A behavioural pseudometric for probabilistic transition systems

Franck van Breugel[a,1], James Worrell[b,*,2]

[a]*Department of Computer Science, York University, 4700 Keele Street, Toronto, Canada M3J 1P3*
[b]*Department of Mathematics, Tulane University, Gibson Hall, 6823 St Charles Avenue, New Orleans LA 70118, USA*

Abstract

Discrete notions of behavioural equivalence sit uneasily with semantic models featuring quantitative data, like probabilistic transition systems. In this paper, we present a pseudometric on a class of probabilistic transition systems yielding a quantitative notion of behavioural equivalence. The pseudometric is defined via the terminal coalgebra of a functor based on a metric on the space of Borel probability measures on a metric space. States of a probabilistic transition system have distance 0 if and only if they are probabilistic bisimilar. We also characterize our distance function in terms of a real-valued modal logic.
© 2004 Elsevier B.V. All rights reserved.

Keywords: Probabilistic transition system; Pseudometric; Probabilistic bisimilarity; Terminal coalgebra; Real-valued modal logic

1. Introduction

The majority of verification methods for concurrent systems only produce qualitative information. Questions like "Does the system satisfy its specification?" and "Do the systems behave the same?" are answered "Yes" or "No". Giacalone et al. [14], Huth and

---

* Corresponding author. Tel.: +1 504 862 3428; fax: +1 504 865 5063.
  E-mail address: jbw@math.tulane.edu (J. Worrell).

Kwiatkowska [19] and Desharnais et al. [12] have pointed out that such Boolean-valued reasoning sit uneasily with semantic models featuring quantitative data, like probabilistic transition systems.

The probabilities occurring in a probabilistic model of a system may be based on statistical sampling. In this case, the model is only an approximate description of a system, and it makes no sense to ask if any two states in the model behave exactly the same. Even if we have a precise description of a system, we may still want to express the idea that two states exhibit almost the same behaviour. Furthermore, the problem of automatically verifying that two states are exactly equivalent will typically require exact real arithmetic. For automatic verification of probabilistic systems, in the conventional setting of floating point arithmetic, it is more reasonable to consider approximate equivalence.

The above observations apply to a number of different semantics for probabilistic systems. In this paper, however, we concentrate on Larsen and Skou's probabilistic bisimulation [24]. Recall that a probabilistic bisimulation is an equivalence relation on the state space of a transition system such that related states have exactly the same probability of making a transition into any equivalence class. Thus, for instance, the states s_0 and s_0^ε of the probabilistic transition system
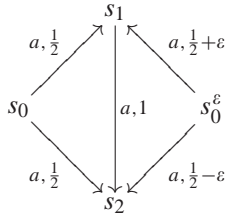
are only probabilistic bisimilar if ε is 0. However, the two states behave almost the same for very small ε different from 0.

In the previous example, varying ε gave different probabilities on the same underlying transition system. (For instance, ε may correspond to a rounding error arising from giving a finite presentation of a real number.) We also want to consider approximate equivalence where the underlying transition systems are different. For example, consider the infinite state system
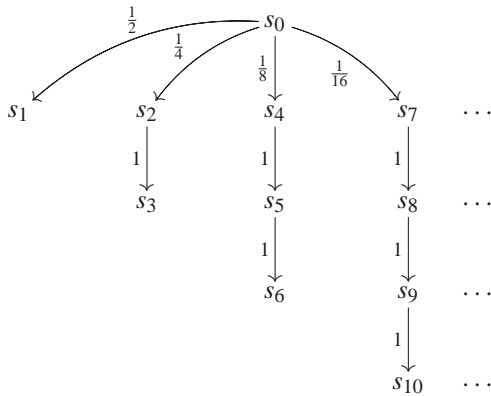
Below we illustrate the finite state system arising by truncating to depth 3.
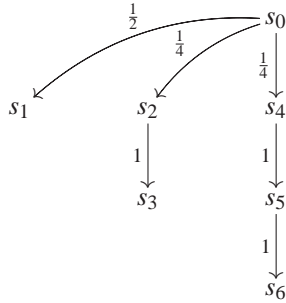
Such truncations are not bisimilar to the original infinite system; nevertheless it is intuitively clear that by truncating at greater and greater depths one gets closer to the original system. It would be useful to formalize and quantify this convergence so that one could safely reason about the infinite state system by examining a suitable finite approximant.

To the best of our knowledge, the earliest attempt to address some of the problems with probabilistic bisimulation as outlined above is the paper of Giacalone et al. [14]. They define a pseudometric on the states of a (restricted type of) probabilistic transition system, yielding a smooth, quantitative notion of behavioural equivalence. A pseudometric differs from an ordinary metric in that different elements, that is, states, can have distance 0. We would like that the distance between states, a real number between 0 and 1, will express the similarity of the behaviour of those states. The smaller the distance, the more alike the behaviour is. In particular, the distance between states is 0 if they are behaviourally indistinguishable.

The present paper is most closely related to the work of Desharnais et al. [11] and De Vink and Rutten [33]. The former introduce a pseudometric on a class of probabilistic transition systems more general than that considered by Giacalone et al. This pseudometric is defined via a nonstandard semantics for a probabilistic modal logic, where formulae get interpreted as measurable functions into the interval [0, 1], rather than as Boolean-valued functions. We show that what we present in this paper is essentially a coinductive account of a closely related pseudometric. We will also discuss some of the advantages conferred by such an account.

The connection between the present paper and De Vink and Rutten [33] is in the modelling of probabilistic systems as coalgebras. Coalgebras offer a simple and uniform categorical notion of transition system, including an account of bisimulation. Rutten [29] shows that many different kinds of transition system can be captured in this framework. Roughly speaking, a coalgebra consists of a carrier set, and a coalgebraic structure determining how elements of the carrier can be decomposed into other elements of the carrier. Thus coalgebras are dual to algebras. For transition systems, the coalgebraic structure is given by the dynamics of the system.

By modelling a restricted class of probabilistic transition systems as coalgebras, and using a standard result from the general theory of coalgebras, De Vink and Rutten established the existence of a terminal object in their category of systems. By definition there is a unique map from an arbitrary system to the terminal one. Furthermore, De Vink and Rutten showed that the kernel of the unique map coincides with probabilistic bisimilarity.

In this paper, we exploit the coalgebraic framework to define a notion of quantitative behavioural equivalence for probabilistic transition systems. In particular, we define a pseudometric on the states of a probabilistic transition system in terms of the terminal coalgebra of an endofunctor P on the category of pseudometric spaces and nonexpansive maps. The definition of P is based on a metric on Borel probability measures. This metric is known as Hutchinson metric [18], Kantorovich metric, Monge–Kantorovich metric, Kantorovich–Rubenstein metric [20], Vaserstein (Wasserstein and even Vasserstein) metric [32], transport metric, earthmover's metric, and match metric. P-coalgebras can be seen as probabilistic transition systems with discrete or continuous state spaces. The terminal P-coalgebra provides for a notion of approximate equivalence similar to the pseudometric of Desharnais et al. mentioned above. In fact, we define a pseudometric on the state space of a probabilistic transition system, seen as a P-coalgebra, as the pseudometric kernel of the unique map to the terminal P-coalgebra. That is, the distance between two states is the distance between their images under the unique map to the terminal P-coalgebra. Moreover states are at distance 0 just in case they are probabilistic bisimilar in the sense of Larsen and Skou.

So far we have motivated our concern for defining a notion of quantitative behavioural equivalence by examples featuring probabilistic transition systems with discrete state spaces. However our framework is sufficiently general to model probabilistic transition systems the state space of which is continuous, like [0, 1]. We refer the reader to [12] for a discussion of the importance of modelling continuous as well as discrete systems.

The rest of this paper is organized as follows. In Section 2, we present some minor variations on results of [1,31] which allow us to prove that a terminal P-coalgebra exists. In Section 3, we present the metric on Borel probability measures, and recall a number of standard results about this metric. In Section 4, we introduce a functorial extension of this metric, and we verify that it satisfies the properties required for the application of the terminal coalgebra theorem. In Section 5, we present the functor P and we show that all discrete probabilistic transition systems and a large class of continuous probabilistic transition systems can be viewed as P-coalgebras. We introduce our pseudometric in Section 6. In Section 7, we introduce a pseudometric defined in terms of a modal logic a la Desharnais et al. Section 8 contains the main result of the paper: we show that the coalgebraic pseudometric, introduced in Section 6, coincides with the logical pseudometric, introduced in Section 7. The proof involves an application of the Stone–Weierstrass approximation theorem for continuous functions. In the concluding section we present related and future work.

The reader is assumed to be familiar with some very elementary category theory, metric space theory and probability theory. For more details, we refer the reader to, for example, the texts of Mac Lane [25], Sutherland [30] and Billingsley [6].

2. A pseudometric terminal coalgebra theorem

In this section, we first introduce coalgebras and then we give the ingredients of our metric coalgebraic framework, to obtain a mild generalization of Rutten and Turi's metric terminal coalgebra theorem [31].

Definition 1. Let $\mathcal{C}$ be a category. Let $F : \mathcal{C} \to \mathcal{C}$ be a functor. An F-coalgebra consists of an object C in $\mathcal{C}$ together with an arrow $f : C \to F(C)$ in $\mathcal{C}$. The object C is called the carrier. The arrow f is called the structure. An F-homomorphism from an F-coalgebra $\langle C, f \rangle$ to an F-coalgebra $\langle D, g \rangle$ is an arrow $\phi : C \to D$ in $\mathcal{C}$ such that $F(\phi) \circ f = g \circ \phi$.

$$
\begin{array}{ccc}
C & \xrightarrow{\ \ \phi\ \ } & D \\
\downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} \\
F(C) & \xrightarrow{\ F(\phi)\ } & F(D)
\end{array}
$$

The F-coalgebras and F-homomorphisms form a category. If this category has a terminal object, then this object is called the terminal F-coalgebra.

For more details about the theory of coalgebras we refer the reader to, for example, Rutten's [29].

In the rest of this section, we restrict our attention to the category $\mathcal{PMet}_1$ of 1-bounded pseudometric spaces and nonexpansive functions. A pseudometric space differs from an ordinary metric space in that different elements can have distance 0. Elements at distance 0 will be considered equivalent. A pseudometric space is 1-bounded if all its distances are bounded by 1. A function $f : X \to Y$ is nonexpansive if it does not increase any distances, that is, $d_Y(f(x_1), f(x_2)) \leqslant d_X(x_1, x_2)$ for all $x_1, x_2 \in X$. We denote the collection of nonexpansive functions from the space X to the space Y by $X \xrightarrow{1} Y$. This collection can be turned into a pseudometric space by endowing the functions with the supremum metric: $d_{X \xrightarrow{1} Y}(f_1, f_2) = \sup_{x \in X} d_X(f_1(x), f_2(x))$.

Let c be a constant in the open interval $(0, 1)$. A function $f : X \to Y$ is c-contractive if it decreases all distances by at least a factor c, that is, $d_Y(f(x_1), f(x_2)) \leqslant c \cdot d_X(x_1, x_2)$ for all $x_1, x_2 \in X$. This notion can be lifted to functors as follows.

Definition 2. A functor $F : \mathcal{PMet}_1 \to \mathcal{PMet}_1$ is locally c-contractive if for all pseudometric spaces X and Y, the function

$$F_{X,Y} : (X \xrightarrow{1} Y) \to (F(X) \xrightarrow{1} F(Y))$$

defined by

$$F_{X,Y}(f) = F(f)$$

is c-contractive.

In the rest of this section, we restrict ourselves to locally contractive functors. Furthermore, we focus on functors which preserve positivity (different elements have a positive distance) and completeness.

Definition 3. A functor $F : \mathcal{PMet}_1 \to \mathcal{PMet}_1$ preserves positivity and completeness if for all complete metric spaces X, $F(X)$ is a complete metric space.

A functor which preserves positivity and completeness can be restricted to a functor on the category of complete metric spaces and nonexpansive functions.

A locally contractive functor F which preserves positivity and completeness has a unique fixed point. That is, there exists a unique space, say $fix(F)$, such that there is an isometry from $fix(F)$ to $F(fix(F))$. Recall that an isometry is a bijection that preserves all distances.

In the rest of this section, we present simple translations of results in [1,31] from metric spaces to pseudometric spaces.

Lemma 4 (*Turi and Rutten [31, Theorem 7.2]*). *For each locally c-contractive functor F : $\mathcal{PMet}_1 \to \mathcal{PMet}_1$ which preserves positivity and completeness, there exists a unique complete metric space $fix(F)$ such that there is an isometry $i : fix(F) \to F(fix(F))$.*

Next, we show that $\langle fix(F), i \rangle$ is a terminal F-coalgebra. For the rest of this section, we fix $\langle X, f \rangle$ to be an F-coalgebra. To characterize the unique F-homomorphism from the F-coalgebra $\langle X, f \rangle$ to the F-coalgebra $\langle fix(F), i \rangle$ we introduce the following function.

Definition 5. The function $\Phi_{\langle X, f \rangle} : (X \xrightarrow{}_1 fix(F)) \to (X \xrightarrow{}_1 fix(F))$ is defined by

$$\Phi_{\langle X, f \rangle}(\phi) = i^{-1} \circ F(\phi) \circ f.$$

$$
\begin{array}{ccc}
X & \xrightarrow{\phi} & fix(F) \\
{\scriptstyle f}\downarrow & & \uparrow{\scriptstyle i^{-1}} \\
F(X) & \xrightarrow{F(\phi)} & F(fix(F))
\end{array}
$$

Since the functor F is locally c-contractive, we have that the function $\Phi_{\langle X, f \rangle}$ is c-contractive. Because $fix(F)$ is a complete metric space, $X \xrightarrow{}_1 fix(F)$ is a complete metric space as well. Obviously, the space $X \xrightarrow{}_1 fix(F)$ is nonempty. Since $\Phi_{\langle X, f \rangle}$ is a contractive function from a nonempty complete metric space to itself, we can conclude from Banach's theorem that it has a unique fixed point $fix(\Phi_{\langle X, f \rangle})$. This function $fix(\Phi_{\langle X, f \rangle})$ is the unique F-homomorphism from the F-coalgebra $\langle X, f \rangle$ to the F-coalgebra $\langle fix(F), i \rangle$.

Theorem 6 (*Turi and Rutten [31, Proposition 7.1]*). *For every locally c-contractive functor F : $\mathcal{PMet}_1 \to \mathcal{PMet}_1$ which preserves positivity and completeness there exists a terminal F-coalgebra $\langle fix(F), i \rangle$ and $fix(F)$ is a complete metric space.*

We will exploit both the terminal F-coalgebra $\langle fix(F), i \rangle$ and the unique F-homomorphism $fix(\Phi_{\langle X, f \rangle})$ when defining our pseudometric.

If the functor F also preserves compactness, then we can conclude that the carrier of the terminal F-coalgebra is a compact metric space.

Theorem 7 (*Alessi et al. [1, Theorem 4.4]*). *For every locally c-contractive functor $F : \mathcal{PMet}_1 \to \mathcal{PMet}_1$ which preserves positivity and compactness there exists a terminal F-coalgebra $\langle fix(F), i \rangle$ and $fix(F)$ is a compact metric space.*

3. A pseudometric on Borel probability measures

The set of Borel probability measures on a space can be turned into a pseudometric space in several ways (see, for example, Rachev's book [28]). In this section, we introduce a pseudometric on Borel probability measures which gives rise to meaningful distances on probabilistic transition systems.

Let X be a (1-bounded) pseudometric space. We denote the set of Borel probability measures on X by $M(X)$.

Definition 8. The distance function $d_{M(X)} : M(X) \times M(X) \to [0, 1]$ is defined by

$$d_{M(X)}(\mu_1, \mu_2) = \sup \left\{ \int_X f \, d\mu_1 - \int_X f \, d\mu_2 \mid f \in X \xrightarrow{1} [0, 1] \right\}.$$

Before presenting an example, let us first check that this distance function is indeed a pseudometric.

Proposition 9. *The distance function $d_{M(X)}$ is a pseudometric.*

Proof. For all nonexpansive functions $f : X \to [0, 1]$,

$$0 = \int_X 0 \, d\mu \leqslant \int_X f \, d\mu \leqslant \int_X 1 \, d\mu = 1.$$

Hence, $d_{M(X)}(\mu_1, \mu_2) \in [0, 1]$.

Obviously, $d_{M(X)}(\mu, \mu) = 0$.

To prove symmetry, it suffices to observe that for each nonexpansive function $f : X \to [0, 1]$, the function $1 - f : X \to [0, 1]$ is nonexpansive as well, and

$$\int_X (1 - f) \, d\mu = 1 - \int_X f \, d\mu.$$

Since for each nonexpansive function $f : X \to [0, 1]$,

$$\int_X f \, d\mu_1 - \int_X f \, d\mu_3 = \left(\int_X f \, d\mu_1 - \int_X f \, d\mu_2 \right) + \left(\int_X f \, d\mu_2 - \int_X f \, d\mu_3 \right)$$
$$\leqslant d_{M(X)}(\mu_1, \mu_2) + d_{M(X)}(\mu_2, \mu_3),$$

the distance function $d_{M(X)}$ satisfies the triangle inequality. $\quad\square$

Example 10. Let the set $\{x_0, x_1\}$ be endowed with the discrete metric, that is, all distances are either 0 or 1. Let μ_ε be the discrete Borel probability measure determined by

$$\mu_\varepsilon(\{x_0\}) = \tfrac{1}{2} + \varepsilon,$$
$$\mu_\varepsilon(\{x_1\}) = \tfrac{1}{2} - \varepsilon.$$

The measures μ_0 and μ_ε have distance ε. This is witnessed by the function mapping x_0 to 0 and x_1 to 1.

Let the set $[0, 1]$ be endowed with the Euclidean metric. For each $\delta \in [0, 1]$, consider the Borel probability measure μ_δ determined by

$$\mu_\delta ([x_\ell, x_r]) = \int_{[x_\ell, x_r]} g_\delta \, dx,$$

where the function g_δ is defined by

$$g_\delta (x) = \begin{cases} 4x\delta - \delta + 1 & \text{if } x \in [0, \frac{1}{2}], \\ -4x\delta + 3\delta + 1 & \text{if } x \in [\frac{1}{2}, 1]. \end{cases}$$

To compute the distance between the measures μ_0 and μ_1 we need to find a nonexpansive function $f : [0, 1] \to [0, 1]$ which maximizes

$$\int_{[0,1]} f \, d\mu_0 - \int_{[0,1]} f \, d\mu_1. \tag{1}$$

Note that the measure μ_0 distributes the probability evenly over the interval $[0, 1]$ whereas the measure μ_1 concentrates its probability around $\frac{1}{2}$. Therefore, a function that maximizes (1) should take its minimum at $\frac{1}{2}$ and its maximum at 0 and 1. Clearly the nonexpansive function $f : [0, 1] \to [0, 1]$ defined by

$$f (x) = \begin{cases} 1 - x & \text{if } x \in [0, \frac{1}{2}], \\ x & \text{if } x \in [\frac{1}{2}, 1] \end{cases}$$

is such a function. One can now easily verify that the distance between the measures μ_0 and μ_1 is $\frac{1}{12}$.

Next, we present some results that will be exploited later. First of all, we note that M preserves positivity.

Proposition 11. *X is a metric space if and only if $M(X)$ is a metric space.*

Proof. See, for example, Edgar's textbook [13, Proposition 2.5.14]. $\quad\square$

In the rest of this paper, we focus on Borel probability measures which are completely determined by their values for the compact subsets of the space X.

Definition 12. A Borel probability measure μ on X is tight if for all $\varepsilon > 0$ there exists a compact subset K_ε of X such that $\mu (X \setminus K_\varepsilon) < \varepsilon$.

Under quite mild conditions on the space, for example, completeness and separability, every measure is tight (see, for example, Parthasarathy's textbook [27, Theorem II.3.2]). Discrete Borel probability measures are tight. All measures presented in Example 10 are tight. We denote the set of tight Borel probability measures on X by $M_t (X)$. We are interested in these tight measures because of the following:

Theorem 13. (1) *X is complete if and only if $M_t(X)$ is complete.*
(2) *X is compact if and only if $M_t(X)$ is compact.*

Proof. Proofs of (1) and (2) can be found in, for example, the texts of Edgar [13, Theorem 2.5.25] and Barnsley [5, Theorem 9.5.1], respectively. $\square$

4. The functor M_t

We extend M_t to a functor on the category $\mathcal{PMet}_1$ of 1-bounded pseudometric spaces and nonexpansive functions. Furthermore, we show that the functor is locally nonexpansive.

Let X and Y be pseudometric spaces. Let $f : X \to Y$ be a nonexpansive function. To extend M_t to a functor we have to define a nonexpansive function $M_t(f)$ from tight measures on X to tight measures on Y.

Definition 14. The function $M_t(f) : M_t(X) \to M(Y)$ is defined by

$$M_t(f)(\mu) = \mu \circ f^{-1}.$$

Next, we prove that the measure $M_t(f)(\mu)$ is tight and that the function $M_t(f)$ is nonexpansive.

Proposition 15. *The measure $M_t(f)(\mu)$ is tight.*

Proof. Let $\varepsilon > 0$. Since μ is tight, there exists a compact subset K_ε of X such that $\mu(X \setminus K_\varepsilon) < \varepsilon$. Because f is nonexpansive, $f(K_\varepsilon)$ is a compact subset of Y. Since $f^{-1}(Y \setminus f(K_\varepsilon))$ is a subset of $X \setminus K_\varepsilon$, we can conclude that $(\mu \circ f^{-1})(Y \setminus f(K_\varepsilon)) < \varepsilon$. Hence, $\mu \circ f^{-1}$ is tight. $\square$

Proposition 16. *The function $M_t(f)$ is nonexpansive.*

Proof. For all $\mu_1, \mu_2 \in M_t(X)$,

$$d_{M_t(Y)}(M_t(f)(\mu_1), M_t(f)(\mu_2))$$

$$= \sup \left\{ \int_Y g \, d(\mu_1 \circ f^{-1}) - \int_Y g \, d(\mu_2 \circ f^{-1}) \mid g \in Y \xrightarrow[1]{} [0, 1] \right\}$$

$$= \sup \left\{ \int_X (g \circ f) \, d\mu_1 - \int_X (g \circ f) \, d\mu_2 \mid g \in Y \xrightarrow[1]{} [0, 1] \right\}$$

$$\leqslant \sup \left\{ \int_X h \, d\mu_1 - \int_X h \, d\mu_2 \mid h \in X \xrightarrow[1]{} [0, 1] \right\}$$

$$= d_{M_t(X)}(\mu_1, \mu_2). \quad \square$$

Clearly, the action of M_t on arrows is functorial.

We conclude this section with a property of M_t which will later allow us to exploit the pseudometric terminal coalgebra theorem.

Proposition 17. *The functor M_t is locally nonexpansive, that is, for all nonexpansive functions f_1, $f_2 \in X \to Y$,*

$$d_{M_t (X) \to M_t (Y)} (M_t (f_1), M_t (f_2)) \leqslant d_{X \to Y} (f_1, f_2).$$

Proof. For all $\mu \in M_t (X)$,

$$
\begin{aligned}
d_{M_t (Y)} &(M_t (f_1)(\mu), M_t (f_2)(\mu)) \\
&= \sup \left\{ \int_Y g \, \mathrm{d}(\mu \circ f_1^{-1}) - \int_Y g \, \mathrm{d}(\mu \circ f_2^{-1}) \mid g \in Y \underset{1}{\to} [0, 1] \right\} \\
&= \sup \left\{ \int_X (g \circ f_1) \, \mathrm{d}\mu - \int_X (g \circ f_2) \, \mathrm{d}\mu \mid g \in Y \underset{1}{\to} [0, 1] \right\} \\
&= \sup \left\{ \int_X (g \circ f_1 - g \circ f_2) \, \mathrm{d}\mu \mid g \in Y \underset{1}{\to} [0, 1] \right\} \\
&\leqslant d_{X \to Y} (f_1, f_2),
\end{aligned}
$$

since for all $g \in Y \underset{1}{\to} [0, 1]$ and $x \in X$,

$$
\begin{aligned}
(g \circ f_1 &- g \circ f_2) (x) \\
&\leqslant |(g \circ f_1) (x) - (g \circ f_2) (x)| \\
&\leqslant d_Y (f_1 (x), f_2 (x)) \quad (g \text{ is nonexpansive}) \\
&\leqslant d_{X \to Y} (f_1, f_2). \qquad \square
\end{aligned}
$$

5. Probabilistic transition systems as coalgebras

In this section, we introduce discrete and continuous probabilistic transition systems. Furthermore, we present a functor P, and show that all discrete probabilistic transition systems and a large class of continuous probabilistic transition systems can be represented as P-coalgebras.

Before considering continuous systems, we first have a look at discrete probabilistic transition systems.

Definition 18. A discrete probabilistic transition system consists of a finite set S of states, a set Act of actions and a transition function $t : S \times Act \times S \to [0, 1]$ such that

$$\text{for all } s \in S \text{ and } a \in Act, \quad \sum_{s' \in S} t_{s,a} (s') \leqslant 1. \tag{2}$$

Such a system is also called a partial labelled Markov chain in the literature (see, for example, [11]). The function $t_{.,a}$ describes the reaction of the system to the action a selected by the environment. This represents a reactive model of probabilistic systems. For a detailed

discussion of this reactive model and other models, and the relationship of the reactive one to those other models we refer the reader to the work of Van Glabbeek et al. [16]. Given the system is in state s and reacts to action a chosen by the environment, $t_{s,a}(s')$ is the probability that the system makes a transition to the state s'. Note that this is a conditional probability. Also notice that we consider subprobabilities as we use $\leqslant$, instead of $=$, in (2). These subprobabilities allow for the possibility that the system may refuse an action. The probability of refusal of the action a given the system is in state s is

$$1 - \sum_{s' \in S} t_{s,a}(s').$$

Example 19. Consider the discrete probabilistic transition system with the states $s_0, s_0^\varepsilon, s_1,$ $s_1^\varepsilon, s_2, s_3$, the action a, and the transitions
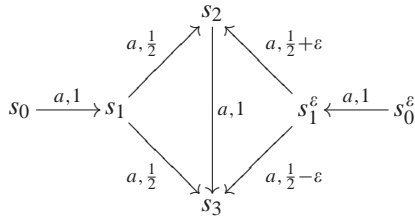
The action a is refused in state s_3 with probability 1.

The key behavioural equivalence on the states of a discrete probabilistic transition system is probabilistic bisimulation. This notion is due to Larsen and Skou [24] and presented in:

Definition 20. An equivalence relation R on the set S of states of a discrete probabilistic transition system is a probabilistic bisimulation if $s_1 \ R \ s_2$ implies $\sum_{s \in E} t_{s_1,a}(s) = \sum_{s \in E} t_{s_2,a}(s)$ for all R-equivalence classes E of states and $a \in Act$.
    States s_1 and s_2 are probabilistic bisimilar if $s_1 \ R \ s_2$ for some probabilistic bisimulation R.

Next, we introduce continuous probabilistic transition systems.

Definition 21. A continuous probabilistic transition system consists of a set S of states, a σ-algebra Σ on S, a finite set Act of actions and a transition function $t : S \times Act \times \Sigma \to [0,1]$ such that
(1) for all $s \in S$ and $a \in Act$, the function $t_{s,a}(\cdot) : \Sigma \to [0,1]$ is a subprobability measure, and
(2) for all $a \in Act$ and $B \in \Sigma$, the function $t_{\cdot,a}(B) : S \to [0,1]$ is measurable.

Such a system is also called a partial labelled Markov process in the literature (see, for example, [11]). Another way of expressing (1) and (2) is that $t_{\cdot,a}$ is a stochastic kernel for each $a \in Act$. The main difference between the definition of discrete and continuous systems is the use of a σ-algebra in the latter. Given the system is in state s and reacts

to action a chosen by the environment, $t_{s,a}(B)$ is the probability that the system makes a transition to a state in the set B.

Example 22. Consider the continuous probabilistic transition system with the set of states $[0, 1]$ and its Borel σ-algebra, actions a and b, and the transition function determined by

$$
\begin{aligned}
t_{s,a} &= \mu_s, \\
t_{s,b}(\{\tfrac{1}{2}\}) &= 1
\end{aligned}
$$

with the measure μ_s as defined in Example 10. Note that the probability of making an a-transition from one state to another is always 0.

A discrete probabilistic transition system is just a special case of a continuous one where the σ-algebra is discrete and the transition subprobability measure is determined by a sub-probability distribution.

Probabilistic bisimulation has been generalized to the continuous setting by Blute et al. [7].

Definition 23. A set B is R-closed if $s_1 \in B$ and $s_1 \, R \, s_2$ implies $s_2 \in B$.

An equivalence relation R on the set S of the states of a continuous probabilistic transition system is a probabilistic bisimulation if $s_1 \, R \, s_2$ implies $t_{s_1,a}(B) = t_{s_2,a}(B)$ for all R-closed measurable sets B and $a \in Act$.

States s_1 and s_2 are probabilistic bisimilar if $s_1 \, R \, s_2$ for some probabilistic bisimulation R.

Next, we introduce a functor $P : \mathcal{PMet}_1 \to \mathcal{PMet}_1$ such that P-coalgebras represent probabilistic transition systems. The functor P is built from a number of functors. Below we only present their action on objects. Their action on arrows can be obtained in a standard way (see, for example, America and Rutten's [2, Section 5]).

- **1** is the terminal object functor. This functor maps each object to the terminal object of $\mathcal{PMet}_1$ which is the singleton space.
- $c \cdot -$ is the scaling functor. The scaling by $c \cdot -$ of an object in $\mathcal{PMet}_1$ leaves the set unchanged and multiplies all distances by c. For the rest of this paper, we fix c to be an arbitrary value in the open interval $(0, 1)$.
- $+$ is the coproduct functor. The coproduct object of the objects X and Y in $\mathcal{PMet}_1$ is the disjoint union of the sets underlying the spaces X and Y endowed with the pseudometric

$$
d_{X+Y}(v, w) = \begin{cases} d_X(v, w) & \text{if } v \in X \text{ and } w \in X, \\ d_Y(v, w) & \text{if } v \in Y \text{ and } w \in Y, \\ 1 & \text{otherwise.} \end{cases}
$$

- M_t is the functor introduced in Section 4.
- $-^{Act}$ is the power functor. For an object X in $\mathcal{PMet}_1$, X^{Act} is the Act-indexed product of copies of X equipped with the supremum metric.

The functor P is defined by

$$P = M_t \, (\mathbf{1} + c \cdot -)^{Act}.$$

This functor can expressed in terms of the auxiliary functors Q and R as follows.

$$
\begin{aligned}
R &= \mathbf{1} + c \cdot -, \\
Q &= M_t \, (R \, (-)), \\
P &= Q \, (-)^{Act}.
\end{aligned}
$$

A P-coalgebra consists of a pseudometric space S together with a nonexpansive function $t : S \to P \, (S)$. The space S corresponds to the set of states of the probabilistic transition system. The nonexpansive function $t : S \to P \, (S)$ characterizes the transitions of the system. Given a state s and an action a, $t_{s,a}$ is a tight Borel probability measure on $R \, (S)$. It captures the reaction on action a of the system in state s. We use $M_t \, (R \, (S))$ to represent subprobabilities on S. The probability of refusal of action a in state s is given by $t_{s,a} \, (\mathbf{1})$. The role of $c \cdot -$ will be discussed later.

Note that each P-coalgebra can be interpreted as a continuous probabilistic transition system, since nonexpansive functions are measurable.

Proposition 24. *Every discrete probabilistic transition system can be represented by a P-coalgebra.*

Proof. We endow the set of states S of the system with the discrete metric. Consequently, every subset of the pseudometric space $R \, (S)$ is a Borel set. For every state s and action a, the Borel probability measure $t_{s,a}$ is the discrete Borel probability measure determined by

$$
\begin{aligned}
t_{s,a} \, (\mathbf{1}) &= \text{probability of refusal of action } a \text{ in state } s, \\
t_{s,a} \, (\{s'\}) &= \text{probability of making an } a\text{-transition from state } s \text{ to state } s'.
\end{aligned}
$$

Obviously, the measure $t_{s,a}$ is tight. Because S is endowed with the discrete metric, the function t from S to $P \, (S)$ is nonexpansive. $\quad\square$

Example 25. The continuous probabilistic transition system of Example 22 can be viewed as a P-coalgebra by endowing its state space with the Euclidean metric.

A continuous probabilistic transition system can be viewed as a P-coalgebra if its set S of states can be endowed with a pseudometric d_S such that
- the Borel σ-algebra induced by the pseudometric d_S coincides with the σ-algebra of the system,
- for all states s and actions a, the system's subprobability measure $t_{s,a}$ is tight, and
- the system's transition function is nonexpansive.

We refer the reader forward to the conclusion for further discussion of these restrictions.

6. A pseudometric on probabilistic transition systems

Next, we present our pseudometric on probabilistic transition systems. Furthermore, we show that states have distance 0 if and only if they are probabilistic bisimilar. The pseudometric on the states of a probabilistic transition system will be defined as a pseudometric kernel.

A function ϕ from a set S to a pseudometric space X defines a distance function d_ϕ on S. We call this distance function the *pseudometric kernel* induced by ϕ. The distance between s_1 and s_2 in S is defined as the distance of their ϕ-images in the pseudometric space X.

Definition 26. Let $\phi : S \to X$. The distance function $d_\phi : S \times S \to [0, 1]$ is defined by

$$d_\phi (s_1, s_2) = d_X (\phi (s_1), \phi (s_2)).$$

One can easily verify that the pseudometric kernel d_ϕ is a pseudometric. Note that s_1 and s_2 have distance 0 if they are mapped by ϕ to the same element in X. For example, if ϕ is a constant function then all distances are 0.

In order to exploit a pseudometric kernel to provide the set S of states of a probabilistic transition system with a pseudometric, we need to introduce the pseudometric space X and the function ϕ. The former will be (the carrier of) the terminal P-coalgebra and the latter will be the unique P-homomorphism from the probabilistic transition system viewed as a P-coalgebra to the terminal P-coalgebra. The details will be provided below.

First, exploiting the pseudometric terminal coalgebra theorem, we prove that there exists a terminal P-coalgebra.

Theorem 27. *There exists a terminal P-coalgebra $\langle \text{fix} (P), i \rangle$.*

Proof. According to America and Rutten's [2, Theorem 5.4], the functors **1** and $+$ are locally nonexpansive and the scaling functor $c \cdot$ is locally contractive. As we have seen in Proposition 17, the functor M_t is locally nonexpansive. As a consequence, the functor P is locally contractive. According to Proposition 11 and Theorem 13, the functor M_t, and hence the functor P, preserves positivity and completeness. Therefore, we can conclude from Theorem 6 that there exists a terminal P-coalgebra $\langle \text{fix} (P), i \rangle$. $\quad\square$

Furthermore, the carrier of the terminal P-coalgebra is a compact metric space. We will exploit this property in Section 8.

Proposition 28. *fix (P) is a compact metric space.*

Proof. By Proposition 11 and Theorem 13, the functor M_t, and hence the functor P, preserves positivity and compactness. Hence, we can conclude from Theorem 7 that the metric space *fix* (P) is compact. $\quad\square$

Since $\langle \mathit{fix}\,(P), i\rangle$ is a terminal P-coalgebra, there exists a unique P-homomorphism ϕ from a P-coalgebra $\langle S, t\rangle$ to the terminal P-coalgebra.

$$
\begin{array}{ccc}
S & \xrightarrow{\ \ \phi\ \ } & \mathit{fix}\,(P) \\
{\scriptstyle t}\downarrow & & \downarrow{\scriptstyle i} \\
P\,(S) & \xrightarrow[P\,(\phi)]{} & P\,(\mathit{fix}\,(P))
\end{array}
$$

The pseudometric kernel d_ϕ induced by ϕ is a pseudometric on the set underlying the carrier S of the P-coalgebra. As each P-coalgebra $\langle S, t\rangle$ represents a continuous probabilistic transition system, having the set underlying S as its set of states, we thus obtain a pseudometric on this set of states. To stress its coalgebraic nature, instead of d_ϕ we will often write $d_{\mathcal{C}}$.

Since the identity map on $\mathit{fix}\,(P)$ is the unique P-homomorphism from the terminal P-coalgebra to itself, we can conclude that the coalgebraic pseudometric $d_{\mathcal{C}}$ on the set underlying the carrier of the terminal P-coalgebra coincides with the metric $d_{\mathit{fix}\,(P)}$ on the carrier of the terminal P-coalgebra.

In order to be able to explicitly compute some coalgebraic distances, we present a characterization of the pseudometric on $Q\,(S)$.

Proposition 29. *For all* $\mu_1, \mu_2 \in Q\,(S)$,

$$
d_{Q\,(S)}\,(\mu_1, \mu_2) = \sup\left\{\left.\int_S f\,\mathrm{d}\mu_1 - \int_S f\,\mathrm{d}\mu_2 \;\right|\; f \in c \cdot S \xrightarrow[1]{} [0, 1]\right\}
$$
$$
+(\mu_1\,(\mathbf{1}) \ominus \mu_2\,(\mathbf{1})),
$$

where

$$
r \ominus r' = \begin{cases} r - r' & \text{if } r \geqslant r', \\ 0 & \text{otherwise}. \end{cases}
$$

Proof.

$$
d_{Q\,(S)}\,(\mu_1, \mu_2)
$$
$$
= \sup\left\{\left.\int_{R\,(S)} f\,\mathrm{d}\mu_1 - \int_{R\,(S)} f\,\mathrm{d}\mu_2 \;\right|\; f \in R\,(S) \xrightarrow[1]{} [0, 1]\right\}
$$
$$
= \sup\left\{\left(r \cdot \mu_1\,(\mathbf{1}) + \int_S f\,\mathrm{d}\mu_1\right) - \left(r \cdot \mu_2\,(\mathbf{1}) + \int_S f\,\mathrm{d}\mu_2\right)\right.
$$
$$
\left. \mid r \in [0, 1] \wedge f \in c \cdot S \xrightarrow[1]{} [0, 1]\right\}
$$
$$
= \sup\left\{\left(\int_S f\,\mathrm{d}\mu_1 - \int_S f\,\mathrm{d}\mu_2\right) + r \cdot (\mu_1\,(\mathbf{1}) - \mu_2\,(\mathbf{1}))\right.
$$
$$
\left. \mid r \in [0, 1] \wedge f \in c \cdot S \xrightarrow[1]{} [0, 1]\right\}
$$
$$
= \sup\left\{\left.\int_S f\,\mathrm{d}\mu_1 - \int_S f\,\mathrm{d}\mu_2 \;\right|\; f \in c \cdot S \xrightarrow[1]{} [0, 1]\right\} + (\mu_1\,(\mathbf{1}) \ominus \mu_2\,(\mathbf{1})). \quad \square
$$

Once we can manage $d_{Q\,(S)}$, we can compute $d_{\mathcal{C}}$ as well.

Example 30. Consider the discrete probabilistic transition system introduced in Example 19. Let ϕ be the unique P-homomorphism from the P-coalgebra representing this system to the terminal P-coalgebra. Then

$$
\begin{aligned}
d_{\mathcal{C}} &(s_2, s_3) \\
&= d_{\text{fix}(P)}(\phi(s_2), \phi(s_3)) \\
&= d_{P(\text{fix}(P))}(i(\phi(s_2)), i(\phi(s_3))) \quad (i \text{ is an isometry}) \\
&= d_{P(\text{fix}(P))}(P(\phi)(t(s_2)), P(\phi)(t(s_3))) \quad (\phi \text{ is a } P\text{-homomorphism}) \\
&= \sup\left\{ \left| \int_{\text{fix}(P)} (f \circ \phi)\, \mathrm{d}t_{s_2,a} - \int_{\text{fix}(P)} (f \circ \phi)\, \mathrm{d}t_{s_3,a} \right| \; f \in c \cdot \text{fix}(P) \xrightarrow[1]{} [0,1] \right\} \\
&\quad + (t_{s_2,a}(\mathbf{1}) \ominus t_{s_3,a}(\mathbf{1})) \quad (\text{Proposition } 29) \\
&= \sup\left\{ f(\phi(s_3)) \mid f \in c \cdot \text{fix}(P) \xrightarrow[1]{} [0,1] \right\} + (0 \ominus 1) \\
&= 1.
\end{aligned}
$$

The rest of the distances can be computed in the same way. All of them are collected in the following table.

	s_0	s_0^ε	s_1	s_1^ε	s_2	s_3
s_0^ε	$c^2\varepsilon$					
s_1	$\frac{c^2+2c}{4}$	$\frac{c^2+2c}{4} + \varepsilon\frac{c^2}{2}$				
s_1^ε	$\frac{c^2+2c}{4} + \varepsilon\frac{c^2-2c}{2}$	$\frac{c^2+2c}{4} + \varepsilon((1+\varepsilon)c^2 - c)$	$c\varepsilon$			
s_2	c	c	$\frac{c}{2}$	$\frac{c}{2} + \varepsilon c$		
s_3	1	1	1	1	1	

The distance between states is a trade-off between the depth of observations needed to distinguish the states and the amount each observation differentiates the states. The relative weight given to these two factors is determined by c lying between 0 and 1: the smaller the value of c the greater the discount on observations made at greater depth. In particular, this is reflected by the fact that $d_{\mathcal{C}}(s_0, s_0^\varepsilon) = c \cdot d_{\mathcal{C}}(s_1, s_1^\varepsilon)$ in the above example.

Example 31. Consider the continuous probabilistic transition system with the set of states $[0, 1]$ and its Borel σ-algebra, a single action a, and the transition function determined by

$$
t_{s,a} = s\mu_0
$$

with the measure μ_0 as defined in Example 10. We have that

$$
d_{\mathcal{C}}(s, 1) = 1 - s.
$$

We conclude this section by showing that our pseudometric contains probabilistic bisimilarity.

Proposition 32. *Let $\langle S, t \rangle$ be a P-coalgebra representing a probabilistic transition system. Let S be an analytic space. States have distance* 0 *if and only if they are probabilistic bisimilar.*

Proof. For all $s_1, s_2 \in S$,

$d_{\mathcal{C}}(s_1, s_2) = 0$

iff $d_{\mathcal{L}}(s_1, s_2) = 0$ (Theorem 42)

iff s_1 and s_2 are probabilistic bisimilar

   (see [10, Corollary 6.1.6 and Theorem 6.1.10]).   □

7. A real-valued modal logic

We present a real-valued modal logic. This logic is closely related to the probabilistic modal logic of Larsen and Skou [24] and to a real-valued modal logic introduced by Desharnais et al. [11]. Along the lines of the latter paper, we define a pseudometric in terms of the logic. In the next section, we show that this pseudometric is the same (up to a fixed multiplying factor) as the one we introduced in Section 6.

Desharnais et al. defined a pseudometric in terms of a real-valued modal logic. Their work builds on ideas of Kozen [21] to generalize logic to handle probabilistic phenomena. In particular, the modality is interpreted as integration. A minor variation on their logic is introduced in the following definition.

Definition 33. The logic $\mathcal{L}$ is defined by

$$\varphi ::= 1 \mid \langle a \rangle\, \varphi \mid \min(\varphi, \varphi) \mid 1 - \varphi \mid \varphi \ominus q$$

where a is an action and q is a rational in [0, 1].

Informally, there is the following correspondence between formulae in $\mathcal{L}$ and formulae in the probabilistic modal logic of Larsen and Skou. True is represented by 1, conjunction is represented by min, negation by $1-$, and the modal connective $\langle a \rangle_q$ decomposes as $\langle a \rangle$ and $\ominus q$.

In analogy to one of De Morgan's laws, max can be expressed in the logic in terms of min and $1-$ as follows:

$$\max(\varphi, \psi) = 1 - \min(1 - \varphi, 1 - \psi).$$

Given a probabilistic transition system represented by the P-coalgebra $\langle S, t \rangle$, each formula φ can be interpreted as a function $\varphi_{\langle S, t \rangle}$ from S to [0, 1] as follows.

Definition 34. For each $\varphi \in \mathcal{L}$, the function $\varphi_{\langle S,t \rangle} : S \to [0, 1]$ is defined by

$$
\begin{aligned}
1_{\langle S,t \rangle}(s) &= 1, \\
(\langle a \rangle \varphi)_{\langle S,t \rangle}(s) &= c \cdot \textstyle\int_S \varphi_{\langle S,t \rangle} \, dt_{s,a}, \\
(\min(\varphi, \psi))_{\langle S,t \rangle}(s) &= \min(\varphi_{\langle S,t \rangle}(s), \psi_{\langle S,t \rangle}(s)), \\
(1 - \varphi)_{\langle S,t \rangle}(s) &= 1 - \varphi_{\langle S,t \rangle}(s), \\
(\varphi \ominus q)_{\langle S,t \rangle}(s) &= \varphi_{\langle S,t \rangle}(s) \ominus q.
\end{aligned}
$$

Next, we verify that for each formula φ, the function $\varphi_{\langle S,t \rangle}$ is c-contractive and hence measurable.

Proposition 35. *For all $\varphi \in \mathcal{L}$, the function $\varphi_{\langle S,t \rangle}$ is c-contractive.*

Proof. By structural induction on φ. We only consider the most interesting case.

$$
\begin{aligned}
&\left| (\langle a \rangle \varphi)_{\langle S,t \rangle}(s_1) - (\langle a \rangle \varphi)_{\langle S,t \rangle}(s_2) \right| \\
&= \left| c \cdot \int_S \varphi_{\langle S,t \rangle} \, dt_{s_1,a} - c \cdot \int_S \varphi_{\langle S,t \rangle} \, dt_{s_2,a} \right| \\
&= c \cdot \left| \int_S \varphi_{\langle S,t \rangle} \, dt_{s_1,a} - \int_S \varphi_{\langle S,t \rangle} \, dt_{s_2,a} \right| \\
&= c \cdot \max \left\{ \int_S \varphi_{\langle S,t \rangle} \, dt_{s_1,a} - \int_S \varphi_{\langle S,t \rangle} \, dt_{s_2,a}, \int_S \varphi_{\langle S,t \rangle} \, dt_{s_2,a} - \int_S \varphi_{\langle S,t \rangle} \, dt_{s_1,a} \right\} \\
&\leqslant c \cdot d_{Q(S)}(t_{s_1,a}, t_{s_2,a}) \\
&\quad \text{(Proposition 29, and } \varphi_{\langle S,t \rangle} \text{ is } c\text{-contractive by induction)} \\
&\leqslant c \cdot d_{P(S)}(t_{s_1}, t_{s_2}) \\
&\leqslant c \cdot d_S(s_1, s_2) \quad (t \text{ is nonexpansive}) \qquad \square
\end{aligned}
$$

The logic $\mathcal{L}$ induces a pseudometric as follows.

Definition 36. The distance function $d_{\mathcal{L}} : S \times S \to [0, 1]$ is defined by

$$
d_{\mathcal{L}}(s_1, s_2) = \sup_{\varphi \in \mathcal{L}} \varphi_{\langle S,t \rangle}(s_1) - \varphi_{\langle S,t \rangle}(s_2).
$$

Clearly, the above-introduced distance function is a pseudometric.

Our logic differs from the one presented by Desharnais et al. [11]. Instead of $\varphi \ominus q$ they write $\lfloor \varphi \rfloor_q$. Furthermore, they introduce $\lceil \varphi \rceil^q$. In the presence of negation, $\lceil \varphi \rceil^q$ is redundant as it is equivalent to $\min(\varphi, 1 - \lfloor 1 \rfloor_q)$. Finally, they introduce a countable supremum over formulae.

The logic considered by Desharnais [10] lacks negation, but does include $\lceil \varphi \rceil^q$ and max.

The presence of negation in our logic has an impact on the distances as is shown in:

Example 37. Consider the following probabilistic transition system.
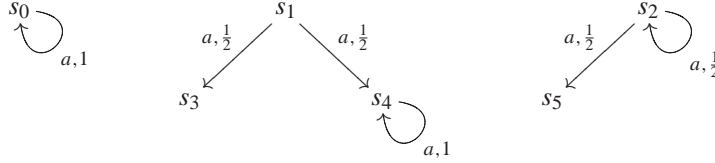
The system in state s_0 terminates with probability 0, in state s_1 with probability $\frac{1}{2}$ and in state s_2 with probability 1. The expected number of transitions to termination starting in state s_0, s_1 and s_2 is ∞, ∞, and 2, respectively. Based on these kind of observations, one may infer that state s_0 behaves more like state s_1 than state s_2. This is reflected by the pseudometric $d_{\mathcal{L}}$. For this example, we fix c to be $\frac{1}{2}$. Then the states s_0 and s_1 are $\frac{1}{8}$ apart, witnessed by $\langle a \rangle \langle a \rangle 1$. The states s_0 and s_2 are at distance $\frac{1}{6}$ which is witnessed by the formulae φ_n defined by

$$\varphi_n = \begin{cases} 1 & \text{if } n = 0, \\ 1 - \left((1 - \langle a \rangle \, \varphi_{n-1}) \ominus \frac{1}{2}\right) & \text{otherwise.} \end{cases}$$

However, in the pseudometric induced by the logic without negation both s_0 and s_1, and s_0 and s_2 are $\frac{1}{8}$ apart. In both cases, $\langle a \rangle \langle a \rangle 1$ is a witness.

To distinguish the set S endowed with the original pseudometric d_S from the set S endowed with the logical pseudometric $d_{\mathcal{L}}$, we denote the former space by $\langle S, d_S \rangle$ and the latter by $\langle S, d_{\mathcal{L}} \rangle$.

The interpretation $\varphi_{\langle S,t \rangle}$ is not only a c-contractive, and hence a nonexpansive, function from $\langle S, d_S \rangle$ to [0, 1] as we have shown in Proposition 35. It is also a nonexpansive function from $\langle S, d_{\mathcal{L}} \rangle$ to [0, 1] as we will show next.

Proposition 38. *For all $\varphi \in \mathcal{L}$, the function $\varphi_{\langle S,t \rangle}$ is nonexpansive with respect to $d_{\mathcal{L}}$.*

Proof. For all $s_1, s_2 \in S$,

$$\begin{aligned}
&|\varphi_{\langle S,t \rangle}(s_1) - \varphi_{\langle S,t \rangle}(s_2)| \\
&= \max \{\varphi_{\langle S,t \rangle}(s_1) - \varphi_{\langle S,t \rangle}(s_2), (1 - \varphi)_{\langle S,t \rangle}(s_1) - (1 - \varphi)_{\langle S,t \rangle}(s_2)\} \\
&\leqslant \sup_{\varphi \in \mathcal{L}} \varphi_{\langle S,t \rangle}(s_1) - \varphi_{\langle S,t \rangle}(s_2) \\
&= d_{\mathcal{L}}(s_1, s_2). \qquad \square
\end{aligned}$$

Each nonexpansive function from $\langle S, d_{\mathcal{L}} \rangle$ to [0, 1] can be approximated by interpretations of formulae of our logic $\mathcal{L}$ provided that the space $\langle S, d_S \rangle$ is compact.

Proposition 39. *If the pseudometric space $\langle S, d_S \rangle$ is compact then the set*

$$\{ \varphi_{\langle S,t \rangle} \mid \varphi \in \mathcal{L} \} \tag{3}$$

is dense in $\langle S, d_{\mathcal{L}} \rangle \underset{1}{\rightarrow} [0, 1]$.

Proof. Let $\delta > 0$ and $f \in \langle S, d_{\mathcal{L}} \rangle \underset{1}{\rightarrow} [0, 1]$. It suffices to show that there exists a formula φ in $\mathcal{L}$ such that f and $\varphi_{\langle S,t \rangle}$ are at most δ apart.

Below, we will exploit the following straightforward variation on the Stone–Weierstrass approximation theorem for continuous functions.

Lemma (*Ash [3, Lemma A.7.2]*). *Let X be a compact pseudometric space. Let A be a subset of $X \underset{1}{\rightarrow} [0, 1]$ such that $f_1, f_2 \in A$ implies $\min (f_1, f_2), \max (f_1, f_2) \in A$. If $f \in X \underset{1}{\rightarrow} [0, 1]$ can be approximated up to δ at each pair of points by functions in A then f itself can also be approximated up to δ by functions in A.*

Since for all $s_1, s_2 \in S$,

$$
\begin{aligned}
d_{\mathcal{L}} \,&(s_1, s_2) \\
&= \sup_{\varphi \in \mathcal{L}} \varphi_{\langle S,t \rangle} (s_1) - \varphi_{\langle S,t \rangle} (s_2) \\
&\leqslant c \cdot d_S (s_1, s_2) \quad \text{(Proposition 35)}
\end{aligned}
$$

and the space $\langle S, d_S \rangle$ is compact, we can conclude that $\langle S, d_{\mathcal{L}} \rangle$ is a compact pseudometric space. According to Proposition 38, the set (3) is a subset of $\langle S, d_{\mathcal{L}} \rangle \underset{1}{\rightarrow} [0, 1]$. Obviously, (3) is closed under min and max. Let $s_1, s_2 \in S$. Hence, according to Ash's lemma, it suffices to show that there exists a formula φ in $\mathcal{L}$ such that $f (s_i)$ and $\varphi_{\langle S,t \rangle} (s_i)$ are at most δ apart.

Without loss of generality, assume that $f (s_1) \geqslant f (s_2)$. Since

$$
\begin{aligned}
\Delta &= f (s_1) - f (s_2) \\
&\leqslant d_{\mathcal{L}} (s_1, s_2) \quad (f \text{ is nonexpansive}) \\
&= \sup_{\varphi \in \mathcal{L}} \varphi_{\langle S,t \rangle} (s_1) - \varphi_{\langle S,t \rangle} (s_2)
\end{aligned}
$$

there exists a formula φ such that $\Delta - \delta \leqslant \varphi_{\langle S,t \rangle} (s_1) - \varphi_{\langle S,t \rangle} (s_2)$. Let p, q and r be rationals in $[0, 1]$ such that

$$
\begin{aligned}
p &\in [\varphi_{\langle S,t \rangle} (s_2) - \delta, \varphi_{\langle S,t \rangle} (s_2)], \\
q &\in [\Delta - \delta, \Delta], \\
r &\in [f (s_2), f (s_2) + \delta].
\end{aligned}
$$

We leave it to the reader to verify that the formula

$$1 - ((1 - \min (\varphi \ominus p, 1 - (1 \ominus q))) \ominus r)$$

has the desired property. $\square$

Note that 1, min, max, $1-$ and $\ominus q$ all play a role in the above proof.

The interpretations of a formula with respect to different P-coalgebras are in general different. But whenever there is a P-homomorphism between P-coalgebras they are related as follows.

Proposition 40. *Let ϕ be a P-homomorphism from a P-coalgebra $\langle S, t \rangle$ to a P-coalgebra $\langle S', t' \rangle$. Then for all formulae φ,*

$$\varphi_{\langle S',t' \rangle} \circ \phi = \varphi_{\langle S,t \rangle}.$$

Proof. By structural induction on φ. We only present the most interesting case. For all $s \in S$,

$$(\langle a \rangle \varphi)_{\langle S',t' \rangle} (\phi(s))$$

$$= c \cdot \int_{S'} \varphi_{\langle S',t' \rangle} \, \mathrm{d}t'_{\phi(s),a}$$

$$= c \cdot \int_{S'} \varphi_{\langle S',t' \rangle} \, \mathrm{d}(P(\phi)(t))_{s,a} \quad [t' \circ \phi = P(\phi) \circ t]$$

$$= c \cdot \int_{S'} \varphi_{\langle S',t' \rangle} \, \mathrm{d}(t_{s,a} \circ \phi^{-1})$$

$$= c \cdot \int_{S} (\varphi_{\langle S',t' \rangle} \circ \phi) \, \mathrm{d}t_{s,a}$$

$$= c \cdot \int_{S} \varphi_{\langle S,t \rangle} \, \mathrm{d}t_{s,a} \quad \text{(induction)}$$

$$= (\langle a \rangle \varphi)_{\langle S,t \rangle} (s). \quad \square$$

Note that once you have the interpretation of a formula with respect to the terminal P-coalgebra, then you can infer it under any other P-coalgebra.

8. Relating the coalgebraic and logical distances

For a large class of probabilistic transition systems we have introduced a coalgebraic distance function $d_{\mathcal{C}}$ and a logical distance function $d_{\mathcal{L}}$. In this section we relate the two pseudometrics. Before considering the general case, we first relate the two distance functions on the set underlying the carrier of the terminal P-coalgebra. Recall that the coalgebraic pseudometric $d_{\mathcal{C}}$ on the set underlying the carrier of the terminal P-coalgebra coincides with the metric $d_{\text{fix}(P)}$ on the carrier of the terminal P-coalgebra.

Proposition 41. *For all $x_1, x_2 \in \text{fix}(P)$,*

$$\frac{d_{\mathcal{L}}(x_1, x_2)}{c} = d_{\mathcal{C}}(x_1, x_2).$$

Proof. Consider the function ι which maps each $x \in \mathit{fix}\,(P)$ to itself. For all $x_1, x_2 \in \mathit{fix}\,(P)$,

$$\frac{d_{\mathcal{L}}\,(\iota\,(x_1),\,\iota\,(x_2))}{c}$$

$$= \frac{d_{\mathcal{L}}\,(x_1, x_2)}{c}$$

$$= \frac{\sup_{\varphi \in \mathcal{L}}\;\varphi_{\langle \mathit{fix}\,(P), i \rangle}\,(x_1) - \varphi_{\langle \mathit{fix}\,(P), i \rangle}\,(x_2)}{c}$$

$$\leqslant d_{\mathcal{C}}\,(x_1, x_2) \quad (\varphi_{\langle \mathit{fix}\,(P), i \rangle} \text{ is } c\text{-contractive by Proposition 35}).$$

Consequently, ι is a nonexpansive function from the space $\langle \mathit{fix}\,(P), d_{\mathcal{C}} \rangle$ to the space $\langle \langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle$.

Next, we introduce a structure t such that $\langle \langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle, t \rangle$ is a P-coalgebra. Because ι is nonexpansive, each Borel set of $R\,\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle$ is also a Borel set of $R\,\langle \mathit{fix}\,(P), d_{\mathcal{C}} \rangle$. Therefore, we can take the function t given by

$$t_{x,a}\,(B) = i_{x,a}\,(B)$$

for $x \in \mathit{fix}\,(P)$, $a \in \mathit{Act}$ and Borel set B of $R\,\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle$. Since the function ι is nonexpansive and the measure $i_{x,a}$ is tight, we can conclude that the measure $t_{x,a}$ is tight as well (cf. Proposition 15).

To conclude that t is the structure of a P-coalgebra with carrier $\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle$, we have left to show that t is nonexpansive. Let $x_1, x_2 \in \mathit{fix}\,(P)$. Then

$$d_{P\,\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle}\,(t_{x_1}, t_{x_2}) = \sup_{a \in \mathit{Act}}\;d_{Q\,\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle}\,(t_{x_1,a}, t_{x_2,a}).$$

Let $a \in \mathit{Act}$. Without loss of generality, assume that $t_{x_1,a}\,(\mathbf{1}) \leqslant t_{x_2,a}\,(\mathbf{1})$. Then,

$$d_{Q\,\langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle}\,(t_{x_1,a}, t_{x_2,a})$$

$$= \sup \left\{ \int_{\mathit{fix}\,(P)} f \; \mathrm{d}t_{x_1,a} - \int_{\mathit{fix}\,(P)} f \; \mathrm{d}t_{x_2,a} \;\Big|\; f \in c \cdot \langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle \xrightarrow{\;1\;} [0,1] \right\}$$

$$\text{(Proposition 29)}$$

$$\leqslant \sup_{\varphi \in \mathcal{L}} \int_{\mathit{fix}\,(P)} \varphi_{\langle \mathit{fix}\,(P), i \rangle} \; \mathrm{d}t_{x_1,a} - \int_{\mathit{fix}\,(P)} \varphi_{\langle \mathit{fix}\,(P), i \rangle} \; \mathrm{d}t_{x_2,a}$$

$$\text{(Proposition 28 and 39)}$$

$$= \frac{\sup_{\varphi \in \mathcal{L}}\,(\langle a \rangle\,\varphi)_{\langle \mathit{fix}\,(P), i \rangle}\,(x_1) - (\langle a \rangle\,\varphi)_{\langle \mathit{fix}\,(P), i \rangle}\,(x_2)}{c}$$

$$\leqslant \frac{d_{\mathcal{L}}\,(x_1, x_2)}{c}.$$

From the definition of t and ι we can easily derive that ι is a P-homomorphism from the P-coalgebra $\langle \langle \mathit{fix}\,(P), d_{\mathcal{C}} \rangle, i \rangle$ to the P-coalgebra $\langle \langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle, t \rangle$. We denote the unique P-homomorphism from the P-coalgebra $\langle \langle \mathit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle, t \rangle$ to the terminal P-coalgebra $\langle \langle \mathit{fix}\,(P), d_{\mathcal{C}} \rangle, i \rangle$ by ϕ.

$$
\begin{array}{ccc}
\langle \textit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle & \xrightleftharpoons[\iota]{\phi} & \langle \textit{fix}\,(P), d_{\mathcal{C}} \rangle \\
\downarrow t & & \downarrow i \\
P\,\langle \textit{fix}\,(P), \frac{d_{\mathcal{L}}}{c} \rangle & \xrightleftharpoons[P\,(\iota)]{P\,(\phi)} & P\,\langle \textit{fix}\,(P), d_{\mathcal{C}} \rangle
\end{array}
$$

Obviously, the identity map on $\textit{fix}\,(P)$ is the unique P-homomorphism from the terminal P-coalgebra $\langle \langle \textit{fix}\,(P), d_{\mathcal{C}} \rangle, i \rangle$ to itself. Since $\phi \circ \iota$ is also such a P-homomorphism, we can conclude that $\phi \circ \iota$ equals the identity map on $\textit{fix}\,(P)$. Therefore, both ϕ and ι are isometries. This observation completes the proof. $\square$

Note that in the above proof the modality $\langle a \rangle$ is used, together with Proposition 39 in whose proof 1, min, max, $1-$ and $\ominus q$ all play a role.

Next, we consider the general case where we have a probabilistic transition system represented by the P-coalgebra $\langle S, t \rangle$. Then we have the following:

Theorem 42. *For all $s_1, s_2 \in S$,*

$$
\frac{d_{\mathcal{L}}\,(s_1, s_2)}{c} = d_{\mathcal{C}}\,(s_1, s_2).
$$

Proof. We denote the unique P-homomorphism from the P-coalgebra $\langle S, t \rangle$ to the terminal P-coalgebra $\langle \textit{fix}\,(P), i \rangle$ by ϕ. For all $s_1, s_2 \in S$,

$$
\frac{d_{\mathcal{L}}\,(s_1, s_2)}{c} \tag{4}
$$

$$
= \frac{d_{\mathcal{L}}\,(\phi\,(s_1), \phi\,(s_2))}{c} \quad \text{(Proposition 40)} \tag{5}
$$

$$
= d_{\mathcal{C}}\,(\phi\,(s_1), \phi\,(s_2)) \quad \text{(Proposition 41)} \tag{6}
$$

$$
= d_{\mathcal{C}}\,(s_1, s_2). \tag{7}
$$

Note that (4) and (5) refer to different logical pseudometrics: the one on S and the one on $\textit{fix}\,(P)$, respectively. Also notice that (6) and (7) refer to different coalgebraic pseudometrics: the one on $\textit{fix}\,(P)$ and the one on S, respectively. $\square$

In [9], we studied a minor variation on the functor P. In that paper, we considered the functor

$$
P' = c \cdot M_t\,(\mathbf{1} + -)^{Act}.
$$

This functor is also locally contractive and preserves positivity and completeness and, therefore, has a terminal coalgebra. The carriers of the terminal P-coalgebra and the terminal P'-coalgebra are related as follows.

Proposition 43. $\mathit{fix}\,(P') = c \cdot \mathit{fix}\,(P)$.

Proof. According to Lemma 4, there exists an isometry i from $\mathit{fix}\,(P)$ to $P\,(\mathit{fix}\,(P))$. Clearly, i is also an isometry from $c \cdot \mathit{fix}\,(P)$ to $c \cdot P\,(\mathit{fix}\,(P))$ which equals $P'\,(c \cdot \mathit{fix}\,(P))$. Using Lemma 4 again, we can conclude that $\mathit{fix}\,(P') = c \cdot \mathit{fix}\,(P)$. □

Consequently, the coalgebraic pseudometric induced by the functor P' coincide with the logical pseudometric.

9. Conclusion

9.1. Related work

As we have already seen in Sections 7 and 8, our coalgebraic pseudometric is closely related to the logical pseudometric of Desharnais et al. [10,11]. In [11], they also introduce a probabilistic process algebra. A number of combinators of the process algebra, including probabilistic choice, are shown to be nonexpansive. This is a quantitative analogue of probabilistic bisimulation being a congruence. It allows for compositional verification of probabilistic transition systems. Since our coalgebraic pseudometric is related to their logical pseudometric, we can conclude that those combinators are also nonexpansive with respect to our pseudometric. Furthermore, Desharnais et al. present a decision procedure for their pseudometric. That is, they provide an algorithm to approximate the logical distances to a prescribed degree of accuracy. The algorithm involves the generation of a representative set of formulae of their real-valued modal logic. They only consider formulae with a restricted number of nested occurrences of the modal connective. Their algorithm approximates the distances in exponential time. In [8], we present an algorithm to approximate our coalgebraic distances. The problem of approximating such distances can be reduced to a particular linear programming problem: the transportation problem. Since the latter problem can be solved in polynomial time, we obtain a polynomial time decision procedure for our distances. We see this practical algorithm as one of the advantages of our coalgebraic approach over the logical approach of Desharnais et al. Another advantage of our approach is that we work within a uniform framework, the theory of coalgebras. We do not know whether there exists a terminal coalgebra of our functor for c equals 1, and hence we cannot use our framework to define a pseudometric when c equals 1. However, the logical approach of Desharnais et al. also works in that case. Furthermore, Desharnais et al. consider a larger class of continuous probabilistic transition systems than we do in this paper. However, we are confident that we can extend our results as we will discuss below. In conclusion, we believe that both approaches have their merits and demerits. The results in Section 8 are very valuable as they allows us to transfer results from one setting to the other.

As far as we know, [14] by Giacalone et al. is the first paper to advocate the use of pseudometric spaces to provide a robust and quantitative notion of behavioural equivalence. They stress the importance of combinators being nonexpansive with respect to the pseudometric, making compositional verification possible. The class of discrete probabilistic transition systems they consider is rather restricted. A decade later, we are able to deal with all

discrete probabilistic transition systems and a large class of continuous probabilistic transition systems.

De Vink and Rutten [33] show that discrete probabilistic transition systems and some continuous probabilistic transition systems can be viewed as coalgebras. Their main contribution is the proof that the kernel of the homomorphism from a coalgebra, representing a probabilistic transition system, to the terminal coalgebra coincides with probabilistic bisimilarity. They only exploit metrics to represent continuous systems as coalgebras. Their metric on the terminal coalgebra only provides qualitative information. For example, in De Vink and Rutten's setting the states s_0 and s_0^ε of the system presented in the introduction are c apart if ε differs from 0. More generally, the distance between two states in their setting is c^n where n is the depth of a probabilistic bisimulation between them. De Vink and Rutten consider the endofunctor

$$(\mathbf{1} + M_c\,(c \cdot -))^{Act}$$

on the category of complete ultrametric spaces and nonexpansive functions. M_c denotes the Borel probability measures with compact support. The main differences between our functor and their functor are the following. First of all, they consider a distance function on Borel probability measures [33, Definition 5.3] different from the one presented in Definition 8. Their distance function only captures qualitative information as the above example illustrates. Secondly, they consider the category of complete ultrametric spaces and nonexpansive functions whereas we consider the considerably larger category of pseudometric spaces and nonexpansive functions. This allows us to capture many more interesting continuous probabilistic transition systems as coalgebras, including systems where the state space is the real interval [0, 1] endowed with the Euclidean metric. Furthermore, they consider Borel probability measures with compact support whereas we consider the more general tight Borel probability measures. Again this allows us to represent more systems as coalgebras. Finally, their model only allows states to refuse transitions with probability 0 or 1. In conclusion, our functor allows to model many more interesting continuous systems, and all the results for their functor in [33, Section 5] [3] can be generalized to our setting.

Baier and Kwiatkowska [4] study a functor which is closely related to the one of De Vink and Rutten. Our work can be compared to theirs in the same way it is compared to the work of De Vink and Rutten in the paragraph above.

In his thesis [17], Den Hartog exploits ultrametric spaces very similar to the terminal coalgebra of De Vink and Rutten. The metric structure is only used to model infinite behaviour. As a consequence, qualitative information suffices. We believe that metrics closely related to the one we present in this paper may be used in his setting as well, possibly providing additional quantitative information about his models.

Kwiatkowska and Norman [22,23,26] present a number of closely related metrics. Like Den Hartog, they use their metric as a means to model recursion. However, their metric is not an ultrametric and contains quantitative information. Let us compare the metric introduced

---

[3] The proof of [33, Theorem 5.8] is incomplete. We also have no proof for this result in our setting.

by Norman in [26, Section 6.1] with our pseudometric. Consider the following probabilistic transition system.
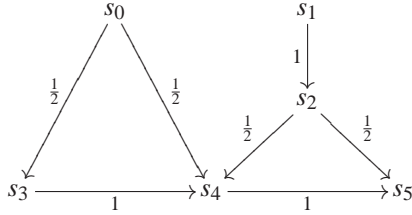
Clearly, the states s_0 and s_1 are not probabilistic bisimilar. In Norman's setting the states have distance 0. In our pseudometric, states only have distance 0 if they are probabilistic bisimilar. In our setting the states are $\frac{c^2}{2}$ apart. This example shows that his distance function gives rise to a topology different from ours. The main differences between his and our approach are the following. First of all, he uses a linear-time model whereas we consider a branching-time model. Secondly, he only handles discrete systems whereas we also consider continuous ones. Finally, we use the usual categorical machinery and various standard constructions whereas his definitions are more ad hoc. We believe however that his metric can also be characterized by means of a terminal coalgebra.

Results similar to the ones in this paper have been presented by the second author in his thesis [34, Chapter 4] in the setting of bimodules and generalized metric spaces. The coalgebraic distance of states s_1 and s_2 can be characterized as the smallest $R(s_1, s_2)$ where R is a bimodule satisfying certain conditions (see [34, Theorem 4.5.12] for the details). This is the quantitative analogue of the characterization of probabilistic bisimilarity as the largest probabilistic bisimulation.

9.2. Future work

Let us isolate two distinct consequences of our use of the pseudometric presented in Section 3. First of all, we can talk about approximate equivalence of states. Secondly, we can model a large class of continuous probabilistic transition systems as coalgebras. An apparent restriction with regard to the latter point is the requirement that the structure of a P-coalgebra, that is, the system's transition function, be nonexpansive. Properly speaking, continuous probabilistic transition systems as formulated in Definition 21 are coalgebras of (a variant of) the Giry monad on the category of measurable spaces and measurable functions [15]. However, we conjecture that the terminal P-coalgebra $\langle fix(P), i \rangle$ is also terminal when seen as a coalgebra of the Giry functor, and that our results can be extended to continuous probabilistic transition systems in general.

In Proposition 28 we have shown that the carrier of our terminal coalgebra is compact and hence separable. Furthermore, we conjecture that the unique homomorphism from the initial algebra of a finitary version of P—this finitary version represents finite discrete probabilistic transition systems with rational probabilities—to the terminal P-coalgebra is a dense embedding. Hence, every continuous system can be approximated by a finite one (see also [12]).

Acknowledgements

References

[1] F. Alessi, P. Baldan, G. Bellè, A fixed-point theorem in a category of compact metric spaces, Theoret. Comput. Sci. 146 (1/2) (1995) 311–320.

[2] P. America, J.J.M.M. Rutten, Solving reflexive domain equations in a category of complete metric spaces, J. Comput. System Sci. 39 (3) (1989) 343–375.

[3] R.B. Ash, Real Analysis and Probability, Probability and Mathematical Statistics, Vol. 11, Academic Press, London, 1972.

[4] C. Baier, M. Kwiatkowska, Domain equations for probabilistic processes, Math. Structures Comput. Sci. 10 (6) (2000) 665–717.

[5] M.F. Barnsley, Fractals Everywhere, second ed., Academic Press, Boston, 1993.

[6] P. Billingsley, Probability and Measure, Series in Probability and Mathematical Statistics, third ed., Wiley, New York, 1995.

[7] R. Blute, J. Desharnais, A. Edalat, P. Panangaden, Bisimulation for labelled Markov processes, in: Proc. 12th Annual IEEE Symp. on Logic in Computer Science, Warsaw, IEEE, New York, June/July 1997, pp. 149–158.

[8] F. van Breugel, J. Worrell, An algorithm for quantitative verification of probabilistic systems, in: K.G. Larsen, M. Nielsen (Eds.), Proc. 12th Internat. Conf. on Concurrency Theory, Lecture Notes in Computer Science, Vol. 2154, Aalborg, Springer, Berlin, August 2001, pp. 336–350.

[9] F. van Breugel, J. Worrell, Towards quantitative verification of probabilistic systems, in: F. Orejas, P.G. Spirakis, J. van Leeuwen (Eds.), Proc. 28th Internat. Colloquium on Automata, Languages and Programming, Lecture Notes in Computer Science, Vol. 2076, Crete, Springer, Berlin, July 2001, pp. 421–432.

[10] J. Desharnais, Labelled Markov Processes, Ph.D. Thesis, McGill University, Montreal, November 1999.

[11] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Metrics for labeled Markov systems, in: J.C.M. Baeten, S. Mauw (Eds.), Proc. 10th Internat. Conf. on Concurrency Theory, Lecture Notes in Computer Science, Vol. 1664, Eindhoven, Springer, Berlin, August 1999, pp. 258–273.

[12] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Approximating labeled Markov processes, in: Proc. 15th Annual IEEE Symp. on Logic in Computer Science, Santa Barbara, IEEE, New York, June 2000, pp. 95–106.

[13] G.A. Edgar, Integral, Probability, and Fractal Measures, Springer, Berlin, 1998.

[14] A. Giacalone, C.-C. Jou, S.A. Smolka, Algebraic reasoning for probabilistic concurrent systems, in: Proc. IFIP WG 2.2/2.3 Working Conf. on Programming Concepts and Methods, Sea of Gallilee, North-Holland, Amsterdam, April 1990, pp. 443–458.

[15] M. Giry, A categorical approach to probability theory, in: B. Banaschewski (Ed.), Proc. Internat. Conf. on Categorical Aspects of Topology and Analysis, Lecture Notes in Mathematics, Vol. 915, Ottawa, Springer, Berlin, August 1981, pp. 68–85.

[16] R.J. van Glabbeek, S.A. Smolka, B. Steffen, Reactive, generative and stratified models of probabilistic processes, Inform. and Comput. 121 (1) (1995) 59–80.

[17] J.I. den Hartog, Probabilistic Extensions of Semantical Models, Ph.D. Thesis, Free University, Amsterdam, October 2002.

[18] J.E. Hutchinson, Fractals and self similarity, Indiana Univ. Math. J. 30 (5) (1981) 713–747.

[19] M. Huth, M. Kwiatkowska, Quantitative analysis and model checking, in: Proc. 12th Annual IEEE Symp. on Logic in Computer Science, Warsaw, IEEE, New York, June/July 1997, pp. 111–122.

[20] L.V. Kantorovich, G.S. Rubinstein, On a functional space and certain extremum problems, Dokl. Akad. Nauk SSSR 115 (1957) 1058–1061.

[21] D. Kozen, A probabilistic PDL, J. Comput. System Sci. 30 (2) (1985) 162–178.

[22] M. Kwiatkowska, G.J. Norman, Probabilistic metric semantics for a simple language with recursion, in: W. Penczek, A. Szalas (Eds.), Proc. 21st Internat. Symp. on Mathematical Foundations of Computer Science, Lecture Notes in Computer Science, Vol. 1113, Cracow, Springer, Berlin, September 1996, pp. 419–430.

[23] M. Kwiatkowska, G.J. Norman, A fully abstract metric-space denotational semantics for reactive probabilistic processes, in: A. Edalat, A. Jung, K. Keimel, M.Z. Kwiatkowska (Eds.) Proc. third Workshop on Computation and Approximation, Electronic Notes in Theoretical Computer Science, Vol. 13, Birmingham, Elsevier, Amsterdam, September 1997.

[24] K.G. Larsen, A. Skou, Bisimulation through probabilistic testing, Inform. and Comput. 94 (1) (1991) 1–28.

[25] S. Mac Lane, Categories for the Working Mathematician, Graduate Texts in Mathematics, Vol. 5, Springer, Berlin, 1971.

[26] G.J. Norman, Metric Semantics for Reactive Probabilistic Systems, Ph.D. Thesis, University of Birmingham, November 1997.

[27] K.R. Parthasarathy, Probability Measures on Metric Spaces, Academic Press, New York, 1967.

[28] S.T. Rachev, Probability Metrics and the Stability of Stochastic Models, Wiley, Chichester, 1991.

[29] J.J.J.M. Rutten, Universal coalgebra: a theory of systems, Theoret. Comput. Sci. 249 (1) (2000) 3–80.

[30] W.A. Sutherland, Introduction to Metric and Topological Spaces, Clarendon Press, Oxford, 1975.

[31] D. Turi, J.J.J.M. Rutten, On the foundations of final semantics: non-standard sets, metric spaces, partial orders, Math. Structures Comput. Sci. 8 (5) (1998) 481–540.

[32] L.N. Vaserstein, Markov processes on countable space products describing large systems of automata, Problems Inform. Transmission 5 (3) (1969) 47–52.

[33] E.P. de Vink, J.J.M.M. Rutten, Bisimulation for probabilistic transition systems: a coalgebraic approach, Theoret. Comput. Sci. 221 (1/2) (1999) 271–293.

[34] J. Worrell, On Coalgebras and Final Semantics, Ph.D. Thesis, Oxford University, October 2000.