On Coalgebras and Final Semantics

James Worrell

Keble College

Trinity 2000

Submitted in partial fulfilment of the requirements for the Doctor of Philosophy in Computation

Oxford University Computing Laboratory
Programming Research Group

Abstract

This thesis is concerned with an analysis of the category $\mathsf{Coalg}(T)$ of *coalgebras of an endofunctor* T. The objects of this category can be thought of as transition systems, and the morphisms can be thought of as functional bisimulations. Thus, $\mathsf{Coalg}(T)$ is a suitable setting for a category-theoretic account of bisimulation and coinduction. We are particularly interested in final coalgebras, viewing them as strongly-extensional universal transition systems and as recursively defined data types.

One of our main results involves the construction of the final coalgebra of an accessible set functor T via an ordinal-indexed sequence of iterates – the so-called final sequence of T. We show that if T is κ-accessible, for some regular cardinal κ, then this sequence converges to a final coalgebra in $\kappa + \kappa$ steps. Thus, the construction of the final coalgebra can be seen as a two-stage process. In particular, for a finitary (ω-accessible) set functor T, the ω-th object in the final sequence of T can be seen, under a natural ultrametric, as the Cauchy completion of the initial T-algebra. It is always a final coalgebra, not necessarily of T, but certainly of a 'completion' of T to an endofunctor on the category of complete ultrametric spaces. Beyond ω, in the second stage of the final coalgebra construction, the final sequence becomes a chain of injections until a final coalgebra is eventually obtained after $\omega + \omega$ steps.

We define a notion of metric bisimulation for coalgebras of endofunctors on the category of metric spaces and non-expansive maps, and we prove that the metric on a final coalgebra is maximum amongst all bisimulations. This result justifies a coinduction proof principle for recursive data types. As suggested by Rutten, and motivated by the idea that metric spaces can be seen as enriched categories, the definition of a metric bisimulation is based on the notion of a bimodule, the enriched categorical counterpart of a relation. Also, we give a condition under which the operational notion of bisimulation agrees with the denotational notion of final semantics. As an application of this theory, we obtain a coinductive characterization of a known metric on the class of probabilistic transitions systems.

We use our analysis of final coalgebras to study more general categorical structure. Given a set functor T, we show that if T preserves wide pullbacks then $\mathsf{Coalg}(T)$ is a category of presheaves; if T is accessible and preserves pullbacks then $\mathsf{Coalg}(T)$ is a category of sheaves; if T is accessible and preserves weak pullbacks then $\mathsf{Coalg}(T)$ is not necessarily a category of sheaves, but its effective completion is a sheaf category. We also characterize products and subobject classifiers in categories of coalgebras, illustrating this characterization in the category of labelled transition systems and functional strong bisimulations.

Acknowledgements

Thanks to Dominic Hughes and Jeff Sanders for proof reading, and to Jiri Adámek, Corina Cîrstea and Peter Johnstone for conversations about coalgebras. Franck van Breugel helped greatly with the presentation of section 4.8 in the final version.

I should mention Damian Catani, Justin Cîrstea, Sadie Creese, Thad Kobylarz, Naci Mehmet, Paul Metzgen, Joel Ouaknine, Silvija Seres, Richard Tolcher and Abhi Sarkar. Had it not been for their good company I would have finished this thesis a lot earlier.

Finally, I would like to thank my examiners Peter Aczel and Luke Ong for their suggestions on improvements for the final version.

Contents

Chapter 1

Introduction

> One of the aims of universal algebra is to extract, whenever possible, the common elements of several seemingly different types of algebraic structures. In achieving this one discovers general concepts, constructions, and results which not only generalize and unify the known special situations, thus leading to an economy of presentation, but, being at a higher level of abstraction, can also be applied to entirely new situations, yielding significant information and giving rise to new directions.
>
> S. Burris and H. Sankappanavar,
> *A course in universal coalgebra.*
> Springer-Verlag, 1981.

1.1 Coalgebra

The category-theoretic notion of a *coalgebra of an endofunctor* provides a convenient meeting place for ideas from concurrency theory and data type theory. We think of coalgebras both as transition systems and as data types consisting of possibly circular, non-well-founded elements. Since the notion is category-theoretic, we immediately have the concept of a *morphism of coalgebras*, which we think of as a *functional bisimulation*. Thus, coalgebras and their morphisms form a category. The 'co' prefix signifies that coalgebras are the duals of algebras. One of the motivations for studying coalgebras in the abstract, from the perspective of computer science, is the aim of developing a general theory of transition systems, bisimulations, coinduction, modal logics, etc. which provides the same benefits as universal algebra in terms of clarity, unity, and conceptual economy. That is, as Rutten [73] puts it, the aim is to develop a theory of 'universal coalgebra'.

We should immediately admit that, in our opinion, it is unlikely that the theory of coalgebras will ever have the richness and importance of universal algebra. Nevertheless, in the limited sphere of semantics of programming languages, coalgebras are, we feel, of almost equal importance. This balance is neatly captured in Freyd's notion of an algebraically compact category, viz. one in which the initial algebra of any endofunctor coincides with the final coalgebra of that functor. This axiom isolates those categories suitable for modelling recursive types, cf. [29].

The notion of algebra we refer to here—an algebra of an endofunctor—is considerably more general than the notion usually considered in universal algebra. The duality between algebra and coalgebra is clearest when one takes an abstract category-theoretic view of the former. Thus, universal coalgebra doesn't just seek to unify what we know from different instances; it also seeks, via this duality, to explain the unfamiliar in terms of the familiar. However, a theory of universal coalgebra cannot be obtained in a trivial way by formally dualizing the theory of algebras. This is because many of the interesting results about algebras assume we are working in a particular base category, like Set. But a category and its dual may have very different properties, for instance $\mathsf{Set}^{\mathrm{op}}$ fails to be locally presentable. Nevertheless, we do use the duals of important results from the (categorical) theory of universal algebra, like Beck's theorem (cf. section 5.3) and Linton's theorem (cf. section 5.4).

The table below lists some important notions from universal algebra, and dual coalgebraic notions, cf. [73]. Sometimes this duality can be expressed in a precise category-theoretic sense, e.g. the concept of a final object in a category is formally dual to the concept of an initial object. Sometimes the duality is intended in a less precise sense (at least at this stage). Thus, modal logics seem to play much the same role for coalgebras as equational logics for algebras: initial algebras are constructed from the terms in an equational logic, whereas final coalgebras can be constructed as ultrafilters of modal logic formulae. Moreover, one has a 'covariety theorem' characterizing those classes of coalgebras which are models of modal logic theories in terms of their closure under universal constructions, cf. Goldblatt [34]. This is comparable to the variety theorem in algebra.

Algebra	Coalgebra
Algebra homomorphism	Coalgebra homomorphism
Initial algebra	Final coalgebra
Congruence	Bisimulation
Induction	Coinduction
Equational logic	Modal logic
Variety	Covariety

1.2 Contributions of this thesis

This thesis is a contribution to the emerging field of universal coalgebra. We are concerned with developing the basic theory, rather than focusing on any particular application—though we use labelled transition systems and probabilistic labelled transition systems as running examples. Our focus is on the first five rows of the table above. We consider the construction of final coalgebras, coinduction proof principles for final coalgebras, and the general structure of categories of coalgebras. Next we briefly explain, chapter by chapter, the particular contributions. More detailed descriptions can be found at the beginning of each chapter.

Chapter 3 is concerned with the construction of final coalgebras—primarily of set functors. The existence of final coalgebras for accessible (or bounded) set functors is known from the theorems of Aczel and Mendler [3] and Barr [13]. However, these results do not give much information about the structure of the final coalgebras: the former use a quotient-of-a-sum construction, whilst the latter is based on some 'abstract nonsense'.

For a more informative account of a final coalgebra of an endofunctor $T\colon \mathcal{C} \to \mathcal{C}$ we consider the so-called final sequence of T: start with the final object of $\mathcal{C}$, then proceed by ordinal iteration, applying T at successor steps and taking limits in the category at limit steps in the iteration. The main result of the chapter, theorem 3.4.7, states that the final sequence of a κ-accessible set functor converges in $\kappa + \kappa$ steps. We also show that if T is an accessible endofunctor on a locally presentable category, then the final sequence of T converges if T preserves monomorphisms.

Theorem 3.4.7 is illustrated with the example of the finite-powerset functor $\wp\colon \mathsf{Set} \to \mathsf{Set}$. The ω-th object in the final sequence of $\wp$ can be seen as a set of trees which are strongly extensional, possibly of infinite depth, and *compactly branching* in a natural ultrametric. This is, in fact, a final coalgebra of the compact powerdomain functor on the category of complete ultrametric spaces. In the next ω steps of the final sequence these trees are 'pruned', one level at a time, until we reach the final $\wp$-coalgebra: the set of *finitely branching*, strongly extensional, possibly infinite-depth trees. In general, for any finitary (ω-accessible) set functor T, the ω-th object in the final sequence of T can always be seen as a final coalgebra—not necessarily of T, but certainly of a 'completion' of T to an endofunctor on the category of complete ultrametric spaces. Beyond ω the final sequence becomes a co-chain of injections, until it eventually stabilizes at $\omega + \omega$.

The ultrametrics on the final coalgebras of finitary set functors measure the 'behavioural distance' of processes: the greater the depth of bisimilarity between them, the closer they are. Bisimilar processes have distance zero. In chapter 4 we take this idea seriously and define the notion of a *metric bisimulation* between coalgebras of endofunctors on the category Met of metric spaces and non-expansive maps. Roughly speaking, a metric bisimulation is a $[0,\infty]$-valued relation, assigning a behavioural distance to each pair of processes.

We prove a coinduction theorem, theorem 4.5.8, characterizing the metric on the carrier of a final coalgebra as the maximum metric bisimulation. As an application of this result we give a coinductive characterization of the metric d_{Proc} on the class of probabilistic processes with continuous states spaces, introduced by Desharnais et al. [25]. This is in contrast to the inductive presentation of d_{Proc} via a real-valued logic, in op. cit. Specifically, we identify a functor $F\colon \mathsf{Met} \to \mathsf{Met}$ such that probabilistic processes correspond to F-coalgebras. Then, given any processes p and q, their distance in the metric d_{Proc} is equal to the maximum value assigned to the pair (p,q) by a metric F-bisimulation. This, in turn, is equal to the distance between the images of p and q under the unique map to the final coalgebra, that is, *the metric d_{Proc} is given by a final coalgebra semantics.*

Chapter 4 is heavily influenced by a paper of Lawvere [53] treating metric spaces as categories enriched in the interval $[0,\infty]$, and preorders as categories enriched in the two-element chain $\underline{2}$. In fact, the work in this chapter applies equally well to preorders. We obtain a notion of simulation for coalgebras of endofunctors on the category of preorders, and a coinduction theorem characterizing the order on a final coalgebra as the maximum simulation. Interestingly, when we apply this theory to a discrete version of the probabilistic processes example referred to above, we get a simple but useful extension of Larsen and Skou's notion of probabilistic bisimulation. This allows us to prove a surprising generalization of a result of theirs on the characterization of bisimilarity by a probabilistic modal logic.

The notion of metric bisimulation is based on the enriched categorical notion of bimodule. This is in contrast to the coalgebraic bisimulations of Aczel and Mendler [3] and Fiore [28], which are based on relations in the underlying category. A pleasant consequence of this is that the condition on the functor F so that the notion of metric F-bisimulation is well-behaved (e.g. so that bisimulations compose and the final coalgebra semantics agrees with metric bisimilarity) is very simple: F must preserve isometric embeddings. In the preorder case the requirement is that F must preserve order embeddings. In practice, this is much simpler to verify than the corresponding condition (F preserves weak pullbacks) for the notions of coalgebraic bisimulations based on relations.

Chapter 5 builds on the analysis of final coalgebras from chapter 3 to investigate the structure of categories of coalgebras. The underlying theme is that there is a close connection between categories of coalgebras and categories of sheaves and presheaves. The principal result in this chapter, corollary 5.2.3, states that the category of coalgebras of a wide-pullback preserving set functor T is a presheaf category. We also show that if T is accessible and preserves pullbacks, then $\mathsf{Coalg}(T)$ is a category of sheaves. If T only preserves weak pullbacks then $\mathsf{Coalg}(T)$ may fail to be a sheaf category, but we show that its effective completion is a category of sheaves.

Categories of sheaves have been described as generalized universes of sets; they support many of the same constructions as Set. For instance, as we have the Booleans in Set so each category of sheaves has a subobject classifier: an object of truth values. We construct products and subobject classifiers in categories of coalgebras as subcoalgebras of *cofree coalgebras*, the duals of free algebras. These constructions are illustrated in the category of unlabelled transition systems and functional strong bisimulations.

1.3 Notes for the reader

Chapter 2 outlines some of the basic ideas about coalgebras. Chapters 3, 4 and 5 contain the contributions of the thesis. They can be read independently, although the examples in chapter 5 assume a familiarity with section 3.6.

Throughout this thesis we will assume familiarity with the basic notions of category theory—category, functor, natural transformation, limits, the Yoneda embedding, and adjoints (see [17], for example)—but we explain more specialized concepts as they arise. Also, we should mention that we occasionally refer to pieces of 'abstract nonsense' in the body of the thesis, e.g. the Special Adjoint Functor Theorem, biequivalences etc. However, these always appear as tangential points, or in references to others' work, and are not crucial to the general development.

Theorems, propositions, lemmas, definitions, remarks, and examples are numbered on the same counter. Thus the number $l.m.n$ denotes the n-th numbered reference in section m of chapter l.

Chapter 2

Coalgebras and final semantics

In this chapter we survey some basic definitions, examples, and theorems from the literature on coalgebras. Almost all of these will be referred to in the following chapters, and here serve to motivate the subsequent development and provide a convenient summary of technical preliminaries. General introductions to the theory of coalgebras include Rutten [73], Gumm [37], and Jacobs and Rutten [42].

2.1 Coalgebras

Given a category $\mathcal{C}$ and an endofunctor $F\colon \mathcal{C} \to \mathcal{C}$, an *$F$-coalgebra* is a pair (A, α) consisting of an object A of $\mathcal{C}$ and a map $\alpha\colon A \to FA$. We call the object A the *carrier* of the coalgebra and α the *structure* of the coalgebra. We emphasize that we are talking about coalgebras *of endofunctors*. There is an older, different notion of coalgebra in the categorical literature, namely *coalgebra of a comonad*. In fact, the two notions of coalgebra are closely related and we utilize this connection in chapter 5.

The simplest examples of coalgebras arise when the category $\mathcal{C}$ above is a preorder and $F\colon \mathcal{C} \to \mathcal{C}$ is a monotone function. Then an F-coalgebra is just a post-fixed point of F, i.e., $x \in \mathcal{C}$ such that $x \sqsubseteq Fx$. The interest in this notion comes from Tarski's fixed point theorem which asserts that if $\mathcal{C}$ is a complete lattice, i.e., a small preorder with all joins (equivalently all meets), then the greatest fixed point of F is given by

$$\bigsqcup\{x \in \mathcal{C} \mid x \sqsubseteq Fx\}.$$

Later we will see that this order-theoretic result is the prototypical example of a *final coalgebra theorem.*

Moving to the richer categorical case, a basic intuition is that an F-coalgebra (A, α) is a generalized transition system: we view A as an object of states and α as giving a transition structure, or dynamics, cf. [4, 73]. Another intuition, particularly in case α is an isomorphism, is that a coalgebra is a data type; in this instance one thinks of α as determining an unfolding of elements of the data type. We are able to capture many familiar notions of transition system and data type as F-coalgebras by choosing the endofunctor F appropriately. We can think of F as providing a signature, or type information for the notion of coalgebra; we may also think of F as defining an interface to the (hidden) state space of the coalgebra.

2.1.1 Streams, Trees and Automata

Let $F_1 \colon \mathsf{Set} \to \mathsf{Set}$ be given by $F_1X = \mathbb{N} \times X$. Then an F_1-coalgebra (A, α) consists of a set A and a tuple of maps $\alpha = \langle \mathsf{head}, \mathsf{tail} \rangle$, where $\mathsf{head} \colon A \to \mathbb{N}$ and $\mathsf{tail} \colon A \to A$. The maps head and tail can be thought of as *destructors* for observing and modifying the hidden state of the coalgebra. For example, take A to be the set of infinite streams of naturals $\mathbb{N}^\omega$, and interpret head and tail in the obvious manner.

By moving to a richer signature we can give a coalgebra structure to the set T of finitely branching trees of possibly infinite depth with nodes labelled by natural numbers. Let $F_2 \colon \mathsf{Set} \to \mathsf{Set}$ be given by $F_2X = \mathbb{N} \times LX$, where LX is the set of finite lists of elements of X. Then the F_2-coalgebra structure on T is given by a tuple of maps $\langle \mathsf{label}, \mathsf{subtrees} \rangle$, where $\mathsf{label} \colon T \to \mathbb{N}$ and $\mathsf{subtrees} \colon T \to LT$ respectively map each tree to the label of its root node and the (ordered) collection of its subtrees.

Let A be a set of input symbols. A *deterministic automaton* is a triple $\mathcal{S} = (S, o, t)$ consisting of a set S of states, an *output function* $o \colon S \to 2$, and a *transition function* $t \colon S \to S^A$. Here 2 is the two element set $\{0, 1\}$, and S^A is the set of functions from A to S. The idea is that the transition function specifies the next state after accepting input $a \in A$, and the output function specifies whether a state s is *terminating* or not, according as $o(s) = 1$ or $o(s) = 0$ respectively. Now, the tuple $\langle o, t \rangle$ has type $S \to 2 \times S^A$, thus deterministic automata may be seen as coalgebras of $F_3 \colon \mathsf{Set} \to \mathsf{Set}$ where $F_3X = 2 \times X^A$.

Here is one important example of an F_3-coalgebra. It has as carrier the set $\mathcal{L} = \mathcal{P}(A^*)$ of languages over alphabet A, where A^* denotes the set of finite words over A and $\mathcal{P}$ denotes powerset. For $L \subseteq A^*$ and $a \in A$ write $L_a = \{v \in A^* \mid av \in L\}$ and $L\downarrow$ if $\varepsilon \in L$. The coalgebra structure $\langle o, t \rangle$ on $\mathcal{L}$ is given by

$$o(L) = \begin{cases} 1 & \varepsilon \in L \\ 0 & \varepsilon \notin L \end{cases} \quad \text{and } t(L)(a) = L_a.$$

Recall that a state s of an automaton is said to *accept a word* $a_1 \ldots a_n \in A^*$ if the state obtained by successively inputing $a_1, a_2, \ldots, a_n$ is terminating. The coalgebra $\mathcal{L}$ has the property that the set of words accepted by $L \in \mathcal{L}$ is L itself.

A *formal power series* is a function $\sigma\colon A^* \to k$ from the set of words over an alphabet A to a semiring k (a ring without subtraction). For example, the characteristic function of a language $L \subseteq A^*$ is a formal power series in the Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$. Putting $A = \{X\}$ and $k = \mathbb{R}$ we get the power series of classical analysis, which are nothing but streams of reals. Motivated by these examples, Rutten [78] considers the set of formal power series over an alphabet A and semiring k as a coalgebra of the set functor $F_4(X) = X^A \times k$. This forms the basis of a coalgebraic approach to some of the basic notions in the theory of formal power series. The semiring structure of k plays a crucial role in this development, even though it is not directly represented in the functor F_4.

2.1.2 Labelled transition systems

Recall the notion of a *labelled transition system* [69], usually arising in connection with the operational semantics of process calculi. Assume a fixed set of labels Act, then a labelled transition system is a pair $(\mathrm{Pr}, \to)$ consisting of a set Pr of processes (or states) and a ternary relation $\to \subseteq \mathrm{Pr} \times \mathrm{Act} \times \mathrm{Pr}$. One reads $(p, a, q) \in \to$ as p makes an a-labelled transition to become (or change state to) q; this is usually denoted $p \xrightarrow{a} q$.

Labelled transition systems can be seen as coalgebras of the functor

$$FX \quad = \quad \mathcal{P}(\mathrm{Act} \times X). \tag{2.1}$$

Here, $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$ is the powerset functor; it takes each set to the collection of its subsets and acts on functions by taking the direct image, i.e., given a function $f\colon X \to Y$ and $S \subseteq X$,

$$\mathcal{P}(f)(S) \quad = \quad \{f(s) \mid s \in S\}. \tag{2.2}$$

$\mathrm{Act} \times X$ denotes Cartesian product in Set, and is made functorial in X in the usual way. Now it is clear that F-coalgebras and labelled transition systems are in 1-1 correspondence via the natural isomorphism defining $\mathcal{P}$, viz.,

$$\mathrm{Sub}(X \times Y) \cong \mathrm{Hom}(X, \mathcal{P}Y)$$

where Sub is the contravariant subobject functor. To the best of our knowledge it was Aczel [4] who first made the connection between labelled transition systems and F-coalgebras.

If we replace $\mathcal{P}$ in (2.1) with the subfunctor $\wp$ sending each set to the collection of its finite subsets, then we obtain finitely branching labelled transition systems as coalgebras.

A more general model of processes is given by *labelled transition systems with divergence* [62]. Divergence, or livelock, models the possibility that a process may perform an infinite sequence of silent, internal actions. Formally, assuming a fixed set of actions Act, a labelled transition system with divergence is a triple

$$(\mathrm{Pr}, \to, \uparrow),$$

where $(\mathrm{Pr}, \to)$ is a labelled transition system, and $\uparrow \subseteq \mathrm{Pr}$ is the set of processes which may diverge. It is natural to model such systems as coalgebras of the functor

$$FX = \mathcal{P}((\mathrm{Act} \times X)_\perp)$$

on the category Preord of preorders and monotone maps. Here, $(-)_\perp$ denotes the lifting functor which adjoins a bottom element to its argument, and we take $\mathcal{P}(X)$ to have the Egli-Milner preorder $\sqsubseteq_{\mathrm{EM}}$, where

$$S \sqsubseteq_{\mathrm{EM}} T \;\text{ iff }\; \forall s \in S\, \exists t \in T\; s \sqsubseteq t \text{ and } \forall t \in T\, \exists s \in S\; s \sqsubseteq t$$

for $S, T \subseteq X$. The reading of a coalgebra (X, α) as a transition system is as before, except that we now say that $x \in X$ diverges iff $\perp \in \alpha(x)$ (think of the preorder on FX as being an *information preorder*, with $\perp$ denoting a complete absence of information about a computation).

2.1.3 Probabilistic transition systems

The labelled transition systems introduced above incorporate *non-determinism* in the sense that a given state may make an a-labelled transition to more than one state. Larsen and Skou [51] consider a notion of transition system where non-determinism is replaced by probabilistic choice; for each state s which can make an a-labelled transition the resultant state is determined by a probability distribution. Here, a discrete probability distribution on a set X is a function $\mu\colon X \to [0,1]$ with countable support, i.e., $\mu(x) \neq 0$ for only countably many $x \in X$, and $\Sigma_{x\in X}\mu(x) = 1$. Given $S \subseteq X$, we write $\mu[S]$ for $\Sigma_{x\in S}\mu(x)$.

Larsen and Skou were concerned with a testing perspective on processes; in particular, they consider whether there exists a realistic testing regime to distinguish *non-bisimilar* processes and give the same results for *bisimilar* processes (of which more later). We consider a slight generalization of their model, allowing states to terminate with probability other than just 0 or 1. Write 1 for the singleton $\{*\}$, where $*$ denotes inaction. A discrete probabilistic transition system over a fixed set of actions Act is a pair (Pr, μ), where Pr is a non-empty set of processes (or states) and $\mu\colon \mathrm{Pr} \times \mathrm{Act} \times (\mathrm{Pr} + 1) \to [0,1]$ is such that

$\mu(p, a, \cdot)$ is a discrete probability distribution. The idea is that $\mu(p, a, *)$ is the probability that p refuses an a-labelled action, and $\mu(p, a, p')$ is the probability that p accepts an a-labelled action and becomes p'.

Discrete probabilistic transition systems are also called *partial labelled Markov chains* in [23]; the adjective 'Markov' refers to the fact that the transition probabilities only depend on the current state and are otherwise history independent. They generalize traditional Markov chains from probability theory by the addition of labels and the possibility of termination.

Write $D(X)$ for the set of discrete probability distributions over a set X. We can extend D to an endofunctor on Set by defining, for a function $f\colon X \to Y$,

$$D(f)(\mu)(y) = \mu[f^{-1}(\{y\})].$$

We also consider the subfunctor D_f, where $D_f(X)$ is the set of finitely supported discrete probability distributions on a set X.

Consider the set functor $FX = D(1 + X)^{\mathrm{Act}}$. It is readily seen that F-coalgebras are in 1-1 correspondence with probabilistic transition systems. This is a slight variation on the approach of de Vink and Rutten [83] who consider coalgebras of the functor $(1 + D_f(-))^{\mathrm{Act}}$ (these correspond to probabilistic transition systems where states terminate with probability 0 or 1). In chapter 4 we will take a coalgebraic approach to probabilistic transition systems with continuous state spaces, like $\mathbb{R}$, where the transition behaviour of states is given by probability measures on the state space.

2.2 Homomorphisms and bisimulations

Let $F\colon \mathcal{C} \to \mathcal{C}$ be an arbitrary endofunctor, and assume two F-coalgebras (A, α) and (B, β). Then a map $f\colon A \to B$ is a *homomorphism of F-coalgebras* if $\beta \cdot f = Ff \cdot \alpha$, i.e., the diagram below commutes.

$$\begin{tikzcd} A \arrow[r, "\alpha"] \arrow[d, "f"'] & FA \arrow[d, "Ff"] \\ B \arrow[r, "\beta"'] & FB \end{tikzcd} \tag{2.3}$$

The identity map on A is a homomorphism $(A, \alpha) \to (A, \alpha)$, and the composition in $\mathcal{C}$ of two homomorphisms is also a homomorphism; thus, we have a category $\mathsf{Coalg}(F)$ of F-coalgebras and F-coalgebra homomorphisms.

Consider Act-labelled transition systems $(\mathrm{Pr}_1, \to_1)$ and $(\mathrm{Pr}_2, \to_2)$ as coalgebras of the functor (2.1). Then a map $f : \mathrm{Pr}_1 \to \mathrm{Pr}_2$ is a coalgebra morphism iff f preserves and reflects transitions, i.e., for all $p \in \mathrm{Pr}_1$,

(i) $\forall a \in \mathrm{Act}, q \in \mathrm{Pr}_1$, if $p \xrightarrow{a}_1 q$ then $f(p) \xrightarrow{a}_2 f(q)$;

(ii) $\forall a \in \mathrm{Act}, r \in \mathrm{Pr}_2$, if $f(p) \xrightarrow{a}_2 r$ then $\exists q \in \mathrm{Pr}_1$ with $p \xrightarrow{a}_1 q$ and $f(q) = r$.

Such functions are sometimes called zig-zag maps.

If (Pr_1, μ_1) and (Pr_2, μ_2) are probabilistic transition systems, i.e., coalgebras of the functor $D(1 + (-))^{\mathrm{Act}}$ introduced above, then $f : \mathrm{Pr}_1 \to \mathrm{Pr}_2$ is a coalgebra map iff for all $p \in \mathrm{Pr}_1, a \in \mathrm{Act}$ and $t \in \mathrm{Pr}_2 + 1$

$$\mu_2(f(p), a, t) = \mu_1(p, a, (f + 1)^{-1}(t)).$$

Next, we introduce the notion of a *coalgebraic bisimulation*, due to Aczel and Mendler [3], a relational generalization of coalgebra homomorphisms (cf. proposition 2.2.1).

In a category $\mathcal{C}$ with binary products, a *binary relation* R between objects A and B is a span $A \xleftarrow{\pi_1} R \xrightarrow{\pi_2} B$ such that π_1 and π_2 are *jointly monic*, i.e., the tuple $\langle \pi_1, \pi_2 \rangle : R \to A \times B$ is monic. Should A and B carry F-coalgebra structures α and β respectively, we say that R is an *F-bisimulation* between (A, α) and (B, β) if there exists a map ρ making the diagram below commute.

$$\begin{tikzcd} A \arrow[d, "\alpha"'] & R \arrow[l, "\pi_1"'] \arrow[r, "\pi_2"] \arrow[d, dashed, "\rho"] & B \arrow[d, "\beta"] \\ FA & FR \arrow[l, "F\pi_1"'] \arrow[r, "F\pi_2"] & FB \end{tikzcd} \tag{2.4}$$

The prototypical example of an F-bisimulation occurs when $F = \mathcal{P}(\mathrm{Act} \times (-))$. In that case, given two F-coalgebras (i.e., labelled transition systems) $(\mathrm{Pr}, \to)$ and $(\mathrm{Pr}', \to)$, $R \subseteq \mathrm{Pr} \times \mathrm{Pr}'$ is an F-bisimulation precisely when for all $p \in \mathrm{Pr}, q \in \mathrm{Pr}'$ and $a \in \mathrm{Act}$, $R(p, q)$ implies that

(i) whenever $p \xrightarrow{a} p'$ then for some q', $q \xrightarrow{a} q'$ and $(p', q') \in S$;

(ii) whenever $q \xrightarrow{a} q'$ then for some p', $p \xrightarrow{a} p'$ and $(p', q') \in S$.

This is precisely the notion of a *strong bisimulation* between labelled transition systems which has been a cornerstone of concurrency theory since its introduction by Park [65]. We say that processes p and q are bisimilar just in case there exists a bisimulation relating them. The relation of bisimilarity is considered the strongest (most discriminating) notion of behavioural equivalence for processes; cf. [61].

A fully general characterization of bisimulations for coalgebras of the functor $D(1 + (-))$ is more problematic, cf. [83]. In chapter 4 we consider a new notion of coalgebraic bisimulation for which this task is much easier.

We remarked on the fundamental place that the notion of bisimilarity has in concurrency theory. This is due in large part to its pleasant mathematical properties: its characterizations via modal logic and as a greatest fixed point, also its connection with non-well-founded set theory. However, from a more pragmatic 'testing' viewpoint it has been questioned whether bisimilarity is a realistic notion of observational equivalence. In particular, Bloom, Estrail and Meyer [16] show that one cannot capture bisimulation as a congruence generated by trace equivalence from any 'reasonable' set of process constructions. It was this work which motivated Larsen and Skou to look at probabilistic transition systems as a model for processes. They showed in [51] that in this model (probabilistic) bisimilarity can indeed be captured as a testing equivalence for a realistic notion of test.

2.2.1 Duality: algebras and congruences

One of the guiding principles in the theory of coalgebras is the duality with algebra. An algebra of an endofunctor $F : \mathcal{C} \to \mathcal{C}$ is a pair (A, α) where A is an object of $\mathcal{C}$ and $\alpha : FA \to A$. Given another F-algebra (B, β), an arrow $f : A \to B$ is a homomorphism of F-algebras if $f \cdot \alpha = \beta \cdot Ff$; this yields a category $\mathsf{Alg}(F)$ of F-algebras and F-algebra homomorphisms. Observe that an F-algebra is the same as a coalgebra of the opposite functor $F^{\mathrm{op}} : \mathcal{C}^{\mathrm{op}} \to \mathcal{C}^{\mathrm{op}}$, and the category $\mathsf{Alg}(F)$ is the same as $(\mathsf{Coalg}(F^{\mathrm{op}}))^{\mathrm{op}}$.

Let Σ be a signature in the sense of universal algebra, i.e., a set of operation symbols, each with a prescribed arity given by $ar : \Sigma \to \mathbb{N}$. Define a functor $T_\Sigma : \mathsf{Set} \to \mathsf{Set}$ on objects by

$$T_\Sigma(X) = \coprod_{\sigma \in \Sigma} X^{ar(\sigma)}$$

(where X^n is the nth Cartesian power of X) and on morphisms by

$$T_\Sigma(f) = \coprod_{\sigma \in \Sigma} f^{ar(\sigma)}$$

(where f^n assigns to each $(x_i)_{i<n}$ the n-tuple $(f(x_i))_{i<n}$). Clearly a T_Σ-algebra is exactly a Σ-algebra in the sense of universal algebra, and a function is a homomorphism of T_Σ-algebras precisely when it is a homomorphisms of Σ-algebras. This example explains why *algebras* (and by duality, coalgebras) of endofunctors are so-called.

An important notion from universal algebra is that of a *congruence*, i.e., the kernel of a homomorphism. Aczel and Mendler [3] carry this definition over to coalgebras: a congruence on an F-coalgebra is the kernel of a coalgebra homomorphism. Below, we summarize some simple properties of bisimulations and congruences, and indicate how these two notions are related.

2.2.2 Some facts on bisimulations and congruences

The following results are taken from [73, 3]. They are proven there for an endofunctor $F\colon \mathsf{Set} \to \mathsf{Set}$, but still hold when Set is replaced by an arbitrary *regular category*; see chapter 5. First we recall some basic notions for relations in Set. Given a function $f\colon A \to B$, the graph of f is the relation $\Gamma_f = \{\langle a, fa\rangle : a \in A\}$. The composition of relations $R \subseteq A \times B$ and $S \subseteq B \times C$ is defined by

$$S \cdot R = \{\langle a, c\rangle \mid \exists b \ \langle a, b\rangle \in R \text{ and } \langle b, c\rangle \in S\}.$$

The identity relation on a set A is $\Delta_A = \{\langle a, a\rangle : a \in A\}$. Finally, the inverse of a relation $R \subseteq A \times B$ is the relation $R^\circ = \{\langle b, a\rangle : \langle a, b\rangle \in R\}$.

Proposition 2.2.1 Given F-coalgebras (A, α) and (B, β), a map $f\colon A \to B$ is a coalgebra homomorphism iff its graph Γ_f is an F-bisimulation. $\square$

Proposition 2.2.2 If R is an F-bisimulation between F-coalgebras (A, α) and (B, β). Then the inverse relation R° is an F-bisimulation between (B, β) and (A, α). $\square$

Recall that a cone over a diagram in a category is said to be a *weak limit* if it satisfies the existence clause but not necessarily the uniqueness clause in the definition of a limit. All the set functors introduced in this chapter preserve weak pullbacks, cf. [73].

Proposition 2.2.3 Suppose that F preserves weak pullbacks, and let R be an F-bisimulation between (A, α) and (B, β), and S an F-bisimulation between (B, β) and (C, γ), then the composite relation $S \cdot R$ is an F-bisimulation between (A, α) and (C, γ). $\square$

Proposition 2.2.4 The union $\bigcup_k R_k$ of a family $\{R_k\}_k$ of F-bisimulations between coalgebras (A, α) and (B, β) is again a bisimulation. The transitive closure of the union of a family of congruences on a coalgebra is again a congruence. $\square$

From the above proposition it follows that on any coalgebra there is a largest congruence and a largest bisimulation; the latter is called *F-bisimilarity*. Should F preserve weak pullbacks then F-bisimilarity will be an equivalence relation by propositions 2.2.2 and 2.2.3.

Proposition 2.2.5 If F preserves weak pullbacks then every congruence is a bisimulation and every bisimulation equivalence is a congruence. □

It follows that, if F preserves weak pullbacks, the relation of F-bisimilarity on any coalgebra is also the largest congruence. We conclude this section by reviewing the construction of the relation of strong bisimilarity on a labelled transition system $(\mathrm{Pr}, \rightarrow)$ by ordinal iteration.

Let $\mathcal{F}$ be the function over binary relations $S \subseteq \mathrm{Pr} \times \mathrm{Pr}$ defined as follows: $(p, q) \in \mathcal{F}S$ iff, for all $a \in \mathrm{Act}$,

(i) Whenever $p \xrightarrow{a} p'$ then for some q', $q \xrightarrow{a} q'$ and $(p', q') \in S$;

(ii) Whenever $q \xrightarrow{a} q'$ then for some p', $p \xrightarrow{a} p'$ and $(p', q') \in S$.

$\mathcal{F}$ is a monotone map on the complete lattice of relations on Pr ordered by inclusion, and a relation is a strong bisimulation precisely when it is a post-fixed point of $\mathcal{F}$. Now consider the decreasing ordinal-indexed sequence of equivalence relations $\approx_\alpha$ on Pr defined by:

(i) $\approx_0 = \mathrm{Pr} \times \mathrm{Pr}$

(ii) $\approx_{\alpha+1} = \mathcal{F}(\approx_\alpha)$

(iii) $\approx_\lambda = \bigcap_{\alpha<\lambda} \approx_\alpha$, for a limit ordinal λ.

The relation $\approx_\alpha$ is called bisimulation up-to α steps. From Hartog's lemma (there is no set into which every ordinal can be injected) it follows that a decreasing ordinal-indexed sequence in a partial order must converge, thus $\approx_\kappa = \mathcal{F}(\approx_\kappa)$ for some κ. Then $\approx_\kappa$ is the greatest strong bisimulation on $(\mathrm{Pr}, \rightarrow)$.

2.2.3 Open maps

Another category-theoretic definition of bisimulation is provided by the work of Joyal, Nielsen and Winskel on *open maps* [46]. This notion gives a uniform definition of bisimulation over a range of different categories of 'transition systems and simulation maps'. Their prototypical example is the category of labelled transitions systems and transition preserving maps. In this approach one identifies a certain subclass of simulation maps, called open maps, satisfying a 'path-lifting' property defined with respect to a specified subcategory of observation objects. Bisimulations are then defined to be spans of open maps. For instance, the coalgebra maps between labelled transition systems and probabilistic transitions systems, characterized above, are examples of open maps. More generally, it is shown in [52] that any category with a fixed class of open maps can be embedded in category of coalgebras and *lax* coalgebra morphisms such that the proper coalgebra morphisms correspond precisely to the open maps.

In the open maps approach, in common with coalgebraic bisimulations in the style of Aczel and Mendler, bisimulations are defined as *spans* of morphisms. In contrast, in chapter 4, we define a notion of bisimulation based on *cospans* of coalgebra maps.

2.3 Final coalgebras

In this section we review some constructions of *final coalgebras* for functors on categories of sets, complete partial orders and metric spaces. An object C in a category is *final* if, for any other object C', there is a unique arrow $C' \to C$. Dually, an object C is *initial* if, for any other object C', there is a unique arrow $C \to C'$. Initial objects and final objects are unique up-to isomorphism; usually one denotes the initial object of a category by 0, and the final object by 1. Final coalgebras are of particular interest as models of recursively defined data types and as domains for interpreting programming languages, like CCS [61] and the pi-calculus [40]. In fact, all the coalgebras presented in section 2.1.1 are examples of final coalgebras.

Earlier we mentioned that the greatest post-fixed point of a monotone function on a partial order is, in fact, a fixed point. The categorical generalization of this is called Lambek's lemma.

Lemma 2.3.1 [50] Let $F\colon \mathcal{C} \to \mathcal{C}$ be an endofunctor. If (A, α) is a final coalgebra then α is an isomorphism; in other words, (A, α) is a *fixed point* or *invariant* of F. □

The dual of this result states that an initial F-algebra is a fixed point of F. Note that we take a fixed point of functor F to be an object X together with a specified isomorphism $i\colon X \cong FX$. Clearly, a fixed point may be regarded as both an F-algebra and an F-coalgebra.

Recall that a *complete partial order* (cpo) is a partial order with joins of all ω-chains, a *pointed* cpo is a cpo with a least element $\perp$, and a *continuous* map of cpos is a monotone function preserving the joins of all ω-chains. As is very well-known, the least fixed point of a continuous function on a pointed cpo may be constructed as $\bigsqcup_n f^n(\perp)$. The categorical generalization of this construction appears as the 'basic lemma' in Smyth and Plotkin [80]. The dual form of this result is as follows.

Theorem 2.3.2 Let $\mathcal{C}$ be a category with final object and limits of all co-chains indexed by the partial order ω^{op}. If $F\colon \mathcal{C} \to \mathcal{C}$ preserves ω^{op}-limits then there is a final F-coalgebra.

Proof. We sketch only the barest details since in the next chapter we give a full treatment of (a transfinite generalization of) this final coalgebra construction. One constructs the carrier of the final coalgebra as the limit L of the ω^{op}-co-chain

$$1 \xleftarrow{!} F1 \xleftarrow{F!} F^2 1 \xleftarrow{F^2!} F^3 1 \longleftarrow \dots$$

where 1 denotes the final object of $\mathcal{C}$. The structure of the final coalgebra is an isomorphism $L \cong FL$ which arises from the universal property of L and of FL (both are ω^{op}-limits). □

The dual of this result constructs an initial F-algebra as a colimit of an ω-chain. Some early uses of theorem 2.3.2 (in dual form) to model recursive data types as initial algebras of functors on Set and on categories of cpos are due to G. Plotkin, M. Smyth and M. Wand [80, 85]. Later, Arbib and Manes [10] modelled recursive data types as final coalgebras of set functors, again constructed by theorem 2.3.2. More recently still, America and Rutten [9] have used a similar construction to generate fixed points of functors on categories of complete metric spaces, extending and generalizing earlier work of de Bakker and Zucker [12] on metric space domains for concurrency.

2.3.1 Set functors

A different construction of final coalgebras for set functors was introduced by Aczel and Mendler [3]. This was inspired by Milner's operational semantics for the process algebra CCS. In particular, [3] introduced the notions of coalgebraic congruences and bisimulations; they obtain a final coalgebra by first constructing a weakly final coalgebra as the coproduct of all 'small' coalgebras and then quotienting this by the maximum congruence (for most functors this is the same as the maximum bisimulation). Strictly speaking, the theorem of Aczel and Mendler was stated for so-called set based endofunctors on the category SET of classes and class functions. Some kind of 'accessibility' hypothesis like this is necessary to ensure the existence of a final coalgebra. For instance, the powerset functor $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$ cannot have a final coalgebra since no set is is isomorphic to its own powerset (Cantor's theorem).

Definition 2.3.3 An endofunctor F on the category SET is said to be *set based* if for every class A and $a \in FA$, there is a set A_0 and an injection $i\colon A_0 \rightarrowtail A$ such that a lies in the image of Fi. □

Theorem 2.3.4 [3] Every set-based endofunctor $\mathsf{SET} \to \mathsf{SET}$ has a final coalgebra. □

In [13] Barr pointed out that theorem 2.3.4 is implicit in some more general work of Makkai and Paré on accessible categories and accessible functors [58]. In particular, one has a version of theorem 2.3.4 for endofunctors on Set which are finitary in some sense. More generally still, one has final coalgebras for finitary functors on so-called *locally presentable* categories, e.g., the category Pos of posets and monotone maps. Below, we give a particularly simple formulation of the final coalgebra theorem for set functors.

Definition 2.3.5 [47, 37] A set functor F is *bounded* if there exists some cardinal κ such that, for every F-coalgebra (A, α) and every $a \in A$, there exists a subcoalgebra, i.e. a mono $m\colon (U_a, \iota) \rightarrowtail (A, \alpha)$, such that a lies in the image of the function m and U_a has cardinal less than κ. □

Theorem 2.3.6 [47, 37] If a set functor is bounded then it has a final coalgebra. □

There are many set functors which satisfy the hypotheses of theorem 2.3.6 but not those of theorem 2.3.2, e.g., the finite-powerset functor $\wp$ and the probability-distributions functor D from section 2.1. It is also worth noting that in practice it is much easier to verify the hypotheses of theorem 2.3.6 than those of theorem 2.3.2.

A radically different approach to the construction of final coalgebras for set functors is provided by Aczel's *Special Final Coalgebra Theorem (SFCT)* [4]. This result is set-theoretic in nature, not category-theoretic. In particular, it assumes that the universe of sets is non-well-founded. Recall that the foundation axiom of set theory prohibits infinitely descending chains of set membership ... $\in x_3 \in x_2 \in x_1$. In non-well-founded set theory one drops foundation and replaces it with an *anti-foundation axiom (AFA)* which guarantees the existence and uniqueness of non-well-founded sets such as $\Omega \stackrel{\text{def}}{=} \{\Omega\}$. A particularly concise formulation of this axiom can be found in [82]. Non-well-founded set theory has quite a long history prior to the work of Aczel; in particular, AFA was originally posited by Forti and Honsell [31] under the less memorable name 'X_1-axiom'. However, SFCT is now seen as one of the major motivations behind the adoption of the anti-foundation axiom.

Theorem 2.3.7 (SFCT) Assume that the universe of sets satisfies AFA. Let F be a set based endofunctor on SET. If F is uniform on maps and preserves inclusions then the greatest fixed point of F (with respect to the inclusion order on SET) is given by the Tarski formula

$$\bigcup\{X \mid X \text{ a set, and } X \subseteq FX\}, \tag{2.5}$$

and is a final F-coalgebra (with structure map the identity). □

A key hypothesis of the SFCT is that the functor in question is 'uniform on maps'. We omit a formal definition of this property and refer the reader to [82]. The condition is usually explained as meaning that, for a map $f\colon X \to Y$, the action of $Ff\colon FX \to FY$

is a substitution of each x from X occurring in FX with fx. For instance, the equation (2.2) shows that the powerset functor is uniform on maps. We would like to draw attention to another important intuition behind the uniformity on maps condition. We see it as a guardedness assumption: all elements x from X occurring in FX are guarded by a $\{...\}$. To illustrate this, observe that the identity functor $X \mapsto X$ is not uniform on maps! Indeed, the final coalgebra is a singleton set but the greatest fixed point is the universe V. However, the naturally isomorphic functor $X \mapsto \{\{x\} : x \in X\}$ is uniform on maps. In fact, the greatest fixed point is the set Ω mentioned above. All this serves to illustrate the set-theoretic (as opposed to category-theoretic) nature of SFCT.

What theorem 2.3.7 provides is set-theoretic *representations* of final coalgebras. For example, the final coalgebra of the functor F_1 from section 2.1.1 can be described as the set of infinite tuples $\langle n_0, \langle n_1, \langle n_2, \langle ... \rangle\rangle\rangle$ where $n_i \in \mathbb{N}$. Note that, assuming the standard Kuratowski encoding of ordered pairs, these tuples are non-well-founded sets. However, adopting a non-standard encoding of ordered pairs, due to Quine, they can be seen as well-founded sets; indeed, Paulson [66] has developed a restricted version of SFCT for well-founded set theory based around just this idea.

2.3.2 Complete partial orders

Earlier we saw some examples of set functors which do not satisfy the hypotheses of theorem 2.3.2. In contrast, for functors on categories of cpos and complete metric spaces one has readily verified and widely satisfied conditions ensuring that the hypotheses of the basic lemma are met. These conditions exploit the *enriched structure* of these categories, cf. [48].

A category $\mathcal{C}$ is called *poset-enriched*, or a Pos-*category*, if each homset can be given the structure of a poset such that, for each triple X, Y, Z of objects of $\mathcal{C}$, the operation of composition in $\mathcal{C}$

$$\circ : \mathcal{C}(Y,Z) \times \mathcal{C}(X,Y) \to \mathcal{C}(X,Z)$$

is a monotone map. An example of a poset-enriched category is Pos itself, where the homsets are given the point-wise order, i.e., given $f, g \colon X \to Y$, $f \sqsubseteq g$ iff $fx \sqsubseteq gx$ for all $x \in X$.

A functor $F \colon \mathcal{C} \to \mathcal{D}$ between Pos-categories is called *locally monotone*, or *a* Pos-*functor*, if for each pair of objects $X, Y \in \mathcal{C}$

$$F_{XY} : \mathcal{C}(X,Y) \to \mathcal{D}(FX,FY)$$

is a monotone map.

In the above definitions, if we replace posets with cpos, and monotone functions with continuous functions, we obtain the concept of a Cpo-*category* and of a Cpo-*functor*, also called a *locally continuous functor*. Just as Pos is poset enriched, so Cpo is cpo-enriched with the pointwise order on homsets. Another important example of a Cpo-category is the category $\mathsf{Cppo}_\perp$ of pointed cpos and *strict* continuous maps, i.e., continuous maps that preserve least elements.

Broadly speaking, local continuity of an endofunctor F on a Cpo-category $\mathcal{C}$ is enough to ensure that F preserves ω^{op}-limits and ω-colimits. Thus the basic lemma (and its dual) can be used to construct a final coalgebras (and initial algebras) of locally continuous endofunctors. Not only that, but we can also use the basic lemma to construct fixed points of a so-called *mixed variance* functor

$$F : \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}$$

(for example the exponentiation functor on cpos).

It is worth spelling out in a little detail how this works. The key idea, due to Smyth and Plotkin [80], is to consider an auxiliary category $\mathcal{C}^{\mathrm{E}}$ with the same objects as $\mathcal{C}$ but with so-called embeddings as morphisms. An arrow $f\colon X \to Y$ of $\mathcal{C}$ is called an *embedding* if there is another arrow $g\colon Y \to X$, called the *projection*, such that $g \cdot f = 1_X$ and $f \cdot g \sqsubseteq 1_Y$. Each half of an embedding-projection pair (f, g) uniquely determines the other, and we write f^R for the projection associated with an embedding f (read R as standing for right adjoint). Now each mixed variance functor F, as above, can be resolved into functor

$$F^{\mathrm{E}} : \mathcal{C}^{\mathrm{E}} \times \mathcal{C}^{\mathrm{E}} \to \mathcal{C}^{\mathrm{E}}$$

by writing $F^{\mathrm{E}}(f, g) = F(f^R, g)$.

Again, the local continuity of F is enough to ensure that F^{E} preserves ω^{op}-limits and ω-colimits. Then, writing $\Delta\colon \mathcal{C}^{\mathrm{E}} \to \mathcal{C}^{\mathrm{E}} \times \mathcal{C}^{\mathrm{E}}$ for the diagonal functor, applying the basic lemma to $F^{\mathrm{E}} \cdot \Delta$ we obtain a fixed point of F:

$$i : \mathrm{Fix}(F) \cong F(\mathrm{Fix}(F), \mathrm{Fix}(F)).$$

It does not make sense to ask if this in an initial algebra or a final coalgebra of F, because F is not an endofunctor. Nevertheless, one can still express the canonicity of this fixed point using the notion of minimal invariance, cf. Pitts [67].

Definition 2.3.8 [67, Defn. 3.2] Let $F\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}$ be a locally continuous functor. An invariant $i : F(D, D) \cong D$ is *minimal* if 1_D is the least fixed point of the continuous function $\delta\colon \mathcal{C}(D, D) \to \mathcal{C}(D, D)$, defined by $\delta(e) = i \cdot F(e, e) \cdot i^{-1}$. $\square$

At first sight the property of minimal invariance seems to be of a different nature than the criterion of initiality and finality for fixed points of endofunctors. In fact there is a close relationship: a minimal invariant of an endofunctor F (which is a trivial instance of a mixed variance functor) is both an initial algebra and a final coalgebra [1]. Not all Cpo-categories posses minimal invariants of locally continuous functors. The classic example of one which does is $\mathsf{Cppo}_{\perp}$, cf. Pitts [67].

The following lemma is crucial for applications of coinduction to minimal invariants. It is a special case of lemma 8.2 in Fiore [28]. That lemma is proved using a categorical generalization of Bekić's theorem, but working with the notion of minimal invariant we can give a more elementary proof using only the partial order version. Let us write $\mathrm{fix}(f)$, or $\mu x.fx$, for the least fixed point of a continuous map f on a cppo.

Lemma 2.3.9 Suppose $F\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}$ has minimal invariant $i\colon F(D,D) \cong D$. Then i is also a minimal invariant of $F(D,-)\colon \mathcal{C} \to \mathcal{C}$.

Proof. Suppose we are given cppos X and Y, and continuous maps $\phi\colon X \times Y \to X$ and $\psi\colon X \times Y \to Y$. Bekić's theorem states that the least fixed point of $\langle \phi, \psi \rangle$ is (x_0, y_0), where $x_0 = \mu x.\phi(x, \mu y.\psi(x,y))$ and $y_0 = \mu y.\psi(x_0, y)$.

For δ in definition 2.3.8, we have that $(\mathrm{fix}(\delta), \mathrm{fix}(\delta))$ is the least fixed point of

$$\theta\colon \mathcal{C}(D,D) \times \mathcal{C}(D,D) \to \mathcal{C}(D,D) \times \mathcal{C}(D,D)$$

given by $\theta(f,g) = (iF(g,f)i^{-1}, iF(f,g)i^{-1})$. It now follows that $1_D = \mu e.iF(D,e)i^{-1}$. □

In particular, if $i\colon F(D,D) \cong D$ is a minimal invariant, then (D, i^{-1}) is a final $F(D,-)$-coalgebra.

2.3.3 Metric Spaces

A *metric space* is a set X together with a *distance function* $d\colon X \times X \to [0,\infty]$ satisfying the following axioms:

(1) $d(x,y) = 0$ iff $x = y$;

(2) $d(x,z) \leq d(x,y) + d(y,z)$;

[1] Turning things on their head, Freyd [29] considers the coincidence of initial algebras and final coalgebras—called *algebraic compactness*—as one of the axioms a category should satisfy in order to qualify as a category of semantic domains. For instance, he shows how one can deduce the existence fixed points of mixed variance functors from this axiom.

(3) $d(x, y) = d(y, x)$.

The second axiom is called the *triangle inequality.* We call X an *ultrametric space* if a stronger version of the triangle inequality holds, viz., $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. If the equivalence in (1) is reduced to a right-to-left implication then we call X a pseudometric space. More radically, in chapter 4, we consider a further weakening of the axioms by dropping symmetry (3), yielding so-called *generalized metric spaces.* Conventionally, ∞ is not included in the range of the distance function of a metric space; our choice to include it is not just a matter of taste, it will be important in chapter 4 to ensure that the set of distances is a complete lattice. For reasons which also become clear in chapter 4 we assume that distances in ultrametric spaces are bounded by 1.

Given metric spaces (X, d_X) and (Y, d_Y), a map $f\colon X \to Y$ is called *non-expansive* if $d_Y(fx, fy) \leq d_X(x, y)$ for all $x, y \in X$. Metric spaces and non-expansive maps form a category Met. Next, we identify two special subclasses of non-expansive maps. If the inequality above is actually an equality we say that f is an *isometry.* We say that f is a *contraction* if there is a real number $0 \leq c < 1$ such that $d_Y(fx, fy) \leq c \cdot d_X(x, y)$ for all $x, y \in X$.

As in the order-theoretic case, metric domain theory is based on notions of approximation, limits, and completeness. The basic notion is that of a *Cauchy sequence.* A sequence $(x_n)_{n\in\mathbb{N}}$ of points of a metric space (X, d) is Cauchy if

$$\forall\varepsilon > 0\, \exists k\, \forall m, n \geq k.\, d(x_m, x_n) < \varepsilon.$$

We say that the sequence (x_n) converges to limit $x \in X$ if

$$\forall\varepsilon > 0\, \exists k \forall n \geq k.\, d(x_n, x) < \varepsilon.$$

A subset $S \subseteq X$ is *dense* in (X, d) if every point of X is the limit of some sequence of points in S. (X, d) is *separable* if it has a countable dense subset, e.g., $\mathbb{Q} \subseteq \mathbb{R}$. We say that $S \subseteq X$ is *compact* if every sequence of points of X contains a subsequence which converges to a point of S. The notion of convergence defined above induces a topology on a metric space (X, d). Formally, the topology has base the set of *open ε-balls* $B_\varepsilon(x) = \{y \in X : d(x, y) < \varepsilon\}$ where $x \in X$ and $\varepsilon > 0$. The notions of density and compactness for metric spaces are instances of the more general notions of the same name for topological spaces.

We say that (X, d) is *complete* if every Cauchy sequence converges to a limit in X. Our interest in this notion is with regard to Banach's fixed point theorem, which says that a contractive map on a complete metric space has a unique fixed point. A metric space (X, d) has a canonical isometric embedding as a dense subset of a complete space, its so-called *Cauchy completion* $(\overline{X}, \overline{d})$[2]. Cauchy completion extends to a functor from Met to the category of complete metric spaces and non-expansive maps, CMet, which is left adjoint to the inclusion $\mathsf{CMet} \hookrightarrow \mathsf{Met}$. In categorical terminology, CMet is a full, reflective subcategory of Met.

A category $\mathcal{C}$ is called *metric-enriched*, or a Met-*category*, if each homset can be given the structure of a metric space such that, for each triple X, Y, Z of objects of $\mathcal{C}$, the operation of composition in $\mathcal{C}$

$$\circ : \mathcal{C}(Y, Z) \otimes \mathcal{C}(X, Y) \to \mathcal{C}(X, Z)$$

is a non-expansive map. Here, the tensor product $(X, d_X) \otimes (Y, d_Y)$ of metric spaces (X, d_X) and (Y, d_Y) has $X \times Y$ as its underlying set, with metric

$$d(\langle x, y\rangle, \langle x', y'\rangle) = d_X(x, x') + d_Y(y, y');$$

cf. Lawvere [53]. Note that this is not the Cartesian product in Met (for which $+$ above would be replaced by max); in fact, Met is not Cartesian closed, but it is closed with respect to the tensor product introduced above. As a result we have that Met is itself a metric-enriched category, where the homsets are given the sup metric, i.e., given $f, g \colon X \to Y$,

$$d_{X \to Y}(f, g) = \sup_{x \in X} d_Y(fx, gx).$$

A functor $F \colon \mathcal{C} \to \mathcal{D}$ between Met-categories is called *locally non-expansive*, or *a* Met-*functor*, if for each pair of objects $X, Y \in \mathcal{C}$

$$F_{XY} : \mathcal{C}(X, Y) \to \mathcal{D}(FX, FY)$$

is a non-expansive map. F is called *locally contractive* if each F_{XY} is a contractive map.

Theorem 2.3.10 [9] Let $\mathcal{C}$ be a CMet-category, and $F \colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}$ a locally contracting functor; then F has a unique fixed point.

Proof. (Sketch) The proof of this result is a metric analogue of the construction of fixed points of Cpo-functors. The idea is to consider a metric variant of the category of embedding-projection pairs, resolving F into a covariant endofunctor F^{E} on this category. Analogously to the Cpo-enriched case, the preservation of ω^{op}-limits by F^{E} follows from the fact that it is locally contracting. Thus one can construct a fixed point using the basic lemma. In fact, one can show that any fixed point of F^{E} is a final coalgebra, so the fixed point constructed is unique up-to isomorphism. $\square$

[2] We have the following construction of $(\overline{X}, \overline{d})$. Define a pseudometric on the set of Cauchy sequences in X by $d((x_n), (y_n)) = \mathrm{Lim}_n d(x_n, y_n)$, then factoring out sequences with distance 0 we obtain a metric space.

Theorem 2.3.11 [82, Thm 7.3] Suppose $\mathcal{C}$ is a CMet-category, and $F\colon \mathcal{C} \to \mathcal{C}$ is a locally contracting functor with (unique) fixed point $i : \mathrm{Fix}(F) \to F(\mathrm{Fix}(F))$. Then $(\mathrm{Fix}(F), i)$ is a final F-coalgebra and $(\mathrm{Fix}(F), i^{-1})$ is an initial F-algebra. □

2.4 Coinduction

The final coalgebra theorems described above make cpos and complete metric spaces convenient for modelling recursive data types. A natural question arising from this is whether the use of metrics and partial orders is just a formal trick to provide us with the fixed point theorems we need to model recursion, or whether they have any significance beyond that.

An answer to this question, for partial orders, is found in the *full abstraction* theorems of Abramsky, Pitts, Fiore, Turi and Rutten [2, 28, 68, 82]. Suppose one has a model of a data type as the final coalgebra (X, α) of a Cpo-endofunctor F. The coalgebra (or unfolding) structure of the data type enables one to think of it as a transition system and to consider notions of simulation on it. For instance, Fiore [28] has a notion of *ordered F-bisimulation*, an order enriched version of the coalgebraic bisimulation notion in section 2.2. In this context a full abstraction theorem asserts that the order (or inequality) relation on X is the maximum simulation. Thus, the domain (X, α) is *abstract* with respect to the operational semantics provided by the notion of simulation. One can also think of such a result as a (*coinductive*) proof principle: to prove that one element approximates another, exhibit a simulation between them.

For final coalgebras of set functors there is no inequality relation to characterize. Instead one has a coinduction result characterizing the *equality* relation on the carrier of a final F-coalgebra as the maximum F-bisimulation, cf. [4]. About the same time that Aczel proved this result for set functors, van Glabbeek and Rutten [33] showed that de Bakker and Zucker's metric space domain for processes was abstract with respect to strong bisimulation. In fact, from the realisation that the (unique) fixed point $\mathrm{Fix}(F)$ of $F\colon \mathsf{CMet} \to \mathsf{CMet}$ is a final F-coalgebra, it easily follows by the same argument as for set functors that the equality relation is the maximum F-bisimulation on $\mathrm{Fix}(F)$, cf. [82]. However, in line with the full abstraction results for partial orders, one might also hope to characterize the *metric* rather than just the equality relation, on $\mathrm{Fix}(F)$ as the maximum bisimulation; we prove such a result in chapter 4.

The name coinduction for this kind of proof principle is due to Milner and Tofte [63]. To be precise, they worked with a set-theoretic formulation of coinduction whereas we are particularly concerned with a category-theoretic formulation—*coalgebraic coinduction*. Connections and differences between the two are spelt out by Lenisa [54]. Some applications of coinduction are: Milner [61], to prove the validity of various algebraic laws in the process algebra CCS; Rutten [74], to prove equalities between certain regular expressions; and Gordon [35], to reason about the equality of programs in higher order functional programming languages.

Coinduction really comes into its own as a proof principle when placed in tandem with *definition by coinduction.* The idea behind the latter is that to define a function from a set Y into a final F-coalgebra (X, α) it suffices to define an F-coalgebra structure β on the set Y. Then the finality of (X, α) yields a map of coalgebras $!\colon Y \to X$ called the *coinductive extension* of β. For instance, Rutten [78] coinductively defines new and known operators on formal power series, like shuffle product, and uses coinduction to prove various properties of these operators.

Coinductively defined functions into final coalgebras are important in the semantics of programming languages. In particular, Aczel [4] and Turi [82] show that the 'denotation function' from the syntax of a programming language into a semantic domain can often be seen as being coinductively defined. From this point of view, an important requirement is that a final coalgebra semantics should correspond to the notion of bisimulation at hand. More precisely, for any F-coalgebra (A, α), the kernel pair[3] of the unique map to the final coalgebra should be the greatest F-bisimulation on (A, α); in other words elements have the same final semantics iff they are bisimilar.

A criterion ensuring that this condition holds is that F preserves weak pullbacks, cf. [3, Prop 6.2]. The reason for this is that if F preserves weak pullbacks then pullbacks of F-coalgebra maps *lift to F-bisimulations.* That is, if $f : (A, \alpha) \to (C, \gamma)$ and $g\colon (B, \beta) \to (C, \gamma)$ are F-coalgebra maps such that f and g have pullback R in the underlying category, then R can be given the structure of an F-coalgebra $\rho\colon R \to F(R)$ such that the projections $\pi_1\colon R \to A$ and $\pi_2\colon R \to B$ are coalgebra maps. To see this, notice that R, with projections $\alpha \cdot F(\pi_1)$ and $\beta \cdot F(\pi_2)$, defines a cone over $F(f)$ and $F(g)$. But $F(R)$ is a weak pullback of $F(f)$ and $F(g)$; thus, we take ρ to be the comparison map $R \to F(R)$.

[3] In a category with pullbacks, the kernel pair of a map f is the pullback of f with itself.

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 [64] 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. [60], in which a domain for approximating hereditarily finite non-well-founded sets is obtained, and Moss [64], 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 $(f^\beta_\gamma : A_\beta \to A_\gamma)_{\gamma \leq \beta}$ 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 $(f^\alpha_\beta : A_\alpha \to A_\beta)_{\beta<\alpha}$ 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^α_λ by giving the composites $f^\lambda_\gamma \cdot f^\alpha_\lambda$, for $\gamma < \lambda$, and proving that they form a cone over the co-chain $(A_\gamma)_{\gamma<\lambda}$. We define $f^\lambda_\gamma \cdot f^\alpha_\lambda = f^{\gamma+1}_\gamma \cdot T(f^\lambda_\gamma)$. This does indeed define a cone over the A_γ since

$$f^\gamma_\delta \cdot f^{\gamma+1}_\gamma \cdot T(f^\lambda_\gamma) = f^{\gamma+1}_\delta \cdot T(f^\lambda_\gamma) = f^{\delta+1}_\delta \cdot f^{\gamma+1}_{\delta+1} \cdot T(f^\lambda_\gamma)$$

$$= f^{\delta+1}_\delta \cdot T(f^\gamma_\delta) \cdot T(f^\lambda_\gamma) = f^{\delta+1}_\delta \cdot T(f^\lambda_\delta).$$

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

$$f^{\beta+1}_{\gamma+1} = f^\alpha_{\gamma+1} = f^\beta_{\gamma+1} \cdot f^\alpha_\beta = f^{\gamma+2}_{\gamma+1} \cdot T(f^\beta_{\gamma+1}) = T(f^{\gamma+1}_\gamma \cdot f^\beta_{\gamma+1}) = T(f^\beta_\gamma).$$

(2b) Case: $\alpha = \beta + 2$ for some β. Set $A_\alpha = T(A_{\beta+1})$, $f^\alpha_\alpha = id$, and $f^\alpha_{\beta+1} = T(f^{\beta+1}_\beta)$. For $\gamma < \beta + 1$, we set $f^\alpha_\gamma = f^{\beta+1}_\gamma \cdot f^\alpha_{\beta+1}$. 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 $f^{\kappa+1}_\kappa$ is an isomorphism. Then $(A_\kappa, (f^{\kappa+1}_\kappa)^{-1})$ is a final T-coalgebra. □

Proposition 3.2.2 Let $\beta < \alpha$ be such that f^α_β is an isomorphism. Then $f^{\beta+1}_\beta$ 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. $\square$

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. $\square$

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. $\square$

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 [82] 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. $\square$

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. ☐

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 $(f_j \stackrel{a_j}{\to} f)_{j\in J}$. 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 $g = h$. ☐

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

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. [83].

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. [81], 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 $\beta < \lambda$. $\square$

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). $\square$

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 [83].

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. $\square$

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. $\square$

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 $\{R_\varepsilon\}_{0\le\varepsilon\le1}$ on X given by

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

(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 $\varepsilon_1 \le \varepsilon_2$. Conversely, given a family of equivalence relations $\{R_\varepsilon\}_{0\le\varepsilon\le1}$ 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 [83], 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 [83] 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

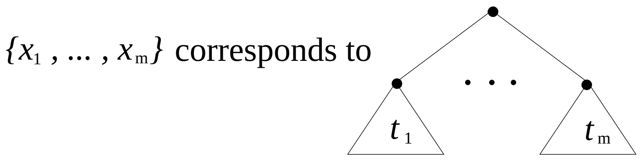

Labelling the leaf nodes of depth n in a tree in A_n by $*$ communicates the intuition that these trees are to be thought of as partial. We can think of the projection map f^{n+1}_n as taking a tree in A_{n+1}, cutting off the depth $n+1$ nodes, taking the extensional quotient, and finally relabelling 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.

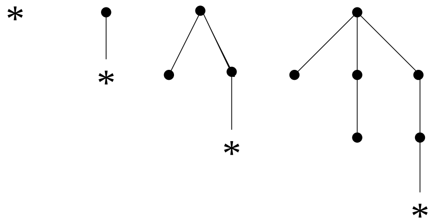

Extending the given sequence in the obvious way we can imagine the trees as projections of the following countably-branching infinite-depth tree, living in the limit A_ω.

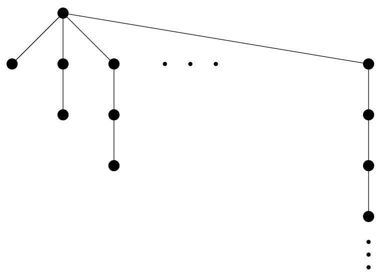

The above picture is due to Turi and Rutten [82], from their explanation of why the final coalgebra of $\wp$ is not given by the ω^{op}-limit $1 \leftarrow \wp 1 \leftarrow ... \leftarrow \wp^\omega 1$. Notice that since the trees in A_ω are potentially infinitely branching, the obvious candidate for 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.

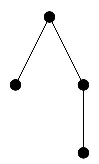

We define a sequence of trees $(t_n)_{n<\omega}$ by defining t_{n+1} to be the tree obtained from t_n by replacing each leaf of t_n with a copy of t_1. The sequence $(t_n)_{n<\omega}$ together with the inclusions $t_n \hookrightarrow t_{n+1}$ forms an ω-chain in Graph. The colimit of this chain is an 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 $\alpha \leq \omega$.

Chapter 4

Coinduction for recursively defined data types

4.1 Introduction

Coinduction is a principle for reasoning about potentially infinite or circular elements of recursive data types, like streams, processes, or formal power series; cf. [73, 28, 68]. The central notions in the category-theoretic account of coinduction are: coalgebra, homomorphism of coalgebras and bisimulation between coalgebras.

Given an endofunctor $F\colon \mathsf{Set} \to \mathsf{Set}$, an F-bisimulation between two F-coalgebras is a relation between their carriers whose projections are both coalgebra maps. This notion was introduced by Aczel and Mendler [3]. The associated coinduction principle says that *the equality relation is the maximum F-bisimulation on the final F-coalgebra.* There is also a notion of *ordered F-bisimulation* and an associated coinduction principle for locally continuous endofunctors on Cpo-categories [28, 82]. In particular, for a locally continuous endofunctor F on the category $\mathsf{Cppo}_\perp$ of pointed cpos and strict continuous maps between them, it says that *the order relation is the maximum ordered F-bisimulation on the final F-coalgebra.*

In this chapter, following a proposal of Rutten [76], we define a notion of simulation[1] for coalgebras which is based on the enriched categorical counterpart of relations, called bimodules. This only makes sense if the objects of the underlying category can themselves be viewed as (enriched) categories. Thus, this approach is to some extent narrower in

[1]The examples of 'metric bisimulations' given in [76] are instances of our notion. However we prefer to talk about simulations, as opposed to bisimulations, to emphasize that the idea we have in mind is in general asymmetric.

scope than the relational approaches described above; however, it applies to the typical categories one has in semantics, i.e., categories of preorders and metric spaces. The associated coinduction theorem, our main result, characterizes the order (respectively the metric) on a final coalgebra of an endofunctor F on the category of preorders (respectively the category of generalized metric spaces) as maximum amongst F-simulations. In fact, following Wagner [84], the main results are stated for an endofunctor F on the category Ω-Cat of all (small) Ω-enriched categories and functors between them, for some quantale Ω; cf. definition 4.3.1. Preorders and generalized metric spaces arise as instances by respectively choosing Ω to be the two element chain $\underline{2}$ and the interval $[0, \infty]$. Ω-bimodules can be seen as Ω-valued relations: a simulation assigns to each pair of processes p and q a value in Ω which can be thought of as the truth-value of the statement q simulates p. Thus, our notion of simulation is quantitative as well as qualitative.

One way to define Aczel–Mendler bisimulations for a set functor T is via a lifting of T to a graph homomorphism $\mathrm{Rel}\, T \colon \mathsf{Rel} \to \mathsf{Rel}$, where Rel is the category of sets and relations. Correspondingly, for an endofunctor $F \colon \Omega\text{-}\mathsf{Cat} \to \Omega\text{-}\mathsf{Cat}$, the notion of F-simulation introduced in the sequel is defined via a lifting of F to a graph homomorphism $\mathrm{Mod}\, F \colon \Omega\text{-}\mathsf{Mod} \to \Omega\text{-}\mathsf{Mod}$, where Ω-Mod is the category of Ω-categories and Ω-bimodules. In general, $\mathrm{Mod}\, F$ is not functorial. However we show that it is lax functorial iff F preserves strong monomorphisms. This condition is important as it ensures that F-simulations compose, and that the final coalgebra semantics is abstract with respect to F-similarity. The corresponding condition on T, ensuring that $\mathrm{Rel}\, T$ is a lax functor, is preservation of weak pullbacks. In practice it is easier to verify preservation of strong monos than preservation of weak pullbacks, and we give an example, involving discrete probabilistic transition systems, where a coalgebraic analysis using the notion of simulation introduced here is much easier than with Aczel–Mendler bisimulations.

Other contributions of this chapter involve variants of the above coinduction principle. In the first of these, which resembles Milner's 'bisimulation up-to bisimilarity', one does not need to verify the bimodular action of simulations; this is important since one wants to do as little work as possible in establishing that a simulation holds. We also show that the coinduction principle applies to invariants of typical functors on the category of complete metric spaces and non-expansive maps. This category is more suited to domain theory than the category GMet of generalized metric spaces and non-expansive maps (a straightforward instance of Ω-Cat). Applying the last two results, we give a couple of examples to illustrate the quantitative aspect to our notion of simulation. In the first of these we consider an ultrametric space 'domain for applicative bisimulation', and show that the metric on this domain corresponds to the depth of experiment required to differentiate between processes. In another example, we consider probabilistic transition systems with continuous state spaces as coalgebras of a certain functor based on the Hutchinson metric on the space of probability measures (measures with total mass 1) on a metric space. We show that the resultant notion of coalgebraic bisimilarity agrees with the metric defined on probabilistic processes by Desharnais et al. in [25].

Finally, a few remarks about the approach we take. The terminology surrounding the notion of Ω-category will probably be unfamiliar to the reader. However, Ω-categories are really very simple entities, and working at this level of generality is no harder than working with, say, generalized metric spaces, since the only relevant property of $\mathbb{R}$ for this work is that it is a commutative unital quantale. Indeed, we find that the separation of concerns enforced by making our work parametric in Ω to be a positive aid to understanding. The enriched categorical view of preorders and metric spaces was first proposed by Lawvere [53]. It was later taken up by Wagner [84] and Rutten [77] to unify the basic notions and constructions of order-theoretic and metric domain theory.

4.2 Coinduction in Set

As in Rutten [76], our approach to coinduction for Ω-categories is motivated by looking at the definition of bisimulations for set functors. For a set functor T we give a characterization of T-bisimulations in terms of a weakly commuting diagram in Rel. The counterpart of this diagram in the Ω-category case will be taken as the primitive definition of simulation.

Let Rel be the category with sets as objects and relations as arrows, with composition of relations as defined in section 2.2.2. Rel is poset-enriched with the partial order on homsets being inclusion of relations. We have an *involution* functor $(-)^\circ \colon \mathsf{Rel} \to \mathsf{Rel}^{\mathrm{op}}$ mapping each set to itself and each relation R to its inverse R°. We also have an embedding of Set into Rel which is the identity on objects and sends a function $f \colon A \to B$ to its graph $\Gamma_f \subseteq A \times B$. We will sometimes write $R(a,b)$ for $\langle a, b\rangle \in R$.

Proposition 4.2.1 [76, Prop 1.2] Given a set functor T, and T-coalgebras (X,α) and (Y,β), a relation $R \subseteq X \times Y$ is a T-bisimulation iff $\langle x, y\rangle \in R$ implies $\langle \alpha x, \beta y\rangle \in T\pi_2 \cdot (T\pi_1)^\circ$ for all $x \in X, y \in Y$. □

Following Carboni, Kelly and Wood [21] who were interested in 'lifting' endofunctors on regular categories to endofunctors on their category of relations, let us write $\mathrm{Rel}\,T$ for the graph homomorphism $\mathsf{Rel} \to \mathsf{Rel}$ with

$$(\mathrm{Rel}\,T)X = TX \quad \text{and} \quad (\mathrm{Rel}\,T)R = T\pi_2 \cdot (T\pi_1)^\circ \tag{4.1}$$

for each set X and relation R with projections π_1, π_2. The following proposition follows easily from proposition 4.2.1.

Proposition 4.2.2 Given a set functor T, and T-coalgebras (X, α) and (Y, β), a relation $R \subseteq X \times Y$ is a T-bisimulation iff the square

$$\begin{tikzcd} X \arrow[r, "\alpha"] \arrow[d, "R"'] \arrow[dr, phantom, "\subseteq"] & TX \arrow[d, "(\mathrm{Rel}\,T)R"] \\ Y \arrow[r, "\beta"'] & TY \end{tikzcd} \tag{4.2}$$

commutes up-to inclusion in Rel. □

In this case we could say that R is a lax morphism of coalgebras, though it is not clear that such 'morphisms' compose unless $\mathrm{Rel}\,T$ is actually a lax functor. It follows from [21] that $\mathrm{Rel}\,T$ is a lax functor iff it is a functor iff T preserves weak pullbacks.

Finally, recall that the coinduction principle says that $R \subseteq \Delta_X$ for any T-bisimulation R on a final T-coalgebra (X, α).

4.3 Ω-categories

In this section we recall the notions of Ω-category, -functor and -bimodule. These are exactly as in enriched (or $\mathcal{V}$-)category theory [48], but specialized to the case in which $\mathcal{V}$ is a quantale (this makes things dramatically simpler). Choosing Ω appropriately, we recover preorders and monotone maps; another choice of Ω yields generalized metric spaces and non-expansive maps, cf. [53, 84, 77].

Definition 4.3.1 A *quantale* is a complete lattice with a binary, associative operation $\otimes$, such that for every element a, both $a \otimes -$ and $- \otimes a$ have right adjoints (equivalently $\otimes$ preserves all joins). A quantale is called commutative whenever its tensor is, and it is called unital if there is an element 1, the unit, such that $1 \otimes a = a = a \otimes 1$ for all a. For commutative quantales we denote the right adjoint to $a \otimes -$ by $a \multimap -$. □

In the sequel we will assume a commutative unital quantale $(\Omega, \otimes, 1)$.

Definition 4.3.2 (i) **Ω-category**. An Ω-category (X_0, X) consists of a set X_0, and a mapping $X \colon X_0 \times X_0 \to \Omega$, called the hom-functor, such that:

$1 \leq X(x, x)$ for all $x \in X_0$ (reflexivity).

$X(x, y) \otimes X(y, z) \leq X(x, z)$ for all $x, y, z \in X_0$ (transitivity).

For short, we usually refer to the Ω-category X rather than (X_0, X).

(ii)Ω-**functor**. Given two Ω-categories X and Y, an Ω-functor $f\colon X \to Y$ is a function $f_0\colon X_0 \to Y_0$, between the underlying sets, such that $X(x, x') \leq Y(fx, fx')$ for all $x, x' \in X_0$. We have a preorder on the set of Ω-functors $X \to Y$, with $f \leq g$ if $1 \leq Y(fx, gx)$ for all $x \in X_0$. □

Example 4.3.3 Let $\Omega = \underline{2}$, the two point chain, with $\otimes = \wedge$. Writing $\underline{2} = \{\bot, \top\}$, where $\bot \leq \top$, a $\underline{2}$-category X (not to be confused with a 2 category in the sense of a Cat enriched category) is a preorder, where $x \leq y$ iff $X(x, y) = \top$ for $x, y \in X_0$. $\underline{2}$-functors are monotone maps, and $\underline{2}$-categories and -functors form an ordinary category Preord. The reflexivity and transitivity laws for $\underline{2}$-categories amount to the laws with the same names for preorders.□

Example 4.3.4 Let $\Omega = [0, 1]$, with the opposite to the usual ordering of reals, and max as $\otimes$. We find that $x \multimap y = y$ if $x < y$, and 0 otherwise. A $[0, 1]$-category is a generalized ultrametric space (gus), and a $[0, 1]$-functor is a non-expansive map; these organize themselves in an ordinary category GUMet. The Ω-category laws amount to requiring that the metric on a space X is reflexive: $X(x, x) = 0$, and satisfies the strong triangle inequality: $X(x, z) \leq \max\{X(x, y), X(y, z)\}$. Thus all ultrametric spaces are generalized ultrametric spaces, but in the latter the metric need not be symmetric and different elements can have distance 0. □

Example 4.3.5 Let $\Omega = [0, \infty]$, with the opposite to the usual ordering of reals, and $+$ as $\otimes$ (where $\infty + r = r + \infty = \infty$). Then $r \multimap s$ is given by $s \dot{-} r$, where $\dot{-}$ is truncated subtraction. A $[0, \infty]$-category is a generalized metric spaces (gms)—just as above, but replacing max with $+$ in the triangle inequality. A $[0, \infty]$-functor is a non-expansive map; generalized metric spaces and non-expansive maps form a category GMet. □

The fact that the hom-functor in a gms need not be symmetric runs contrary to the distance intuition for ordinary metric spaces. Rather, we think of it as measuring the *degree of approximation* between two elements. Thus gms's combine the qualitative notion of elements approximating each other, from partial orders, with the quantitative notion of elements being close to each other, from metric spaces. To avoid confusion we never use the term 'metric space' as an abbreviation for 'generalized metric space'.

Example 4.3.6 Ω can itself be made into an Ω-category by taking $\multimap$ as hom-functor. The reflexivity and transitivity laws follow from the $\multimap$-$\otimes$ adjunction. Note that if $f, g\colon X \to \Omega$ are Ω-functors, then $f \leq g$ iff as $fx \leq gx$ for all x, cf. definition 4.3.2 (ii). □

We saw that each preorder can be thought of as an Ω-category. Conversely, for each Ω-category X there is an underlying preorder $(X_0, \leq_X)$, where $x \leq_X y$ iff $1 \leq X(x, y)$.

Definition 4.3.7 The category of Ω-categories and Ω-functors is denoted Ω-Cat. Definition 4.3.2 (ii) gives a preorder on each homset of Ω-Cat making it a preorder enriched (or Preord-)category. □

Some constructions of Ω-categories, cf. [48]:

Definition 4.3.8 (i) **Tensor**. Given two Ω-categories X and Y, their tensor product $X \otimes Y$ has $X_0 \times Y_0$ as its underlying set, with hom-functor defined by

$$(X \otimes Y)((x, y), (x', y')) = X(x, x') \otimes Y(y, y').$$

(ii) **Dual**. The Ω-category X^{op} has the same underlying set as X, and hom-functor $X^{\mathrm{op}}(x, x') = X(x', x)$.
(iii) **Hom**. The set of Ω-functors $X \to Y$ can be given the structure of an Ω-category, denoted $(X \to Y)$, by defining

$$(X \to Y)(f, g) = \bigwedge_{x \in X_0} Y(fx, gx).$$

It is straightforward that Ω-Cat forms a symmetric monoidal closed category. □

The tensor here agrees with the tensor product of metric spaces from section 2.3.3; for *ultra*metric spaces and preorders it is the Cartesian product. In the case of preorders (iii) gives the pointwise preorder on monotone functions, and for generalized metric spaces it gives the supremum metric on non-expansive functions.

We recall from [79] the construction of the quotient of an Ω-category. Let X be an Ω-category, Y a set and $q \colon X \to Y$ a surjective map. We construct the minimal hom on Y such that q is an Ω-functor. Given $y, y' \in Y$, define a q-path from y to y' to be a finite sequence $x_0, ..., x_n$ of points of X such that $q(x_0) = y$ and $q(x_n) = y'$. Define the *cost* of the q-path $\pi = x_0, ..., x_n$ to be the tensor product

$$c(\pi) = \bigotimes \{X(x_i, x_{i+1}) \mid q(x_i) \neq q(x_{i+1})\}.$$

For example, for $\Omega = [0, \infty]$ this gives the total cost of travelling from x_0 to x_n along the path π, where travelling between points with the same q-image is for free. Now we can define

$$Y(y, y') = \bigvee \{c(\pi) \mid \pi \text{ is a } q\text{-path from } y \text{ to } y'\}.$$

4.3.1 Bimodules

Definition 4.3.9 Given Ω-categories X and Y, an *Ω-bimodule (or profunctor, or distributor)* $R\colon Y \nrightarrow X$ is an Ω-functor $R\colon X^{\mathrm{op}} \otimes Y \to \Omega$. [2] $\square$

An Ω-bimodule $R\colon Y \nrightarrow X$ can be thought of as an Ω-valued relation from X to Y, with $R(x,y)$ giving the truth value of the relatedness of x to y. In case $\Omega = [0,\infty]$, an Ω-bimodule is sometimes called a fuzzy relation; if $\Omega = \underline{2}$ then an Ω-bimodule is just a relation $R \subseteq X \times Y$ such that $x' \leq x$, $R(x,y)$ and $y \leq y'$ together imply $R(x', y')$.

Analogous to composition of relations we have composition of bimodules, given by the following 'matrix-product' or 'least-cost' formula.

Definition 4.3.10 Suppose $R\colon Y \nrightarrow X$ and $S\colon Z \nrightarrow Y$ are Ω-bimodules, then the formula

$$(R \cdot S)(x,z) \stackrel{\mathrm{def}}{=} \bigvee_{y \in Y_0} R(x,y) \otimes S(y,z) \tag{4.3}$$

defines a bimodule $Z \nrightarrow X$—the composite of R with S. $\square$

It is straightforward to verify that composition is associative and that the identity bimodule on X is the hom-functor $X(-,-)\colon X^{\mathrm{op}} \otimes X \to \Omega$. Thus we may make the following definition.

Definition 4.3.11 Ω-Mod is the Preord-category of Ω-categories and Ω-bimodules. For each pair of Ω-categories X and Y, the ordering on Ω-$\mathsf{Mod}(Y,X)$ is the ordering of Ω-functors $X^{\mathrm{op}} \otimes Y \to \Omega$. $\square$

Since Ω-Cat and Ω-Mod are Preord-categories, we can define adjointness for arrows in each. In both cases an arrow $f\colon X \to Y$ is left adjoint to $g\colon Y \to X$ if $f \cdot g \leq 1_Y$ and $1_X \leq g \cdot f$. As we now explain, one of the motivations behind the introduction of Ω-Mod is to give an adjoint to every Ω-functor.

Definition 4.3.12 An Ω-functor $f\colon Y \to X$ gives rise to a bimodule $f_*\colon Y \nrightarrow X$ defined by $f_*(x,y) = X(x, fy)$, and a bimodule $f^*\colon X \nrightarrow Y$ defined by $f^*(y,x) = X(fy, x)$. It is easily seen that f_* is left adjoint to f^*. We have an embedding Ω-$\mathsf{Cat} \to \Omega$-Mod, which is locally a preorder embedding, taking each Ω-category to itself, and each Ω-functor f to f_*. Similarly $(-)^*$ yields an embedding from the horizontal, vertical opposite of Ω-Cat to Ω-Mod. We will usually omit the $*$ subscript when regarding functors as bimodules. $\square$

[2] The convention that R has source Y and target X, rather than vice-versa, is unfortunately inconsistent with the convention for relations. However, this approach is standard and avoids inelegant contravariant behaviour in the subsequent theory.

Remark 4.3.13 The functoriality of $R\colon Y \nrightarrow X$, viz.

$$X(x', x) \otimes Y(y, y') \leq R(x, y) \multimap R(x', y'),$$

is equivalent to the following inequalities:

$$X(x', x) \otimes R(x, y) \leq R(x', y) \quad \text{and} \quad R(x, y) \otimes Y(y, y') \leq R(x, y').$$

From these it follows that for Ω-functors $W \xrightarrow{f} X$ and $Z \xrightarrow{g} Y$, and bimodule $R\colon Y \nrightarrow X$, we have $(f^* \cdot R \cdot g)(w, z) = R(fw, gz)$. □

4.3.2 Topology

We conclude this section with a very brief discussion of the topologies one might put on a gms X. First, observe that the usual notion of a metric topology, i.e. that generated by the base of ε-balls $B(x, \varepsilon) = \{y \mid X(x, y) < \varepsilon\}$, still makes sense for gms's. Smyth [79] calls this the standard topology. Note that non-expansive maps of gms's are continuous in the standard topology. Later, we will consider endofunctors on GMet which make use of topological notions, e.g. the compact powerdomain functor and the functor of Borel probability measures: these are all with regard to the standard topology. However, we should mention that this is far from the only topology one could put on a gms. In fact, Reilly et al. [71] identify seven distinct notions of Cauchy sequence and convergence for a gms. Rutten [75] gives an enriched-categorical treatment of convergence for so-called *forward Cauchy* sequences in terms of *weighted colimits*.

Example 4.3.14 [79, Ex. 4] Consider the gms Σ^∞ of finite and infinite strings from an alphabet Σ with the generalized metric

$$\Sigma^\infty(x, y) = \inf\{2^{-n} \mid x[n] \leq y[n]\}$$

where $x[n]$ denotes the truncation of x to the first n terms. The metric topology on Σ^∞ is the *Scott topology*, i.e. that generated by the base of opens $\uparrow \sigma$, for σ a finite string. In fact, the Scott topology on any ω-algebraic cpo can be seen as the metric topology with respect to some generalized ultrametric [79]. □

4.4 Lifting endofunctors on Ω-Cat to Ω-Mod

Throughout this section let F denote a locally monotone endofunctor $\Omega\text{-}\mathsf{Cat} \to \Omega\text{-}\mathsf{Cat}$, i.e., F preserves the preorder on each homset. This general case includes as specific instances:

locally monotone endofunctors of Preord, and locally non-expansive endofunctors of GMet. We define a lifting of F to a graph homomorphism $\Omega\text{-}\mathsf{Mod} \to \Omega\text{-}\mathsf{Mod}$.

Observe that every morphism R in Rel can be *tabulated* by a span of functions $X \xleftarrow{\pi_1} R \xrightarrow{\pi_2} Y$ in the sense that $R = \pi_2 \cdot \pi_1^\circ$ and $\pi_1^\circ \cdot \pi_1 \cap \pi_2^\circ \cdot \pi_2 = 1$—just take π_1 and π_2 to be the projections associated to R. Correspondingly, we recall from [19] how Ω-bimodules can be tabulated by a cospan of Ω-functors.

Definition 4.4.1 [19] Suppose $R\colon Y \nrightarrow X$ is a bimodule. The *collage* of R, denoted $\mathrm{Coll}(R)$, is the Ω-category with the disjoint union of Y_0 and X_0 as its underlying set, and such that for all $y, y' \in Y_0$ and $x, x' \in X_0$,

- $\mathrm{Coll}(R)(y, y') = Y(y, y')$;
- $\mathrm{Coll}(R)(x, x') = X(x, x')$;
- $\mathrm{Coll}(R)(x, y) = R(x, y)$;
- $\mathrm{Coll}(R)(y, x) = \bot_\Omega$.

The satisfaction of the reflexivity and transitivity laws in $\mathrm{Coll}(R)$ follows from the corresponding laws for X and Y, and remark 4.3.13. Note that the coproduct injections in Set give us a cospan of Ω-functors $Y \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} X$, and we have $R = r_2^* \cdot r_1$. □

Definition 4.4.2 Define a graph homomorphism $\mathrm{Mod}\, F\colon \Omega\text{-}\mathsf{Mod} \to \Omega\text{-}\mathsf{Mod}$ by

$$(\mathrm{Mod}\, F)X = FX \quad \text{and} \quad (\mathrm{Mod}\, F)R = (Fr_2)^* \cdot Fr_1 \colon FY \nrightarrow FX$$

for each Ω-category X and each bimodule $R\colon Y \nrightarrow X$ with collage $Y \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} X$. It is helpful to recall that, by remark 4.3.13,

$$(\mathrm{Mod}\, F)(R)(x, y) = F(\mathrm{Coll}(R))((Fr_2)x, (Fr_1)y)$$

for $x \in FX, y \in FY$. To be concise we will henceforth denote the lifting of a functor by the same name, but in calligraphic font. □

The following proposition provides a sanity check for this definition. It says that $\mathcal{F}\colon \Omega\text{-}\mathsf{Mod} \to \Omega\text{-}\mathsf{Mod}$ really is an *lifting* of F.

Proposition 4.4.3 Suppose that $f\colon X \to Y$ is an Ω-functor; then $\mathcal{F}f = Ff$ and $\mathcal{F}(f^*) = (Ff)^*$.

Proof. We prove the second equality, the proof of the first being similar. If $Y \xrightarrow{r_1} \mathrm{Coll}(f^*) \xleftarrow{r_2} X$ is the collage of f^*, let $h\colon \mathrm{Coll}(f^*) \to Y$ be the Ω-functor whose underlying function is given by the cotuple $[1, f]$. We verify that h is left adjoint to r_1 in Ω-Cat. Clearly $h \cdot r_1 = 1$; also $1 \le r_1 \cdot h$ since, for $x \in X_0, y \in Y_0$

$$1 \le Y(fx, fx) = \mathrm{Coll}(f^*)(x, fx) = \mathrm{Coll}(f^*)(x, r_1 hx) \text{ and} \tag{4.4}$$

$$1 \le Y(y, y) = \mathrm{Coll}(f^*)(y, y) = \mathrm{Coll}(f^*)(y, r_1 hy). \tag{4.5}$$

It follows that Fh is left adjoint to Fr_1 in Ω-Cat, and hence also in Ω-Mod; thus,

$$\mathcal{F}(f^*) = (Fr_2)^* \cdot Fr_1 = (Fr_2)^* \cdot (Fh)^* = (Fh \cdot Fr_2)^* = (Ff)^*.$$

$\square$

Example 4.4.4 Here are some examples of liftings of an endofunctor $G\colon \mathsf{GUMet} \to \mathsf{GUMet}$. In the following $R\colon Y \nrightarrow X$ is a bimodule with collage $Y \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} X$.

(i) Let $G(X) = A \otimes X$ for a fixed space A.

$$\begin{aligned} \mathcal{G}R(\langle a, x\rangle, \langle b, y\rangle) &= G(\mathrm{Coll}(R))((Gr_2)\langle a, x\rangle, (Gr_1)\langle b, y\rangle) \\ &= \max\{A(a, b), R(x, y)\}. \end{aligned}$$

In particular, if A has the discrete metric,

$$\mathcal{G}R(\langle a, x\rangle, \langle b, y\rangle) = \begin{cases} R(x, y) & \text{if } a = b \\ 1 & \text{otherwise.} \end{cases}$$

(ii) Let G be given by $G(X) = (A \to X)$, for fixed A. Then, given non-expansive functions $f\colon A \to X$ and $g\colon A \to Y$,

$$\begin{aligned} \mathcal{G}R(f, g) &= G(\mathrm{Coll}(R))((Gr_2)f, (Gr_1)g) \\ &= \sup\{\mathrm{Coll}(R)(fa, ga) \mid a \in A_0\} \\ &= \sup\{R(fa, ga) \mid a \in A_0\}. \end{aligned}$$

(iii) Let G be given by $G(X) = A + X$ for fixed A, where $+$ is coproduct in GMet (this is calculated by taking the coproduct of the underlying sets, giving elements in different summands a distance of 1, and with the obvious distance for elements in the same summand). Then, for $x \in (A + X)_0$ and $y \in (A + Y)_0$,

$$\begin{aligned} \mathcal{G}R(x, y) &= G(\mathrm{Coll}(R))((Gr_2)x, (Gr_1)y) \\ &= \begin{cases} R(x, y) & \text{if } x \in X_0 \text{ and } y \in Y_0 \\ A(x, y) & \text{if } x, y \in A_0 \\ 1 & \text{otherwise.} \end{cases} \end{aligned}$$

(iv) Suppose G is the compact powerdomain functor on generalized metric spaces, assigning to a space X the space of all compact subsets of X in the standard metric topology, equipped with the Hausdorff distance:

$$G(X)(V,W) = \max\{\sup_{v\in V}\inf_{w\in W} X(v,w), \sup_{w\in W}\inf_{v\in V} X(v,w)\}.$$

Then we have that for compact $V \subseteq X$ and $W \subseteq Y$,

$$\begin{aligned} \mathcal{G}R(V,W) &= G(\mathrm{Coll}(R))((Gr_2)V,(Gr_1)W) \\ &= \max\{\sup_{v\in V}\inf_{w\in W} R(v,w), \sup_{w\in W}\inf_{v\in V} R(v,w)\}. \end{aligned}$$

□

A locally monotone graph homomorphism between Preord-categories $G\colon \mathcal{C} \to \mathcal{D}$ is called a *lax functor* if $1_{GC} \leq G1_C$ for each object $C \in \mathcal{C}$, and $G\phi \cdot G\psi \leq G(\phi\cdot\psi)$ for each pair of composable arrows ϕ and ψ. Next we show a condition on F which will ensure that $\mathcal{F}$ is a lax functor. Recall the definition of strong monos in a category from section 3.2.2. It is easy to see that the strong monos in Ω-Cat are precisely those maps $m\colon X \to Y$ such that $X(x,x') = Y(mx,mx')$ for all $x,x' \in X_0$; we call these maps *embeddings*.[3] Note that all the functors in example 4.4.4 preserve embeddings.

Theorem 4.4.5 $\mathcal{F}$ is a lax functor iff F preserves embeddings.

Proof. ($\Leftarrow$) Consider the commuting diagram

$$\begin{tikzcd}
 & & P & & \\
 & \mathrm{Coll}(S) \arrow[ur, "i_1"] & \mathrm{Coll}(R\cdot S) \arrow[u, "m"] & \mathrm{Coll}(R) \arrow[ul, "i_2"'] & \\
Z \arrow[ur, "s_1"] \arrow[urr, "f_1"'] & & Y \arrow[ul, "s_2"'] \arrow[ur, "r_1"] & & X \arrow[ul, "r_2"'] \arrow[ull, "f_2"]
\end{tikzcd} \tag{4.6}$$

showing the collages of the bimodules R, S and $R\cdot S$, and with P the pushout in Ω-Cat of s_2 and r_1. P has underlying set given by the pushout of s_2 and r_1 in Set; the Ω-category structure is the quotient structure determined by the cotuple $[i_1,i_2]$. Observe that the evident inclusion m is an embedding.

[3]In the present context the class of strong monos coincides with the class of extremal monos and the class of regular monos.

Applying F to the diagram we have the following chain of inequalities, where the second step holds because $Fr_1 \cdot (Fs_2)^* \leq (Fi_2)^* \cdot Fi_1$ follows from $Fi_1 \cdot Fs_2 = Fi_2 \cdot Fr_1$ by a simple manipulation of adjoints.

$$\begin{aligned}
\mathcal{F}(R) \cdot \mathcal{F}(S) &= (Fr_2)^* \cdot Fr_1 \cdot (Fs_2)^* \cdot Fs_1 \\
&\leq (Fr_2)^* \cdot (Fi_2)^* \cdot Fi_1 \cdot Fs_1 \\
&= (Ff_2)^* \cdot (Fm)^* \cdot Fm \cdot Ff_1 \quad ((-)^* \text{ is contravariantly functorial}) \\
&= (Ff_2)^* \cdot Ff_1 \quad (Fm \text{ is an embedding}) \\
&= \mathcal{F}(R \cdot S).
\end{aligned}$$

($\Rightarrow$) If m is an embedding then $Fm^* \cdot Fm \leq \mathcal{F}(m^* \cdot m) = F(1) = 1$, so Fm is an embedding. □

It is worth comparing theorem 4.4.5 with a result of Carboni, Kelly and Wood [21]. They use formula (4.1) to define a lifting of a functor $T : \mathcal{C} \to \mathcal{D}$ between regular categories to a graph homomorphism $\mathrm{Rel}\, T : \mathrm{Rel}\, \mathcal{C} \to \mathrm{Rel}\, \mathcal{D}$ between their respective categories of relations. This lifting is based on the decomposition of arrows of $\mathrm{Rel}\, \mathcal{C}$ into spans of *maps* (a *map* in Preord-category is an arrow with a right adjoint). On the other hand, the lifting considered in this section is based on a decomposition of an arrow of Ω-Mod into a *cospan* of maps. Correspondingly, theorem 4.4.5 looks like a horizontal, vertical dual of a result of [21], stating that $\mathrm{Rel}\, T$ is a colax functor iff T preserves strong epis. Another result of [21] is that $\mathrm{Rel}\, T$ is a lax functor iff T *nearly preserves pullbacks*, where the latter means that the comparison map in the diagram for pullback preservation is strong epic. As Rutten [76] has already pointed out, this explains the importance of the 'weak pullback' preservation condition for set functors in the theory of coalgebras.

4.5 F-simulations and coinduction

Again, let F denote a Preord-functor Ω-Cat $\to$ Ω-Cat. We are now ready to define the concept of an F-simulation, but first we consider two special cases of this notion.

Definition 4.5.1 For F-coalgebras (X, α) and (Y, β), we say that $f : Y \to X$ is a *lax F-coalgebra morphism* $(Y, \beta) \to (X, \alpha)$ iff the following diagram commutes up-to inequality in Ω-Cat. If the opposite inequality holds then we say that f is a *colax morphism of F-coalgebras.* □

$$\begin{tikzcd} Y \arrow[r, "\beta"] \arrow[d, "f"'] \arrow[dr, phantom, "\leq"] & FY \arrow[d, "Ff"] \\ X \arrow[r, "\alpha"] & FX \end{tikzcd} \tag{4.7}$$

Definition 4.5.2 Given F-coalgebras (X, α) and (Y, β) we say that the bimodule $R\colon Y \nrightarrow X$ is an *F-simulation* iff the following diagram commutes up-to inequality in Ω-Mod. □

$$\begin{tikzcd} Y \arrow[r, "\beta"] \arrow[d, "R"'] \arrow[dr, phantom, "\leq"] & FY \arrow[d, "\mathcal{F}R"] \\ X \arrow[r, "\alpha"'] & FX \end{tikzcd} \tag{4.8}$$

From the adjunction $\alpha \dashv \alpha^*$ it is easily seen that the inequality (4.8) is equivalent to $R \leq \alpha^* \cdot \mathcal{F}R \cdot \beta$, that is,

$$R(x, y) \leq \mathcal{F}R(\alpha x, \beta y) \quad \text{for all } x \in X_0, y \in Y_0. \tag{4.9}$$

Recalling that $\mathcal{F}R(\alpha x, \beta y) = F(\mathrm{Coll}(R))((Fr_2)\alpha x, (Fr_1)\beta y)$, yet another way to characterize F-simulations is as those R such that the unique function making the diagram below commute in Set is actually an Ω-functor.

$$\begin{tikzcd} Y \arrow[r, "r_1"] \arrow[d, "\beta"'] & \mathrm{Coll}(R) \arrow[d, dotted] & X \arrow[l, "r_2"'] \arrow[d, "\alpha"] \\ FY \arrow[r, "Fr_1"] & F(\mathrm{Coll}(R)) & FX \arrow[l, "Fr_2"'] \end{tikzcd} \tag{4.10}$$

4.5.1 Some universal coalgebra

In his 'universal coalgebra' Rutten [73] shows the following two facts about T-bisimulations for a set functor T: a function f is a map of T-coalgebras iff the graph of f is a T-bisimulation; the union of a family of T-bisimulations is a T-bisimulation. Under the additional hypothesis that T preserves weak pullbacks he also shows that: T-bisimulations are closed under composition, the pullback of two T-coalgebra maps is a bisimulation, and the image of a tuple of T-coalgebra maps with common domain is a bisimulation. Below, we prove corresponding facts for our notion of F-simulation. In this correspondence the role of the condition 'T preserves weak pullbacks' is taken by 'F preserves embeddings'.

Proposition 4.5.3 Given F-coalgebras (X, α) and (Y, β), a map $f\colon Y \to X$ is a lax coalgebra morphism iff f_* is an F-simulation, and f is a colax coalgebra morphism iff f^* is an F-simulation. Thus f is a coalgebra map iff both f^* and f_* are simulations. In particular, the identity bimodule $X(-,-)\colon X \nrightarrow X$ is an F-simulation on (X, α).

Proof. This follows from proposition 4.4.3 and the fact that functors $(-)_*$ and $(-)^*$ of definition 4.3.12 are locally preorder embeddings. $\square$

Proposition 4.5.4 The join (computed pointwise) of a family of F-simulations $\{R_k\}_k$ between F-coalgebras (X, α) and (Y, β) is again an F-simulation. The join of all F-simulations on a coalgebra is called F-similarity.

Proof. The result follows by the following chain of inequalities.

$$\begin{aligned} \bigvee_k R_k &\leq \bigvee_k (\alpha^* \cdot \mathcal{F} R_k \cdot \beta) \quad (R_k \text{ an } F\text{-simulation}) \\ &\leq \alpha^* \cdot (\bigvee_k \mathcal{F} R_k) \cdot \beta \\ &\leq \alpha^* \cdot \mathcal{F}(\bigvee_k R_k) \cdot \beta \quad (\mathcal{F} \text{ is locally monotone}). \end{aligned}$$

Proposition 4.5.5 Suppose that F preserves embeddings, and the bimodules R and S, shown below, are F-simulations; then the composite $R \cdot S$ is also an F-simulation.

Proof. By pasting, the outer rectangle in (4.11) commutes up-to inequality; then the result follows since $\mathcal{F}$ is a lax functor. $\square$

$$\begin{tikzcd} Z \arrow[r, "\gamma"] \arrow[d, "S"'] \arrow[dr, phantom, "\leq"] & FZ \arrow[d, "\mathcal{F}S"] \\ Y \arrow[r, "\beta"] \arrow[d, "R"'] \arrow[dr, phantom, "\leq"] & FY \arrow[d, "\mathcal{F}R"] \\ X \arrow[r, "\alpha"'] & FX \end{tikzcd} \tag{4.11}$$

In a regular category, the pullback of maps f and g with common target is the relation $g^\circ \cdot f$, and the image of a tuple $\langle h, k \rangle$, where h and k have common domain, is the relation $k \cdot h^\circ$. In the light of the aforementioned universal coalgebra this motivates:

Proposition 4.5.6 Suppose F preserves embeddings, and f, g are F-coalgebra morphisms $(Y, \beta) \to (X, \alpha)$, then $g^* \cdot f$ is an F-simulation on (Y, β) and $g \cdot f^*$ is an F-simulation on (X, α).

Proof. By propositions 4.5.3 and 4.5.5. $\square$

Remark 4.5.7 For functors F that preserve embeddings, proposition 4.5.5 also gives us a characterization of F-simulations which doesn't mention the notion of a collage or the lifting Mod F. Given F-coalgebras (X, α) and (Y, β), $R\colon Y \nrightarrow X$ is an F-simulation iff there is a cospan of coalgebra maps

$$Y \xrightarrow{f} Z \xleftarrow{g} X$$

such that $R = g^* \cdot f$. This observation is also of interest because, in categorical approaches, bisimulations are typically defined by spans: spans of coalgebra maps in [3], and spans of so-called open maps in [46]. □

4.5.2 Coinduction

Next, we prove a coinduction result for final F-coalgebras. This specializes to a coinduction principle for final coalgebras of endofunctors on Preord, GUMet and GMet; it also, pretty obviously, works for variants like Pos. Note that the proof of theorem 4.5.8 uses both the existence and uniqueness of homomorphisms to the final coalgebra. In contrast, the proof of coinduction for Aczel–Mendler style bisimulations uses only uniqueness, and Fiore's coinduction result for final coalgebras of Cpo-endofunctors of Cpo-categories [28] utilises the fact that such final coalgebras are *maximally final* (so does not yield a coinduction principle for endofunctors on Pos, cf. section 4.9).

Theorem 4.5.8 If (X, α) is a final F-coalgebra, and $x, y \in X_0$, then

$$X(x, y) = \bigvee \{R(x, y) \mid R \text{ is an } F\text{-simulation on } (X, \alpha)\}.$$

Proof. Suppose $R\colon X \nrightarrow X$ is an F-simulation on (X, α).

$$\begin{tikzcd} X \arrow[r, "r_1"] \arrow[d, "\alpha"'] & \mathrm{Coll}(R) \arrow[d, dotted] & X \arrow[l, "r_2"'] \arrow[d, "\alpha"] \\ FX \arrow[r, "Fr_1"] & F(\mathrm{Coll}(R)) & FX \arrow[l, "Fr_2"'] \end{tikzcd} \tag{4.12}$$

In diagram (4.12), r_1 and r_2 have a common left inverse by the finality of (X, α). The function underlying this map must be the cotuple $[id_X, id_X]$, but the Ω-functoriality of this map says that $R(x, y) \le X(x, y)$ for all $x, y \in X_0$. Proposition 4.5.3 tells us that the identity bimodule on X is an F-simulation, thus the result holds. □

The fact that simulations have to be Ω-bimodules would seem to add an onerous burden in employing the above coinduction theorem. We go on to show how this problem can be circumvented by a technique resembling Milner's use of 'bisimulation up-to bisimilarity'[4]; this is demonstrated by an example in section 4.6.1.

[4]Recall that in CCS, a relation R between processes is called a strong bisimulation up-to $\sim$, where $\sim$ denotes bisimilarity, if $\sim R \sim$ is a strong bisimulation; cf. [61].

Proposition 4.5.9 Given an Ω-category X and a function $f\colon X_0 \to \Omega$, then

$$\hat{f}(x) \stackrel{\text{def}}{=} \bigvee\{X(x', x) \otimes f(x') \mid x' \in X_0\}$$

defines an Ω-functor $X \to \Omega$.

Proof. By the $\multimap$-$\otimes$ adjunction for Ω it will suffice to show that $X(x, y) \otimes \hat{f}(x) \leq \hat{f}(y)$ for all $x, y \in X_0$; this is straightforward once it is recalled that $a \otimes -$, being a left adjoint, preserves joins. $\square$

Given a set X_0, let us call a function $R\colon X_0 \times X_0 \to \Omega$ an Ω-relation on X_0 (so that Ω-bimodules are the functorial Ω-relations). Observe that (4.3) serves to define a composition of Ω-relations, and that for an Ω-relation R on an Ω-category X

$$\hat{R} = X(-,-) \cdot R \cdot X(-,-)$$

is an Ω-bimodule by the previous proposition.

Let (X, α) be a final F-coalgebra. We call an Ω-relation R on X_0 an *F-simulation up-to similarity* if

$$R(x, y) \leq \mathcal{F}\hat{R}(\alpha x, \alpha y) \quad \text{for all } x, y \in X_0.$$

Observe that every F-simulation is an F-simulation up-to similarity, and, if R is an F-simulation up-to similarity, then $\hat{R}$ is an F-simulation since $\hat{R}$ is the smallest bimodule which is (pointwise) greater than R. Thus, we have:

Corollary 4.5.10 If (X, α) is a final F-coalgebra, then

$$X(x, y) = \bigvee\{R(x, y) \mid R \text{ is an } F\text{-simulation up-to similarity}\}$$

for all $x \in X_0, y \in Y_0$. $\square$

Example 4.5.11 As a simple illustration of this result we recover a known coinduction principle for proving inclusions of languages; cf. [74]. We consider an endofunctor on Preord which is a counterpart to the set functor F_3 for deterministic automata from section 2.1.1. Let $G\colon \mathsf{Preord} \to \mathsf{Preord}$ be defined by $G(X) = \underline{2} \times (A \to X)$, where A is a fixed set of input symbols. The final G-coalgebra is the set of languages over the alphabet A, ordered by inclusion; we can carry over the description of the coalgebra structure on this set from loc. cit.

A simulation up-to similarity on the final coalgebra is a relation R such that for each pair of languages L and L',

$$L\,R\,L' \text{ implies } \begin{cases} L\!\downarrow \Rightarrow L'\!\downarrow \text{ and} \\ L_a\,\hat{R}\,L'_a. \end{cases}$$

Note that in this case $\hat{R} = \leq \cdot R \cdot \leq$, where $\leq$ is the order relation on the final coalgebra.

Thus, in order to prove an inclusion of regular languages one only has to exhibit a simulation up-to similarity between them. For instance, Rutten [74] uses this principle to prove inclusions of languages produced by regular expressions. His justification for the principle is ad hoc—in contrast to the rest of that paper which is described as an exercise in universal coalgebra. □

4.5.3 Final semantics

Any F-coalgebra can be interpreted in the final one via the unique homomorphism (denoted for an F-coalgebra (Y, β) by $[\![-]\!]_\beta$)—this yields a so-called final semantics for F-coalgebras. Should F preserve embeddings we have that this final semantics is abstract with respect to F-simulations. The following theorem makes this precise.

Theorem 4.5.12 Suppose F preserves embeddings, (X, α) is a final F-coalgebra and (Y, β) an F-coalgebra; then for all $y, y' \in Y_0$,

$$X([\![y]\!]_\beta, [\![y']\!]_\beta) = \bigvee \{R(y, y') \mid R \text{ is an } F\text{-simulation on } (Y, \beta)\}.$$

Proof. Suppose R is an F-simulation on (Y, β) with collage $Y \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} Y$; then there exists an F-coalgebra structure ρ on $\mathrm{Coll}(R)$ such that r_1 and r_2 are F-coalgebra maps. Thus

$$X([\![y]\!]_\beta, [\![y']\!]_\beta) = X([\![r_1(y)]\!]_\rho, [\![r_2(y')]\!]_\rho) \geq \mathrm{Coll}(R)(r_1(y), r_2(y')) = R(y, y').$$

Also, by proposition 4.5.6, we have that $X([\![-]\!]_\beta, [\![-]\!]_\beta) = [\![-]\!]_\beta^* \cdot [\![-]\!]_\beta$ is an F-simulation on (Y, β), and so the result holds. □

4.6 Complete metric spaces

In metric domain theory, to facilitate a uniform construction of invariants for a wide class of functors, one works with complete metric spaces. Below, we apply the ideas of section 4.5 to derive coinduction principles for invariants of functors on the category CMet of complete metric spaces and non-expansive maps. Perhaps it is surprising that we obtain a coinduction principle for CMet. The collage of a bimodule between two pseudometric spaces is, in general, only a *generalized* metric space. However, the kind of functors arising in metric semantics—product, sum, powerdomain and function space—are restrictions of functors on GMet. This fact enables us to prove a coinduction theorem for invariants of such functors by directly appealing to theorem 4.5.8. In these applications the notion of F-similarity is symmetric, but since this depends on particular properties of the functor F, and to avoid confusion with Aczel–Mendler bisimulations, we resist the temptation to talk about F-*bi*similarity.

Thus far we have only considered invariants of endofunctors, however the need to model 'mixed-variance' domain constructions, like function space, leads us to consider invariants of functors $\mathsf{GMet}^{\mathrm{op}} \times \mathsf{GMet} \to \mathsf{GMet}$. This class of functors includes endofunctors of GMet by taking them to be dummy in the first variable. Recall that theorem 2.3.10 tells us that mixed-variance functors on CMet have unique invariants as long as they are locally contractive.

Theorem 4.6.1 Suppose that $F\colon \mathsf{GMet}^{\mathrm{op}} \times \mathsf{GMet} \to \mathsf{GMet}$ is locally contractive and restricts to a functor $\mathsf{CMet}^{\mathrm{op}} \times \mathsf{CMet} \to \mathsf{CMet}$. If (the restriction of) F has invariant $i\colon \mathrm{Fix}(F) \cong F(\mathrm{Fix}(F), \mathrm{Fix}(F))$, then $\forall x, y \in \mathrm{Fix}(F)_0$

$$\mathrm{Fix}(F)(x, y) = \inf\{R(x, y) \mid R \text{ is an } F(\mathrm{Fix}(F), -)\text{-simulation on } \mathrm{Fix}(F)\}.$$

Proof. We show that $(\mathrm{Fix}(F), i)$ is a final $F(\mathrm{Fix}(F), -)$-coalgebra; then the result holds by theorem 4.5.8. To this end, let (A, α) be an $F(\mathrm{Fix}(F), -)$-coalgebra. The completeness of $\mathrm{Fix}(F)$ implies that $(A \to \mathrm{Fix}(F))$ is also complete; it is also non-empty since it contains constant functions. Now, a map $f\colon A \to \mathrm{Fix}(F)$ is a morphism of coalgebras $(A, \alpha) \to (\mathrm{Fix}(F), i)$ iff it is a fixed point of

$$\phi\colon (A \to \mathrm{Fix}(F)) \to (A \to \mathrm{Fix}(F)),$$

given by $\phi(g) = i^{-1} \cdot F(\mathrm{Fix}(F), g) \cdot \alpha$. Since F is locally contractive it follows that ϕ is a contractive map; but Banach's theorem states that a contractive map on a non-empty complete metric space has a unique fixed point. $\square$

4.6.1 Example: metric applicative bisimulation

Let the functor $F\colon \mathsf{GUMet}^{\mathrm{op}} \times \mathsf{GUMet} \to \mathsf{GUMet}$ be defined by

$$F(X, Y) = 1 + \tfrac{1}{2} \cdot (X \to Y) \tag{4.13}$$

where $1 = \{*\}$ is the ultrametric space with one element, and $\frac{1}{2} \cdot X$ has the same set of elements as X but with all distances halved. Observe that F restricts to a locally contractive functor $\mathsf{CUMet}^{\mathrm{op}} \times \mathsf{CUMet} \to \mathsf{CUMet}$. This restriction has a unique fixed point—written $\mathrm{Fix}(F)$. We aim to prove a so-called internal full abstraction theorem for $\mathrm{Fix}(F)$—considered as an applicative transition system—using theorem 4.6.1. We define a notion of *metric applicative bisimulation*, and, by relating it to our categorical notion of simulation, we show that the metric on $\mathrm{Fix}(F)$ is the maximum metric applicative bisimulation. This is patterned after some earlier work of Pitts [68] and Abramsky [1] in an order-theoretic setting. They considered the minimal solution of a domain equation for the untyped lazy λ-calculus in $\mathsf{Cppo}_\perp$, and showed that the order relation on their solution is the maximum applicative simulation.

Fix an object A of CUMet, and suppose that (X, α) and (Y, β) are $F(A, -)$-coalgebras. We think of such coalgebras as applicative transition systems; cf. [1]. If $x \in X_0$ has $\alpha x = *$ we say that x becomes inert and write $x \Downarrow *$; if $\alpha x = f \in F(A, X)$ we write $x \Downarrow f$.

Definition 4.6.2 Since $F(A, -)$-coalgebras are applicative transition systems, we have the following definition. A relation $R \subseteq X_0 \times Y_0$ is an *applicative bisimulation* between (X, α) and (Y, β) if $\forall x \in X_0, y \in Y_0$,

$$\begin{aligned} R(x, y) \ \wedge \ x \Downarrow * \quad &\Rightarrow \quad y \Downarrow * \\ R(x, y) \ \wedge \ x \Downarrow f \quad &\Rightarrow \quad \exists g :\ y \Downarrow g \ \wedge (\forall a \in A_0)\ R(fa, ga). \end{aligned}$$

We also make the following new definition: a function $R\colon X_0 \times Y_0 \to [0, 1]$ is a *metric applicative bisimulation* if $\forall x \in X_0, y \in Y_0$,

$$\begin{aligned} R(x, y) < 1 \ \wedge \ x \Downarrow * \quad &\Rightarrow \quad y \Downarrow * \\ R(x, y) < 1 \ \wedge \ x \Downarrow f \quad &\Rightarrow \quad \exists g :\ y \Downarrow g \ \wedge (\forall a \in A_0)\ R(fa, ga) \leq 2.R(x, y). \end{aligned}$$

The idea is that $R(x, y) < 2^{-n}$ implies x and y are applicatively bisimilar up-to $n + 1$ steps. □

Remark 4.6.3 If R is a metric applicative bisimulation, then $R^{-1}\{0\} \subseteq X_0 \times Y_0$ is an applicative bisimulation. Also, note that metric applicative bisimulations are just functions—we do not require them to be non-expansive maps. □

By definition of $\mathrm{Fix}(F)$, we have an isomorphism $i\colon \mathrm{Fix}(F) \cong F(\mathrm{Fix}(F), \mathrm{Fix}(F))$. Then, as in the proof of theorem 4.6.1, $(\mathrm{Fix}(F), i)$ is a final coalgebra of $F(\mathrm{Fix}(F), -)$.

Proposition 4.6.4 For $(\mathrm{Fix}(F), i)$ as above,

$$\mathrm{Fix}(F)(x, y) = \inf\{R(x, y) \mid R \text{ is a metric applicative bisimulation on } \mathrm{Fix}(F)\}$$

Proof. By definition, a bimodule R on $\mathrm{Fix}(F)$ is an $F(\mathrm{Fix}(F), -)$-simulation iff $R(x, y) \geq F(\mathrm{Fix}(F), \mathrm{Coll}(R))$ for all $x, y \in \mathrm{Fix}(F)$. It is readily seen that each $F(\mathrm{Fix}(F), -)$-simulation on $(\mathrm{Fix}(F), i)$ is a metric applicative bisimulation, and each metric applicative bisimulation is an $F(\mathrm{Fix}(F), -)$-simulation up-to similarity. The statement of the proposition now follows from corollary 4.5.10. □

4.7 Probabilistic bisimulation I

De Vink and Rutten [83] have given an analysis of probabilistic bisimulation, first introduced by Larsen and Skou [51], using the Aczel–Mendler notion of coalgebraic bisimulation. One striking aspect of their work is that they use a well-known result of graph theory, the *max-flow min-cut theorem*, to prove that their final coalgebra semantics for probabilistic transition systems is abstract with respect to probabilistic bisimulation. In this section we give an alternative analysis of probabilistic bisimulation using the notion of F-simulation introduced above. It turns out that we consider the same final coalgebra as [83], but full abstraction is trivial in our case. Also, a direct characterization of F-simulations for the relevant functor F is immediate (see (4.15)), and generalizes Larsen and Skou's original definition of a probabilistic bisimulation. Next, this generalization motivates a new characterization of probabilistic bisimilarity in terms of a probabilistic modal logic.

4.7.1 Bisimulation

Recall the discrete-distributions functor $D\colon \mathsf{Set} \to \mathsf{Set}$ from section 2.1.3. If X is a partially ordered set, then DX can be ordered in the following way:

$$\mu \leq \nu \text{ iff } \int_X f d\mu \leq \int_X f d\nu \;\text{ for all monotone } f\colon X \to \underline{2}. \tag{4.14}$$

Clearly this is the same as saying that $\mu \leq \nu$ iff $\mu[U] \leq \nu[U]$ for all upward closed $U \subseteq X$. Also, observe that if X has the discrete partial order then so does DX.

We can extend D to an endofunctor on Pos by defining $D(f)(\mu)(x) = \mu[f^{-1}(\{x\})]$ for each monotone map f. Now think of Set as a full, reflective subcategory of Pos by giving each set the discrete partial order (the reflection of a poset $(X, \leq)$ is the quotient of X by the smallest equivalence relation containing $\leq$). Assume a fixed set Act of actions and let 1 denote the singleton set $\{*\}$; then the functor

$$F = D(1 + (-))^{\mathrm{Act}} : \mathsf{Pos} \to \mathsf{Pos}$$

preserves the subcategory Set. Indeed, the restriction of F to an endofunctor on Set is the functor considered in [83]. Moreover, from lemma 4.7.1 it follows that the final coalgebra of the restriction of F to Set, whose existence follows from the accessibility of F, and which was considered in [83], is in fact the final F-coalgebra.

Lemma 4.7.1 (The reflective subcategory lemma [29]) If $\mathcal{A}$ is a reflective subcategory of $\mathcal{B}$, and if $T\colon \mathcal{B} \to \mathcal{B}$ is an endofunctor which has $\mathcal{A}$ as an invariant subcategory, then a final coalgebra for the restriction of T to an endofunctor on $\mathcal{A}$ remains a final coalgebra for $T\colon \mathcal{B} \to \mathcal{B}$. □

As we already pointed out in section 2.1.3, F-coalgebras whose carriers have the discrete order can be seen as probabilistic transition systems, also called partial labelled Markov chains (plMc*s*). Thus, we regard 'plMc' as shorthand for 'F-coalgebra whose carrier has the discrete order'. To lessen the notational clutter, given an F-coalgebra (Pr, μ), $p \in \mathrm{Pr}$ and $a \in \mathrm{Act}$, we write $\mu_{p,a}$ for $\mu(p)(a)$.

Larsen and Skou define a probabilistic bisimulation on a plMc to be a partition of the state space into blocks, such that for any two states in the same block the probability of moving to any other block is the same. Formally:

Definition 4.7.2 A *probabilistic bisimulation* on a plMc (Pr, μ) is an equivalence relation $\equiv$ on Pr such that

$$p \equiv q \Rightarrow \mu_{p,a}[E] = \mu_{q,a}[E]$$

for all $a \in \mathrm{Act}$ and $E \in \mathrm{Pr}/\equiv$. Two processes p and q are said to be *probabilistically bisimilar* if some probabilistic bisimulation contains the pair (p, q). □

Observe that a $\underline{2}$-bimodule between discrete partial orders is just a relation. Let (Pr, μ) and (Pr', μ') be plMc*s* and $R \subseteq \mathrm{Pr} \times \mathrm{Pr}'$ a relation. Write $\mathrm{Pr}' \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} \mathrm{Pr}$ for the collage of R, then inspecting definition 4.5.2 we find that R is an F-simulation iff for all $a \in \mathrm{Act}$, $p, q \in \mathrm{Pr}$ and upward closed $U \subseteq 1 + \mathrm{Coll}(R)$:

$$pRq \Rightarrow \mu_{p,a}[(1 + r_2)^{-1}U] \leq \mu'_{q,a}[(1 + r_1)^{-1}U].$$

With a little thought one sees that this is equivalent to requiring that

$$pRq \Rightarrow \begin{cases} \mu_{p,a}(*) = \mu'_{q,a}(*) \text{ and} \\ \forall O \subseteq \mathrm{Pr} : \mu_{p,a}[O] \leq \mu'_{q,a}[R(O)], \end{cases} \tag{4.15}$$

where $R(O)$ is the image of O under R. Observe that, as a consequence of the inequality $\mu_{p,a}[\mathrm{Pr}] \leq \mu_{q,a}[R(\mathrm{Pr})] \leq \mu_{q,a}[\mathrm{Pr}']$, and the fact the distributions have total mass 1, we write an equality on the upper right-hand side, rather than an inequality.

If R is an equivalence relation on the carrier of a plMc then the notion of F-simulation is in agreement with the notion of probabilistic bisimulation from definition 4.7.2. First, suppose that R is an F-simulation on (Pr, μ), and $p, q \in \mathrm{Pr}$ are processes such that pRq. Then, for each R-equivalence class O we have $R(O) = O$, so $\mu_{p,a}[E] \leq \mu_{q,a}[E]$ for each $a \in \mathrm{Act}$ where E is either an R-equivalence class or $\{*\}$. But probability distributions must sum to 1 over any partition of $\mathrm{Pr} + 1$, hence we actually get equality above. Conversely, if R is a probabilistic bisimulation, then for all $O \subseteq \mathrm{Pr}$ it holds that $\mu_{p,a}[O] \leq \mu_{p,a}[R(O)] = \mu_{q,a}[R(O)]$ since $R(O)$ is a union of equivalence classes.

From the characterization in (4.15) it is readily verified that the composition of two F-simulations is an F-simulation. Moreover, if R is an F-simulation between (Pr, μ) and (Pr', μ') then so is R^{-1} since $qR^{-1}p$ implies that for all $O \subseteq \mathrm{Pr}'$, $a \in \mathrm{Act}$,

$$
\begin{array}{ll}
 & R(X \setminus R^{-1}(O)) \subseteq Y \setminus O \\
\Rightarrow & (R \text{ a bisimulation}) \\
 & \mu_{p,a}[X \setminus R^{-1}(O)] \leq \mu'_{q,a}[Y \setminus O] \\
\Rightarrow & \\
 & \mu'_{q,a}[O] \leq \mu_{p,a}[R^{-1}(O)].
\end{array}
$$

It follows that F-similarity on (Pr, μ)—the union of all F-simulations—is an equivalence relation, and so agrees with probabilistic bisimilarity as it is given in definition 4.7.2.

Given an order embedding $\iota : X \hookrightarrow Y$ and an upward closed $O \subseteq X$, O is equal to the inverse image under ι of the upward closure of $\iota(O)$ in Y. From this it easily follows that F preserves order embeddings, and so, by theorem 4.5.12, the largest F-simulation on a F-coalgebra is f^*f, where f is the unique map to the final F-coalgebra. Since the order on the latter is discrete, the relation f^*f is just the kernel of f. Thus, any $p, q \in \mathrm{Pr}$ are probabilistically bisimilar iff they have the same denotation in the final F-coalgebra. The proof of this result in [83] required a lemma saying that the restriction of F to Set preserves weak pullbacks—this is where the max-flow min-cut theorem was used. This last theorem is used by Jones [45] to prove the 'splitting lemma', and it seems to us that the reason for the simplicity of the current approach is that the definition of F-simulation refers directly to an inequality of measures, whereas span-based definitions of probabilistic bisimulation express this inequality indirectly through splittings of the relevant measures (in this respect the treatment of probabilistic bisimulations as spans of open maps in [22] is similar).

The generalization of the definition of probabilistic bisimulation (4.15), allowing for arbitrary relations rather than just equivalence relations, is new as far as we know. Though the notion of probabilistic *bisimilarity* is unchanged, the more permissive definition of *bisimulation* yields a more powerful coinduction principle. At the very least this enables us to avoid certain technical circumlocutions in the literature: like having to combine two systems into one before one talking about bisimulations between their states, and having to take transitive closures of relations.

Finally, we remark on a surprising aspect of the analysis above. Ultimately we were only interested in coalgebras of a set functor F as plMc*s*, with bisimilarity on any plMc being given by the kernel of the unique homomorphism to the final plMc. However, we found it useful to extend F to an endofunctor on Pos *purely to define the notion of F-simulation*!

4.7.2 Modal logic

Our generalization of the notion of probabilistic bisimulation to relations which are not necessarily equivalences is reminiscent of the notion, due to Bloom et. al [16], of a ready simulation on a labelled transition system. This motivates a characterization of probabilistic bisimilarity in terms of a probabilistic version of the Limited Modal Logic of [16], strengthening a result of [51]. We conclude by comparing this result to a theorem of Desharnais et. al [23].

Let $(\mathrm{Pr}, \rightarrow)$ be an Act-labelled transition system. A relation R on Pr is called a *ready simulation* (also a 2/3-bisimulation) if for all $p, q \in \mathrm{Pr}$ and $a \in \mathrm{Act}$,

$$pRq \Rightarrow \begin{cases} p \stackrel{a}{\rightarrow} p' \text{ implies } \exists q'.q \stackrel{a}{\rightarrow} q' \wedge p'Rq', \text{ and} \\ q \stackrel{a}{\rightarrow} \text{ implies } p \stackrel{a}{\rightarrow} \end{cases}$$

If $O \subseteq \mathrm{Pr}$ write $p \stackrel{a}{\rightarrow} O$ if there exists $q \in O$ such that $p \stackrel{a}{\rightarrow} q$. We may characterize ready simulations as those relations R such that for all $p, q \in \mathrm{Pr}$ and $a \in \mathrm{Act}$,

$$pRq \Rightarrow \begin{cases} p \stackrel{a}{\rightarrow} \text{ iff } q \stackrel{a}{\rightarrow} \text{, and} \\ \forall O \subseteq \mathrm{Pr} : p \stackrel{a}{\rightarrow} O \text{ implies } p \stackrel{a}{\rightarrow} R(O). \end{cases}$$

We say that p and q are *ready similar* if there exists a ready simulation R such that pRq. Despite the obvious likeness between ready simulations and probabilistic bisimulations, observe that the inverse of a ready simulation need not be a ready simulation. Thus the notion of ready similarity is not symmetric. The proof of the corresponding fact for probabilistic bisimulations used the fact the probabilities must sum to 1. The intersection the relation of ready similarity and its inverse is called *GSOS trace congruence* in [16]; they show that this relation is the congruence generated by trace equivalence by the so-called GSOS process operations.

The formulas of *Limited Modal Logic (LML)* are given by the syntax

$$\phi ::= \mathrm{tt} \mid \phi_1 \wedge \phi_2 \mid \langle a \rangle \phi \mid [a]\mathrm{ff}.$$

The relation $R := \{(p, q) : p \vDash \phi$ implies $q \vDash \phi$ for all LML formulae $\phi\}$ is a ready simulation. Thus p and q are ready similar just in case it holds that whenver p satisfies an LML formula then q also satisfies the same formula. Correspondingly, we show how to characterize probabilistic bisimilarity by a probabilistic version of Limited Modal Logic which we call $\mathcal{L}$.

The formulas of $\mathcal{L}$ are given by the following syntax:

$$\phi ::= \mathrm{tt} \mid \phi_1 \wedge \phi_2 \mid \Delta_{a,r} \mid \langle a \rangle_r \phi$$

where $a \in \mathrm{Act}$ and $r \in [0,1] \cap \mathbb{Q}$.

As usual, we define the *modal depth* of a formula $\phi \in \mathcal{L}$ as follows.

$$\begin{aligned} \mathrm{md}(\mathrm{tt}) &= 0 \\ \mathrm{md}(\Delta_{a,r}) &= 1 \\ \mathrm{md}(\phi_1 \wedge \phi_2) &= \max(\mathrm{md}(\phi_1), \mathrm{md}(\phi_2)) \\ \mathrm{md}(\langle a \rangle_r \phi) &= \mathrm{md}(\phi) + 1. \end{aligned}$$

Given a plMc (Pr, μ) we define the satisfaction relation $\vDash$ between processes and formulae by induction on ϕ. We write $[\![\phi]\!]$ for the set of processes in Pr which satisfy ϕ. The definition of satisfaction for the propositional connectives is the obvious one. We have $p \vDash \langle a \rangle_r \phi$ just in case $\mu_{p,a}[\![\phi]\!] > r$, and $p \vDash \Delta_{a,r}$ just in case $\mu_{p,a}(*) > r$. Henceforth assume a fixed plMc (Pr, μ).

Proposition 4.7.3 Let $R_n := \{(p,q) : \forall \phi \in \mathcal{L}$ s.t. $\mathrm{md}(\phi) \leq n, p \vDash \phi \Rightarrow q \vDash \phi\}$ for $n < \omega$. Then for all $n < \omega$ (i) R_n is an equivalence, and (ii) $R_{n+1} = \{(p,q) : \mu_{p,a}[E] = \mu_{q,a}[E]$ for all $a \in \mathrm{Act}, E \in \mathrm{Pr}/R_n\}$.

Proof. Obviously (i) holds for $n = 0$, and if (i) and (ii) hold for $n = n_0$ then (i) holds for $n = n_0 + 1$. It suffices to prove that, for each $n < \omega$, (i) implies (ii).

Suppose that R_n is an equivalence relation. Let E be an R_n-equivalence class, $e \in E$, and $\phi_1, \phi_2, ...$ an enumeration of $\{\phi \in \mathcal{L} : e \vDash \phi$ and $\mathrm{md}(\phi) \leq n\}$. From the definition of R_n (note that the asymmetry is crucial here) we have

$$E = [\![\phi_1]\!] \cap [\![\phi_1 \wedge \phi_2]\!] \cap [\![\phi_1 \wedge \phi_2 \wedge \phi_3]\!] \cap ... \tag{4.16}$$

We prove the equality in (ii) by showing that each side includes the other. Suppose $(p,q) \in R_{n+1}$ and $a \in \mathrm{Act}$. Then $\mu_{p,a}[[\![\phi]\!]] \leq \mu_{q,a}[[\![\phi]\!]]$ for each $\phi \in \mathcal{L}$ with $\mathrm{md}(\phi) \leq n$. From the equation (4.16) it follows that $\mu_{p,a}[E] \leq \mu_{q.a}[E]$ for each R_n-equivalence class E; cf. (refcontinuity). But, from the presence of the formulas $\langle a \rangle_r \mathrm{tt}$ and $\Delta_{a,r}$ we have that $\mu_{p,a}(*) = \mu_{q,a}(*)$, and since $\mu_{p,a}$ and $\mu_{q,a}$ have total mass 1 we conclude that $\mu_{p,a}[E] = \mu_{q,a}[E]$ for each R_n-equivalence class E. It remains to prove the converse—that $(p,q) \in R_{n+1}$ if $\mu_{p,a}[E] = \mu_{q,a}[E]$ for each R_n-equivalence class E and $a \in \mathrm{Act}$. Note first that if $\mathrm{md}(\phi) \leq n$ then $[\![\phi]\!]$ is a disjoint union of R_n-equivalence classes. Thus, if $\mu_{p,a}$ and $\mu_{q,a}$ are identical on R_n-equivalence classes, then $\mu_{p,a}[[\![\phi]\!]] = \mu_{q,a}[[\![\phi]\!]]$. So $p \models \langle a \rangle_r \phi$ iff $q \models \langle a \rangle_r \phi$ and (ii) is established. □

Corollary 4.7.4 The relation $R_\omega := \{(p,q) : \forall \phi \in \mathcal{L}, p \models \phi \Rightarrow q \models \phi\}$ coincides with probabilistic bisimilarity on (Pr, μ).

Proof. Clearly $R_\omega = \bigcap_{n<\omega} R_n$, so, by the previous proposition, R_ω is an equivalence relation. Writing $[p]_\alpha$ for the equivalence class of p under R_α, we have $[p]_\omega = \bigcap_{n<\omega} [p]_n$. Furthermore, suppose that $pR_\omega q$, then $\mu_{p,a}[[s]_n] = \mu_{q,a}[[s]_n]$ for each $n < \omega$, $s \in \mathrm{Pr}$ by proposition 4.7.3; so $\mu_{p,a}[[s]_\omega] = \mu_{q,a}[[s]_\omega]$ and R_ω is a probabilistic bisimulation. □

Remark 4.7.5 Let (Pr, μ) be a plMc and E a bisimulation equivalence class. Repeating an idea from the proof of proposition 4.7.3, let $e \in E$ and let $\phi_1, \phi_2, ...$ be an enumeration of the set $\{\phi \in \mathcal{L} : e \models \phi\}$, then, we can write

$$E = [\![\phi_1]\!] \cap [\![\phi_1 \wedge \phi_2]\!] \cap [\![\phi_1 \wedge \phi_2 \wedge \phi_3]\!] \cap ... \tag{4.17}$$

Immediately we see that if Pr is finite then $E = [\![\phi_1 \wedge ... \wedge \phi_n]\!]$ for some n. Thus, for a finite plMc, not only does the logic $\mathcal{L}$ characterize bisimilarity, but for each bisimulation equivalence class there is a single formula which is satisfied only by the states in that class. This is not the case for an arbitrary plMc, but if we extend the language $\mathcal{L}$ to include countable conjunctions then we do again get formulas characterizing bisimulation equivalence classes. □

In [51] Larsen and Skou introduce a probabilistic version of Hennessy-Milner logic; this is basically the logic $\mathcal{L}$ plus disjunction. They show that if (Pr, μ) is a plMc satisfying the *minimal-deviation assumption*, then two processes satisfy the same formulas of their logic just in case they are probabilistically bisimilar. The minimal-deviation assumption says that the transition probabilities in (Pr, μ) cannot come arbitrarily close to each other: there exists $\varepsilon > 0$ such that for all $p, p' \in \mathrm{Pr}, a \in \mathrm{Act}$, $\mu_{p,a}(p')$ is a multiple of ε. In particular, the number of a-labelled transitions from any given state in Pr is bounded above by $\lceil 1/\epsilon \rceil$.

Our result generalizes Larsen and Skou's in a number of ways. We omit the minimal deviation assumption, we have a logic without disjunction, and we do not require two processes to satisfy the same formulas—just that one should satisfy at least as many formulas as the other.

Recently, Desharnais, Edalat and Panangaden [23] have considered bisimulation for discrete and continuous probabilistic transition systems. They show that in both cases two processes are bisimilar just in case they satisfy the same formulas of the sublanguage of PML generated by tt, $\phi_1 \wedge \phi_2$ and $\langle a \rangle_r \phi$—another generalization of the result of Larsen and Skou. The proof uses the $\lambda - \pi$ theorem of measure theory. Intuitively, in the discrete case, their result corresponds to corollary 4.7.4 where one has a richer language, but one only needs to verify that one process satisfies at least as many formulas as the other. Our result does not immediately seem to imply (the discrete case of) theirs or vice-versa. However, remark 4.7.5 strengthens a couple of propositions from loc. cit.—the corresponding characterizability results therein require a language with negation.

Similarly to [23] we observe that $\mathcal{L}$ characterizes bisimulation for possibly countably branching processes even though it features only finitary conjunction. In contrast, to characterize strong bisimulation for countably branching non-probabilistic processes in Hennessy-Milner logic one needs countable conjunctions.

4.8 Probabilistic bisimulation II

One interesting example of a quantitative notion of equivalence between processes is the metric on the class of partial labelled Markov processes, introduced by Desharnais et. al. in their work on compositional reasoning principles for concurrent interacting probabilistic systems [25, 24]. Generalizing the approach of Larsen and Skou [51], to accommodate examples arising from probabilistic specification languages and control software interacting with physical systems, they consider both discrete and continuous systems. The notion of probabilistic bisimulation turns out to be too fragile for their purposes: a slight perturbation of the probabilities is enough to make two systems non-bisimilar. Thus, they consider a more robust notion of approximate equivalence given by a metric (properly speaking, a pseudometric) on the class of processes, and they focus on the intersubstitutivity of *approximately* equivalent processes. For instance, corresponding to the usual congruence proofs for bisimulation they show that various process combinators are non-expansive with respect to the metric. In addition, they are concerned with approximating arbitrary systems by finite state systems, and reasoning about systems in terms of their finite approximants. The equivalence on processes yielded by the notion of zero distance coincides the with probabilistic bisimilarity as given in [24].

In this section we show that the metric of Desharnais et. al is the same as F-similarity in case F is a functor based on the Hutchinson metric on the set of probability measures on a metric space. Until now, the examples we have considered have all involved *ultra*metric spaces, with the distance in the final coalgebra X being giving by the formula

$$X(x, y) = \inf\{2^{-n} \mid x, y \text{ are bisimilar up-to } n \text{ steps}\};$$

the example in section 4.6.1 is typical. In the present section we consider probabilistic processes, where a one step observation of a process yields a value in the interval $[0, 1]$. The distance in the final coalgebra is a trade-off between the depth of observations needed to distinguish processes, and the amount each observation differentiates the processes. The relative weight given to these two factors is determined by the contraction introduced in the definition of the functor F. In any case, the distance in the final coalgebra is a metric, but not an ultrametric. The development in this section is very clearly analogous to that in section 4.7, except that we have changed the logic from one based on the quantale $(\underline{2}, \wedge, 1)$ to one based on the quantale $([0, \infty], +, 0)$.

4.8.1 The Hutchinson Metric

By a measure on a gms X we mean a probability measure on the σ-algebra of Borel sets generated by the metric topology on X. Such a measure μ is said to be *tight*, or *almost compactly supported*, if for each $\varepsilon > 0$ there exists a compact set $K_\varepsilon \subseteq X$ such that $\mu(X \setminus K_\varepsilon) < \varepsilon$. Under quite mild conditions on a metric space X, e.g., completeness and separability, every measure on X is tight, cf. [26]. Complete, separable metric spaces are usually called *Polish spaces*.

Denote the set of tight measures on a gms X by $\mathcal{M}X$. Write $[0, 1]$ for the subspace of the interval $[0, \infty]$ equipped with the generalized metric from example 4.3.5. So $[0, 1](x, y) = y \dot{-} x$. In passing, we remark that this metric generates the Scott topology on the continuous cpo $([0, 1], \geq)$, cf. [79]. We equip $\mathcal{M}X$ with the following generalized metric—a metric analogy of (4.14)

$$\mathcal{M}X(\mu, \rho) \quad = \quad \sup\left\{[0, 1]\left(\int_X f d\mu, \int_X f d\rho\right) \mid f : X \to [0, 1] \text{ n.e.}\right\}. \tag{4.18}$$

Suppose X is a pseudometric space. Then $f : X \to [0, 1]$ is non-expansive precisely when $|fx - fy| \leq X(x, y)$ for all $x, y \in X$. It follows that $1 - f$ is a non-expansive map $X \to [0, 1]$; thus, in this case (4.18) yields the so-called Hutchinson metric on $\mathcal{M}X$ [26], viz.,

$$\mathcal{M}X(\mu, \rho) \quad = \quad \sup\left\{\left|\int_X f d\mu - \int_X f d\rho\right| \mid f : X \to [0, 1], \forall x, y\ |fx - fy| \leq X(x, y)\right\}.$$

The Hutchinson metric metrizes the weak topology on the space of measures. This topology is discussed in the standard introductions to probability theory, and is particularly important in the theory of fractals [11].

Theorem 4.8.1 [26, Thm 2.5.25] If X is a complete metric space then so is $\mathcal{M}X$.

Proof. (Sketch) Suppose μ_n is a Cauchy sequence of measures. Define a linear functional Φ on the space of bounded Lipschitz functions on X by $\Phi(f) = \mathrm{Lim}_n \int_X f d\mu_n$. By the Riesz representation theorem there is a measure μ such that $\Phi(f) = \int_X f d\mu$. Then one shows that μ is tight and $\mu_n \Rightarrow \mu$. □

$\mathcal{M}$ extends to a locally non-expansive endofunctor on the category GMet of generalized metric spaces and non-expansive maps by defining, for $f\colon X \to Y$,

$$\mathcal{M}(f)(\mu) = \mu \cdot f^{-1}.$$

This measure is still tight. Given $\varepsilon > 0$, find $K_\varepsilon \subseteq X$ such that $\mu(X \setminus K_\varepsilon) < \varepsilon$. Then $f(K_\varepsilon)$ is compact, being the continuous image of a compact set, and $\mathcal{M}(f)(\mu)(Y \setminus f(K_\varepsilon)) \leq \mu(X \setminus K_\varepsilon) < \varepsilon$.

Proposition 4.8.2 Suppose $i\colon X \rightarrowtail Y$ is an isometric embedding of generalized metric spaces; then any non-expansive map $f\colon X \to [0,\infty]$ can be extended to a non-expansive map $Y \to [0,\infty]$.

Proof. Here are two solutions. In fact, they are the largest and smallest ones; in the enriched categorical terminology of [53] they are respectively the left and right Kan extensions of f along i.

$$\begin{aligned} \overline{f}(y) &= \inf_{x\in X}[f(x) + Y(i(x), y)] \\ \underline{f}(y) &= \sup_{x\in X}[f(x) \dot{-} Y(y, i(x))]. \end{aligned}$$

Notice that if f maps into $[0,1]$ then so does $\underline{f}$. □

Proposition 4.8.3 The functor $\mathcal{M}\colon \mathsf{GMet} \to \mathsf{GMet}$ preserves isometric embeddings.

Proof. Suppose that $i\colon X \rightarrowtail Y$ is an isometric embedding, and $\mu, \rho \in \mathcal{M}X$. Then we have the following chain of equalities.

$$\begin{aligned} &\mathcal{M}Y(\mathcal{M}(i)(\mu), \mathcal{M}(i)(\rho)) \\ &= \sup\left\{[0,1]\left(\int_Y fd(\mu\cdot i^{-1}), \int_Y fd(\rho\cdot i^{-1})\right) \mid f : Y \to [0,1] \text{ non-exp.}\right\} \\ &= \sup\left\{[0,1]\left(\int_X (f\cdot i)d\mu, \int_X (f\cdot i)d\rho\right) \mid f : Y \to [0,1] \text{ non-exp.}\right\} \\ &= \sup\left\{[0,1]\left(\int_X fd\mu, \int_X fd\rho\right) \mid f : X \to [0,1] \text{ non-exp.}\right\} \text{ by prop. 4.8.2} \\ &= \mathcal{M}X(\mu, \rho). \end{aligned}$$

□

Let us assume in the rest of this section a constant $c \in (0,1)$, a set Act of labels, and, as usual, write 1 for the one element metric space $\{*\}$.

Theorem 4.8.4 Let $F\colon \mathsf{GMet} \to \mathsf{GMet}$ be defined by $FX = [\mathrm{Act} \to c \cdot \mathcal{M}(1+X)]$; then there is a final F-coalgebra.

Proof. F preserves the subcategory of complete metric spaces by theorem 4.8.1. Since $\mathcal{M}$ is locally non-expansive, F is locally contractive, and so the restriction of F to CMet has a final coalgebra by theorem 2.3.10. By the argument in the proof of theorem 4.6.1 it follows that this coalgebra remains final for $F\colon \mathsf{GMet} \to \mathsf{GMet}$. Alternatively, this last fact holds by Freyd's reflective subcategory lemma, because symmetrization then Cauchy completion defines a reflection of GMet in the subcategory CMet (adjunctions compose). □

Van Breugel [18] has already shown that the functor $1 + c\cdot\mathcal{M}(-)\colon \mathsf{CMet} \to \mathsf{CMet}$ has a final coalgebra, however his note does not contain an analysis of this domain.

Proposition 4.8.5 Let X be a generalized metric space; then the metric on $\mathcal{M}(1+X)$ may be characterized as follows:

$$\begin{aligned}&\mathcal{M}(1+X)(\mu,\rho) = \\ &\qquad \sup\left\{[0,1]\left(\int_X f d\mu, \int_X f d\rho\right) \mid f : X \to [0,1] \text{ n.e.}\right\} + [0,1](\mu(*),\rho(*)).\end{aligned} \tag{4.19}$$

Proof. Equation (4.19) follows from the inequality

$$[0,1](\textstyle\int_{X+1} f d\mu, \int_{X+1} f d\rho) \leq [0,1](\int_X f d\mu, \int_X f d\rho) + [0,1](\mu(*),\rho(*)).$$

□

4.8.2 Partial labelled Markov processes

Desharnais et. al [23] introduce the notion of a *partial labelled Markov process*. Below we recapitulate their definition, though, as we later explain, the term will have a more limited scope in this chapter.

Definition 4.8.6 [23, Defn 4.2] A *partial labelled Markov process* (plMp) over a set of actions Act is a triple (Pr,Σ,μ), where Pr is an analytic space[5] of processes or states, Σ is the Borel σ-algebra on $\mathrm{Pr}+1$, and $\mu\colon \mathrm{Pr}\times\mathrm{Act}\times\Sigma \to [0,1]$ is a function such that for

[5]Analytic spaces generalize Polish spaces; we refer the reader to [23] for details.

each $p \in \mathrm{Pr}, a \in \mathrm{Act}$, $\mu_{p,a} := \lambda A.\mu(p, a, A)$ is a measure, and for each $a \in \mathrm{Act}$, $A \in \Sigma$, $\lambda p.\mu(p, a, A)$ is measurable. □

In the terminology of probability theory μ is an Act-indexed set of *conditional probability measures*. Note that each plMc can be seen as a plMp where Σ is taken to be the powerset of the state space.

We think of plMp*s* as probabilistic transition systems with a continuous state space, for instance $\mathbb{R}$. Unlike in the discrete case, we cannot define transition probabilities between states since these will typically be zero. Instead, the result state after each transition is given by a probability measure. If $(\mathrm{Pr}, \Sigma, \mu)$ is a plMp and $p \in \mathrm{Pr}$ then $\mu_{p,a}(*)$ is thought of as the probability that p refuses an a-labelled transition, and, if $A \in \Sigma$, then $\mu_{p,a}(A)$ is thought of as the probability that p accepts an a-labelled transition and ends up in a state in A.

Observe that if the carrier of an F-coalgebra is a metric space then it can be seen as a plMp according to definition 4.8.6. In fact, we are only really ever interested in F-coalgebras whose carriers are metric spaces. This is just as in the previous section—the extension of F to an (isometric embedding preserving) endofunctor on GMet merely allows us to define F-simulations and prove that they have good properties! In light of this we adopt the following convention: by a plMp we mean an F-coalgebra whose carrier is a (proper) metric space. This is in keeping with our convention on partial labelled Markov *chains* in the previous section. Theorem 4.8.4 tells us that there is a final plMp. Note that our use of term plMp is slightly more restrictive than that of [23] because the structure map of an F-coalgebra is required to be non-expansive.

If (X, α) is a plMp, $a \in \mathrm{Act}$ and $x \in X$, then write $x \xrightarrow{a} \mu$ in case $\alpha(x)(a) = \mu \in \mathcal{M}(1 + X)$ (this is a slight departure from the notation of the previous section, but it makes things more manageable). Many times we refer to plMp*s* just by the name of their carrier sets, suppressing the coalgebra structure.

Given two plMp*s* X and Y, a non-expansive map $f\colon X \to Y$ is an F-coalgebra homomorphism (also called a zig-zag morphism) iff for all $x \in X$, $a \in \mathrm{Act}$,

$$x \xrightarrow{a} \mu \Rightarrow fx \xrightarrow{a} \rho \text{ where } \rho = \mu \cdot (1 + f)^{-1}.$$

Furthermore, by (4.19), a bimodule $R\colon Y \nrightarrow X$ is an F-simulation if for all $x \in X, y \in Y, a \in \mathrm{Act}$ and $f\colon \mathrm{Coll}(R) \to [0,1]$ it holds that

$$R(x, y) \;\geq\; c \cdot [0,1]\left(\int_X f \cdot r_2\, d\mu, \int_Y f \cdot r_1\, d\rho\right) + c \cdot [0,1](\mu(*), \rho(*)) \qquad (4.20)$$

where $x \xrightarrow{a} \mu, y \xrightarrow{a} \rho$ and R has collage $Y \xrightarrow{r_1} \mathrm{Coll}(R) \xleftarrow{r_2} X$. This is a metric version of formula (4.15).

Example 4.8.7 Assume that Act is the singleton $\{a\}$, i.e. we consider unlabelled processes. Figure 4.1 shows four simple plMp*s* from [24]. Their state spaces are taken to have the discrete metric, and each transition is labelled by the probability of its occurrence: the difference between 1 and the sum of the probabilities of the transitions leaving any one node is the probability that the given node refuses a transition.

Suppose $R\colon T \nrightarrow S$ is an F-simulation and set $\lambda_{i,j} := R(s_i, t_j)$ for $0 \leq i \leq 3$ and $0 \leq j \leq 4$. What is the minimum value that $\lambda_{0,0}$ can possibly take? That is, how behaviourally equivalent are s_0 and t_0? It is straightforward to see that the minimum possible values of $\lambda_{1,1}$, $\lambda_{1,2}$, $\lambda_{2,1}$ and $\lambda_{2,2}$ are $\frac{1}{2}c$, c, $\frac{1}{2}c$ and 0 respectively. Now suppose $f\colon \mathrm{Coll}(R) \to [0,1]$ is a non-expansive map; set $\alpha_i := f(s_i)$ and $\beta_j := f(t_j)$. Then, in particular,

$$\beta_j \dot{-} \alpha_i \leq \lambda_{i,j} \quad \text{for } 1 \leq i,j \leq 2 \tag{4.21}$$

$$0 \leq \alpha_i, \beta_j \leq 1 \quad \text{for } 1 \leq i,j \leq 2 \tag{4.22}$$

From (4.20) we see that the minimum possible value of $\lambda_{0,0}$ is equal to the minimum possible value of

$$c.((\frac{1}{2}\beta_1 + \frac{1}{2}\beta_2) \dot{-} (\frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2))$$

where the α_i and β_j are subject subject to the inequalities (4.21) and (4.22). Now this is nothing but a linear programming problem. We can solve this by the simplex method (or by inspection in this simple example) to obtain $\lambda_{0,0} = \frac{1}{4}c^2$.

Referring again to figure 4.1, suppose $R\colon Y \nrightarrow X$ is an F-simulation and write $\lambda := R(x,y)$. Then a little reflection on (4.20) shows that λ must satisfy the inequality $\lambda \geq c.(1 - \frac{1}{2}(1-\lambda))$. From this we conclude that the minimum possible value for λ is $\frac{c}{2-c}$. Observe that the smaller the value of c, i.e. the more the future is discounted, the smaller λ becomes. Also, if we replace the transition values 1 and $\frac{1}{2}$ in this example with variables p and q, then the behavioural distance of x and y is a continuous function in p and q. This is what we mean when we say that our model is smooth. □

Remark 4.8.8 Since the functor F preserves isometric embeddings, F-similarity is determined by a final semantics in the sense of theorem 4.5.12: the measure of F-similarity between any two states of a plMp is equal to the distance of their images in the final coalgebra. □

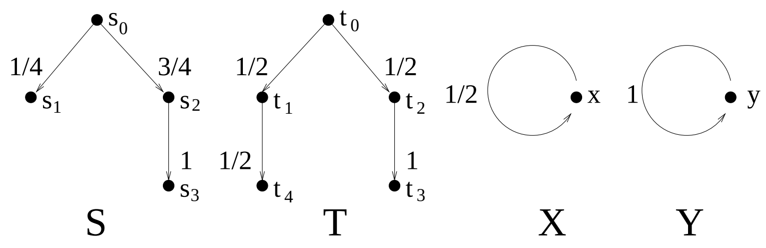

Figure 4.1: Four plMps

4.8.3 Functional expressions and F-(bi)similarity

The definition of the metric on the class of plMps from [25] is motivated by modal logic characterizations of probabilistic bisimilarity. Just as each formula of PML gets interpreted as a Boolean-valued expression on the state space of a plMc, so *op. cit.* defines a syntactic class of *functional expressions* which get interpreted as $[0,1]$-valued functions on the state space of a plMp. This builds on an idea of Kozen [49] to generalize logic to handle probabilistic phenomena, summarized in the table below which is reproduced from [25].

Classical logic	Generalization
Truth values $\{0,1\}$	Interval $[0,1]$
Propositional function	Measurable function
State	Measure
Evaluation of prop. functions	Integration

The set $\mathcal{F}$ of functional expressions is generated by the syntax:

$$f ::= 0 \mid \max(f_1, f_2) \mid 1 - f \mid \langle a \rangle f \mid f \dot{-} q,$$

where $a \in \mathrm{Act}$ and $q \in [0,1] \cap \mathbb{Q}$. We also use the following abbreviations for functional expressions: $\min(f_1, f_2)$ for $1 - \max(1 - f_1, 1 - f_2)$, $[a]f$ for $c - \langle a \rangle (1 - f)$, and $f \dot{+} r$ for $1 - ((1 - f) \dot{-} r)$.

Informally, there is the following correspondence between functional expressions and formulas of PML: tt is represented by 0; conjunction by max; negation by $1 - (-)$; and the connective $\langle a \rangle_q$ decomposes as the prefixing construct $\langle a \rangle$ and truncated subtraction $f \dot{-} q$.

Given a plMp $X = (X, \alpha)$, to each functional expression f there corresponds a map $f_X : X \to [0,1]$ defined by induction on f. We define 0_X to be $\lambda x.0$; the clause for $\langle a \rangle f$ is

$$(\langle a \rangle f)_X(x) = c \left(\int_X f_X d\mu + \mu(*) \right) \quad \text{where } x \xrightarrow{a} \mu;$$

the other clauses are interpreted in the obvious manner. Note that we have

$$([a]f)_X(x) = c\int_X f_X d\mu \quad \text{where } x \xrightarrow{a} \mu.$$

The authors of [25] consider plMp*s* with distinguished initial states. Given two such processes X, Y, with respective initial states x, y, they define their distance as follows:

$$d^c(X,Y) = \sup_{f\in\mathcal{F}} |f_X(x) - f_Y(y)|. \tag{4.23}$$

Example 4.8.9 Consider the processes S and T from example 4.8.7. Then the functional expression

$$f := \langle a\rangle \min(\langle a\rangle 0, [a]1)$$

satisfies $|f_S(s_0) - f_T(t_0)| = \frac{1}{4}c^2$.

Consider the processes X and Y from example 4.8.7. We define a sequence (f_n) of functional expressions as follows. $f_0 = 0$ and $f_{n+1} = \langle a\rangle f_n$. Then, writing $\lambda_n := |(f_n)_X(x) - (f_n)_Y(y)|$, we find that $\lambda_{n+1} = \frac{1}{2}c(1+\lambda_n)$. It follows that λ_n converges to $\frac{2}{2-c}$ from below as $n \to \infty$. Moreover, an easy induction shows that $|f_X(x) - f_Y(y)| \leq \frac{2}{2-c}$ for all functional expressions f; thus $d^c(X,Y) = \frac{2}{2-c}$. □

Comparing examples 4.8.7 and 4.8.9 we see that the metric (4.23) assigns the same distance to the processes X and Y as F-similarity does. We should mention that there are a couple of differences between example 4.8.9 and the corresponding example from [25]. Firstly they interpret the functional expression $\langle a\rangle$ as we interpret $[a]$, but this doesn't make any difference to the metric d^c in the presence of negation. Also they obtain too small a value for $d^c(S,T)$ by not considering enough functional expressions. Next, we prove that the metric d^c corresponds to F-similarity for arbitrary plMp*s*.

Proposition 4.8.10 Let (X, α) be a plMp; then for each functional expression $f \in \mathcal{F}$ the map $f_X : X \to [0,1]$ is non-expansive.

Proof. The proof is by induction on $f \in \mathcal{F}$; the only non-trivial case is for $f = \langle a\rangle g$. Let $x, y \in X$ and $a \in \mathrm{Act}$ be such that $x \xrightarrow{a} \mu$ and $y \xrightarrow{a} \rho$. Then

$$\begin{aligned}
|((\langle a\rangle g)_X(x) - ((\langle a\rangle g)_X(y)| &= c.\left|\int_X g_X d\mu + \mu(*) - \int_X g_X d\rho - \rho(*)\right| \\
&\leq c\cdot\mathcal{M}(1+X)(\mu,\rho) \ (g_X \text{ is non-expansive}) \\
&\leq X(x,y) \ (\text{the coalgebra structure on } X \text{ is non-expansive}).
\end{aligned}$$

□

Proposition 4.8.11 Suppose (X, α) and (Y, β) are plMp*s*, $\varphi\colon X \to Y$ is a zig-zag map, and $f \in \mathcal{F}$ is a functional expression; then $f_X = f_Y \cdot \varphi$.

Proof. The proof is by induction on $f \in \mathcal{F}$. Again, the only non-trivial case is for $f = \langle a \rangle g$. Suppose that $x \in X$, $x \xrightarrow{a} \mu$ and $\varphi(x) \xrightarrow{a} \rho$; then $\rho = \mu \cdot (1 + \varphi)^{-1}$ by the zig-zag condition. Now

$$\begin{aligned}
(\langle a \rangle g)_X(x) &= c \cdot \left(\int_X g_X d\mu + \mu(*) \right) \\
&= c \cdot \left(\int_X g_Y \cdot \varphi d\mu + \mu(*) \right) \quad \text{(inductive hypothesis)} \\
&= c \cdot \left(\int_Y g_Y d\rho + \rho(*) \right) \quad \text{(zig-zag condition)} \\
&= (\langle a \rangle g)_Y(\varphi(x)).\square
\end{aligned}$$

Theorem 4.8.12 Suppose (X, α) is a plMp and let (Y, β) be the final plMp. If $[\![-]\!]\colon X \to Y$ is the unique zig-zag map, then for all $x, y \in X$

$$Y([\![x]\!], [\![y]\!]) = \sup_{f \in \mathcal{F}} |f_X(x) - f_X(y)|.$$

That is, the metric (4.23) is given by a final coalgebra semantics.

Proof. Let $f \in \mathcal{F}$ be a functional expression and $x, y \in X$. By proposition 4.8.11 we have that $|f_X(x) - f_X(y)| = |f_Y([\![x]\!]) - f_Y([\![y]\!])|$. Furthermore, by proposition 4.8.10, this last is not greater than $Y([\![x]\!], [\![y]\!])$. Thus

$$\sup_{f \in \mathcal{F}} |f_X(x) - f_X(y)| \leq Y([\![x]\!], [\![y]\!]).$$

It remains to prove the converse inequality. To this end, define an equivalence relation $\equiv$ on the state space of the plMp (X, α) by

$$x \equiv y \ \text{ iff } \ f_X(x) = f_X(y) \ \text{ for all } f \in \mathcal{F}.$$

Write $[x]$ for the equivalence class of $x \in X$ under $\equiv$. We define a metric on the set Q of $\equiv$-equivalence classes by

$$Q([x], [y]) = \sup_{f \in \mathcal{F}} |f_X(x) - f_X(y)|$$

(clearly this is well-defined). Since f_X is non-expansive for each $f \in \mathcal{F}$, we have that the quotient map $q\colon X \to Q$ is non-expansive. We claim that the following defines a plMp structure on Q:

$$\text{if } x \xrightarrow{a} \mu \text{ then } [x] \xrightarrow{a} \mu \cdot (1 + q)^{-1}. \tag{4.24}$$

Once this claim is proven it follows that q is a zig-zag map. Thus the unique zig-zag map $[\![-]\!]\colon X \to Y$ factors through q and it follows that for all $x, y \in X$,

$$Y([\![x]\!], [\![y]\!]) \leq Q(qx, qy) = \sup_{f \in \mathcal{F}} |f_X(x) - f_X(y)|.$$

This completes the proof of the theorem. $\square$

It remains to be shown that the transition structure (4.24) is well-defined and induces a *non-expansive* map from Q to $[\mathrm{Act} \to c \cdot \mathcal{M}(1+Q)]$. The key to the proof is to show that the maps f_X can approximate, arbitrarily closely, any non-expansive map $X \to [0,1]$ which factors through q in Met.

Proposition 4.8.13 Suppose $f\colon X \to [0,1]$ is a non-expansive map factoring through q, and let $x, y \in X$. Then there is a sequence g_n of functional expressions in $\mathcal{F}$ such that $(g_n)_X(x) \to f(x)$ and $(g_n)_X(y) \to f(y)$.

Proof. Without loss of generality assume that $f(y) - f(x) = l \geq 0$ and let $\varepsilon > 0$ be given. Since f factors through q there exists a functional expression $g \in \mathcal{F}$ such that $|g_X(y) - g_X(x)| > l - \varepsilon$. By replacing g by $1 - g$ if necessary we may assume that $g_X(y) - g_X(x) > l - \varepsilon$. Clearly we can now choose $r \in \mathbb{Q} \cap [0,1]$ such that $g' \stackrel{\mathrm{def}}{=} \min(g, r)$ satisfies $l - \varepsilon < g'_X(y) - g'_X(x) < l + \varepsilon$. Finally, we can define the functional expression g'' by 'vertically shifting' g'—either using truncated subtraction or truncated addition—so that $|g''_X(x) - f(x)| < \varepsilon$ and $|g''_X(y) - f(y)| < \varepsilon$. $\square$

Next, we need a standard result about spaces of continuous functions. Let K be a compact Hausdorff topological space, and write $C(K)$ for the space of continuous real-valued functions on K equipped with the sup metric: $d(f, g) = \sup\{|fx - gx| : x \in K\}$. The following lemma is from Ash [11] where it is used to prove the Stone-Weierstrass theorem.

Lemma 4.8.14 Let A be a subset of $C(K)$ that is closed under the lattice operations; that is, if $f, g \in A$, then $\max(f, g) \in A$ and $\min(f, g) \in A$. If $f \in C(K)$ can be approximated at each pair of points by functions in A [if $x \neq y$, there is a sequence of functions $f_n \in A$ with $f_n(x) \to f(x)$ and $f_n(y) \to f(y)$], then f lies in the closure of A. $\square$

Combining lemma 4.8.14 and proposition 4.8.13 we conclude that, for any compact $K \subseteq X$, each non-expansive $f\colon X \to [0,1]$ factoring through q can be uniformly approximated on K by maps g_X, where $g \in \mathcal{F}$. In particular, if μ is a tight measure on X then for a given $\varepsilon > 0$ there exists $g \in \mathcal{F}$ such that $\int_X |f - g_X| d\mu < \varepsilon$. Now let $x, y \in X$, $x \xrightarrow{a} \mu$,

$y \xrightarrow{a} \rho$, and suppose without loss of generality that $\rho(*) \leq \mu(*)$; then

$$\begin{aligned}
& c \cdot \mathcal{M}(1+Q)(\mu \cdot (1+q)^{-1}, \rho \cdot (1+q)^{-1}) \\
= \;& c.\sup\left\{\left|\int_{1+Q} f d(\mu \cdot (1+q)^{-1}) - \int_{1+Q} f d(\rho \cdot (1+q)^{-1})\right| \mid f : 1+Q \to [0,1] \text{ n.e.}\right\} \\
= \;& c.\sup\left\{\left|\int_X (f \cdot q) d\mu - \int_X (f \cdot q) d\rho\right| \mid f : Q \to [0,1] \text{ n.e.}\right\} \quad \text{by (4.19)} \\
= \;& c.\sup_{f \in \mathcal{F}} \left|\int_X f_X d\mu - \int_X f_X d\rho\right| \quad \text{(see above)} \\
= \;& \sup_{f \in \mathcal{F}} |([a]f)_X(x) - ([a]f)_X(y)| \\
\leq \;& Q([x],[y]).
\end{aligned}$$

This proves that the coalgebra structure we defined on Q is both well-defined and non-expansive. Thus we have discharged the obligations arising from the proof of theorem 4.8.12.

Corollary 4.8.15 Given plMp*s* (X, α) and (Y, β) with distinguished initial states x and y respectively, then $d^c(X, Y) = R(x, y)$, where R denotes F-similarity. □

Proof. This follows from applying theorem 4.8.12 to the coproduct of (X, α) and (Y, β).

As [25] explains, the exact numerical values conferred by the metric on plMp*s* are not as relevant as the induced notions of convergence and of zero distance. With regard to the latter they show that two processes are a distance zero apart iff there exists a probabilistic bisimulation between them. The notion of probabilistic bisimulation refered to here is a continuous extension of Larsen and Skou's notion for discrete systems.

Remark 4.8.16 As we have already said, the definition of the metric d^c was motivated by the characterization of probabilistic bisimilarity by a modal logic. In view of corollary 4.7.4 we might ask how necessary it is to include the negation construct in the grammar for functional expressions. Informally, the class of functional expressions corresponding to the modal logic $\mathcal{L}$ is that generated by the grammar

$$f ::= 0 \mid \max(f_1, f_2) \mid \langle a \rangle f \mid f \dot{-} q \mid [a]1.$$

Note that the functional expressions considered in example 4.8.9 belong to this class. □

4.8.4 Related work

Another metric space domain for probabilistic bisimilarity is obtained by de Vink and Rutten [83]. But they consider a functor on the category of complete *ultra*metric spaces, and the distance in the final coalgebra reflects the number of steps of *exact* observational equivalence of processes. A slight perturbation in the probabilities associated to a process can result in a large difference in its denotation in the final coalgebra. In fact, as we showed in chapter 3, this domain is closely related to the final coalgebra of the corresponding set functor. Obviously, working with coalgebras on ultrametric spaces also excludes many of the examples covered by the present approach. Finally, *op. cit.* works with Aczel–Mendler style bisimulations, so they are interested in whether their functor preserves weak pullbacks (an open question) as opposed to whether it preserves isometric embeddings.

4.9 Coinduction for cpos

The main focus of this chapter has been on functors on Preord and GMet. However, since the most common type of structure used for modelling recursive types is some kind of cpo, we briefly sketch how our approach applies to cpos, and we compare it to Fiore's notion of an ordered categorical bisimulation.

We saw above that the basic coinduction result, theorem 4.5.8, could be directly applied to fixed points of typical functors on the category of complete metric spaces. This relied on the fact that functors arising in metric semantics tend to be restrictions of functors $\mathsf{GMet}^{\mathrm{op}} \times \mathsf{GMet} \to \mathsf{GMet}$. We cannot take the same approach for cpos, in particular the exponentiation functor mapping a pair of cpos X and Y to the cpo of continuous functions $X \to Y$ does not arise as the restriction of a functor $\mathsf{Pos}^{\mathrm{op}} \times \mathsf{Pos} \to \mathsf{Pos}$.

In general, given cpos X and Y, the collage of bimodule $R\colon Y \nrightarrow X$ is not a cpo. In fact, $\mathrm{Coll}(R)$ is a cpo iff R satisfies the following axiom for all ω-chains (x_n):

$$x \;=\; \bigsqcup_{n<\omega} x_n \;\text{ and }\; \forall n.R(x_n, y) \;\text{ implies }\; R(x, y). \tag{4.25}$$

We call this condition *admissibility.* This is in agreement with the notion of an admissible relation in Pitts [67]. Indeed, his notion of a relational structure is general enough to include bimodules, and (4.25) is an instance of his general definition of an admissible relation. Now, routinely following the development in this chapter, we can define a lifting of a functor $F\colon \mathsf{Cpo} \to \mathsf{Cpo}$ to an action on admissible bimodules. From this we obtain a notion of F-simulation, and a coinduction theorem characterizing the order on a final F-coalgebra as the maximum F-simulation.

We note that Fiore [28], in the spirit of Smyth and Plotkin [80], formulates his coinduction principle for Cpo-enriched categories, thus avoiding a commitment to any specific category of domains. His approach is based on the notion of an ordered F-bisimulation, generalizing an earlier definition of Rutten and Turi [82].

Definition 4.9.1 [28, Definition 5.1] Let (X, α) and (Y, β) be coalgebras of a locally monotone endofunctor F on a Pos-category. A relation $r\colon R \rightarrowtail X \times Y$ *lifts to an ordered categorical bisimulation* if there exists an F-coalgebra structure ρ on R such that

$$\begin{tikzcd} X \arrow[d, "\alpha"'] \arrow[dr, phantom, "\leq"] & R \arrow[l, "r_1"'] \arrow[r, "r_2"] \arrow[d, "\rho"] & Y \arrow[d, "\beta"] \\ FX & FR \arrow[l, "Fr_1"] \arrow[r, "Fr_2"'] & FY \end{tikzcd} \tag{4.26}$$

where $r_i \stackrel{def}{=} \pi_i.r (i = 1, 2)$. □

The associated coinduction principle says that the order relation of a final coalgebra of a locally continuous endofunctor F of a Cpo-category is the maximum ordered F-bisimulation. Moreover, if F maps bistrict lax-kernel pairs to weak right-strict lax-kernel pairs then the final semantics is abstract with respect to ordered F-bisimulation. The proof of the coinduction principle relies on the fact that final coalgebras of locally continuous endofunctors are *maximally final*, that is, for any coalgebra, any lax coalgebra map to the final coalgebra is less than or equal to the unique (proper) coalgebra map to the final coalgebra. In particular, the coinduction principle does not apply to Pos, where final coalgebras need not be maximally final [28].

In summary, an advantage of Fiore's approach is its robustness. The formulation in terms of Cpo-enriched categories means that it is directly applicable to most of the different categories of domains that are based around cpos. On the other hand, the advantages of our notion of simulation over relational notions seem even more apparent when we are working in Cpo rather than Set. In particular, in specific instances it seems much easier to characterize F-simulations then ordered F-bisimulations. Also, the condition that F preserve embeddings, ensuring that the final semantics is abstract with respect to the notion of simulation, is much easier to verify than the corresponding condition in Fiore's approach.

Chapter 5

On the structure of $\mathsf{Coalg}(F)$

In this chapter we investigate the structure of the category of coalgebras of an endofunctor $F\colon \mathcal{E} \to \mathcal{E}$. All of our results assume enough about F to ensure the existence of a final F-coalgebra (and, more generally, cofree F-coalgebras; cf. section 5.3). Beyond that they also tend to follow the same pattern: a given categorical property lifts from $\mathcal{E}$ to $\mathsf{Coalg}(F)$ if F preserves a certain kind of limit.

Recall that a presheaf category is one which is equivalent to a functor category $[\mathcal{C}^{op}, \mathsf{Set}]$ for some small category $\mathcal{C}$. We show that the category of coalgebras of a wide-pullback preserving endofunctor F on a presheaf category is itself a presheaf category. This is inspired by a result of Carboni and Johnstone [20] that the category obtained by Artin gluing along a limit preserving functor between presheaf categories is also a presheaf category. The proof works by constructing a freely generated path category $\mathcal{C}$ from the functor F, such that F-coalgebras correspond to presheaves on $\mathcal{C}$.

Should the functor F above only preserve pullbacks then we no longer say that $\mathsf{Coalg}(F)$ is a category of presheaves, however we can show that it is a category of sheaves on a site, that is, a so-called Grothendieck topos (a category of presheaves is a special case of this)[1]. Moreover, this remains true if the functor we start with is an endofunctor on a category of sheaves. This is not too surprising; what is more surprising is that the same conclusion 'almost holds' under the yet weaker assumption that F only preserves weak pullbacks. In this case $\mathsf{Coalg}(F)$ satisfies all the smallness and exactness conditions required to be a category of sheaves except, possibly, effectiveness of equivalence relations. Embedding $\mathsf{Coalg}(F)$ in its effective completion we obtain a Grothendieck topos.

It is remarkable that there are three 'elementary' axioms[2] which characterize categories

[1] For completeness we define these concepts in Appendix A; however, the definitions are not really needed in what follows because we rely on a characterization of Grothendieck topoi, due to Giraud.

[2] That is, axioms that can be formulated in a first order theory of categories, assuming no set theory.

of sheaves amongst those categories with small coproducts and a set of generators (cf. definition A.0.2). These are the axioms for an elementary topos. They arise from categorical characterizations of familiar constructions from Set, namely Cartesian products, function spaces and the set of truth values: a category is an elementary topos if it has products, exponentials and an object of truth values; cf. definition 5.4.8. In section 5.4 we study the elementary topos structure of $\mathsf{Coalg}(F)$—in particular, we characterize the truth object and products in the category of coalgebras of the finite-powerset functor.

We refer the reader to Barr and Wells [15] for an interpretation of a topos as a 'generalized universe of sets' and on sheaves as 'variable sets'—their picture of a sheaf as a set evolving through time is particularly pertinent to the work below. In a similar vein, Goguen's sheaf semantics for concurrent interacting objects [36] is also relevant. However, this last reference adopts a linear model of time, corresponding to trace equivalence of processes, whereas we are working within a finer branching-time model of processes, corresponding to bisimulation. For further discussion we refer the reader to chapter 6, where we also suggest a link between a sheaf theoretic semantics for temporal logic and the coalgebraic approach of Jacobs [41].

Excepting the results on wide-pullback preserving functors, which are unpublished, this chapter is a development of [86]. That paper overlapped with some independent work of Power and Watanabe [70] who also studied subobject classifiers in categories of coalgebras. The two papers were combined and extended in the joint publication [44]. The present chapter has greatly benefited in both style and content from the collaboration with the authors of [44]. In particular, I wish to acknowledge the contribution of Peter Johnstone. He made me aware of the more general version of Linton's theorem (cf. theorem 5.4.2) and suggested, with example 5.1.5, distinguishing between functors preserving weak pullbacks and those preserving pullbacks of monos. He also first proposed using the effectivization construction employed in section 5.5.

5.1 Limit preservation properties

Definition 5.1.1 A *wide pullback*, also called a fibre product, is the limit of a diagram indexed by a poset with a greatest element in which all maximal chains have length 2. □

Clearly, ordinary pullbacks are a special case of wide pullbacks; so a functor preserving the latter also preserves the former. One important class of wide-pullback preserving functors $\mathsf{Set}^n \to \mathsf{Set}$ is the class of *data-functors*; this is inductively defined as the least class closed under the following constructs, cf. [72].

(i) Projection functors $\Pi^n_i \colon \mathsf{Set}^n \to \mathsf{Set} : (X_1, ..., X_n) \mapsto X_i$;

(ii) constant functors $C \colon \mathsf{Set}^n \to \mathsf{Set} : (X_1, ..., X_n) \mapsto C$;

(iii) the product functor $\times \colon \mathsf{Set}^2 \to \mathsf{Set} : (X_1, X_2) \mapsto X_1 \times X_2$;

(iv) the coproduct functor $+ \colon \mathsf{Set}^2 \to \mathsf{Set} : (X_1, X_2) \mapsto X_1 + X_2$;

(v) the exponent functor $(-)^E \colon \mathsf{Set} \to \mathsf{Set} : X \mapsto X^E$;

(vi) composition of data-functors—i.e. if $U \colon \mathsf{Set}^m \to \mathsf{Set}$ is a data-functor and, for $1 \leq i \leq m$, $T_i \colon \mathsf{Set}^n \to \mathsf{Set}$ are data-functors, then $U \circ \langle T_1, ..., T_m \rangle : \mathsf{Set}^n \to \mathsf{Set}$ is a data-functor;

(vii) the initial algebra carrier $\mu X.T(-, X) \colon \mathsf{Set}^n \to \mathsf{Set}$, where $T \colon \mathsf{Set}^{n+1} \to \mathsf{Set}$ is a data-functor.

(viii) the final coalgebra carrier $\nu X.T(-, X) \colon \mathsf{Set}^n \to \mathsf{Set}$, where $T \colon \mathsf{Set}^{n+1} \to \mathsf{Set}$ is a data-functor.

For instance, the finite-list functor is given by $L(X) = \mu Y.(Y \mapsto 1 + X \times Y)$, and is thus a data-functor.

We are also interested in much weaker limit preservation properties of functors. In particular, we consider functors—like the finite-powerset functor $\wp$—which do not preserve wide pullbacks (or even ordinary pullbacks) but do preserve weak pullbacks (we remind the reader that for a weak limit one requires only existence, not uniqueness, of comparison maps). The following straightforward characterization of weak limits is useful.

Proposition 5.1.2 Suppose $F \colon D \to C$ is a functor with limit $\mathrm{Lim}\, F$. Then an object P is a weak limit of F iff the comparison map $P \to \mathrm{Lim}\, F$ is a split epi. In particular, a limit of F is also a weak limit of F. $\square$

We identify the following hierarchy of limit preservation properties of functors, with each condition stronger than the one below:

(i) preservation of wide pullbacks;

(ii) preservation of pullbacks;

(iii) preservation of weak pullbacks;

(iv) preservation of pullbacks of monos.

Recall that the pullback of a mono $m \colon A \rightarrowtail B$ along $f \colon A \to B$ is usually thought of as the inverse image of m under f (in Set, if $A \subseteq B$, then the pullback of m along f is $f^{-1}(A)$). In the special case that f is also a mono we call the pullback the *intersection* of m and f.

Also, we note that a functor satisfying any of the conditions above preserves monos since a map $m\colon A \to B$ is mono iff

$$\begin{tikzcd} A \arrow[r, equal] \arrow[d, equal] & A \arrow[d, "m"] \\ A \arrow[r, "m"'] & B \end{tikzcd} \tag{5.1}$$

is a pullback.

The following examples show that the each of the conditions (i)-(iv) above is strictly stronger than the one below, and that the last condition is non-trivial even for set functors (the example of the finite-powerset functor already serves to show that (ii) is strictly stronger than (iii), cf. Rutten [73]).

Example 5.1.3 Here is a functor which preserves pullbacks but does not preserve wide pullbacks. Let $F\colon \mathsf{Set} \to \mathsf{Set}$ be the subfunctor of the exponential functor $(-)^{\mathbb{N}}$ consisting of the 'eventually constant functions'. More precisely, we define

$$F(X) = \{f \in X^{\mathbb{N}} : \exists m \forall n, n' \geq m \; f(n) = f(n')\}.$$

It is routine to verify that F preserves pullbacks (indeed it preserves all finite limits); however, it does not preserve the infinite product $P = \mathbb{N} \times \mathbb{N} \times \mathbb{N} \times ...$ since there is no eventually constant map $\mathbb{N} \to P$ which corresponds to the tuple $(f_n\colon \mathbb{N} \to \mathbb{N})_{n \in \mathbb{N}}$ where $f_n(x) = \min(x, n)$. □

Example 5.1.4 Consider the functor $F\colon \mathsf{Set} \to \mathsf{Set}$ given on objects by

$$F(X) = \{(x_1, x_2, x_3) \in X^3 : |\{x_1, x_2, x_3\}| \leq 2\},$$

and on maps $f\colon X \to Y$ by

$$F(f)(x_1, x_2, x_3) = (f(x_1), f(x_2), f(x_3)).$$

It is shown in [3] that this functor does not preserve weak pullbacks; however it is readily seen that it preserves pullbacks of monos (i.e. inverse images). □

Example 5.1.5 Here is functor F which fails even to preserve pullbacks of monos. However, F does satisfy an even weaker limit-preserving condition: preservation of equalizers of coreflexive pairs, cf. section 5.4. Let $F\colon \mathsf{Set} \to \mathsf{Set}$ be defined on objects by $F(X) = X \times X + 1$ and on functions by

$$F(f)(x, y) = \begin{cases} (f(x), f(y)) & \text{if } x = y \;\vee\; f(x) \neq f(y) \\ * & \text{otherwise} \end{cases}$$

and $F(f)(*) = *$ for $f\colon X \to Y$. Then F does not preserve the pullback

$$\begin{tikzcd} \{a\} \arrow[r, equal] \arrow[d] & \{a\} = A \arrow[d] \\ \{a,b,c\} \arrow[r, "g"] & \{a,b\} \end{tikzcd} \tag{5.2}$$

where $g(a) = a$, $g(b) = b$ and $g(c) = b$. For instance, the pair (b, c) lies in $(Fg)^{-1}(FA)$ but not in $F(g^{-1}(A))$; thus the comparison map from the latter to the former is not an epi. □

In the remainder of this section we consider a 'reduction' of a wide-pullback preserving endofunctor on a category $\mathcal{C}$ to a limit preserving endofunctor on a slice of $\mathcal{C}$. Suppose we are given a functor $T : \mathcal{A} \to \mathcal{A}$, where $\mathcal{A}$ has all pullbacks, and an arrow $\alpha : A \to TA$ of $\mathcal{A}$. T has an obvious lifting to a functor T_A on slice categories $\mathcal{A}/A \to \mathcal{A}/TA$, and composing this with the pullback functor $\alpha^* : \mathcal{A}/TA \to \mathcal{A}/A$ we obtain an endofunctor

$$T_\alpha : \mathcal{A}/A \to \mathcal{A}/A.$$

So for an object f of $\mathcal{A}/A$, $T_\alpha f$ is defined by the pullback below.

$$\begin{tikzcd} \bullet \arrow[r] \arrow[d, "T_\alpha f"'] & TA' \arrow[d, "Tf"] \\ A \arrow[r, "\alpha"'] & TA \end{tikzcd} \tag{5.3}$$

Lemma 5.1.6 If $T : \mathcal{A} \to \mathcal{A}$ is a wide-pullback preserving endofunctor on a complete category, and $\alpha : A \to TA$, then T_α is continuous—it preserves all limits.

Proof. Observing that wide pullbacks in a slice category $\mathcal{C}/C$ are created by the forgetful functor $\mathcal{C}/C \to \mathcal{C}$, it is easy to see that T_A inherits the property of preserving wide pullbacks from T. Furthermore, T_A clearly preserves final objects, thus it preserves all limits (since any limit may be constructed from final objects and wide pullbacks). The functor α^* is right adjoint (see below), and thus preserves all limits. It follows that $T_\alpha = \alpha^* \cdot T_A$ is continuous. □

The left adjoint to α^* above is $\Sigma_\alpha : \mathcal{A}/A \to \mathcal{B}/TA$—defined on objects by $\Sigma_\alpha(f') = \alpha \cdot f'$, and on morphisms by $\Sigma_\alpha(h) = h$. Thus we have the following isomorphism, natural in $f, f' \in \mathcal{A}/A$:

$$\mathcal{A}/A(f', T_\alpha f) \cong \mathcal{A}/TA(\Sigma_\alpha f', T_A f). \tag{5.4}$$

Proposition 5.1.7 If $\mathcal{A}$ has all pullbacks, $T : \mathcal{A} \to \mathcal{A}$ in an endofunctor and (A, α) is a T-coalgebra, then $\mathsf{Coalg}(T_\alpha)$ is isomorphic to the slice category $\mathsf{Coalg}(T)/(A, \alpha)$.

Proof. The correspondence (5.4) yields a bijection between the objects of the two categories. That this extends to an isomorphism of categories (acting as the identity on homsets) follows from the naturality of (5.4). □

Finally, we have the following proposition.

Proposition 5.1.8 Suppose $\mathcal{A}$ is a complete category, then every wide-pullback preserving endofunctor $T : \mathcal{A} \to \mathcal{A}$ has a final coalgebra.

Proof. [20, Lemma 2.1] tells us that a wide-pullback preserving functor whose domain category is complete actually preserves all connected limits. It follows that T preserves limits indexed by the co-chain ω^{op}. Thus the final coalgebra of T may be constructed as the limit of the ω^{op}-co-chain

$$1 \leftarrow T1 \leftarrow T^2 1 \leftarrow \ldots$$

in the standard manner. □

5.2 Coalgebras as presheaves

We spell out a biequivalence of the bicategory[3] of small categories and bimodules and the bicategory of Pshf presheaf categories and continuous functors, cf. [53, 20].

Definition 5.2.1 A *bimodule* (also called a profunctor or distributor) from a category $\mathcal{A}$ to a category $\mathcal{B}$, written $\phi : \mathcal{A} \nrightarrow \mathcal{B}$, is a functor

$$\phi : \mathcal{B}^{\mathrm{op}} \times \mathcal{A} \to \mathsf{Set}.$$

We write $\mathsf{Mod}(\mathcal{A}, \mathcal{B})$ for the category of bimodules $\mathcal{A} \nrightarrow \mathcal{B}$ and natural transformations between them. One can define a composition of bimodules; then small categories, bimodules, and natural transformations form a bicategory Mod. □

Note the similarity with definition 4.3.9. In fact, both definitions are instances of the enriched categorical notion of bimodule from Kelly [48].

[3]It is convenient to use the language of bicategories here; however we don't include a precise definition as an informal picture will suffice for the present work, but see [17]. Roughly speaking, a bicategory $\mathcal{C}$ is a category where each hom-set is itself a category (whose morphisms are called 2-cells), and where the associativity axiom and the identity axiom for composition in $\mathcal{C}$ need only hold up-to isomorphism of 2-cells. A biequivalence of bicategories is an equivalence of the underlying categories where the normal 'isomorphism of hom-sets' becomes an 'equivalence of hom-categories'.

Observe that a bimodule $\phi : \mathcal{A} \nrightarrow \mathcal{B}$ may be regarded, by exponential transposition, as a functor $\mathcal{A} \to [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}]$. But the category of such functors is equivalent to the category of cocontinuous functors $[\mathcal{A}^{\mathrm{op}}, \mathsf{Set}] \to [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}]$: the two components of the equivalence being, respectively, restriction and left Kan extension along the Yoneda embedding $\mathsf{y}_{\mathcal{A}} : \mathcal{A} \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Set}]$ (since the Yoneda embedding is full and faithful this really does define an equivalence, cf. [56, p. 235]). Finally, by the Special Adjoint Functor Theorem there is a 1–1 correspondence between cocontinuous functors $[\mathcal{A}^{\mathrm{op}}, \mathsf{Set}] \to [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}]$ and continuous functors $[\mathcal{B}^{\mathrm{op}}, \mathsf{Set}] \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Set}]$ (every presheaf category is well-powered, co-well-powered, has a generator and a cogenerator). Thus, for each pair of small categories $\mathcal{A}$ and $\mathcal{B}$, there is a functor $\Delta_{\mathcal{AB}}$ giving an equivalence between $\mathsf{Mod}(\mathcal{A}, \mathcal{B})$ and the category of continuous functors $[\mathcal{B}^{\mathrm{op}}, \mathsf{Set}] \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Set}]$ and natural transformations between them. In fact, we obtain a biequivalence of bicategories Δ from Mod to the bicategory of presheaf categories and continuous functors. We don't need to verify this last fact, but we do need an explicit calculation of $\Delta_{\mathcal{AB}}$.

Given a bimodule $\phi : \mathcal{A} \nrightarrow \mathcal{B}$, writing $\overline{\phi}$ for the functor $\lambda a\, \lambda b\, \phi(b, a) : \mathcal{A} \to [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}]$, consider $\mathrm{Lan}_{\mathsf{y}_{\mathcal{A}}}\overline{\phi}$, the left Kan extension of $\overline{\phi}$ along the Yoneda embedding $\mathsf{y}_{\mathcal{A}} : \mathcal{A} \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Set}]$. This is a cocontinuous functor which is given by the formula

$$\mathrm{Lan}_{\mathsf{y}_{\mathcal{A}}}\overline{\phi}(P) = \mathrm{Colim}(\mathrm{Elts}(P) \xrightarrow{U} \mathcal{A} \xrightarrow{\overline{\phi}} [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}])$$

for a presheaf $P : \mathcal{A}^{\mathrm{op}} \to \mathsf{Set}$, where $\mathrm{Elts}(P)$ is the comma category $(1 \downarrow P)$.

Fixing a presheaf $Q : \mathcal{B}^{\mathrm{op}} \to \mathsf{Set}$, a morphism $\mathrm{Lan}_{\mathsf{y}_{\mathcal{A}}}\overline{\phi}(P) \Rightarrow Q$ corresponds to a cocone from the diagram $\overline{\phi} \cdot U$ to Q. The data for such a cocone is, for each pair $(a, x) \in \mathrm{Elts}(P)$, a choice of a natural transformation $\alpha_{(a,x)} : \overline{\phi}(a) \Rightarrow Q$—this choice being natural in (a, x). This amounts to a morphism $P \Rightarrow \Delta_{\mathcal{AB}}(\phi)(Q)$, where the functor

$$\Delta_{\mathcal{AB}}(\phi) : [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}] \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Set}]$$

is defined by

$$\Delta_{\mathcal{AB}}(\phi)(Q)(a) = [\mathcal{B}^{\mathrm{op}}, \mathsf{Set}](\phi(-, a), Q),$$

and is by this definition right adjoint to $\mathrm{Lan}_{\mathsf{y}_{\mathcal{A}}}\overline{\phi}$.

We are now ready to introduce the construction which underlies the proof of the main result in this section, corollary 5.2.3.

Suppose $\mathcal{A}$ is a small category and $\phi : \mathcal{A} \nrightarrow \mathcal{A}$. Let $\mathcal{G}$ be the graph with

(i) nodes: the set of objects of $\mathcal{A}$;

(ii) edges: for each arrow $f\colon a \to b$ of $\mathcal{A}$ an edge $f\colon a \to b$ of $\mathcal{G}$, and, for each pair of objects a, b of $\mathcal{A}$ and each $e \in \phi(b, a)$, an edge $e\colon b \to a$ of $\mathcal{G}$.

From the graph $\mathcal{G}$ we freely generate a category, which we denote $\mathcal{C}$, subject to the following equations on composition in $\mathcal{C}$ (written as $\cdot_{\mathcal{C}}$).

(a) For composable morphisms f, g of $\mathcal{A}$, $f \cdot_{\mathcal{C}} g = f \cdot g$;

(b) If $e \in \phi(b, a)$ and $f\colon b' \to b$ is an arrow of $\mathcal{A}$, then $e \cdot_{\mathcal{C}} f = \phi(f, a)e$;

(c) If $e \in \phi(b, a)$ and $f\colon a \to a'$ is an arrow of $\mathcal{A}$, then $f \cdot_{\mathcal{C}} e = \phi(b, f)e$.

A graph homomorphism $P : \mathcal{G}^{op} \to \mathsf{Set}$ consists of a graph homomorphism $P_0 : \mathcal{A}^{op} \to \mathsf{Set}$ plus a family of mappings, indexed over pairs of objects $a, b \in \mathcal{A}$,

$$\alpha(b, a) : \phi(b, a) \to P_0(b)^{P_0(a)}.$$

By exponential transposition this last datum amounts to a family of mappings,

$$\overline{\alpha}(b, a) : P_0(a) \to P_0(b)^{\phi(b,a)}.$$

The graph homomorphism P will be a functor if it preserves identities in $\mathcal{C}$ and the three types of composition (a-c) above. Preservation of identities and composites of type (a) is equivalent to P_0 being a functor $\mathcal{A}^{op} \to \mathsf{Set}$. Given this, P preserves composites of type (b) precisely when, for each $x \in P_0(a)$, $\overline{\alpha}(-, a)x$ is a natural transformation $\phi(-, a) \to P_0$, i.e.,

$$\overline{\alpha}(-, a) : P_0(a) \to \Delta(\phi)(P_0)(a).$$

In addition, preservation of composites of type (c) is the same as requiring that the above family of maps is natural in $a \in \mathcal{A}$. Thus we have shown that a presheaf P on $\mathcal{C}$ amounts to a pair $(P_0, \overline{\alpha})$, where P_0 is a presheaf on $\mathcal{A}$ and

$$\overline{\alpha} : P_0 \to \Delta(\phi)(P_0)$$

is a natural transformation.

Let us suppose we have another presheaf Q on $\mathcal{C}$, consisting of a presheaf Q_0 on $\mathcal{A}$ and a family of maps $\beta(b, a) : \phi(b, a) \to Q_0(b)^{Q_0(a)}$, indexed by objects $b, a \in \mathcal{A}$. A natural transformation $\Xi : P \Rightarrow Q$ is precisely a natural transformation $\Xi_0 : P_0 \Rightarrow Q_0$ such that the left hand diagram, below, commutes for each pair of objects $a, b \in \mathcal{A}$. But, by exponential transposition, this is just the same as requiring that the right hand diagram commutes.

$$\begin{tikzcd} \phi(b,a) \arrow[r, "\alpha(b,a)"] \arrow[d, "\beta(b,a)"'] & P_0(b)^{P_0(a)} \arrow[d, "{\Xi_{0,b}}^{P_0(a)}"] \\ Q_0(b)^{Q_0(a)} \arrow[r, "Q_0(b)^{\Xi_{0,a}}"'] & Q_0(b)^{P_0(a)} \end{tikzcd} \qquad \begin{tikzcd} P_0 \arrow[r, "\overline{\alpha}"] \arrow[d, "\Xi_0"'] & \Delta(\phi)(P_0) \arrow[d, "\Delta(\phi)(\Xi_0)"] \\ Q_0 \arrow[r, "\overline{\beta}"'] & \Delta(\phi)(Q_0) \end{tikzcd} \tag{5.5}$$

It is now clear that we have an isomorphism of categories between $[\mathcal{C}^{op}, \mathsf{Set}]$ and the category of coalgebras of $\Delta(\phi)$. Since any continuous endofunctor $T : [\mathcal{A}^{op}, \mathsf{Set}] \to [\mathcal{A}^{op}, \mathsf{Set}]$ is of the form $\Delta(\phi)$, for some ϕ, we obtain the following result.

Theorem 5.2.2 Let T be a continuous endofunctor on a presheaf category; then $\mathsf{Coalg}(T)$ is itself a presheaf category.

Corollary 5.2.3 Suppose T is a wide-pullback preserving endofunctor on a presheaf category; then $\mathsf{Coalg}(T)$ is itself a presheaf category.

Proof. Proposition 5.1 tells us that there is a final T-coalgebra (A, α). From proposition 5.1.7 it follows that $\mathsf{Coalg}(T) \cong \mathsf{Coalg}(T)/(A, \alpha) \cong \mathsf{Coalg}(T_\alpha)$. Since T_α is continuous, and the property of being a presheaf category is preserved by taking slices [43, Corl 2.18], the result follows. $\square$

5.3 Cofree coalgebras and comonads

Generalizing the concept of an initial F-algebra, we have the notion of a free F-algebra over an object X. The paradigm example arises when F is the functor T_Σ arising from an algebraic signature Σ (cf. section 2.2.1); then the free T_Σ-algebra over a set X is the set of Σ-terms with variables in X. Dually, we have the notion of a cofree coalgebra over an object X; this time X is envisioned as an object of 'colours', cf. Rutten [73].

Definition 5.3.1 Given an endofunctor $F \colon \mathcal{E} \to \mathcal{E}$ and an object X of $\mathcal{E}$, an F-coalgebra (S_X, σ_X) together with a 'colouring' map $\varepsilon_X : S_X \to X$ is *cofree over* X if, for any other F-coalgebra (A, α) and colouring $f \colon A \to X$, there exists a unique coalgebra map $g \colon A \to S_X$ making the diagram below commute. $\square$

$$\begin{tikzcd} A \arrow[d, dashed, "g"'] \arrow[dr, "f"] & \\ S_X \arrow[r, "\varepsilon_X"'] & X \end{tikzcd} \tag{5.6}$$

Notice that, assuming $\mathcal{E}$ has binary products, a cofree F-coalgebra over an object X is nothing but a final coalgebra of the endofunctor $F_X = F(-) \times X$. Explicitly, the F_X-coalgebra structure on S_X is the tuple $\langle \sigma_X, \varepsilon_X \rangle$. In particular, a coalgebra is final iff it is cofree over 1.

Example 5.3.2 In case F is the finite-powerset functor $\wp\colon \mathsf{Set} \to \mathsf{Set}$, we can modify the treatment of the final F-coalgebra in section 3.6 to describe a cofree F-coalgebra over a set A. The carrier S_A is the set of finitely-branching trees of possibly infinite depth, with nodes labelled by elements of A, and which are *strongly extensional as A-labelled trees.* We say that an A-labelled tree t is strongly extensional if, viewing t as an unlabelled transition system, no two children of any node are related by a label-preserving strong bisimulation (i.e. a strong bisimulation which only relates nodes with the same label). The colouring map $\varepsilon_A\colon S_A \to A$ maps each tree to the label of its root node. Given any A-labelled tree, the quotient by the largest depth-preserving label-preserving strong bisimulation (i.e. a strong bisimulation which only relates nodes with the same depth and label) is still an A-labelled tree and is strongly extensional. □

Write U for the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$, mapping each coalgebra to its carrier. The existence of a cofree F-coalgebra (S_X, σ_X) over each object X of $\mathcal{E}$ is equivalent to the existence of a right adjoint $R\colon \mathcal{E} \to \mathsf{Coalg}(F)$ to the forgetful functor U. This follows from the Pointwise Adjointness Theorem, cf. [15, Thm 9.3.5]. The action of R on objects is $X \mapsto S_X$; the action of R on arrows is determined by the universal property of (S_X, σ_X). The colouring maps ε_X now determine a natural transformation $\varepsilon\colon UR \Rightarrow 1$: the counit of the adjunction.

Next, we review a very general condition for the existence of cofree coalgebras, due to Barr [13]: an accessible!functor accessible functor on a locally presentable category has all cofree coalgebras. This condition is a simple consequence of a result of Makkai and Paré on accessible categories.[4] Suppose $F\colon \mathcal{E} \to \mathcal{E}$ is an endofunctor on a locally presentable category $\mathcal{E}$ (e.g. a Grothendieck topos, cf. [58]), and let U denote the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$. Observe that there is a natural transformation $U \Rightarrow FU$, whose value at a coalgebra (A, α) is $\alpha\colon A \to FA$. In other words, $\mathsf{Coalg}(F)$ is an *oplax cone* over the diagram

$$ {}^{F}\circlearrowright \mathcal{E} \tag{5.7}$$

[4]The definition of a κ-accessible category is obtained from the definition of a locally κ-presentable category, definition 3.2.4, by requiring only the existence of κ-filtered colimits rather than all colimits.

in the meta-2-category CAT of categories, functors and natural transformations. In fact, it is final amongst such oplax cones. But [58, Thm 5.1.8] says that the oplax limit in CAT of a diagram of accessible categories and functors is itself accessible. Thus, if F is an accessible functor, Coalg(F) is both accessible and cocomplete (the forgetful functor U creates colimits), i.e. it is locally presentable. Being locally presentable, Coalg(F) has a set of generators and is co-well-powered, cf. [8], so it follows from the Special Adjoint Functor Theorem [17] that U has a right adjoint.

Remark 5.3.3 If $\mathcal{E}$ is a presheaf category and $T\colon \mathcal{E} \to \mathcal{E}$ preserves wide pullbacks, then, by corollary 5.2.3, Coalg(T) is an oplax limit of (5.7) regarded as a diagram in Pshf. Since Δ from section 5.2 defines a biequivalence of Mod and Pshf (and thus preserves oplax limits) it follows that the 'path category' constructed in section 5.2 is an oplax limit of a diagram in Mod of the same shape as (5.7). □

Definition 5.3.4 A *comonad* (or *cotriple*) over a category $\mathcal{E}$ is an ordered triple $\mathbb{G} = (G, \varepsilon, \delta)$ consisting of a functor $G\colon \mathcal{E} \to \mathcal{E}$, and natural transformations $\varepsilon : G \Rightarrow 1$ and $\delta\colon G \Rightarrow G^2$, such that $G\delta \cdot \delta = \delta G \cdot \delta$ and $G\varepsilon \cdot \delta = \varepsilon G \cdot \delta = 1$. □

Definition 5.3.5 Suppose $\mathbb{G} = (G, \varepsilon, \delta)$ is a comonad on a category $\mathcal{E}$. A G-coalgebra (A, α) is called a $\mathbb{G}$-*coalgebra* (i.e., a coalgebra for the comonad $\mathbb{G}$) if the following two diagrams commute:

$$\begin{tikzcd} A \arrow[r, "\alpha"] \arrow[d, "\alpha"'] & GA \arrow[d, "G\alpha"] \\ GA \arrow[r, "\delta_A"'] & G^2A \end{tikzcd} \qquad \begin{tikzcd} GA \arrow[r, "\varepsilon_A"] & A \\ A \arrow[u, "\alpha"] \arrow[ur, equal] & \end{tikzcd} \tag{5.8}$$

A coalgebra for a comonad is sometimes called an *Eilenberg-Moore coalgebra*. The full subcategory of Coalg(G) consisting of the $\mathbb{G}$-coalgebras is denoted Coalg($\mathbb{G}$). The forgetful functor $U_{\mathbb{G}}\colon$ Coalg($\mathbb{G}$) $\to \mathcal{E}$ has a right adjoint $R_{\mathbb{G}}$ mapping C in $\mathcal{E}$ to (GC, δ_C). □

The adjunction

$$\mathsf{Coalg}(F) \underset{R}{\overset{U}{\rightleftarrows}} \mathcal{E} \qquad\qquad U \dashv R \tag{5.9}$$

generates a comonad $\mathbb{G} = (G, \varepsilon, \delta)$ in the standard way: i.e., writing η and ε for the unit and counit of the adjunction, we set $G = UR$ and $\delta = U\eta R\colon UR \Rightarrow URUR$; cf. [56, p. 134]. Also, from the general theory of comonads, we have a 'comparison functor' $K\colon$ Coalg(F) $\to$ Coalg($\mathbb{G}$) defined by

$$K(A, \alpha) = (A, U\eta_{(A,\alpha)}\colon A \to GA)$$

and making the diagram

$$\begin{tikzcd} \mathsf{Coalg}(F) \arrow[r, "K"] \arrow[d, shift left, "U"] & \mathsf{Coalg}(\mathbb{G}) \arrow[d, shift left, "U_{\mathbb{G}}"] \\ \mathcal{E} \arrow[u, shift left, "R"] \arrow[r, equal] & \mathcal{E} \arrow[u, shift left, "R_{\mathbb{G}}"] \end{tikzcd} \tag{5.10}$$

commute.

The forgetful functor U creates equalizers for those parallel pairs f, g in $\mathsf{Coalg}(F)$ for which Uf and Ug have an absolute equalizer[5] in $\mathcal{E}$ (suppose Uf and Ug have absolute equalizer E; since F preserves this equalizer it easily follows that there is a unique coalgebra structure on E making it an equalizer of f, g in $\mathsf{Coalg}(F)$). By (the dual of) Beck's theorem [56] this implies that the comparison functor K is an isomorphism—in other words, U is *comonadic*. Thus, we may view F-coalgebras as coalgebras (with the same carrier) for the corresponding comonad $\mathbb{G}$ and vice-versa.

Example 5.3.6 Let $\mathbb{G} = (G, \varepsilon, \delta)$ be the comonad generated by the finite-powerset functor $\wp$, so that GA is the set of strongly extensional A-labelled trees for each set A. The comparison functor K maps a $\wp$-coalgebra (A, α) to a $\mathbb{G}$-coalgebra $(A, \overline{\alpha})$, where $\overline{\alpha}(a)$ is coinductively defined to be the tree with root labelled a and maximal proper subtrees $\overline{\alpha}(a_1), ..., \overline{\alpha}(a_n)$ where $\alpha(a) = \{a_1, ..., a_n\}$. □

5.4 Elementary topos structure

In general, limits in categories of coalgebras are much less straightforward than colimits. However, there are some simple sub-cases; in particular, if $\mathcal{E}$ has all limits indexed by some category $\mathcal{I}$, and if $F\colon \mathcal{E} \to \mathcal{E}$ preserves $\mathcal{I}$-indexed limits, then the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$ creates $\mathcal{I}$-indexed limits. Similarly, for a comonad $\mathbb{G} = (G, \varepsilon, \delta)$ on $\mathcal{E}$, if $\mathcal{E}$ has $\mathcal{I}$-indexed limits and G preserves $\mathcal{I}$-indexed limits, then the forgetful functor $\mathsf{Coalg}(\mathbb{G}) \to \mathcal{E}$ creates $\mathcal{I}$-indexed limits.

Proposition 5.4.1 If F preserves weak pullbacks then the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$ also preserves weak pullbacks.

[5] An equalizer in a category $\mathcal{C}$ is absolute if it is preserved by all functors with domain $\mathcal{C}$.

Proof. Suppose that $f\colon (C,\gamma) \to (E,\varepsilon)$ and $g\colon (D,\delta) \to (E,\varepsilon)$ have weak pullback (A,α) with projections p_1 and p_2. Let P be the pullback of f and g in $\mathcal{E}$, as shown below.

$$\begin{tikzcd}
P \arrow[dr, shift left, "i"] \arrow[drr, bend left, "\pi_1"] \arrow[ddr, bend right, "\pi_2"'] & & \\
& A \arrow[ul, shift left, "e"] \arrow[r, "p_1"] \arrow[d, "p_2"'] & C \arrow[d, "f"] \\
& D \arrow[r, "g"] & E
\end{tikzcd}$$

By the universal property of P we have a comparison map $e\colon A \to P$. Since F preserves weak pullbacks, P can be given an F-coalgebra structure such that π_1 and π_2 are coalgebra maps. Thus, we obtain a comparison map $i\colon P \to A$. It follows that e is split epi, and so, by proposition 5.1.2, A is a weak pullback of f and g in $\mathcal{E}$. $\square$

It is pertinent to say something here about subcoalgebras. If (A,α) is an F-coalgebra, then by a *subcoalgebra* we mean a mono $m\colon (B,\beta) \rightarrowtail (A,\alpha)$ in the category $\mathsf{Coalg}(F)$. Note that the forgetful functor $U : \mathsf{Coalg}(F) \to \mathcal{E}$ need not preserve monos, see example 5.5.4, thus Um is not necessarily a mono in $\mathcal{E}$. However, if F preserves weak pullbacks, then, from the characterization of a mono as a map m making (5.1) a pullback, and the fact that the U preserves weak pullbacks (cf. prop. 5.4.1), it follows that U preserves monos.

On the other hand, if we start with a mono $m\colon B \to A$ in $\mathcal{E}$ then, assuming F preserves monos, there is only one possible coalgebra structure on B making it a subcoalgebra of (A,α); thus we identify the subcoalgebra with its carrier. This last statement also applies to coalgebras of a comonad $\mathbb{G} = (G,\varepsilon,\delta)$ where G preserves monos; furthermore, in this case, a sub-G-coalgebra of a $\mathbb{G}$-coalgebra is automatically a $\mathbb{G}$-coalgebra, cf. [57]. Finally, we observe that if (A,α) is a $\mathbb{G}$-coalgebra and $m\colon B \rightarrowtail A$ is a mono in $\mathcal{E}$, then B is a subcoalgebra iff the coalgebra structure $\beta\colon B \to GB$ on B actually makes

$$\begin{tikzcd}
A \arrow[r, "\alpha"] & GA \\
B \arrow[u, "m"] \arrow[r, dashed, "\beta"'] & GB \arrow[u, "Gm"']
\end{tikzcd} \tag{5.11}$$

a pullback in $\mathcal{E}$ (use the fact that $\varepsilon_A \cdot \alpha = 1_A$ and $\varepsilon_B \cdot \beta = 1_B$).

Recall that a parallel pair of maps $f, g\colon A \to B$ in a category is *coreflexive* if there is a map $h\colon B \to A$ such that $h \cdot f = h \cdot g = 1_A$. In particular, both maps in a coreflexive pair are split monos and their equalizer, should it exist, is also the pullback of one along the other, i.e. their intersection. Note that since all set functors preserve non-empty finite intersections (cf. [81]) they also preserve equalizers of coreflexive pairs. Now (the dual of) Linton's theorem on coproducts of algebras tells us how to construct products of coalgebras.

Theorem 5.4.2 (Linton.) Suppose $\mathbb{G} = (G, \varepsilon, \delta)$ is a comonad on category $\mathcal{E}$ with all finite limits. If G preserves equalizers of coreflexive pairs then $\mathsf{Coalg}(\mathbb{G})$ has all finite limits.

Proof. Since G preserves equalizers of coreflexive pairs, the forgetful functor $\mathsf{Coalg}(\mathbb{G}) \to \mathcal{E}$ creates them. But, by the dual of [55, Cor. 3], it follows that $\mathsf{Coalg}(\mathbb{G})$ has all limits, with the product of $\mathbb{G}$-coalgebras (A, α) and (B, β) being the subcoalgebra of the cofree coalgebra $(G(A \times B), \delta_{A\times B})$ determined by the equalizer of the (coreflexive) pair

$$\begin{tikzcd} G(A \times B) \arrow[r, "\delta_{A\times B}"'] \arrow[rr, bend left, "G(\alpha\times\beta)"] & G^2(A \times B) \arrow[r, "G\langle G\pi_1, G\pi_2\rangle"'] & G(GA \times GB) \end{tikzcd} \tag{5.12}$$

where π_1 and π_2 are the projections associated with the product $A \times B$. □

Remark 5.4.3 If, in theorem 5.4.2, we strengthen the hypotheses by assuming that G preserves pullbacks of monos, then we claim that the equalizer of the pair (5.12) is also the equalizer of the pair

$$\begin{tikzcd} G(A \times B) \arrow[r, "\varepsilon_{A\times B}"] \arrow[rr, bend right, "\langle G\pi_1, G\pi_2\rangle"'] & A \times B \arrow[r, "\alpha\times\beta"] & GA \times GB\,. \end{tikzcd} \tag{5.13}$$

By adjointness, a coalgebra map $(C, \gamma) \to (G(A \times B), \delta_{A\times B})$ equalizes the pair (5.12) iff it equalizes the pair (5.13); thus it will suffice to prove that the equalizer of (5.13) determines a subcoalgebra of $(G(A \times B), \delta_{A\times B})$.

Observe that, by naturality of ε, the pair (5.13) is the same as the pair

$$\begin{tikzcd} G(A \times B) \arrow[r, "\langle G\pi_1, G\pi_2\rangle"] \arrow[rrr, bend right, "\langle G\pi_1, G\pi_2\rangle"'] & GA \times GB \arrow[r, "\varepsilon_A\times\varepsilon_B"] & A \times B \arrow[r, "\alpha\times\beta"] & GA \times GB\,. \end{tikzcd} \tag{5.14}$$

Now, $\alpha \times \beta$ is the equalizer of $(\alpha \times \beta) \cdot (\varepsilon_A \times \varepsilon_B)$ and $1_{GA\times GB}$ since $\varepsilon_A \cdot \alpha = 1_A$ and $\varepsilon_B \cdot \beta = 1_B$. It follows that the equalizer of (5.14) is the pullback of $\alpha \times \beta$ along $\langle G\pi_1, G\pi_2\rangle$, i.e. the intersection of the pullback of α along $G\pi_1$ with the pullback of β along $G\pi_2$. But this is a subcoalgebra of $(G(A \times B), \delta_{A\times B})$ since the forgetful functor $U_{\mathbb{G}}$ creates pullbacks of monos. □

Example 5.4.4 We describe products in the category of $\wp$-coalgebras. Let $\mathbb{G} = (G, \varepsilon, \delta)$ be the comonad generated by the functor $\wp$. Suppose that (A, α) and (B, β) are $\wp$-coalgebras with corresponding $\mathbb{G}$-coalgebras $(A, \overline{\alpha})$ and $(B, \overline{\beta})$; their product is the subcoalgebra of the cofree coalgebra on the set $A \times B$ determined by the equalizer of the pair in (5.13). Thus, an $(A \times B)$-labelled tree $t \in G(A \times B)$—with root labelled (a, b)—lies in the product if: (i) $\overline{\alpha}(a)$ is the same as the A-labelled obtained by applying the first projection to each label of t and then taking the strongly extensional quotient; (ii) $\overline{\beta}(b)$ is the same as the B-labelled tree obtained by applying the second projection to each label of t and then taking the strongly extensional quotient. □

Definition 5.4.5 A *subobject classifier* in a category $\mathcal{E}$ consists of an object Ω and a map $t\colon 1 \to \Omega$, such that for each monomorphism $m\colon A' \to A$ in $\mathcal{E}$ there is a unique map $\chi_m\colon A \to \Omega$ (the classifying map of m) making

$$\begin{tikzcd} A' \arrow[r] \arrow[d, "m"'] & 1 \arrow[d, "t"] \\ A \arrow[r, "\chi_m"'] & \Omega \end{tikzcd} \tag{5.15}$$

a pullback diagram. □

The paradigm example of a subobject classifier is the two-element set of Booleans in Set. In that case, the classifying map of a subset is better known as its characteristic function.

Theorem 5.4.6 Suppose $\mathcal{E}$ has finite limits and a subobject classifier, $F\colon \mathcal{E} \to \mathcal{E}$ preserves weak pullbacks, and the forgetful functor $U\colon \mathsf{Coalg}(F) \to \mathcal{E}$ has a right adjoint R. Then $\mathsf{Coalg}(F)$ has a subobject classifier.

Proof. Write $\mathbb{G} = (G, \varepsilon, \delta)$ for the comonad generated by the adjunction $U \dashv R$. Proposition 5.4.1 tells us that U preserves weak pullbacks, thus $G(= UR)$ preserves weak pullbacks and U also preserves monos. It now remains to verify that the construction of a subobject classifier in $\mathsf{Coalg}(\mathbb{G})$ in the case that G preserves all finite limits [43, Thm 2.32] still works in the present context.

Let $t\colon 1 \to \Omega$ be the subobject classifier in $\mathcal{E}$. Write $\tau\colon G\Omega \to \Omega$ for the classifying map of the subobject $Gt\colon G1 \rightarrowtail G\Omega$ (G preserves monos). We claim that the subobject classifier in $\mathsf{Coalg}(\mathbb{G})$ is given by the equalizer of the following parallel pair in $\mathsf{Coalg}(\mathbb{G})$ (where all objects shown have the structure of cofree coalgebras).

$$\begin{tikzcd} G\Omega \arrow[r, "\delta_\Omega"'] \arrow[rr, bend left, "1_\Omega"] & G^2\Omega \arrow[r, "G\tau"'] & G\Omega \end{tikzcd} \tag{5.16}$$

Given a $\mathbb{G}$-coalgebra (A, α), we show a correspondence between subcoalgebras of (A, α) and maps $(A, \alpha) \to (G\Omega, \delta_\Omega)$ which equalize the pair above.

Recall that, since G preserves monos, a subcoalgebra of (A, α) is precisely a mono $m\colon B \rightarrowtail A$ in $\mathcal{E}$ such that there is a G-coalgebra structure β on B making (5.11) a pullback. In other words, m is a subcoalgebra iff $m = \alpha^*(Gm)$. In terms of classifying maps: $m\colon B \rightarrowtail A$ is a subcoalgebra iff χ_m also classifies $\alpha^*(Gm)$.

Write $\overline{\chi_m}\colon A \to G\Omega$ for the cofree extension of χ_m. Then $\overline{\chi_m}$ equalizes the pair (5.16), i.e. $\overline{\chi_m} = G\tau \cdot \delta_\Omega \cdot \overline{\chi_m}$, if and only if $\overline{\chi_m} = G\tau \cdot \delta_\Omega \cdot G\chi_m \cdot \alpha$, since $\overline{\chi_m} = G\chi_m \cdot \alpha$. Taking transposes, this holds iff $\chi_m = \varepsilon_\Omega \cdot G\tau \cdot \delta_\Omega \cdot G\chi_m \cdot \alpha$. Next, by the naturality of ε and the comonad equation $\varepsilon_{G\Omega} \cdot \delta_\Omega = 1$ this is equivalent to $\chi_m = \tau \cdot G\chi_m \cdot \alpha$. But we claim that $\tau \cdot G\chi_m$ classifies the subobject Gm, hence this last says that χ_m classifies $\alpha^*(Gm)$.

It remains to verify the claim above. Observe that both squares in (5.17) are pullbacks; for the left-hand square this is because G preserves pullbacks of monos. Thus, the outer rectangle is a pullback and $\tau \cdot G(\chi_m)$ classifies the subobject Gm.

$$\begin{tikzcd} GB \arrow[r, "G(!_B)"] \arrow[d, "Gm"'] & G1 \arrow[r] \arrow[d, "Gt"] & 1 \arrow[d, "t"] \\ GA \arrow[r, "G\chi_m"'] & G\Omega \arrow[r, "\tau"'] & \Omega \end{tikzcd} \tag{5.17}$$

□

Example 5.4.7 We describe the subobject classifier in the category of $\wp$-coalgebras. The proof of theorem 5.4.6 shows that it is the subobject of the cofree coalgebra on the set $\Omega = \{\mathrm{tt}, \mathrm{ff}\}$ determined by the equalizer of (5.16). The lower map in (5.16) acts on a tree t as follows: first, each node n of t is relabelled by the minimum from the set of labels of n and all the nodes descended from n, then one takes the strongly extensional quotient. Thus, the subobject classifier consists of those Ω-labelled trees such that if a node is labelled tt then all of its descendants are also labelled tt—'once true always true'. This is reasonable, because if a state of a $\wp$-coalgebra lies in a subcoalgebra then so do all its successor states.

The classifying map of a subobject $(A', \alpha') \rightarrowtail (A, \alpha)$ maps $a \in A$ to the Ω-labelled tree obtained from $\overline{\alpha}(a)$—as defined in example 5.3.6—by first relabelling each node with the truth value of the assertion that the original label of the node lies in A', and then taking the strongly extensional quotient. □

As we explained in the introduction, the following definition provides a link between the work contained in section 5.2 and the work in the present section.

Definition 5.4.8 A category $\mathcal{E}$ is an *elementary topos* (henceforth, just a topos) if it possesses finite limits, a subobject classifier and exponentials. □

A central result of topos theory asserts that if $\mathbb{G} = (G, \varepsilon, \delta)$ is a comonad on a topos, and if G preserves all finite limits, then $\mathsf{Coalg}(\mathbb{G})$ is itself a topos; cf. [43, Thm 2.32]. In the joint paper [44] the result is generalized by weakening the hypothesis that G preserves all finite limits to G preserves pullbacks. This is achieved by a 'reduction' of a comonad $\mathbb{G}$

whose functor part preserves pullbacks to a comonad $\mathbb{G}'$ whose functor part preserves all finite limits such that $\mathsf{Coalg}(\mathbb{G}) \cong \mathsf{Coalg}(\mathbb{G}')$. We give a direct proof of this last result in appendix A.

If $\mathcal{E}$ is a topos and $F\colon \mathcal{E} \to \mathcal{E}$ preserves pullbacks then the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$ preserves pullbacks. It follows that if U has a right adjoint then, in the generated comonad $\mathbb{G} = (G, \varepsilon, \delta)$, G also preserves pullbacks. Thus, in this case, $\mathsf{Coalg}(F)$ is a topos by theorem A.2.2. In short, a smallness condition plus a limit preservation condition on the functor F imply that $\mathsf{Coalg}(F)$ is a topos.

5.5 Coalgebras as sheaves

In the remaining part of this chapter we investigate what can be said if we weaken the limit preservation assumption to 'F preserves weak pullbacks'. Recall that we already know that if F preserves weak pullbacks, and if all cofree coalgebras exist, then $\mathsf{Coalg}(F)$ has finite limits and a subobject classifier. Our analysis is based on version of Giraud's theorem (cf. [57, p. 591]) characterizing Grothendieck toposes in terms of so-called exactness conditions and a smallness condition. First let us explain some of the relevant terminology.

A coproduct $C = \coprod_\alpha C_\alpha$ in $\mathcal{E}$ is said to be *disjoint* if, for each $\alpha \neq \beta$, the pullback of the coproduct injections $i_\alpha\colon C_\alpha \to C$ and $i_\beta\colon C_\beta \to C$ is the initial object in $\mathcal{E}$. The same coproduct is said to be *stable under pullbacks* if, for each map $B \to C$, the pullbacks $B \times_C C_\alpha \to B$ of the coproduct injections $i_\alpha\colon C_\alpha \to C$ yield an isomorphism $\coprod_\alpha (B \times_C C_\alpha) \cong B$.

If $\mathcal{E}$ is a category with finite limits then an *equivalence relation* on an object X is a relation $R \rightarrowtail X \times X$ satisfying categorical versions of the reflexivity, symmetry and transitivity axioms, cf. definition A.0.1. If $f\colon X \to Y$ is any morphism then the kernel pair of f is always an equivalence relation on X. We say that an equivalence relation is *effective* if it arises as the kernel pair of some morphism. If $\mathcal{E}$ has coequalizers, then an equivalence relation is effective iff it is the kernel pair of the coequalizer of its projections (i.e. its quotient).

Theorem 5.5.1 (Giraud) A category $\mathcal{E}$ with small hom-sets and all finite limits is a Grothendieck topos iff it has the following properties:

1. $\mathcal{E}$ has small coproducts, they are disjoint and stable under pullback;
2. every epi in $\mathcal{E}$ is the coequaliser of its kernel pair;
3. epis in $\mathcal{E}$ are stable under pullback;
4. every equivalence relation in $\mathcal{E}$ is effective and has a quotient;

5. there is a small set of objects on $\mathcal{E}$ which generate $\mathcal{E}$. $\square$

Proposition 5.5.2 Suppose that $\mathcal{E}$ is a Grothendieck topos and $F\colon \mathcal{E} \to \mathcal{E}$ is an accessible functor preserving weak pullbacks. Then $\mathsf{Coalg}(F)$ satisfies all the hypotheses of Giraud's theorem except, possibly, effectiveness of equivalence relations.

Proof. 1. $\mathsf{Coalg}(F)$ has all coproducts; indeed, the forgetful functor U creates them. Also, the coproduct injections are monos (in both $\mathcal{E}$ and $\mathsf{Coalg}(F)$) since they are always monos in a topos. Since U also creates pullbacks of monos, and coproducts are disjoint in $\mathcal{E}$, it follows that coproducts are also disjoint in $\mathsf{Coalg}(F)$. Similarly, these same limit/colimit creation properties of the forgetful functor ensure that the property of having colimits which are stable under pullbacks lifts from $\mathcal{E}$ to $\mathsf{Coalg}(F)$.

2. Suppose $q\colon (A,\alpha) \to (B,\beta)$ is epi in $\mathsf{Coalg}(F)$, with kernel pair π_1 and π_2. Let p_1 and p_2 be the kernel pair of q in $\mathcal{E}$; then there is a comparison map $i\colon P \to C$ since U preserves weak pullbacks—see below. Now, q is the coequalizer of p_1 and p_2 in $\mathcal{E}$ since the latter category satisfies all the hypotheses of Giraud's theorem. It easily follows that q is the coequalizer of π_1 and π_2 in $\mathcal{E}$. But U reflects coequalizers, so q is also the coequalizer of π_1 and π_2 in $\mathsf{Coalg}(F)$.

$$\begin{tikzcd}
P \arrow[dr, "i"] \arrow[drr, bend left, "p_1"] \arrow[ddr, bend right, "p_2"'] & & \\
& C \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & A \arrow[d, "q"] \\
& A \arrow[r, "q"'] & B
\end{tikzcd} \tag{5.18}$$

3. Let the square below be a pullback, and suppose that $q\colon (A,\alpha) \to (B,\beta)$ is an epi.

$$\begin{tikzcd}
(E,\varepsilon) \arrow[r, "\pi_1"] \arrow[d, "\pi_2"'] & (A,\alpha) \arrow[d, "q"] \\
(C,\gamma) \arrow[r, "f"'] & (B,\beta)
\end{tikzcd} \tag{5.19}$$

Since U preserves weak pullbacks we find that π_2 factors through the pullback of q along f in $\mathcal{E}$ via a split epi. Since the pullback q along f in $\mathcal{E}$ is epi, it follows that π_2 is epi in $\mathcal{E}$ and hence also in $\mathsf{Coalg}(F)$.

5. We have already remarked on the existence of a set of generators in section 5.3. $\square$

Example 5.5.3 Here is an example of a non-effective equivalence relation in the category of $\wp$-coalgebras (unlabelled transitions systems). Consider the $\wp$-coalgebra (X,α), depicted

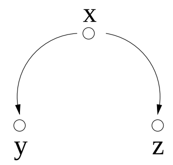

Figure 5.1: A transition system

as a transition system in fig. 5.1. The product of (X, α) with itself is depicted in fig. 5.2. The states of this transition system are drawn as rectangles, and the transitions are drawn as arrows between rectangles. The rectangles themselves are labelled with $(X \times X)$-labelled trees (recall the description of products as subcoalgebras of cofree coalgebras in example 5.4.4).

Let us take the equivalence relation (R, ρ) on (X, α) to be the subcoalgebra of the product consisting of all the states drawn as solid rectangles. Then the quotient of (R, ρ) is the map $q\colon (X, \alpha) \to (X, \alpha)/(R, \rho)$ identifying y and z. But the pullback of q with itself is the whole product $(X, \alpha) \times (X, \alpha)$. □

Intuitively, equivalence relations failed to be effective in the example above because bisimilar pairs of elements could appear 'with multiplicity' in the product. In particular, the product of two coalgebras was not the greatest bisimulation between them.

Example 5.5.4 Gumm and Schröder [39] show the functor $F\colon \mathsf{Set} \to \mathsf{Set}$ from example 5.1.4 preserves jointly monic spans (also called mono sources), and that, as a consequence, the product of two coalgebras is just the greatest bisimulation between them. However, F does not preserve weak pullbacks and $\mathsf{Coalg}(F)$ fails to satisfy a number of the Giraud hypotheses. Crucially, the forgetful functor $\mathsf{Coalg}(F) \to \mathsf{Set}$ does not preserve monos.

Consider the F-coalgebra (R, ρ), where $R = \{r, s\}$, $\rho(r) = (r, r, s)$, and $\rho(s) = (r, s, s)$. The carrier of the final F-coalgebra is a singleton, so the unique map f from (R, ρ) to this coalgebra is an epimorphism. Now it is clear that $r \in R$ is not bisimilar to $s \in R$, thus the greatest bisimulation on (R, ρ) is the identity relation. It follows that the kernel pair of f in $\mathsf{Coalg}(F)$ is the identity span $(R, \rho) \leftarrow (R, \rho) \to (R, \rho)$. But the coequalizer of the two projections in this span is just $1\colon (R, \rho) \to (R, \rho)$; thus f is not the coequalizer of its kernel pair. □

The examples above prompt the question of whether preservation of pullbacks by F is necessary for $\mathsf{Coalg}(F)$ to be a topos. A partial answer to this question is given in the next theorem which appears in the joint paper [44], but which is not due to the present author. The theorem applies to example 5.5.3, but not example 5.5.4.

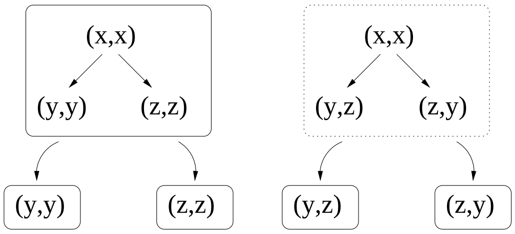

Figure 5.2: A non-effective equivalence relation

Theorem 5.5.5 Let $\mathcal{E}$ be a topos, $F\colon \mathcal{E} \to \mathcal{E}$ a functor which preserves weak pullbacks, and assume that the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$ has a right adjoint. Suppose F decomposes as $1 + G$ for some functor G and where 1 denotes the constant functor with value 1. Then $\mathsf{Coalg}(F)$ is a topos iff F preserves pullbacks. $\square$

Recall that a category $\mathcal{E}$ is said to be *regular* if it has finite limits, strong-epi mono factorizations of maps, and if strong epis are stable under pullback.

Proposition 5.5.6 Suppose that $\mathcal{E}$ is a regular category, $F\colon \mathcal{E} \to \mathcal{E}$ preserves weak pullbacks, and assume that the forgetful functor $\mathsf{Coalg}(F) \to \mathcal{E}$ has a right adjoint. Then $\mathsf{Coalg}(F)$ is regular.

Proof. By theorem 5.4.2 $\mathsf{Coalg}(F)$ has finite limits. Let $f\colon (A,\alpha) \to (B,\beta)$ be a homomorphism of F-coalgebras, and suppose that f has image factorization $A \xrightarrow{q} C \stackrel{i}{\rightarrowtail} B$ in $\mathcal{E}$. Now Fi is a mono and q is a strong epimorphism, so there exists a unique map $\gamma\colon C \to FC$ making (5.20) commute. It is easy to see that q is a strong epi in $\mathsf{Coalg}(F)$; thus q and i provide an image factorization for f regarded as a map in $\mathsf{Coalg}(F)$.

$$\begin{tikzcd} A \arrow[r, "q"] \arrow[d, "\alpha"] & C \arrow[r, "i"] \arrow[d, "\gamma"] & B \arrow[d, "\beta"] \\ FA \arrow[r, "Fq"] & FC \arrow[r, "Fi"] & FB \end{tikzcd} \tag{5.20}$$

Recalling that split epis are strong, and that the composition of two strong epis is also a strong epi, the stability of strong epis under pullbacks in $\mathsf{Coalg}(F)$ follows exactly as in the proof of (3) in proposition 5.5.2. $\square$

Given a regular category $\mathcal{C}$, McLarty [59] describes an embedding of $\mathcal{C}$ into a regular category with effective equivalence relations, denoted $\mathsf{Map}(\mathcal{C})$. The construction is universal in the sense that any regular functor (i.e. functor preserving finite limits and image factorizations) from $\mathcal{C}$ to a regular category with effective equivalence relations must factor through the embedding $I\colon \mathcal{C} \to \mathsf{Map}(\mathcal{C})$.

The objects of $\mathsf{Map}(\mathcal{C})$ are pairs (A, R) where A is an object of $\mathcal{C}$ and $R \rightarrowtail A \times A$ is an equivalence relation. An arrow $F\colon (A, R) \to (A', R')$ in $\mathsf{Map}(\mathcal{C})$ is a relation $F \rightarrowtail A \times A'$ in $\mathcal{C}$ such that:

- F *respects* R *and* R', i.e. $F \cdot R \equiv F \equiv R' \cdot F$;
- F is *functional*, i.e. $R \leq F^{\circ} \cdot F$ and $F \cdot F^{\circ} \leq R'$.

Composition in $\mathsf{Map}(\mathcal{C})$ is the usual composition of relations in $\mathcal{C}$, and the identity on an object (A, R) is the relation R itself. The embedding $I\colon \mathcal{C} \to \mathsf{Map}(\mathcal{C})$ sends each object A of $\mathcal{C}$ to the pair (A, Δ_A), where $\Delta_A = \langle 1_A, 1_A \rangle$ is the diagonal relation on A, and each map $f\colon A \to B$ of $\mathcal{C}$ to the graph of f, i.e. the pullback of f and 1_B.

Corollary 5.5.7 Suppose that $\mathcal{E}$ is a Grothendieck topos and $F\colon \mathcal{E} \to \mathcal{E}$ is an accessible functor preserving weak pullbacks. Then $\mathsf{Map}(\mathsf{Coalg}(F))$ is a Grothendieck topos.

Proof. Suppose $\mathcal{C}$ is a regular category. Freyd and Scedrov [30] show that $\mathsf{Map}(\mathcal{C})$ inherits the property of having disjoint pullback-stable coproducts from $\mathcal{C}$. Furthermore, if $\langle \pi_1, \pi_2 \rangle\colon R \rightarrowtail A \times A$ is an equivalence relation in $\mathcal{C}$, then $R\colon (A, \Delta_A) \to (A, R)$ is the quotient of π_1 and π_2 in $\mathsf{Map}(\mathcal{C})$. It follows that the image under $I\colon \mathcal{C} \to \mathsf{Map}(\mathcal{C})$ of a set of generators for $\mathcal{C}$ is a set of generators for $\mathsf{Map}(\mathcal{C})$.

Applying the above in case $\mathcal{C}$ is $\mathsf{Coalg}(F)$ we see that $\mathsf{Map}(\mathsf{Coalg}(F))$ satisfies all the hypotheses of Giraud's theorem. $\square$

Since we have an inclusion I of $\mathsf{Coalg}(F)$ into a topos $\mathsf{Map}(\mathsf{Coalg}(F))$, and since we know that the former already has a subobject classifier, it is natural to ask whether this is preserved by I. In fact, the answer to this question, in the affirmative, is implicit in the following proposition.

Proposition 5.5.8 [59, Lemma 25.23] Suppose that $1 \overset{t}{\rightarrowtail} (A, R)$ is a mono in $\mathsf{Map}(\mathcal{E})$, and that for any mono $m\colon A \rightarrowtail B$ in $\mathcal{E}$ there is a unique arrow $F\colon I(B) \to (A, R)$ such that $I(m)$ is the pullback of $1 \overset{t}{\rightarrowtail} (A, R)$ along F. Then $1 \overset{t}{\rightarrowtail} (A, R)$ is a subobject classifier for $\mathsf{Map}(\mathcal{E})$. $\square$

This says that if $1 \overset{t}{\rightarrowtail} (A, R)$ classifies subobjects from $\mathcal{C}$ then it is the subobject classifier for $\mathsf{Map}(\mathcal{C})$. Now the embedding $I\colon \mathcal{C} \rightarrowtail \mathsf{Map}(\mathcal{C})$ preserves finite limits, and, being full and faithful, it also reflects finite limits. It follows that if $\mathcal{C}$ has a subobject classifier, the image under I still classifies $\mathcal{C}$ subobjects in $\mathsf{Map}(\mathcal{C})$. Thus, by the proposition, it is the subobject classifer for $\mathsf{Map}(\mathcal{C})$.

Chapter 6

Evaluation and future work

In this thesis we have investigated the theory of universal coalgebra, considered as a general framework for studying transition systems, bisimulations, and coinduction.

In chapter 3 we considered the iterative construction of final coalgebras of accessible endofunctors via their final sequences. We gave a condition ensuring the convergence of this sequence, and, for set functors, a bound on number of steps needed for convergence. This led us to identify a simple relationship between existing set-theoretic and metric space models of non-deterministic and probabilistic processes.

In chapter 4 we defined a notion of simulation for coalgebras of endofunctors on categories of preorders and generalized metric spaces, and we proved a coinduction theorem characterizing the order (respectively, the metric) on a final coalgebra as the maximum simulation. In application, this led to a coinductive characterization of a known metric on the class of probabilistic processes. Also we indentified an easily verified condition on endofunctors, preservation of strong monos, ensuring that simulations compose and that the final semantics is abstract with respect to similarity.

Finally, chapter 5 was concerned with the structure of categories of coalgebras. We described products, subobject classifiers and exponentials in this category; we gave conditions under which it is a category of presheaves, a category of sheaves, and an elementary topos; finally we gave a couple of examples where the category of coalgebras fails to be a topos.

Below we present, chapter by chapter, an evaluation of our results and suggestions for future research.

6.1 Final sequences

The example in section 3.6 shows that the $\omega + \omega$ bound on the convergence of the final sequence of a finitary set functor is tight. However, we do not have an example of an ω_1-accessible set functor whose terminal sequence takes fully $\omega_1 + \omega_1$ steps to converge. As we mentioned earlier, if the functor preserves wide pullbacks of monos then one has convergence in $\omega_1 + \omega$ steps; this case seems to cover all of the 'reasonable' ω_1-accessible set functors one comes across, e.g. the countable-powerset functor.

Recently, Adámek [5] has posed the question of whether the behaviour one observes of the final sequence of a finitary set functor extends to finitary endofunctors on other locally finitely presentable categories. He obtains a partial answer to this question, showing that on certain locally finitely presentable categories, like $\mathsf{Preord}, \mathsf{Vect}$, and Graph, finitary endofunctors preserving strong monos and bimorphisms also have final sequences which converge in $\omega + \omega$ steps. The proof of this result uses theorem 3.4.7. In the absence of any counter-examples, the question of whether one can drop any of the side conditions on the functor remains open. As it stands, the result already applies, for instance, to endofunctors on Preord built up from constants, sums, products and finite-powerset (with either the Smyth, Hoare, or Egli-Milner preorders).

6.2 Metric coinduction

It is worth remarking that our coinduction principle for metric spaces used generalized metric spaces in an essential way, since the collage of a bimodule between (proper) metric spaces is typically only a generalized metric space. We see this as providing some evidence of the utility of the enriched categorical view of metric spaces—at least for the limited purposes of metric semantics. In this respect the most interesting examples are those to do with probabilistic bisimilarity, since the collage construction is visible even in the elementary characterizations of probabilistic bisimulations. The second of these examples also introduced some interesting connections between the 'generalized logics' of Lawvere [53] and Kozen [49].

A general task for future work is to reexamine some of the issues of the paper of Desharnais et. al. [25] in the light of our coinductive presentation of their metric on processes (as opposed to the inductive presentation via functional expressions). For instance, their calculations of the distance between processes involve defining formula, often by recursion, to witness that two processes are a particular distance apart, cf. example 4.8.9. On the other hand, the coinductive calculation of the same quantity (in simple cases at least) seems to boil down to some case analysis and the solution of inequalities, cf. example 4.8.7. In fact, example 4.8.9 corrects [25, example 6.7]. We were alerted to this by the coinductive calculation of the same quantity.

We also remind the reader of the question, raised in remark 4.8.16, about the necessity of including the negation construct in the grammar for the set of functional expressions.

6.3 Coalgebras, sheaves and toposes

Two possibilities for future work are immediate. One would like to see if it is possible to strengthen theorem 5.5.5, concerning sufficient and necessary conditions for $\mathsf{Coalg}(F)$ to be a topos, by removing the assumption that F decomposes as $1 + G$ for some G. For instance, such an extension would apply to the functor sending each set to the collection of its finite non-empty subsets (whose category of coalgebras fails to have effective equivalence relations). Also, an 'elementary' version of corollary 5.5.7 would be desirable, i.e. a version applying to elementary toposes, avoiding the set-theoretic concepts involved in the notion of a Grothendieck topos.

More speculatively, we would like to pin down the connections between the notion of a sheaf as a variable set—we are thinking of temporal variation—and the intuition of a coalgebra as a transition system. One starting point for such an investigation is with the work of Ghilardi and Meloni [32] on an interpretation of modal logic based on presheaves rather than Kripke models. They consider a temporal logic with two modal operators, respectively interpreted as 'generated sub-presheaf' and 'co-generated sub-presheaf'.

Let $P\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ be a presheaf on a small category $\mathcal{C}$. By a predicate ϕ on P we mean a family of sets (ϕ_C), indexed by the set of objects $\mathcal{C}_0$ of $\mathcal{C}$, such that $\phi_C \subseteq P(C)$ for all $C \in \mathcal{C}_0$. We say that a predicate ϕ is a sub-presheaf of P if ϕ can be given the structure of functor such that the inclusions $(\phi(C) \hookrightarrow P(C))_{C \in \mathcal{C}_0}$ form a natural transformation.

For each predicate ϕ we have the co-generated sub-presheaf $\underline{\phi}$, i.e. the maximum sub-presheaf of P contained in ϕ. This is given by the formula

$$\underline{\phi}(C) = \{x \ : \forall B \in \mathcal{C}_0, \forall f\colon B \to C.\ \ P(f)(x) \in \phi(B)\}.$$

We also have the generated sub-presheaf $\overline{\phi}$, i.e. the minimum sub-presheaf of P containing ϕ. This is given by the formula

$$\overline{\phi}(C) = \{y \ : \exists B \in \mathcal{C}_0, \exists g\colon C \to B, \exists x \in \phi(B).\ \ y = P(g)(x)\}.$$

Suppose T is a wide-pullback preserving set functor and $\mathcal{C}$ is the category constructed in section 5.2 such that $\mathsf{Coalg}(T)$ and $[\mathcal{C}^{\mathrm{op}}, \mathsf{Set}]$ are isomorphic. Examining the details of this isomorphism we find that if a presheaf P corresponds to a coalgebra (A, α), then predicates ϕ on P are in 1-1 correspondence with predicates (i.e. subsets) S on A. Under this correspondence, the generated sub-presheaf $\overline{\phi}$ becomes the smallest subcoalgebra of (A, α) containing S, and the co-generated sub-presheaf $\underline{\phi}$ becomes the largest subcoalgebra of (A, α) contained in S. Jacobs [41] uses the concepts of generated and co-generated subcoalgebras as the basis for a temporal logic on coalgebras. It would therefore be of interest to compare his coalgebraic semantics for modal logics with the analysis of Ghilardi and Meloni.

It is possible to generalize the ideas of Ghilardi and Meloni to sheaves on a site $(\mathcal{C}, J)$. That is, for a predicate ϕ on a sheaf P we have a generated subsheaf $\overline{\phi}$ and a co-generated subsheaf $\underline{\phi}$. We would like to see if these correspond to generated and co-generated subcoalgebras under the coalgebras-as-sheaves correspondence drawn in section 5.5 for coalgebras of weak-pullback preserving functors. In general, a Grothendieck topos is equivalent to a category of sheaves on many different sites, and it seems to us that the key to solving this problem is to find right the site for $\mathsf{Map}(\mathsf{Coalg}(T))$, i.e. a site such that predicates on sheaves correspond to predicates on coalgebras. We speculate that the gluing condition for sheaves on this site will be related to the reconciliation of the different possible evolutions of a non-deterministic process.

Appendix A

Sheaves and toposes

Definition A.0.1 Let $\mathcal{E}$ be a category with finite limits. A relation $\langle \pi_1, \pi_2 \rangle \colon R \rightarrowtail X \times X$ is an *equivalence relation* on X if

(i) R is reflexive, i.e. the diagonal subobject $\Delta_X \colon X \rightarrowtail X \times X$ factors through $\langle \pi_1, \pi_2 \rangle$.

(ii) R is symmetric, i.e. there exists a map $\tau \colon R \to R$ such that $\pi_1 \cdot \tau = \pi_2$ and $\pi_2 \cdot \tau = \pi_1$.

(iii) R is transitive, i.e. in the pullback

$$\begin{tikzcd} T \arrow[r, "q"] \arrow[d, "p"'] & R \arrow[d, "\pi_1"] \\ R \arrow[r, "\pi_2"'] & X \end{tikzcd} \tag{A.1}$$

the tuple $\langle \pi_1 q, \pi_2 p \rangle \colon T \to X \times X$ factors through $\langle \pi_1, \pi_2 \rangle$. $\square$

Definition A.0.2 Let $\mathcal{E}$ be a category and $\mathcal{G}$ a set of objects of $\mathcal{E}$. We say that $\mathcal{G}$ is a *set of generators* for $\mathcal{E}$ if, whenever we have arrows $f, g \colon X \to Y$ with $f \neq g$, there exists $G \in \mathcal{G}$ and $h \colon G \to X$ such that $f \cdot h \neq g \cdot h$. $\square$

A.1 Topologies and sheaves

Definition A.1.1 A *sieve* on an object C of a category $\mathcal{C}$ is a family S of arrows in $\mathcal{C}$ with codomain C such that $f \in S$ implies $f \cdot g \in S$ for all composable f and g. If S is a sieve on C and $h \colon D \to C$ then

$$h^*(S) = \{g \mid \mathrm{cod}(g) = D \text{ and } h \cdot g \in S\}$$

is a sieve on D. $\square$

Definition A.1.2 [57, Defn III.2.1] Let $\mathcal{C}$ be a small category. A *Grothendieck topology* on $\mathcal{C}$ is a function J assigning to each object C in $\mathcal{C}$ a collection of sieves on C such that:

(i) the maximal sieve $t_C = \{f \mid \mathrm{cod}(f) = C\}$ is in $J(C)$;

(ii) (stability axiom) if $S \in J(C)$, then $h^*(S) \in J(D)$ for any arrow $h\colon D \to C$;

(iii) (transitivity axiom) if $S \in J(C)$ and R is any sieve on C such that $h^*(R) \in J(D)$ for all $h\colon D \to C$ in S, then $R \in J(C)$.

The pair $(\mathcal{C}, J)$ is called a *site*. □

Definition A.1.3 Let J be a Grothendieck topology on a small category $\mathcal{C}$. If $P\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ is a presheaf and $S \in J(C)$ is a sieve, a *matching family* for S of elements of P is a function assigning to each $f\colon D \to C \in S$ an element $x_f \in P(D)$, such that

$$P(g)(x_f) = x_{f\cdot g} \quad \text{for all } g\colon E \to D \in \mathcal{C}.$$

An *amalgamation* of such a matching family is an element $x \in P(C)$ such that $P(f)(x) = x_f$ for all $f \in S$. Then we say that P is a *J-sheaf* iff each matching family has a unique amalgamation. We write $\mathsf{Sh}(\mathcal{C}, J)$ for the full subcategory of $[\mathcal{C}^{\mathrm{op}}, \mathsf{Set}]$ consisting of the J-sheaves. A category is a *Grothendieck topos* if it is equivalent to $\mathsf{Sh}(\mathcal{C}, J)$ for some site $(\mathcal{C}, J)$. □

A.2 The topos of coalgebras

A basic result of topos theory is the theorem – due to Lawvere and Tierney – that the category of coalgebras of a comonad (G, ϵ, δ) on a topos $\mathcal{E}$ is itself a topos if G preserves finite limits. In the joint paper [44] the hypothesis 'G preserves finite limits' is weakened to 'G preserves pullbacks'. This achieved by a reduction of the latter case to an instance of the first case on a slice of $\mathcal{E}$. However, once we observe that toposes have partial map classifiers (defined below) we can modify the proof for the finite-limit preserving case to get a direct proof for the pullback preserving case which is 'only one diagram longer'.

Definition A.2.1 A *partial map* from A to B consists of a monic $i\colon I \rightarrowtail A$ and an arrow $f\colon I \to B$. We call i the domain of definition of the partial arrow. A *partial map classifier* for an object B is a monic $\eta_B\colon B \rightarrowtail \widetilde{B}$ such that, for every partial map (i, f) from any object A to B, there is a unique arrow f' which makes the following a pullback:

$$\begin{tikzcd} A \arrow[r, "f'"] & \widetilde{B} \\ I \arrow[u, "i"] \arrow[r, "f"] & B \arrow[u, "\eta_B"'] \end{tikzcd} \tag{A.2}$$

□

In a topos every object has a partial map classifier [57]. For any arrow $f : B \to C$ we have a partial arrow $(\eta_B, f) : \widetilde{B} \to C$, let us denote its total counterpart $\widetilde{f} : \widetilde{B} \to \widetilde{C}$. Now it is straightforward to see that we have a functor $\widetilde{(-)} : \mathcal{E} \to \mathcal{E}$.

Theorem A.2.2 If $\mathbb{G} = (G, \epsilon, \delta)$ is a comonad on a topos $\mathcal{E}$, and if G preserves pullbacks, then $\mathsf{Coalg}(\mathbb{G})$ is itself a topos. □

We have already remarked that the forgetful functor $U_{\mathbb{G}} : \mathsf{Coalg}(\mathbb{G}) \to \mathcal{E}$ creates those limits which are preserved by G. Also, from the proof of theorem 5.4.6 we see that $\mathsf{Coalg}(\mathbb{G})$ has a subobject classifier. Thus, it remains to prove that $\mathsf{Coalg}(\mathbb{G})$ has exponentials.

Let (A, α) and (B, β) be $\mathbb{G}$-coalgebras and let (F, ψ) be the final $\mathbb{G}$-coalgebra. We denote by $A \times_F B$ the pullback in $\mathcal{E}$ of the unique coalgebra maps $!_A \colon A \to F$ and $!_B \colon B \to F$. We also write $GA \times_{GF} GB$ for the pullback in $\mathcal{E}$ of $G(!_A)$ and $G(!_B)$. Since G preserves pullbacks we have an isomorphism

$$\phi : GA \times_{GF} GB \to G(A \times_F B).$$

By the universal property of $GA \times_{GF} GB$, α and β induce a map $A \times_F B \to GA \times_{GF} GB$, written $\alpha \times_{GF} \beta$. Then the product of (A, α) and (B, β) is the $\mathbb{G}$-coalgebra

$$A \times_F B \xrightarrow{\alpha \times_{GF} \beta} GA \times_{GF} GB \xrightarrow{\phi} G(A \times_F B)$$

Before constructing exponentials let us first fix some more notation. The projections π_A and π_B, from $A \times_F B$ to A and B respectively, give us a monic $i = \langle \pi_A, \pi_B \rangle : A \times_F B \rightarrowtail A \times B$.

Given a map $f \colon A \times_F B \to C$ in $\mathcal{E}$, (i, f) is a partial map which has total counterpart $f' \colon A \times B \to \widetilde{C}$. In turn, f' has an exponential transpose $f'' \colon A \to \widetilde{C}^B$. Finally, f'' has a cofree extension to a map of coalgebras $f''' \colon (A, \alpha) \to (G(\widetilde{C}^B), \delta_{\widetilde{C}^B})$. The strategy for constructing exponentials is to convert the conditions on f which make it a map of coalgebras to conditions on f''', the final condition being that f''' factors through a certain sub-coalgebra of $(G(\widetilde{C}^B), \delta_{\widetilde{C}^B})$ which will turn out to be $(C, \gamma)^{(B,\beta)}$.

Given objects C, D of $\mathcal{E}$, since G preserves pullbacks we have an isomorphism $G(C \times D) \cong GC \times_{G1} GD$. This gives rise to a family of maps

$$\lambda_{C,D} : GC \times GD \to G(\widetilde{C \times D})$$

natural in C and D.

An arrow $f : A \times_F B \to C$ in $\mathcal{E}$ is a coalgebra map from $(A, \alpha) \times (B, \beta)$ to (C, γ) iff the left-hand rectangle in (A.3) commutes. But the right-hand square commutes and and $G(\eta_C)$ is monic, so the left-hand rectangle commutes iff the outer rectangle commutes.

$$\begin{tikzcd}
A \times_F B \arrow[r, "\alpha \times_{GF} \beta"] \arrow[d, "f"] & GA \times_{GF} GB \arrow[r, "\phi"] & G(A \times_F B) \arrow[r, "G(i)"] \arrow[d, "G(f)"'] & G(A \times B) \arrow[d, "G(f')"'] \\
C \arrow[rr, "\gamma"] & & GC \arrow[r, "G(\eta_C)"] & G\widetilde{C}
\end{tikzcd} \tag{A.3}$$

The top and the bottom leg of (A.3) and the monic i gives us two partial maps from $A \times B$ to $G\widetilde{C}$. Taking total counterparts we have that (A.3) commutes iff the diagram below commutes.

$$\begin{tikzcd}
A \times B \arrow[r, "\alpha \times \beta"] \arrow[d, "f'"] & GA \times GB \arrow[r, "\lambda_{A,B}"] & G\widetilde{(A \times B)} \arrow[d, "\widetilde{G(f')}"] \\
\widetilde{C} \arrow[r, "\widetilde{\gamma}"'] & \widetilde{GC} \arrow[r, "\widetilde{G(\eta_C)}"'] & \widetilde{G\widetilde{C}}
\end{tikzcd} \tag{A.4}$$

The next step of the proof is to take the exponential transpose in $\mathcal{E}$ of the top and bottom legs of (A.4). The transpose of the bottom leg, from $A \times B$ to $\widetilde{G\widetilde{C}}$, is:

$$A \xrightarrow{f''} \widetilde{C}^B \xrightarrow{\widetilde{\gamma}^B} (\widetilde{GC})^B \xrightarrow{\widetilde{G(\eta_C)}^B} (\widetilde{G\widetilde{C}})^B \tag{A.5}$$

Now the top leg of (A.4) is the top row of (A.6) which is equal to the bottom row of (A.6) since this diagram commutes.

$$\begin{tikzcd}
A \times B \arrow[r, "\alpha \times 1"] \arrow[d, equal] & GA \times B \arrow[r, "1 \times \beta"] \arrow[d, "G(f'') \times 1"'] & GA \times GB \arrow[r, "\lambda_{A,B}"] \arrow[d, "G(f'') \times 1"'] & G\widetilde{(A \times B)} \arrow[r, "\widetilde{G(f')}"] \arrow[d, "\widetilde{G(f'' \times 1)}"'] & \widetilde{G\widetilde{C}} \arrow[d, equal] \\
A \times B \arrow[r, "f''' \times 1"'] & G(\widetilde{C}^B) \times B \arrow[r, "1 \times \beta"'] & G(\widetilde{C}^B) \times GB \arrow[r, "\lambda_{\widetilde{C}^B, B}"'] & G(\widetilde{\widetilde{C}^B \times B}) \arrow[r, "\widetilde{Ge}"'] & \widetilde{G\widetilde{C}}
\end{tikzcd} \tag{A.6}$$

Let us denote the exponential transpose of $\widetilde{Ge} \cdot \lambda_{\widetilde{C}^B, B}$ by $\rho : G(\widetilde{C}^B) \to (\widetilde{G\widetilde{C}})^{GB}$. Then, by naturality of transposition, the bottom row of (A.6) has exponential transpose

$$A \xrightarrow{f'''} G(\widetilde{C}^B) \xrightarrow{\rho} (\widetilde{G\widetilde{C}})^{GB} \xrightarrow{(\widetilde{G\widetilde{C}})^{\beta}} (\widetilde{G\widetilde{C}})^B \,. \tag{A.7}$$

The arrows (A.5) and (A.7) will be equal iff their transposes along the adjunction between $\mathcal{E}$ and $\mathsf{Coalg}(\mathbb{G})$ are the same. We write the respective transposes below recalling that, as f''' is a map of coalgebras, $G(f''') \cdot \alpha = \delta_{\widetilde{C}^B} \cdot f'''$.

$$A \xrightarrow{f'''} G(\widetilde{C}^B) \xrightarrow{G(\widetilde{\gamma}^B)} G((\widetilde{G\widetilde{C}})^B) \xrightarrow{G(\widetilde{G\eta_C}^{\,B})} G((\widetilde{G\widetilde{C}})^B) \tag{A.8}$$

$$\begin{tikzcd}
A \arrow[r, "f'''"] & G(\widetilde{C}^B) \arrow[d, "\delta"] & & \\
& G^2(\widetilde{C}^B) \arrow[r, "G\rho"] & G((\widetilde{G\widetilde{C}})^{GB}) \arrow[r, "G((\widetilde{G\widetilde{C}})^{\beta})"] & G((\widetilde{G\widetilde{C}})^B)
\end{tikzcd}$$

Ignoring f''' we have in (A.8) a parallel pair from $G(\widetilde{C}^B)$ to $G((\widetilde{G\widetilde{C}})^B)$. We may regard this as a parallel pair of maps of cofree $\mathbb{G}$-coalgebras. But their equalizer in $\mathsf{Coalg}(\mathbb{G})$ is the exponential $(C, \gamma)^{(B,\beta)}$ since the two maps in (A.8) are equal iff f''' factors this equalizer.

Bibliography

[1] S. Abramsky. (1990) The lazy lambda calculus. In: Turner, D. (ed.) *Research topics in functional programming*, Addison-Wesley, pp. 65–116.

[2] S. Abramsky. (1991) A domain equation for bisimulation. *Information and Computation*, 92(2), pp. 161–218.

[3] P. Aczel and P.F. Mendler. (1989) A final coalgebra theorem, in *Category Theory and Computer Science*, Lecture Notes in Comp. Sci. vol. 389, Springer–Verlag, pp. 357–365.

[4] P. Aczel. (1988) *Non-well-founded sets.* Number 14 in CLSI Lecture Notes. Center for the Study of Languages and Information, Stanford.

[5] J. Adámek. On final coalgebras of continuous functors. Preprint, available from `http://www.iti.cs.tu-bs.de/`.

[6] J. Adámek and V. Koubek. (1995) On the greatest fixed point of a set functor, *Theoret. Comput. Sci.*, 150, pp. 57–75.

[7] J. Adámek and V. Koubek. (1979) Least fixed point of a functor. *Journal of Computer and System Sciences*, 19, pp. 163–178.

[8] J. Adámek and J. Rosický. (1994) *Locally Presentable and Accessible Categories.* LMS Lecture notes 189, Cambridge University Press.

[9] P. America and J. Rutten. (1989) Solving reflexive domain equations in a category of complete metric spaces. *Journal of computer and system sciences*, 39(3), pp. 343–375.

[10] M. Arbib and E.G. Manes. (1982) Parametrized data types do not need highly constrained parameters. *Information and Control*, 52, pp. 139–158.

[11] R.B. Ash. (1972) *Real Analysis and Probability.* Academic Press.

[12] J. de Bakker and J. Zucker. (1982) Processes and the denotational semantics of concurrency. *Information and Control*, 54, pp. 70–120.

[13] M. Barr. (1993) Terminal coalgebras in well-founded set theory, *Theoret. Comput. Sci.*, 114, pp. 299–315.

[14] M. Barr. (1992) Algebraically compact functors. *Journal for Pure and Applied Algebra*, 82, pp. 211–231.

[15] M. Barr and C. Wells. (1990) *Category theory for computer science.* International Series In Computer Science. Prentice Hall.

[16] B. Bloom, S. Estrail and A. Meyer. (1995) Bisimulation can't be traced. *Journal of the ACM*, 42(1), pp. 232–268.

[17] Francis Borceux. (1994) *Handbook of Categorical Algebra (vol. 1)*. Cambridge University Press.

[18] F. van Breugel. (1999) A note on the Hutchinson metric. Available from `http://www.cs.yorku.ca/~franck/`.

[19] A. Carboni, S. Kasangian and R. Walters. (1987) An axiomatics for bicategories of modules. *Journal of Pure and Applied Algebra*, 45, pp. 127–141.

[20] A. Carboni and P.T. Johnstone. (1995) Connected limits, familial representability and Artin gluing, *Math. Struct. Comp. Sci.*, 5, pp. 441–459.

[21] A. Carboni, G.M. Kelly and R.J. Wood. (1991) A 2-categorical approach to change of base and geometric morphisms 1. *Cah. de Top. et Geo. Diff. Cat.*, 32(1), pp. 47–95.

[22] A. Cheng and M. Nielsen. (1995) Observing Behaviour Categorically. In Thiagarajan, editor, *Foundations of Software Technology and Theoretical Computer Science: 15th Conference*, LNCS 1026, pp. 263-278.

[23] J. Desharnais, A. Edalat and P. Panangaden. (1997) Bisimulation for labelled Markov processes. *In Proc. 12th IEEE Int. Symp. on Logic in Computer Science*, pp. 149–159. IEEE Computer Society Press. Journal version to appear in *Information and Computation.*

[24] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. (1999) Metrics for Labelled Markov Systems. In J.C.M. Baeten and S. Mauw, editors, *Proceedings of the*

10th International Conference on Concurrency Theory, volume 1664 of Lecture Notes in Computer Science. Springer-Verlag.

[25] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. (2000) Approximating continuous Markov processes. *In Proc. 15th IEEE Int. Symp. on Logic in Computer Science*, pp. 95–106. IEEE Computer Society Press.

[26] G.A. Edgar. (1988) *Integral, Probability, and Fractal Measures.* Springer.

[27] A. Edalat. (1999) Semi-pullbacks in categories of Markov processes. *Mathematical Structures in Computer Science*, 9(5), pp. 523–543.

[28] M. Fiore. (1996) A coinduction principle for recursive data types based on bisimulation. *Information and computation*, 127(2), pp. 186–198.

[29] P. Freyd. (1991) Algebraically complete categories. In: Carboni, A., Perdicchio, M. and Rosolini, G. (eds.) Category theory – Proc. of the Int'l Conf. held in Como, Italy, July 1990. *Lecture Notes in Mathematics*, 1488, pp. 95–104. Springer-Verlag.

[30] P. Freyd and A. Scedrov. (1990) *Categories, Allegories.* North-Holland Mathematical Library, vol. 39, Elsevier.

[31] M. Forti and F. Honsell. (1983) Set theory with free construction principles. *Annali Scuola Normale Superiore*, Pisa X (3), pp. 493–522.

[32] S. Ghilardi and G. Meloni. (1988) Modal and tense predicate logic: models in presheaves and categorical conceptualization. *Categorical algebra and its applications,* number 1348, LNM, pp. 130–142, Springer, Berlin.

[33] R.J. van Glabbeek and J.J.M.M. Rutten. (1989) The processes of De Bakker and Zucker represent bisimulation equivalence classes. In *J.W. de Bakker, 25 jaar semantiek*, pp. 243–246. CWI, Amsterdam.

[34] R. Goldblatt. (2000) What is the coalgebraic analogue of Birkhoff's variety theorem? To appear in *Theoretical Computer Science.* Preprint available at `http://www.mcs.vuw.ac.nz/~rob/`.

[35] A. Gordon. (1995) Bisimilarity as a theory of functional programming. BRICS Notes Series NS-95-3, BRICS, University of Aarhus.

[36] J. Goguen. (1992) Sheaf semantics for concurrent interacting systems. *Math. Struct. in. Comp. Sci*, 2, pp. 159–191.

[37] H.P. Gumm. (1999) Elements of the general theory of coalgebras. LUATCS'99, Rand Africaans University, Johannesburg, South Africa.

[38] H.P. Gumm. Functors for coalgebras. To appear, *Algebra Universalis*. Preprint at `http://www.mathematik.uni-marburg.de/~gumm/Papers/publ.html`.

[39] H.P. Gumm and T. Schröder. Products of coalgebras. To appear, *Algebra Universalis*. Preprint at `http://www.mathematik.uni-marburg.de/~gumm/Papers/publ.html`.

[40] F. Honsell, M. Lenisa, U. Montanari and M. Pistore. (1998) Final semantics for the pi-calculus. *Programming Concepts and Methods '98*, D. Gries et al. eds., Chapman and Hall. pp. 225-243.

[41] B. Jacobs. The temporal logic of coalgebras via Galois algebras. Technical Report CSI-R9906, Computing Science Institute, University of Nijmegen.

[42] B. Jacobs and J.J.M.M. Rutten. (1997) A tutorial on (co)algebras and (co)induction. EATCS Bulletin, no. 62, pp. 222–259.

[43] P.T. Johnstone. (1977) *Topos Theory*, L.M.S. Mathematical Monographs no. 10, Academic Press.

[44] P.T. Johnstone, A.J. Power, H. Watanabe, T. Tsujishita and J. Worrell. On the structure of categories of coalgebras. To appear, *Theoretical Computer Science*, 260.

[45] C. Jones. (1990) Probabilistic non-determinism. PhD Thesis, University of Edinburgh. Also published as technical report No. CST-63-90.

[46] A. Joyal, M. Nielsen and G. Winskel. (1996) Bisimulation from open maps. *Information and Computation*, 127(2), pp. 164–185.

[47] Y. Kawahara and M. Mori. (2000) A small final coalgebra theorem. *Theoretical computer science*, 233(1-2), pp. 129-145.

[48] G.M. Kelly. (1982) *Basic concepts of enriched category theory*, Cambridge University Press.

[49] D. Kozen. (1985) A probabilistic PDL. *Journal of Computer and System Sciences*, 30(2), pp. 162–178.

[50] J. Lambek. (1968) A fixpoint theorem for complete categories. *Math. Zeitschr.*, 103, pp. 151–161.

[51] K. Larsen and A. Skou. (1991) Bisimulation through probabilistic testing. *Information and Computation*, 94, pp. 1-28.

[52] S. Lasota. (1999) Coalgebra morphisms subsume open maps. *Proceedings of coalgebraic methods in computer science '99*, Elect. Notes Theor. Comp. Sci. 19, pp. 23–38.

[53] F.W. Lawvere. (1973) Metric spaces, generalized logic, and closed categories. *Rendiconti del Seminario Matematico e Fisico di Milano* 43, pp. 135–166.

[54] M. Lenisa. (1999) From set-theoretic coinduction to coalgebraic coinduction: some results, some open problems. *Proceedings of coalgebraic methods in computer science '99*, Elect. Notes Theor. Comp. Sci. vol 19, pp. 1–22.

[55] F. Linton. (1969) Coequalizers in categories of algebras. *Seminar on Triples and Categorical Homology theory*, Lecture Notes in Mathematics, volume 80, pp. 75–90. Springer-Verlag.

[56] S. Mac Lane. (1971) *Categories for the Working Mathematician*, Graduate Texts in Mathematics no. 5. Springer–Verlag.

[57] S. Mac Lane and I. Moerdijk. (1992) *Sheaves in Geometry and Logic.* Springer-Verlag.

[58] M. Makkai and R. Paré. (1989) *Accessible Categories: the Foundations of Categorical Model Theory*, Contemporary Mathematics vol. 104 (Amer. Math. Soc.)

[59] C. McLarty (1992) *Elementary Categories, Elementary Toposes.* Oxford University Press.

[60] Michael W. Mislove, Lawrence S. Moss and Frank J. Oles. (1991) Non-well-founded Sets Modelled as Ideal Fixed Points. *Information and Computation*, 93(1), pp. 16-54.

[61] R. Milner. (1989) *Communication and Concurrency.* Prentice Hall.

[62] R. Milner. (1981) A modal characterization of observable machine behaviour, in *Proceedings CAAP 81*, G. Astesiano and C. Bohm, eds., volume 112 of LNCS, pp. 25–34. Springer-Verlag.

[63] R. Milner and M. Tofte. (1991) Co-induction in relational semantics. *Theoretical computer science.* 87, pp. 209–220.

[64] Lawrence S. Moss. (1999) Coalgebraic Logic. *Annals of pure and applied logic*, 99, no. 1-3, pp. 241-259.

[65] D.M.R. Park. (1981) Concurrency and automata on infinite sequences. In P. Deussen, editor, *Proceedings 5th GI Conference on Theoretical Computer Science*, volume 104 of LNCS, pp. 15–32. Springer, Berlin.

[66] L.C. Paulson. (1999) Final coalgebras as greatest fixed points in ZF set theory. *Mathematical Structures in Computer Science*, 9(5), pp. 545–567.

[67] A. Pitts. (1996) Relational properties of domains. *Information and Computation*, 127(2), pp. 66–90.

[68] A. Pitts. (1994) A coinduction principle for recursively defined domains. *Theoretical Computer Science*, 124, pp. 195–219.

[69] G.D. Plotkin. (1981) A structural approach to operational semantics. Report DAIMI FN-19, Aarhus University, Aarhus.

[70] J. Power and H. Watanabe. (1998) An axiomatics for categories of coalgebras. *Proceedings of coalgebraic methods in computer science '98.* Elect. Notes Theor. Comp. Sci. vol 11, pp. 161–178.

[71] I.L. Reilly, P.V. Subrahmanyam and M.K. Vamanamurthy. (1982) Cauchy sequences in quasi-pseudo-metric spaces. *Monatshefte fur Mathematik*, 93, pp. 127-120.

[72] M. Rössiger. (1999) Languages for coalgebras on datafunctors. *Proceedings of coalgebraic methods in computer science '99*, Elect. Notes Theor. Comp. Sci., 19., pp. 55–76.

[73] J.J.M.M. Rutten. (2000) Universal coalgebra: a theory of systems, *Theor. Comp. Sci.*, 249(1), pp. 3–80. Also Technical report CS-R9652, Centrum voor Wiskunde en Informatica, Amsterdam, 1996.

[74] J.J.M.M. Rutten. (1998) Automata and coinduction (an exercise in coalgebra). *Proceedings of CONCUR 98*, LNCS 1466, pp. 194–218.

[75] J.J.M.M. Rutten. (1998) Weighted colimits and formal balls in generalized metric spaces. *Topology and its applications*, volume 89, pp. 179–202.

[76] J.J.M.M. Rutten. (1998) Relators and metric bisimulations. *Proceedings of coalgebraic methods in computer science '98*, Elec. Notes in Theor. Comp. Sci., 11, pp. 257–263.

[77] J.J.M.M. Rutten. (1996) Elements of generalized ultrametric domain theory. *Theoretical Computer Science*, 170, pp. 349–381.

[78] J.J.M.M. Rutten. (1999) Automata, power series and coinduction: taking input derivatives seriously. In *Proceedings of ICALP 99*. LNCS 1644, pp. 645–654.

[79] M.B. Smyth. (1990) Topology. In S. Abramsky, D. Gabbay and T.S.E. Maibaum (Eds.) *Handbook of Logic in Computer Science*. Oxford University Press.

[80] M. Smyth and G. Plotkin, (1992) The category theoretic solution of recursive domain equations, *SIAM journal of computing*, 11(4), pp. 761–783.

[81] V. Trnkova. (1969) Some properties of set functors, *Comment. Math. Univ. Carolinae*, 10, pp. 323–352.

[82] D. Turi and J. Rutten. (1998) On the foundations of final semantics: Non-standard sets, metric spaces, partial orders. *Mathematical structures in computer science*, 8(5), pp. 481–540.

[83] E. de Vink and J. Rutten. (1999) Bisimulation for probabilistic transition systems: a coalgebraic approach. *Theoretical Computer Science*, 221(1-2), pp. 271–293.

[84] K. Wagner. (1997) Liminf convergence for Ω-categories. *Theoretical Computer Science*, 184, pp. 61–104.

[85] M. Wand. (1974) On the recursive specification of data types, in *Category theory applied to computation and control*, E.G. Manes, ed., Springer Lecture Notes in Computer Science, 25.

[86] J.B. Worrell. (1998) Toposes of coalgebras and hidden algebras. *Proceedings of coalgebraic methods in computer science '98*, Elect. Notes Theor. Comp. Sci. 11, pp. 215–234.

[87] J.B. Worrell. (1999) Terminal sequences for accessible endofunctors. *Proceedings of coalgebraic methods in computer science '99*, Elect. Notes Theor. Comp. Sci. 19, pp. 39–54.

[88] J.B. Worrell. (2000) Coinduction for recursive data types: partial orders, metric spaces and Ω-categories. *Proceedings of coalgebraic methods in computer science '00*, Elect. Notes Theor. Comp. Sci. 33, pp. 367–386.

[89] G. Wraith. (1974) Artin Gluing, *Journal of Pure and Applied Algebra*, 4, pp. 345-348.

Index