write $A_s, B_s, C_s, \dots$ for the domain of sort s.

In this article, graphs will always be simple, directed, and labelled. More precisely, a graph is a structure of the form $\mathfrak{G} = \langle V, (E_c)_{c \in A}, (P_i)_{i \in \Sigma}, v_0 \rangle$ where the E_c are binary edge relations, the P_i are unary predicates, and v_0 is a distinguished vertex. We assume that the edge relations E_c are *deterministic*, that is, every vertex in V has at most one outgoing E_c-edge, for every c. We will write E for the union $\bigcup_c E_c$ of all edge relations.

A *path* of a graph $\mathfrak{G}$ is a (finite or infinite) sequence $(e_n)_n$ of edges such that, for all n, the second vertex of e_n coincides with the first vertex of e_{n+1}. We write $w : x \to y$ if w is a path from the vertex x to the vertex y. Similarly, we write $\beta : x \to \infty$ if β is an infinite path starting at the vertex x.

Definition 2.1. (a) A *tree domain* is a subset $T \subseteq \omega^{<\omega}$ that is prefix-closed and finitely branching, i.e.,

- $u \leq v \in T$ implies $u \in T$ and
- for every $w \in T$, there are only finitely many $d < \omega$ with $wd \in T$.

(b) A *tree* is a graph $\mathfrak{T} = \langle T, (E_d)_{d<\omega} \rangle$ where T is a tree domain and $E_d := \{ \langle w, wd \rangle \mid wd \in T \}$ is the d-th *successor relation.* The *predecessor* of a vertex $w \in T$ is the vertex $w' \in T$ such that $w = w'd$, for some $d < \omega$. We denote it by $\operatorname{prec}(w)$. Note that $\operatorname{prec}(\langle\rangle)$ is undefined.

Definition 2.2. (a) An *ω-semigroup* is a tuple $\langle S, S_\omega, \alpha, \pi \rangle$ where S is a semigroup, S_ω is a set, $\alpha : S \times S_\omega \to S_\omega$ is a (left-)action of S on S_ω, and $\pi : S^\omega \to S_\omega$ is a function satisfying the following two associative laws: for every sequence $(s_n)_n \in S^\omega$ and all strictly increasing sequences $0 = k_0 < k_1 < \dots$ of natural numbers, we have

$$\pi(s_0, s_1, s_2, \dots) = \alpha(s_0, \pi(s_1, s_2, \dots)),$$
$$\pi(s_0, s_1, s_2, \dots) = \pi\big((s_{k_0} \cdots s_{k_1-1}), (s_{k_1} \cdots s_{k_2-1}), (s_{k_2} \cdots s_{k_3-1}), \dots\big).$$

Usually, we omit α from the notation and simply write su instead of $\alpha(s,u)$. Also we write $\prod_n s_n$ instead of $\pi(s_0, s_1, \dots)$.

(b) An *homomorphism* between ω-semigroups $\langle S, S_\omega, \alpha, \pi \rangle$ and $\langle T, T_\omega, \beta, \rho \rangle$ consists of two functions $h : S \to T$ and $h_\omega : S_\omega \to T_\omega$ such that

- h is a homomorphism of semigroups,

- $\beta(hs, h_\omega u) = h_\omega(\alpha(s,u))$, for all $s \in S$ and $u \in S_\omega$, and
- $h_\omega(\pi(s_0, s_1, s_2, \dots)) = \rho(hs_0, hs_1, hs_2, \dots)$, for all $(s_n)_n \in S^\omega$.

In [1] we have introduced certain algebras called *ω-hyperclones* that can be used to give an algebraic characterisation of regular languages of infinite trees. The main result of that article is a theorem (see Theorem 2.8 below) stating that a set of (possibly infinite) terms is regular if, and only if, it is recognised by a morphism into a finitary ω-hyperclone with certain additional properties.

Intuitively, an ω-hyperclone is an algebra where each element can be thought of as a tuple of objects each of which has a number of *ports*. Each port is labelled by a natural number. For instance, the objects could be terms where each occurrence of a variable x_i corresponds to a port with label i. In particular, the ports are arranged in a left-to-right fashion and there may be several ports with the same label. For simplicity, we assume that each object has only finitely many ports. Hence, to each object we can associate a finite tuple of natural numbers. We can depict an element of an ω-hyperclone consisting of four objects with ports $\langle 0, 2 \rangle$, $\langle 1, 0, 1 \rangle$, $\langle 2 \rangle$, and $\langle 2, 0 \rangle$, respectively, as in the following diagram.

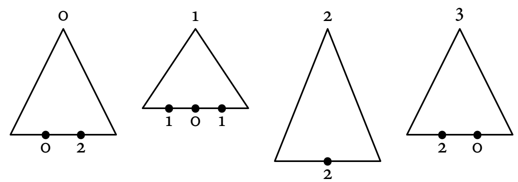

An ω-hyperclone has three main operations: a finite horizontal product $\oplus$, a finite vertical product $\cdot$, and an infinite vertical product π. The horizontal product $\oplus$ is just the concatenation of tuples.

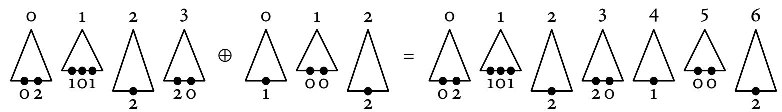

For the vertical product $\cdot$ we plug in the i-th object of the second tuple into every port of the first tuple with label i. For instance, if the objects are terms then the vertical product might correspond in substituting the i-th term of the second tuple for the variable x_i in each term of the first tuple.

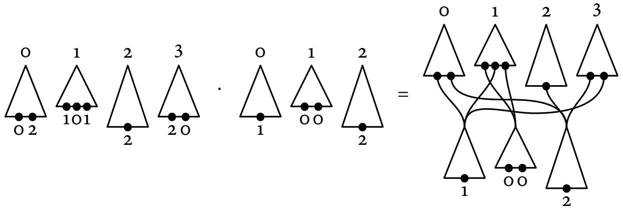

For technical reasons, we will not use the basic vertical composition $\cdot$, but a more complicated operation $:_{I,\tau}$ where we can plug in objects into only some of the ports, while changing the numbers of the remaining ports.

An important example of an ω-hyperclone is the *free ω-hyperclone* $\mathfrak{F}_\omega[\Sigma]$ where the elements are tuples of terms.

Definition 2.3. (a) Let Σ be a signature and $X = \{x_0, x_1, \dots\}$ a countable set of variables. We denote by $T_\omega[\Sigma]$ the set of all (possibly infinite) terms with operations from Σ and variables from X with only finitely many occurrences of variables.

For a term $t \in T_\omega[\Sigma]$, we denote by $\mathrm{var}(t) \in \omega^{<\omega}$ the sequence of all (indices of) variables appearing in t in left-to-right order. Formally, we define var as the unique function $T_\omega[\Sigma] \to \omega^{<\omega}$ satisfying the following equations:

$$\begin{aligned} \mathrm{var}(t) &= \langle\rangle\,, \quad \text{if } t \text{ does not contain a variable,}\\ \mathrm{var}(x_i) &= \langle i\rangle\,,\\ \text{and}\quad \mathrm{var}(f(t_0,\dots,t_{n-1})) &= \mathrm{var}(t_0)\dots\mathrm{var}(t_{n-1})\,. \end{aligned}$$

(b) The *free ω-hyperclone* over Σ is the many-sorted structure

$$\mathfrak{F}_\omega[\Sigma] := \big\langle (F_{\bar u}[\Sigma])_{\bar u \in (\omega^{<\omega})^{<\omega}}, \oplus, \mathrm{o}, (\lambda_\sigma)_\sigma, (:_{I,\sigma})_{I,\sigma}, \pi, \leq \big\rangle$$

with set of sorts $(\omega^{<\omega})^{<\omega}$ where the domain of sort $\bar u = \langle u_0, \dots, u_{n-1}\rangle \in (\omega^{<\omega})^{<\omega}$ is the set

$$F_{\bar u}[\Sigma] := \big\{ \langle t_0,\dots,t_{n-1}\rangle \bigm| t_i \in T_\omega[\Sigma] \smallsetminus \{x_0, x_1, \dots\} \text{ with } \mathrm{var}(t_i) = u_i \big\}$$

of all finite tuples of non-trivial terms such that the i-th term has variables u_i. It has the following operations:

- $\oplus$ is the concatenation of tuples.

- $:_{I,\sigma}$, for $I \subseteq [n]$ and $\sigma : [k] \to \omega$, is defined as

$$\langle s_0,\dots,s_{m-1}\rangle :_{I,\sigma} \langle t_0,\dots,t_{n-1}\rangle := \langle u_0,\dots,u_{m-1}\rangle\,,$$

  where u_l is obtained from s_l by replacing every occurrence of a variable x_i with $i \in I$ by the term t_i. Variables x_i with $i \notin I$ are replaced by the variable $x_{\sigma(i)}$ instead.
- λ_σ reorders its argument according to $\sigma : [m] \to [n]$:

$$\lambda_\sigma\langle t_0,\dots,t_{n-1}\rangle = \langle t_{\sigma(0)},\dots,t_{\sigma(m-1)}\rangle\,.$$

  Note that σ need neither be injective, nor surjective.
- $\pi(a_0, a_1, \dots)$ is the limit of the terms a_0, $(a_0 \cdot a_1)$, $(a_0 \cdot a_1 \cdot a_2)$,…, where, for an m-tuple $\bar s$ and an n-tuple $\bar t$, the simple version of the vertical composition is defined by $\bar s \cdot \bar t := \bar s :_{[n],\mathrm{id}} \bar t$.
- o denotes the empty tuple of terms.
- The order $\leq$ is trivial: $a \leq b$ iff $a = b$.

An arbitrary ω-hyperclone $\mathfrak{C}$ is a homomorphic image of such a free ω-hyperclone with the additional requirement that all operations are monotone with respect to the ordering. (The ordering $\leq$ will be needed for the definition of path-continuity below. It plays no other role in the theory of ω-hyperclones.) We refer the reader to [1] for an axiomatic definition of ω-hyperclones and a proof that the free ω-hyperclone is actually free. Let us recall some definitions from [1].

Definition 2.4. Let $\mathfrak{C}$ be an ω-hyperclone, $a \in C_{\bar u}$ an element of sort $\bar u = \langle u_0,\dots,u_{m-1}\rangle$ and $b \in C_{\bar v}$ an element of sort $\bar v = \langle v_0,\dots,v_{n-1}\rangle$.

(a) We introduce the following abbreviations:

$$\begin{aligned} a :_I b &:= a :_{I,\mathrm{id}} b\,,\\ a \cdot b &:= a :_J b\,, && \text{where } J = u_0 \cup \dots \cup u_{m-1} \text{ is the set of all ports of } a\,,\\ \rho_\sigma(a) &:= a :_{\varnothing,\sigma} c\,, && \text{for an arbitrary element } c\,. \end{aligned}$$

(The axioms of an ω-hyperclone imply that $a :_{\varnothing,\sigma} c$ does not depend on c.) To simplify notation, we further set

$$ab := a \cdot b\,, \quad \sigma a := \lambda_\sigma(a)\,, \quad a\sigma := \rho_\sigma(a)\,,$$

(b) The *decomposition* of a is the tuple $\langle a_0, \dots, a_{m-1}\rangle$ of elements $a_i \in C_{u_i}$ such that $a = a_0 \oplus \dots \oplus a_{m-1}$. Its *width* is the number m. (By the axioms of an ω-hyperclone, this decomposition is unique.)

(c) We say that $a \in C_{\bar{u}}$ is *in separation normal form* if there are numbers $0 = k_0 \leq \dots \leq k_m < \omega$ such that

$$u_i = \langle k_i, k_i + 1, \dots, k_{i+1} - 1\rangle\,, \quad \text{for all } i < m\,.$$

(d) The *normal form* of a is an element b in separation normal form such that $a = \rho_\sigma(b)$, for some σ. We denote the normal form of a by $\operatorname{sep}(a)$.

Example. In the free ω-hyperclone we have

$$\operatorname{sep}\big\langle f(x_0, f(x_2, x_0)),\ f(x_0, x_2)\big\rangle = \big\langle f(x_0, f(x_1, x_2)),\ f(x_3, x_4)\big\rangle\,.$$

Remark. (a) When viewing elements of an ω-hyperclone as objects with ports, a vertical product $a \cdot b \cdot c \cdot d$ produces a directed acyclic graph of objects. The main property of elements in separation normal form is that their product produces a tree instead.

(b) Every element $a \in C_{\bar{u}}$ has a unique normal form. Hence, $\operatorname{sep}(a)$ is well-defined.

Let us introduce some concepts and terminology that are useful when dealing with infinite products.

Definition 2.5. Let $\mathfrak{C}$ be an ω-hyperclone.

(a) We use the notation $a^\square$ for a sequence $(a^n)_{n<\omega}$ that can be multiplied, i.e., the sorts of a^n and a^{n+1} are such that the product $a^n \cdot a^{n+1}$ is defined for every $n < \omega$.

(b) We say that a sequence $a^\square = (a^n)_{n<\omega}$ is *in separation normal form* if every a^n is in separation normal form.

(c) The *unravelling* of an arbitrary sequence $a^\square$ is a sequence $\hat{a}^\square$ in separation normal form that is defined as follows. We simultaneously define functions σ_n and elements $\hat{a}^n$, by induction on n. We start with $\sigma_0 := \mathrm{id}$. If σ_n is already defined, we set

$$\hat{a}^n := \operatorname{sep}(\lambda_{\sigma_n}(a^n))\,,$$

and we choose σ_{n+1} such that

$$\lambda_{\sigma_n}(a^n) = \rho_{\sigma_{n+1}}(\hat{a}^n)\,.$$

Given a sequence $a^\square$, the product can be viewed as a directed acyclic graph consisting of infinitely many levels that correspond to the elements a^n. Formally, if $a^n = a^n_0 \oplus \dots \oplus a^n_{m(n)-1}$ is the decomposition of a^n, we construct the graph whose vertices are pairs of indices $\langle n, i\rangle$ (representing a^n_i), for all $n < \omega$ and $i < m(i)$, where we add an edge from $\langle n, i\rangle$ to $\langle n+1, k\rangle$ if k is a port of a^n_i. The tree-unfolding of this graph is called the *branch tree* of $a^\square$ (which is actually a forest). The sequence $b^\square$ obtained from $a^\square$ by this tree-unfolding operation is the unravelling of $a^\square$.

Example. In the free ω-hyperclone, consider the sequence $(a^n)_{n<\omega}$ where $a^n = \langle f(x_0, x_1),\ g(x_0)\rangle$, for every n. The first levels of the branch tree of this sequence are

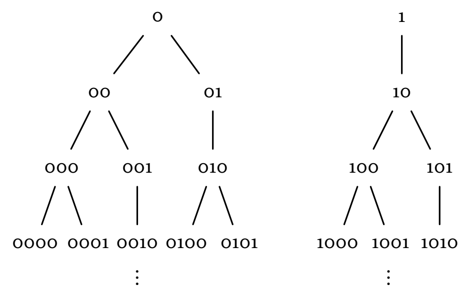

There is a canonical function from the branch tree of $a^\square$ to the original sequence which we denote by μ. This function can be extended to map finite, connected subsets of the branch tree to the products of the corresponding elements of $a^\square$. The precise definitions are as follows.

Definition 2.6. Let $\mathfrak{C}$ be an ω-hyperclone.

(a) Let $a^\square = (a^n)_{n<\omega}$ be a sequence in $\mathfrak{C}$ and $a^n = a^n_0 \oplus \dots \oplus a^n_{m(n)-1}$ the decomposition of $a^n \in C_{\bar{u}^n}$ into elements $a^n_i \in C_{u^n_i}$. The *branch tree* of $a^\square$ is the forest

$$\begin{aligned}\Lambda(a^\square) := \big\{\, \eta \in \omega^{<\omega} \;\big|\; |\eta| > 0,\ \eta(0) < m(0),\ \text{and} \\ \eta(n+1) \in u^n_{\eta(n)},\ \text{for all } n \,\big\}\,.\end{aligned}$$

(Here we write $x \in u$ for a tuple u to state that x is one of the components of u.) A *branch* of $a^\square$ is a sequence $\beta \in \omega^{\leq\omega}$ of maximal length such that every finite prefix $\eta \preceq \beta$ belongs to $\Lambda(a^\square)$.

(b) We define a function $\mu : \Lambda(a^\square) \to C$ by

$$\mu(\eta) := a^{n-1}_{\eta(n-1)}\,, \quad \text{where } n := |\eta|\,.$$

A subset $U \subseteq \Lambda(a^\square)$ is *connected* if there exists an element $\eta \in U$ such that

$$\zeta \in U \quad \text{implies} \quad \eta \preceq \zeta \text{ and } \xi \in U \text{ for all } \eta \preceq \xi \preceq \zeta\,.$$

We extend μ to finite connected subsets $U \subseteq \Lambda(a^\square)$ as follows. We define an element $\mu[U] \in C$ by induction on $|U|$. For $U = \{\eta\}$, we set

$$\mu[\{\eta\}] := \operatorname{sep}(\mu(\eta))\,.$$

For $|U| > 1$, let η be the minimal element of U and let $\zeta_0, \dots, \zeta_{m-1}$ be its successors in $\Lambda(a^\square)$. We define

$$\mu[U] := \operatorname{sep}(b :_I c)\,.$$

where

$$\begin{aligned} I &:= \{\, i < m \mid \zeta_i \in U \,\}\,, \\ b &:= \mu[\{\eta\}]\,, \\ c &:= \mu[U_0] \oplus \dots \oplus \mu[U_{m-1}] \quad \text{for} \quad U_i := \{\zeta_i\} \cup \{\, \xi \in U \mid \xi \succeq \zeta_i \,\}\,. \end{aligned}$$

Remark. Let $\hat{a}^\square$ be the unravelling of $a^\square$. If $\eta_0, \dots, \eta_m$ is an enumeration, from left to right, of all vertices of level $n+1$ of the branch tree $\Lambda(a^\square)$, then

$$\hat{a}^n = \operatorname{sep}(\mu(\eta_0) \oplus \dots \oplus \mu(\eta_m))\,.$$

We are mainly interested in ω-hyperclones that are *finitary* and *path-continuous.*

Definition 2.7. Let $\mathfrak{C}$ be an ordered ω-hyperclone and J a set of elements of $\mathfrak{C}$.

(a) $\mathfrak{C}$ is *finitary* if it has only finitely many elements of each sort $\bar{u}$.

(b) A sequence $(b^n)_{n<\omega}$ where $b^n \in C_{\langle\varnothing,\dots,\varnothing\rangle}$ is *locally consistent* with $a^\square$ if $b_n = a_n b_{n+1}$, for every $n < \omega$. We denote by $\mathrm{LC}_J(a^\square)$ the set of all locally consistent sequences $(b^n)_{n<\omega}$ with $b^n \in J$, for all n.

(c) Let $a^\square$ be a sequence, $a^n = a^n_0 \oplus \dots \oplus a^n_{m(n)-1}$ the decomposition of a^n, suppose that $(b^n)_n \in \mathrm{LC}_J(a^\square)$ is locally consistent with $a^\square$, and let β be a branch of $a^\square$. The *trace* of $(b^n)_n$ along β is the sequence $c^\square$ with

$$c^n := \operatorname{sep}(a^n_{\beta(n)} :_{I_n} b^{n+1})\,,$$

where

$$I_n := \begin{cases} [m(n+1)] \setminus \{\beta(n+1)\} & \text{if } n+1 \in \operatorname{dom}(\beta)\,, \\ [m(n+1)] & \text{otherwise}\,. \end{cases}$$

We set

$$\begin{aligned} \mathrm{Tr}_J(\beta) &:= \{\, \pi(c^\square) \mid c^\square \text{ the trace of some } (b^n)_n \in \mathrm{LC}_J(a^\square) \text{ along } \beta \,\}\,, \\ \mathrm{BT}_J(a^\square) &:= \big\{\, \mathrm{Tr}_J(\beta) \bigm| \beta \text{ a branch of } \Lambda(a^\square) \,\big\}\,. \end{aligned}$$

(d) An *ideal* of $\mathfrak{C}$ is a sub-ω-hyperclone $\mathfrak{J} \subseteq \mathfrak{C}$ where each domain $J_{\bar{u}}$ is downward closed, i.e., $\mathfrak{J}$ is a substructure of $\mathfrak{C}$ that is an ω-hyperclone such that $a \leq b \in J_{\bar{u}}$ implies $a \in J_{\bar{u}}$.

(e) $\mathfrak{C}$ is *path-continuous* if there exists an ideal $\mathfrak{J}$ such that

- for every sequence $a^\square$ in separation normal form, we have

$$\pi(a^\square) = \sup\big\{\, \pi(b^\square) \bigm| b^\square \leq a^\square \text{ a sequence in } J \,\big\}\,,$$

- and the product $\pi(a^\square)$ of a sequence $a^\square$ in J is uniquely determined by the set $\mathrm{BT}_J(a^\square)$.

With these definitions we can state the main theorem of [1].

Theorem 2.8. *Let L be a set of (possibly infinite) terms over the signature Σ. The set L is regular if, and only if, there exists a morphism $\varphi : \mathfrak{F}_\omega[\Sigma] \to \mathfrak{C}$ into a finitary path-continuous ω-hyperclone $\mathfrak{C}$ and a subset $P \subseteq C_{\langle\rangle}$ such that $L = \varphi^{-1}[P]$.*

In order to apply this theorem to decision problems, we would like to be able to compute the ω-hyperclone $\mathfrak{C}$ from a description of L. To do so, we need a finite representation of ω-hyperclones. Even for a finitary ω-hyperclone $\mathfrak{C}$, there are still infinitely many sorts and the description of the infinite product $\pi : C^\omega \to C$ can be infinite. The first problem is easily solved: instead of the whole ω-hyperclone, it will turn out to be sufficient to compute its restriction to finitely many sorts. For the second problem, we use the same solution as in the theory of ω-semigroups: we replace the infinite product $\pi : C^\omega \to C$ by a unary operation ${}^\omega : C \to C$ computing ω-th powers. The main part of this article consists in the proof that this power-operation uniquely determines the infinite product.

3 Additive labellings

In this section we introduce the combinatorial machinery that will be used in the remainder of the article to represent an infinite product π by its associated power operation $^\omega$. The main result is the statement that, for every (infinite) tree, we can find a regular tree that is equivalent to the first one in a sense made precise below.

Note that one well-known consequence of Rabin's Theorem is the fact that every nonempty MSO-definable class of infinite trees contains a regular tree. Our alternative proof of this theorem follows the opposite direction. We first prove the existence of certain regular labellings of the binary tree. Then we derive Rabin's Theorem from this result.

Definition 3.1. Let $\langle S, S_\omega, \alpha, \pi \rangle$ be a finite ω-semigroup and $\mathfrak{G} = \langle V, \bar{E}, \bar{P}, v_0 \rangle$ a graph.

(a) An *additive labelling* of $\mathfrak{G}$ is a function $\lambda : E \to S$. Such a labelling can be extended to a function mapping every finite or infinite path of $\mathfrak{G}$ to an element of $S \cup S_\omega$:

$$\lambda((e_n)_n) := \prod_n \lambda(e_n)\,, \quad \text{for every nonempty path } (e_n)_n\,.$$

(b) Let λ be an additive labelling of $\mathfrak{G}$. The *limit set* of λ is

$$\lim \lambda := \{\, \lambda(\alpha) \mid \alpha : v_0 \to \infty \,\}\,.$$

(c) For a tree $\mathfrak{G} = \mathfrak{T}$ and vertices $u < v$ in T, we also write $\lambda(u, v)$ for the value of $\lambda(w)$ where w is the unique path $w : u \to v$.

(d) Two additive labellings λ and λ', not necessarily of the same graph, are *equivalent* if $\lim \lambda = \lim \lambda'$.

We will show that, for every additive labelling of the binary tree, we can find an equivalent additive labelling that is regular. The main combinatorial tool we will use is based on the following Ramsey-like factorisations of labelled trees.

Definition 3.2. Let $\mathfrak{T} = \langle T, \bar{E} \rangle$ be a tree, $\sigma : T \to [n]$ a function, and let λ be an additive labelling of $\mathfrak{T}$.

(a) For $u, v \in T$, we define $u \sqsubseteq_\sigma v$ if

- $u \leq v$,
- $\sigma(u) = \sigma(v)$,
- $\sigma(w) \geq \sigma(v)$, for all $w \in T$ with $u \leq w \leq v$.

(b) The function σ is a *forward Ramseyan split* of λ if we have

$$\lambda(u, v) = \lambda(u, v) \cdot \lambda(x, y)\,, \qquad \text{for all } u \sqsubset_\sigma v \text{ and } u \sqsubseteq_\sigma x \sqsubset_\sigma y \text{ such that } v \sqsubseteq_\sigma y \text{ or } y \sqsubseteq_\sigma v\,.$$

Theorem 3.3 (Colcombet [2]). *Every additive labelling λ of a tree $\mathfrak{T} = \langle T, \bar{E} \rangle$ has a forward Ramseyan split $\sigma : T \to [n]$ where $n := |S|$.*

To make use of the fact that a function σ is a forward Ramseyan split, we need positions that are related by the relation $\sqsubseteq_\sigma$. The next lemma shows that every sufficiently long path contains such positions.

Lemma 3.4. *Let $\mathfrak{T} = \langle T, \bar{E} \rangle$ be a tree, $\sigma : T \to [n]$ a function, and $u, v \in T$ vertices with $u \leq v$. If $|v| - |u| + 1 \geq m^n$ then there exist vertices $u \leq x_0 < \cdots < x_{m-1} \leq v$ such that $x_0 \sqsubset_\sigma \cdots \sqsubset_\sigma x_{m-1}$.*

Proof. We prove the claim by induction on n. If $n = 1$ and $|v| - |u| + 1 \geq m$ then we can choose vertices $u = x_0 < \cdots < x_{m-1} = v$. Since $\sigma(x_i) = 1$, for all i, it follows that $x_0 \sqsubset_\sigma \cdots \sqsubset_\sigma x_{m-1}$.

Suppose that $n > 1$ and $|v| - |u| + 1 \geq m^n$. Let $x_0 < \cdots < x_{l-1}$ be an enumeration of all vertices $u \leq x \leq v$ with $\sigma(x) = 1$. We distinguish four cases.

(1) If $l \geq m$ then we are done.

(2) If there is some index i such that

$$(|x_{i+1}| - 1) - (|x_i| + 1) + 1 \geq m^{n-1}$$

then, by induction hypothesis, we can find a sequence $x_i < y_0 < \cdots < y_{m-1} < x_{i+1}$ such that $y_0 \sqsubset_\sigma \cdots \sqsubset_\sigma y_{m-1}$.

(3) If

$$(|x_0| - 1) - |u| + 1 \geq m^{n-1} \quad \text{or} \quad |v| - (|x_{l-1}| + 1) + 1 \geq m^{n-1},$$

we can obtain a sequence

$$u \leq y_0 \sqsubset_\sigma \cdots \sqsubset_\sigma y_{m-1} < x_0 \quad \text{or} \quad x_{l-1} < y_0 \sqsubset_\sigma \cdots \sqsubset_\sigma y_{m-1} \leq v$$

as in (2).

(4) Hence, we may assume that $l < m$,

$$|x_{i+1}| - |x_i| < m^{n-1} + 1\,,$$
$$|x_0| - |u| < m^{n-1},$$
and $$|v| - |x_{l-1}| < m^{n-1}.$$

It follows that

$$\begin{aligned}|v| - |u| + 1 &= |v| - |x_{l-1}| + \sum_{i=0}^{l-2}\big(|x_{i+1}| - |x_i|\big) + |x_0| - |u| + 1 \\ &\leq m^{n-1} - 1 + \sum_{i=0}^{l-2} m^{n-1} + m^{n-1} - 1 + 1 = (l+1)m^{n-1} - 1 \\ &\leq m^n - 1\,.\end{aligned}$$

A contradiction □

We will represent regular trees and their labellings by finite trees with *back-edges.*

Definition 3.5. (a) A *pseudo-tree* is a graph $\mathfrak{T} = \langle T, (E_d)_{d<\omega}, \langle\rangle\rangle$ where

- $T \subseteq \omega^{<\omega}$ is a tree domain,
- every vertex $v \in T$ has at most one E_d-successor, for every $d < \omega$,
- if $vd \in T$, for $d < \omega$, then $\langle v, vd\rangle \in E_d$,
- if $\langle v, w\rangle \in E_d$, for $d < \omega$, then either $w = vd$, or $w \leq v$ and $vd \notin T$. Edges of the latter form are called *back-edges.*

(b) A *branch* of $\mathfrak{T}$ is an infinite path $\beta : \langle\rangle \to \infty$.
(c) The *height* $\mathrm{ht}(\mathfrak{T})$ of a pseudo-tree $\mathfrak{T}$ is the minimal ordinal $\alpha \leq \omega$ such that $T \subseteq \omega^{<\alpha}$.
(d) Let $\mathfrak{T}$ be a pseudo-tree. The tree domain $S := \{\, w \mid w : \langle\rangle \to x,\ \text{for some } x \in T \,\}$ is called the *unravelling* of $\mathfrak{T}$.
(e) Every additive labelling of a pseudo-tree $\mathfrak{T}$ induces an additive labelling λ of its unravelling $\mathfrak{S}$. In the following we will not distinguish between these labellings.

In order to show that every labelling λ of a tree domain T is equivalent to a regular one we construct a finite pseudo-tree and a corresponding labelling.

Definition 3.6. Let T be a tree domain with additive labelling λ, σ a forward Ramseyan split of λ, and $P \subseteq T$ a prefix-closed subset of T.
(a) An infinite, strictly increasing sequence $u_0 \prec u_1 \prec u_2 \prec \dots$ in T is called *homogeneous* if $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ and

$$\lambda(u_i, u_k) = \lambda(u_0, u_1) \quad \text{for all } i < k\,.$$

A homogeneous sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ is *maximal* if there is no homogeneous sequence $v_0 \sqsubset_\sigma v_1 \sqsubset_\sigma \dots$ such that $\{\, u_i \mid i < \omega \,\} \subset \{\, v_i \mid i < \omega \,\}$.
(b) The *contraction* of λ induced by σ and P is the pseudo-tree $\mathfrak{C}^P_\sigma(\lambda)$ obtained from $\mathfrak{T}$ in the following way. For every maximal homogeneous sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$, we consider the minimal index $k \geq 2$ with $u_k \notin P$ and we replace the edge $\langle \mathrm{prec}(u_k), u_k\rangle$ by a back-edge $\langle \mathrm{prec}(u_k), u_1\rangle$. The pseudo-tree $\mathfrak{C}^P_\sigma(\lambda)$ consists of all vertices of the resulting graph that are reachable from the root. We denote by λ^P_σ the labelling of $\mathfrak{C}^P_\sigma(\lambda)$ induced by λ, that is, if e is an edge of both $\mathfrak{T}$ and $\mathfrak{C}^P_\sigma(\lambda)$ then we set $\lambda^P_\sigma(e) := \lambda(e)$ and, if e is a back-edge $\langle \mathrm{prec}(u_k), u_1\rangle$ introduced as replacement for the edge $\langle \mathrm{prec}(u_k), u_k\rangle$ then we set $\lambda^P_\sigma(e) := \lambda(\mathrm{prec}(u_k), u_k)$.

Remark. Note that the definition implies that $P \subseteq C^P_\sigma(\lambda)$. Hence, we can use the set $P \subseteq T$ to mark vertices we want to keep in the contraction.

Below we will make frequent use of the following simple observations.

Lemma 3.7. *Let $\mathfrak{T}$ be a tree, λ an additive labelling of $\mathfrak{T}$, σ a forward Ramseyan split of λ, and $x \leq u \leq y \leq v$ vertices of T.*

(a) *If $x \sqsubseteq_\sigma y$ and $u \sqsubseteq_\sigma v$, then $x \sqsubseteq_\sigma u \sqsubseteq_\sigma y \sqsubseteq_\sigma v$.*

(b) *If $x \sqsubseteq_\sigma y$ and $\langle v, u\rangle$ is a back-edge of $\mathfrak{C}^P_\sigma(\lambda)$, then $x \sqsubseteq_\sigma u \sqsubseteq_\sigma y$.*

Proof. (a) $u \leq y \leq v$ and $u \sqsubseteq_\sigma v$ implies $\sigma(y) \geq \sigma(u) = \sigma(v)$, while $x \leq u \leq y$ and $x \sqsubseteq_\sigma y$ implies $\sigma(u) \geq \sigma(x) = \sigma(y)$. Hence, $\sigma(x) = \sigma(y) = \sigma(u) = \sigma(v)$ and $x \sqsubseteq_\sigma u \sqsubseteq_\sigma y \sqsubseteq_\sigma v$.
(b) Let $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ be the homogeneous sequence inducing the back-edge $\langle v, u\rangle$. Then $u = u_1$ and $v = \mathrm{prec}(u_k)$ for some $k \geq 2$. Since $u \sqsubset_\sigma u_k$, it follows by (a) that $x \sqsubseteq_\sigma u \sqsubseteq_\sigma y \sqsubseteq_\sigma u_k$. □

Let us show that $\lim \lambda^P_\sigma \subseteq \lim \lambda$ and that, for finite P, the labelling λ^P_σ is regular, i.e., that the contraction $\mathfrak{C}^P_\sigma(\lambda)$ is finite. In general, the converse inclusion does not hold, but we shall show below that we can choose the set P such that we obtain an equality $\lim \lambda^P_\sigma = \lim \lambda$. To prove the regularity of λ^P_σ we show that $\mathfrak{C}^P_\sigma(\lambda)$ is finite.

Lemma 3.8. *Let λ be an additive labelling of T, σ a forward Ramseyan split of λ, and $P \subseteq T$ prefix-closed and finite. Every branch of $\mathfrak{C}^P_\sigma(\lambda)$ contains a back-edge.*

Proof. If $\mathfrak{C}^P_\sigma(\lambda)$ had an infinite branch α without back-edges then there would exist an infinite sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma \dots$. By Ramsey's Theorem, we can find an infinite set $I \subseteq \omega$ such that the subsequence $(u_i)_{i \in I}$ is homogeneous. Since P is finite, we have $u_n \notin P$, for all sufficiently large n. Hence, in the above construction we would have created a back-edge $\langle \mathrm{prec}(u_k), u_i \rangle$ for some $i < k$, and all u_l with $l \geq k$ would have been removed from $\mathfrak{C}^P_\sigma(\lambda)$. Contradiction. □

By König's Lemma it follows that $\mathfrak{C}^P_\sigma(\lambda)$ is finite.

Corollary 3.9. *For finite P, $\mathfrak{C}^P_\sigma(\lambda)$ is finite.*

The proof that $\lim \lambda^P_\sigma \subseteq \lim \lambda$ is more involved. We start by showing that we can replace certain paths by paths without back-edges. In the next lemma we deal with finite paths and in the lemma below with infinite paths. Let us introduce some terminology. We call a path *straight* if it contains no back-edges. Note that, for vertices $u \leq v$, there is a unique straight path $u \to v$. For a finite path $w : u \to v$, we denote by $\mathrm{st}(w)$ the corresponding straight path $u \to v$, if it exists.

Lemma 3.10. *Let λ be an additive labelling of a tree $\mathfrak{T}$ and $\mathfrak{C} := \mathfrak{C}^P_\sigma(\lambda)$ a contraction of λ.*

(a) *Let u, v, x, y, z be vertices with $u \leq z \sqsubset_\sigma y < x, v$ such that $e := \langle x, y \rangle$ is a back-edge, and let $w : u \to x$ and $w' : y \to v$ be straight paths. Then*

$$\lambda^P_\sigma(\mathrm{st}(wew')) = \lambda^P_\sigma(wew') .$$

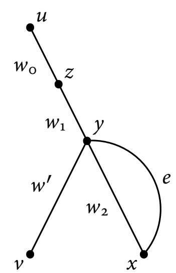

(b) *For every path $w : \langle\rangle \to v$,*

$$\lambda^P_\sigma(\mathrm{st}(w)) = \lambda^P_\sigma(w) .$$

(c) *For every path $w : u \to v$ with $u < v$ such that*

- *w contains no edge $\langle x, y \rangle$ with $y = u$ and*
- *$\sigma(x) \geq \sigma(u)$, for all vertices x of w,*

$$\lambda^P_\sigma(\mathrm{st}(w)) = \lambda^P_\sigma(w) .$$

Proof. (a) We factorise $w = w_0 w_1 w_2$ where $w_0 : u \to z$, $w_1 : z \to y$, and $w_2 : y \to x$. Suppose that $e \in E_c$ and set $x^+ := xc$. By definition of $\mathfrak{C}^P_\sigma(\lambda)$ we have $z \sqsubset_\sigma y \sqsubset_\sigma x^+$. As σ is a forward Ramseyan split, it follows that

$$\lambda^P_\sigma(w_1 w_2 e) = \lambda(z, x^+) = \lambda(z, y) = \lambda^P_\sigma(w_1) .$$

Hence,

$$\begin{aligned} \lambda^P_\sigma(\mathrm{st}(wew')) &= \lambda^P_\sigma(w_0 e w') \\ &= \lambda^P_\sigma(w_0) \cdot \lambda^P_\sigma(w_1) \cdot \lambda^P_\sigma(w') \\ &= \lambda^P_\sigma(w_0) \cdot \lambda^P_\sigma(w_1 w_2 e) \cdot \lambda^P_\sigma(w') = \lambda^P_\sigma(wew') . \end{aligned}$$

(b) Let $w = w_0 e_0 w_1 e_1 \dots w_{n-1} e_{n-1} w_n$ where $e_0, \dots, e_{n-1}$ are the back-edges. We prove the claim by induction on n. Suppose that $e_{n-1} = \langle x, y \rangle$ and set $w' := w_0 e_0 \dots e_{n-2} w_{n-1}$. By induction hypothesis, we have

$$\lambda^P_\sigma(\mathrm{st}(w')) = \lambda^P_\sigma(w') ,$$

which implies that

$$\begin{aligned} \lambda^P_\sigma(\mathrm{st}(w') e_{n-1} w_n) &= \lambda^P_\sigma(\mathrm{st}(w')) \cdot \lambda^P_\sigma(e_{n-1} w_n) \\ &= \lambda^P_\sigma(w') \cdot \lambda^P_\sigma(e_{n-1} w_n) \\ &= \lambda^P_\sigma(w) . \end{aligned}$$

By definition of $\mathfrak{C}^P_\sigma(\lambda)$, there exists a vertex $z < y$ with $z \sqsubset_\sigma y$. Setting $u := \langle\rangle$ it follows by (a) that

$$\lambda^P_\sigma(\mathrm{st}(w)) = \lambda^P_\sigma(\mathrm{st}(w') e_{n-1} w_n) = \lambda^P_\sigma(w) .$$

(c) Let $w = w_0 e_0 w_1 e_1 \dots w_{n-1} e_{n-1} w_n$ where $e_0, \dots, e_{n-1}$ are the back-edges. We prove the claim by induction on n. Suppose that $e_{n-1} = \langle x, y \rangle$ and set $w' := w_0 e_0 \dots e_{n-2} w_{n-1}$. By induction hypothesis,

$$\lambda^P_\sigma(\mathrm{st}(w')) = \lambda^P_\sigma(w') ,$$

which implies that

$$\begin{aligned}\lambda^P_\sigma(\mathrm{st}(w')e_{n-1}w_n) &= \lambda^P_\sigma(\mathrm{st}(w'))\cdot\lambda^P_\sigma(e_{n-1}w_n)\\ &= \lambda^P_\sigma(w')\cdot\lambda^P_\sigma(e_{n-1}w_n)\\ &= \lambda^P_\sigma(w)\,.\end{aligned}$$

If $\sigma(y) > \sigma(u)$ then, by definition of $\mathfrak{C}^P_\sigma(\lambda)$, there exists a vertex $u < z < y$ with $z \sqsubset_\sigma y$. Hence, it follows by (a) that

$$\lambda^P_\sigma(\mathrm{st}(w)) = \lambda^P_\sigma(\mathrm{st}(w')e_{n-1}w_n) = \lambda^P_\sigma(w)\,.$$

Similarly, if $\sigma(y) = \sigma(u)$ then we have $u \sqsubset_\sigma y$. Hence, setting $z := u$ it follows by (a) that

$$\lambda^P_\sigma(\mathrm{st}(w)) = \lambda^P_\sigma(\mathrm{st}(w')e_{n-1}w_n) = \lambda^P_\sigma(w)\,. \qquad \square$$

Lemma 3.11. *Let λ be an additive labelling of a tree $\mathfrak{T}$ and $\mathfrak{C} := \mathfrak{C}^P_\sigma(\lambda)$ a contraction of λ.*

(a) *If $e = \langle x, y\rangle$ is a back-edge of $\mathfrak{C}$ and if $w : \langle\rangle \to y$ and $w' : y \to x$ are straight paths then*

$$\lambda^P_\sigma(w)\cdot\big(\lambda^P_\sigma(w'e)\big)^\omega \in \lim\lambda\,.$$

(b) *For every branch $\alpha : \langle\rangle \to \infty$, there exist a back-edge $e = \langle x, y\rangle$ of $\mathfrak{C}$ and straight paths $w : \langle\rangle \to y$ and $w' : y \to x$ such that*

$$\lambda^P_\sigma(\alpha) = \lambda^P_\sigma(w)\cdot\big(\lambda^P_\sigma(w'e)\big)^\omega.$$

Proof. (a) By definition of $\mathfrak{C}^P_\sigma(\lambda)$ there exists an infinite homogeneous sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ and an index k such that $y = u_1$ and $x = \mathrm{prec}(u_k)$. It follows that

$$\begin{aligned}\lambda^P_\sigma(w)\cdot(\lambda^P_\sigma(w'e))^\omega &= \lambda(\langle\rangle, u_1)\cdot(\lambda(u_1, u_k))^\omega\\ &= \lambda(\langle\rangle, u_0)\cdot(\lambda(u_0, u_1))^\omega\\ &= \lambda(\langle\rangle, u_0)\cdot\prod_{n<\omega}\lambda(u_n, u_{n+1}) \in \lim\lambda\end{aligned}$$

(b) Let k be the minimal number such that there exists a factorisation $\alpha = w_0e_0w_1e_1\dots$ where each $e_i = \langle x_i, y_i\rangle$ is a back-edge with $\sigma(y_i) = k$. We assume that this factorisation is chosen such that, for $i > 0$, the path w_i contains no vertices x with $\sigma(x) < k$. By the Pigeon Hole Principle, we can choose an infinite subset $I \subseteq \omega$ such that $y_i = y_l$, for all $i, l \in I$. By choosing the factorisation of α suitably we may w.l.o.g. assume that $I = \omega$, that $|y_i|$ is minimal, and that the paths w_i, for $i > 0$, only contain vertices z with $y_i \leq z$. Furthermore, by Ramsey's Theorem, there exist an element $s \in S$ and an infinite subset $I \subseteq \omega$ such that

$$\lambda^P_\sigma(w_{i+1}e_{i+1}\dots w_le_l) = s\,, \quad \text{for all } i < l \text{ in } I\,.$$

Again, we may assume w.l.o.g. that $I = \omega$, that is, $\lambda^P_\sigma(w_ie_i) = s$, for all $i > 0$. By Lemma 3.10 (b), the straight path $w'_0 := \mathrm{st}(w_0e_0)$ satisfies $\lambda^P_\sigma(w'_0) = \lambda^P_\sigma(w_0e_0)$. Similarly, it follows by Lemma 3.10 (c) that the straight paths $w'_i := \mathrm{st}(w_i)$, $i > 0$, satisfy $\lambda^P_\sigma(w'_i) = \lambda^P_\sigma(w_i)$. Hence,

$$\begin{aligned}\lambda^P_\sigma(\alpha) &= \lambda^P_\sigma(w_0e_0)\cdot\prod_{n<\omega}\lambda^P_\sigma(w_{n+1}e_{n+1})\\ &= \lambda^P_\sigma(w'_0)\cdot\prod_{n<\omega}\lambda^P_\sigma(w'_{n+1}e_{n+1})\\ &= \lambda^P_\sigma(w'_0)\cdot\big(\lambda^P_\sigma(w'_1e_1)\big)^\omega. \qquad \square\end{aligned}$$

The preceding lemma implies that $\lim\lambda^P_\sigma$ is a subset of $\lim\lambda$.

Proposition 3.12. *Let λ be an additive labelling of a tree $\mathfrak{T}$ and $\mathfrak{C}^P_\sigma(\lambda)$ a finite contraction of λ. Then $\lim\lambda^P_\sigma \subseteq \lim\lambda$.*

Proof. Given a branch α of $\mathfrak{T}$, we can use Lemma 3.11 (b) to find a back-edge $e = \langle x, y\rangle$ and straight paths $w : \langle\rangle \to y$ and $w' : y \to x$ such that

$$\lambda^P_\sigma(\alpha) = \lambda^P_\sigma(w)\cdot\big(\lambda^P_\sigma(w'e)\big)^\omega.$$

By Lemma 3.11 (a), it follows that $\lambda^P_\sigma(\alpha) \in \lim\lambda$. $\square$

Next we prove that, by choosing P suitably, we can ensure that $\lim\lambda^P_\sigma = \lim\lambda$.

Lemma 3.13. *Let λ be an additive labelling of a tree $\mathfrak{T}$ and σ a forward Ramseyan split of λ. There exists a finite prefix-closed subset $P \subseteq T$ such that $\lim\lambda^P_\sigma = \lim\lambda$.*

Proof. For every $s \in \lim\lambda$, we choose a branch α^s of T with $\lambda(\alpha^s) = s$. Let $u^s_0 \sqsubset_\sigma u^s_1 \sqsubset_\sigma u^s_2 \sqsubset \dots$ be the maximal homogeneous sequence with $u^s_n \leq \alpha^s$, for all n. We claim that

$$P := \{\, x \in T \mid x < u^s_2 \text{ for some } s \in \lim\lambda \,\}$$

is the desired set. Setting $v^s := \mathrm{prec}(u_2^s)$ it follows that $v^s \in P \subseteq C^P_\sigma(\lambda)$ and $e^s := \langle v^s, u_1^s \rangle$ is a back-edge of $\mathfrak{C}^P_\sigma(\lambda)$. Consequently,

$$\begin{aligned} s &= \lambda(\langle\rangle, u_1^s) \cdot \prod_{n<\omega} \lambda(u_n^s, u_{n+1}^s) \\ &= \lambda(\langle\rangle, u_1^s) \cdot \big(\lambda(u_1^s, u_2^s)\big)^\omega \\ &= \lambda(\langle\rangle, u_1^s) \cdot \big(\lambda(u_1^s, v^s) \cdot \lambda(v^s, u_2^s)\big)^\omega \\ &= \lambda^P_\sigma(\langle\rangle, u_1^s) \cdot \big(\lambda^P_\sigma(u_1^s, v^s) \cdot \lambda^P_\sigma(e^s)\big)^\omega \in \lim \lambda^P_\sigma . \end{aligned}$$ □

Summarising we have obtained the following result.

Corollary 3.14. *For every additive labelling λ of a tree $\mathfrak{T}$, there exists a finite pseudo-tree $\mathfrak{T}'$ and an additive labelling λ' of $\mathfrak{T}'$ such that* $\lim \lambda' = \lim \lambda$.

4 Finding small contractions

Corollary 3.14 does not include a bound on the size of the pseudo-tree $\mathfrak{T}'$. In this section we prove two theorems that improve Corollary 3.14 by providing bounds on the height of the resulting pseudo-tree. To compute such a bound, we show that every pseudo-tree whose height is larger than a certain threshold can be reduced in size. The following proposition introduces the operation we use to shrink a pseudo-tree.

Proposition 4.1. *Let λ be an additive labelling of a tree $\mathfrak{T}$ and $\mathfrak{C}^P_\sigma(\lambda)$ a contraction of λ. Suppose that $x, y, z \in C^P_\sigma(\lambda)$ are vertices such that $x \sqsubset_\sigma y \sqsubset_\sigma z$ and there is no back-edge $\langle v, u\rangle$ with $y \leq u < z \leq v$. Let $\mathfrak{T}'$ be the tree obtained from $\mathfrak{T}$ by replacing the subtree rooted at y by the one rooted at z and let λ', σ', and P' be obtained from, respectively, λ, σ, and P in the same way. (In particular, if $y \in P$, but $z \notin P$, we have $y \notin P'$.)*

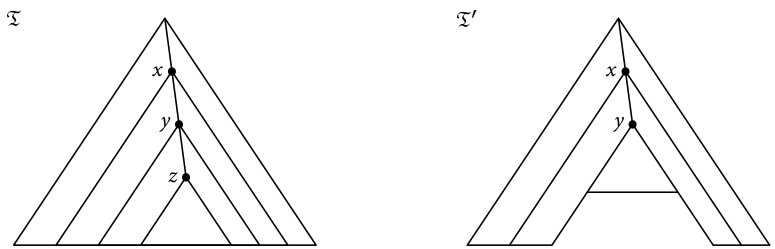

(a) *σ' is a forward Ramseyan split of λ' such that*

$$\lim (\lambda')^{P'}_{\sigma'} \subseteq \lim \lambda^P_\sigma \quad \textit{and} \quad |C^{P'}_{\sigma'}(\lambda')| < |C^P_\sigma(\lambda)| .$$

(b) *If, for every back-edge $e = \langle v, u\rangle$ with $z \nleq v$, there is some back-edge $\tilde{e} = \langle \tilde{v}, \tilde{u}\rangle$ with $y \nleq \tilde{v}$ such that*

$$\lambda(\langle\rangle, \tilde{u}) \cdot (\lambda(\tilde{u}, \tilde{v}) \cdot \lambda(\tilde{e}))^\omega = \lambda(\langle\rangle, u) \cdot (\lambda(u, v) \cdot \lambda(e))^\omega ,$$

then $\lim (\lambda')^{P'}_{\sigma'} = \lim \lambda^P_\sigma$.

Proof. (a) Let $h : T' \to T$ be the function where

$$h(u) := \begin{cases} zv & \text{if } u = yv , \\ u & \text{if } y \nleq u . \end{cases}$$

We start by showing that $\sigma' = \sigma \circ h$ is a forward Ramseyan split of λ'. Clearly, $hu \sqsubset_\sigma hv$ implies $u \sqsubset_{\sigma'} v$. Since $y \sqsubset_\sigma z$, it also follows that $u \sqsubset_{\sigma'} v$ implies $hu \sqsubset_\sigma hv$. Hence, to show that σ' is a Ramseyan split it is sufficient to prove that

$$u \sqsubset_{\sigma'} v \quad \text{implies} \quad \lambda'(u, v) = \lambda(hu, hv) .$$

If $y \leq u$ or $y \nless v$, then the claim follows by definition of λ'. Therefore, we may assume that $u < y < v$. If $u < x$ then

$$\begin{aligned} \lambda'(u, v) &= \lambda'(u, x) \cdot \lambda'(x, y) \cdot \lambda'(y, v) \\ &= \lambda(hu, x) \cdot \lambda(x, y) \cdot \lambda(z, hv) \\ &= \lambda(hu, x) \cdot \lambda(x, z) \cdot \lambda(z, hv) = \lambda(hu, hv) . \end{aligned}$$

If $x \leq u$ then Lemma 3.7 (a) implies $x \sqsubseteq_\sigma u \sqsubseteq_\sigma y \sqsubseteq_\sigma z$. Since σ is a forward Ramseyan split, it follows that

$$\begin{aligned} \lambda'(u, v) &= \lambda'(u, y) \cdot \lambda'(y, v) \\ &= \lambda(hu, y) \cdot \lambda(z, hv) \\ &= \lambda(hu, z) \cdot \lambda(z, hv) = \lambda(hu, hv) . \end{aligned}$$

We have shown that σ' is a forward Ramseyan split of λ'. Next we prove that $h[C^{P'}_{\sigma'}(\lambda')] \subseteq C^P_\sigma(\lambda)$. Since there is no back-edge $\langle v, u\rangle$ with $y < u < z \leq v$, it is sufficient to show that, if $\langle hv, hu\rangle$ is a back-edge of $\mathfrak{C}^P_\sigma(\lambda)$ with label d, then

$\langle v, u\rangle$ is a back-edge of $\mathfrak{C}^{P'}_{\sigma'}(\lambda')$ with label d. Hence, suppose that $\langle hv, hu\rangle$ is a back-edge of $\mathfrak{C}^{P}_{\sigma}(\lambda)$ and let $w_0 \sqsubset_\sigma w_1 \sqsubset_\sigma w_2 \sqsubset_\sigma \dots$ be the maximal homogeneous sequence inducing it. If there is no n with $y \leq w_n$, then $w_0 \sqsubset_{\sigma'} w_1 \sqsubset_{\sigma'} w_2 \sqsubset_{\sigma'} \dots$ is also a maximal homogeneous sequence for λ' and we are done. Otherwise, let i be the maximal index with $w_i \prec y$ and let k be the minimal index with $z \leq w_k$. Since z is a vertex of the contraction, we have $hu = w_1$ and $hv = w_m$, for some $m \geq k$. It follows that the sequence $w_0 \sqsubset_{\sigma'} w_1 \sqsubset_{\sigma'} \cdots \sqsubset_{\sigma'} w_i \sqsubset_{\sigma'} h^{-1}w_k \sqsubset_{\sigma'} h^{-1}w_{k+1} \sqsubset_{\sigma'} \dots$ is a maximal homogeneous for λ'. This sequence induces the back-edge $\langle h^{-1}w_m, w_1\rangle = \langle v, u\rangle$.

We have shown that $h[C^{P'}_{\sigma'}(\lambda')] \subseteq C^{P}_{\sigma}(\lambda)$. Since $y \notin h[C^{P'}_{\sigma'}(\lambda)]$ this inclusion is strict. Consequently, $|C^{P'}_{\sigma'}(\lambda')| < |C^{P}_{\sigma}(\lambda)|$.

It remains to prove that $\lim\,(\lambda')^{P'}_{\sigma'} \subseteq \lim \lambda^P_\sigma$. Let α be a branch of $\mathfrak{C}^{P'}_{\sigma'}(\lambda')$. By Lemma 3.11 (b), we may assume that $\alpha = w(w'e)^\omega$, where $e = \langle v, u\rangle$ is a back-edge and w, w' are straight paths. If α does not contain the vertex y then α is a branch of $\mathfrak{C}^{P}_{\sigma}(\lambda)$ and we are done.

If $y \leq u$, we have

$$\begin{aligned}(\lambda')^{P'}_{\sigma'}(\alpha) &= \lambda'(\langle\rangle, x)\cdot\lambda'(x,y)\cdot\lambda'(y,u)\cdot(\lambda'(u,v)\cdot\lambda'(e))^\omega\\ &= \lambda(\langle\rangle, x)\cdot\lambda(x,y)\cdot\lambda(z,hu)\cdot(\lambda(hu,hv)\cdot\lambda(he))^\omega\\ &= \lambda(\langle\rangle, x)\cdot\lambda(x,z)\cdot\lambda(z,hu)\cdot(\lambda(hu,hv)\cdot\lambda(he))^\omega\\ &= \lambda(\langle\rangle, hu)\cdot(\lambda(hu,hv)\cdot\lambda(he))^\omega \in \lim \lambda^P_\sigma\,.\end{aligned}$$

Similarly, if $u \leq x \leq y \leq v$,

$$\begin{aligned}(\lambda')^{P'}_{\sigma'}(\alpha) &= \lambda'(\langle\rangle, u)\cdot(\lambda'(u,x)\cdot\lambda'(x,y)\cdot\lambda'(y,v)\cdot\lambda'(e))^\omega\\ &= \lambda(\langle\rangle, u)\cdot(\lambda(u,x)\cdot\lambda(x,y)\cdot\lambda(z,hv)\cdot\lambda(he))^\omega\\ &= \lambda(\langle\rangle, u)\cdot(\lambda(u,x)\cdot\lambda(x,z)\cdot\lambda(z,hv)\cdot\lambda(he))^\omega\\ &= \lambda(\langle\rangle, u)\cdot(\lambda(u,hv)\cdot\lambda(he))^\omega \in \lim \lambda^P_\sigma\,.\end{aligned}$$

It remains to consider the case that $x \prec u \prec y \leq v$. Then Lemma 3.7 implies $x \sqsubset_\sigma u \sqsubset_\sigma y \sqsubseteq_\sigma z \sqsubset_\sigma hv'$, where v' is the successor of v in $\mathfrak{T}'$ inducing the back-edge e. Since σ is a forward Ramseyan split, it follows that

$$\lambda(u, hv') = \lambda(u,y)\cdot\lambda(y,hv') = \lambda(u,y) = \lambda(u,y)\cdot\lambda(z,hv')\,.$$

Let $\tilde{w}'$ be the path in $\mathfrak{C}^{P}_{\sigma}(\lambda)$ obtained from w' by inserting the path $y \to z$ at the appropriate place. It follows that

$$\lambda^P_\sigma(\tilde{w}'he) = \lambda(u,hv') = \lambda(u,y)\cdot\lambda(z,hv') = \lambda^{P'}_{\sigma'}(w'e)\,.$$

Consequently, $(\lambda')^{P'}_{\sigma'}(w(w'e)^\omega) = \lambda^P_\sigma(w(\tilde{w}'e)^\omega) \in \lim \lambda^P_\sigma$.

(b) By (a) it remains to show that $\lim \lambda^P_\sigma \subseteq \lim\,(\lambda')^{P}_{\sigma'}$. Let α be a branch of $\mathfrak{C}^{P}_{\sigma}(\lambda)$. Again we may assume that $\alpha = w(w'e)^\omega$ where $e = \langle v, u\rangle$ is a back-edge and w, w' are straight paths. Let v' be the successor of v inducing e.

If α does not contain the vertex y then α is a branch of $\mathfrak{C}^{P'}_{\sigma'}(\lambda')$ and $\lambda(\alpha) = \lambda'(\alpha) \in \lim \lambda'$. Hence, we may assume that $y \leq v$.

Let us first consider the case that $z \nleq v$. By assumption, there is some back-edge $\tilde{e} = \langle \tilde{v}, \tilde{u}\rangle$ with $y \nleq \tilde{v}$ such that

$$\lambda(\langle\rangle, \tilde{u})\cdot(\lambda(\tilde{u},\tilde{v})\cdot\lambda(\tilde{e}))^\omega = \lambda(\langle\rangle, u)\cdot(\lambda(u,v)\cdot\lambda(e))^\omega\,.$$

It follows that

$$\begin{aligned}\lambda^P_\sigma(\alpha) &= \lambda(\langle\rangle, u)\cdot(\lambda(u,v)\cdot\lambda(e))^\omega\\ &= \lambda(\langle\rangle, \tilde{u})\cdot(\lambda(\tilde{u},\tilde{v})\cdot\lambda(\tilde{e}))^\omega\\ &= \lambda'(\langle\rangle, \tilde{u})\cdot(\lambda'(\tilde{u},\tilde{v})\cdot\lambda'(\tilde{e}))^\omega \in \lim(\lambda')^{P'}_{\sigma'}\,.\end{aligned}$$

Hence, it remains to consider the case that $z \leq v$. We distinguish several cases, depending on the position of u. By assumption, there is no back-edge $\langle v_0, u_0\rangle$ with $y \leq u_0 \prec z \leq v_0$. Therefore, either $u \prec y$ or $z \leq u$.

If $z \leq u$, we have

$$\begin{aligned}\lambda^P_\sigma(\alpha) &= \lambda(\langle\rangle, x)\cdot\lambda(x,z)\cdot\lambda(z,u)\cdot\lambda(u,v')^\omega\\ &= \lambda(\langle\rangle, x)\cdot\lambda(x,y)\cdot\lambda(z,u)\cdot\lambda(u,v')^\omega\\ &= \lambda'(\langle\rangle, x)\cdot\lambda'(x,y)\cdot\lambda'(y,h^{-1}u)\cdot\lambda'(h^{-1}u,h^{-1}v')^\omega\\ &= \lambda'(\langle\rangle, h^{-1}u)\cdot\lambda'(h^{-1}u,h^{-1}v')^\omega \in \lim(\lambda')^{P'}_{\sigma'}\,.\end{aligned}$$

Similarly, if $u \leq x$,

$$\begin{aligned}\lambda^P_\sigma(\alpha) &= \lambda(\langle\rangle, u)\cdot(\lambda(u,x)\cdot\lambda(x,z)\cdot\lambda(z,v'))^\omega\\ &= \lambda(\langle\rangle, u)\cdot(\lambda(u,x)\cdot\lambda(x,y)\cdot\lambda(z,v'))^\omega\\ &= \lambda'(\langle\rangle, u)\cdot(\lambda'(u,x)\cdot\lambda'(x,y)\cdot\lambda'(y,h^{-1}v'))^\omega\\ &= \lambda'(\langle\rangle, u)\cdot\lambda'(u,h^{-1}v')^\omega \in \lim(\lambda')^{P'}_{\sigma'}\,.\end{aligned}$$

Hence, it remains to consider the case that $x \prec u \prec y$. As above, Lemma 3.7

implies that $x \sqsubset_\sigma u \sqsubset_\sigma y \sqsubset_\sigma z \sqsubset_\sigma v'$. Therefore,

$$\begin{aligned}
\lambda^P_\sigma(\alpha) &= \lambda(\langle\rangle, u) \cdot (\lambda(u,z) \cdot \lambda(z,v'))^\omega \\
&= \lambda(\langle\rangle, u) \cdot (\lambda(u,y) \cdot \lambda(z,v'))^\omega \\
&= \lambda'(\langle\rangle, u) \cdot (\lambda'(u,y) \cdot \lambda'(y,h^{-1}v'))^\omega \\
&= \lambda'(\langle\rangle, u) \cdot \lambda'(u,h^{-1}v')^\omega \in \lim(\lambda')^{P'}_{\sigma'} .
\end{aligned}$$

□

It remains to compute the size of a minimal pseudo-tree. In order to apply the operation from the proposition, we try to find vertices $x \sqsubset_\sigma y \sqsubset_\sigma z$ such that no back-edge from below z ends between y and z. In the following sequence of technical lemmas we show how to find such vertices.

Lemma 4.2. *Suppose that $e := \langle y,x\rangle$ is a back-edge in $\mathfrak{C}^P_\sigma(\lambda)$ and set $a := \lambda(x,y)\cdot\lambda(e)$.*

(a) *There is exactly one vertex $w \sqsubset_\sigma x$ such that $\lambda(w,x) = a$.*

(b) *For every vertex z with $x \sqsubset_\sigma z \preceq y$ we have $\lambda(x,z) = a$.*

Proof. Let $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ be the homogeneous sequence inducing the back-edge e. Then $x = u_1$ and $y = \mathrm{prec}(u_k)$ for some $k \geq 2$.

(a) Setting $w := u_0$ we have $w \sqsubset_\sigma x$ and $\lambda(w,x) = \lambda(x,u_k) = \lambda(x,y)\cdot\lambda(e) = a$. Hence, it remains to show that there is not a second such vertex. For a contradiction, suppose that there is some $w \neq u_0$ with $w \sqsubset_\sigma x$ and $\lambda(w,x) = a$.

If $w < u_0$ then we have $w \sqsubset_\sigma u_0$ and $\lambda(w,u_0) = \lambda(w,x) = a$. Therefore,

$$w \sqsubset_\sigma u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$$

is homogeneous. Hence, the sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ is not maximal. A contradiction.

Similarly, if $u_0 < w$ then $u_0 \sqsubset_\sigma w \sqsubset_\sigma x$ and

$$\lambda(u_0,w) = \lambda(u_0,x) = a \quad \text{and} \quad \lambda(w,x) = a .$$

Therefore, $u_0 \sqsubset_\sigma w \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ is homogeneous and, again, the sequence $u_0 \sqsubset_\sigma u_1 \sqsubset_\sigma u_2 \sqsubset_\sigma \dots$ is not maximal.

(b) Since $u_1 = x \sqsubset_\sigma z < u_k$ we have $x \sqsubset_\sigma z \sqsubset_\sigma u_k$. Consequently, $\lambda(x,z) = \lambda(x,u_k) = a$. □

Lemma 4.3. *Let $x_0 \sqsubset_\sigma \dots \sqsubset_\sigma x_n$ be vertices in $\mathfrak{C}^P_\sigma(\lambda)$.*

(a) *There are at most $|S|$ indices $i < n$ such that there exists a back-edge $\langle y, x_i\rangle$ with $x_{i+1} \preceq y$.*

(b) *If $n \geq (m+1)(|S|+1)$ then we can find an index $i \leq n-m$ such that there is no back-edge $\langle v,u\rangle$ such that*

$$x_i \preceq u < x_k \preceq v, \quad \textit{for any } k \textit{ with } i < k \leq i+m .$$

Proof. (a) Suppose that there are more than $m := |S|$ such indices. By the Pigeon Hole Principle, we can then find indices $i < k$ and back-edges $e_0 = \langle y,x_i\rangle$ and $e_1 = \langle z,x_k\rangle$ with

$$x_{i+1} \preceq y, \quad x_{k+1} \preceq z, \quad \text{and} \quad \lambda(x_i,y)\cdot\lambda(e_0) = \lambda(x_k,z)\cdot\lambda(e_1) .$$

Let $a := \lambda(x_i,y)\cdot\lambda(e_0)$. By Lemma 4.2 (a) there are unique vertices $u \sqsubset_\sigma x_i$ and $v \sqsubset_\sigma x_k$ such that

$$\lambda(u,x_i) = a \quad \text{and} \quad \lambda(v,x_k) = a .$$

Note that, since $u \sqsubset_\sigma x_i \sqsubset_\sigma x_k$ and σ is a forward Ramseyan split, we have

$$\lambda(u,x_k) = \lambda(u,x_i) = a .$$

Consequently, $u = v$. Furthermore, $x_i \sqsubset_\sigma x_{i+1} \preceq y$ implies by Lemma 4.2 (b) that

$$\lambda(x_i,x_{i+1}) = \lambda(x_i,y)\cdot\lambda(e_0) = a .$$

Therefore, $\lambda(x_i,x_k) = \lambda(x_i,x_{i+1}) = a$. By uniqueness of v, it follows that $v = x_i$. But $u = v = x_i$ contradicts $u < x_i$.

(b) Let $u_0 < \dots < u_{l-1}$ be an enumeration of all vertices u such that, for some i, there is a back-edge $\langle v,u\rangle$ with $x_i \preceq u < x_{i+1} \preceq v$. For each u_j, we fix one such back-edge $e_j = \langle v_j,u_j\rangle$.

By Lemma 3.7 (b), $x_i \preceq u_j < x_{i+1} \preceq v_j$ implies $\sigma(u_j) = \sigma(x_i) = \sigma(x_0)$ and, hence,

$$u_0 \sqsubset_\sigma \dots \sqsubset_\sigma u_{l-1} .$$

Consequently, we obtain an increasing chain

$$y_0 \sqsubset_\sigma \dots \sqsubset_\sigma y_s \quad \text{with} \quad \{y_0,\dots,y_s\} = \{x_0,\dots,x_n,u_0,\dots,u_{l-1}\} .$$

Given an index $j < l$, let $k \leq s$ and $i < n$ be the indices such that $y_k = u_j$ and $x_i \leq u_j < x_{i+1}$. Then $y_k = u_j < y_{k+1} \leq x_{i+1} \leq v_j$. Hence, there are at least l indices k such that there exists a back edge $\langle v, u\rangle$ with $u = y_k$ and $y_{k+1} \leq v$. Applying (a) to the sequence $y_0 \sqsubset_\sigma \dots \sqsubset_\sigma y_s$, it follows that $l \leq |S|$.

Since $n \geq (m+1)(|S|+1) \geq (m+1)(l+1)$ we can find an index $i \leq n - m$ such that there is no j with $x_i \leq u_j < x_{i+m}$. This index has the desired properties since, if $\langle v, u\rangle$ is a back-edge with $x_i \leq u < x_k \leq v$, for some $i < k \leq i + m$, then $u = u_j$, for some j. By choice of i, this is not possible. □

Combining the preceding results we obtain the desired reduction from arbitrary trees to regular ones.

Theorem 4.4. *Let $\langle S, S_\omega\rangle$ be a finite ω-semigroup. There exists a finite number $m < \omega$ with the following property: for every additive labelling λ of a tree $\mathfrak{T}$, there exists a finite pseudo-tree $\mathfrak{T}'$, an additive labelling λ' of $\mathfrak{T}'$, and a forward Ramseyan split σ' of λ' such that*

$$\lim \lambda' = \lim \lambda \quad \textit{and} \quad \mathrm{ht}(\mathfrak{C}^{P}_{\sigma'}(\lambda')) < m^{|S|}.$$

Proof. Set $l := |S|$ and $m := (l+1)(2^l + 2) + 2$. By Corollary 3.14, there exists a finite pseudo-tree of the form $\mathfrak{C}^{P'}_{\sigma'}(\lambda')$ such that $\lim \lambda' = \lim \lambda$. We choose $\mathfrak{C}^{P'}_{\sigma'}(\lambda')$ such that its number of vertices is minimal. We claim that

$$\mathrm{ht}(\mathfrak{C}^{P'}_{\sigma'}(\lambda')) < m^{|S|}.$$

For a contradiction, suppose otherwise. Then there exists a vertex $z \in C^{P'}_{\sigma'}(\lambda')$ of length $|z| = m^{|S|} - 1$. We can use Lemma 3.4 to find a sequence $x_0 \sqsubset_\sigma \dots \sqsubset_\sigma x_{m-1} \leq z$. Applying Lemma 4.3 (b) to the sequence $x_1 \sqsubset_\sigma \dots \sqsubset_\sigma x_{m-1}$ we obtain an index $i > 0$ such that there is no back-edge $\langle v, u\rangle$ with

$$x_i \leq u < x_k \leq v\,, \quad \text{for some } i < k \leq i + \frac{m-2}{l+1} - 1\,.$$

Set $j := i + \frac{m-2}{l+1} - 1$. To each index $i \leq s < j$, we associate the set

$$X_s := \big\{\, \lambda(\langle\rangle, u) \cdot (\lambda(u,v) \cdot \lambda(e))^\omega \;\big|\; e = \langle v, u\rangle \text{ a back-edge with } x_s \not\leq v \,\big\}\,.$$

Since $j - i = \frac{m-2}{l+1} - 1 = 2^l + 1 > 2^l$ we can use the Pigeon Hole Principle to find indices $i \leq r < s < j$ with $X_r = X_s$. Hence, we can apply Proposition 4.1 to the sequence $x_0 \sqsubset_\sigma x_r \sqsubset_\sigma x_s$ to construct a pseudo-tree $\mathfrak{C}^{P''}_{\sigma''}(\lambda'')$ with $|C^{P''}_{\sigma''}(\lambda'')| < |C^{P'}_{\sigma'}(\lambda')|$ and $\lim \lambda'' = \lim \lambda'$. This contradicts the minimality of $|C^{P'}_{\sigma'}(\lambda')|$. □

In this theorem we are given an arbitrary tree and we produce an equivalent regular one, i.e., the unravelling of a finite pseudo-tree. If the input tree is already regular, the tree we obtain is not only regular, but we can also describe its relation to the input tree. This is the content of the following theorem.

Definition 4.5. Let $\mathfrak{T}$ be the unravelling of a finite pseudo-tree $\mathfrak{T}_0$ and let $p : T \to T_0$ be the canonical projection. A set $\mathcal{L}$ of additive labellings of $\mathfrak{T}$ is *closed under p-compatible substitutions* if, for all labellings $\lambda \in \mathcal{L}$ and all vertices $x, y \in T$ with $p(x) = p(y)$, the set $\mathcal{L}$ also contains the labelling obtained from λ by replacing the subtree rooted at x by the subtree rooted at y.

Theorem 4.6. *Let $\langle S, S_\omega\rangle$ be a finite ω-semigroup and $\mathfrak{T}_0$ a finite pseudo-tree. There exists a finite number $N < \omega$ with the following property: suppose that $\mathfrak{T}$ is the unravelling of $\mathfrak{T}_0$, $p : T \to T_0$ the canonical projection, and $\mathcal{L}$ a set of additive labellings of $\mathfrak{T}$ that is closed under p-compatible substitutions; for every $\lambda \in \mathcal{L}$, there is a labelling $\lambda' \in \mathcal{L}$, a Ramseyan split σ of λ', and a set $P \subseteq T$ such that*

$$\lim \lambda = \lim (\lambda')^P_\sigma \quad \textit{and} \quad \mathrm{ht}(\mathfrak{C}^P_\sigma(\lambda')) < N\,.$$

Proof. The proof follows the same lines as that of Theorem 4.4. Set $l := |S|$, $b := |T_0|$, and $m := (l+1)(2^l+2)+2$. We choose a labelling $\lambda' \in \mathcal{L}$, a Ramseyan split σ of λ', and a set $P \subseteq T$ such that $\lim(\lambda')^P_\sigma = \lim \lambda$ and the size of $\mathfrak{C}^P_\sigma(\lambda')$ is minimal. Note that this is possible since, choosing $\lambda' := \lambda$ and a Ramseyan split σ of λ', we can use Lemma 3.13 to find a subset $P \subseteq T$ such that $\lim(\lambda')^P_\sigma = \lim \lambda$. We claim that

$$\mathrm{ht}(\mathfrak{C}^P_\sigma(\lambda')) < (bm)^l\,.$$

For a contradiction, suppose otherwise. Then there exists a vertex $z \in C^P_\sigma$ with $|z| = (bm)^l - 1$. We can use Lemma 3.4 to find a sequence $y_0 \sqsubset_\sigma \dots \sqsubset_\sigma y_{bm-1} \leq z$. Being an element of T_0, the function value $p(y_i)$ can take at most $|T_0| = b$ possible values. Hence, the sequence $y_0 \sqsubset_\sigma \dots \sqsubset_\sigma y_{bm-1}$ contains a subsequence $x_0 \sqsubset_\sigma \dots \sqsubset_\sigma x_{m-1}$ with $p(x_0) = \dots = p(x_{m-1})$. Applying Lemma 4.3 (b) to the sequence $x_1 \sqsubset_\sigma \dots \sqsubset_\sigma x_{m-1}$ we obtain an index $i > 0$ such that there is no back-edge $\langle v, u\rangle$ with

$$x_i \leq u < x_k \leq v\,, \quad \text{for some } i < k \leq i + \frac{m-2}{l+1} - 1\,.$$

Set $j := i + \frac{m-2}{l+1} - 1$. To each index $i \leq s < j$, we associate the set

$$X_s := \big\{\, \lambda(\langle\rangle, u) \cdot (\lambda(u,v) \cdot \lambda(e))^\omega \;\big|\; e = \langle v, u\rangle \text{ a back-edge with } x_s \not\leq v \,\big\}\,.$$

Since $j - i = \frac{m-2}{l+1} - 1 = 2^l + 1 > 2^l$ we can use the Pigeon Hole Principle to find indices $i \leq r < s < j$ with $X_r = X_s$. Hence, we can apply Proposition 4.1 to the sequence $x_0 \sqsubset_\sigma x_r \sqsubset_\sigma x_s$ to construct a tree $\mathfrak{T}''$, a labelling λ'' of $\mathfrak{T}'$ and a contraction $\mathfrak{C}^{P''}_{\sigma''}(\lambda'')$ with $|C^{P''}_{\sigma''}(\lambda'')| < |C^P_\sigma(\lambda')|$ and $\lim \lambda'' = \lim \lambda'$. Note that $p(x_r) = p(x_s)$ implies that $\mathfrak{T}''$ is isomorphic to $\mathfrak{T}$ and that $\lambda'' \in \mathcal{L}$. Therefore, we obtain a contradiction to the minimality of $|C^P_\sigma(\lambda')|$. □

5 Power-hyperclones

Having developed the combinatorial machinery of Theorems 4.4 and 4.6, we turn to our original problem of finding a finite representation for ω-hyperclones. This is the same situation as with ω-semigroups. For an ω-semigroup $\mathfrak{S}$, the solution is to replace the infinite product $\pi : S^\omega \to S$ by a unary operation ${}^\omega : S \to S$ computing the ω-th power of its argument. The resulting type of algebra is called a *Wilke algebra*. The same thing can be done for an ω-hyperclone $\mathfrak{C}$. We replace the infinite product $\pi : C^\omega \to C$ by a unary operation ${}^\omega : C \to C$ computing ω-th powers.

Definition 5.1. Let $S := (\omega^{<\omega})^{<\omega}$. A *power-hyperclone* is a structure of the form

$$\mathfrak{C} = \big\langle (C_{\bar{u}})_{\bar{u} \in S},\ \oplus,\ \cdot,\ (\,:_{I,\sigma})_{I,\sigma},\ 0,\ (\lambda_\sigma)_{\sigma \in \omega^{<\omega}},\ \leq,\ {}^\omega \big\rangle$$

that is obtained from an ω-hyperclone by replacing the infinite product π by a (family of) unary operations

$${}^\omega : C_{\bar{u}} \to C_{\langle \varnothing, \dots, \varnothing \rangle}\,, \quad \text{for } \bar{u} \in S \text{ such that } u_i \subseteq [|\bar{u}|] \text{ for all } i,$$

such that, for every $0 < n < \omega$ and all suitable a, b,

$$(a^n)^\omega = a^\omega\,, \quad (ab)^\omega = a(ba)^\omega\,,$$
$$\text{and} \quad a \leq b \quad \text{implies} \quad a^\omega \leq b^\omega\,.$$

Of course, we can associate with every ω-hyperclone a power-hyperclone by setting $a^\omega := \pi(a, a, a, \dots)$. The main result of this section states that, conversely, every (finitary, path-continuous) power-hyperclone can be turned into an ω-hyperclone.

Definition 5.2. An infinite product π is *compatible* with a power-hyperclone $\mathfrak{C}$ if $\langle \mathfrak{C}, \pi \rangle$ is an ω-hyperclone satisfying

$$\pi(a, a, a, \dots) = a^\omega\,, \quad \text{for all } a \in C\,.$$

First, we deal with an easy special case. If we have a sequence $a^\square$ with $a^n \in C_{\langle 0 \rangle}$, for every n, we can use the Theorem of Ramsey to define the infinite product $\pi(a^\square)$, as in the case of ω-semigroups.

Definition 5.3. Let $\mathfrak{C}$ be a finitary power-hyperclone.

(a) For $a^0, a^1, \dots \in C_{\langle 0 \rangle}$, we define

$$\pi(a^0, a^1, \dots) := b \cdot c^\omega\,,$$

where $b, c \in C_{\langle 0 \rangle}$ are chosen such that there are indices $0 < k_0 < k_1 < \dots$ with

$$b = a^0 \cdot \dots \cdot a^{k_0 - 1} \quad \text{and} \quad c = a^{k_i} \cdot \dots \cdot a^{k_{i+1}-1}\,, \quad \text{for all } i < \omega\,.$$

(Note that such elements exists by the Theorem of Ramsey. Furthermore, the axioms of a power-hyperclone ensure that the product $b \cdot c^\omega$ is uniquely determined by these conditions, even if b and c are not.)

(b) The *trace semigroup* of $\mathfrak{C}$ is the semigroup $\langle S, S_\omega \rangle$ where

$$S := \wp(C_{\langle 0 \rangle}) \quad \text{and} \quad S_\omega := \wp(C_{\langle\rangle})\,.$$

We define multiplication as follows. For $P, Q \in S$ and $U \in S_\omega$, we set

$$P \cdot Q := \{\, st \mid s \in P,\ t \in Q \,\}\,,$$
$$P \cdot U := \{\, su \mid s \in P,\ u \in U \,\}\,.$$

We define the infinite product of $P_0, P_1, \dots \in S$ by

$$\pi(P_0, P_1, \dots) := \{\, \pi(a_0, a_1, \dots) \mid a_n \in P_n \,\}\,,$$

where the product $\pi(a_0, a_1, \dots)$ is defined as in (a).

To extend this definition of π to arbitrary sequences, we employ the machinery of Section 3. With every sequence $a^\square$ we associate an additive labelling λ of its branch tree $\Lambda(a^\square)$. Choosing a suitable contraction λ^P_σ we can construct elements b and c such that we can set

$$\pi(a^\square) := b \cdot c^\omega\,.$$

Definition 5.4. Let $\mathfrak{C}$ be a power-hyperclone, $\mathfrak{J}$ an ideal of $\mathfrak{C}$, $a^\square$ a sequence in $\mathfrak{C}$, and let $\hat{a}^\square$ be the unravelling of $a^\square$.

(a) The *trace labelling* of $a^\square$ (with respect to $\mathfrak{J}$) is the following additive labelling λ on the branch tree $\Lambda(a^\square)$ of $a^\square$. Let $\mathfrak{S} = \langle S, S_\omega \rangle$ be the trace semigroup of $\mathfrak{C}$ and let $\mu : \Lambda(a^\square) \to C$ be the function from Definition 2.6. For a vertex $x \in \Lambda(a^\square)$ of length $m := |x|$ with l immediate successors $y_0, \dots, y_{l-1}$, we define

$$\lambda(x, y_k) := \left\{ \operatorname{sep}(\operatorname{sep}(\mu(x)) :_{I_k} b^m) \;\middle|\; (b^n)_n \in \mathrm{LC}_J(\hat{a}^\square) \right\} \in S$$

where $I_k := [l] \setminus \{k\}$.

This uniquely determines an additive labelling of $\Lambda(a^\square)$ since, if $x \prec y$ and y is not the immediate successor of x, then

$$\lambda(x, y) = \lambda(z_0, z_1) \cdot \dots \cdot \lambda(z_{m-1}, z_m)\,,$$

where $z_0 \prec \dots \prec z_m$ is the path from x to y.

(b) Let λ^P_σ be a finite contraction of the trace labelling λ with $\lim \lambda^P_\sigma = \lim \lambda$. The *regular factorisation* of $a^\square$ induced by λ^P_σ is the pair b, c defined as follows. Let $e_0 = \langle v_0, u_0 \rangle, \dots, e_{n-1} = \langle v_{n-1}, u_{n-1} \rangle$ be all back-edges of $\mathfrak{C}^P_\sigma(\lambda)$, ordered by their end-vertex u_i from left to right. We decompose $\mathfrak{C}^P_\sigma(\lambda)$ as follows. Set

$$U_* := \{\, x \in C^P_\sigma(\lambda) \mid u_i \npreceq x \text{ for all i} \,\}\,,$$
$$U_i := \{\, x \in C^P_\sigma(\lambda) \mid u_i \preceq x \,\}\,, \quad \text{for } i < n\,.$$

Identifying U_* and U_i with the corresponding subsets of $\Lambda(a^\square)$ we set

$$b := \mu[U_*] \quad \text{and} \quad c := \mu[U_0]\sigma_0 \oplus \dots \oplus \mu[U_{n-1}]\sigma_{n-1}\,,$$

where $\sigma_i : [l_i] \to [n]$ is the function such that $v_{\sigma_i(0)}, \dots, v_{\sigma_i(l_i-1)}$ is an enumeration (in left-right order) of all vertices v_j with $v_j \succeq u_i$.

Note that it follows from Lemma 3.13 that every sequence $a^\square$ in a finitary power-hyperclone does have a regular factorisation. We use the trace labelling and regular factorisations to define products of arbitrary sequence. First, let us record a basic property of the trace labelling.

Lemma 5.5. *Let $\mathfrak{C}$ be a finitary power-hyperclone, $\mathfrak{J}$ an ideal of $\mathfrak{C}$, and π an infinite product that is compatible with $\mathfrak{C}$. Suppose that λ is the trace labelling of a sequence $a^\square$ with respect to $\mathfrak{J}$, and let λ^P_σ be a contraction of λ such that* $\lim \lambda^P_\sigma = \lim \lambda$. *Then*

$$\lim \lambda = \mathrm{BT}_J(a^\square) \quad \textit{and} \quad \lim \lambda^P_\sigma = \mathrm{BT}_J(b, c, c, c, \dots)\,,$$

where $b \cdot c^\omega$ is the regular factorisation of $a^\square$ induced by λ^P_σ.

Proof. For the first equation, let β be a branch of $\Lambda(a^\square)$. By definition of λ, it follows that $\lambda(\beta) = \mathrm{Tr}_J(\beta)$. Hence,

$$\lim \lambda = \{\, \lambda(\beta) \mid \beta \text{ a branch} \,\} = \{\, \mathrm{Tr}_J(\beta) \mid \beta \text{ a branc} \,\} = \mathrm{BT}_J(a^\square)\,.$$

A similar argument works for the second equation. ◻

Since we are interested in obtaining a path-continuous ω-hyperclone, we need a corresponding restriction on power-hyperclones.

Definition 5.6. An ordered power-hyperclone $\mathfrak{C}$ is *path-continuous* if there exists an ideal $\mathfrak{J} \subseteq \mathfrak{C}$ such that

- for every sequence $a^\square$ in separation normal form with regular factorisation $x \cdot y^\omega$, we have

$$x \cdot y^\omega = \sup \left\{ u \cdot v^\omega \;\middle|\; u \cdot v^\omega \text{ is a regular factorisation of some sequence } b^\square \le a^\square \text{ in } \mathfrak{J} \right\},$$

- and the value of an expression $a \cdot b^\omega$ with $a, b \in \mathfrak{J}$ is uniquely determined by the set $\mathrm{BT}_{\mathfrak{J}}(a, b, b, b, \dots)$.

It follows from the next lemma that the power-hyperclone $\langle \mathfrak{C}, {}^\omega \rangle$ associated with a path-continuous ω-hyperclone $\langle \mathfrak{C}, \pi \rangle$ by setting $a^\omega := \pi(a, a, a, \dots)$ is path-continuous. We shall prove below that, conversely, every finitary, path-continuous power-hyperclone can be turned into a path-continuous ω-hyperclone in a unique way. We start by showing that there exists at most one infinitary product compatible with a given power-hyperclone.

Lemma 5.7. *Let $\mathfrak{C}$ be a finitary power-hyperclone, $\mathfrak{J}$ an ideal of $\mathfrak{C}$, and let π be an infinite product that is compatible with $\mathfrak{C}$ and such that $\langle \mathfrak{C}, \pi \rangle$ is path-continuous with respect to the ideal $\mathfrak{J}$. For every sequence $a^\square$ in $\mathfrak{J}$, we have*

$$\pi(a^\square) = b \cdot c^\omega\,,$$

where $b \cdot c^\omega$ is any regular factorisation of $a^\square$.

Proof. By path-continuity there is a function p such that

$$\pi(a^\square) = p(\mathrm{BT}_J(a^\square))\,, \quad \text{for every sequence } a^\square \text{ in } J\,.$$

Let λ^P_σ be a contraction of the trace labelling λ with $\lim \lambda^P_\sigma = \lim \lambda$ and let $b \cdot c^\omega$ be the regular factorisation of $a^\square$ induced by λ^P_σ. By Lemma 5.5 it follows that

$$\begin{aligned}\pi(a^\square) = p(\mathrm{BT}_J(a^\square)) = p(\lim \lambda) &\\ &= p(\lim \lambda^P_\sigma)\\ &= p(\mathrm{BT}_J(b, c, c, \dots)) = \pi(b, c, c, \dots) = b \cdot c^\omega,\end{aligned}$$

as desired. ☐

Corollary 5.8. *Let $\mathfrak{C}$ be a finitary power-hyperclone and $\mathfrak{J}$ an ideal of $\mathfrak{C}$. There exists at most one infinite product π that is compatible with $\mathfrak{C}$ and such that $\langle \mathfrak{C}, \pi\rangle$ is path-continuous with respect to the ideal $\mathfrak{J}$.*

Proof. Let π, π' be two infinite products that are compatible with $\mathfrak{C}$ and such that $\langle \mathfrak{C}, \pi\rangle$ and $\langle \mathfrak{C}, \pi'\rangle$ are path-continuous. It is sufficient to prove that π and π' agree on every sequence $a^\square$ in separation normal form in $\mathfrak{J}$. But this follows from the preceding lemma, since

$$\pi(a^\square) = b \cdot c^\omega = \pi'(a^\square),$$

for any regular factorisation $b \cdot c^\omega$ induced by a contraction of the trace labelling λ. ☐

For path-continuous power-hyperclones, it follows that the function mapping a sequence $a^\square$ to the value $b \cdot c^\omega$ of its regular factorisation is well-defined.

Lemma 5.9. *Let $\mathfrak{C}$ be a path-continuous power-hyperclone, and $b_0 \cdot c_0^\omega$, $b_1 \cdot c_1^\omega$ two regular factorisations of a sequence $a^\square$. Then*

$$b_0 \cdot c_0^\omega = b_1 \cdot c_1^\omega.$$

Proof. By path-continuity, we have

$$\begin{aligned} b_0 \cdot c_0^\omega &= \sup \{\, u \cdot v^\omega \mid u \cdot v^\omega \text{ is a reg. fact. of some } a_0^\square \leq a^\square \text{ in } \mathfrak{J} \,\}\\ &= b_1 \cdot c_1^\omega.\end{aligned}$$ ☐

Existence of π is more involved. First, we show that, if a compatible product exists, the resulting ω-hyperclone is path-continuous.

Lemma 5.10. *Let $\mathfrak{C}$ be a finitary path-continuous power-hyperclone with ideal $\mathfrak{J}$ and let π be an infinite product that is compatible with $\mathfrak{C}$. Then $\langle \mathfrak{C}, \pi\rangle$ is path-continuous with the same ideal $\mathfrak{J}$.*

Proof. Let $a^\square$ be a sequence in separation normal form. We have to show that

$$\pi(a^\square) = \sup \{\, \pi(b^\square) \mid b^\square \leq a^\square \text{ in } J \,\}.$$

Let $x \cdot y^\omega$ be the regular factorisation of $a^\square$. Since $\mathfrak{C}$ is path-continuous, we have

$$\begin{aligned} x \cdot y^\omega = \sup \Big\{\, u \cdot v^\omega \;\Big|\; & u \cdot v^\omega \text{ is a regular factorisation of some sequence}\\ & b^\square \leq a^\square \text{ in } J \,\Big\}.\end{aligned}$$

By Lemma 5.7, it follows that

$$\begin{aligned}\pi(a^\square) &= x \cdot y^\omega\\ &= \sup \Big\{\, u \cdot v^\omega \;\Big|\; u \cdot v^\omega \text{ a reg. fact. of } b^\square \leq a^\square \text{ in } J \,\Big\}\\ &= \sup \Big\{\, \pi(b^\square) \;\Big|\; b^\square \leq a^\square \text{ in } J \,\Big\},\end{aligned}$$

as desired.

For the second condition, we have to show that the value of an infinite product $\pi(a^\square)$ with $a^\square$ in $\mathfrak{J}$ only depends on $\mathrm{BT}_J(a^\square)$. Consider two sequence $a^\square$ and $b^\square$ in $\mathfrak{J}$ with $\mathrm{BT}_J(a^\square) = \mathrm{BT}_J(b^\square)$. We claim that $\pi(a^\square) = \pi(b^\square)$. Let $u \cdot v^\omega$ and $x \cdot y^\omega$ be regular factorisations of, respectively, $a^\square$ and $b^\square$ that are induced by contractions of the respective trace labellings. Then it follows by Lemma 5.5 that

$$\mathrm{BT}_J(u, v, v, \dots) = \mathrm{BT}_J(a^\square) = \mathrm{BT}_J(b^\square) = \mathrm{BT}_J(x, y, y, \dots).$$

By path-continuity of $\mathfrak{C}$ this implies that $u \cdot v^\omega = x \cdot y^\omega$. Hence,

$$\pi(a^\square) = u \cdot v^\omega = x \cdot y^\omega = \pi(b^\square),$$

as desired. ☐

Finally, we can prove the existence of compatible products. Together with their uniqueness, it follows that every finitary, path-continuous power-hyperclone can be extended in a unique way to an ω-hyperclone.

Theorem 5.11. *Let $\mathfrak{C}$ be a finitary, path-continuous power-hyperclone with ideal $\mathfrak{J}$. There exists exactly one infinite product π that is compatible with $\mathfrak{C}$ and such that $\langle \mathfrak{C}, \pi\rangle$ is path-continuous with respect to $\mathfrak{J}$.*

Proof. By Corollary 5.8, we only need to prove the existence of π. According to Lemma 5.7, if $a^\square$ is a sequence in $\mathfrak{J}$ in separation normal form, we have no choice but to define

$$\pi(a^\square) := b \cdot c^\omega,$$

where $b \cdot c^\omega$ is a regular factorisation induced by some contraction λ^P_σ of the trace labelling λ. In fact, we use this definition for all sequences $a^\square$ in separation normal form. Note that, by Lemma 5.9, this definition does not depend on the regular factorisations we choose.

We extend the definition to sequences $a^\square$ not in separation normal form by setting

$$\pi(a^\square) := \pi(b^\square), \quad \text{where } b^\square \text{ is the unravelling of } a^\square.$$

It remains to prove that π turns $\mathfrak{C}$ into a path-continuous ω-hyperclone. Path-continuity follows from Lemma 5.10. Hence, we only have to check that $\langle \mathfrak{C}, \pi \rangle$ is an ω-hyperclone and that π is compatible with $\mathfrak{C}$.

For compatibility, consider a constant sequence $a^\square = (a, a, a, a, \dots)$. The function $\sigma : \Lambda(a^\square) \to [1]$ mapping every element to 0 is a Ramseyan split of the trace labelling λ. It follows that $a \cdot a^\omega$ is a regular factorisation of $a^\square$. Consequently,

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

It remains to verify that π satisfies the axioms of an ω-hyperclone. According to [1], there are five axioms to check:

(11) For all $k_0 < k_1 < k_2 < \dots < \omega$, we have the associative laws

$$\pi(a^0, a^1, a^2, \dots) = a^0 \cdot \pi(a^1, a^2, \dots),$$
$$\pi(a^0, a^1, a^2, \dots) = \pi((a^0 \cdots a^{k_0 - 1}), (a^{k_0} \cdots a^{k_1 - 1}), (a^{k_1} \cdots a^{k_2 - 1}), \dots).$$

(12) Let $b^\square$ be the factorisation (see below) of $a^\square$ induced by some $H \subseteq \Lambda(a^\square)$. Then

$$\pi(b^\square) = \pi(a^\square).$$

(13) Let $m(n) := |\bar{u}^n|$ be the width of a^n and set $b^n := \pi(a^n, a^{n+1}, \dots)$. For all $I_n \subseteq [m(n)]$, we have

$$\pi(a^0, a^1, \dots) = \pi(a^0 :_{I_1} b^1,\ a^1 :_{I_2} b^2,\ a^2 :_{I_3} b^3, \dots).$$

(14) If $\sigma_n : [|\bar{u}^n|] \to [|\bar{u}^n|]$, $n < \omega$, are functions such that

$$\sigma_n \pi(a^n, a^{n+1}, \dots) = \pi(a^n, a^{n+1}, \dots),$$

then

$$\pi(\sigma_0 a^0, \sigma_1 a^1, \sigma_2 a^2, \dots) = \pi(a^0, a^1, a^2, \dots).$$

(15) For all sequences $a^\square$ and $b^\square$,

$$a^n \leq b^n, \text{ for all } n < \omega, \quad \text{implies} \quad \pi(a^0, a^1, \dots) \leq \pi(b^0, b^1, \dots).$$

We start with a proof of (12). Let us first recall the definition of a factorisation from [1]. A set $H \subseteq \Lambda(a^\square)$ induces a factorisation of $a^\square$ if, for every $v \in \Lambda(a^\square)$, there is some $u \preceq v$ in H. For each $u \in H$, we set

$$T_u := \{ v \in \Lambda(a^\square) \mid u \preceq v \text{ and there is no } w \in H \text{ with } u \prec w \preceq v \}.$$

Then H induces a factorisation of $a^\square$ if, and only if, the sets T_u form a partition of $\Lambda(a^\square)$. By definition, this factorisation is obtained from $a^\square$ by replacing all elements in T_u by the single element $\mu[T_u]$, where μ is the function from Definition 2.6.

Returning to the proof, let $b^\square$ be the factorisation of $a^\square$ induced by $H \subseteq \Lambda(a^\square)$. We have to show that $\pi(b^\square) = \pi(a^\square)$.

First, let us consider the case that $a^\square$ and $b^\square$ are in separation normal form. Let λ and λ_0 be the trace labellings of, respectively, $a^\square$ and $b^\square$ and fix a Ramseyan split σ_0 for λ_0. Note that λ_0 is the restriction of λ to H. Therefore, we can extend σ_0 to a Ramseyan split σ for λ in the following way. For each vertex $u \in H$, we fix a Ramseyan split σ_u for the restriction of λ to the subtree T_u defined above. Let $n < \omega$ be some number such that $\operatorname{rng} \sigma_0 \subseteq [n]$. We define σ by

$$\sigma(x) := \begin{cases} \sigma_0(x) & \text{if } x \in H, \\ n + \sigma_u(x) & \text{if } x \notin H \text{ and } x \in T_u, \end{cases}$$

Then σ is a Ramseyan split of $\Lambda(a^\square)$ whose restriction to H coincides with σ_0.

By the definition of σ, the canonical embedding $i : H \to \Lambda(a^\square)$ induces an embedding of every contraction $(\lambda_0)^{P_0}_{\sigma_0}$ of λ_0 into the corresponding contraction λ^P_σ of λ. It follows that the regular factorisations induced by $(\lambda_0)^{P_0}_{\sigma_0}$ and λ^P_σ coincide. Therefore, the products $\pi(a^\square)$ and $\pi(b^\square)$ have the same value.

It remains to prove the claim for arbitrary sequences $a^\square$ and $b^\square$. Let $\hat{a}^\square$ and $\hat{b}^\square$ be the unravellings of $a^\square$ and $b^\square$, respectively. By definition, we have $\pi(a^\square) = \pi(\hat{a}^\square)$ and $\pi(b^\square) = \pi(\hat{b}^\square)$. Furthermore, $\hat{b}^\square$ is a factorisation of $\hat{a}^\square$ induced by some set $\hat{H} \subseteq \Lambda(\hat{a}^\square)$. Since $\hat{a}^\square$ and $\hat{b}^\square$ are in separation normal form, it follows by what we have shown above that

$$\pi(b^\square) = \pi(\hat{b}^\square) = \pi(\hat{a}^\square) = \pi(a^\square) .$$

(11) The second equation follows from (12). Hence, it remains to prove that

$$\pi(a^0, a^1, \dots) = a^0 \cdot \pi(a^1, a^2, \dots) , \quad \text{for every sequence } a^\square .$$

By definition of π, it is sufficient to consider sequences $a^\square$ in separation normal form. Let $a^\square$ be such a sequence and let $a^1 = a^1_0 \oplus \dots \oplus a^1_{m-1}$ be the decomposition of the element a^1. We denote by $b^\square_i$ the sequence $(a^1_i, a^2, a^3, \dots)$. Then

$$\pi(a^1, a^2, \dots) = \pi(b^\square_0) \oplus \dots \oplus \pi(b^\square_{m-1}) .$$

Let λ and λ_i, $i < m$, be the trace labellings of, respectively, $a^\square$ and $b^\square_i$. We fix Ramseyan splits σ_i of λ_i and sets $P_i \subseteq \Lambda(b^\square_i)$ such that $\lim(\lambda_i)^{P_i}_{\sigma_i} = \lim \lambda_i$. Let $x_i \cdot y_i^\omega$ be the corresponding regular factorisation of $b^\square_i$. Then

$$\begin{aligned} a_0 \cdot \pi(a^1, a^2, \dots) &= a_0 \cdot (\pi(b^\square_0) \oplus \dots \oplus \pi(b^\square_{m-1})) \\ &= a_0 \cdot (x_0 \cdot y_0^\omega \oplus \dots \oplus x_{m-1} \cdot y_{m-1}^\omega) . \end{aligned}$$

We denote by $h_i : \Lambda(b^\square_i) \to \Lambda(a^\square)$ the tree embedding mapping each position of $b^\square_i$ to the corresponding one of $a^\square$. Let $P \subseteq \Lambda(a^\square)$ be the prefix-closure of the set

$$h_0[P_0] \cup \dots \cup h_{m-1}[P_{m-1}] .$$

We define a Ramseyan split σ of λ by setting

$$\sigma(w) := \begin{cases} \sigma_i(u) & \text{if } w = h_i(u) \text{ for some } u \text{ and } i , \\ |w| + k & \text{otherwise} , \end{cases}$$

where k is a constant larger than any number in $\operatorname{rng} \sigma_0 \cup \dots \cup \operatorname{rng} \sigma_{m-1}$. It follows that the contraction λ^P_σ of λ induced by P and σ consists of the prefix of $\Lambda(a^\square)$ corresponding to a^0 together with the contractions $(\lambda_i)^{P_i}_{\sigma_i}$. Therefore, the regular factorisation induced by λ^P_σ is of the form

$$\begin{aligned} &\operatorname{sep}(a^0 \cdot (x_0 \oplus \dots \oplus x_{m-1})) \cdot (y_0 \tau_0 \oplus \dots \oplus y_{m-1} \tau_{m-1})^\omega \\ &= a^0 \cdot (x_0 \cdot y_0^\omega \oplus \dots \oplus x_{m-1} \cdot y_{m-1}^\omega) , \end{aligned}$$

for suitable functions $\tau_0, \dots, \tau_{m-1}$. It follows that

$$\pi(a^\square) = a^0 \cdot (x_0 \cdot y_0^\omega \oplus \dots \oplus x_{m-1} \cdot y_{m-1}^\omega) = a^0 \cdot \pi(a^1, a^2, \dots) .$$

(13) We may again assume that $a^\square$ is in separation normal form. Let $c^\square$ be the sequence $(a^0 :_{I_1} b^1,\ a^1 :_{I_2} b^2,\ a^2 :_{I_3} b^3, \dots)$. Note that $\Lambda(c^\square)$ is a subset of $\Lambda(a^\square)$ obtained by replacing some subtrees by leaves. Let $h : \Lambda(c^\square) \to \Lambda(a^\square)$ be the corresponding embedding. Let λ be the trace labelling of $\Lambda(a^\square)$ and let $u \cdot v^\omega$ be a regular factorisation of $a^\square$ induced by some contraction λ^P_σ of λ. Note that $\lambda_0(x, y) := \lambda(h(x), h(y))$ is the trace labelling of $\Lambda(c^\square)$. Hence, $\sigma_0 := \sigma \circ h$ is a Ramseyan split of λ_0. The corresponding factorisation of $c^\square$ is $u_0 \cdot (v_0)^\omega$ where u_0 and v_0 are of the form

$$u_0 = \operatorname{sep}(u :_I w) \quad \text{and} \quad v_0 = \rho_v(v :_{J,\tau} w) , \quad \text{for suitable } I, J, v, \tau ,$$

where $w := v^\omega \oplus \dots \oplus v^\omega$. It follows that $u_0 \cdot (v_0)^\omega = u \cdot v^\omega$.

(14) can be checked in a similar way. W.l.o.g. let $a^\square$ be in separation normal form and let $b^\square$ be the sequence $(\sigma_0 a^0, \sigma_1 a^1, \sigma_2 a^2, \dots)$. There exists an isomorphism $\pi : \Lambda(a^\square) \to \Lambda(b^\square)$. Using π we can transform a Ramseyan split σ for the trace labelling of $\Lambda(a^\square)$ into a Ramseyan split for $\Lambda(b^\square)$. For the regular factorisations $x \cdot y^\omega$ and $u \cdot v^\omega$ of $a^\square$ and $b^\square$, respectively, it follows that

$$u = x\tau \quad \text{and} \quad v = \tau^{-1} y \tau , \qquad \text{for a suitable bijection } \tau .$$

It follows that $u \cdot v^\omega = x\tau(\tau^{-1} y \tau)^\omega = x \cdot y^\omega$, as desired.

(15) Suppose that $a^\square \leq b^\square$. Let $\hat{a}^\square$ and $\hat{b}^\square$ be the unravellings of $a^\square$ and $b^\square$ and let $u \cdot v^\omega$ and $x \cdot y^\omega$ be regular factorisations of $\hat{a}^\square$ and $\hat{b}^\square$, respectively. Then $\hat{a}^\square \leq \hat{b}^\square$ and, by path-continuity,

$$\begin{aligned} \pi(a^\square) = \pi(\hat{a}^\square) &= u \cdot v^\omega \\ &= \sup \{\, w \cdot z^\omega \mid w \cdot z^\omega \text{ a reg. fact. of some } c^\square \leq \hat{a}^\square \text{ in } \mathfrak{J} \,\} \\ &\leq \sup \{\, w \cdot z^\omega \mid w \cdot z^\omega \text{ a reg. fact. of some } c^\square \leq \hat{b}^\square \text{ in } \mathfrak{J} \,\} \\ &= x \cdot y^\omega = \pi(\hat{b}^\square) = \pi(b^\square) . \end{aligned}$$

□

6 The Theorem of Rabin

In this short section we present an application of our results. We give an alternative proof of the Theorem of Rabin on the decidability of the monadic second-order theory of the binary tree. The reason why we introduced power-hyperclones is that, *a priori,* we need an infinite amount of information to specify the infinite product of an ω-hyperclone. But even after having replaced the infinite product by the unary power operation, we still deal with an infinite structure since there are infinitely many sorts. To actually compute with power-hyperclones we have therefore to choose a finite subset of sorts and to only deal with the reduct of the power-hyperclone to this set of sorts.

Definition 6.1. An *effective presentation* of a power-hyperclone $\mathfrak{C}$ is a computable function g that, given a finite set of sorts $S_0 \subseteq (\omega^{<\omega})^{<\omega}$ computes the reduct $\mathfrak{C}|_{S_0}$.

We have shown in [1] that the class of languages recognised by finitary path-continuous ω-hyperclones is closed under boolean operations and projections.

Theorem 6.2 ([1]). *Let $p : \Sigma \to \Gamma$ be a surjective, arity preserving function between functional signatures. If $L, L' \subseteq F_{\langle\varnothing\rangle}[\Sigma]$ are recognised by path-continuous ω-hyperclones then so are $L \cap L'$, $L \cup L'$, $F_{\langle\varnothing\rangle}[\Sigma] \setminus L$, and $p[L]$.*

The Theorem of Rabin will follow from an effective version of this theorem. To obtain such an effective version, we replace ω-hyperclones by power-hyperclones and we show that the two operations on ω-hyperclones underlying the theorem are effective for power-hyperclones. These two operations are the direct product of two power-hyperclones and the power set operation.

Definition 6.3. Let $\mathfrak{C}$ and $\mathfrak{C}'$ be path-continuous power-hyperclones.

(a) The *product* $\mathfrak{C} \times \mathfrak{C}'$ is the power-hyperclone $\mathfrak{D}$ where the domain of sort $\bar{u}$ is

$$D_{\bar{u}} := C_{\bar{u}} \times C'_{\bar{u}} ,$$

and where all operations are defined component-wise.

(b) We define the power-hyperclone $\mathcal{P}(\mathfrak{C})$ as follows. The domain of sort $\bar{u} = \langle u_0, \dots, u_{m-1}\rangle$ is

$$\wp(C_{u_0}) \times \dots \times \wp(C_{u_{m-1}}) .$$

(For $m = 0$, we take the empty product $\{\langle\rangle\}$.) To simplify notation we will identify an element $a = a_0 \oplus \dots \oplus a_{m-1} \in C_{\bar{u}}$ with the m-tuple $\bar{a} = \langle a_0, \dots, a_{m-1}\rangle$, and we write $\bar{a} \in \bar{A}$, for $\bar{a} \in C_{\bar{u}}$ and $\bar{A} \in \wp(C_{u_0}) \times \dots \times \wp(C_{u_{m-1}})$, if we have $a_i \in A_i$, for all i. For elements $\bar{A}$ and $\bar{B}$ of $\mathcal{P}(\mathfrak{C})$, we define the operations by

$$\begin{aligned}
\bar{A} \oplus \bar{B} &:= \bar{A}\bar{B} ,\\
0 &:= \langle\rangle ,\\
\lambda_\sigma(\bar{A}) &:= \langle A_{\sigma(0)}, \dots, A_{\sigma(m-1)}\rangle ,\\
\bar{A} :_{I,\sigma} \bar{B} &:= \bar{D} , \qquad \text{where } D_i := \{\, a :_{I,\sigma} \bar{b} \mid a \in A_i,\ \bar{b} \in \bar{B} \,\} .
\end{aligned}$$

The power operation is defined as

$$\bar{A}^\omega := \{\, \pi(\bar{a}^\square) \mid \bar{a}^n \in \bar{A}^n_0 , \text{ for all } n \,\} ,$$

where $\bar{A}^\square_0$ is the unravelling of the sequence $\bar{A}, \bar{A}, \bar{A}, \dots$ and π is the unique infinite product compatible with $\mathfrak{C}$. Finally, the ordering is defined by

$$\begin{aligned}
\bar{A} \leq \bar{B} \quad &:\text{iff} \quad \text{there exist injections } \varphi_i : A_i \to B_i \text{ such that}\\
&\qquad\qquad a \leq \varphi_i(a), \text{ for all } a \in A_i .
\end{aligned}$$

Obviously, we can compute an effective representation of $\mathfrak{C} \times \mathfrak{C}'$ from effective representations of $\mathfrak{C}$ and $\mathfrak{C}'$. It is less clear that we can compute effective representations of $\mathcal{P}(\mathfrak{C})$ from those of $\mathfrak{C}$.

Proposition 6.4. *Given an effective representation of a finitary path-continuous power-hyperclone $\mathfrak{C}$, we can compute an effective representation of the power-hyperclone $\mathcal{P}(\mathfrak{C})$.*

Proof. Except for the ω-th power, all operations of $\mathcal{P}(\mathfrak{C})$ are obviously computable. Hence, it remains to show how to compute the value of $\bar{A}^\omega$. Let $\bar{A}^\square_0$ be the unravelling of the sequence $\bar{A}, \bar{A}, \bar{A}, \dots$. Note that $\bar{A}^\omega$ is the set of all values $\bar{u} \cdot \bar{v}^\omega$, where $\bar{u} \cdot \bar{v}^\omega$ is a regular factorisation of some sequence $\bar{a}^\square$ with $\bar{a}^n \in \bar{A}^n_0$. By Theorem 4.6 there is a (computable) bound $k < \omega$ such that we only need to consider regular factorisations that are induced by a contraction $\mathfrak{C}^p_\sigma(\lambda)$ of the trace labelling of $\bar{a}^\square$ such that $\mathfrak{C}^p_\sigma(\lambda)$ is contained in the first k elements of the sequence $\bar{A}^\square_0$. Since there are only finitely many such contractions, we can enumerate all of them to compute $\bar{A}^\omega$. ☐

Theorem 6.5. *For every* MSO*-sentence φ, one can compute an effective representation of a finitary, path-continuous power-hyperclone $\mathfrak{C}$ and a set $P \subseteq C_{\langle\rangle}$ such that there exists a morphism $h : \mathfrak{F}_\omega[\Sigma] \to \mathfrak{C}$ with*

$$\mathfrak{T} \vDash \varphi \quad \text{iff} \quad h(\mathfrak{T}) \in P .$$

Proof. We compute the representation of $\mathfrak{C}$ and the set P, by induction on φ. The power-hyperclones for atomic formulae given in Theorem 6.9 of [1] are computable from φ. For boolean operations, we perform a direct product of the given power-hyperclones, which is obviously effective. For existential quantifiers, we use Proposition 6.4 to compute $\mathcal{P}(\mathfrak{C})$. In each case, the corresponding description of the set P given in [1] is clearly effective. □

The Theorem of Rabin will follow from this result, once we have shown that one can compute the image of the morphism h.

Lemma 6.6. *Given an effective representation of a finitary, path-continuous power-hyperclone $\mathfrak{C}$, a finite set $X \subseteq C$ and an element $a \in C_{\langle\rangle}$, one can decide whether or not the element a is contained in the sub-power-hyperclone $\langle X \rangle$ of $\mathfrak{C}$ generated by X.*

Proof. Let π be the infinite product associated with $\mathfrak{C}$. Let X^+ be the minimal set containing X such that $a \oplus b \in X^+$ implies $a, b \in X^+$, and $a \in X^+$ implies $\mathrm{sep}(a) \in X^+$. Note that X^+ is still finite.

Then $a \in \langle X \rangle$ if, and only if, $a = \pi(b^\square)$, for some sequence $b^\square$ where each element is of the form $b^n = x^n_0 \oplus \cdots \oplus x^n_{m(n)-1}$ with $x^n_i \in X^+$. By Theorem 4.4, there is a computable constant N such that each $b^\square$ has a regular factorisation $u \cdot v^\omega$ where u and v are of the form

$$u = u^0 \cdots u^k \quad \text{and} \quad v = v^0 \cdots v^l$$

with $k, l < N$ where

$$u_i = x^i_0 \oplus \cdots \oplus x^i_{m(i)-1} \quad \text{and} \quad v_i = y^i_0 \oplus \cdots \oplus y^i_{n(i)-1} \quad \text{for } x^i_j, y^i_j \in X^+\,.$$

Since $\pi(b^\square) = u \cdot v^\omega$ it is therefore sufficient to check these finitely many regular factorisations to determine whether or not $a \in \langle X \rangle$. □

Corollary 6.7. *It is decidable whether a given* MSO*-formula* φ *has a tree model.*

Proof. Given an MSO-sentence φ, we construct the corresponding power-hyperclone $\mathfrak{C}$ and the set $P \subseteq C_{\langle\rangle}$ as in the theorem. Let $h : \mathfrak{F}_\omega[\Sigma] \to \mathfrak{C}$ be the corresponding morphism. Then

$$\varphi \text{ has a tree model} \quad \text{iff} \quad h^{-1}[P] \neq \varnothing \quad \text{iff} \quad P \cap \mathrm{rng}\, h \neq \varnothing\,.$$

Since $\mathfrak{F}_\omega[\Sigma]$ is finitely generated (by the elements of Σ), it follows by Lemma 6.6 that the latter condition is decidable. □

Corollary 6.8 (Rabin)**.** *The* MSO*-theory of the full binary tree* $\mathfrak{T}_2 := \langle 2^{<\omega}, E_0, E_1 \rangle$ *is decidable.*

Proof. The binary tree $\mathfrak{T}_2$ can be axiomatised by a single MSO-sentence θ. It follows that

$$\mathfrak{T}_2 \vDash \varphi \quad \text{iff} \quad \theta \wedge \varphi \text{ has a tree model.}$$

The latter question is decidable by the preceding corollary. □

References

[1] A. Blumensath, *Recognisability for algebras of infinite trees*, Theoretical Computer Science, 412 (2011), pp. 3463–3486.

[2] T. Colcombet, *On Factorization Forests*, Tech. Rep. hal-00125047, Irisa Rennes, 2007.

[3] E. Grädel, W. Thomas, and T. Wilke, *Automata, Logic, and Infinite Games*, LNCS 2500, Springer-Verlag, 2002.

[4] D. E. Muller and P. E. Schupp, *Simulating alternating tree automata by nondeterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra*, Theoretical Computer Science, 141 (1995), pp. 69–107.

[5] S. Shelah, *The Monadic Second Order Theory of Order*, Annals of Mathematics, 102 (1975), pp. 379–419.

[6] W. Thomas, *Ehrenfeucht Games, the Composition Method, and the Monadic Theory of Ordinal Words*, LNCS, 1261 (1997), pp. 118–143.