



An Algorithm for Quantitative Verification of Probabilistic Transition Systems

Franck van Breugel and James Worrell

Technical Report CS-2001-01

April 2001

Department of Computer Science

4700 Keele Street, Toronto, Ontario M3J 1P3, Canada

An Algorithm for Quantitative Verification of Probabilistic Transition Systems

Franck van Breugel*

York University, Department of Computer Science
4700 Keele Street, Toronto, Canada M3J 1P3

franck@cs.yorku.ca


James Worrell

Tulane University, Department of Mathematics
6823 St Charles Avenue, New Orleans LA 70118

jbw@math.tulane.edu

Introduction

In this paper, we consider *probabilistic transition systems*. We view a probabilistic transition system as consisting of a set of states and a set of labels. At any one time the system is in a particular state. When the environment selects an action (indicated by a label) the system reacts by either refusing the action or making a transition to a new state. This reaction—whether to terminate or which state to move to—is determined by a probability measure. We restrict ourselves to probabilistic transition systems without labels to simplify the presentation. However, all our results can easily be generalized to a setting with labels. Some details will be provided in the concluding section. Consider the following system.




Observe, for example, that in state s_1^0 the transitions to states s_2^0 and s_3^1 are each taken with probability $\frac{1}{2}$. Also, note that the probability of termination in, for example, state s_2^0 is $\frac{1}{2}$.

Larsen and Skou [LS91] introduced *probabilistic bisimulation* as a notion of behavioural equivalence for such systems. Briefly, a probabilistic bisimulation is an equivalence relation such that for any two related states their probability of making a transition to any equivalence class is equal. In the above system, states with the same colour/subscript are behaviourally equivalent. One problem with this notion is that it is not robust: two states are either bisimilar or they are not bisimilar, and a slight change in the probabilities associated to a system can cause bisimilar states to become non-bisimilar and vice-versa. Consider, for

* Supported by Natural Sciences and Engineering Research Council of Canada.

example, the following system.



The states s_0^0 and s_0^ϵ are only bisimilar if ϵ is 0. However, the states give rise to almost the same behaviour for very small ϵ different from 0.

To address this problem Giacalone, Jou and Smolka [GJS90] suggested to define a *pseudometric* on the states of a probabilistic transition system. This yields a smooth, *quantitative* notion of behavioural equivalence. A pseudometric differs from an ordinary metric in that different elements, i.e., states, can have distance 0. The distance between states, a real number between 0 and 1, can be used to express the similarity of the behaviour of the system started in those states. The smaller the distance, the more alike the behaviour is. In particular, the distance between states is 0 if they are indistinguishable, i.e., probabilistic bisimilar. In [BW01], we presented such a pseudometric. For example, in our pseudometric the states s_0^0 and s_0^ϵ are $CF \cdot \epsilon$ apart where CF is a constant between 0 and 1 (we will discuss the role of this constant later). Desharnais, Gupta, Jagadeesan and Panangaden [DGJP99, Des99], Giacalone et al. [GJS90], and Kwiatkowska and Norman [KN96, Nor97, KN98] also presented pseudometrics for probabilistic systems. For a comparison of these pseudometrics with ours, we refer the reader to [BW01].

At CONCUR'99, Desharnais et al. [DGJP99] presented an algorithm to calculate their distances up to a prescribed degree of accuracy. We have shown in [BW01] that our pseudometric is very closely related to theirs. In this paper, we present an algorithm to approximate our pseudometric. Whereas their algorithm approximates the distances in exponential time, we will show that our algorithm calculates the distances in polynomial time. The key ingredients of our pseudometric and algorithm are coalgebras, the Hutchinson metric and linear programming. We will discuss these ingredients next.

Many different kinds of transition system can be viewed as *coalgebras*; Rutten [Rut00] provides numerous examples. De Vink and Rutten [VR99] have shown that probabilistic transition systems correspond to P' -coalgebras, where P' is an endofunctor on the category of 1-bounded complete ultrametric spaces and nonexpansive functions. Furthermore, they have proved that the functor P' is locally contractive. Hence, according to Rutten and Turi's (ultra)metric terminal coalgebra theorem [RT92], there exists a *terminal* P' -coalgebra. By definition, there is a unique map from an arbitrary P' -coalgebra, i.e., a probabilistic transition system, to the terminal P' -coalgebra. De Vink and Rutten have also shown that the *kernel* of this unique map is probabilistic bisimilarity on the states of the probabilistic transition system. That is, two states are mapped to the same element in the terminal P' -coalgebra by the unique map if and only if they are probabilistic bisimilar.

In this paper, we study a variation on the endofunctor P' . Our endofunctor P on the category \mathcal{CMet}_1 of 1-bounded complete metric spaces and nonexpansive function is based on the *Hutchinson metric* on probability measures. This metric arises in very different contexts including statistics and fractal geometry, and under different names including the Kantorovich metric and the Wasserstein metric. Like P' -coalgebras, also P -coalgebras can be seen as probabilistic transition systems, as we will show. Furthermore, we will prove that the functor P is locally contractive as well. Hence, there exists a terminal P -coalgebra. Since the terminal P -coalgebra carries a metric, we can also consider the *metric kernel* of the unique map from a P -coalgebra to the terminal P -coalgebra. This is a pseudometric on the carrier of the P -coalgebra. The distance between two states of a P -coalgebra, i.e., a probabilistic transition system, is the distance in the terminal P -coalgebra of their images under the unique map. Since our functor is similar to the one considered by De Vink and Rutten, we still have that two states are bisimilar if and only if they are mapped to the same element in the terminal P -coalgebra and hence have distance 0.

As Rutten and Turi [RT92] have shown, the unique map from an F -coalgebra to the terminal F -coalgebra, where F is a locally contractive endofunctor on the category \mathcal{CMet}_1 , can be defined as the unique fixed point $fix(\Phi)$ of a function Φ from a complete metric space to itself. Since the functor F is locally contractive, the function Φ is contractive. Hence, according to Banach's fixed point theorem, Φ has a unique fixed point

$fix(\Phi)$. This fixed point can be approximated by a sequence of functions $(\phi_n)_n$. The function ϕ_0 is an arbitrary function from the F -coalgebra to the terminal F -coalgebra and the other functions are defined by $\phi_n = \Phi(\phi_{n-1})$. Not only the metric kernel of the unique map $fix(\Phi)$ defines a pseudometric $d_{fix(\Phi)}$ on the carrier of the F -coalgebra. Also the metric kernels of the approximations ϕ_n induce pseudometrics d_{ϕ_n} . We will show that the pseudometric $d_{fix(\Phi)}$ can be approximated by the pseudometrics d_{ϕ_n} . In particular, to calculate the $d_{fix(\Phi)}$ -distances to a prescribed degree of accuracy α , we only have to calculate the $\phi_1, \dots, \phi_{\log_{\text{CF}(\frac{\alpha}{2})}}$ -distances.

Next, we discuss how to compute the distance $d_{\phi_n}(s_i, s_j)$, where s_i and s_j are elements of the carrier of the P -coalgebra, i.e., states of the probabilistic transition system. We will show that this problem can be reduced a particular linear programming problem: the *transshipment problem*. For a detailed discussion of this problem and algorithms which can solve this problem in polynomial time we refer the reader to, for example, Chvátal's textbook [Chv83]. The transshipment problem is to find the cheapest way to ship a prescribed amount of a commodity from specified origins to specified destinations through a concrete transportation network. This network is represented by a directed graph. There is a demand for some commodity at some nodes and a supply (or negative demand) of some commodity at other nodes. In our setting, the total supply equals the total demand. This allows for a simpler and more elegant solution. With each edge, we associate the cost of shipping a unit amount along the edge.

The rest of this paper is organized as follows. In Section 1, we introduce probabilistic transition systems and probabilistic bisimilarity. Our algorithm is presented in Section 2. The other sections contain its correctness proof. Coalgebras and Rutten and Turi's terminal coalgebra theorem are presented in Section 3. In Section 4, we study metric kernels. The Hutchinson metric and its extension to a functor are the topics of Section 5. In Section 6, we focus on metric kernels for P -coalgebras. The reduction of the calculation of the d_{ϕ_n} -distances to the transshipment problem is presented in Section 7. In Section 8, we compare our algorithm with the one of Desharnais et al. and we discuss some future work.

Acknowledgements

The first author would like to thank Jeff Edmonds for stimulating discussions about linear programming.

1 Probabilistic Transition Systems

In this section, we introduce the probabilistic analog of transition systems. Furthermore, we present a behavioural equivalence on the states of these systems.

Probabilistic transition systems are introduced in

DEFINITION 1 A *probabilistic transition system* consists of a finite set S of states together with a transition function $\pi : S \times S \rightarrow [0, 1]$ such that $\sum_{s' \in S} \pi(s, s') \leq 1$ for all $s \in S$. \square

The transition function π is a conditional probability: $\pi(s, -)$ assigns to each state s' the probability of making a transition to state s' given that it was in state s before the transition. On the transition function we impose the restriction $\sum_{s' \in S} \pi(s, s') \leq 1$ instead of the more common, but also more restrictive, condition $\sum_{s' \in S} \pi(s, s') = 1$ or 0—the latter corresponding to termination. We interpret $1 - \sum_{s' \in S} \pi(s, s')$ as the probability that the system terminates in state s . To simplify our presentation we add a special state $\mathbf{0}$ for termination: $\pi(s, \mathbf{0}) = 1 - \sum_{s' \in S} \pi(s, s')$.

Larsen and Skou [LS91] adapted bisimulation for probabilistic transition systems as follows.

DEFINITION 2 Let $\langle S, \pi \rangle$ be a probabilistic transition system. An equivalence relation \mathcal{R} on the set of states S is a *probabilistic bisimulation* if $s_1 \mathcal{R} s_2$ implies $\sum_{s' \in E} \pi(s_1, s') = \sum_{s' \in E} \pi(s_2, s')$ for all \mathcal{R} -equivalence classes E . States s_1 and s_2 are *probabilistic bisimilar* if $s_1 \mathcal{R} s_2$ for some probabilistic bisimulation \mathcal{R} . \square

This notion has been shown to be canonical both from the coalgebraic point of view, as demonstrated by De Vink and Rutten in [VR99], and in the setting of open maps, as illustrated by Cheng and Nielsen [CN95].

2 The Algorithm

We present our algorithm to calculate the distances between the states of a probabilistic transition system up to a prescribed degree of accuracy.

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 the constant CF lying between 0 and 1: the smaller the value of CF the greater the discount on observations made at greater depth. The constant CF corresponds to the contraction factor of our functor P which we will introduce in Section 5. This constant also appears in the work of Desharnais et al. It plays a crucial role in the metrics used to model programming languages as well (see, for example, the textbook [BV96] of De Bakker and De Vink).

For the second system presented in the introduction, the distances are given in the table below. Since the states with the same colour/subscript (different from 0) are probabilistic bisimilar, they have distance 0 and are therefore identified.

	s_0^0	s_0^ϵ	s_1
s_0^ϵ	$\epsilon \cdot CF$		
s_1	$\frac{1}{2} \cdot CF$	$(\frac{1}{2} + \epsilon) \cdot CF$	
s_2	1	1	1

Next, we present the algorithm

Distance(*probability, accuracy*)

Input:

- probability matrix: for all $0 \leq i < N$ and $0 \leq j < N$, *probability*[i][j] is the probability of going from state s_i to state s_j ;
- accuracy of the distance function to be computed: the actual distances are greater than or equal to the computed distances and differ by at most *accuracy* from the computed distances.

Output:

- distance matrix: for all $0 \leq i < N$ and $0 \leq j < N$, *distance*[i][j] is the computed distance from state s_i to state s_j .

```

(1)  for k ← 0 to N do
(2)    for l ← 0 to N2 + N - 1 do
(3)      if k = (l + 1 + (l div (N + 1))) mod (N + 1) then
(4)        matrix[k][l] ← 1
(5)      else if k = l div N then
(6)        matrix[k][l] ← -1
(7)      else
(8)        matrix[k][l] ← 0
(9)  for i ← 0 to N - 1 do
(10)   stop[i] ← 1
(11)  for j ← 0 to N - 1 do
(12)   stop[i] ← stop[i] - probability[i][j]
(13)  for i ← 0 to N - 1 do
(14)   for j ← 0 to N - 1 do
(15)    for k ← 0 to N - 1 do
(16)     demand[i][j][k] ← probability[i][k] - probability[j][k]
```

```

(17)      demand[i][j][N] ← stop[i] − stop[j]
(18)  for i ← 0 to N − 1 do
(19)      for j ← 0 to N − 1 do
(20)          distance[i][j] ← 0
(21)  for n ← 0 to ⌈logCF(accuracy/2)⌉ do
(22)      for l ← 0 to N2 + N − 1 do
(23)          cost[l] ← CF · distance[(l + 1 + (l div (N + 1))) mod (N + 1)][l div N]
(24)      for i ← 0 to N − 1 do
(25)          for j ← 1 to i − 1 do
(26)              distance[i][j], distance[j][i] ← Transshipment(matrix, cost, demand[i][j])
(27)  return distance

```

In (1)–(8) we initialize the matrix which represents the graph of the transshipment problem. In (9)–(12) we set $stop[i]$ to the probability of termination in state s_i . In (13)–(16), we assign to $demand[i][j]$ the demand vector used to calculate the distance between states s_i and s_j . In (18)–(20) we initialize all distances to 0. As we will see, this corresponds to the pseudometric d_{ϕ_0} . In (21)–(26) we compute the pseudometrics $d_{\phi_1}, \dots, d_{\phi_{\lceil \log_{CF}(accuracy/2) \rceil}}$.

Transshipment(*matrix*, *cost*, *demand*)

Input:

- an $m \times n$ -matrix *matrix* with each column consisting of one 1, one -1 and 0's;
- an n -vector *cost*;
- an m -vector *demand*.

Output:

- minimal value of $cost \cdot vector^T$ for n -vector *vector* satisfying $matrix \cdot vector^T = demand^T$ and $0 \leq vector$.

For an algorithm which solves the transshipment problem in polynomial time we refer the reader to Chvátal's textbook [Chv83] for example. The remaining sections of this paper contain the correctness proof of the algorithm Distance.

3 A Metric Terminal Coalgebra Theorem

In this section, we introduce coalgebras and Rutten and Turi's metric terminal coalgebra theorem [RT92]. For more details about the theory of coalgebra we refer the reader to, for example, the tutorial [JR97] of Jacobs and Rutten.

DEFINITION 3 Let \mathcal{C} be a category. Let $F : \mathcal{C} \rightarrow \mathcal{C}$ be a functor. An F -coalgebra consists of an object C in \mathcal{C} together with an arrow $f : C \rightarrow F(C)$ in \mathcal{C} . The object C is called the *carrier*. An F -homomorphism from F -coalgebra $\langle C, f \rangle$ to F -coalgebra $\langle D, g \rangle$ is an arrow $\phi : C \rightarrow D$ in \mathcal{C} such that $F(\phi) \circ f = g \circ \phi$.

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

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*. ┘

We restrict our attention to the category \mathcal{CMet}_1 of 1-bounded complete metric spaces and nonexpansive functions. A metric space is 1-bounded if all its distances are bounded by 1. A function is nonexpansive if it does not increase any distances. We denote the collection of nonexpansive functions from the space X to the space Y by $X \rightarrow_1 Y$. This collection can be turned into a metric space by endowing the functions with the supremum metric.

For the rest of this paper, we fix c to be a constant between 0 and 1. This constant is the contraction factor of the functor we will consider and of the function that it induces. A function is c -contractive if it decreases all distances by at least a factor c . As we will see later, the constant c corresponds to the constant CF used in our algorithm.

DEFINITION 4 A functor $F : \mathcal{CMet}_1 \rightarrow \mathcal{CMet}_1$ is *locally c -contractive* if for all 1-bounded complete metric spaces X and Y , the function $F_{X,Y} : (X \rightarrow_1 Y) \rightarrow (F(X) \rightarrow_1 F(Y))$ defined by

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

is c -contractive. ┘

In the rest of this section, we restrict ourselves to locally c -contractive functors. For these functors, we have

THEOREM 5 *There exists a terminal F -coalgebra $\langle \text{fix}(F), \iota \rangle$.*

PROOF See [RT92, Theorem 4.8]. □

For the rest of this section, we fix $\langle X, \mu \rangle$ to be an F -coalgebra. To characterize the unique map from the F -coalgebra $\langle X, \mu \rangle$ to the terminal F -coalgebra $\langle \text{fix}(F), \iota \rangle$ we introduce the following function.

DEFINITION 6 The function $\Phi : (X \rightarrow_1 \text{fix}(F)) \rightarrow (X \rightarrow_1 \text{fix}(F))$ is defined by

$$\Phi(\phi) = \iota^{-1} \circ F(\phi) \circ \mu.$$

$$\begin{array}{ccc} X & \xrightarrow{\phi} & \text{fix}(F) \\ \mu \downarrow & & \uparrow \iota^{-1} \\ F(X) & \xrightarrow{F(\phi)} & F(\text{fix}(F)) \end{array}$$

Since the functor F is locally c -contractive, we have that the function Φ is c -contractive. ┘

PROPOSITION 7 *The function Φ is c -contractive.*

PROOF See proof of [RT92, Theorem 4.5]. □

Since Φ is a contractive function from a complete metric space to itself, we can conclude from Banach's theorem that it has a unique fixed point $\text{fix}(\Phi)$.

PROPOSITION 8 *The function $\text{fix}(\Phi)$ is the unique F -homomorphism from the F -coalgebra $\langle X, \mu \rangle$ to the terminal F -coalgebra $\langle \text{fix}(F), \iota \rangle$.*

PROOF See proof of [RT92, Theorem 4.5]. □

We conclude this section by showing that the unique map $\text{fix}(\Phi)$ can be approximated by the maps ϕ_n .

DEFINITION 9 Let $\phi_0 : X \rightarrow_1 \text{fix}(F)$ be some constant function. For $n > 0$, the function $\phi_n : X \rightarrow_1 \text{fix}(F)$ is defined by

$$\phi_n = \Phi(\phi_{n-1}).$$

PROPOSITION 10 *For all $n \geq 0$,*

$$d_{X \rightarrow_1 \text{fix}(F)}(\phi_n, \text{fix}(\Phi)) \leq c^n.$$

PROOF By induction on n . □

4 Metric Kernels

Our pseudometric on the states of a probabilistic transition system will be defined as the so-called metric kernel induced by the unique map from the probabilistic transition system, viewed as a coalgebra, to the terminal coalgebra. In this section, we introduce metric kernels. Furthermore, we show that the metric kernel induced by $\mathit{fix}(\Phi)$ can be approximated by the metric kernels induced by ϕ_n .

A function ϕ from the space X to the space $\mathit{fix}(F)$ defines a distance function d_ϕ on X . We call this distance function the metric kernel induced by ϕ . The distance between x_1 and x_2 in X is defined as the distance of their ϕ -images in the metric space $\mathit{fix}(F)$.

DEFINITION 11 Let $\phi \in X \rightarrow \mathit{fix}(F)$. The distance function $d_\phi : X \times X \rightarrow [0, 1]$ is defined by

$$d_\phi(x_1, x_2) = d_{\mathit{fix}(F)}(\phi(x_1), \phi(x_2)).$$

□

One can easily verify that the metric kernel d_ϕ is a pseudometric. Note that x_1 and x_2 have distance 0 only if they are mapped by ϕ to the same element in $\mathit{fix}(F)$.

The pseudometric $d_{\mathit{fix}(\Phi)}$ can be approximated by the pseudometrics d_{ϕ_n} as is shown in

PROPOSITION 12 For all $n \geq 0$ and $x_1, x_2 \in X$,

$$|d_{\phi_n}(x_1, x_2) - d_{\mathit{fix}(\Phi)}(x_1, x_2)| \leq 2 \cdot c^n.$$

PROOF

$$\begin{aligned} & |d_{\phi_n}(x_1, x_2) - d_{\mathit{fix}(\Phi)}(x_1, x_2)| \\ &= |d_{\mathit{fix}(F)}(\phi_n(x_1), \phi_n(x_2)) - d_{\mathit{fix}(F)}(\mathit{fix}(\Phi)(x_1), \mathit{fix}(\Phi)(x_2))| \\ &\leq d_{\mathit{fix}(F)}(\phi_n(x_1), \mathit{fix}(\Phi)(x_1)) + d_{\mathit{fix}(F)}(\phi_n(x_2), \mathit{fix}(\Phi)(x_2)) \quad [\text{triangle inequality}] \\ &\leq 2 \cdot d_{X \rightarrow \mathit{fix}(F)}(\phi_n, \mathit{fix}(\Phi)) \\ &\leq 2 \cdot c^n \quad [\text{Proposition 10}] \end{aligned}$$

□

To compute the $d_{\mathit{fix}(\Phi)}$ -distances up to accuracy α , it suffices to calculate the $d_{\phi_{\lceil \log_c(\alpha/2) \rceil}}$ -distances.

PROPOSITION 13 For all $0 < \alpha < 1$ and $x_1, x_2 \in X$,

$$|d_{\phi_{\lceil \log_c(\alpha/2) \rceil}}(x_1, x_2) - d_{\mathit{fix}(\Phi)}(x_1, x_2)| \leq \alpha.$$

PROOF

$$\begin{aligned} & |d_{\phi_{\lceil \log_c(\alpha/2) \rceil}}(x_1, x_2) - d_{\mathit{fix}(\Phi)}(x_1, x_2)| \\ &\leq 2 \cdot c^{\lceil \log_c(\alpha/2) \rceil} \quad [\text{Proposition 12}] \\ &\leq 2 \cdot c^{\log_c(\alpha/2)} \\ &= \alpha. \end{aligned}$$

□

The above proposition is reflected in line (21) of our algorithm.

5 The Hutchinson Functor

We introduce the Hutchinson metric on Borel probability measures of a metric space and its extension to a functor. From this Hutchinson functor we build the functor P using standard constructions. We show that every probabilistic transition system can be represented as a P -coalgebra. Furthermore, we prove that the Hutchinson functor is locally nonexpansive. This implies that the functor P is locally contractive. And,

hence, a terminal P -coalgebra exists according to Theorem 5. The metric kernel of the unique map to the terminal P -coalgebra defines a pseudometric on the carrier of a P -coalgebra and hence on the states of a probabilistic transition system. In Section 6 and 7 we will show that this is the pseudometric approximated by the algorithm presented in Section 2.

In [Hut81], Hutchinson introduced a metric on the set of Borel probability measures on a metric space. We restrict ourselves to spaces in which the distances are bounded by 1, since they serve our purpose. Let X be a 1-bounded metric space. We denote the set of Borel probability measures on X by $M(X)$. The Hutchinson distance on $M(X)$ is introduced in

DEFINITION 14 The distance function $d_{M(X)} : M(X) \times M(X) \rightarrow [0, 1]$ is defined by

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

For a proof that $d_{M(X)}$ is a 1-bounded metric, we refer the reader to, for example, Edgar's textbook [Edg98, Proposition 2.5.14].

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

Under quite mild conditions on the space, for example, completeness and separability, every measure is tight (see, for example, Parthasarathy's textbook [Par67, Theorem II.3.2]). In particular, all probabilistic transition systems can be represented using tight measures as we will see in Example 21. 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 16 X is complete if and only if $M_t(X)$ is complete.

PROOF See, for example, [Edg98, Theorem 2.5.25]. □

Using a standard construction, M_t can be extended to an endofunctor on the category \mathcal{CMet}_1 . Let X and Y be 1-bounded complete metric spaces. Let $f : X \rightarrow Y$ be a nonexpansive function. We have to define a nonexpansive function $M_t(f) : M_t(X) \rightarrow M_t(Y)$ satisfying the usual axioms.

DEFINITION 17 The function $M_t(f) : M_t(X) \rightarrow M_t(Y)$ is defined by

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

PROPOSITION 18 The Borel probability measure $M_t(f)(\mu)$ is tight.

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

PROPOSITION 19 The function $M_t(f)$ is nonexpansive.

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

$$\begin{aligned} & d_{M_t(Y)}(M_t(f)(\mu_1), M_t(f)(\mu_2)) \\ &= \sup \left\{ \int_Y g d(\mu_1 \circ f^{-1}) - \int_Y g d(\mu_2 \circ f^{-1}) \mid g \in Y \xrightarrow{1} [0, \infty) \right\} \\ &= \sup \left\{ \int_X (g \circ f) d\mu_1 - \int_X (g \circ f) d\mu_2 \mid g \in Y \xrightarrow{1} [0, \infty) \right\} \\ &\leq \sup \left\{ \int_X h d\mu_1 - \int_X h d\mu_2 \mid h \in X \xrightarrow{1} [0, \infty) \right\} \\ &= d_{M_t(X)}(\mu_1, \mu_2). \end{aligned}$$

□

For a proof that $M_t(f)$ satisfies the usual axioms we refer the reader to, for example, Giry's paper [Gir81, Section 1]. Next, we will show that the functor M_t is locally nonexpansive.

PROPOSITION 20 *The functor M_t is locally nonexpansive.*

PROOF We have to show that for all $f_1, f_2 \in X \xrightarrow{1} Y$,

$$d_{M_t(X) \rightarrow M_t(Y)}(M_t(f_1), M_t(f_2)) \leq d_{X \rightarrow Y}(f_1, f_2).$$

This immediate follows from the fact that we have for all $\mu \in M_t(X)$,

$$\begin{aligned} & d_{M_t(Y)}(M_t(f_1)(\mu), M_t(f_2)(\mu)) \\ &= \sup \left\{ \int_Y g d(\mu \circ f_1^{-1}) - \int_Y g d(\mu \circ f_2^{-1}) \mid g \in Y \xrightarrow{1} [0, \infty) \right\} \\ &= \sup \left\{ \int_X (g \circ f_1) d\mu - \int_X (g \circ f_2) d\mu \mid g \in Y \xrightarrow{1} [0, \infty) \right\} \\ &= \sup \left\{ \int_X (g \circ f_1 - g \circ f_2) d\mu \mid g \in Y \xrightarrow{1} [0, \infty) \right\} \\ &\leq d_{X \rightarrow Y}(f_1, f_2), \end{aligned}$$

since for all $x \in X$,

$$\begin{aligned} & (g \circ f_1 - g \circ f_2)(x) \\ &\leq d_Y(f_1(x), f_2(x)) \quad [g \text{ is nonexpansive}] \\ &\leq d_{X \rightarrow Y}(f_1, f_2). \end{aligned}$$

□

Now, we are ready to present the functor P . But first we introduce the functor T which models termination:

$$T = \mathbf{1} + c \cdot - : \mathcal{CMet}_1 \rightarrow \mathcal{CMet}_1,$$

where $\mathbf{1}$ is the terminal object¹ functor, $+$ is the coproduct² functor, and $c \cdot$ is the scaling³ functor. The functor P is defined by

$$P = M_t \circ T : \mathcal{CMet}_1 \rightarrow \mathcal{CMet}_1.$$

Every probabilistic transition system can be seen as a P -coalgebra as is demonstrated in

EXAMPLE 21 Let $\langle S, \pi \rangle$ be a probabilistic transition system. We endow the set of states S with the discrete metric. Consequently, every subset of the 1-bounded complete metric space $T(S)$ is a Borel set. For every state s , the Borel probability measure μ_s is the discrete Borel probability measure determined by

$$\begin{aligned} \mu_s(\mathbf{1}) &= \pi(s, \mathbf{0}) \\ \mu_s(\{s'\}) &= \pi(s, s') \end{aligned}$$

Obviously, the measure μ_s is tight. Because S is endowed with the discrete metric, the function μ mapping the state s to the measure μ_s is nonexpansive. Hence, every probabilistic transition system can be viewed as a P -coalgebra. ┘

¹The terminal object of \mathcal{CMet}_1 is the singleton space $\mathbf{1}$ whose single element we denote by $\mathbf{0}$.

²The coproduct object of the objects X and Y in \mathcal{CMet}_1 is the disjoint union of the sets underlying the spaces X and Y endowed with the metric

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

³The scaling by $c \cdot$ of an object in \mathcal{CMet}_1 leaves the set unchanged and multiplies all distances by c .

To exploit the metric terminal coalgebra theorem of Section 3 and the metric kernels of Section 4, we have to show that the functor P is locally contractive.

PROPOSITION 22 *The functor P is locally c -contractive.*

PROOF Since the functor $c \cdot$ is locally c -contractive and the functors $+$ and M_t are locally nonexpansive, the functor P is locally c -contractive. \square

According to Theorem 5, there exists a terminal P -coalgebra. Our pseudometric on a probabilistic transition system is defined as the metric kernel $d_{fix(\Phi)}$ where $fix(\Phi)$ is the unique map from the probabilistic transition system, viewed as a P -coalgebra, to the terminal P -coalgebra. In this pseudometric, states have distance 0 only if they are probabilistic bisimilar.

PROPOSITION 23 *For all states s_1, s_2 ,*

$$s_1 \text{ and } s_2 \text{ are probabilistic bisimilar if and only if } d_{fix(\Phi)}(s_1, s_2) = 0.$$

PROOF

s_1 and s_2 are probabilistic bisimilar

$$\text{iff } fix(\Phi)(s_1) = fix(\Phi)(s_2) \quad [\text{Corollary 5.12 of [VR99]}]$$

$$\text{iff } d_{fix(\Phi)}(s_1, s_2) = 0.$$

\square

6 Metric Kernels for P -coalgebras

Our pseudometric on a probabilistic transition system is defined as the metric kernel $d_{fix(\Phi)}$ where $fix(\Phi)$ is the unique map from the probabilistic transition system, viewed as a P -coalgebra, to the terminal P -coalgebra. As we have already seen in Section 4, $d_{fix(\Phi)}$ can be approximated by the metric kernels d_{ϕ_n} . In this section, we present a characterization of the pseudometrics d_{ϕ_n} for P -coalgebras in general and for probabilistic transition systems in particular. Furthermore, we will show that the d_{ϕ_n} -distances are smaller than or equal to the $d_{fix(\Phi)}$ -distances.

To prove our characterizations, we need the following

PROPOSITION 24 *Let $\phi \in X \rightarrow_1 fix(P)$. Then composition with $T(\phi)$ defines a surjection between $T(fix(P)) \rightarrow_1 [0, \infty)$ and $T\langle X, d_\phi \rangle \rightarrow_1 [0, \infty)$.*

PROOF ϕ may be regarded as an isometric embedding of the pseudometric space $\langle X, d_\phi \rangle$ in $fix(P)$. Thus $T(\phi)$ is an isometric embedding of $T\langle X, d_\phi \rangle$ into $T(fix(P))$. Now [Law73, Corollary on page 162] tells us that any nonexpansive map $f : T\langle X, d_\phi \rangle \rightarrow [0, \infty)$ has an extension $g : T(fix(P)) \rightarrow [0, \infty)$ in the sense that $g \circ T(\phi) = f$. \square

We can characterize the pseudometric d_{ϕ_n} on the carrier of a P -coalgebra $\langle X, \mu \rangle$ as follows.

THEOREM 25 *For all $x_1, x_2 \in X$,*

$$d_{\phi_0}(x_1, x_2) = 0.$$

For all $n > 0$ and $x_1, x_2 \in X$,

$$d_{\phi_n}(x_1, x_2) = \sup \left\{ \int_{T(X)} g d\mu_{x_1} - \int_{T(X)} g d\mu_{x_2} \mid g \in T\langle X, d_{\phi_{n-1}} \rangle \rightarrow_1 [0, \infty) \right\}.$$

PROOF Obviously,

$$\begin{aligned}
d_{\phi_0}(x_1, x_2) &= d_{f_{ix}(P)}(\phi_0(x_1), \phi_0(x_2)) \\
&= 0 \quad [\phi_0 \text{ is a constant function}]
\end{aligned}$$

For all $n > 0$,

$$\begin{aligned}
d_{\phi_n}(x_1, x_2) &= d_{f_{ix}(P)}(\phi_n(x_1), \phi_n(x_2)) \\
&= d_{f_{ix}(P)}(\Phi(\phi_{n-1})(x_1), \Phi(\phi_{n-1})(x_2)) \\
&= d_{f_{ix}(P)}((\iota^{-1} \circ P(\phi_{n-1}) \circ \mu)(x_1), (\iota^{-1} \circ P(\phi_{n-1}) \circ \mu)(x_2)) \\
&= d_{P(f_{ix}(P))}((P(\phi_{n-1}) \circ \mu)(x_1), (P(\phi_{n-1}) \circ \mu)(x_2)) \quad [\iota \text{ is isometric}] \\
&= \sup \left\{ \int_{T(f_{ix}(P))} f d((P(\phi_{n-1}) \circ \mu)(x_1)) - \int_{T(f_{ix}(P))} f d((P(\phi_{n-1}) \circ \mu)(x_2)) \mid f \in T(f_{ix}(P)) \xrightarrow{1} [0, \infty) \right\} \\
&= \sup \left\{ \int_{T(X)} (f \circ T(\phi_{n-1})) d\mu_{x_1} - \int_{T(X)} (f \circ T(\phi_{n-1})) d\mu_{x_2} \mid f \in T(f_{ix}(P)) \xrightarrow{1} [0, \infty) \right\} \\
&= \sup \left\{ \int_{T(X)} g d\mu_{x_1} - \int_{T(X)} g d\mu_{x_2} \mid g \in T\langle X, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\} \quad [\text{Proposition 24}]
\end{aligned}$$

□

Next, we refine the above characterization for the case that the P -coalgebra represents a probabilistic transition system $\langle S, \pi \rangle$.

PROPOSITION 26 For all $n > 0$ and $s_i, s_j \in S$,

$$d_{\phi_n}(s_i, s_j) = \sup \left\{ \sum_{s' \in T(S)} (\pi(s_i, s') - \pi(s_j, s')) \cdot g(s') \mid g \in T\langle S, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\}.$$

PROOF

$$\begin{aligned}
d_{\phi_n}(s_i, s_j) &= \sup \left\{ \int_{T(S)} g d\mu_{s_i} - \int_{T(S)} g d\mu_{s_j} \mid g \in T\langle S, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\} \quad [\text{Theorem 25}] \\
&= \sup \left\{ \sum_{s' \in T(S)} g(s') \cdot \mu_{s_i}(\{s'\}) - \sum_{s' \in T(S)} g(s') \cdot \mu_{s_j}(\{s'\}) \mid g \in T\langle S, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\} \\
&= \sup \left\{ \sum_{s' \in T(S)} g(s') \cdot \pi(s_i, s') - \sum_{s' \in T(S)} g(s') \cdot \pi(s_j, s') \mid g \in T\langle S, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\} \\
&= \sup \left\{ \sum_{s' \in T(S)} (\pi(s_i, s') - \pi(s_j, s')) \cdot g(s') \mid g \in T\langle S, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\}.
\end{aligned}$$

□

We conclude this section with a proof that the d_{ϕ_n} -distances are smaller than or equal to the $d_{f_{ix}(\Phi)}$ -distances. This result is proved for P -coalgebras in general.

PROPOSITION 27 For all $n \geq 0$,

$$d_{\phi_n} \leq d_{\phi_{n+1}}.$$

PROOF By induction on n . The case $n = 0$ is trivial. Let $n > 0$. For all $x_1, x_2 \in X$,

$$\begin{aligned} & d_{\phi_n}(x_1, x_2) \\ &= \sup \left\{ \int_{T(X)} g d\mu_{x_1} - \int_{T(X)} g d\mu_{x_2} \mid g \in T \langle X, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \right\} \quad [\text{Theorem 25}] \\ &\leq \sup \left\{ \int_{T(X)} g d\mu_{x_1} - \int_{T(X)} g d\mu_{x_2} \mid g \in T \langle X, d_{\phi_n} \rangle \xrightarrow{1} [0, \infty) \right\} \\ &\quad [\text{by induction } d_{\phi_{n-1}} \leq d_{\phi_n} \text{ and hence } T \langle X, d_{\phi_{n-1}} \rangle \xrightarrow{1} [0, \infty) \subseteq T \langle X, d_{\phi_n} \rangle \xrightarrow{1} [0, \infty)] \\ &= d_{\phi_{n+1}}(x_1, x_2) \quad [\text{Theorem 25}] \end{aligned}$$

□

PROPOSITION 28 For all $n \geq 0$,

$$d_{\phi_n} \leq d_{f_{ix}(\phi)}.$$

PROOF Immediate consequence of Proposition 27. □

In our algorithm, we compute $d_{\phi_{\lceil \log_c(\alpha/2) \rceil}}$. From the above result we can conclude that the computed $d_{\phi_{\lceil \log_c(\alpha/2) \rceil}}$ -distances are smaller than or equal to the actual $d_{f_{ix}(\phi)}$ -distances.

7 Linear Programming

In this section, we show that the computation of the d_{ϕ_n} -distances can be reduced to the transshipment problem. We already discussed this linear programming problem in the introduction.

Let $\langle \{s_0, \dots, s_{N-1}\}, \pi \rangle$ be a probabilistic transition system. According to Proposition 26, to compute $d_{\phi_n}(s_i, s_j)$ we need to

$$\text{maximize } \sum_{0 \leq k < N} (\pi(s_i, s_k) - \pi(s_j, s_k)) \cdot y_k + (\pi(s_i, \mathbf{0}) - \pi(s_j, \mathbf{0})) \cdot y_N \quad (1)$$

such that for all $0 \leq k < N$ and $0 \leq l < N$ with $k \neq l$,

$$y_k - y_l \leq c \cdot d_{\phi_{n-1}}(s_k, s_l) \quad (2)$$

and for all $0 \leq k < N$,

$$y_k - y_N \leq 1 \text{ and } y_N - y_k \leq 1 \quad (3)$$

and for all $0 \leq k \leq N$,

$$y_k \geq 0. \quad (4)$$

Note that y_k , for $0 \leq k < N$, and y_N play the role of $g(s_k)$ and $g(\mathbf{0})$, respectively. Constraints (2) and (3) reflect that the function g is nonexpansive. The fact that g maps states and $\mathbf{0}$ to nonnegative real numbers is captured by constraint (4). The above is a linear programming problem. If all y_k 's are 0 then the constraints (2), (3) and (4) are satisfied. Hence, the linear programming problem has a feasible origin. As we have already seen, (1) is bounded by 1. According to the fundamental theorem of linear programming (see [Chv83, Theorem 3.4] for example), the linear programming problem has an optimal solution.

The dual of this linear programming problem is

$$\text{minimize } \sum_{0 \leq k < N \wedge 0 \leq l < N \wedge k \neq l} c \cdot d_{\phi_{n-1}}(s_k, s_l) \cdot z_{k,l} + \sum_{0 \leq k < N} z_{k,N} + z_{N,k} \quad (5)$$

such that for all $0 \leq k < N$,

$$\sum_{0 \leq l \leq N \wedge k \neq l} z_{k,l} - z_{l,k} \geq \pi(s_i, s_k) - \pi(s_j, s_k) \quad (6)$$

and

$$\sum_{0 \leq l < N} z_{N,l} - z_{l,N} \geq \pi(s_i, \mathbf{0}) - \pi(s_j, \mathbf{0}) \quad (7)$$

and for all $0 \leq k \leq N$ and $0 \leq l \leq N$ with $k \neq l$,

$$z_{k,l} \geq 0. \quad (8)$$

Since the sum of the left hand sides of the constraints (6) and (7)

$$\sum_{0 \leq k \leq N} \sum_{0 \leq l \leq N \wedge k \neq l} z_{k,l} - z_{l,k} = 0$$

and the sum of the right hand sides of the constraints (6) and (7)

$$\sum_{0 \leq k < N} (\pi(s_i, s_k) - \pi(s_j, s_k)) + (\pi(s_i, \mathbf{0}) - \pi(s_j, \mathbf{0})) = 0,$$

we may simplify the constraints (6) and (7) to

$$\sum_{0 \leq l \leq N \wedge k \neq l} z_{k,l} - z_{l,k} = \pi(s_i, s_k) - \pi(s_j, s_k)$$

for all $0 \leq k < N$, and

$$\sum_{0 \leq l < N} z_{N,l} - z_{l,N} = \pi(s_i, \mathbf{0}) - \pi(s_j, \mathbf{0}).$$

According to the duality theorem of linear programming (see [Chv83, Theorem 5.1] for example), this dual problem also has an optimal solution with the same optimal value for (5) as the primal problem has for (1). The above dual problem can be formulated in terms of matrices and vectors as follows.

$$\text{minimize } \mathit{cost} \cdot \mathit{vector}^T$$

such that

$$\mathit{matrix} \cdot \mathit{vector}^T = \mathit{demand}^T$$

and

$$\mathit{vector} \geq 0$$

where the $(N^2 + N)$ -vector cost is defined by

$$\mathit{cost}_l = \begin{cases} 1 & \text{if } l \bmod N = 0 \text{ or } N^2 \leq l < N^2 + N \\ c \cdot d_{\phi_{n-1}}(s_{(l+1+(l \operatorname{div}(N+1))) \bmod (N+1)}, s_{l \operatorname{div} N}) & \text{otherwise} \end{cases}$$

and the $(N + 1)$ -vector demand is defined by

$$\mathit{demand}_k = \begin{cases} \pi(s_i, s_k) - \pi(s_j, s_k) & \text{if } 0 \leq k < N \\ \pi(s_i, \mathbf{0}) - \pi(s_j, \mathbf{0}) & \text{if } k = N \end{cases}$$

and the $(N + 1) \times (N^2 + N)$ -matrix matrix is defined by

$$\mathit{matrix}_{k,l} = \begin{cases} 1 & \text{if } k = (l + 1 + (l \operatorname{div}(N + 1))) \bmod (N + 1) \\ -1 & \text{if } k = l \operatorname{div} N \\ 0 & \text{otherwise} \end{cases}$$

For $N = 3$, the matrix looks as follows.


$$\begin{pmatrix} 1 & 1 & 1 & -1 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 1 & 1 & 1 & 0 & -1 & 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 & -1 & 0 & 1 & 1 & 1 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 & -1 & 1 & 1 & 1 \end{pmatrix}$$

Note that each column consists of one 1 and one -1 and the rest 0's. Consequently, the above is an instance of the transshipment problem.

8 Related and Future Work

At CONCUR'99, Desharnais et al. presented a pseudometric for probabilistic transition systems. Their pseudometric is defined by means of a real-valued modal logic where, in particular, the modal connective is interpreted by integration. Also in their setting, states have distance 0 if and only if they are probabilistic bisimilar.

We believe that our distances are more intuitive than theirs. For example, consider the following system.



The probability of termination when started in state s_0 , s_1 , s_2 and s_3 is 1, $\frac{6}{10}$, 1 and 0, respectively. The probability that the system makes, for example, at most three transitions when started in state s_0 , s_1 , s_2 and s_3 is $\frac{113}{125}$, $\frac{6}{10}$, 1 and 0, respectively. Based on these kind of observations, one may infer that (the system started in) state s_3 behaves more like state s_1 than state s_0 . This is reflected by our pseudometric, since the states s_3 and s_1 are $\frac{3}{10}$ apart (if $CF = \frac{1}{2}$) whereas the states s_3 and s_0 are at distance $\frac{3}{8}$. However, in the pseudometric introduced by Desharnais et al. both s_3 and s_1 and also s_3 and s_0 are $\frac{3}{20}$ apart. However, by adding negation⁴ to their real-valued model logic, the two pseudometrics differ by a factor $\frac{1}{2}$ (cf. [BW01]).

Desharnais et al. present an algorithm to calculate their distances up to a prescribed degree of accuracy. Their algorithm involves the generation of a representative set of formulas of their logic. They only consider formulas with a restricted number of nested occurrences of the modal connective. This corresponds to our approximation of $d_{fix}(\phi)$ by d_{ϕ_n} . Both restrict the depth at which observations are considered. Their algorithm calculates the distances in exponential time, whereas our algorithm computes them in polynomial time. Furthermore, it is not clear to us whether their algorithm can be adapted (in a straight forward manner) for the logic with negation (so that it can compute the more intuitive distances).

We have implemented our algorithm (see <http://www.cs.yorku.ca/~franck>). We are confident that we can improve our algorithm a little by exploiting the fact that d_{ϕ_n} is a pseudometric and hence satisfies the triangle inequality.

Many process combinators, like parallel composition, prefixing and probabilistic choice, can be shown to be nonexpansive with respect to our pseudometric. This quantitative analogue of congruence allows for compositional verification (see also [GJS90, DGJP99]).

In this paper, we considered probabilistic transition systems without labels to simplify the presentation. However, all our results can easily be generalized to a setting with labels. The coalgebras of the functor

$$L \rightarrow P : \mathcal{C}Met_1 \rightarrow \mathcal{C}Met_1,$$

where L is the finite set of labels, represent labelled probabilistic transition systems. To compute the d_{ϕ_n} -distance between states of a labelled system, for each label we consider only the transitions with that label and compute the d_{ϕ_n} -distance between the states, and take the maximum of all the computed distances.

⁴In a draft version, but not in the final version, of [DGJP00] negation was considered.

References

- [BV96] J.W. de Bakker and E.P. de Vink. *Control Flow Semantics*. Foundations of Computing Series. The MIT Press, Cambridge, 1996.
- [BW01] F. van Breugel and J. Worrell. Towards Quantitative Verification of Probabilistic Systems. To appear in *Proceedings of 28th International Colloquium on Automata, Languages and Programming*. Available at <http://www.cs.yorku.ca/~franck>, January 2001.
- [Chv83] V. Chvátal. *Linear Programming*. W.H. Freeman and Company, New York/San Francisco, 1983.
- [CN95] A. Cheng and M. Nielsen. Observing Behaviour Categorically. In P.S. Thiagarajan, editor, *Proceedings of the 15th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *Lecture Notes in Computer Science*, pages 263–278, Bangalore, December 1995. Springer-Verlag.
- [Des99] J. Desharnais. *Labelled Markov Processes*. PhD thesis, McGill University, Montreal, November 1999.
- [DGJP99] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for Labeled Markov Systems. In J.C.M. Baeten and S. Mauw, editors, *Proceedings of the 10th International Conference on Concurrency Theory*, volume 1664 of *Lecture Notes in Computer Science*, pages 258–273, Eindhoven, August 1999. Springer-Verlag.
- [DGJP00] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating Labeled Markov Processes. In *Proceedings of the 15th Annual IEEE Symposium on Logic in Computer Science*, pages 95–106, Santa Barbara, June 2000. IEEE.
- [Edg98] G.A. Edgar. *Integral, Probability, and Fractal Measures*. Springer-Verlag, 1998.
- [Gir81] M. Giry. A Categorical Approach to Probability Theory. In B. Banaschewski, editor, *Proceedings of the International Conference on Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85, Ottawa, August 1981. Springer-Verlag.
- [GJS90] A. Giacalone, C.-C. Jou, and S.A. Smolka. Algebraic Reasoning for Probabilistic Concurrent Systems. In *Proceedings of the IFIP WG 2.2/2.3 Working Conference on Programming Concepts and Methods*, pages 443–458, Sea of Galilee, April 1990. North-Holland.
- [Hut81] J.E. Hutchinson. Fractals and Self Similarity. *Indiana University Mathematics Journal*, 30(5):713–747, 1981.
- [JR97] B. Jacobs and J. Rutten. A Tutorial on (Co)Algebras and (Co)Induction. *Bulletin of the EATCS*, 62, June 1997.
- [KN96] M. Kwiatkowska and G.J. Norman. Probabilistic Metric Semantics for a Simple Language with Recursion. In W. Penczek and A. Szalas, editors, *Proceedings of the 21st International Symposium on Mathematical Foundations of Computer Science*, volume 1113 of *Lecture Notes in Computer Science*, pages 419–430, Cracow, September 1996. Springer-Verlag.
- [KN98] M. Kwiatkowska and G.J. Norman. A Testing Equivalence for Reactive Probabilistic Processes. In I. Castellani and C. Palamidessi, editors, *Proceedings of the 5th International Workshop on Expressiveness in Concurrency*, volume 16 of *Electronic Notes in Theoretical Computer Science*, Nice, September 1998. Elsevier.
- [Law73] F.W. Lawvere. Metric Spaces, Generalized Logic, and Closed Categories. *Rendiconti del Seminario Matematico e Fisico di Milano*, 43:135–166, 1973.
- [LS91] K.G. Larsen and A. Skou. Bisimulation through Probabilistic Testing. *Information and Computation*, 94(1):1–28, September 1991.
- [Nor97] G.J. Norman. *Metric Semantics for Reactive Probabilistic Systems*. PhD thesis, University of Birmingham, November 1997.
- [Par67] K.R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, 1967.
- [RT92] J.J.M.M. Rutten and D. Turi. On the Foundations of Final Semantics: non-standard sets, metric spaces, partial orders. In J.W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Proceedings of the REX Workshop on Semantics: Foundations and Applications*, volume 666 of *Lecture Notes in Computer Science*, pages 477–530, Beekbergen, June 1992. Springer Verlag.
- [Rut00] J.J.J.M. Rutten. Universal Coalgebra: a Theory of Systems. *Theoretical Computer Science*, 249(1):3–80, October 2000.
- [VR99] E.P. de Vink and J.J.M.M. Rutten. Bisimulation for Probabilistic Transition Systems: a Coalgebraic Approach. *Theoretical Computer Science*, 221(1/2):271–293, June 1999.