Chapter 3

Final sequences for accessible endofunctors

3.1 Introduction

The theorems of Aczel and Mendler (theorem 2.3.4), and Barr (cf. section 5.3), guarantee the existence of final coalgebras for a wide class of endofunctors on Set. In each case one can, in theory, derive a general recipe for constructing final coalgebras from the proof. However, in practice it is very difficult to learn much about the structure of the final coalgebra of a specific endofunctor. On this question Moss [65] says:

> Ever since people looked at final coalgebras as models of intensional phenomena, there was a question of getting representations of the final coalgebras in terms of some sort of entities that served as approximations.

Aczel and Mendler obtain a final coalgebra as a quotient (by bisimilarity) of a coproduct of a set of coalgebras. Barr shows that if a set functor T is *accessible* (cf. section 3.2.2), then $\mathsf{Coalg}(T)$ has a set of generators, and uses the Special Adjoint Functor Theorem (whose proof also involves a quotient-of-a-sum construction) to derive the existence of a final coalgebra. Work on the problem of more concrete constructions of final coalgebras includes Mislove et al. [61], in which a domain for approximating hereditarily finite non-well-founded sets is obtained, and Moss [65], where elements of final coalgebras are approximated by formulas of a *coalgebraic logic*.

In this chapter we adopt the approach of Adámek and Koubek [6] (and of Barr in another paper [14]). We consider an endofunctor T on a category $\mathcal{C}$ with limits of co-chains, and define the *final sequence* of T: an ordinal-indexed sequence $(A_\alpha)_\alpha$ of objects of $\mathcal{C}$, with arrows $f^\alpha_\beta : A_\alpha \to A_\beta$ for $\beta \leq \alpha$. Briefly, this is defined by $A_{\alpha+1} = T(A_\alpha)$, and $A_\lambda = \mathrm{Lim}_{\alpha<\lambda} A_\alpha$ for λ a limit ordinal. Fuller details are given in the next section. It is shown in [6] that if this sequence stabilizes at some α, in the sense that $f^{\alpha+1}_\alpha$ is an isomorphism, then $(A_\alpha, (f^{\alpha+1}_\alpha)^{-1})$ is a final T-coalgebra. This generalizes the iterative construction of the greatest fixed point of a monotone function f on a complete lattice as the stabilizing value of the ordinal-indexed sequence $(a_\alpha)_\alpha$, where $a_\alpha = \sqcap_{\beta<\alpha} f(a_\beta)$. On the other hand, the theorems of [3, 13] generalize the construction of the greatest fixed point of f as the join of all post-fixed points (the formula (2.5) shows that Aczel's Special Final Coalgebra Theorem can also be seen in this way).

For set functors accessibility seems to be the common denominator amongst the hypotheses involved in the various final coalgebra theorems in the literature, e.g., being bounded (definition 2.3.5), or set based (definition 2.3.3). In a recent contribution to the category theory mailing list Adámek has shown that the assumption of boundedness of a set functor is equivalent to accessibility. Aczel and Mendler actually consider endofunctors on the category of classes and class functions which are set based, but, as Barr observes, this basically amounts to assuming an inaccessible cardinal[1] κ and considering κ-accessible set functors preserving the subcategory of sets of size less than or equal to κ. Finally, Barr's final coalgebra theorem, though stated for set functors, is not specific to Set: every accessible endofunctor on a locally presentable category has a final coalgebra, cf. section 5.3. Thus, we are led to study the final sequences of accessible endofunctors on locally presentable categories.

A good proportion of the literature deals with coalgebras of ω^{op}-continuous endofunctors. We saw in chapter 2 that for endofunctors on the category of complete metric spaces ω^{op}-continuity is guaranteed by quite a mild (and readily verified) criterion, namely being locally contractive. Similarly, for endofunctors on the category of cpos, ω^{op}-continuity is guaranteed by local continuity. There are, however, set functors which are ω-accessible, so have a final coalgebra, but whose final sequences do not stabilize at ω; one example is the finite-powerset functor $\wp$, cf. [6].

Adámek and Koubek [6] show that for set functors the mere existence of a final coalgebra is sufficient to ensure convergence of the final sequence. They do not in general give bounds for convergence, although in the case of $\wp$ they show convergence at ω_1—the first uncountable ordinal. The main point we wish to make in this chapter is that for finitary (ω-accessible) endofunctors on Set the construction of the final coalgebra via the final sequence is a two-stage process, each of which is finitary. More precisely, the final sequence stabilizes

[1] A cardinal $\kappa > \omega$ is inaccessible if $\lambda < \kappa$ implies that $2^\lambda < \kappa$. The fact that we talk about accessibility for cardinals on the one hand, and functors and categories on the other, is purely coincidental!

at $\omega + \omega$. A corresponding result holds if we replace ω with any regular cardinal κ.

The first stage of the final coalgebra construction can be seen as taking a Cauchy completion of the initial algebra, whilst the next stage can be seen as pruning this to obtain the final coalgebra. We show that, in general, the ω-th iterate in the final sequence of a finitary set functor is always a final coalgebra—not necessarily of T, but certainly of the 'completion' of T to an endofunctor on the category of complete ultrametric spaces and non-expansive maps. We give two examples to support these intuitions. In particular, for the finite-powerset functor $\wp$, the first ω steps in the final sequence construct the set of *compactly branching*, strongly extensional, possibly infinite-depth trees. In the next ω steps of the final sequence, these trees are 'pruned', one level at a time, until we reach the set of *finitely branching*, strongly extensional, possibly infinite-depth trees.

The bound on the convergence of the final sequence of a κ-accessible set functor relies on the fact that injectives with non-empty domain are split monic in Set. Thus, the argument given here doesn't generalize to arbitrary locally presentable categories in a straightforward manner (though it does also apply to Vect, the category of vector spaces over some fixed field and linear maps between them, where monos also split). However, we do prove that if T is an accessible endofunctor on a locally presentable category, and if T preserves monos, then the final sequence converges.

3.2 Final sequences and local presentability

In this section we define the final sequence of an endofunctor, and go on to recall the definition of a locally presentable category. The former is defined in [14], whilst [8] is our reference for accessibility and local presentability.

3.2.1 Final Sequences

Let $\mathcal{C}$ be a category with limits of all co-chains, and T an endofunctor on $\mathcal{C}$. The final sequence of T is an ordinal-indexed sequence of objects $(A_\alpha)_\alpha$, with maps $(f^\beta_\gamma : A_\beta \to A_\gamma)_{\gamma \leq \beta}$, uniquely defined by the following conditions (where $\delta \leq \gamma \leq \beta$):

- FS-1 $A_{\beta+1} = T(A_\beta)$;
- FS-2 $f^{\beta+1}_{\gamma+1} = T(f^\beta_\gamma)$;
- FS-3 $f^\beta_\beta = \mathrm{id}$;
- FS-4 $f^\beta_\delta = f^\gamma_\delta \cdot f^\beta_\gamma$;

- FS-5 if β is a limit ordinal, the cone $(f_\gamma^\beta : A_\beta \to A_\gamma)_{\gamma \leq \beta}$ is limiting.

We define the sequence by ordinal induction, checking at each stage that conditions [FS-1] to [FS-5] hold for the portion of the sequence already defined. For clarity we divide the successor case in the induction into two subcases.

(1) Case: α is a limit ordinal. We define $(f_\beta^\alpha : A_\alpha \to A_\beta)_{\beta<\alpha}$ to be the limit of the co-chain $(A_\beta)_{\beta<\alpha}$, and we set $f_\alpha^\alpha = id$. Conditions [FS-1] to [FS-5] are easily verified.

(2a) Case: $\alpha = \lambda + 1$, where λ is a limit ordinal. We define $A_\alpha = T(A_\lambda)$. A_λ is the limit of the co-chain $(A_\gamma)_{\gamma<\lambda}$, so we can define f_λ^α by giving the composites $f_\gamma^\lambda \cdot f_\lambda^\alpha$, for $\gamma < \lambda$, and proving that they form a cone over the co-chain $(A_\gamma)_{\gamma<\lambda}$. We define $f_\gamma^\lambda \cdot f_\lambda^\alpha = f_\gamma^{\gamma+1} \cdot T(f_\gamma^\lambda)$. This does indeed define a cone over the A_γ since

$$f_\delta^\gamma \cdot f_\gamma^{\gamma+1} \cdot T(f_\gamma^\lambda) = f_\delta^{\gamma+1} \cdot T(f_\gamma^\lambda) = f_\delta^{\delta+1} \cdot f_{\delta+1}^{\gamma+1} \cdot T(f_\gamma^\lambda)$$

$$= f_\delta^{\delta+1} \cdot T(f_\delta^\gamma) \cdot T(f_\gamma^\lambda) = f_\delta^{\delta+1} \cdot T(f_\delta^\lambda).$$

For $\gamma < \lambda$ we set $f_\gamma^\alpha = f_\gamma^\lambda \cdot f_\lambda^\alpha$. The requisite conditions are all immediate apart from [FS-2] for $\beta = \lambda$; but this follows from

$$f_{\gamma+1}^{\beta+1} = f_{\gamma+1}^\alpha = f_{\gamma+1}^\beta \cdot f_\beta^\alpha = f_{\gamma+1}^{\gamma+2} \cdot T(f_{\gamma+1}^\beta) = T(f_\gamma^{\gamma+1} \cdot f_{\gamma+1}^\beta) = T(f_\gamma^\beta).$$

(2b) Case: $\alpha = \beta + 2$ for some β. Set $A_\alpha = T(A_{\beta+1})$, $f_\alpha^\alpha = id$, and $f_{\beta+1}^\alpha = T(f_\beta^{\beta+1})$. For $\gamma < \beta + 1$, we set $f_\gamma^\alpha = f_\gamma^{\beta+1} \cdot f_{\beta+1}^\alpha$. Again, it is easy to verify the conditions [FS-1] to [FS-5].

The following results can be found in [6, 14].

Theorem 3.2.1 Suppose the final sequence of T stabilizes at κ—in the sense that $f_\kappa^{\kappa+1}$ is an isomorphism. Then $(A_\kappa, (f_\kappa^{\kappa+1})^{-1})$ is a final T-coalgebra. □

Proposition 3.2.2 Let $\beta < \alpha$ be such that f_β^α is an isomorphism. Then $f_\beta^{\beta+1}$ is an isomorphism. □

3.2.2 Accessibility and Local Presentability

Let f be a monotone function on a complete lattice L. The proof of the convergence of the sequence $a_\alpha = \sqcap_{\beta<\alpha} f(a_\beta)$ depends on the fact that L is a set; this result does not generalize without caveat to final sequences of endofunctors. In particular, the final sequence of the powerset functor on Set will never stabilize, since this functor does not have any fixed points.

In seeking an analogous result for endofunctors on a complete category $\mathcal{C}$ we have to take account of the fact that $\mathcal{C}$ may in general be a large category. Now there are large categories on which every endofunctor has a fixed point: the so-called *algebraically complete* categories. Examples include the category of countable sets and the category of countably dimensional vector spaces. However this class of categories is quite limited; in particular, an algebraically complete category with all small limits must be a preorder, cf. [7].

We take a different approach: we consider categories with good properties and good endofunctors on them. A useful hypothesis in this situation is that $\mathcal{C}$ is, in a suitable sense, generated by a small subcategory. This is the content of our assumption of local presentability of $\mathcal{C}$. The concepts of filtered colimit, presentable object, locally presentable category, and accessible functor, below, are respectively category-theoretic generalizations of the notions of directed join, compact element (of a partial order), algebraic lattice, and continuous function between partial orders.

Definition 3.2.3 Let κ be an infinite regular cardinal (i.e, κ is strictly greater than the sum of less than κ smaller cardinals). Then a small category $\mathcal{I}$ is *κ-filtered* if for any category $\mathcal{J}$ of cardinality less than κ, any diagram $D : \mathcal{J} \to \mathcal{I}$ has a cocone over it. For short we call a colimit of a diagram whose domain is κ-filtered a κ-filtered colimit. □

Definition 3.2.4 An object K of a category $\mathcal{K}$ is *κ-presentable* if the functor $\mathrm{Hom}(K, -) : \mathcal{K} \to \mathsf{Set}$ preserves κ-filtered colimits. We say that $\mathcal{K}$ is *locally κ-presentable* (locally finitely presentable if $\kappa = \omega$) if it is cocomplete and there is a set $\mathcal{A}$ of κ-presentable objects such that every object is a κ-filtered colimit of objects from $\mathcal{A}$. Finally $\mathcal{K}$ is *locally presentable* if it is locally κ-presentable for some κ. □

Spelt out in elementary terms, an object K is κ-presentable precisely when, given a κ-filtered diagram

$$D : \mathcal{D} \to \mathcal{K}$$

with colimiting cocone $(D_d \xrightarrow{c_d} C)$, and given a morphism $f : K \to C$, then: (i) f factors through some c_d, i.e., $f = c_d \cdot g$ for some $g : K \to D_d$; (ii) if f also factors through $c_{d'}$ via g', then there exist arrows $d \to d''$ and $d' \to d''$ of $\mathcal{D}$ such that $D(d \to d'') \cdot g = D(d' \to d'') \cdot g'$.

Definition 3.2.5 Let $\mathcal{K}$ and $\mathcal{L}$ be locally κ-presentable categories. A functor $T : \mathcal{K} \to \mathcal{L}$ is *κ-accessible* if it preserves κ-filtered colimits. We say T is *accessible* if it is κ-accessible for some κ. □

It follows easily that in a locally κ-presentable category $\mathcal{K}$ there is, up-to isomorphism, only a set of κ-presentable objects. We will assume a set of representatives of all κ-presentable objects and denote it $\mathrm{Pres}_\kappa\mathcal{K}$.

Examples of locally finitely presentable categories include Set, Pos (posets and monotone maps), Graph (graphs and graph homomorphisms), and Vect (vector spaces and linear transformations over some fixed field). A category which fails to be locally finitely presentable is Cpo, cf. [8]. In Set the ω-presentable (or finitely presentable) objects are precisely the finite sets; Set is locally finitely presentable since every set is the union of the filtered diagram of all its finite subsets. The ω-accessible (or finitary) endofunctors on Set are those that are completely determined by their behaviour on the subcategory of finite cardinals. A simple criterion for a set functor to be finitary, from [7], is:

- for every set X and $x \in FX$, there exists a finite $Y \subseteq X$ and $y \in FY$ such that $x = F(Y \hookrightarrow X)y$.

Finally, recall that an epimorphism e is said to be *strong* if whenever we have a commuting square as below, with m a monomorphism, then there is a diagonal fill-in map $h\colon Q \to A$ such that $f = h \cdot e$.

$$\begin{tikzcd} P \arrow[r, "e"] \arrow[d, "f"'] & Q \arrow[d, "g"] \\ A \arrow[r, "m"'] & B \end{tikzcd} \tag{3.1}$$

Every arrow in a locally presentable category can be factored as a strong epi followed by a mono.

3.3 Final sequences in locally presentable categories

In this section we show that if T is an accessible endofunctor on a locally presentable category $\mathcal{K}$, and if T preserves monos, then its final sequence $\{A_\alpha, f^\alpha_\beta\}$ converges. The first step is to prove that for some cardinal λ the map $f^{\lambda+1}_\lambda$ is monic. This is reminiscent of the proof of a result in [6], stating that if a set functor has two fixed points of successor cardinalities, then its final sequence converges.

Example 3.3.1 To motivate the general development we consider the final sequence $\{A_\alpha, f^\alpha_\beta\}$ of the finite-powerset functor $\wp\colon \mathsf{Set} \to \mathsf{Set}$. Both [6] and [83] show that, even though this functor is finitary, the final coalgebra cannot be constructed by the usual ω^{op}-limit, i.e. the final sequence doesn't stop in ω steps. In fact, the projection $f^{\omega+1}_\omega$ is not surjective.

To see this, notice that the sequence (x_n), where $x_{n+1} = A_n$, lies in A_ω. Suppose for a contradiction that $(x_n) = f_\omega^{\omega+1}(y)$ for some $y \in A_{\omega+1}$, then

$$A_n = f_{n+1}^\omega((x_n)) = f_{n+1}^\omega f_\omega^{\omega+1}(y) = \wp(f_n^\omega)(y).$$

But this implies that $|A_n| \leq |y|$ for all $n < \omega$, which is impossible.

Thus, $f_\omega^{\omega+1}$ is not surjective. However, it is easily seen to be injective. Let $S = \{d_1, ..., d_l\} \subseteq A_\omega$, $T = \{e_1, ..., e_m\} \subseteq A_\omega$ and suppose $f_\omega^{\omega+1}(S) = f_\omega^{\omega+1}(T)$; then

$$\wp(f_n^\omega)(S) = f_{n+1}^{\omega+1}(S) = f_{n+1}^\omega\, f_\omega^{\omega+1}(S) = f_{n+1}^\omega\, f_\omega^{\omega+1}(T) = f_{n+1}^{\omega+1}(T) = \wp(f_n^\omega)(T)$$

for all $n < \omega$. Now pick $d_i \in S$; since T is finite there exists $e_j \in T$ such that $f_n^\omega(d_i) = f_n^\omega(e_j)$ for infinitely many n; thus $d_i = e_j$. This proves that $S \subseteq T$ and the the converse follows by symmetry. □

Proceeding now to the general theory we recall a standard embedding of locally presentable categories in categories of *relational structures*. Let S be a set of sorts, and Σ a finitary S-sorted relational signature, that is, an S^*-sorted set. The category $\mathsf{Rel}(S,\Sigma)$ has objects: relational Σ-structures A, consisting of an underlying S-sorted set $(A_s)_{s\in S}$ and, for each $\sigma \in \Sigma_{s_1,...,s_n}$, a relation $\sigma_A \subseteq A_{s_1} \times ... \times A_{s_n}$. A morphism $f : A \to B$ is an S-sorted function from (A_s) to (B_s) preserving the relations, i.e. such that $(x_1, ..., x_n) \in \sigma_A$ implies $(f_{s_1}(x_1), ..., f_{s_n}(x_n)) \in \sigma_B$ for each σ of arity $[s_1, ..., s_n]$.

$\mathsf{Rel}(S,\Sigma)$ is locally finitely presentable; a relational structure A is κ-presentable iff it has less than κ vertices (i.e., $\Sigma_{s\in S}\, |A_s| < \kappa$) and less than κ edges (i.e., $\Sigma_{\sigma\in\Sigma}|\sigma_A| < \kappa$). It is a standard result that for any κ-presentable category $\mathcal{K}$ there is a signature (S,Σ) and a κ-filtered colimit preserving embedding of $\mathcal{K}$ as a full reflective subcategory of $\mathsf{Rel}(S,\Sigma)$, cf. [8].

Theorem 3.3.2 Let $T : \mathcal{K} \to \mathcal{K}$ be an accessible endofunctor on a locally presentable category. If T preserves monos then the final T-sequence converges, necessarily to a final T-coalgebra.

Proof. We can choose a regular cardinal κ such that $\mathcal{K}$ is locally κ-presentable and T is κ-accessible. Let the functor I embed $\mathcal{K}$ as a reflective subcategory of the locally finitely presentable category $\mathsf{Rel}(S,\Sigma)$ for some signature (S,Σ). Since every locally presentable category is co-well-powered[2], we may assume that there is a regular cardinal λ such that the image under I of the strong quotient of a κ-presentable object in $\mathcal{K}$ is λ-presentable in

[2] A category $\mathcal{C}$ is said to be well-powered if the collection of subobjects of any given object of $\mathcal{C}$ forms a set; $\mathcal{C}$ is said to be co-well-powered if $\mathcal{C}^{\mathrm{op}}$ is well-powered.

$\mathsf{Rel}(S,\Sigma)$. We will show that in the final T-sequence $\{A_\alpha, f^\alpha_\beta\}$ the projection $f^{\lambda+1}_\lambda$ is monic. It then follows that, for all $\alpha \geq \lambda$, f^α_λ is monic; the statement of the theorem follows by proposition 3.2.2 and the fact that locally presentable categories are well-powered.

A_λ is a colimit $(in_j : D_j \to A_\lambda)_{j\in J}$ of a κ-filtered diagram $D : J \to \mathcal{K}$ of κ-presentable objects. We can factor each arrow in_j as a strong epimorphism $e_j : D_j \to D'_j$ followed by a monomorphism $m_j : D'_j \to A_\lambda$. From the fact that e_j is strong, for each arrow $D_{j_1} \to D_{j_2}$ we get a corresponding arrow $D'_{j_1} \to D'_{j_2}$ (shown below as the dotted line), and A_λ is the colimit of the resulting diagram $D'\colon J \to \mathcal{K}$, with cocone $(m_j : D'_j \to A_\lambda)_{j\in J}$.

$$\begin{tikzcd}
D_{j_1} \arrow[r, "e_{j_1}"] \arrow[d] & D'_{j_1} \arrow[dd, "m_{j_1}"] \arrow[ddl, dashed] \\
D_{j_2} \arrow[d, "e_{j_2}"'] & \\
D'_{j_2} \arrow[r, "m_{j_2}"'] & A_\lambda
\end{tikzcd} \tag{3.2}$$

The functor I preserves limits, being a right adjoint, so we have that $IA_\lambda = \mathrm{Lim}_{\alpha<\lambda} IA_\alpha$. Now, fix $j \in J$. I preserves monos since it preserves limits (cf. section 5.1), so ID'_j is a subobject of IA_λ. Since λ is regular and ID'_j is λ-presentable—so has less than λ vertices—and since limits in $\mathsf{Rel}(S,\Sigma)$ are created by the forgetful functor to Set^S, there exists $\alpha < \lambda$ such that $If^\lambda_\alpha \cdot Im_j$ is monic in $\mathsf{Rel}(S,\Sigma)$. It follows that $f^\lambda_\alpha \cdot m_j$ is monic in $\mathcal{K}$, whence $Tf^\lambda_\alpha \cdot Tm_j = f^{\lambda+1}_{\alpha+1} \cdot Tm_j$ is also monic. Thus, for all $j \in J$, $f^{\lambda+1}_\lambda \cdot Tm_j$ is monic.

Now, $(Tm_j : TD'_j \to TA_\lambda)$ is a colimit of $T \cdot D'$, thus the colimit of the diagram $j \mapsto f^{\lambda+1}_\lambda \cdot Tm_j$ in the slice category $\mathcal{K}/A_\lambda$ is $f^{\lambda+1}_\lambda$ (since colimits in $\mathcal{K}/A$ are created by the forgetful functor to $\mathcal{K}$). But this map must be monic by the lemma below. $\square$

Lemma 3.3.3 Let $\mathcal{K}$ be locally κ-presentable and $B \in \mathcal{K}$. Then the colimit of a κ-filtered diagram in $\mathcal{K}/B$, each of whose vertices is monic, is itself monic.

Proof. Let us suppose a κ-filtered diagram of monos in $\mathcal{K}/B$ with colimit $(f_j \xrightarrow{a_j} f)_{j\in J}$. In the diagram below suppose that $f \cdot g = f \cdot h$, then we want to prove $g = h$. Without loss of generality we may assume that K is κ-presentable since $\mathrm{Pres}_\kappa\mathcal{K}$ is a generator. But now there exists j such that g, h both factor through a_j via g', h'; then $f_j \cdot g' = f_j \cdot h'$, so $g' = h'$ and $g = h$. $\square$

$$\begin{tikzcd}
K \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & A \arrow[r, "h"] & B \\
& A_j \arrow[u, "a_j"] \arrow[ur, "f_j"'] &
\end{tikzcd} \tag{3.3}$$

For a given locally presentable category $\mathcal{K}$ we may be able to prove the appropriate instance of proposition 3.3.2 directly, without first passing to $\mathsf{Rel}(S, \Sigma)$. The advantage of this is that we may be able to get a bound on the cardinal λ from the proof of that result. For instance, if κ is a regular cardinal, whenever we have a κ^{op}-limit in Set

$$A_0 \leftarrow A_1 \leftarrow A_2 \leftarrow \ldots \leftarrow \mathrm{Lim}_{\alpha<\kappa} A_\alpha$$

with projections π_α, then for any subset S of the limit, with $|S| < \kappa$, there exists $\alpha < \kappa$ such that $\pi_\alpha \cdot (S \hookrightarrow \mathrm{Lim}_\alpha A_\alpha)$ is injective. By a similar proof to the theorem above it follows that for a κ-accessible set functor T, with final sequence $\{A_\alpha, f^\alpha_\beta\}$, the projection $f^{\kappa+1}_\kappa$ is injective. For example, from the criterion for a set functor to be finitary, given at the end of section 3.2.2, we easily see that the finite-powerset functor $\wp$ is finitary, and we recover the observation of example 3.3.1. Here is another example, arising from the coalgebraic semantics of probabilistic non-determinism, cf. [84].

Example 3.3.4 Recall the functor $D_f \colon \mathsf{Set} \to \mathsf{Set}$ of finitely supported discrete probability distributions from section 2.1.3. From the criterion in section 3.2.2 one easily sees that the functor $D_f(1 + (-))$ is finitary, and hence that the projection $f^{\omega+1}_\omega$ in the final sequence of $D_f(1 + (-))$ is injective. □

3.4 Final sequences in Set

From now on we concentrate on the final sequence $\{A_\alpha, f^\alpha_\beta\}$ of a κ-accessible endofunctor T on Set. Without loss of generality we assume that T is not the constant $\emptyset$ functor. With this assumption we know that there is at least one T-coalgebra with a non-empty carrier, and since this coalgebra can be extended to a cone over the final T-sequence (see below) it follows that $A_\alpha \neq \emptyset$ for all α. We have already seen that $f^{\kappa+1}_\kappa$ is injective; we can now say that it is split monic, since injectives with non-empty domain are split monic in Set. Moreover, since T preserves split monos, and the projections from the limit of a co-chain of injectives are themselves injective, a simple induction establishes that f^α_β is injective for all $\alpha \geq \beta \geq \kappa$.

Many set functors preserve arbitrary intersections, that is, they preserve wide pullbacks of monos, cf. section 5.1. For such functors (an example is $\wp$), once we know that $f^{\kappa+1}_\kappa$ is injective it follows that the final sequence stabilizes in $\kappa + \omega$ steps since $A_{\kappa+\omega}$ is just the intersection of the 'decreasing' co-chain $(A_{\kappa+n})_{n<\omega}$. It is known that all set functors preserve finite non-empty intersections, cf. [82], however there are set functors which do not preserve infinite non-empty intersections, cf. [38]. Nevertheless, we are able to prove that the final sequence of any ω-accessible set functor converges in $\omega + \omega$ steps. The proof uses the fact that $f_\omega^{\omega+1}$ is split monic, and thus yields a coalgebra structure on A_ω. This coalgebra turns out to be weakly final , and we obtain a final coalgebra as a retract. More generally, the final sequence of any κ-accessible set functor converges in $\kappa + \kappa$ steps.

Proposition 3.4.1 Let (E, e) be a T-coalgebra; then we can extend (E, e) to a cone $(v_\alpha : E \to A_\alpha)_\alpha$ over the final T-sequence such that $v_{\alpha+1} = T(v_\alpha) \cdot e$.

Proof. We define the v_α by transfinite induction, verifying at each stage that we have $(\forall \beta \leq \alpha) f_\beta^\alpha \cdot v_\alpha = v_\beta$. The successor clause is $v_{\alpha+1} \stackrel{\text{def}}{=} T(v_\alpha) \cdot e$. Then

$$\begin{aligned} f_\beta^{\alpha+1} \cdot v_{\alpha+1} &= f_\beta^{\beta+1} \cdot f_{\beta+1}^{\alpha+1} \cdot T(v_\alpha) \cdot e \\ &= f_\beta^{\beta+1} \cdot T(f_\beta^\alpha \cdot v_\alpha) \cdot e \\ &= f_\beta^{\beta+1} \cdot T(v_\beta) \cdot e \\ &= f_\beta^{\beta+1} \cdot v_{\beta+1} \\ &= v_\beta. \end{aligned}$$

For a limit ordinal λ, we define v_λ by $f_\beta^\lambda \cdot v_\lambda = v_\beta$ for $\beta < \lambda$. □

Assume T is κ-accessible. Recall that, by our assumption $T \neq \emptyset$, A_α is non-empty for all α. Thus, we may choose a left inverse $l : A_\kappa \to A_{\kappa+1}$ to the injection $f_\kappa^{\kappa+1}$.

Proposition 3.4.2 The T-coalgebra (A_κ, l) is weakly final.

Proof. Let (E, e) be a T-coalgebra. (E, e) extends to a cone $(v_\alpha : E \to A_\alpha)_\alpha$ over the final sequence of T such that $v_\kappa = f_\kappa^{\kappa+1} \cdot T(v_\kappa) \cdot e$. It follows that $l \cdot v_\kappa = T(v_\kappa) \cdot e$, that is v_κ is a T-coalgebra morphism from (E, e) to (A_κ, l). □

Proposition 3.4.3 Extend the coalgebra (A_κ, l) to a cone $(u_\alpha \colon A_\kappa \to A_\alpha)_\alpha$ over the final sequence of T. Then u_κ is an idempotent map of coalgebras $(A_\kappa, l) \to (A_\kappa, l)$.

Proof. We prove by induction on $\alpha \leq \kappa$ that $u_\alpha \cdot u_\kappa = u_\alpha$. The successor case is:

$$\begin{aligned} u_{\alpha+1} \cdot u_\kappa &= T(u_\alpha) \cdot l \cdot u_\kappa \\ &= T(u_\alpha) \cdot T(u_\kappa) \cdot l \\ &= T(u_\alpha \cdot u_\kappa) \cdot l \\ &= T(u_\alpha) \cdot l \\ &= u_{\alpha+1}. \end{aligned}$$

For α a limit ordinal:

$$f^{\alpha}_{\beta} \cdot u_{\alpha} \cdot u_{\kappa} = u_{\beta} \cdot u_{\kappa} = u_{\beta} = f^{\alpha}_{\beta} \cdot u_{\alpha}$$

for all $\beta < \alpha$. Thus $u_{\alpha} \cdot u_{\kappa} = u_{\alpha}$, and the proposition holds. □

All idempotents in Set split, so we can write $u_{\kappa} = i \cdot q$, where $q : A_{\kappa} \to G$ and $i : G \to A_{\kappa}$ satisfy $q \cdot i = 1_G$. Notice that the pair $T(q) \cdot l$ and $f^{\kappa+1}_{\kappa} \cdot T(i)$ is also a splitting of u_{κ}; thus, by the uniqueness of splittings of idempotents, we have an isomorphism g making the diagram below commute.

$$\begin{tikzcd} A_{\kappa} \arrow[r, "l"] \arrow[d, "q"'] & A_{\kappa+1} \arrow[d, "T(q)"] \\ G \arrow[r, "g"] \arrow[d, "i"'] & TG \arrow[d, "T(i)"] \\ A_{\kappa} & A_{\kappa+1} \arrow[l, "f^{\kappa+1}_{\kappa}"] \end{tikzcd} \tag{3.4}$$

Proposition 3.4.4 The T-coalgebra (G, g) is final.

Proof. Clearly (G, g) is weakly final, since (A_{κ}, l) is weakly final. Suppose h and k are two coalgebra maps $(E, e) \to (G, g)$. (E, e) extends to a cone $(v_{\alpha} : E \to A_{\alpha})_{\alpha}$ over the final T-sequence; we show by induction that $f^{\kappa}_{\alpha} \cdot i \cdot h = v_{\alpha} = f^{\kappa}_{\alpha} \cdot i \cdot k$ for all $\alpha < \kappa$. The case for α a limit ordinal is trivial, the successor case is:

$$\begin{aligned} f^{\kappa}_{\alpha+1} \cdot i \cdot h &= f^{\kappa}_{\alpha+1} \cdot i \cdot g^{-1} \cdot T(h) \cdot e \\ &= f^{\kappa}_{\alpha+1} \cdot f^{\kappa+1}_{\kappa} \cdot T(i) \cdot T(h) \cdot e \\ &= f^{\kappa+1}_{\alpha+1} \cdot T(i \cdot h) \cdot e \\ &= T(f^{\kappa}_{\alpha} \cdot i \cdot h) \cdot e \\ &= T(v_{\alpha}) \cdot e \\ &= v_{\alpha+1}. \end{aligned}$$

Similarly one proves that $f^{\kappa}_{\alpha+1} \cdot i \cdot k = v_{\alpha+1}$. It follows that $i \cdot h = i \cdot k$, and thence that $h = k$. □

Remark 3.4.5 We have obtained a final T-coalgebra as a subalgebra of the T-algebra $(A_{\kappa}, f^{\kappa+1}_{\kappa})$. Taking G above to be a subset of A_{κ}, it is not difficult to prove that $a \in G$ iff there is a T-coalgebra (E, e) such that a lies in the image of v_{κ}, where $(v_{\alpha} : E \to A_{\alpha})_{\alpha}$ is the cone over the final sequence generated by (E, e). We consider a T-coalgebra to be a *guard*, and $G \subseteq A_{\kappa}$ to be the subset of *guarded elements*. □

Proposition 3.4.6 Let u_κ be as in proposition 3.4.3; then $u_\kappa \cdot f_\kappa^{\kappa+\kappa} = f_\kappa^{\kappa+\kappa}$.

Proof. We show by induction on $\alpha \leq \kappa$ that

$$(\forall \alpha \leq \kappa)\, u_\alpha \cdot f_\kappa^{\kappa+\alpha} = f_\alpha^{\kappa+\alpha};$$

this implies the proposition to be proven.
Case: α a limit ordinal. The desired conclusion follows from

$$\begin{aligned} f_\beta^\alpha \cdot u_\alpha \cdot f_\kappa^{\kappa+\alpha} &= u_\beta \cdot f_\kappa^{\kappa+\alpha} \quad \text{for } \beta < \alpha \\ &= u_\beta \cdot f_\kappa^{\kappa+\beta} \cdot f_{\kappa+\beta}^{\kappa+\alpha} \\ &= f_\beta^{\kappa+\beta} \cdot f_{\kappa+\beta}^{\kappa+\alpha} \\ &= f_\beta^{\kappa+\alpha} \\ &= f_\beta^\alpha \cdot f_\alpha^{\kappa+\alpha}. \end{aligned}$$

Case: $\alpha = \beta + 1$.

$$\begin{aligned} u_\alpha \cdot f_\kappa^{\kappa+\alpha} &= T(u_\beta) \cdot l \cdot f_\kappa^{\kappa+\alpha} \\ &= T(u_\beta) \cdot l \cdot f_\kappa^{\kappa+1} \cdot f_{\kappa+1}^{\kappa+\alpha} \\ &= T(u_\beta) \cdot f_{\kappa+1}^{\kappa+\alpha} \\ &= T(u_\beta \cdot f_\kappa^{\kappa+\beta}) \\ &= T(f_\beta^{\kappa+\beta}) \\ &= f_\alpha^{\kappa+\alpha}. \end{aligned}$$

□

Theorem 3.4.7 Let T be a κ-accessible endofunctor on Set with final sequence $\{A_\alpha, f_\beta^\alpha\}$; then $f_{\kappa+\kappa}^{\kappa+\kappa+1}$ is an isomorphism.

Proof. The final T-coalgebra (G, g), as constructed in proposition 3.4.4, extends to a cone $(v_\alpha : G \to A_\alpha)_\alpha$ over the final T-sequence. By definition of this cone, the top square in diagram (3.5) commutes. The bottom two squares of this diagram are just the bottom two squares of (3.4)—where q, i split the idempotent u_κ—thus they also commute. The second top square commutes by definition of the final T-sequence, thus the whole diagram commutes.

Consider the rectangle consisting of the middle two squares. The map $q \cdot f_\kappa^{\kappa+\kappa}$ is injective as its composition with i gives $u_\kappa \cdot f_\kappa^{\kappa+\kappa}$ which, by proposition 3.4.6, equals (the injective map) $f_\kappa^{\kappa+\kappa}$. By finality of (G, g) we have that $q \cdot f_\kappa^{\kappa+\kappa} \cdot v_{\kappa+\kappa}$ is the identity, so $q \cdot f_\kappa^{\kappa+\kappa}$ is also surjective and hence an isomorphism of sets. It follows that three sides of the middle rectangle are isomorphisms, thus the fourth side—$f^{\kappa+\kappa+1}_{\kappa+\kappa}$—is also an isomorphism. □

$$
\begin{tikzcd}
G \arrow[r, "g"] \arrow[d, "v_{\kappa+\kappa}"'] & TG \arrow[d, "Tv_{\kappa+\kappa}"] \\
A_{\kappa+\kappa} \arrow[d, "f^{\kappa+\kappa}_{\kappa}"'] & TA_{\kappa+\kappa} \arrow[l] \arrow[d, "Tf^{\kappa+\kappa}_{\kappa}"] \\
A_{\kappa} \arrow[d, "q"'] & TA_{\kappa} \arrow[l] \arrow[d, "Tq"] \\
G \arrow[r, "g"] \arrow[d, "i"'] & TG \arrow[d, "Ti"] \\
A_{\kappa} & TA_{\kappa} \arrow[l, "f^{\kappa+1}_{\kappa}"]
\end{tikzcd}
\tag{3.5}
$$

Without too much effort we can extend this result to functors which may be neither κ-accessible nor κ^{op}-continuous. An example, in the case that $\kappa = \omega$, is the set functor $\wp(-)^A$ where A is infinite. The coalgebras of this functor are the so-called *image-finite* transition systems (those transition systems such that for each state s, and label $a \in A$, the set of states reachable from s by a one-step a-labelled transition is finite). First we recall the following simple proposition which says that coproducts commute with κ^{op}-limits in Set.

Proposition 3.4.8 Suppose we have a family of κ^{op}-limit co-chains in Set

$$X_{i0} \leftarrow X_{i1} \leftarrow \ldots \leftarrow X_{i\alpha} \leftarrow X_{i(\alpha+1)} \leftarrow \ldots \leftarrow X_{i\kappa}$$

indexed over $i \in I$; then

$$\coprod X_{i0} \leftarrow \coprod X_{i1} \leftarrow \ldots \leftarrow \coprod X_{i\alpha} \leftarrow \coprod X_{i(\alpha+1)} \leftarrow \ldots \leftarrow \coprod X_{i\kappa}$$

is also a limit co-chain for each i. □

Theorem 3.4.9 The class of endofunctors on Set such that the arrow $f^{\kappa+1}_{\kappa}$ in their final sequence is injective is closed under:

1 κ-accessible functors;

2 κ^{op}-continuous functors;

3 composition of functors;

4 arbitrary coproducts of functors;

5 $\mathcal{I}$-indexed-limits of functors, for any small category $\mathcal{I}$.

Proof. Consider the following property of an endofunctor T: for all κ^{op}-limits

$$B_0 \leftarrow B_1 \leftarrow \ldots \leftarrow B_\alpha \leftarrow \ldots \leftarrow B_\kappa,$$

with B_κ non-empty, the comparison map $\mathrm{Lim}_{\alpha<\kappa}T(B_\alpha) \leftarrow T(B_\kappa)$ is injective. We prove that this property is closed under [1-5] above.

Closure under 2 is trivial, whilst closure under 1 was shown in the previous section. Closure under 4 holds by proposition 3.4.8, and the fact that in Set, as in any topos, coproducts preserve monos. For 3 suppose $S, T : \mathsf{Set} \to \mathsf{Set}$, and assume the property in question holds of S and T. Either T is the constant $\emptyset$ functor, in which case $S \cdot T$ is constant, or $T(X) \neq \emptyset$ for $X \neq \emptyset$, in which case the composite

$$\mathrm{Lim}_{\alpha<\kappa}ST(B_\alpha) \leftarrow S(\mathrm{Lim}_{\alpha<\kappa}T(B_\alpha)) \leftarrow ST(B_\kappa)$$

is injective as S preserves injectives with non-empty domain. Finally, for closure under 5, suppose $T = \mathrm{Lim}_{I\in\mathcal{I}}T_I$; then $T(B_\kappa) \cong \mathrm{Lim}_{I\in\mathcal{I}}T_I(B_\kappa)$, and

$$\mathrm{Lim}_{\alpha<\kappa}T(B_\alpha) \cong \mathrm{Lim}_{\alpha<\kappa}\mathrm{Lim}_{I\in\mathcal{I}}T_I(B_\alpha) \cong \mathrm{Lim}_{I\in\mathcal{I}}\mathrm{Lim}_{\alpha<\kappa}T_I(B_\alpha)$$

since limits commute with each other in a complete category. By assumption, for each $I \in \mathcal{I}$, the comparison map $\mathrm{Lim}_{\alpha<\kappa}T_I(B_\alpha) \leftarrow T_I(B_\kappa)$ is injective; moreover the functor $\mathrm{Lim}_{I\in\mathcal{I}}(-) : [\mathcal{I}, \mathsf{Set}] \to \mathsf{Set}$ preserves monos as it is a right adjoint, so the comparison map $\mathrm{Lim}_{\alpha<\kappa}T(B_\alpha) \leftarrow T(B_\kappa)$ is injective.

The theorem now holds since an endofunctor on Set is either the constant $\emptyset$, or each set in its final sequence is non-empty. $\square$

Corollary 3.4.10 The class of endofunctors on Set whose final sequences converge in at most $\kappa + \kappa$ steps is closed under [1-5] in theorem 3.4.9.

Proof. This follows from theorem 3.4.9 and the proof of theorem 3.4.7. $\square$

Example 3.4.11 Consider the functor $T = D(1+(-))$, where D is the functor of probability distributions (of potentially infinite support). T is not finitary, and doesn't immediately appear to be covered by corollary 3.4.10 above. However it is still the case that the final sequence $\{A_\alpha, f^\alpha_\beta\}$ of T converges in $\omega + \omega$ steps. For suppose $\mu, \rho \in D(1 + A_\omega)$ and

$f_\omega^{\omega+1}(\mu) = f_\omega^{\omega+1}(\rho)$. Then for each $\vec{x} \in A_\omega$, since $\mu[-]$ is a measure on A_ω, using (3.8) we have

$$\mu(\vec{x}) = \mu[\bigcap_n (f_n^\omega)^{-1}(x_n)] = \lim_n \mu[(f_n^\omega)^{-1}(x_n)] = \lim_n T(f_n^\omega)(\mu)(x_n),$$

and similarly for ρ. But $T(f_n^\omega)(\mu) = (f_{n+1}^\omega \cdot f_\omega^{\omega+1})(\mu) = (f_{n+1}^\omega \cdot f_\omega^{\omega+1})(\rho) = T(f_n^\omega)(\rho)$. Thus $\mu = \rho$. $\square$

3.5 Set-theoretic and metric final semantics

In chapter 2 we referred to the fact that the (final) solution of the type equation $X \cong \wp(\mathrm{Act} \times X)$ in the category Set serves as a domain in which one can interpret finitary CCS terms (over a fixed set of actions Act), which is fully abstract with respect to strong bisimulation. Another denotational semantics for the finite part of CCS, again fully abstract with respect to strong bisimulation, is obtained by Abramsky [2, Thms 5.20,5.21] by solving the 'corresponding' type equation in the category SFP of SFP domains and continuous functions, viz.,

$$X \cong (\mathbf{1})_\perp \oplus \mathcal{P}[\Sigma_{a\in\mathrm{Act}} X].$$

Here $\mathbf{1}$ is the one element domain, $(.)_\perp$ is lifting, $\oplus$ coalesced sum, Σ disjoint sum, and $\mathcal{P}$ the Plotkin power-domain; cf. [2].

Abramsky notes that the operational models appropriate to CCS, labelled transition systems, are *coalgebras* of the functor $\wp(\mathrm{Act} \times -)$ on Set, and that the solution of the above domain equation, obtained by the standard methods, is both an initial algebra *and a final coalgebra* of the corresponding functor $X \mapsto (1)_\perp \oplus \mathcal{P}[\Sigma_{a\in\mathrm{Act}} X]$ on SFP. His comment:

> Thus our results can in some sense be seen as concerning the interpretation and 'best approximation' of Set-based structures in topological ones. Clearly some general theory is called for here.

Yet another approach to the problem of finding a domain for strong bisimulation is to work in the category CUMet of complete ultrametric spaces and non-expansive maps, cf. [12]. In this case one has the compact power-domain functor $\mathcal{P}_{\mathrm{co}} : \mathsf{CUMet} \to \mathsf{CUMet}$, and one solves the equation $X \cong \mathcal{P}_{\mathrm{co}}(\mathrm{Act} \times \frac{1}{2} \cdot X)$. Here $\mathcal{P}_{\mathrm{co}}(X)$ is the space of compact subsets of a space X with the so-called Hausdorff metric d_H, where

$$d_H(V, W) = \max\{\sup_{v\in V} \inf_{w\in W} d_X(v, w), \sup_{w\in W} \inf_{v\in V} d_X(v, w)\}$$

for compact $V, W \subseteq X$, and where all sups and infs are taken over the interval $[0,1]$. $\frac{1}{2} \cdot -$ is a scaling functor, mapping a space X to a space with the same set of points but where all distances are halved, and the set of actions Act is given the discrete metric.

Abramsky [2] mentions a general result concerning the solutions of 'corresponding' domain equations in SFP and CUMet. It turns out that for domain equations involving the sum, product, and powerdomain constructs, the space of maximal elements of the (canonical) solution in SFP is homeomorphic to the solution in CUMet. In this section we conduct a similar comparison between final semantics in Set and CUMet. For each set functor T we define a corresponding locally contractive endofunctor $\overline{T}$ on CUMet, for example, the finite-powerset functor $\wp$ gives rise to the compact powerdomain functor above. Theorem 2.3.10 implies that $\overline{T}$ has a unique fixed point which is both an initial algebra and a final coalgebra. It turns out that, for a finitary set functor T, the ω-th iterate in the final sequence of T, with a natural ultrametric, is the unique fixed point of $\overline{T}$. Thus, the final T-coalgebra can be seen as a subspace (indeed, a sub-T-algebra) of the final $\overline{T}$-coalgebra. For instance, in the next section the final coalgebra of $\mathcal{P}_{\mathrm{co}}(\frac{1}{2} \cdot -)$ is described as the coalgebra of strongly extensional, *compactly branching trees*, and the final coalgebra of $\wp$ is described as the coalgebra of strongly extensional, *finitely branching* trees. In another example we compare two domains for probabilistic bisimulation from [84].

3.5.1 Metrics on final coalgebras

Let $T \colon \mathsf{Set} \to \mathsf{Set}$ have final sequence $\{A_\alpha, f^\alpha_\beta\}$. Then A_ω has a natural topology as an inverse limit of discrete spaces; in fact it is an ultrametric space, with distance

$$d_\omega(\vec{x}, \vec{y}) = \inf\{2^{-n} \mid x_n = y_n\}. \tag{3.6}$$

Barr [13] observed that, if T is both ω-accessible and ω^{op}-continuous, and $T(\emptyset) \neq \emptyset$, then the final T-coalgebra is the Cauchy completion of the initial T-algebra. From theorem 3.4.9 we observe that, assuming only ω-accessibility of T, we may still define a natural topology on the final T-coalgebra, arising from the restriction of the ultrametric (3.6) to $A_{\omega+\omega}$. The following lemma, which is the key to the subsequent development, shows that this subspace is dense in A_ω.

Lemma 3.5.1 Suppose T is a set functor with final sequence $\{A_\alpha, f^\alpha_\beta\}$; then, for $\lambda \geq \omega$, the image of f^λ_ω is a dense subset of A_ω.

Proof. Without loss of generality assume T is not the constant functor $\emptyset$. We define a family of maps

$$g^n_\alpha : A_n \to A_\alpha$$

for $n < \min(\omega, \alpha + 1)$ such that

- G1 $g^m_\alpha \cdot g^n_m = g^n_\alpha \quad n \leq m \leq \alpha$;
- G2 $g^{n+1}_{\alpha+1} = T(g^n_\alpha)$;
- G3 $f^\alpha_\beta \cdot g^n_\alpha = g^n_\beta \quad n \leq \beta \leq \alpha$;
- G4 $g^n_n = 1_{A_n}$.

The definition of g^n_α proceeds by induction on α. First we define g^0_1 to be a right inverse to $f^1_0 : T1 \to 1$. Suppose α is a successor ordinal greater than 1. In the subcase $n > 0$ we use clause G2 to define g^n_α, otherwise we define $g^0_\alpha = g^1_\alpha \cdot g^0_1$, where g^1_α is defined according to the first subcase. For α a limit ordinal g^n_α is defined to be the unique map making clause G3 hold.

Given $\vec{x} \in A_\omega$ and $n < \omega$ we have that

$$f^\omega_n f^\lambda_\omega g^n_\lambda(x_n) = f^\lambda_n g^n_\lambda(x_n) = x_n$$

hence $d_\omega(f^\lambda_\omega g^n_\lambda(x_n), \vec{x}) \leq 2^{-n}$. Since n can be chosen to be arbitrarily large the lemma is proved. □

Next, we show how to derive an endofunctor on CUMet from a set functor T via a lifting of T to equivalence relations.

3.5.2 Lifting set functors to CUMet

Associated to each equivalence relation R on a set X we have the cospan $X \xrightarrow{q} X/R \xleftarrow{q} X$, with X/R the set of equivalence classes and q the quotient map. Observe that $R = q^\circ q$, where q° is the inverse relation to q. Applying T we get a cospan $TX \xrightarrow{Tq} T(X/R) \xleftarrow{Tq} TX$, and we define the equivalence relation $(\mathrm{Eq}\ T)R$ on TX to be $(Tq)^\circ Tq$. Ordering equivalence relations on X by inclusion, Eq T acts as a monotone operator since $R \subseteq S$ iff the cospan associated to S factors through the cospan associated to R. Also, taking R and S to be an equivalence relations on sets X and Y respectively, and given a function $f : X \to Y$, we have that $R \subseteq f^\circ \cdot S \cdot f$ iff the cospan

$$X \xrightarrow{f} Y \to Y/S \leftarrow Y \xleftarrow{f} X$$

factors through the cospan associated to R. Thus,

$$R \subseteq f^\circ \cdot S \cdot f \text{ implies } (\mathrm{Eq}\ T)R \subseteq (Tf)^\circ \cdot (\mathrm{Eq}\ T)S \cdot Tf. \tag{3.7}$$

Example 3.5.2 Assume an equivalence relation R on a set X. For a fixed set A, let T be the exponentiation functor $(-)^A : \mathsf{Set} \to \mathsf{Set}$; then

$$(\mathrm{Eq}\ T)R \subseteq X^A \times X^A = \{\langle f, g\rangle \mid \forall a \in A\ \langle fa, ga\rangle \in R\}.$$

If T is the finite-powerset functor, then $(\mathrm{Eq}\ T)R$ is the so-called Egli-Milner extension of R, viz.

$$\begin{aligned}(\mathrm{Eq}\ T)R \subseteq \wp X \times \wp X = \{\langle U, V\rangle \mid{} & \forall u \in U\ \exists v \in V \langle u, v\rangle \in R \\ & \forall v \in V\ \exists u \in U \langle u, v\rangle \in R\}.\end{aligned}$$

If T is the functor of finitely supported probability distributions from example 3.3.4, then

$$(\mathrm{Eq}\ T)R \subseteq D_f X \times D_f X = \{\langle \mu, \rho\rangle \mid \forall x \in X\ \mu[[x]_R] = \rho[[x]_R]\},$$

where $[x]_R$ is the equivalence class of $x \in X$ under R. □

Remark 3.5.3 In the next chapter, in connection with the notion of T-bisimulation, we consider a lifting of T to act on arbitrary relations. This lifting is based on the span (of projections) associated to each relation. It turns out that the two notions of lifting agree if T preserves weak pullbacks, but not in general. This last is a reflection of the fact that final semantics is not always abstract with respect to bisimulation. □

Recall our assumption that distances in an ultrametric space are bounded by 1. Associated with each ultrametric space (X, d) one has a family of equivalence relations $\{R_\varepsilon\}_{0 \leq \varepsilon \leq 1}$ on X given by

$$R_\varepsilon = \{\langle x, y\rangle \mid d(x, y) \leq \varepsilon\}$$

(this family is typically introduced as a base for a filter of relations making X into a uniform space). Clearly we have that $R_0 = \Delta_X$, the diagonal relation on X, and $R_{\varepsilon_1} \subseteq R_{\varepsilon_2}$ for $\varepsilon_1 \leq \varepsilon_2$. Conversely, given a family of equivalence relations $\{R_\varepsilon\}_{0 \leq \varepsilon \leq 1}$ satisfying these two conditions, one gets an ultrametric d by defining

$$d(x, y) = \inf\{\varepsilon \in [0, 1] \mid \langle x, y\rangle \in R_\varepsilon\}.$$

Let UMet be the category of ultrametric spaces and non-expansive maps and T a set functor. We define a locally contractive endofunctor $\widetilde{T} : \mathsf{UMet} \to \mathsf{UMet}$, corresponding to T, as follows. For an ultrametric space (X, d), determining a family of equivalence relations

$\{R_\varepsilon\}_{0\leq\varepsilon\leq 1}$, $\widetilde{T}(X,d)$ has underlying set of points TX, with ultrametric determined by the family of equivalence relations

$$\{(\mathrm{Eq}\ T)R_{\min(2\varepsilon,1)}\}_{0\leq\varepsilon\leq 1}.$$

If $f\colon (X,d)\to(Y,d')$ is a non-expansive map then it follows from (3.7) that Tf is a non-expansive map from $\widetilde{T}(X,d)$ to $\widetilde{T}(Y,d')$.

The category of complete ultrametric spaces CUMet is a full, reflective subcategory of UMet; the left adjoint to the inclusion is given by Cauchy completion $\mathsf{UMet}\to\mathsf{CUMet}$. We now define $\overline{T}$ to be the composite

$$\mathsf{CUMet}\hookrightarrow\mathsf{UMet}\xrightarrow{\widetilde{T}}\mathsf{UMet}\to\mathsf{CUMet}.$$

Example 3.5.4 In case T is the finite-powerset functor, $\widetilde{T}$ takes a space X to the space of its finite subsets with the Hausdorff metric (and scaling factor 1/2). Since this space is dense in $\mathcal{P}_{\mathrm{co}}(\frac{1}{2}\cdot X)$—see below—it follows that $\overline{T}$ is given by $\mathcal{P}_{\mathrm{co}}(\frac{1}{2}\cdot -)$. In case T is the finite distributions functor D_f, it follows from proposition 3.5.5 that $\overline{T}$ is the functor M of Borel probability distributions with compact support (plus a contraction factor), introduced in [84], and which we now describe. □

Describing the extension of the functor D_f to CUMet requires some measure-theoretic terminlogy, which we briefly review. Recall that a *σ-algebra* on a set X is a family of subsets of X containing $\emptyset$ and closed under complements and countable unions. The pair (X,Σ) is called a *measurable space*. We say that $A\subseteq X$ is *(Σ-)measurable* if $A\in\Sigma$. A function $f\colon(X,\Sigma)\to(X',\Sigma')$ between measurable spaces is said to be a *measurable function* if $f^{-1}(A)\in\Sigma$ for each $A\in\Sigma'$. If X is a topological space then the *Borel σ-algebra* is the smallest σ-algebra on X containing all the open subsets of X. Finally, a function $\mu\colon\Sigma\to[0,1]$ is a *probability measure* on (X,Σ) if $\mu(X)=1$ and $\mu(\bigcup_n A_n)=\sum_n\mu(A_n)$ for any disjoint countable family of measurable sets $\{A_n\}$. In particular, this condition implies that if $\{A_n\}$ is a sequence of measurable sets then

$$A_{n+1}\subseteq A_n \text{ for all } n<\omega \text{ implies } \mu(A_n)\to\mu(\cap_n A_n). \tag{3.8}$$

We say that a Borel measure μ on an ultrametric space X has compact support if there exists a compact set $K\subseteq X$ such that for all Borel sets U, $U\cap K=\emptyset$ implies $\mu(U)=0$. Let $M(X)$ denote the set of Borel probability measures on X with compact support. For $\varepsilon>0$ put $\mathcal{O}_\varepsilon=\{O\subseteq X\mid\forall x\in O: B_\varepsilon(x)\subseteq O\}$. Let the metric d on $M(X)$ be given by

$$d(\mu,\rho)=\inf\{\varepsilon>0\mid\forall O\in\mathcal{O}_\varepsilon:\mu(O)=\rho(O)\}.$$

In [84] it is shown that $(M(X), d)$ is a complete ultrametric space. M is turned into an endofunctor on CUMet by defining $M(f)(\mu)(O) = \mu(f^{-1}(O))$ for a non-expansive map $f : X \to Y$. For each ultrametric space X we have an inclusion of sets $D_f(X) \subseteq M(X)$, where for $\mu \in D_f(X)$ and each Borel set E, $\mu(E) \stackrel{def}{=} \mu[E]$. In this way we may regard $D_f(X)$ as an ultrametric space.

Proposition 3.5.5 Let X be a complete ultrametric space. Then $\wp(X)$ is a dense subspace of $\mathcal{P}_{\mathrm{co}}(X)$, and $D_f(X)$ is a dense subspace of $M(X)$.

Proof. Let $K \subseteq X$ be compact and let $\varepsilon > 0$. There is a finite set $K' = \{x_1, ..., x_m\} \subseteq K$ such that $\{B_\varepsilon(x) \mid x \in K'\}$ covers X. It is clear that the distance between K' and K in the Hausdorff metric is less than ε, proving the first assertion above. For the second, let μ be a Borel measure with compact support K above. Observe that without loss of generality we may assume that the $B_\varepsilon(x_i)$ are pairwise disjoint, since in an ultrametric space two open ε-balls are either disjoint or equal. Define the distribution ρ by support$(\rho) = \{x_1, ..., x_m\}$ and $\rho(x_j) = \mu(B_\varepsilon(x_j))$ for $1 \leq j \leq m$. Then the distance between μ and ρ in $M(X)$ is less that ε since if $O \in \mathcal{O}_\varepsilon$ then each $B_\varepsilon(x_j)$ is either a subset of O or doesn't meet O, so

$$\mu(O) = \sum_{x_j \in O} \mu(B_\varepsilon(x_j)) = \sum_{x_j \in O} \rho[B_\varepsilon(x_j)] = \rho(O).$$

□

Returning to the general development we have:

Lemma 3.5.6 Let $T : \mathsf{Set} \to \mathsf{Set}$ be finitary with final sequence $\{A_\alpha, f^\alpha_\beta\}$, and give A_ω the metric d_ω in (3.6). Then $f^{\omega+1}_\omega$ is an isometric embedding of $\widetilde{T}(A_\omega, d_\omega)$ into (A_ω, d_ω).

Proof.

$$\begin{aligned} d_\omega(f^{\omega+1}_\omega(x), f^{\omega+1}_\omega(y)) &= \inf\{2^{-n} \mid f^{\omega+1}_n(x) = f^{\omega+1}_n(y)\} \\ &= \inf\{2^{-(n+1)} \mid f^{\omega+1}_{n+1}(x) = f^{\omega+1}_{n+1}(y)\} \quad \text{(inf taken in [0,1])} \\ &= \inf\{2^{-(n+1)} \mid T(f^\omega_n)(x) = T(f^\omega_n)(y)\} \\ &= \inf\{2^{-(n+1)} \mid \langle x, y\rangle \in (\mathrm{Eq}\ T)(\mathrm{kernel}(f^\omega_n))\} \\ &= d(x, y) \text{ in } \widetilde{T}A_\omega. \end{aligned}$$

□

Putting lemmas 3.5.1 and 3.5.6 together we get that $f^{\omega+1}_\omega$ is a dense isometric embedding of $\widetilde{T}A_\omega$ into A_ω. Hence, the unique extension to a non-expansive map $\overline{T}A_\omega \to A_\omega$ is an isomorphism, and is thus the final coalgebra/initial algebra of $\overline{T}$. In summary, the ω-th iterate in the final sequence of a finitary set functor T is always a final coalgebra—not necessarily of T, but certainly of the 'completion' of T to an endofunctor on CUMet. In the next section we give an example showing how the next ω steps in the final sequence can be seen as 'pruning' A_ω to obtain a final T-coalgebra.

3.6 An example: non-well-founded sets

In this section we sketch an application of the final coalgebra construction from the proof of theorem 3.4.7 to the case of the finite-powerset functor $\wp$. This example illustrates well why the final sequence of a finitary set functor converges in $\omega + \omega$ steps. It also shows the relation between Aczel's relative consistency proof for his anti-foundation axiom (i.e., the construction of a final coalgebra for a powerset functor) and the construction of Forti and Honsell [31] of a model of a set theory with a universal set.

3.6.1 The final sequence of $\wp$

For our purposes a tree t is a directed graph with a distinguished node, denoted $\mathrm{root}(t)$, such that every node is reachable from $\mathrm{root}(t)$ by a unique finite path; thus, each element of a tree has finite depth. We consider trees which are isomorphic as directed graphs with distinguished nodes to be the same, and denote the class of trees by $\mathcal{T}$. Given a node n of t, by *the maximal subtree rooted at* n we mean the maximal subgraph of t which is a tree with root n. A tree is called *extensional* if, for any two children x and y of any node, whenever the maximal subtrees rooted at x and y are isomorphic, then $x = y$. Given a tree t of finite depth n we can compute an extensional tree, the extensional quotient of t, as follows. First, identify the maximal subtrees rooted at any two sibling nodes x, y at level n in t if they are isomorphic; perform the same operation on the level $n-1$ nodes of the resultant tree, and continue until the top of the tree is reached. We call a tree t *strongly extensional* if, viewing t as a transition system where nodes make transitions to their children, no two distinct children of any node are strongly bisimilar (this is in keeping with the usage of the term strongly extensional for sets, cf. [4]). Clearly, any strongly extensional tree is extensional, although the converse doesn't hold (see below). Writing $\approx_n$ for bisimilarity up to n steps, we define a pseudometric[3] $d_{\mathcal{T}}$ on the set of strongly extensional trees by

$$d_{\mathcal{T}}(t, t') = \inf\{2^{-n} \mid \mathrm{root}(t) \approx_n \mathrm{root}(t')\}.$$

[3]In fact, when we restrict to the compactly branching trees, $d_{\mathcal{T}}$ is actually a metric.

Write $\{*\}$ for the final object in Set, and $\{A_\alpha, f^\alpha_\beta\}$ for the final sequence of $\wp$. For each $n < \omega$ there is an isomorphism between A_n and the set of finitely branching extensional trees of depth $\leq n$ whose depth n nodes are labelled by $*$. This goes by induction: if $x_1, ..., x_m \in A_n$ correspond to trees $t_1, ..., t_m$ then

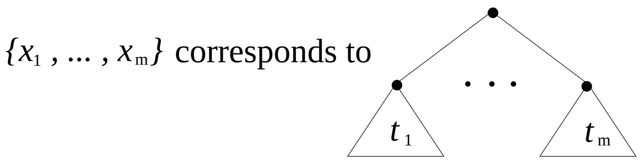

Labelling the leaf nodes of depth n in a tree in A_n by $*$ communicates the intuition that these trees are to be thought of as partial. We can think of the projection map f^{n+1}_n as taking a tree in A_{n+1}, cutting off the depth $n+1$ nodes, taking the extensional quotient, and finally relabelling the depth n nodes by $*$. Below, we draw a row of trees taken from the first four objects of the final sequence, with each tree projecting down to the tree to the left.

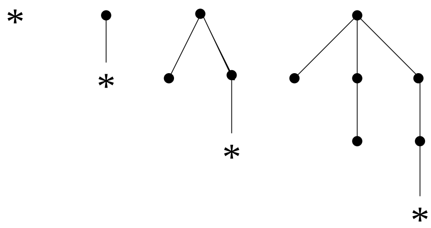

Extending the given sequence in the obvious way we can imagine the trees as projections of the following countably-branching infinite-depth tree, living in the limit A_ω.

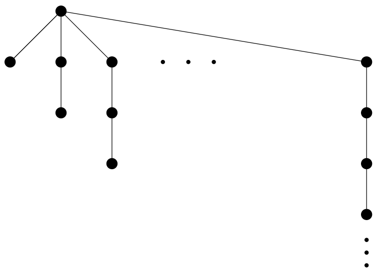

The above picture is due to Turi and Rutten [83], from their explanation of why the final coalgebra of $\wp$ is not given by the ω^{op}-limit $1 \leftarrow \wp 1 \leftarrow ... \leftarrow \wp^\omega 1$. Notice that since the trees in A_ω are potentially infinitely branching, the obvious candidate for a coalgebra structure on A_ω, i.e. the map sending a tree t to the set of its maximal proper subtrees, does not work. Concerning this example Barr [13] also remarks that 'the inverse limit is apparently not a coalgebra'. However, in the last section we showed that A_ω can be given a coalgebra structure, indeed the structure of a final coalgebra, but of the compact powerdomain functor, not $\wp$. In fact, the limit A_ω can be depicted as the set of strongly extensional trees that are compactly branching in the sense that, for each node, the set of maximal subtrees rooted at its children is compact with respect to the metric $d_{\mathcal{T}}$. The projection f^ω_n is given by: cut a tree to depth n, take the extensional quotient and relabel the depth n nodes by $*$. The injection $f^{\omega+1}_\omega$ can be thought of as the inclusion of the subset of those trees which are finitely branching at the root. More generally, the injection $f^{\omega+n}_\omega$ can be thought of as the inclusion of the subset of those trees which are finitely branching up-to depth $n < \omega$. Finally, the carrier of the final coalgebra, $A_{\omega+\omega}$, is the set of finitely branching, strongly extensional trees. The coalgebra structure on this set sends each tree t to the (finite) set of its maximal proper subtrees.

Barr has already pointed out that the set of finitely-branching extensional trees of possibly infinite depth is not a final coalgebra of $\wp$; he constructs a final coalgebra as a set of equivalence classes of such trees. The strongly extensional trees are canonical representatives of these equivalence classes; this slightly more concrete presentation of the final coalgebra pays dividends when we come to characterize products and the subobject classifier in the category of $\wp$-coalgebras, cf. section 5.4. From Barr's example of two extensional trees that must be identified in the final coalgebra we get an example of an extensional tree that is not strongly extensional. To describe this, let t_1 be the tree below.

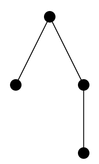

We define a sequence of trees $(t_n)_{n<\omega}$ by defining t_{n+1} to be the tree obtained from t_n by replacing each leaf of t_n with a copy of t_1. The sequence $(t_n)_{n<\omega}$ together with the inclusions $t_n \hookrightarrow t_{n+1}$ forms an ω-chain in Graph. The colimit of this chain is an infinite-depth tree t which is extensional but not strongly extensional (e.g., the relation R on the nodes of t given by $(n, n') \in R$ iff $\mathrm{depth}(n) = \mathrm{depth}(n')$ is easily seen to be a strong bisimulation).

3.6.2 Non-well-founded sets

Via the interpretation of the 'is a child of' relation on trees as set membership, both A_ω and $A_{\omega+\omega}$ are interesting universes of sets. In the latter case it is well-known that the final coalgebra of $\wp$ forms a model for a set theory where the axiom of foundation is replaced by an axiom of *anti-foundation*, guaranteeing the existence (and uniqueness) of sets like $\Omega = \{\Omega\}$. With regard to the former case we observe that A_ω is isomorphic to Forti and Honsell's [31] construction of a so-called *ω-hyperuniverse*. Such hyperuniverses are models of set theories including not only the anti-foundation axiom but also other axioms not satisfied by Aczel's model, like the existence of a universal set. Forti and Honsell see a model of set theory V as a transition system, with each set in V making a transition to its elements, and consider the sequence of equivalences $(\approx_n)_{n<\omega}$ on V given by n-step bisimilarity. They construct an ω-hyperuniverse $\mathcal{N}_\omega$ as the quotient set $V/\approx_\omega$, with the membership relation $\in_\alpha$ on $V/\approx_\alpha$, $\alpha \leq \omega$, defined by

$$[x] \in_\alpha [y] \text{ iff } (\forall n < \alpha, \exists z \in y \;\; x \approx_n z).$$

Now, a straightforward induction shows that $V/\approx_\alpha$ is isomorphic (as a graph) to $\wp^\alpha(1)$ for all $\alpha \leq \omega$.