Chapter 4

Coinduction for recursive data types: partial orders, metric spaces and Ω-categories

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. [74, 28, 69]. 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, 83]. 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 [77], we define a notion of simulation[1] for coalgebras which is based on the enriched categorical counterpart of relations, called

[1]The examples of 'metric bisimulations' given in [77] 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.

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 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 [85], 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 [54]. It was later taken up by Wagner [85] and Rutten [78] to unify the basic notions and constructions of order-theoretic and metric domain theory.

4.2 Coinduction in Set

As in Rutten [77], 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 [77, 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 $x \in X, y \in Y$. $\square$

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. [54, 85, 78].

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 $x \in X_0$. $\square$

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.$\square$

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. $\square$

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. $\square$

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). $\square$

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. $\square$

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. $\square$

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 [80] 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] □

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 $S : Z \nrightarrow Y$ are Ω-bimodules, then the formula

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

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

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 $X^{\mathrm{op}} \otimes Y \to \Omega$. □

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.

[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.

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 $\Omega\text{-}\mathsf{Cat} \to \Omega\text{-}\mathsf{Mod}$, 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 $\Omega\text{-}\mathsf{Cat}$ to $\Omega\text{-}\mathsf{Mod}$. We will usually omit the $*$ subscript when regarding functors as bimodules. $\square$

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 $(f^* \cdot R \cdot g)(w,z) = R(fw, gz)$. $\square$

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 [80] 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. [72] identify seven distinct notions of Cauchy sequence and convergence for a gms. Rutten [76] gives an enriched-categorical treatment of convergence for so-called *forward Cauchy* sequences in terms of *weighted colimits*.

Example 4.3.14 [80, 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 [80]. $\square$

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 $X \overset{\pi_1}{\leftarrow} R \overset{\pi_2}{\rightarrow} Y$ 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 $Y \overset{r_1}{\rightarrow} \mathrm{Coll}(R) \overset{r_2}{\leftarrow} X$, and we have $R = r_2^* \cdot r_1$. $\square$

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 $Y \overset{r_1}{\rightarrow} \mathrm{Coll}(R) \overset{r_2}{\leftarrow} X$. 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. $\square$

The following proposition provides a sanity check for this definition. It says that $\mathcal{F}$: $\Omega\text{-}\mathsf{Mod} \to \Omega\text{-}\mathsf{Mod}$ 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 $1 \leq r_1 \cdot h$ since, for $x \in X_0, y \in Y_0$

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

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

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 $W \subseteq Y$,

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

$\square$

A locally monotone graph homomorphism between Preord-categories $G\colon \mathcal{C} \to \mathcal{D}$ is called a *lax functor* if $1_{GC} \le G1_C$ for each object $C \in \mathcal{C}$, and $G\phi \cdot G\psi \le G(\phi\cdot\psi)$ 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

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

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

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.

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.

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

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

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 $T\colon \mathcal{C} \to \mathcal{D}$ 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 [77] 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-$\mathsf{Cat} \to \Omega$-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 $f\colon Y \to X$ 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.* $\square$

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

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

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

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 [74] 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 $f: Y \to X$ 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 $X(-,-): X \nrightarrow X$ 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. □

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), $R \colon Y \looparrowright X$ 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 $R \colon X \looparrowright X$ 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.

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. [75]. We consider an endofunctor on Preord

[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. [62].

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',

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

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 [75] 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. $\square$

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. $\square$

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 [69] 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 $\forall x \in X_0, y \in Y_0$,

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

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 [84] have given an analysis of probabilistic bisimulation, first introduced by Larsen and Skou [52], 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 [84], 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 [84]. 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 [84], 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 $O \subseteq \mathrm{Pr}'$, $a \in \mathrm{Act}$,

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

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 [84] 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 [52]. 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 $R := \{(p, q) : p \vDash \phi \text{ implies } q \vDash \phi \text{ for all LML formulae } \phi\}$ 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 $R_n := \{(p, q) : \forall \phi \in \mathcal{L} \text{ s.t. } \mathrm{md}(\phi) \leq n, p \vDash \phi \Rightarrow q \vDash \phi\}$ for $n < \omega$. Then for all $n < \omega$ (i) R_n is an equivalence, and (ii) $R_{n+1} = \{(p, q) : \mu_{p,a}[E] = \mu_{q,a}[E] \text{ for all } a \in \mathrm{Act}, E \in \mathrm{Pr}/R_n\}$.

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 $\phi_1, \phi_2, \ldots$ an enumeration of $\{\phi \in \mathcal{L} : e \models \phi \text{ and } \mathrm{md}(\phi) \leq n\}$. From the definition of R_n (note that the asymmetry is crucial here) we have

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

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 $\phi_1, \phi_2, \ldots$ be an enumeration of the set $\{\phi \in \mathcal{L} : e \models \phi\}$, then, we can write

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

Immediately we see that if Pr is finite then $E = [\![\phi_1 \wedge \ldots \wedge \phi_n]\!]$ 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 [52] 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 $\mathrm{tt}, \phi_1 \wedge \phi_2$ 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 [52], 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. [80]. We equip $\mathcal{M}X$ with the following generalized metric—a metric analogy of (4.14)

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

Suppose X is a pseudometric space. Then $f\colon X \to [0,1]$ is non-expansive precisely when $|fx - fy| \le X(x,y)$ 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.,

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

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 $\mathcal{M}(f)(\mu)(Y \setminus f(K_\varepsilon)) \le \mu(X \setminus K_\varepsilon) < \varepsilon$.

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 [54] 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.

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

□

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). □

In an unpublished note, van Breugel [18] shows that the functor $1 + c\cdot\mathcal{M}(-)\colon \mathsf{CMet} \to \mathsf{CMet}$ has a final coalgebra. However, he does not undertake any analysis of this domain as we do for the final coalgebra in theorem 4.8.4.

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

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

Proof. Equation (4.19) follows from the inequality

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

□

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 $(\mathrm{Pr}, \Sigma, \mu)$, 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 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.

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

Given two plMps 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 $R\colon Y \looparrowright X$ 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

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

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 plMps 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 $R\colon T \looparrowright S$ 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,

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

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

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 $R\colon Y \looparrowright X$ 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 $\frac{c}{2-c}$.

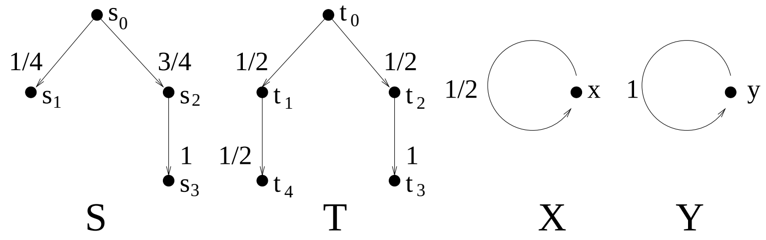

Figure 4.1: Four plMps

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. $\square$

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. $\square$

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 $a \in \text{Act}$ be such that $x \xrightarrow{a} \mu$ and $y \xrightarrow{a} \rho$. Then

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

□

Proposition 4.8.11 Suppose (X, α) and (Y, β) are plMp*s*, $\varphi : X \to Y$ 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 $[\![-]\!] : X \to Y$ 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. □

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 $g' \stackrel{\text{def}}{=} \min(g, r)$ 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 $|g''_X(y) - f(y)| < \varepsilon$. □

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. □

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

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

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 [84]. 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.

Yet another metric-space domain $\mathcal{D}$ for reactive probabilistic processes is described in Kwiatowska and Norman [50]. In contrast to the ultrametric domain introduced in [84], $\mathcal{D}$ is claimed to be a smooth domain; however we present below a counter-example to this claim. In fact we only consider a subspace of $\mathcal{D}$—the domain $\mathcal{P}$ of *probabilistic computations*—which Kwiatowska and Norman introduce as a precursor to the full domain.

The construction of $\mathcal{P}$ follows the method of de Bakker and Zucker [12], using the metric

$$d(\mu, \rho) \quad = \quad \frac{1}{2} \cdot \sum_{x \in X} |\mu(x) - \rho(x)| \tag{4.25}$$

on the set of distributions on a set X.

Fix a set Act of actions,let 1 denote the singleton set $\{*\}$ and once more let $D(X)$ stand for the set of discrete probability distributions on X. Consider the inductively defined collection of sets $(P_n)_{n<\omega}$ given by $P_0 = 1$ and $P_{n+1} = 1 + \mathrm{Act} \times D(P_n)$. The set

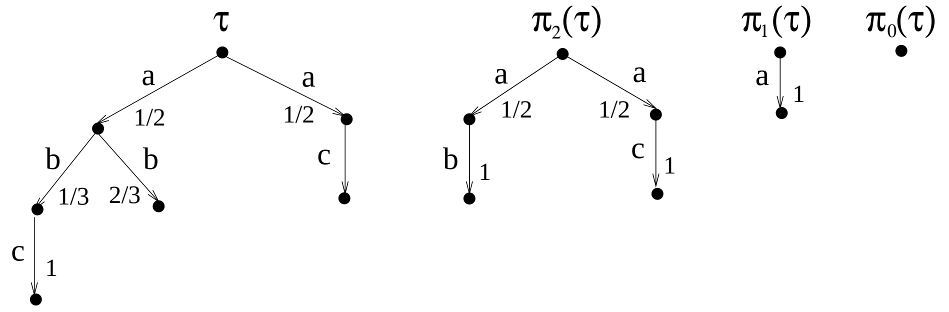

Figure 4.2: Truncations of a tree τ.

P_ω of bounded probabilistic computations is taken to be $\mathrm{Colim}_{n<\omega} P_n$, so we have an isomorphism $\iota\colon P_\omega \to 1 + \mathrm{Act} \times D(P_\omega)$. Elements of P_ω are finite depth strongly-extensional[6] probabilistic synchronization trees.

Next, define projections $\pi_n : P_\omega \to P_n$ for all $n < \omega$ as follows. π_0 is the unique map $P_\omega \to 1$ and

$$\pi_{n+1}(x) = \begin{cases} * & \text{if } \iota(x) = * \\ (a, D(\pi_n)(\mu)) & \text{if } \iota(x) = (a, \mu) \end{cases}$$

for all $x \in P_\omega$.

The projection π_n acts as follows. First it truncates its argument to get a tree of depth n, then it quotients the result to make it strongly-extensional. For each set P_n, write d_n for the metric on the space of distributions on P_n as described by the formula (4.25). Now define the metric d_ω on P_ω by putting $d_\omega(*, *) = 0$, $d_\omega(*, x) = d_\omega(x, *) = 1$ if $x \neq *$, and, if $x, y \in P_\omega$ are such that $\iota(x) = (a, \mu)$ and $\iota(y) = (a', \rho)$, then

$$d_\omega(x, y) = \begin{cases} 1 & \text{if } a \neq a' \\ \sum_n 2^{-n} d_n(D(\pi_n)(\mu), D(\pi_n)(\rho)) & \text{otherwise.} \end{cases}$$

The metric space $\mathcal{P}$ is taken to be Cauchy completion of the space (P_ω, d_ω).

Example 4.8.17 Consider the trees τ and σ drawn in figure 2, where $0 < \epsilon < 1$; then $d_\omega(\tau, \sigma) = 0.25$. However, for $\epsilon = 0$ we have $d_\omega(\tau, \sigma) = 0$.

There is nothing cunning about this choice of counter-example, it just reflects the fact that the metric (4.25) is not smooth. By this last remark we mean that for a metric space X the

[6]We say that a probabilistic synchronization tree is *strongly-extensional* if, viewing the tree as a probabilistic transition system, no two distinct children of any node are bisimilar.

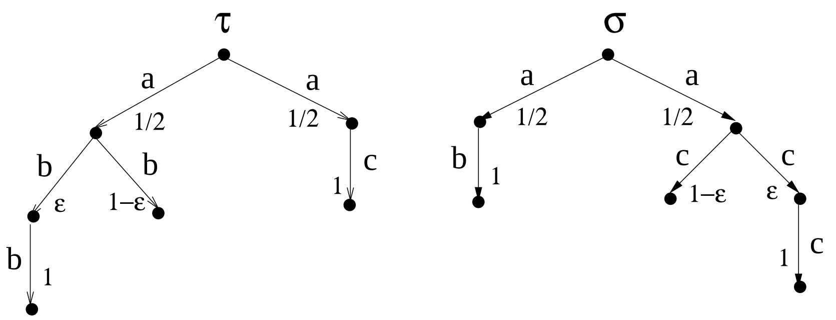

Figure 4.3: Counter-example to smoothness.

map $x \mapsto \delta_x : X \to D(X)$ is not continuous, where δ_x is the Dirac distribution concentrated at $x \in X$. In contrast, if we equip $D(X)$ with the Hutchinson metric then this map is an isometry.

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):

$$x \quad = \quad \bigsqcup_{n<\omega} x_n \;\text{ and }\; \forall n.R(x_n, y) \;\text{ implies }\; R(x, y). \tag{4.26}$$

We call this condition *admissibility*. This is in agreement with the notion of an admissible relation in Pitts [68]. Indeed, his notion of a relational structure is general enough to include bimodules, and (4.26) 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 [81], 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 [83].

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

$$\begin{tikzcd} X \arrow[d, "\alpha"'] & R \arrow[l, "r_1"'] \arrow[r, "r_2"] \arrow[d, "\rho"] \arrow[dl, phantom, "\leq"] & Y \arrow[d, "\beta"] \\ FX & FR \arrow[l, "Fr_1"] \arrow[r, "Fr_2"'] & FY \end{tikzcd} \tag{4.27}$$

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.