

ON FINAL COALGEBRAS OF POWER-SET FUNCTORS AND SATURATED TREES

JIŘÍ ADÁMEK, PAUL LEVY, STEFAN MILIUS, LAWRENCE S. MOSS AND LURDES
SOUSA

To George Janelidze on the occasion of his sixtieth birthday

ABSTRACT: The final coalgebra for the finite power-set functor was described by Worrell who also proved that the final chain converges in $\omega + \omega$ steps. We describe the step ω as the set of saturated trees, a concept equivalent to the modally saturated trees introduced by K. Fine in the 1970s in his study of modal logic. And for the bounded power-set functors \mathcal{P}_λ , where λ is an infinite regular cardinal, we prove that the construction needs precisely $\lambda + \omega$ steps. We also generalize Worrell’s result to M -labeled trees for a commutative monoid M , yielding a final coalgebra for the corresponding functor \mathcal{M}_f studied by H.-P. Gumm and T. Schröder. We describe the final chain of the power-set functor by introducing the concept of i -saturated tree for all ordinals i , and then prove that for i of cofinality ω , the i -th step in the final chain consists of all i -saturated, strongly extensional trees.

1. Introduction

Coalgebras for power-set functors, e.g. the full one \mathcal{P} or the finite power-set functor \mathcal{P}_f , are of major importance in modal logic (Kripke structures), set theory (non-wellfounded sets) and process algebra. Final coalgebras serve, in general, as Rutten’s fundamental study [27] demonstrated, as a basis for analyses of numerous systems. Although we know that \mathcal{P} does not have a final coalgebra, it is of interest to describe the steps \mathcal{P}^i of its final chain. For that we introduce the concept of an i -saturated tree (for every ordinal i) and prove that the case $i = \omega$ is nothing else than the classical concept of modally saturated tree due to K. Fine [15]. We then describe the final chain \mathcal{P}^i of \mathcal{P} as the set of all i -saturated, strongly extensional trees. For the finite power-set functor \mathcal{P}_f two beautiful descriptions of the final coalgebra exist: as the set of all hereditarily finite sets in the non-wellfounded set theory due to P. Aczel [2] and as the set of all strongly extensional, finitely

Received October 24, 2012. Revised August 12, 2014.

The 2nd author is supported by EPSRC Advanced Research Fellowship EP/E056091/1.

The 4th author was partially supported by Simons Foundation grant #245591.

The 5th author is partially supported by CMUC/FCT (Portugal) and the FCT Grant PTDC/MAT/120222/2010.

branching trees* due to J. Worrell [30]. He used metric spaces and described the limit $\mathcal{P}_f^\omega 1$ of the final chain of \mathcal{P}_f as the set of all strongly extensional, compactly branching trees. From that he derived the above description of the final coalgebra. We give below new descriptions that do not need topology, one combinatorial and one using modal logic. We prove that the limit $\mathcal{P}_f^\omega 1$ consists of all saturated, strongly extensional trees, a concept we introduce and prove to be equivalent to modally saturated trees. Another description of $\mathcal{P}_f^\omega 1$ we present is as the set of all maximal consistent theories of the modal logic \mathbf{K} . A related description of the final coalgebra of \mathcal{P}_f is as the set of all hereditarily finite maximal consistent theories in \mathbf{K} . Other descriptions were previously given by S. Abramsky [1], A. Kurz and D. Pattinson [20] and by J. Rutten [26, Theorem 7.4].

We also present a generalization in two directions: one uses finite multisets with multiplicities drawn from a given commutative monoid M , as introduced by H.-P. Gumm and T. Schröder [17]. Form the functor \mathcal{M}_f of all such finite multisets; its coalgebras are labeled transition systems with actions labeled by $M \setminus \{0\}$. We prove a direct generalization for all monoids for which \mathcal{M}_f preserves weak pullbacks: the final coalgebra for \mathcal{M}_f consists of all finitely branching, strongly extensional M -labeled trees. For general monoids this result is not true, but we prove that the final coalgebra for \mathcal{M}_f is the coalgebra of finitely branching M -labeled trees modulo an equivalence generalizing M. Barr's equivalence for \mathcal{P}_f , see [10].

The other direction of generalization concerns moving from \mathcal{P}_f to \mathcal{P}_λ , the functor of all subsets of power less than λ . (Here and below, λ is an infinite cardinal.) This functor \mathcal{P}_λ has the final coalgebra consisting of all strongly extensional λ -branching trees, as proved by D. Schwencke [28]. We present a new and much simpler proof. We also prove that the final chain converges precisely at $\lambda + \omega$ if λ is a regular cardinal, else it converges at the cardinal successor of λ .

Returning to the power-set functor \mathcal{P} we present a partial description of the final chain $\mathcal{P}^i 1$. We describe $\mathcal{P}^i 1$ as the set of all strongly extensional

*Throughout the paper trees are directed graphs with a distinguished node called the root from which every other node can be reached by a unique directed path, and they are always considered up to isomorphism. Strong extensionality for trees is recalled in Section 2.2 below.

i -saturated trees, provided i has cofinality ω . We do not have a description of $\mathcal{P}^i 1$ for general ordinals i .

Our related paper. This is a substantially extended version of the paper [7] presented at the conference Computer Science Logic (CSL 2011). Besides containing all full proofs, the main new results in the present version are Theorems 3.17 on the final chain of the power-set functor, and Theorem 4.4 describing the final coalgebra of \mathcal{P}_λ . The claim made in [7] that i -saturated trees form the i -th step of the final chain even for uncountable ordinals is withdrawn: we have found a mistake in the proof of Theorem 5.11 of that paper, and we no longer believe the theorem is true. In addition, we have removed mistaken claims about the final chain of the countable power-set functor \mathcal{P}_c .

2. Background on graphs and trees, and on the final chain of \mathcal{P}

This section presents background on graphs and trees. Although the real work of the paper begins in the next section with our results on the final coalgebra of the finite power set functor \mathcal{P}_f , some of our work in this section is new.

2.1. Graphs. By a *graph* in this paper, we mean a coalgebra of the power set functor \mathcal{P} on **Set**. So a graph is a structure $G = (A, e)$, where $e : A \rightarrow \mathcal{P}(A)$. One recovers graphs in the ordinary sense (a set with an arbitrary relation on it), by taking the set of children of a node x to be $e(x)$. We work with graphs as coalgebras because the notion of morphism that is relevant in this paper is that of *coalgebra morphism* rather than morphism of relational structures. In plainer terms, coalgebra homomorphisms $f : A \rightarrow B$ are the maps f which preserve edges, and for every edge from $f(a)$ to b in B there exists an edge from a to some a' in A with $b = f(a')$.

Definition 2.1. Let G and H be graphs. A *bisimulation* between G and H is a relation R between the vertices of G and those of H such that if $x R y$, then every child of x is related by R to some child of y ; and vice-versa.

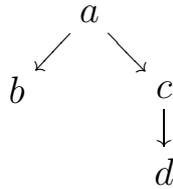
Remark 2.2. A bisimulation between a graph $e : G \rightarrow \mathcal{P}(G)$ and itself is called a bisimulation *on* G . A graph G is called *strongly extensional* if distinct nodes are not bisimilar, i.e., for any bisimulation $R \subseteq G \times G$ we have $R \subseteq \Delta_G$, where Δ_G is the diagonal relation on G . It follows from Aczel [2] that every

G has a largest bisimulation R , and R is an equivalence relation. The set G/R of equivalence classes hosts a graph structure $f : G/R \rightarrow \mathcal{P}(G/R)$ by setting

$$f([x]) = \{[y] : \text{for some } x' \in [x] \text{ there is some } y' \in [y] \text{ such that } y' \in e(x')\}.$$

G/R is called the *strongly extensional quotient* of G .

Example 2.3. Consider the graph G shown below:



The relation $\{(b, d)\}$ is a bisimulation. The largest bisimulation is $\Delta \cup \{(b, d), (d, b)\}$.

2.2. Trees. In this paper, a *tree* is a graph with a distinguished vertex called the *root*, and with the additional property that for each point x in the graph, there is a unique (finite) path from the root to x . So our trees are unordered.

Given a tree t , the subtree of t rooted at the node x is denoted by t_x .

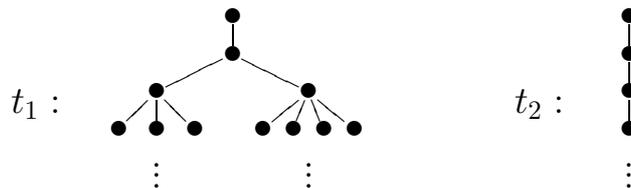
A *morphism of trees* is a coalgebra morphism which preserves the roots. Morphisms of trees preserve all distances to the root. Also, the requirement that every point is reachable from the root implies that every morphism of trees is surjective.

Definition 2.4 (J. Worrell [30]). For trees t and s a *tree bisimulation* is a graph bisimulation $R \subseteq t \times s$ such that the roots are related, and two related child nodes are either both roots or both children with related parents.

Please note that a tree bisimulation is not the same thing as a graph bisimulation on trees, due to the requirement that two related child nodes have related parents.

Henceforth, we always identify isomorphic trees.

Example 2.5. The following trees are tree-bisimilar, where t_1 has, by breadth-first search, n children of the n -th node:



The following relation is a tree bisimulation: relate all nodes of a given level on the left with the node of the same level on the right.

We might also note that for all nodes x of the tree on the right, $t_x = t$.

Definition 2.6 (J. Worrell [30]). A tree t is called *strongly extensional* if distinct children of any node are not tree bisimilar. Equivalently, every tree bisimulation $R \subseteq t \times t$ satisfies $R \subseteq \Delta_t$.

Example 2.7. The tree in Example 2.3 is strongly extensional. (However, as a graph, G is not strongly extensional.) The infinite path t_2 in Example 2.5 is a strongly extensional tree. This is the only strongly extensional tree without leaves because for every tree t without leaves the relation

$$x R y \quad \text{iff } x \text{ and } y \text{ have the same depth}$$

is a tree bisimulation.

Remark 2.8. The terminology “strongly extensional” stems from calling a tree *extensional* if two different children of any node yield different subtrees. Observe that this is, indeed, a weaker condition than strong extensionality (since the relation “yield the same subtree” is a tree bisimulation). For trees of finite depth, the two notions are clearly equivalent.

Remark 2.9. It is easy to see that tree bisimulations are total and closed under composition, unions and opposite relations. Consequently:

(a) For every tree t there is a largest tree bisimulation $R \subseteq t \times t$. It is an equivalence relation. The corresponding quotient graph t/R (see Remark 2.2), is called *the strongly extensional quotient* of t . It is the least tree quotient of t .

(b) Given two trees t and u , there is a tree bisimulation from t to u iff t and u have the same strongly extensional quotient. Consequently if two strongly extensional trees are tree bisimilar they are equal (up to isomorphism).

2.3. Tree expansions of pointed graphs. A *pointed graph* is a graph with a designated vertex. Let Gra_p be the category of pointed graphs and coalgebra morphisms which preserve the designated vertex. Let Tree be the full subcategory of trees.

Notation 2.10. Let P be the “paths functor”

$$P : \text{Gra}_p \rightarrow \text{Tree}$$

taking a pointed graph G to the tree $\mathbf{P}(G)$ of all finite (non-empty) sequences \mathbf{u} which begin with the distinguished vertex and follow the edge relation in G . The children of a vertex \mathbf{u} in $\mathbf{P}(G)$ are the one-edge extensions of it. The functoriality is easy to check; however we never use this fact. $\mathbf{P}(G)$ is said to be the *tree expansion* of the pointed graph G .

Let G be a pointed graph, let g be the distinguished vertex, and let $t = \mathbf{P}(G)$. Let $\text{last} : t \rightarrow G$ take a finite non-empty sequence to its last element. Thus, last is a function. Indeed, t and G are coalgebras of the power set functor, and last is a coalgebra morphism. However, we also consider last (or rather its graph) as a relation between G and t .

Proposition 2.11. *The following hold for every pointed graph G :*

- (1) *As a relation, last is a graph bisimulation.*
- (2) *If \equiv is a tree bisimulation on $\mathbf{P}(G)$, then*

$$\{(\text{last}(\mathbf{u}), \text{last}(\mathbf{v})) : \mathbf{u} \equiv \mathbf{v}\}$$

is a graph bisimulation on G .

- (3) *If G is strongly extensional, then $t = \mathbf{P}(G)$ is a strongly extensional tree.*
- (4) *The strongly extensional quotient of $\mathbf{P}(G)$ as a graph is (isomorphic to) the strongly extensional quotient of G .*

Proof: Write t for $\mathbf{P}(G)$, and let g be the distinguished vertex of G . The first two parts are straightforward. For the third, let \equiv be a tree bisimulation on t . We show by induction on n that if \mathbf{u} and \mathbf{v} are sequences of length n and $\mathbf{u} \equiv \mathbf{v}$, then $\mathbf{u} = \mathbf{v}$. For $n = 1$, \mathbf{u} and \mathbf{v} must be the one-point sequence g . Assuming our result for n , let \mathbf{u} and \mathbf{v} be sequences of length $n + 1$. Since \equiv is a tree bisimulation, the first n terms in \mathbf{u} and \mathbf{v} are identical. And by part (2) above, the last terms are also the same. The last part of this result is also easy. ■

Corollary 2.12. *A tree t is strongly extensional iff there is a strongly extensional pointed graph G with $t = \mathbf{P}(G)$.*

2.4. Background on final coalgebras and final chains. Let $H : \mathbf{Set} \rightarrow \mathbf{Set}$ be any endofunctor. Recall that an H -coalgebra is a set A together with a morphism $a : A \rightarrow HA$. A coalgebra homomorphism into $b : B \rightarrow HB$ is a morphism $f : A \rightarrow B$ with $b \cdot f = Hf \cdot a$. A coalgebra A is *final* if for

every coalgebra B , there is a unique coalgebra morphism $f : B \rightarrow A$. The final coalgebra, if it exists, is denoted by νH . By Lambek's Lemma [22] the coalgebra structure of a final coalgebra is an isomorphism $\nu H \xrightarrow{\cong} H(\nu H)$. This implies that the power-set functor \mathcal{P} has no final coalgebra.

Dualizing the initial chain of [4], M. Barr [10] defined the *final chain* for a **Set**-endofunctor H . Let **Ord** be the class of all ordinals with the usual linear order considered as a category. The final chain is the chain $W : \mathbf{Ord}^{\text{op}} \rightarrow \mathbf{Set}$ determined (uniquely up-to natural isomorphism) by its objects $W_i, i \in \mathbf{Ord}$, and connecting morphisms $w_{i,j} : W_i \rightarrow W_j$ ($i \geq j$) as follows: $W_0 = 1$, $W_{i+1} = HW_i$, and $W_i = \lim_{j < i} W_j$ for limit ordinals i and $w_{i+1,j+1} = Hw_{i,j}$, whereas $(w_{i,j})_{j < i}$ is a limit cone for limit ordinals i . If this chain *converges* at some ordinal i , i.e., the connecting map $HW_i \rightarrow W_i$ is an isomorphism, then its inverse yields the final coalgebra for H . The finite steps of the final chain of H are called the *final ω^{op} -chain of H* .

We are primarily interested in this for $H = \mathcal{P}$ and $H = \mathcal{P}_\lambda$, where λ is an infinite cardinal number.

Notation 2.13. For every graph $e : G \rightarrow \mathcal{P}G$ and every ordinal number α , we have a *canonical morphism* $\kappa_\alpha^G : G \rightarrow \mathcal{P}^\alpha 1$. κ_0^G is the unique map $G \rightarrow 1$. Given $\kappa_\alpha^G : G \rightarrow \mathcal{P}^\alpha 1$, we set

$$\kappa_{\alpha+1}^G = G \xrightarrow{e} \mathcal{P}(G) \xrightarrow{\mathcal{P}\kappa_\alpha^G} \mathcal{P}^{\alpha+1} 1$$

For a limit ordinal λ , one checks that $(\kappa_\alpha^G)_{\alpha < \lambda}$ is a cone, and so there is a unique $\kappa_\lambda^G : G \rightarrow \mathcal{P}^\lambda 1$ such that for all $\alpha < \lambda$, $w_{\lambda,\alpha} \cdot \kappa_\alpha^G = \kappa_\lambda^G$.

It is easy to check by induction that if $f : G \rightarrow H$ is a coalgebra morphism, then for all α , $\kappa_\alpha^H \cdot f = \kappa_\alpha^G$.

Notation 2.14. For every tree t denote by $\partial_n t$ the strongly extensional tree obtained by cutting t at level n (i.e. deleting all nodes of depth $> n$) and forming the strongly extensional quotient. For all trees t and u , we write $t \sim_n u$ to mean that $\partial_n t = \partial_n u$ (remember that we identify isomorphic trees).

Remark 2.15. The final chain of \mathcal{P} begins with the following ω^{op} -chain:

$$\begin{aligned} \mathcal{P}^n 1 &= \text{all strongly extensional trees of depth } \leq n \\ &\text{with connecting maps } \partial_n : \mathcal{P}^{n+1} 1 \rightarrow \mathcal{P}^n 1. \end{aligned}$$

Indeed, the unique element of 1 can be taken to be the root-only tree. Given a set $M \subseteq \mathcal{P}^n 1$, we identify it with the tree tupling of its elements

and obtain a tree in $\mathcal{P}^{n+1}1$. The first connecting map from $\mathcal{P}1$ to 1 is obviously ∂_0 , and given that the n -th connecting map is $\partial_n: \mathcal{P}^{n+1}1 \rightarrow \mathcal{P}^n1$, it follows that the next connecting map, $\mathcal{P}\partial_n$, is (with the above tree tupling identification) precisely ∂_{n+1} .

Remark 2.16. Besides the power-set functor \mathcal{P} we also treat the finite power-set functor \mathcal{P}_f . Its final chain starts with the same ω^{op} -chain which we mentioned in Remark 2.15 just above. We describe the final chains and the final coalgebras below. Let us recall Barr's description for \mathcal{P}_f .

Notation 2.17. The set B of all finitely branching extensional trees is a coalgebra for \mathcal{P}_f : the coalgebra map is the inverse of tree tupling. This coalgebra is weakly final, and a final coalgebra can be described as its strongly extensional quotient, as we recall in Theorem 2.18 next (see Notation 3.1 for the definition of Barr equivalence.)

Theorem 2.18 (M. Barr [10]). *The final coalgebra for \mathcal{P}_f can be described as the quotient B/\sim_ω of the coalgebra of all finitely branching, extensional trees modulo the relation \sim_ω of Barr equivalence.*

3. Saturated trees and the final chain of the power-set functor

In this section we discuss the final chain of the power-set functor \mathcal{P} . We introduce the concept of an i -saturated tree for all ordinals i and prove that \mathcal{P}^i1 is the set of all i -saturated, strongly extensional trees for all ordinals of cofinality ω .

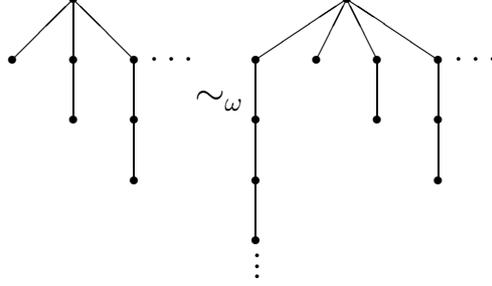
Notation 3.1. Recall that the subtree of t rooted at the node x is denoted by t_x . We define equivalences \sim_i on the class of all trees for every ordinal i by transfinite induction:

$$\begin{array}{ll} s \sim_0 t & \text{holds for all pairs } s, t; \\ s \sim_{i+1} t & \text{holds iff for every child } x \text{ of the root} \\ & \text{of } s \text{ there is a child } y \text{ of the root of } t \\ & \text{with } s_x \sim_i t_y, \text{ and vice versa} \end{array}$$

and for limit ordinals i , $s \sim_i t$ means $s \sim_j t$ for all $j < i$.

Following [6], we call trees t and u *Barr equivalent* if $t \sim_\omega u$.

Example 3.2. The trees in the picture below are Barr equivalent trees.



These trees are not related by $\sim_{\omega+1}$.

As shown by Malitz [24], there exist, for every ordinal i , trees s and t with $s \sim_i t$ but $s \not\sim_{i+1} t$. For a different proof, see [6].

The astute reader may have noticed that we already defined \sim_n for natural numbers n in Notation 2.14. The next result tells us that this earlier usage is consistent with our present, more general definition.

Proposition 3.3. *For all finite n , and all trees t and u , $t \sim_n u$ iff, in the sense of Notation 2.14, $\partial_n t = \partial_n u$.*

Proof: By induction on n . The case for $n = 0$ is clear: $\partial_0 t$ is a one-point tree for all t , and \sim_0 relates all pairs of trees. Assume our result for n . Suppose that $t \sim_{n+1} u$, so that every tree t_x is related by \sim_n to some tree u_y ; and vice-versa. By induction hypothesis, we see that every tree $\partial_n t_x$ is equal to some tree $\partial_n u_y$; and vice-versa. Now $\partial_{n+1} t$ is the tree-tupling of the trees $\partial_n t_x$ (without repetition); and the same holds for $\partial_{n+1} u$. So these trees $\partial_{n+1} t$ and $\partial_{n+1} u$ are isomorphic (equal). The converse is similar. ■

Proposition 3.4. *Let t and u be trees, considered as coalgebras for \mathcal{P} , and consider the canonical morphisms $\kappa_\alpha^t : t \rightarrow \mathcal{P}^\alpha 1$ and $\kappa_\alpha^u : u \rightarrow \mathcal{P}^\alpha 1$ (see Notation 2.13). Then $t \sim_\alpha u$ iff $\kappa_\alpha^t(\text{root}(t)) = \kappa_\alpha^u(\text{root}(u))$.*

Proof: By induction on α . The cases $\alpha = 0$ and α a limit ordinal are easy. Assuming our result for α , we prove it for $\alpha + 1$. Then $t \sim_{\alpha+1} u$ means, by induction hypothesis, that for every child x of the root of t there is a child y of the root of u with $\kappa_\alpha^{t_x}(x) = \kappa_\alpha^{u_y}(y)$, and vice-versa. The last equality shows that $\kappa_\alpha^t(x) = \kappa_\alpha^u(y)$: use the fact that the inclusion $t_x \rightarrow t$ is a coalgebra morphism. The above holds iff $\kappa_{\alpha+1}^t(\text{root}(t)) = \kappa_{\alpha+1}^u(\text{root}(u))$, as required. ■

Definition 3.5. We define the concept of *i-saturated tree* for every ordinal i by transfinite induction: A tree t is i -saturated iff

- (a) $i = 0$: t consists of the root only[†],
- (b) $i = j + 1$: t_x is j -saturated for every child x of the root, and
- (c) i a limit ordinal: given a tree s and a node x of t having children x_j with $s \sim_j t_{x_j}$ ($j < i$), then x has a child y with $s \sim_i t_y$.

- Examples 3.6.** (a) For i finite, a tree is i -saturated iff it has height at most i .
 (b) An example of an $(\omega + 1)$ -saturated tree which is not ω -saturated is the left-hand tree in Example 3.2.
 (c) In contrast to the preceding example, a concatenation of finitely many ω -saturated trees is ω -saturated.

Remark 3.7. If t and u are bisimilar trees, then they are equivalent under all of the above equivalences \sim_i . This is easy to see by transfinite induction.

Also, if t is i -saturated, then every tree bisimilar to t is i -saturated. In particular, the strongly extensional quotient of a tree t is i -saturated whenever t is.

Lemma 3.8. *For all i , if s and t are i -saturated strongly extensional trees, and if $s \sim_i t$, then $s = t$.*

Proof: By induction on i . The steps for $i = 0$ and for successor ordinals are easy. When i is a limit ordinal, we show that the relation $R \subseteq s \times t$, defined recursively as follows, is a tree bisimulation: $x R y$ iff x and y are the roots or have R -related parents and $s_x \sim_i t_y$.

For this, suppose that $s_x \sim_i t_y$, and let x^* be a child of x in s . For all $j < i$, $x \sim_{j+1} y$, so there is some child y_j^* of y such that $x^* \sim_j y_j^*$. This for all j together with the fact that t is i -saturated implies that there is a fixed child y^* of y such that for all j , $x^* \sim_j y^*$. Since i is a limit ordinal, $x^* \sim_i y^*$. The converse assertion is proved the same way, and we thus have proved that R indeed is a tree bisimulation. Consequently, $s = t$ (see Remark 2.9). ■

Theorem 3.9. *For every tree t and every ordinal α , there is a unique α -saturated, strongly extensional tree t_α^* with $t_\alpha^* \sim_\alpha t$.*

The proof is a bit technical and may be found in the Appendix.

[†]An alternative would be to say that a 0-saturated tree is one with no leaves. Theorem 3.17 below would still hold, but Example 3.6(a) would be lost.

Definition 3.10. The α -saturation of a tree t is the unique α -saturated tree u such that $t \sim_\alpha u$ with u strongly extensional. Its existence was stated in Theorem 3.9 just above.

Remark 3.11. Even though a tree t is i -saturated, it might not be strongly extensional and so might not be its own i -saturation.

We shall use α -saturations of trees later, in our description of the final chain of \mathcal{P} in Theorem 3.17. That result makes use of a certain graph structure on the sets $\mathcal{P}^\lambda 1$.

Notation 3.12. For each limit ordinal λ , the relation \rightsquigarrow on $\mathcal{P}^\lambda 1$ is given as follows:

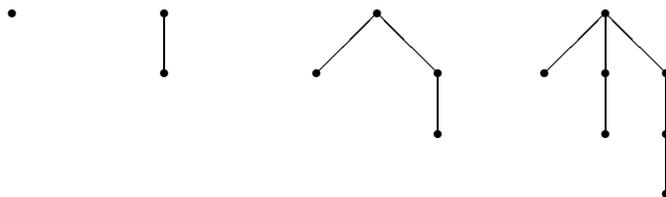
$$x \rightsquigarrow y \quad \text{iff for all } \alpha < \lambda, w_{\lambda,\alpha}(y) \in w_{\lambda,\alpha+1}(x).$$

We use G_λ to denote the graph $(\mathcal{P}^\lambda 1, \rightsquigarrow)$. As in Notation 2.13, we let $\kappa_\alpha^{G_\lambda} : \mathcal{P}^\lambda 1 \rightarrow \mathcal{P}^\alpha 1$ be the canonical map. For any $x \in \mathcal{P}^\lambda 1$, we let

$$\hat{x} = P(G_\lambda, x).$$

This is the tree expansion of $\mathcal{P}^\lambda 1$ starting from x , using the relation \rightsquigarrow . Note that $\text{root}(\hat{x})$ is (x) .

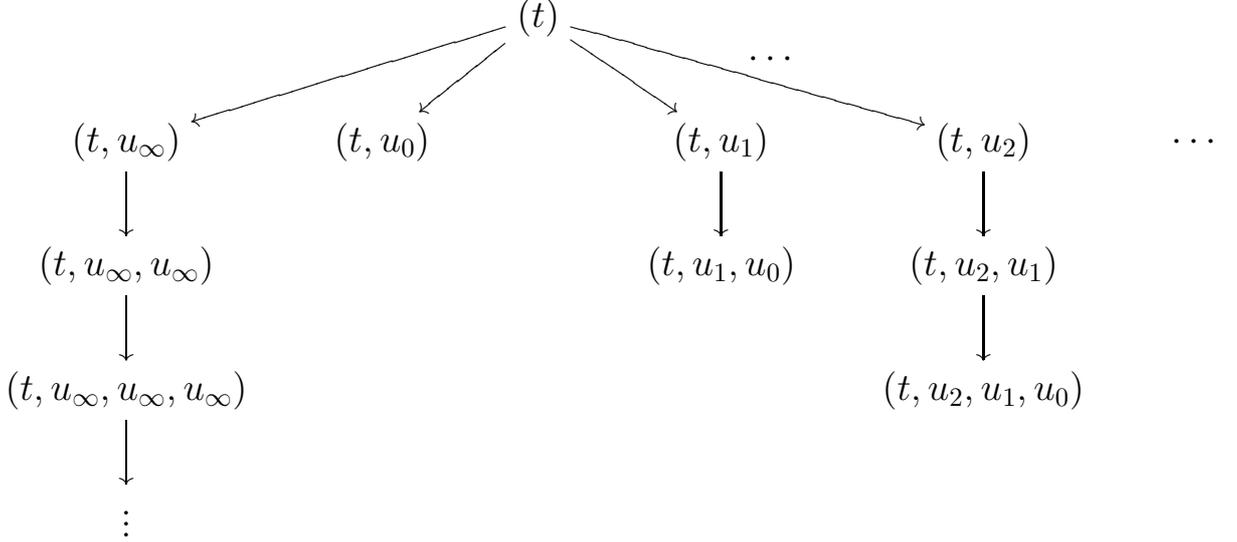
Example 3.13. Let t_0, t_1, t_2 , and t_3 be as shown below:



More generally, let t_n have a root and n children, say $1, \dots, n$, with the j th child the root of a chain of length j . Then t_n is an extensional tree of depth n . Moreover, $w_{n,m}(t_n) = t_m$, for all $m \leq n$. Thus $n \mapsto t_n$ is an element of $\mathcal{P}^\omega 1$. We call this element t .

For all n , let c_n be a chain of length n . Also, let $u_n \in \mathcal{P}^\omega 1$ be $m \mapsto c_{\min(n,m)}$. (Note that for $m \geq n$, $w_{\omega,m}u_n = c_n$.) Let u_∞ be $m \mapsto c_m$. Then $u_0, u_1, \dots, u_\infty$ all belong to $\mathcal{P}^\omega 1$. Moreover, \hat{t} , the tree expansion of the

graph $G_\omega = (\mathcal{P}^\omega 1, \rightsquigarrow)$ determined by t above looks like



(We omit the details on how this tree expansion was calculated. Although they are interesting, they lead us to a different branch of our main tree.) Notice that this is the *second* tree shown in Example 3.2. This tree is ω -saturated. The point is that t itself is the image of the *first* tree in Example 3.2 in $\mathcal{P}^\omega 1$. In our notation, this is $\kappa_\omega^{G_\omega}(\text{root}(t))$. So \widehat{t} , the tree expansion of the image of t inside $\mathcal{P}^\omega 1$, is the “ ω -saturation” of t . Our results below show that this is the case for all trees.

Observation 3.14. *For all $x \in \mathcal{P}^\lambda 1$ and all vertices \mathbf{u} of \widehat{x} , $\widehat{x}_{\mathbf{u}} = \widehat{\text{last}(u)}$. Also, for all $x \in \mathcal{P}^\lambda 1$ and all ordinals α , $\kappa_\alpha^{\widehat{x}}(\text{root}(\widehat{x})) = \kappa_\alpha^{G^\lambda}(x)$ for all ordinals $\alpha < \lambda$.*

Proof: The function **last** may be regarded as a coalgebra morphism $\text{last} : \widehat{x} \rightarrow \mathcal{P}^\lambda 1$. Note also that **last** takes the one point sequence (x) to x itself. Also canonical morphisms are preserved by coalgebra morphisms. \blacksquare

Before presenting our main results, we remind the reader that an ordinal λ has cofinality ω when there is a strictly increasing function $f : \omega \rightarrow \lambda$ such that $\lambda = \sup_n f(n)$.

Lemma 3.15. *For every ordinal λ of cofinality ω , and every $\alpha < \lambda$, the connecting map $w_{\lambda, \alpha} : \mathcal{P}^\lambda 1 \rightarrow \mathcal{P}^\alpha 1$ satisfies:*

$$w_{\lambda, \alpha+1}(x) = \{w_{\lambda, \alpha}(y) : x \rightsquigarrow y\}.$$

Proof: The inclusion \supseteq follows from the definition of \rightsquigarrow : we have $w_{\lambda,\alpha}(y) \in w_{\lambda,\alpha+1}(x)$.

To prove the inclusion \subseteq , choose $r_0 \in w_{\lambda,\alpha+1}(x)$. We must find $x \rightsquigarrow y$ with $r_0 = w_{\lambda,\alpha}(y)$. For that, first express λ as a supremum of an increasing sequence α_n of ordinals with $\alpha_0 = \alpha$. We are going to present, by induction, elements $r_n \in w_{\lambda,\alpha_{n+1}}(x)$ which are compatible; i.e., the connecting map from $\mathcal{P}^{\alpha_{n+1}}1$ to $\mathcal{P}^{\alpha_n}1$ takes r_{n+1} to r_n .

Given $r_n \in w_{\lambda,\alpha_{n+1}}(x)$, use the commutative triangle with $i = \alpha_{n+1}$ and $j = \alpha_n$:

$$\begin{array}{ccc} & \mathcal{P}^{\lambda}1 & \\ w_{\lambda,j+1} \swarrow & & \searrow w_{\lambda,i+1} \\ \mathcal{P}^{j+1}1 = \mathcal{P}(\mathcal{P}^j1) & \xleftarrow{w_{i+1,j+1} = \mathcal{P}w_{ij}} & \mathcal{P}(\mathcal{P}^i1) = \mathcal{P}^{i+1}1 \end{array}$$

Since $r_n \in w_{\lambda,j+1}(x)$, there exists $r_{n+1} \in w_{\lambda,i+1}(x)$ with $w_{i,j}(r_{n+1}) = r_n$, as requested.

The unique element y of $\mathcal{P}^{\lambda}1 = \lim_{n < \omega} \mathcal{P}^{\alpha_n}1$ with $r_n = w_{\lambda,\alpha_n}(y)$ satisfies $r_0 = w_{\lambda,\alpha}(y)$. And from $\lambda = \bigvee_n \alpha_n$ and $w_{\lambda,\alpha_n}(y) \in w_{\lambda,\alpha_{n+1}}(x)$, we conclude $x \rightsquigarrow y$. \blacksquare

Proposition 3.16. *Let λ be a limit ordinal of cofinality ω .*

- (1) *The canonical maps of G^{λ} are given by $\kappa_{\alpha}^{G^{\lambda}} = w_{\lambda,\alpha}$.*
- (2) *For all $x \in \mathcal{P}^{\lambda}1$, \hat{x} is strongly extensional and λ -saturated.*
- (3) *Every λ -saturated, strongly extensional tree t is of the form \hat{x} for a unique $x \in \mathcal{P}^{\lambda}1$.*

Proof: Part 1 is proved by induction on α . The case for 0 and limit ordinals is easy. Assume that $\kappa_{\alpha}^{G^{\lambda}} = w_{\lambda,\alpha}$, and fix x . Then $\kappa_{\alpha+1}^{G^{\lambda}}(x) = \{\kappa_{\alpha}^{G^{\lambda}}(y) : x \rightsquigarrow y\}$. We then apply Lemma 3.15 and the induction hypothesis.

For part 2, we use Corollary 2.12 and prove that the graph $(\mathcal{P}^{\lambda}1, \rightsquigarrow)$ is strongly extensional. Let R be a bisimulation on this graph. For all α , R is a subrelation of the kernel relation of $\kappa_{\alpha}^{G^{\lambda}}$. When $\lambda = \alpha$, $\kappa_{\alpha}^{G^{\lambda}} = w_{\lambda,\lambda} = \text{id}$, and so its kernel relation is the diagonal. Hence R is a subset of the diagonal.

We next prove that \hat{x} is λ -saturated. Fix a node $x \in \mathcal{P}^{\lambda}1$, and also fix a node \mathbf{u} in the tree \hat{x} . We show that

If t is a tree with the property that for all $\alpha < \lambda$ there is some child \mathbf{z}_{α} of \mathbf{u} in \hat{x} such that $\hat{x}_{\mathbf{z}_{\alpha}} \sim_{\alpha} t$, then there is some fixed child \mathbf{z} of \mathbf{u} in \hat{x} such that $\hat{x}_{\mathbf{z}} \sim_{\alpha} t$ for all $\alpha < \lambda$.

Before doing this, it is useful to make a few observations in order to simplify the notation a bit. The nodes in \widehat{x} are sequences starting from x . Let $u = \text{last}(\mathbf{u})$. The children of \mathbf{u} in \widehat{x} are sequences of the form \mathbf{u}, y , where $u \rightsquigarrow y$. Moreover, for such a sequence \mathbf{u}, y , we have $\widehat{x}_{\mathbf{u}, y} = \widehat{y}$ (see Observation 3.14). With this in mind, we can reformulate what we need to show:

If t is a tree with the property that for all $\alpha < \lambda$ there is some $y_\alpha \in \mathcal{P}^\lambda 1$ such that $u \rightsquigarrow y_\alpha$ and $\widehat{y}_\alpha \sim_\alpha t$, then there is some fixed $y \in \mathcal{P}^\lambda 1$ such that $u \rightsquigarrow y$ and $\widehat{y} \sim_\alpha t$ for all $\alpha < \lambda$.

Here is the proof: Consider t as a \mathcal{P} -coalgebra, and let $y = \kappa_\lambda^t(\text{root}(t))$. We claim that for all $\alpha < \lambda$, $w_{\lambda, \alpha}(y) \in w_{\lambda, \alpha+1}(u)$. Fix α , and also $y_\alpha \in \mathcal{P}^\lambda 1$ such that $u \rightsquigarrow y_\alpha$ and $\widehat{y}_\alpha \sim_\alpha t$. We calculate:

$$\begin{aligned}
w_{\lambda, \alpha}(y) &= w_{\lambda, \alpha} \kappa_\lambda^t(\text{root}(t)) && \text{by definition of } y \\
&= \kappa_\alpha^t(\text{root}(t)) && \text{since } \kappa^t \text{ is a cone} \\
&= \kappa_\alpha^{\widehat{y}_\alpha}(\text{root}(\widehat{y}_\alpha)) && \text{since } \widehat{y}_\alpha \sim_\alpha t: \text{ see Proposition 3.4} \\
&= \kappa_\alpha^{G_\lambda}(y_\alpha) && \text{by Observation 3.14} \\
&= w_{\lambda, \alpha}(y_\alpha) && \text{by part 1}
\end{aligned} \tag{3.1}$$

And since $u \rightsquigarrow y_\alpha$, we indeed see that $w_{\lambda, \alpha}(y) \in w_{\lambda, \alpha+1}(u)$. This claim for all $\alpha < \lambda$ implies that $u \rightsquigarrow y$. To finish part 2 of this lemma, we show that for all $\alpha < \lambda$, $\widehat{y} \sim_\alpha t$. For this,

$$\begin{aligned}
\kappa_\alpha^{\widehat{y}_\alpha}(\text{root}(\widehat{y}_\alpha)) &= \kappa_\alpha^{G_\lambda}(y_\alpha) && \text{by Observation 3.14} \\
&= w_{\lambda, \alpha}(y_\alpha) && \text{by part 1} \\
&= w_{\lambda, \alpha}(y) && \text{by (3.1)} \\
&= \kappa_\alpha^{G_\lambda}(y) && \text{by part 1} \\
&= \kappa_\alpha^{\widehat{y}}(\text{root}(\widehat{y})) && \text{by Observation 3.14}
\end{aligned}$$

and so by Proposition 3.4, $\widehat{y} \sim_\alpha \widehat{y}_\alpha$. Since $\widehat{y}_\alpha \sim_\alpha t$, we are done.

For the last part, let t be strongly extensional and λ -saturated. Let $x = \kappa_\lambda^t(\text{root}(t))$. To show that $\widehat{x} = t$, it is sufficient to see that $t \sim_\lambda \widehat{x}$, and so $t = \widehat{x}$ by Lemma 3.8. But $\kappa_\lambda^t(\text{root}(t)) = x = w_{\lambda, \lambda}(x) = \kappa_\lambda^{G_\lambda}(x) = \kappa_\lambda^{\widehat{x}}(\text{root}(\widehat{x}))$, with the last equality using Observation 3.14. Then, again by Proposition 3.4, $t \sim_\lambda \widehat{x}$. For the uniqueness of x , suppose that $y \in \mathcal{P}^\lambda 1$ has $\widehat{y} = t$. Then, again by Observation 3.14,

$$y = w_{\lambda, \lambda}(y) = \kappa_\lambda^{G_\lambda}(y) = \kappa_\lambda^{\widehat{y}}(\text{root}(\widehat{y})) = \kappa_\lambda^t(\text{root}(t)) = x$$

This completes the proof. ■

The following is our main result in this section.

Theorem 3.17. *Let α be an ordinal which is either 0 or of cofinality ω . The map $x \mapsto \widehat{x}$ is injective on $\mathcal{P}^\alpha 1$. Identifying each $x \in \mathcal{P}^\alpha 1$ with the corresponding tree \widehat{x} , we have the following facts:*

- (1) $\mathcal{P}^\alpha 1 =$ all α -saturated strongly extensional trees.
- (2) For all natural numbers n , $\mathcal{P}^{\alpha+n} 1 =$ all $(\alpha + n)$ -saturated strongly extensional trees.
- (3) If G is any graph, then the canonical cone κ_α^G assigns to every vertex g the α -saturation of the tree expansion of g in G .
- (4) If $\beta > \alpha$ is also of cofinality ω , then the connecting map $w_{\beta,\alpha}$ is given by α -saturation of trees.

Proof: Suppose $\widehat{x} = \widehat{y}$. By Proposition 3.16, part (2), these trees are λ -saturated and strongly extensional, and by part (3) of the same proposition each is uniquely of the form \widehat{z} for some $z \in \mathcal{P}^\lambda 1$. Thus $x = y$.

We consider the parts of this result in turn. The first part is just a re-statement of Proposition 3.16, parts (2) and (3). The second part follows from the first by induction on $n \in \omega$. The induction step follows from the definition of $(\alpha + n + 1)$ -saturated and from the fact that the tree tupling of a set of (distinct) trees is a strongly extensional tree iff each tree in the set is itself strongly extensional.

For the next part, fix a graph G and a vertex g . Write t the tree expansion of g in G . Let $x = \kappa_\alpha^t(\text{root}(t))$. We show that $\widehat{x} \sim_\alpha t$, and then we are done since \widehat{x} is a strongly extensional α -saturated tree. Regarding t as a \mathcal{P} -coalgebra, we see that

$$\begin{aligned} \kappa_\alpha^G(g) &= \kappa_\alpha^t(\text{root}(t)) \quad \text{since } \text{last} : t \rightarrow G \text{ is a coalgebra morphism} \\ &= x \\ &= \kappa_\alpha^{G_\alpha}(x) \quad \text{by Proposition 3.16, part (1), } \kappa_\alpha^{G_\alpha} = w_{\alpha,\alpha} = \text{id} \\ &= \kappa_\alpha^{\widehat{x}}(\text{root}(\widehat{x})) \quad \text{by Observation 3.14} \end{aligned}$$

Thus $t \sim_\alpha \widehat{x}$. Identifying x with \widehat{x} , this means that the α -saturation of t is x .

The last part follows from part (3) and Proposition 3.16, part (1). For the graph $G_\beta = (\mathcal{P}^\beta 1, \rightsquigarrow)$, $w_{\beta,\alpha} = \kappa_\alpha^{G_\beta}$. ■

Corollary 3.18. *The connecting maps $w_{\beta,\alpha} : \mathcal{P}^\beta 1 \rightarrow \mathcal{P}^\alpha 1$ are epimorphisms whenever $\alpha \leq \beta$ are limit ordinals of cofinality ω .*

Indeed, every α -saturated tree $t \in \mathcal{P}^\alpha 1$ is β -saturated. Thus, this is an element of $\mathcal{P}^\beta 1$ sent by $w_{\beta,\alpha}$ to itself.

Other limit ordinals, such as ω_1 , present a greater difficulty: see related work in the last section.

Turning from \mathcal{P} to the finite power-set functor \mathcal{P}_f , we prove a result of J. Worrell. In the following we shall say that a tree is “saturated” if it is ω -saturated. We also remind the reader of the description of $\mathcal{P}^n 1$ in Remark 2.15. Inspired by J. Worrell’s proof in [30] that the final chain of \mathcal{P}_f converges in $\omega + \omega$ steps, we have a description of the sets $\mathcal{P}_f^\alpha 1$ in this chain:

Corollary 3.19. *The final chain of \mathcal{P}_f converges in $\omega + \omega$ steps. We may describe the sets $\mathcal{P}_f^\alpha 1$ for $\alpha \leq \omega + \omega$ as follows:*

$$\begin{aligned} \mathcal{P}_f^\omega 1 &= \text{all saturated, strongly extensional trees,} \\ \mathcal{P}_f^{\omega+n} 1 &= \text{all saturated, strongly extensional trees finitely branching up to level } n - 1 \\ \mathcal{P}_f^{\omega+\omega} 1 &= \text{all finitely branching, strongly extensional trees.} \end{aligned}$$

Proof: First, note that $\mathcal{P}_f^n 1 = \mathcal{P}^n 1$ for all finite n , and so $\mathcal{P}_f^\omega 1 = \mathcal{P}^\omega 1$. Thus, the description of $\mathcal{P}_f^\omega 1$ follows from Theorem 3.17. For $n = 1$ we have $\mathcal{P}_f^{\omega+1} 1 = \mathcal{P}_f(\mathcal{P}_f^\omega 1)$ and we identify, again, every finite set $M \subseteq \mathcal{P}_f^\omega 1$ of saturated, strongly extensional trees with its tree-tupling. This is, by Example 3.6(c), a saturated, strongly extensional tree which is finitely branching at the root—and conversely, every such tree is a tree tupling of a finite subset of $\mathcal{P}_f^\omega 1$. Analogously for $n = 2$: we have $\mathcal{P}_f^{\omega+2} 1 = \mathcal{P}_f(\mathcal{P}_f^{\omega+1} 1)$ and the resulting trees are precisely those trees in $\mathcal{P}_f^\omega 1$ that are finitely branching at levels 0 and 1, etc. The connecting maps are the inclusion maps. The limit $\mathcal{P}_f^{\omega+\omega} 1 = \lim_{n < \omega} \mathcal{P}_f^{\omega+n} 1$ is the intersection of these subsets of $\mathcal{P}_f^\omega 1$ which consists of all finitely branching, strongly extensional trees: they are saturated, see Example 3.6. ■

4. On the final chain of \mathcal{P}_λ

Let λ be an uncountable cardinal number. We do not have a concrete description of the final chain of \mathcal{P}_λ except for $\lambda = \omega$; this is our work on \mathcal{P}_f above. But we know from D. Schwencke [28] that the final coalgebra for \mathcal{P}_λ is carried by the set of all λ -branching, strongly extensional trees (where λ -branching means that every vertex has less than λ children). The proof in [28] used the theory of coequations developed there. Here we give a simple direct proof:

Theorem 4.1. *For every infinite cardinal λ , the final coalgebra for \mathcal{P}_λ can be described as the set T_λ of all strongly extensional λ -branching trees. The coalgebra structure $T_\lambda \rightarrow \mathcal{P}_\lambda T_\lambda$ is the inverse of tree-tupling.*

Proof: Firstly, T_λ is a set because each λ -branching tree has at most λ vertices and thus is isomorphic to a tree whose set of vertices is a subset of λ . As a \mathcal{P}_λ -coalgebra (i. e. a λ -branching graph), it is strongly extensional because two strongly extensional trees are bisimilar iff they are isomorphic, and we identify isomorphic trees. For any strongly extensional \mathcal{P}_λ -coalgebra G , we have a \mathcal{P}_λ -coalgebra morphism from G to T_λ given by

$$g \mapsto \mathbf{P}(G, g),$$

where \mathbf{P} is the paths functor from Notation 2.10. Therefore, for a \mathcal{P}_λ -coalgebra H , we have a \mathcal{P}_λ -coalgebra morphism

$$H \longrightarrow H/\sim \longrightarrow T_\lambda,$$

where \sim is the largest bisimulation on H . This morphism into T_λ is unique by strong extensionality. ■

Remark 4.2. (a) For regular cardinals λ it follows from Worrell's paper [30] that the ω steps in the final chain after λ are all monomorphisms:

$$\mathcal{P}_\lambda^\lambda 1 \longleftarrow \mathcal{P}_\lambda^{\lambda+1} \longleftarrow \mathcal{P}_\lambda^{\lambda+2} \longleftarrow \dots$$

Consequently, the final chain for \mathcal{P}_λ stops after at most $\lambda + \omega$ steps.

We are going to prove in the next theorem that this is the best possible result.

(b) Given an infinite cardinal λ , denote by λ^+ the cardinal successor. This is always a regular cardinal (a consequence of the axiom of choice, see e.g. [18]). Somewhat surprisingly, when λ is singular, the final chain of \mathcal{P}_λ does not converge before λ^+ steps. The key technical tool is Lemma 4.3 below.

Lemma 4.3. *For all cardinals λ , the functor \mathcal{P}_λ preserves κ^{op} -limits for all regular $\kappa > \lambda$.*

Proof: Consider a κ^{op} -chain A_i , and denote by $a_{i,j} : A_i \rightarrow A_j$ the connecting maps for $j \leq i < \kappa$. Let $X_i \in \mathcal{P}_\lambda A_i$ be a compatible family. This means that for $i \geq j$, $a_{i,j}[X_i] = X_j$. So $a_{i,j}$ restricts to a surjection of X_i onto X_j . Our goal will be to find an ordinal $k < \kappa$ such that whenever $k \leq j$, $a_{j,k}$ restricts

to a bijection of X_k with X_j . It then clearly follows that \mathcal{P}_λ preserves the given limit because every compatible family is eventually constant.

Define an increasing sequence $(g(p))_{p < \kappa}$ of ordinals $< \kappa$ by recursion:

- $g(0) = 0$;
- $g(p + 1)$ is the least i in the range $g(p) \leq i < \kappa$ such that $a_{i,g(p)}$ is not injective on X_i —if there is no such i then $g(p)$ has the desired property so we stop;
- for limit ordinals p , $g(p) = \sup_{q < p} g(q)$, which is $< \kappa$ by regularity.

For $r < \kappa$ such that $g(r)$ is defined, we shall show

$$r \leq \text{card } X_{g(r)} \times \text{card } X_{g(r)} < \lambda \times \lambda. \quad (4.1)$$

The special case $r = \lambda \times \lambda$ (which is $< \kappa$) implies that $g(\lambda \times \lambda)$ is undefined. Therefore there is some $p < \lambda \times \lambda$ such that $g(p)$ is defined and $g(p + 1)$ is not, i.e. $g(p)$ has the desired property.

To prove (4.1), recall that for $p < r$, $X_{g(p)} = a_{g(r),g(p)}[X_{g(r)}]$. Since $p + 1 \leq r$, $g(p + 1)$ is defined, so there is a pair $h(p)$ of elements of $X_{g(r)}$ whose projections to $A_{g(p+1)}$ are different but are merged by $a_{g(p+1),g(p)}$. This gives a map $h : r \rightarrow X_{g(r)} \times X_{g(r)}$. Let us check that h is injective. Let $p < q$. Since $p + 1 \leq q$, $g(p + 1) \leq g(q)$. Now $h(q)$ is a pair whose projections to $A_{g(q+1)}$ are merged by $a_{g(q+1),g(q)}$, so the projections of this pair $h(q)$ to $A_{g(p+1)}$ are not different. But the projections of the pair $h(p)$ to $A_{g(p+1)}$ are different, whence $h(q) \neq h(p)$. ■

Note that the case $\lambda = \omega$ of Lemma 4.3 appears in [5, Example 3(iv)].

For any set functor F , we say that an ordinal i is the *convergence ordinal* of the final chain of F if $F^i 1 = F^{i+1} 1$, and i is the smallest ordinal with this property.

Theorem 4.4. *For λ an infinite cardinal, the convergence ordinal of the final chain of \mathcal{P}_λ is precisely*

- (a) $\lambda + \omega$, if λ is regular, and
- (b) λ^+ , if λ is singular.

Proof: (a) Let λ be regular. In view of Worrell's result all we need to prove about the final chain $W_i = \mathcal{P}_\lambda^i 1$ is that the connecting map after λ steps

$$w_{\lambda+1,\lambda} : W_{\lambda+1} \rightarrow W_\lambda$$

is not an epimorphism. It easily follows that the next one, $w_{\lambda+2,\lambda+1} = \mathcal{P}_\lambda w_{\lambda+1,\lambda}$ is also not an epimorphism, etc. Thus, convergence before $\lambda + \omega$ is impossible.

We define elements $a_j^i \in W_i$ for all ordinals $j \leq i \leq \lambda$ and prove that

$$a_\lambda^\lambda \text{ does not lie in the image of } w_{\lambda+1,\lambda}.$$

At the same time, we verify that for $j, k \leq i$,

$$w_{i,k}(a_j^i) = a_{\min(j,k)}^k. \quad (4.2)$$

These elements a_j^i are defined by transfinite recursion on i : For $i = 0$, we take a_0^0 to be the element of $1 = W_0$. Obviously we have $w_{0,k}(a_j^0) = a_{\min(j,k)}^k$ for $j = k = 0$.

Suppose that $i < \lambda$ and that we have $a_j^i \in W_i$ for all $j \leq i$. Let

$$a_j^{i+1} = \{a_l^i; l < j\}, \quad j \leq i + 1.$$

Since $i < \lambda$, this is an element of $\mathcal{P}_\lambda W_i = W_{i+1}$. To verify (4.2), we use induction on k . For $k = 0$, this point is again trivial. Assuming (4.2) for $k \leq i$, we see that for all $j \leq i + 1$,

$$\begin{aligned} w_{i+1,k+1}(a_j^{i+1}) &= \mathcal{P}_\lambda w_{i,k} \{a_l^i : l < j\} \\ &= \{w_{i,k}(a_l^i); l < j\} \\ &= \{a_{\min(k,l)}^k; l < j\} \\ &= \{a_l^k; l < \min(j, k + 1)\} \\ &= a_{\min(j,k+1)}^{k+1} \end{aligned}$$

When $k \leq i$ is a limit ordinal, note that for $l < k$,

$$\begin{aligned} w_{k,l} \cdot w_{i+1,k}(a_j^{i+1}) &= w_{i+1,l}(a_j^{i+1}) \\ &= a_{\min(j,l)}^l \\ &= a_{\min(l,\min(j,k))}^l \\ &= w_{k,l}(a_{\min(j,k)}^k) \end{aligned}$$

This for all $l < k$ shows that $w_{i+1,k}(a_j^{i+1}) = a_{\min(j,k)}^k$.

This concludes the definition of the elements a_j^{i+1} . If $i \leq \lambda$ is a limit ordinal, given a_l^k for all $l \leq k < i$, define

$$a_j^i = (a_0^0, a_1^1, \dots, a_j^j, a_j^{j+1}, a_j^{j+2}, \dots), \quad j \leq i.$$

(That is, for $k \leq i$, the k -th term in this sequence is $a_{\min(j,k)}^k$.) This sequence is easily seen to be compatible with the chain morphisms, using (4.2) below i . Further, (4.2) holds at i , by definition.

At this point, we have elements $a_j^i \in W_i$ for all $j \leq i \leq \lambda$. It is easy to check by induction on i that if $j \neq k \leq i$, then $a_j^i \neq a_k^i$. We are ready to prove that a_λ^λ does not lie in the image of $w_{\lambda+1,\lambda}$. Assuming the contrary,

$$w_{\lambda+1,\lambda}(b) = a_\lambda^\lambda \text{ for some } b \in \mathcal{P}_\lambda W_\lambda,$$

we derive a contradiction: the set $b \subseteq W_\lambda$ must have cardinality at least λ . To see this, we prove that the following subsets

$$c_i = \{y \in b; w_{\lambda,i+1}(y) = a_i^{i+1}\}, \quad i < \lambda,$$

are nonempty and pairwise disjoint.

We check first that $c_i \neq \emptyset$: the subset a_{i+2}^{i+2} of W_{i+1} contains clearly the element a_i^{i+1} . From (4.2) we get

$$\begin{aligned} a_{i+2}^{i+2} &= w_{\lambda,i+2}(a_\lambda^\lambda) \\ &= w_{\lambda,i+2} \cdot w_{\lambda+1,\lambda}(b) \\ &= \mathcal{P}_\lambda w_{\lambda,i+1}(b) \\ &= \{w_{\lambda,i+1}(y); y \in b\}. \end{aligned}$$

Thus, for $a_i^{i+1} \in a_{i+2}^{i+2}$ there exists $y \in b$ with $a_i^{i+1} = w_{\lambda,i+1}(y)$.

Second, $c_i \cap c_j = \emptyset$ for all $i < j$: consider $y \in c_j$, and observe that (4.2) implies

$$w_{\lambda,i+1}(y) = w_{j+1,i+1} \cdot w_{\lambda,j+1}(y) = w_{j+1,i+1}(a_j^{j+1}) = a_{\min(j,i+1)}^{i+1} = a_{i+1}^{i+1}.$$

Since the elements a_i^{i+1} and a_{i+1}^{i+1} are distinct, $y \notin c_i$.

(b) Let λ be singular. By Lemma 4.3, \mathcal{P}_λ preserves limits of λ^+ -chains. Thus the final chain converges in at most λ^+ steps. In order to prove that it does not converge before λ^+ , we can argue precisely as in [6, Theorems 5.5 and 5.8]. (Let us remark that the equivalence \sim_i in that paper is precisely $\sim_{i+\omega}$ of Notation 3.1 above. For the rest of the present proof we use the notation of [6].) In that argument it is sufficient to find, for every ordinal $i < \lambda^+$, a pair t_i, s_i of λ -branching trees satisfying $t_i \sim_i s_i$ but not $t_i \sim_{i+1} s_i$. For that, given a limit ordinal $j < \lambda^+$, choose a cofinal subset Q_j of j containing 0 and having cardinality less than λ . (This is possible since λ is singular, hence, the first cardinal with cofinality larger or equal to λ is λ^+ .) Now the construction of the trees t_i and s_i presented in [6] works provided that

we perform one simple modification: where the trees for limit ordinals j are defined in Definition 5.7, the running index k ranges (instead of through all ordinals smaller than j) through Q_j . ■

Remark 4.5. We also have precise information about the convergence ordinal of the initial chain of \mathcal{P}_λ . It is the smallest regular cardinal $\geq \lambda$. Call this κ , so that

$$\kappa = \begin{cases} \lambda & \text{if } \lambda \text{ is regular,} \\ \lambda^+ & \text{if } \lambda \text{ is singular.} \end{cases}$$

Here is the proof. We assume that the initial chain is constructed so that the connecting maps are inclusions and the limit steps are unions. For each $i < \kappa$, fix a cofinal set Q_i of size $< \lambda$.

Define sets \widehat{Q}_i by recursion:

$$\widehat{Q}_i = \{\widehat{Q}_j : j \in Q_i\}.$$

By induction on i , we see that $\widehat{Q}_i \in \mathcal{P}_\lambda^{i+1}(0)$. Assuming this for $j < i$,

$$\widehat{Q}_i \subseteq \bigcup_{j < i} \mathcal{P}_\lambda^{j+1}(0) \subseteq \mathcal{P}_\lambda^i(0).$$

Since $\text{card}(\widehat{Q}_i) < \lambda$, $\widehat{Q}_i \in \mathcal{P}_\lambda^{i+1}(0)$. But also, we show that for $i \leq j < \kappa$, $\widehat{Q}_j \notin \mathcal{P}_\lambda^i(0)$. The proof is by induction on i . For $i = 0$, this is because $\mathcal{P}_\lambda^0(0) = 0$. Assume that for $i \leq j$, $\widehat{Q}_j \notin \mathcal{P}_\lambda^i(0)$. Fix $j \geq i + 1 > i$. Since Q_j is cofinal in j , let $k \geq i$ belong to Q_j . So $\widehat{Q}_k \in \widehat{Q}_j$. By induction hypothesis, $\widehat{Q}_k \notin \mathcal{P}_\lambda^i(0)$. And so $\widehat{Q}_j \notin \mathcal{P}_\lambda^{i+1}(0)$. Finally, let i be a limit ordinal and assume that for $k < i$: if $k \leq j$, then $\widehat{Q}_j \notin \mathcal{P}_\lambda^k(0)$. Fix $j \geq i$. Then $\widehat{Q}_j \notin \bigcup_{k < i} \mathcal{P}_\lambda^k(0) = \mathcal{P}_\lambda^i(0)$.

The upshot is that for $i < \kappa$, $\widehat{Q}_i \in \mathcal{P}_\lambda^{i+1}(0) \setminus \mathcal{P}_\lambda^i(0)$. This shows that the ordinal of convergence of the initial chain of \mathcal{P}_λ is at least κ . Since \mathcal{P}_λ is κ -accessible, it is at most κ .

Note that in [9] it is shown that the initial chain of any endofunctor on **Set**, if it converges at all, converges at either 0, 1, 2, 3 or a regular infinite cardinal.

Remark 4.6. We have only treated \mathcal{P}_λ for infinite cardinals. We can also consider the functors \mathcal{P}_n for $1 < n < \omega$. \mathcal{P}_n gives the set of subsets of at most $n - 1$ elements. \mathcal{P}_n preserves colimits of ω -chains, and so the initial

chain converges at ω . Lemma 4.3 shows that the final chain of \mathcal{P}_n also converges at ω .

5. Modally saturated trees and modal structures

At this point, we have concluded our general results on the final chain of \mathcal{P} and related functors such as \mathcal{P}_λ . We turn back to the finite power set functor \mathcal{P}_f , and we present another characterization of the trees in $\mathcal{P}^\omega 1 = \mathcal{P}_f^\omega 1$; as we know, these are the ω -saturated trees. As before, we simplify terminology and refer to them as *saturated*.

K. Fine [14] introduced the concept of modal saturatedness for Kripke structures in modal logic. In this section, we review all of the needed definitions, and we prove that modally saturated trees are the same as saturated trees.

- (a) We work with modal logic formulated *without* atomic propositions. The sentences φ of modal logic are then given by

$$\varphi ::= \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi$$

We use the usual abbreviations:

$$\perp = \neg\top \quad \varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi) \quad \varphi \rightarrow \psi = \neg\varphi \vee \psi \quad \Diamond\varphi = \neg\Box\neg\varphi.$$

A sentence has *depth* n if n is the maximum of nested \Box in it.

- (b) We interpret modal logic on Kripke structures. Since we have no atomic sentences, our Kripke structures are just graphs $G = (G, \rightarrow)$, where \rightarrow is a binary relation on the set G . The main semantic relation is the *satisfaction* relation \models between the vertex set of a given graph and the sentences of the logic. This is defined as follows:

$$\begin{aligned} a \models \top & \quad \text{always} \\ a \models \neg\varphi & \quad \text{iff it is not the case that } a \models \varphi \\ a \models \varphi \wedge \psi & \quad \text{iff } a \models \varphi \text{ and } a \models \psi \\ a \models \Box\varphi & \quad \text{iff for all neighbors } b \text{ of } a, b \models \varphi \end{aligned}$$

Given a tree t we write $t \models \varphi$ if the root satisfies φ .

- (c) A *theory* is a set S of sentences. We write $a \models S$ if $a \models \varphi$ for all $\varphi \in S$ and call a a *model* of S .
- (d) Turning to the proof system, the modal logic \mathbf{K} extends the propositional logic (Hilbert's style) by one axiom $\Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$, called \mathbf{K} ,

and one deduction rule: if $\varphi \in \mathbf{K}$ then $\Box\varphi \in \mathbf{K}$. We write $\vdash \varphi$ if φ can be derived in this logic.

This logic is sound and complete. That is, $\vdash \varphi$ holds iff for every vertex a of any graph, $a \models \varphi$.

- (e) A theory S is *inconsistent* if for some finite $\{\varphi_1, \dots, \varphi_n\} \subseteq S$, $\vdash \neg \bigwedge \varphi_i$. S is *consistent* if S is not inconsistent. Or, equivalently, S has a model. If, moreover, $S \cup \{\varphi\}$ is inconsistent for every sentence $\varphi \notin S$, then S is *maximal consistent*.
- (f) $\Box S$ denotes the theory $\{\Box\varphi : \varphi \in S\}$, and $\Box^k S = \Box(\Box^{k-1} S)$ for $k \geq 2$. We use the notation $\Diamond S$ similarly.

Definition 5.1. We define *canonical sentences* χ of depth n by recursion on n , as follows:

- (a) \top is the only canonical sentence of depth 0, and
- (b) canonical sentences of depth $n + 1$ are precisely the sentences

$$\nabla S = (\bigwedge \Diamond S) \wedge \Box \bigvee S,$$

where S is a set of canonical sentences of depth n .

We use the conventions that $\bigwedge \emptyset = \top$, $\bigvee \emptyset = \perp$, and we often identify sentences φ and ψ when $\vdash \varphi \leftrightarrow \psi$ in \mathbf{K} .

Example 5.2. We have two canonical sentences of depth 1.

$$\nabla \emptyset = \top \wedge \Box \perp = \Box \perp \quad \text{and} \quad \nabla \{\top\} = \Diamond \top \wedge \Box \top = \Diamond \top$$

distinguishing whether the given vertex has a neighbor or not.

Theorem 5.3 (K. Fine [14] and L. Moss [25]). *For every vertex a of a graph and every $n \in \mathbb{N}$ there exists a unique canonical sentence χ of depth n satisfied by a . Moreover, for every canonical sentence χ of depth n and every sentence ψ of depth at most n , either $\vdash \chi \rightarrow \psi$ or $\vdash \chi \rightarrow \neg\psi$.*

Corollary 5.4. *The sentences of depth at most n form a finite set (up to logical equivalence in \mathbf{K}).*

Proof: Observe first that there are only finitely many canonical sentences of depth n . Let ψ be any sentence of depth n . Let A be the set of all canonical sentences χ of depth n with $\vdash \chi \rightarrow \psi$ and let B be the canonical sentences χ of depth n with $\vdash \chi \rightarrow \neg\psi$. So we have $\vdash \bigvee A \rightarrow \psi$ and $\vdash \bigvee B \rightarrow \neg\psi$. By Theorem 5.3, we have

$$\vdash \bigvee A \vee \bigvee B.$$

So by propositional logic we have $\vdash \bigvee A \leftrightarrow \psi$. Thus, every sentence of depth n is equivalent to a disjunction of canonical sentences of depth n , from which the desired result follows. ■

Notation 5.5. (a) For every tree t we denote by $\chi_n(t)$ the unique canonical sentence of depth n satisfied in the root. It is easy to prove that

$$\chi_{n+1}(t) = \nabla \{ \chi_n(t_x) : x \text{ child of the root of } t \}.$$

(b) For any graph G , and any $a \in G$, we denote by S_a the set of all sentences φ with $a \models \varphi$ in G . For a tree t , we similarly denote by S_t the set of sentences satisfied by the root of t .

(c) Recall from [11] that the *canonical model* of \mathbf{K} is the graph C whose vertices are the maximal consistent theories, and with $S \rightarrow S'$ iff $\diamond S' \subseteq S$ (equivalently, $\Box S \subseteq S'$). The *Truth Lemma* (see [11, Lemma 4.21]) is the statement that for all $S \in C$,

$$\{ \varphi : S \models \varphi \text{ in } C \} = S.$$

This lemma is easy to check by induction on φ .

Corollary 5.6. *For two trees t and s we have $t \sim_n s$ iff $t \models \chi_n(s)$. Consequently, $t \sim_\omega s$ iff $S_t = S_s$.*

Proposition 5.7. *The limit $\mathcal{P}_f^\omega 1$ can be described as the set C of all maximal consistent theories in \mathbf{K} .*

Proof: We have described $\mathcal{P}_f^\omega 1$ as the set of all saturated, strongly extensional trees. We prove that $t \mapsto S_t$ is a bijection between this set and C . This finishes the proof. (a) For every $t \in \mathcal{P}_f^\omega 1$ the theory S_t is maximal consistent. Indeed, it is obviously consistent. Given $\varphi \notin S$ of depth n , we have $t \not\models \varphi$ and $t \models \chi_n(t)$, thus, $\not\vdash \chi_n(t) \rightarrow \varphi$. By Theorem 5.3, $\vdash \chi_n(t) \rightarrow \neg\varphi$. Therefore, $S_t \cup \{\varphi\}$ is inconsistent. (b) By the Truth Lemma, every maximal consistent theory S is of the form S_t for some t : let t be the expansion of the canonical graph C at S . Moreover, t can be taken as saturated and strongly extensional, since the saturation operation on trees preserves modal theories (see Corollary 5.6). ■

Definition 5.8. A theory S is called *hereditarily finite* if it is maximal consistent and for every $k \in \mathbb{N}$ there exist only finitely many maximal consistent theories S' with $\diamond^k S' \subseteq S$.

Theorem 5.9. *The set of all hereditarily finite theories is a final coalgebra for \mathcal{P}_f via the coalgebra map $S \mapsto \{S' : \diamond S' \subseteq S\}$.*

Proof: We prove that the bijection $t \mapsto S_t$ of Proposition 5.7 has the property that for $t \in \mathcal{P}_f^\omega 1$ we have that t is finitely branching iff S_t hereditarily finite. From that our theorem follows, since the coalgebra map above corresponds to the coalgebra map of $\nu \mathcal{P}_f$. Indeed:

(a) If S_t is hereditarily finite, then t is finitely branching. It is sufficient to verify that t is finitely branching at the root. Given a node x of depth k , we then apply this to t_x : the theory of this subtree is also hereditarily finite, since $\diamond^k S_{t_x} \subseteq S_t$ (indeed: if $t_x \models \varphi$ then $t \models \diamond^k \varphi$).

Every child a of the root of t fulfils $\diamond S_{t_a} \subseteq S_t$. Thus, there are only finitely many such theories S_{t_a} . Now let a and b be children of the root of t with $S_{t_a} = S_{t_b}$, whence $t_a \sim_\omega t_b$ by Corollary 5.6. So since t_a and t_b are saturated and strongly extensional, we have $t_a = t_b$ by Lemma 3.8. Therefore, the root has only finitely many children.

(b) If t is finitely branching, then S_t is hereditarily finite. Indeed, for every maximal consistent theory S' with $\diamond^k S' \subseteq S_t$ let s be a tree with $S' = S_s$ (see Proposition 5.7). Then for every $n \in \mathbb{N}$ we have $t \models \diamond^k \chi_n(s)$, i.e., some node of t of depth k satisfies $\chi_n(s)$. Since we have only finitely many such nodes, one of them, say a , satisfies $\chi_n(s)$ for all n . That is, $t_a \sim_n s$ for $n \in \mathbb{N}$, hence, $S_{t_a} = S'$, see Corollary 5.6. Since we have only finitely many nodes a of depth k , we see that S_t is hereditarily finite. ■

Definition 5.10 (see [14]). A graph is called *modally saturated* if for every node a , given a theory S such that

$$a \models \diamond \bigwedge S_0 \quad \text{for every finite } S_0 \subseteq S \quad (5.1)$$

there exists a neighbor b of a satisfying S .

Theorem 5.11. *A tree is saturated iff it is modally saturated.*

Proof: (a) Let t be modally saturated. Let a be a node in t , and let s be a tree with the property that there exist children x_n of a with $s \sim_n t_{x_n}$ ($n < \omega$). We prove $s \sim_\omega t_b$ for some child b . The theory S_s fulfils (5.1): given $S_0 \subseteq S_s$ finite, let n be the maximum of the depths of all $\psi \in S_0$; then $\vdash \chi_n(s) \rightarrow \psi$ for all $\psi \in S_0$ (see Theorem 5.3). By Corollary 5.6, $s \sim_n t_{x_n}$ iff $x_n \models \chi_n(s)$, and this implies $x_n \models \psi$ for all $\psi \in S_0$. Thus, $a \models \diamond \bigwedge S_0$. Let b be a neighbor of a satisfying S_s . Then $t_b \models \chi_n(s)$ for all n ; i.e., $s \sim_\omega t_b$ by Corollary 5.6.

(b) Let t be saturated. Let a be a node of t and S be a theory satisfying (5.1). For every natural number n define S_n to be a set of representatives of all $\psi \in S$ of depth at most n modulo logical equivalence in \mathbf{K} . By Corollary 5.4 the sentences of depth n form a finite set (up to logical equivalence). As a corollary of Theorem 5.3 one readily proves that there is only a finite set of sentences of depth at most n (up to logical equivalence). So we have that S_n is finite. By (5.1) we see that for every n , there exists a child b_n of a such that

$$b_n \models \psi \quad \text{for all } \psi \in S_n.$$

It is our task to prove that a has a child b satisfying S .

Let v be the graph whose nodes are all canonical sentences χ of depth any $n = 0, 1, 2, \dots$ such that $a \models \diamond\chi$ and $\vdash \chi \rightarrow \psi$ for all $\psi \in S_n$. We make v a graph using the *converse* of logical implication in \mathbf{K} . So the neighbors of the node χ are all the nodes χ' of depth $n + 1$ with $\vdash \chi' \rightarrow \chi$. The root is \top , and every node χ' of v has indeed a unique parent (so v is a tree): since $a \models \diamond\chi'$, we have a child c of a with $c \models \chi'$ which by Theorem 5.3 implies $\chi' = \chi_{n+1}(t_c)$. Put $\chi = \chi_n(t_c)$, then $\vdash \chi' \rightarrow \chi$. (This is because $\vdash \chi' \rightarrow \neg\chi$ cannot happen due to $c \models \chi'$ and $c \models \chi$. Now use Theorem 5.3). Consequently, χ is a parent of χ' . And the uniqueness of the parent is obvious: suppose $\vdash \chi \rightarrow \chi'$ where $\chi' \in v$ has depth n , then $t_c \models \chi'$, therefore $\chi' = \chi_n(t_c)$.

The tree v is obviously finitely branching. And since each $\chi_n(t_{b_n})$ lies in v and each of these formulas has a different depth, they form an infinite set of nodes of v . By König's Lemma, v has an infinite branch

$$\top = \chi_0 \leftarrow \chi_1 \leftarrow \chi_2 \dots$$

Each $S \cup \{\chi_n\}$ is consistent. Indeed, by compactness it is sufficient to verify that $S_k \cup \{\chi_n\}$ is consistent for every $k \geq n$: due to $a \models \diamond\chi_k$ we have a child c of a satisfying χ_k , then t_c is a model of S_k (due to $\vdash \chi_k \rightarrow \psi$ for all $\psi \in S_k$) and of χ_n (due to $\vdash \chi_k \rightarrow \chi_n$). Consequently, $S \cup \{\chi_0, \chi_1, \chi_2, \dots\}$ is consistent: use compactness again. Let s be a tree which is a model of the last theory. Then $s \models \chi_n$ which by Theorem 5.3 implies $\chi_n = \chi_n(s)$ for every n . On the other hand, since $a \models \diamond\chi_n$, we have a child c_n of a with $c_n \models \chi_n$, thus, $\chi_n = \chi_n(t_{c_n})$. By Corollary 5.6 this proves $s \sim_n t_{c_n}$. Since t is saturated, there exists a child b of a with $s \sim_\omega t_b$. Then $S \subseteq S_s = S_{t_b}$ which concludes the proof: b satisfies S . \blacksquare

5.1. Modal structures. Our previous section related $\mathcal{P}^\omega 1$, the set of all strongly extensional saturated trees, to the set of modally saturated trees. The point is that we have uncovered a definition of a structure which happens to be isomorphic to $\mathcal{P}^\omega 1$ and which was proposed for other reasons. We continue in this vein. The following definition comes from R. Fagin and M. Vardi [13], changing the notation a little but not the ideas. The changes are partly due to our decision to work without atomic propositions; these are of little importance in what we do. Further, we also work without “agents”, since the concerns of epistemic logic are even farther from this paper.

Definition 5.12. We define the sets \mathcal{S}_n of n -ary worlds as follows. \mathcal{S}_0 is a singleton set, say $\{*\}$. \mathcal{S}_{n+1} is the set of functions f with domain $n+1 = \{0, \dots, n\}$ meeting certain conditions; we write f_i for the value of f at i . These f_i are functions in the sense of being sets of ordered pairs; their codomains are their ranges. We require that the following conditions hold:

- (1) $f_0 = *$.
- (2) For $1 \leq i < n+1$, $f_i \subseteq \mathcal{S}_i$.
- (3) If $n > 1$, then for all $(n-1)$ -ary worlds g , $g \in f_{n-1}$ iff there is some set g_n^* of n -ary worlds such that $g \cup \{(n, g_n^*)\} \in f_n$.

As an illustration of what is going on, for every graph G , we get maps $k_n : G \rightarrow \mathcal{S}_n$ by recursion. The map k_0 is the constant f_0 . The inductive step defines $k_n(x)(i)$ for all vertices x in G and all $0 \leq i < n+1$ by $k_{n+1}(x)(0) = *$, and for $i > 0$,

$$k_{n+1}(x)(i) = \{k_i(y) : y \text{ is a child of } x \text{ in } G\}.$$

We shall prove that n -ary modal worlds correspond to elements of $\mathcal{P}^n 1$, that is, to extensional trees of depth at most n . To do this, we build a family of bijections $b_n : \mathcal{S}_n \rightarrow \mathcal{P}^n 1$. For this, we need some maps based on the definition above. There is a map

$$\rho_n : \mathcal{S}_{n+1} \rightarrow \mathcal{S}_n$$

taking an $(n+1)$ -ary world and restricting it to $n = \{0, 1, \dots, n-1\}$. (For $n = 0$, the map is the constant.) There is also a map

$$\eta_n : \mathcal{S}_{n+1} \rightarrow \mathcal{P}\mathcal{S}_n$$

taking an $(n+1)$ -ary world and returning its value on n .

Our formulation of point (3) above is that for all $n \geq 1$, we have a pullback square

$$\begin{array}{ccc} \mathcal{S}_{n+2} & \xrightarrow{\eta_{n+1}} & \mathcal{P}\mathcal{S}_{n+1} \\ \rho_{n+1} \downarrow & & \downarrow \mathcal{P}\rho_n \\ \mathcal{S}_{n+1} & \xrightarrow{\eta_n} & \mathcal{P}\mathcal{S}_n \end{array} \quad (5.2)$$

Lemma 5.13. *For all n , ρ_n is surjective and η_n is bijective.*

Proof: We show by induction on n that ρ_n and η_n are surjective. For $n = 0$, we verify directly. Obviously, ρ_0 is surjective. As for η_1 , $\mathcal{P}(\mathcal{S}_0) = \{\emptyset, \{*\}\}$, and

$$\begin{aligned} \eta_1(\{(0, *), (1, \{*\})\}) &= \{*\} \\ \eta_1(\{(0, *), (1, \emptyset)\}) &= \emptyset \end{aligned}$$

Assuming that ρ_n and η_n are surjective, we get the same property for $n + 1$ using the pullback square in (5.2).

We are left with the verification that each η_n is injective. Suppose that $\eta_n(f) = \eta_n(f')$. Then by (5.2), $\eta_{n-1} \circ \rho_n(f) = \eta_{n-1} \circ \rho_n(f')$. Precomposing with $\mathcal{P}\rho_{n-2}$ and using (5.2) again, we see that

$$\eta_{n-2} \circ \rho_{n-1} \circ \rho_n(f) = \eta_{n-2} \circ \rho_{n-1} \circ \rho_n(f').$$

Continuing in this manner, we see that for $1 \leq j \leq n$,

$$\eta_j \circ \rho_{j+1} \circ \cdots \circ \rho_n(f) = \eta_j \circ \rho_{j+1} \circ \cdots \circ \rho_n(f').$$

This is a long-winded way to say that the functions f and f' are identical. ■

In the statement below, we remind the reader of Remark 2.15, a description of the first ω -terms in the final chain of \mathcal{P} in terms of extensional trees of finite depth and the “cutoff” functions ∂_n .

Lemma 5.14. *There is a family of bijections $b_n : \mathcal{S}_n \rightarrow \mathcal{P}^n \mathbf{1}$ such that the following diagrams commute:*

$$\begin{array}{ccc} \mathcal{S}_{n+1} & \xrightarrow{b_{n+1}} & \mathcal{P}^{n+1} \mathbf{1} \\ \rho_n \downarrow & & \downarrow \partial_n \\ \mathcal{S}_n & \xrightarrow{b_n} & \mathcal{P}^n \mathbf{1} \end{array} \quad (5.3)$$

Proof: We define the maps b_n by recursion. For $n = 0$, \mathcal{S}_0 and $\mathcal{P}^0 1$ are both singleton sets, so b_0 is determined. The main part of the definition is $b_{n+1} = \mathcal{P}b_n \cdot \eta_n$. An easy induction using Lemma 5.13 shows that the maps b_n are bijections. We check that the squares in (5.3) commute by induction on n . For $n = 0$, we use the fact that $\mathcal{P}^0 1$ is a singleton. Assuming that (5.3) commutes for n , we show that it commutes for $n + 1$ by examining the diagram below:

$$\begin{array}{ccc}
 \mathcal{S}_{n+2} & \xrightarrow{b_{n+2}} & \mathcal{P}^{n+2} 1 \\
 \downarrow \rho_{n+1} & \searrow \eta_{n+1} \quad \nearrow \mathcal{P}b_{n+1} & \downarrow \partial_{n+1} = \mathcal{P}\partial_n \\
 & \mathcal{P}\mathcal{S}_{n+1} & \\
 & \downarrow \mathcal{P}\rho_n & \\
 & \mathcal{P}\mathcal{S}_n & \\
 \nearrow \eta_n & & \searrow \mathcal{P}b_n \\
 \mathcal{S}_{n+1} & \xrightarrow{b_{n+1}} & \mathcal{P}^{n+1} 1
 \end{array}$$

The top and bottom triangles are the definitions of b_{n+2} and b_{n+1} , respectively. The region on the right is the induction hypothesis, with \mathcal{P} added. The region on the left is (5.2). \blacksquare

At this point, we know that the set \mathcal{S}_n of n -ary worlds correspond to $\mathcal{P}^n 1$. The main definition in [13] is that of a *modal structure*:

Definition 5.15. A *modal structure* is a function f with domain ω such that for all $n \geq 1$, the restriction $r_n(f)$ of f to n is an n -ary world. Let \mathcal{MS} be the set of modal structures.

It is clear that the \mathcal{MS} together with the maps $(r_n)_{n \in \omega}$ is a limit cone for

$$\mathcal{S}_0 \xleftarrow{\rho_0} \mathcal{S}_1 \xleftarrow{\rho_1} \mathcal{S}_2 \xleftarrow{\rho_2} \dots$$

By Lemma 5.14, this diagram is isomorphic to

$$\mathcal{P}^0 1 \xleftarrow{\partial_0} \mathcal{P}^1 1 \xleftarrow{\partial_1} \mathcal{P}^2 1 \xleftarrow{\partial_2} \dots$$

This proves the following result:

Theorem 5.16. *The set \mathcal{MS} of modal structures is isomorphic to the set $\mathcal{P}^\omega 1$.*

In fact, features of modal structures mentioned in [13] are now immediate corollaries. Let us mention two of them: There are canonical maps $\kappa_n^G : G \rightarrow \mathcal{S}_n$ for all graphs G ; we have seen these in the beginning of this section. And there is a way to turn \mathcal{MS} into a graph, and the result is strongly extensional.

Summary. The set $\mathcal{P}^\omega 1$ may be described in various ways:

- (1) the limit of the first ω terms in the final chain of \mathcal{P} ,
- (2) the set of strongly extensional saturated trees,
- (3) the set of modally saturated trees,
- (4) the set of maximal consistent sets in \mathbf{K} , and
- (5) the set of modal structures.

6. Finite multisets with multiplicities in a commutative monoid

Here we continue the project initiated by H.-P. Gumm and T. Schröder [17] of investigating finitely branching Kripke structures with transitions having weights from a given commutative monoid $(M, +, 0)$. These are the coalgebras for the functor $\mathcal{M}_f : \mathbf{Set} \rightarrow \mathbf{Set}$ (denoted by \mathcal{M}_ω in [17]) assigning to every set X the set $\mathcal{M}_f X$ of all finite multisets in X , i.e. all functions $A : X \rightarrow M$ with $A^{-1}[M \setminus \{0\}]$ finite. Given a function $h : X \rightarrow Y$, the map $\mathcal{M}_f h$ assigns to every finite multiset $A : X \rightarrow M$ the finite multiset $\mathcal{M}_f h(A)$ sending $y \in Y$ to $\sum_{x \in X, h(x)=y} A(x)$.

Example 6.1. The Boolean monoid $P = \{0, 1\}$ yields the finite power-set functor \mathcal{P}_f . The cyclic group $C = \{0, 1\}$ yields a functor \mathcal{C}_f which coincides with \mathcal{P}_f on objects but is very different on morphisms. Using the natural numbers $(N, +, 0)$, we obtain the usual notion of a multiset.

Definition 6.2. By an M -labeled graph G is meant a graph whose edges are labeled in $M \setminus \{0\}$. We denote by $w_G : G \times G \rightarrow M$ the corresponding “weight” function with $w_G(x, y) \neq 0$ iff y is a neighbor of x .

Remark 6.3. (a) The coalgebras for \mathcal{M}_f are precisely the finitely branching M -labeled graphs. Indeed, given such a graph G , define the coalgebra structure $G \rightarrow \mathcal{M}_f G$ by assigning to every vertex x the finite multiset $w_G(x, -) : G \rightarrow M$. Conversely, every finitely branching M -labeled graph is obtained from precisely one coalgebra for \mathcal{M}_f .

- (b) Coalgebra homomorphisms between two finitely branching M -labeled graphs G and H are precisely the functions $f: G \rightarrow H$ between the vertex sets such that

$$w_H(f(x), y) = \sum_{x' \in G, f(x')=y} w_G(x, x') \quad \text{for all } x \in G, y \in H. \quad (6.1)$$

- (c) We identify, once again, two M -labeled trees whenever they are isomorphic (as coalgebras for \mathcal{M}_f).

Definition 6.4. An M -labeled tree is *extensional* if distinct children of any vertex define non-isomorphic M -labeled subtrees.

The *extensional modification* of a finite M -labeled tree t is obtained by successively performing the following operation, from the leaves towards the root: Given two vertices x and y of t with the same parent z , and such that $t_x = t_y$, if $w_t(z, x) + w_t(z, y) \neq 0$, identify t_x with t_y and put $w_t(z, \bar{x}) = w_t(z, x) + w_t(z, y)$, where \bar{x} is the identification of x with y ; otherwise remove t_x and t_y and the edges (z, x) and (z, y) . Since t is finite, this process certainly stops, and the resulting tree is extensional.

We use \sim_n and \sim_ω in an obvious analogy to Notations 2.14 and 3.1.

- Remark 6.5.** (a) Given a set functor $F: \mathbf{Set} \rightarrow \mathbf{Set}$, let (A, α) and (B, β) be coalgebras for F . Following [3], a bisimulation from (A, α) to (B, β) is a relation $R \subseteq A \times B$ for which there is a structure map $\delta: R \rightarrow FR$ making the projection maps $\pi_1: R \rightarrow A$ and $\pi_2: R \rightarrow B$ coalgebra homomorphisms. If F weakly preserves pullbacks, then a congruence in a coalgebra (A, α) is just an equivalence relation which is a bisimulation.
- (b) As observed in [17, Lemma 5.5], it follows that, in the particular case of the functor \mathcal{M}_f , a relation $R \subseteq A \times B$ is a bisimulation iff, for every aRb , there is a matrix:

$$m: A \times B \rightarrow M$$

with only a finite number of non-zero entries, such that:

- (1) $w_A(a, a') = \sum_{b' \in B} m(a', b')$ for all $a' \in A$,
- (2) $w_B(b, b') = \sum_{a' \in A} m(a', b')$ for all $b' \in B$, and
- (3) $m(a', b') \neq 0$ implies $a'Rb'$.

In the following, B denotes the coalgebra of all finitely branching M -labeled trees with the structure map $\beta : B \rightarrow \mathcal{M}_f B$ assigning to each tree t with root x_0 , the map $\beta_t : B \rightarrow M$ defined by $\beta_t(t') = w_B(t, t') = \sum_{t'=t_x} w_t(x_0, x)$ with x running through the vertices of t .

Theorem 6.6. *Let M be a commutative monoid. The coalgebra B/\sim_ω of all finitely branching M -labeled trees modulo Barr equivalence is final for \mathcal{M}_f .*

Proof: (1) B is weakly final. Indeed, for every finitely branching M -labeled graph (A, α) we define a coalgebra homomorphism $h : A \rightarrow B$ by assigning to every vertex $a \in A$ the tree expansion of a , t_a . Recall that the vertices of the tree expansion of a are the paths $a_0 a_1 \dots a_k$ of A starting in $a = a_0$, including the empty path, a , which is the root. A child of $a_0 a_1 \dots a_k$ is any extension $a_0 a_1 \dots a_k a_{k+1}$ and its weight in the tree expansion of a is $w_G(a_k, a_{k+1})$, see Definition 6.2. We need to prove that the square

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & \mathcal{M}_f A \\ h \downarrow & & \downarrow \mathcal{M}_f h \\ B & \xrightarrow{\beta} & \mathcal{M}_f B \end{array}$$

commutes, that is, for all $a \in A$ and $s \in B$, it holds $w_B(h(a), s) = \sum_{\substack{a' \in A \\ h(a')=s}} w_A(a, a')$

(see Remark 6.3(b)). Indeed, let a be a vertex of A . Then

$$w_B(h(a), s) = w_B(t_a, s) = \sum_{t_{a'}=s} w_{t_a}(a, aa') = \sum_{\substack{a' \in A \\ h(a')=s}} w_A(a, a').$$

(2) The final coalgebra is obtained from B by the quotient modulo the largest congruence: in general the quotient of a weakly final object by the largest congruence on it yields a final object.

(3) The Barr equivalence is a congruence on B . That is, the quotient B/\sim_ω carries a coalgebra structure for \mathcal{M}_f such that the quotient map $q : B \rightarrow B/\sim_\omega$ is a coalgebra homomorphism. To prove this, all we need to verify is

that given two trees



with $t \sim_\omega u$, then the multiset given by $[t_i]$ and m_i is the same one as the given by $[u_j]$ and n_j . This means that for every $s \in B$ it holds the following equality:

$$\sum_{\substack{i=1 \\ s \sim_\omega t_i}}^k m_i = \sum_{\substack{j=1 \\ s \sim_\omega u_j}}^l n_j.$$

Indeed, we have $\partial_n t = \partial_n u$ for all $n \in \omega$, that is, the cutting of t and u at the level n have the same extensional modification (see Definition 6.4). We can obtain $\partial_n t$ from t by first transforming each t_i into $\partial_{n-1} t_i$ for all i and finally identifying all those trees $\partial_{n-1} t_i$ which are equal, with the corresponding weights given as described in Definition 6.4. Analogously for the tree u . Thus, it is clear that, for all $n \in \omega$, we have

$$\sum_{\substack{i=1 \\ s \sim_n t_i}}^k m_i = \sum_{\substack{j=1 \\ s \sim_n u_j}}^l n_j.$$

For a fixed tree s of B , let A_n be the set of all trees t_i , $i = 1, \dots, k$, and all trees u_j , $j = 1, \dots, l$ which are \sim_n -related with s . Of course $A_m \subseteq A_n$ for $m \geq n$. Consequently, since the sets A_n have cardinality not greater than $k + l$, there is some $n_0 \in \omega$, from which on all A_n are equal, and then A_{n_0} is just the set of all t_i and u_j which are \sim_ω -related to s . Consequently,

$$\sum_{\substack{i=1 \\ s \sim_\omega t_i}}^k m_i = \sum_{\substack{i=1 \\ s \sim_{n_0} t_i}}^k m_i = \sum_{\substack{j=1 \\ s \sim_{n_0} u_j}}^l n_j = \sum_{\substack{j=1 \\ s \sim_\omega u_j}}^l n_j.$$

(4) Every congruence \approx on B is contained in \sim_ω . That is, our task is to prove the implication

$$t \approx t' \quad \text{implies} \quad \partial_n t = \partial_n t' \text{ for all } n \in \mathbb{N}.$$

As observed in (3), to be a congruence means that for every pair $t \approx t'$ of trees of the form

$$t = \begin{array}{c} m_1 \qquad \qquad m_k \\ \diagdown \qquad \diagup \\ \dots \\ \diagdown \qquad \diagup \\ t_1 \qquad \qquad t_k \end{array} \qquad t' = \begin{array}{c} m'_1 \qquad \qquad m'_\ell \\ \diagdown \qquad \diagup \\ \dots \\ \diagdown \qquad \diagup \\ t'_1 \qquad \qquad t'_\ell \end{array}$$

for every tree $s \in B$ the two sums below are equal:

$$\sum_{\substack{i=1 \\ s \approx t_i}}^k m_i = \sum_{\substack{j=1 \\ s \approx t'_j}}^{\ell} m'_j. \quad (6.2)$$

From this we derive $\partial_n t = \partial_n t'$ as follows.

Case $n = 0$ is trivial: $\partial_n t$ is the root-only tree.

Case $n = 1$. We are to prove $m_1 + \dots + m_k = m'_1 + \dots + m'_\ell$. For every $s = t_{i_0}$, $i_0 = 1, \dots, k$, we have the equality in (6.2). In case the left-hand sum is nonzero, we thus have some j with $t_{i_0} \approx t'_j$. And we can express $m_1 + \dots + m_k$ as the sum of all non-zero sums

$$\sum_{s \approx t_{i_0}} m_i$$

where i_0 ranges over a set of representatives (for \approx) of all indexes $1, \dots, k$ making the left-hand sum in (6.2) nonzero. By symmetry, this yields $m_1 + \dots + m_k = m'_1 + \dots + m'_\ell$, as desired.

Analogously for $n = 2$: here we take any t_{i_0} with $\sum_{t_{i_0} \approx t_i} m_i \neq 0$ and find a corresponding $t'_j \approx t_{i_0}$ (and vice versa). Then, by applying the case $n = 1$

to the pairs t_{i_0}, t'_j , we conclude that, for each t_{i_0} , $\sum_{\substack{i=1 \\ t_{i_0} \approx_1 t_i}}^k m_i = \sum_{\substack{j=1 \\ t_{i_0} \approx_1 t'_j}}^{\ell} m'_j$, then

$t \sim_2 t'$. Etc. ■

Definition 6.7 (See [17]). A commutative monoid M is called

- (a) *positive* if $a + b = 0$ implies $a = 0 = b$ and
- (b) *refinable* if $a_1 + a_2 = b_1 + b_2$ implies that there exists a 2×2 matrix with row sums a_1 and a_2 , respectively, and column sums b_1 and b_2 , respectively.

Theorem 6.8. *The following conditions on a commutative monoid M are equivalent:*

- (a) The functor \mathcal{M}_f weakly preserves pullbacks,
- (b) M is positive and refinable, and
- (c) whenever $a_1 + \dots + a_n = b_1 + \dots + b_k$, there exists an $n \times k$ -matrix whose vector of row sums is a_1, \dots, a_n and the vector of column sums is b_1, \dots, b_k .

In [17] this theorem is proved, except that in lieu of (a) weak preservation of *non-empty* pullbacks is requested. However, the functor \mathcal{M}_f has a unique distinguished point in the sense of V. Trnková [29], namely, the empty set $\emptyset \in \mathcal{M}_f X$. Since $\mathcal{M}_f \emptyset = \{\emptyset\}$, it follows from the result in [29] that \mathcal{M}_f preserves weak pullbacks iff it preserves the nonempty ones. Now for (a) \iff (b), see [17, Theorem 5.13], and concerning (b) \iff (c), Proposition 5.10 of *loc. cit.* states that refinability is equivalent to condition (c) with $n, k > 1$, and positivity of M is equivalent to condition (c) with $n > 1$ and $k = 0$. For $n = 1$, condition (c) is trivial.

Example 6.9 (See [17]). The Boolean monoid $P = \{0, 1\}$ and the monoids $(\mathbb{N}, +, 0)$ and $(\mathbb{N}, \cdot, 1)$ are positive and refinable. The cyclic group $\mathcal{C} = \{0, 1\}$ is refinable but not positive. For every lattice L the monoid $\mathcal{L} = (L, \vee, 0)$ is positive, and it is refinable iff L is a distributive lattice.

Remark 6.10. For M a positive and refinable monoid, the concepts of tree bisimulation and strong extensionality (see Definitions 2.4 and 2.6) immediately generalize to M -labeled trees. It is clear that tree bisimulations are closed under unions. Thus, for each M -labeled tree t there is a largest tree bisimulation on t . Since, by Theorem 6.8, \mathcal{M}_f preserves weak pullbacks, the largest bisimulation is an equivalence, and the corresponding quotient is a strongly extensional M -labeled tree \bar{t} .

Theorem 6.11. *Let M be a positive and refinable monoid. The coalgebra B_s of all strongly extensional, finitely branching M -labeled trees is final for \mathcal{M}_f .*

Remark. For the coalgebra B of M -labeled trees, all strongly extensional trees clearly form a subcoalgebra $m: B_s \hookrightarrow B$. We prove that the composite of m with the quotient homomorphism $q: B \rightarrow B/\sim_\omega$ is an isomorphism $q \cdot m: B_s \rightarrow B/\sim_\omega$. This proves that B_s is final.

Proof: Since $q \cdot m$ is a homomorphism of coalgebras, it is sufficient to prove that it is a bijection, then it is an isomorphism. In other words: we are to prove that B_s is a choice class of \sim_ω on the set B .

(1) Every tree t in B is Barr equivalent to its strongly extensional quotient tree \bar{t} (see Remark 6.10). Indeed, since the roots of t and \bar{t} are bisimilar, the two unique homomorphisms into the final coalgebra B/\sim_ω map them to the same element of B/\sim_ω . But for every tree t we know that the unique coalgebra homomorphism $f : t \rightarrow B/\sim_\omega$ takes the root of t to the \sim_ω -equivalence class $[t]$. Consequently, $t \sim_\omega \bar{t}$.

(2) If two strongly extensional trees are Barr equivalent, then they are equal. Instead, we prove in items (3) and (4) below that given extensional trees $t, s \in B$ then

if $t \sim_\omega s$ then t is tree bisimilar to s .

Thus, since \mathcal{M}_f weakly preserves pullbacks, this proves in case t and s are strongly extensional, that they are equal (up to isomorphism), see Remark 6.10.

(3) We consider the given trees $t \sim_\omega s$ as elements of the coalgebra B . We know that \sim_ω is the greatest congruence, hence, the greatest bisimulation on B . By Remark 6.5, there exists a matrix

$$m : B \times B \rightarrow M$$

such that

$$(a) \quad w_B(t, t') = \sum_{s' \in B} m(t', s') \quad \text{for all } t' \in B$$

$$(b) \quad w_B(s, s') = \sum_{t' \in B} m(t', s') \quad \text{for all } s' \in B, \text{ and}$$

$$(c) \quad m(t', s') \neq 0 \text{ implies } t' \sim_\omega s'.$$

Since M is positive, whenever $m(t', s') \neq 0$ we have $w_B(t, t') \neq 0$, that is, there exists a child x of the root x_0 of t with

$$t' = t_x \quad \text{and} \quad w_B(t, t') = w_t(x_0, x).$$

Analogously, $m(t', s') \neq 0$ implies $s' = s_y$ for some child y of the root y_0 of s with

$$w_B(s, s') = w_s(y_0, y).$$

Since t and s are extensional, the trees $t' \in B$ with $w_B(t, t') \neq 0$ are in bijective correspondence with the children x of x_0 in t via $x \mapsto t_x$. Analogously for s . Therefore we can translate (a)–(c) as follows:

$$(a^*) \quad w_t(x_0, x) = \sum_{y \in s} m(t_x, t_y) \quad \text{for all } x \in t$$

(b*) $w_s(y_0, y) = \sum_{x \in t} m(t_x, t_y)$ for all $y \in s$, and

(c*) $m(t', s') \neq 0$ implies that there exists a unique child x of x_0 in t and a unique child y of y_0 in s with $t_x \sim_\omega t_y$, $t' = t_x$ and $s' = s_y$.

(4) We prove that given trees $\bar{t}, \bar{s} \in B$ with $\bar{t} \sim_\omega \bar{s}$, it follows that the relation $R \subseteq \bar{t} \times \bar{s}$ defined recursively by

$x R y$ iff $\bar{t}_x \sim_\omega \bar{s}_y$ and x and y are roots or have R -related parents

is a tree bisimulation. If $x R y$ then put $t := \bar{t}_x$ and $s := \bar{s}_y$ and let $\bar{m}: \bar{t} \times \bar{s} \rightarrow M$ be the following matrix

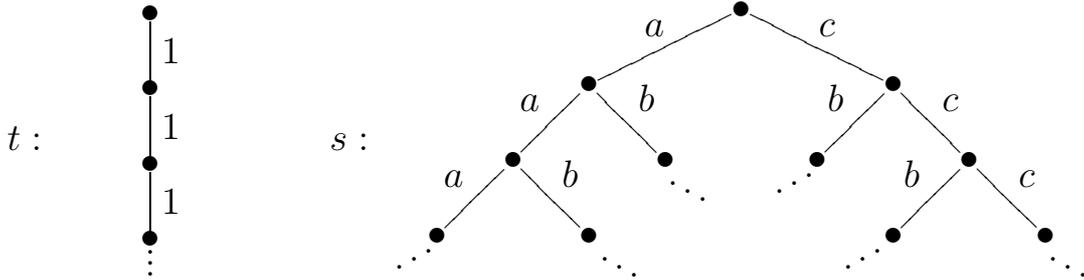
$$\bar{m}(x', y') = \begin{cases} m(t_{x'}, s_{y'}) & \text{if } x' \text{ is a child of } x \text{ and } y' \text{ a child of } y \\ 0 & \text{else} \end{cases}$$

The property (c*) tells us that \bar{m} is obtained from the matrix m by removing all zero columns and zero rows. Therefore, (a*) and (b*) imply that \bar{m} has the desired row and column sums:

$$\begin{aligned} w_{\bar{t}}(x, x') &= \sum_{y' \in \bar{s}} \bar{m}(x', y') && \text{for all } x \in \bar{t} \\ w_{\bar{s}}(y, y') &= \sum_{x' \in \bar{t}} \bar{m}(x', y') && \text{for all } y' \in \bar{s}. \end{aligned}$$

Moreover, by definition, $\bar{m}(x', y') \neq 0$ only if x' and y' have R -related parents (or are the roots) and $m(t_{x'}, t_{y'}) \neq 0$; and, by (c), $m(t_{x'}, t_{y'}) \neq 0$ implies that $t_{x'} \sim_\omega t_{y'}$. Therefore $\bar{m}(x', y') \neq 0$ implies $x' R y'$. ■

Example 6.12. The above theorem does not generalize to all positive monoids. To see this, consider the monoid $\mathcal{L} = (L, \vee, 0)$ for the lattice $\{0, a, b, c, 1\}$ where a, b, c are pairwise incomparable. Then strongly extensional finitely branching \mathcal{L} -labeled trees do not form a final coalgebra, since they are not a choice class of the Barr equivalence. The following trees are easily seen to be Barr equivalent:



Here s has as vertices the binary words, and the weights are, for all $x \in \{0, 1\}^*$, defined by $w_s(x0, x00) = a$, $w_s(x1, x11) = c$ and $w_s(x0, x01) = b = w_s(x1, x10)$. It is obvious that t is strongly extensional. To prove that so is s , let $R \subseteq s \times s$ be a tree bisimulation. Using the conditions (a)–(c) in the preceding proof it is easy to verify that $R \subseteq \Delta_s$.

7. Conclusions and related work

A new description of the final coalgebra and/or the final chain of the power-set functor and its “relatives” has been presented in our paper. For example, for the finite power-set functor \mathcal{P}_f we have given a short proof of Worrell’s description of the final coalgebra as the set of all finitely branching, strongly extensional trees (cf. [30]). And we provided an alternative description as the set of all hereditarily finite modal theories. We also described the step ω , $\mathcal{P}_f^\omega 1$, of the final chain as the set of all saturated, strongly extensional trees (which is related to Worrell’s description as all compactly branching, strongly extensional trees). Related descriptions were provided by S. Abramsky [1], A. Kurz and D. Pattinson [20] and by J. Rutten [26, Theorem 7.4].

The above saturated trees were also proved to precisely correspond to the modally saturated trees of K. Fine [14]. We generalized saturatedness to α -saturatedness and proved that the final chain of the power-set functor \mathcal{P} can be described for all α of cofinality ω by saying that $\mathcal{P}^\alpha 1$ consists of all strongly extensional α -saturated trees; for such ordinals, the connecting maps $w_{\alpha, \beta}$ were proved to be surjective, and the canonical maps from graphs were proved to be given by the α -saturation of the tree expansions. We also proved that for all infinite regular cardinals λ , the smallest ordinal for the convergence of the final chain of \mathcal{P}_λ is $\lambda + \omega$, and it is λ^+ in the case where λ is an infinite singular cardinal.

General ordinals present a difficulty. Thus we have at present no analogous description of $\mathcal{P}^{\omega_1} 1$. We leave to future work an analysis of this issue in the light of the results concerning completeness in Forti and Honsell [16, especially Lemma 2.2], and also R. Lazić and A. Roscoe [21, Theorems 14 and 15].

Another direction generalizing the functor \mathcal{P}_f was taken by H.-P. Gumm and T. Schröder [17]. They introduced the functor \mathcal{M}_f of finite multisets with multiplicities from a given commutative monoid. We have described its final coalgebra: it consists of all finitely branching strongly extensional M -labeled trees. This holds for all positive and refinable monoids. Our proof

is substantially different from Worrell's, since it is based on congruences on the coalgebra of all extensional trees. We would like to generalize our work on saturated trees to the case of functors \mathcal{M}_f . And we plan to apply our methods to probabilistic transition systems.

References

- [1] Abramsky, S.: A Cook's tour of the finitary non-wellfounded sets. In S. Artemov et al (eds), *We Will Show Them: Essays in honour of Dov Gabbay*, College Publications, Vol. 1, 1–18, 2005
- [2] Aczel, P.: *Non-well-founded sets*. Lecture Notes 14, CSLI, Stanford 1988
- [3] Aczel, P., Mendler, P. F.: A final coalgebra theorem. *Lecture Notes in Comput. Sci.* **389**, Springer 1989, 357–365
- [4] Adámek, J.: Free algebras and automata realizations in the language of categories. *Comment. Math. Univ. Carolinae* **15** (1974) 589–609
- [5] Adámek, J., Koubek, V.: On the greatest fixed point of a set functor. *Theoret. Comput. Sci.* **150** (1995) 57–75
- [6] Adámek, J., Milius, S., Velebil, J.: On coalgebra based on classes. *Theoret. Comput. Sci.* **316** (2004) 685–705
- [7] Adámek, J., Milius, S., Moss, L. S., Sousa, L.: Power-set Functors and Saturated Trees. *Proc. CSL'11, LIPICS* **12** (2011), 5–19
- [8] Adámek, J., Trnková, V.: *Automata and algebras in categories*. Kluwer Academic Publ., Dordrecht 1990
- [9] Adámek, J., Trnková, V.: Initial algebras and terminal coalgebras in many-sorted sets. *Math. Str. Comput. Sci.* **21** (2011), 481–509
- [10] Barr, M.: Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.* **114** (1993) 299–315
- [11] Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge University Press 2001
- [12] Borceux, F.: *Handbook of categorical algebra 2*. Cambridge University Press 1994
- [13] Fagin, R., Vardi, M. Y.: An Internal Semantics for Modal Logic: Preliminary Report. In *Proceedings of STOC'85*, 305–315 (1985)
- [14] Fine, K.: Normal forms in modal logic. *Notre Dame J. Formal Logic* **16** (1975), 229–237
- [15] Fine, K.: Some Connections Between Elementary and Modal Logic, in S. Kanger (ed), *Proc. of the Third Scandinavian Logic Symposium*. Amsterdam: North Holland, 1975, 15–31
- [16] Forti, M., Honsell, F.: Models of self-descriptive set theories. in *Partial Differential Equations and the Calculus of Variations*, Vol. I, Birkhäuser Boston, 1989, 473–518.
- [17] Gumm, H.-P., Schröder T.: Monoid-labelled transition systems. *Electron. Notes Theor. Comput. Sci.* **44** (2001)
- [18] Jech, T.: *Set theory*, Academic Press, Boston, 1978.
- [19] Joyal, A.: Foncteurs analytiques et espèce de structure. *Lecture Notes in Math.* **1234**, Springer 1986, 126–159
- [20] Kurz, A., Pattinson, D.: Coalgebraic modal logic of finite rank. *Math. Structures Comput. Sci.* **15:3** (2005), 453–473
- [21] Lazić, R. S., Roscoe, A. W.: On transition systems and non-well-founded sets. *Papers on general topology and applications*, Ann. New York Acad. Sci., 1996, 238264.
- [22] Lambek, J.: A fixpoint theorem for Complete Categories, *Math. Z.* **103** (1968), 151–161
- [23] Mac Lane, S.: *Categories for the Working Mathematician*, Second Ed. Springer-Verlag, 1998.

- [24] Malitz, J.: *Set theory in which the axiom of foundation fails*, PhD thesis, University of California, Berkeley, 1976.
- [25] Moss, L. S.: Finite models constructed from canonical formulas, *J. Philos. Logic* **36** (2007), 605–640
- [26] Rutten, J.: A calculus of transition systems, *CSLI lecture notes* **53**, 1995.
- [27] Rutten, J.: Universal coalgebra: a theory of systems, *Theoret. Comput. Sci.* **249** (2000), 3–80
- [28] Schwencke, D.: Coequational logic for accessible functors. *Inform. and Comput.* **208** (2010), 1469–1489
- [29] Trnková, V.: On descriptive classification of set functors I, *Comment. Math. Univ. Carolinae* **12** (1971), 143–174
- [30] Worrell, J.: On the final sequence of a finitary set functor. *Theoret. Comput. Sci.* **338** (2005) 184–199

8. Appendix: saturations of trees for all ordinals

We prove Theorem 3.9: for each tree t and each ordinal α , there is a unique α -saturation of t .

Proposition 8.1. *For every ordinal α , there is a set S_α such that every tree is \sim_α to some tree in S_α .*

Proof: By induction on α . S_0 is the singleton of a one-point tree. Given S_α , let $S_{\alpha+1}$ be the set of tree-tuplings of sets of trees from S_α . Given S_α for $\alpha < \lambda$, first let $T = \prod_{\alpha < \lambda} S_\alpha$. Then let

$$T' = \{f \in T : \text{there is a tree } t \text{ so that } t \sim_\alpha f(\alpha) \text{ for all } \alpha < \lambda\}.$$

Finally, let S_λ be any set of trees with the property that for every $f \in T'$ there is some $t \in S_\lambda$ such that for all $\alpha < \lambda$, $t \sim_\alpha f(\alpha)$. To check that this works, let t be any tree. Let $f \in T$ be such that for all $\alpha < \lambda$, $t \sim_\alpha f(\alpha)$. By construction $f \in T'$. So for some $t' \in S$, $t' \sim_\alpha f(\alpha)$ for all $\alpha < \lambda$. Then $t \sim_\alpha t'$ for all $\alpha < \lambda$, so $t \sim_\lambda t'$. ■

Definition 8.2. Let s and t be trees. Let \mathcal{T} denote the category of trees and maps that preserve the root and the edges. (Of course, \mathcal{T} has more morphisms than the category `Tree` used in Notation 2.10.) A \mathcal{T} -morphism $f : s \rightarrow t$ is an \sim_α -embedding if f is injective on vertices, and for all $x \in s$, $s_x \sim_\alpha t_{f(x)}$.

An \sim_α -chain is a functor

$$F : \beta \rightarrow \mathcal{T}$$

for some ordinal β , such that for every $\gamma < \delta < \beta$, $Ff_{\gamma,\delta}$ is an \sim_α -embedding, where $f_{\gamma,\delta} : \gamma \rightarrow \delta$ is the morphism in β . We write t^δ for $F(\delta)$.

Lemma 8.3. *Let $F : \beta \rightarrow \mathcal{T}$ be an \sim_α -chain, and let t^* be its colimit, with injections $i_\gamma : t^\gamma \rightarrow t^*$. Then each morphism i_γ is an \sim_α -embedding.*

Proof: It is easy to check that the colimit t^* is formed as in **Set**, by taking the disjoint union of the trees t^γ and making a tree in the natural way, by putting an edge from x to y if for some $\gamma < \beta$, x and y both belong to t^γ , and x has y as a child in t^γ . Moreover, each colimit injection i_γ is an inclusion.

We check by induction on $\delta \leq \alpha$ that each map $i_\gamma : t^\gamma \rightarrow t^*$ is a \sim_δ -embedding. For 0, this is obvious, and the induction step for limit ordinals is trivial. So assume our result for $\delta < \alpha$, and then let us check that for all γ and all $x \in t^\gamma$, $t_x^\gamma \sim_{\delta+1} t_{i_\gamma(x)}^*$. First, let y be a child of x in t^γ . Then by induction hypothesis, $t_y^\gamma \sim_\delta t_{i_\gamma(y)}^*$. Since $i_\gamma(y)$ is a child of $i_\gamma(x)$ in t^* , we have verified half of what we need. For the other half, let y be any child of $i_\gamma(x)$ in t^* . For some $\gamma' < \beta$, and some $x', y' \in t^{\gamma'}$, $i_{\gamma'}(x') = i_\gamma(x)$, y' is a child of x' in $t^{\gamma'}$, and $i_{\gamma'}(y') = y$. We may assume that $\gamma < \gamma'$, since the case $\gamma' \leq \gamma$ is similar. Recall that $i_{\gamma, \gamma'}$ is an \sim_α -embedding. Thus $t_x^\gamma \sim_\alpha t_{i_{\gamma, \gamma'}(x)}^{\gamma'}$. And since $\delta + 1 \leq \alpha$, $t_x^\gamma \sim_{\delta+1} t_{i_{\gamma, \gamma'}(x)}^{\gamma'}$. Moreover, $i_{\gamma, \gamma'}(x) = x'$; this is because $i_{\gamma'}(i_{\gamma, \gamma'}(x)) = i_\gamma(x) = i_{\gamma'}(x')$, and $i_{\gamma'}$ is injective. So we have $t_x^\gamma \sim_{\delta+1} t_{x'}^{\gamma'}$. Thus there is some child y of x in t^γ such that

$$t_y^\gamma \sim_\delta t_{y'}^{\gamma'} \sim_\delta t_{i_{\gamma'}(y')}^*.$$

(The last equality uses the induction hypothesis.) And $t_{i_{\gamma'}(y')}^* = t_y^*$. This completes the proof. \blacksquare

With this lemma, we prove Theorem 3.9.

Theorem 3.9: The uniqueness comes from Lemma 3.8, and so we only argue for the existence. We use induction on α . For $\alpha = 0$, we take a one-point tree. Assuming our result for α , we get it for $\alpha + 1$ by tree-tupling the trees $t_{x_\alpha}^*$ as x ranges over the children of the root of t .

For the limit step, we first need to define a certain infinite cardinal κ . As an ordinal κ will be the length of a certain \sim_α -chain. Let S_α be as in Proposition 8.1, and let ρ be the (ordinal) maximum of α and the cardinality of S_α . Let $\kappa = \aleph_{\rho+1}$. Here are the properties of κ which we shall need:

- (1) Being a successor cardinal, κ is regular.
- (2) $\kappa \geq \rho + 1 > \alpha$.
- (3) $\kappa \geq \rho + 1 > |S_\alpha|$.

Fix a map

$$c : \kappa \rightarrow S_\alpha$$

which is surjective in the strong sense that for every tree $u \in S_\alpha$, $c^{-1}(u)$ is unbounded.

Fix a tree t . Define an \sim_α -chain of trees t^β by recursion on $\beta < \kappa$ as follows. First, $t^0 = t$. Fix $\beta \leq \kappa$, and assume that we have trees t^β and \sim_β -embeddings $f_{\gamma,\beta} : t^\gamma \rightarrow t^\beta$ for $\gamma \leq \beta$. If β is a limit ordinal, then we take t^β to be the colimit; by Lemma 8.3 the evident injections are \sim_α -embeddings. The main work concerns the successor case, defining $t^{\beta+1}$ from t^β .

We consider the tree $c(\beta) \in S_\alpha$. Consider all nodes x in t^β such that for all $\gamma < \alpha$ there is some child y of x in t^β so that $c(\beta) \sim_\gamma t_y^\beta$. If there are no such nodes, set $t^{\beta+1} = t^\beta$. Otherwise, for each such x , add a new subtree below x isomorphic to $c(\beta)$. This is $t^{\beta+1}$. We take $f_{\beta,\beta+1}$ to be the inclusion of trees.

We claim that $f_{\beta,\beta+1}$ is an \sim_α -embedding. For this, we show by induction on $\delta < \alpha$ that $f_{\beta,\beta+1}$ is a \sim_δ -embedding. Because this map $f_{\beta,\beta+1}$ is an inclusion, we shall drop it from the notation. The steps for 0 and for limit ordinals are immediate, and we are left with the successor case. Assume that the inclusion is a \sim_δ -embedding; we show that it is a $\sim_{\delta+1}$ -embedding. Fix a node $x \in t^\beta$. Half is easy: for every child y of x in t^β , there is a child y of x in $t^{\beta+1}$ such that $t_y^\beta \sim_\delta t_y^{\beta+1}$. (This is by the induction hypothesis.) In the more interesting direction, let y be a child of x in $t^{\beta+1}$. If y is a node in t^β , then we are easily done by the induction hypothesis. Otherwise, the tree $s = c(\beta)$ has the property that for all $\gamma < \alpha$ there is some child z of x so that $s \sim_\gamma t_z^\beta$, and moreover our construction arranged that, $t_y^{\beta+1} = s$. Taking γ to be δ , we see that there is some child z of x in t^β with $t_y^{\beta+1} \sim_\delta t_z^\beta$, just as desired. This concludes our inductive step.

The claim shown, this concludes the definition of our chain. The tree that we are after is its colimit, t^κ . Let x be the root of $t = t_0$. Since morphisms of our category \mathcal{T} preserve the root, $t = t_x \sim_\alpha t_{i_0(x)}^\kappa = t^\kappa$. Moreover, t^κ need not be strongly extensional, but once we show that it is α -saturated, its strongly extensional quotient will also have this property, and we shall be done.

To check that t^κ is α -saturated, let s be any tree, let x be a node in t^κ and assume that for all $\beta < \alpha$ there is some child y_β of x such that $s \sim_\beta t_{y_\beta}^\kappa$. We might as well assume that s belongs to S_α : it is \sim_α to some tree in S_α , and then replacing s by this other tree and carrying out the coming argument will show what we want. By regularity of κ , there is some $\delta < \kappa$ so that for

all $\beta < \alpha$ there is some child y_β of x such that y_β is a node in t^δ and $s \sim_\beta t_{y_\beta}^\kappa$. (In more detail, for each $\beta < \alpha$, there is a least ordinal ρ_β containing some y_β with the property we want. Then by regularity, $\sup_\beta \rho_\beta < \kappa$.) By Lemma 8.3, $t_{y_\beta}^\kappa \sim_\alpha t_{y_\beta}^\delta$ for all $\beta < \alpha$. By definition of c , let γ be such that $\delta < \gamma < \kappa$ and $c(\gamma) = s$. Using Lemma 8.3 again, we have that for all $\beta < \alpha$, $t_{y_\beta}^\delta \sim_\alpha t_{y_\beta}^\gamma$. Since $\alpha > \beta$, $s \sim_\beta t_{y_\beta}^\gamma$ for all $\beta < \alpha$. Our construction has arranged that $t^{\gamma+1}$ contains a child y of x such that $t_y^{\gamma+1} \sim_\alpha s$. By Lemma 8.3 one last time, $t_y^\kappa \sim_\alpha t_y^{\gamma+1} \sim_\alpha s$. ■

JIŘÍ ADÁMEK

INSTITUT FÜR THEORETISCHE INFORMATIK, TECHNISCHE UNIVERSITÄT BRAUNSCHWEIG, GERMANY

E-mail address: `adamek@iti.cs.tu-bs.de`

PAUL LEVY

SCHOOL OF COMPUTER SCIENCE, UNIVERSITY OF BIRMINGHAM, UK

E-mail address: `P.B.Levy@cs.bham.ac.uk`

STEFAN MILIUS

INSTITUT FÜR THEORETISCHE INFORMATIK, TECHNISCHE UNIVERSITÄT BRAUNSCHWEIG, GERMANY

E-mail address: `mail@stefan-milius.eu`

LAWRENCE S. MOSS

DEPARTMENT OF MATHEMATICS, INDIANA UNIVERSITY, BLOOMINGTON, IN, USA

E-mail address: `lsm@cs.indiana.edu`

LURDES SOUSA

POLYTECHNIC INSTITUTE OF VISEU, PORTUGAL & CENTRE FOR MATHEMATICS OF THE UNIVERSITY OF COIMBRA, PORTUGAL

E-mail address: `sousa@mat.estv.ipv.pt`