Characterising Probabilistic Processes Logically
(Extended Abstract)

Yuxin Deng[1,2]⋆ and Rob van Glabbeek[3,4]

[1] Dept. Comp. Sci. & Eng. and MOE-Microsoft Key Lab for Intell. Comp. & Syst., Shanghai Jiao Tong University, China
[2] State Key Lab of Comp. Sci., Inst. of Software, Chinese Academy of Sciences
[3] NICTA, Australia
[4] University of New South Wales, Australia

Abstract. In this paper we work on (bi)simulation semantics of processes that exhibit both nondeterministic and probabilistic behaviour. We propose a probabilistic extension of the modal mu-calculus and show how to derive characteristic formulae for various simulation-like preorders over finite-state processes without divergence. In addition, we show that even without the fixpoint operators this probabilistic mu-calculus can be used to characterise these behavioural relations in the sense that two states are equivalent if and only if they satisfy the same set of formulae.

1 Introduction

In concurrency theory, behavioural relations such as equivalences and refinement preorders form a basis for establishing system correctness. Usually both specifications and implementations are expressed as processes within the same framework, in which a specification describes some high-level behaviour and an implementation gives the technical details for achieving the behaviour. Then one chooses an equivalence or preorder to verify that the implementation realises the behaviour required by the specification.

A great many behavioural relations are defined on top of labelled transition systems, which offer an operational model of systems. For finitary (i.e. finite-state and finitely branching) systems, these behavioural relations can be computed in a mechanical way, and thus may be incorporated into automatic verification tools. In recent years, probabilistic constructs have been proven useful for giving quantitative specifications of system behaviour. The first papers on probabilistic concurrency theory [12,2,20] proceed by *replacing* nondeterministic with probabilistic constructs. The reconciliation of nondeterministic and probabilistic constructs starts with [13] and has received a lot of attention in the literature [35,31,21,30,16,22,1,18,25,3,34,23,9,7,4]. We shall also work in a framework that features the co-existence of probability and nondeterminism.

Among the behavioural relations that have proven useful in probabilistic concurrency theory are various types of *simulation* and *bisimulation* relations. Axiomatisations for bisimulations have been investigated in [1,10]. Logical characterisations of bisimulations and simulations have been studied in [31,27]. For

⋆ Deng was supported by the National Natural Science Foundation of China (60703033).

example, in [31] the probabilistic computation tree logic (PCTL) [14] is used and it turns out that two states are bisimilar if and only if they satisfy the same set of PCTL formulae.

In the nonprobabilistic setting, there is a line of research on characteristic formulae. The goal is to seek a particular formula φ_s for a given state s such that a necessary and sufficient condition for any state t being bisimilar to s is to satisfy φ_s [32]. This is a very strong property in the sense that to check if t is bisimilar to s it suffices to consider the single formula φ_s and see if it can be satisfied by t. It offers a convenient method for equivalence or preorder checking.

In this paper we partially extend the results of [32] to a probabilistic setting that admits both probabilistic and nondeterministic choice; to make the main ideas neat we do not consider divergence. We present a probabilistic extension of the modal mu-calculus [19] (pMu), where a formula is interpreted as the set of probability distributions satisfying it. This is in contrast to the probabilistic semantics of the mu-calculus as studied in [16,22,23] where formulae denote lower bounds of probabilistic evidence of properties, and the semantics of the generalised probabilistic logic of [3] where a mu-calculus formula is interpreted as a set of deterministic trees that satisfy it.

We shall provide characteristic formulae for strong and weak probabilistic (bi)simulation as introduced in [31,30], as well as forward simulation [30] and failure simulation [7]. The results are obtained in two phases, which we illustrate by taking strong probabilistic bisimilarity $\sim$ as an example. Given a finite-state probabilistic labelled transition system with state space $\{s_1, ..., s_n\}$, we first construct an equation system E of modal formulae in pMu.

$$\begin{aligned} E : X_{s_1} &= \varphi_{s_1} \\ &\vdots \\ X_{s_n} &= \varphi_{s_n} \end{aligned}$$

A solution of the equation system is a function ρ that assigns to each variable X_{s_i} a set of distributions $\rho(X_{s_i})$. The greatest solution of the equation system, denoted by ν_E, has the property that $s_i \sim s_j$ if and only if the point distribution $\overline{s_j}$ is an element of $\nu_E(X_{s_i})$. In the second phase, we apply three transformation rules upon E in order to obtain a pMu formula $\varphi^{\sim}_{s_i}$ whose meaning $[\![\varphi^{\sim}_{s_i}]\!]$ is exactly captured by $\nu_E(X_{s_i})$. As a consequence, we derive a characteristic formula for s_i such that $s_i \sim s_j$ if and only if $\overline{s_j} \in [\![\varphi^{\sim}_{s_i}]\!]$.

Without the fixpoint operators pMu gives rise to a probabilistic extension of the Hennessy-Milner logic [15]. In analogy to the nonprobabilistic setting, it characterises (bi)simulations in the sense that $s \sim t$ if and only if the two states s, t satisfy the same set of formulae.

The paper is organised as follows. In Section 2 we recall the definitions of several (bi)simulations defined over probabilistic labelled transition systems. In Section 3 we introduce the syntax and semantics of pMu. In Section 4 we build characteristic equation systems and derive from them characteristic formulae for all our (bi)simulations. In Section 5 we consider the fixpoint-free fragment of pMu which characterises a state by the class of formulae it satisfies.

In this extended abstract many proofs are omitted; they can be found in [6].

2 Probabilistic (bi)simulations

In this section we recall several probabilistic extensions of simulation and bisimulation [24] that appeared in the literature.

We begin with some notation concerning probability distributions. A *(discrete) probability distribution* over a set S is a function $\Delta : S \to [0,1]$ with $\sum_{s\in S}\Delta(s) = 1$; the *support* of Δ is given by $\lceil\Delta\rceil = \{ s \in S \mid \Delta(s) > 0 \}$. We write $\mathcal{D}(S)$, ranged over by Δ, Θ, for the set of all distributions over S. We also write $\overline{s}$ to denote the point distribution assigning probability 1 to s and 0 to all others, so that $\lceil\overline{s}\rceil = \{s\}$. If $p_i \geq 0$ and Δ_i is a distribution for each i in some index set I, and $\sum_{i\in I} p_i = 1$, then the probability distribution $\sum_{i\in I} p_i \cdot \Delta_i \in \mathcal{D}(S)$ is given by $(\sum_{i\in I} p_i \cdot \Delta_i)(s) = \sum_{i\in I} p_i \cdot \Delta_i(s)$; we will sometimes write it as $p_1 \cdot \Delta_1 + \ldots + p_n \cdot \Delta_n$ when $I = \{1, \ldots, n\}$.

Definition 1. A *finite state probabilistic labelled transition system* (pLTS) is a triple $\langle S, \mathsf{Act}_\tau, \to\rangle$, where

1. S is a finite set of states
2. Act_τ is a set of external actions Act augmented with an internal action $\tau \notin \mathsf{Act}$
3. $\to\ \subseteq\ S \times \mathsf{Act}_\tau \times \mathcal{D}(S)$.

We usually write $s \xrightarrow{a} \Delta$ for $(s, a, \Delta) \in \to$, $s \xrightarrow{a}$ for $\exists\Delta : s \xrightarrow{a} \Delta$, and $s \stackrel{a}{\not\to}$ for the negation of $s \xrightarrow{a}$. We write $s \stackrel{A}{\not\to}$ with $A \subseteq \mathsf{Act}$ when $\forall a \in A \cup \{\tau\} : s \stackrel{a}{\not\to}$, and $\Delta \stackrel{A}{\not\to}$ when $\forall s \in \lceil\Delta\rceil : s \stackrel{A}{\not\to}$. A pLTS is *finitely branching* if, for each state s, the set $\{(a, \Delta) \mid s \xrightarrow{a} \Delta\}$ is finite. A pLTS is *finitary* if it is finite-state and finitely branching.

To define probabilistic (bi)simulations, it is often necessary to lift a relation over states to one over distributions.

Definition 2. Given two sets S and T and a relation $\mathcal{R} \subseteq S \times T$. We lift $\mathcal{R}$ to a relation $\mathcal{R}^\dagger \subseteq \mathcal{D}(S) \times \mathcal{D}(T)$ by letting $\Delta\ \mathcal{R}^\dagger\ \Theta$ whenever

1. $\Delta = \sum_{i\in I} p_i \cdot \overline{s_i}$, where I is a countable index set and $\sum_{i\in I} p_i = 1$
2. for each $i \in I$ there is a state t_i such that $s_i\ \mathcal{R}\ t_i$
3. $\Theta = \sum_{i\in I} p_i \cdot \overline{t_i}$.

Note that in the decomposition of Δ, the states s_i are not necessarily distinct: that is, the decomposition is not in general unique, and similarly for the decomposition of Θ. For example, if $\mathcal{R} = \{(s_1, t_1), (s_1, t_2), (s_2, t_3), (s_3, t_3)\}$, $\Delta = \frac{1}{2}\overline{s_1} + \frac{1}{4}\overline{s_2} + \frac{1}{4}\overline{s_3}$, and $\Theta = \frac{1}{3}\overline{t_1} + \frac{1}{6}\overline{t_2} + \frac{1}{2}\overline{t_3}$, then $\Delta\ \mathcal{R}^\dagger\ \Theta$ holds because of the decompositions $\Delta = \frac{1}{3}\overline{s_1} + \frac{1}{6}\overline{s_1} + \frac{1}{4}\overline{s_2} + \frac{1}{4}\overline{s_3}$ and $\Theta = \frac{1}{3}\overline{t_1} + \frac{1}{6}\overline{t_2} + \frac{1}{4}\overline{t_3} + \frac{1}{4}\overline{t_3}$.

From the above definition, the next two properties follow [5]. In fact, they are sometimes used as alternative methods of lifting relations (see e.g. [31,20]).

Proposition 1. 1. Let Δ and Θ be distributions over S and T, respectively. Then $\Delta\ \mathcal{R}^\dagger\ \Theta$ iff there exists a weight function $w : S \times T \to [0,1]$ such that
   (a) $\forall s \in S : \sum_{t\in T} w(s,t) = \Delta(s)$
   (b) $\forall t \in T : \sum_{s\in S} w(s,t) = \Theta(t)$
   (c) $\forall (s,t) \in S \times T : w(s,t) > 0 \Rightarrow s\ \mathcal{R}\ t$.

2. Let Δ, Θ be distributions over S and $\mathcal{R}$ be an equivalence relation. Then $\Delta \; \mathcal{R}^\dagger \; \Theta$ iff $\Delta(C) = \Theta(C)$ for all equivalence classes $C \in S/\mathcal{R}$, where $\Delta(C)$ stands for the accumulated probability $\sum_{s \in C} \Delta(s)$. □

In a similar way, following [9], we can lift a relation $\mathcal{R} \subseteq S \times \mathcal{D}(T)$ to a relation $\mathcal{R}^\dagger \subseteq \mathcal{D}(S) \times \mathcal{D}(T)$, by letting $\Delta \; \mathcal{R}^\dagger \; \Theta$ whenever

1. $\Delta = \sum_{i \in I} p_i \cdot \overline{s_i}$, where I is a countable index set and $\sum_{i \in I} p_i = 1$
2. for each $i \in I$ there is a distribution Θ_i such that $s_i \; \mathcal{R} \; \Theta_i$
3. $\Theta = \sum_{i \in I} p_i \cdot \Theta_i$.

The above lifting constructions satisfy the following two useful properties.

Proposition 2. Suppose $\mathcal{R} \subseteq S \times S$ or $S \times \mathcal{D}(S)$ and $\sum_{i \in I} p_i = 1$. Then

1. $\Delta_i \; \mathcal{R}^\dagger \; \Theta_i$ for all $i \in I$ implies $(\sum_{i \in I} p_i \cdot \Delta_i) \; \mathcal{R}^\dagger \; (\sum_{i \in I} p_i \cdot \Theta_i)$.
2. If $(\sum_{i \in I} p_i \cdot \Delta_i) \; \mathcal{R}^\dagger \; \Theta$ then $\Theta = \sum_{i \in I} p_i \cdot \Theta_i$ for some set of distributions Θ_i such that $\Delta_i \; \mathcal{R}^\dagger \; \Theta_i$ for all $i \in I$. □

We write $s \xrightarrow{\hat{\tau}} \Delta$ if either $s \xrightarrow{\tau} \Delta$ or $\Delta = \overline{s}$, and $s \xrightarrow{\hat{a}} \Delta$ iff $s \xrightarrow{a} \Delta$ for $a \in \mathsf{Act}$. For any $a \in \mathsf{Act}_\tau$, we know that $\xrightarrow{\hat{a}} \subseteq S \times \mathcal{D}(S)$, so we can lift it to be a transition relation between distributions. With a slight abuse of notation we simply write $\Delta \xrightarrow{\hat{a}} \Theta$ for $\Delta \; (\xrightarrow{\hat{a}})^\dagger \; \Theta$. Then we define weak transitions $\stackrel{\hat{a}}{\Longrightarrow}$ by letting $\stackrel{\hat{\tau}}{\Longrightarrow}$ be the reflexive and transitive closure of $\xrightarrow{\hat{\tau}}$ and writing $\Delta \stackrel{\hat{a}}{\Longrightarrow} \Theta$ for $a \in \mathsf{Act}$ whenever $\Delta \stackrel{\hat{\tau}}{\Longrightarrow}\xrightarrow{\hat{a}}\stackrel{\hat{\tau}}{\Longrightarrow} \Theta$.

Definition 3. A *divergence* is a sequence of states s_i and distributions Δ_i with $s_i \xrightarrow{\tau} \Delta_i$ and $s_{i+1} \in \lceil \Delta_i \rceil$ for $i \geq 0$.

The above definition of $\stackrel{\hat{a}}{\Longrightarrow}$ is sensible only in the absence of divergence. In general, one would need a more complicated notion of $\stackrel{\hat{a}}{\Longrightarrow}$, such as proposed in [8]. Therefore, from here on we restrict attention to divergence-free pLTSs.

Definition 4. A relation $\mathcal{R} \subseteq S \times S$ is a *strong probabilistic simulation* if $s \; \mathcal{R} \; t$ and $a \in \mathsf{Act}_\tau$ implies

– if $s \xrightarrow{a} \Delta$ then there exists some Θ such that $\overline{t} \xrightarrow{a} \Theta$ and $\Delta \; \mathcal{R}^\dagger \; \Theta$

If both $\mathcal{R}$ and $\mathcal{R}^{-1}$ are strong probabilistic simulations, then $\mathcal{R}$ is a *strong probabilistic bisimulation*. A state s is related to another state t via *strong probabilistic similarity* (resp. *bisimilarity*), denoted $s \prec t$ (resp. $s \sim t$), if there exists a strong probabilistic simulation (resp. bisimulation) $\mathcal{R}$ such that $s \; \mathcal{R} \; t$. *Weak probabilistic similarity* ($\precsim$) and *weak probabilistic bisimilarity* ($\approx$) are defined in the same manner just by using $\overline{t} \stackrel{\hat{a}}{\Longrightarrow} \Theta$ in place of $\overline{t} \xrightarrow{a} \Theta$.

All four (bi)simulations above stem from [31,30]. There they were proposed as improvements over the strong bisimulation of [13] and the strong simulation of [17], both of which can be defined as the strong *probabilistic* (bi)simulation above, but using $t \xrightarrow{a} \Theta$ in place of $\overline{t} \xrightarrow{a} \Theta$. Logical characterisations for strong (bi)simulations are similar to those contributed here for strong probabilistic (bi)simulations, but require a state-based interpretation of the modalities $\langle a \rangle$ and $[a]$; see [6] for details. Other definitions of simulation have also appeared in the literature. Here we consider two typical ones: forward simulation [30] and failure simulation [7].

Definition 5. A relation $\mathcal{R} \subseteq S \times \mathcal{D}(S)$ is a *failure simulation* if $s \mathrel{\mathcal{R}} \Theta$ implies

1. if $s \xrightarrow{a} \Delta$ with $a \in \mathsf{Act}_\tau$ then $\exists \Theta'$ such that $\Theta \stackrel{\hat{a}}{\Longrightarrow} \Theta'$ and $\Delta \mathrel{\mathcal{R}^\dagger} \Theta'$;
2. if $s \stackrel{A}{\not\longrightarrow}$ with $A \subseteq \mathsf{Act}$ then $\exists \Theta'$ such that $\Theta \stackrel{\hat{\tau}}{\Longrightarrow} \Theta'$ and $\Theta' \stackrel{A}{\not\longrightarrow}$.

We write $s \lhd_{FS} \Theta$ if there is some failure simulation $\mathcal{R}$ such that $s \mathrel{\mathcal{R}} \Theta$.

Similarly, we define a forward simulation and $s \lhd_{S} \Theta$ by dropping the second clause in Definition 5.

3 The Probabilistic Modal mu-Calculus

Let Var be a countable set of variables. We define a class $\mathcal{L}^{\text{raw}}$ of modal formulae by the following grammar:

$$\varphi := \bigwedge_{i\in I} \varphi_i \mid \bigvee_{i\in I} \varphi_i \mid \neg\varphi \mid \langle a\rangle\varphi \mid [a]\varphi \mid \bigoplus_{i\in I} \varphi_i \mid \bigoplus_{i\in I} p_i \cdot \varphi_i \mid {\downarrow}\varphi \mid X \mid \mu X.\varphi \mid \nu X.\varphi$$

where I is an index set, $a \in \mathsf{Act}_\tau$ and $\sum_{i\in I} p_i = 1$. The *probabilistic modal mu-calculus* (pMu) is given by the subclass $\mathcal{L}$, obtained by imposing the syntactic condition that in $\mu X.\varphi$ and $\nu X.\varphi$ the variable X may occur in φ only within the scope of an even number of negations. The above syntax is obtained by adding a variant of the probabilistic construct $\bigoplus_{i\in I} p_i \cdot \varphi_i$, introduced in [7] in the context of a less expressive logic without fixpoint operators, as well as the novel modalities $\bigoplus_{i\in I} \varphi_i$ and ${\downarrow}\varphi$, to the syntax of the non-probabilistic mu-calculus [19]. As usual, one has $\bigwedge_{i\in\emptyset} \varphi_i = \texttt{true}$ and $\bigvee_{i\in\emptyset} \varphi_i = \texttt{false}$.

The two fixpoint operators μX and νX bind the respective variable X. We apply the usual terminology of free and bound variables in a formula and write $fv(\varphi)$ for the set of free variables in φ. A formula φ is *closed* if $fv(\varphi) = \emptyset$.

For any set Ω, write $\mathcal{P}(\Omega)$ for the power set of Ω. We use *environments*, which bind free variables to sets of distributions, in order to give semantics to formulae. Let

$$\mathsf{Env} = \{\, \rho \mid \rho : \mathsf{Var} \to \mathcal{P}(\mathcal{D}(S)) \,\}$$

be the set of all environments and ranged over by ρ. For a set $V \subseteq \mathcal{D}(S)$ and a variable $X \in \mathsf{Var}$, we write $\rho[X \mapsto V]$ for the environment that maps X to V and Y to $\rho(Y)$ for all $Y \neq X$.

The semantics of a formula φ in an environment ρ is given as the set of distributions $[\![\varphi]\!]_\rho$ satisfying it. This leads to a semantic functional $[\![\ ]\!] : \mathcal{L} \to \mathsf{Env} \to \mathcal{P}(\mathcal{D}(S))$ defined inductively in Table 1. As the meaning of a closed formula φ does not depend on the environment, one writes $[\![\varphi]\!]$ for $[\![\varphi]\!]_\rho$ where ρ is an arbitrary environment. In that case one also writes $\Delta \models \varphi$ for $\Delta \in [\![\varphi]\!]$.

Following [19,29] we give a *strong* and a *weak* semantics of the probabilistic modal mu-calculus. Both are the same as those of the modal mu-calculus [19,29] except that distributions of states are taking the roles of states. The power set of $\mathcal{D}(S)$, $\mathcal{P}(\mathcal{D}(S))$, may be viewed as the complete lattice $(\mathcal{P}(\mathcal{D}(S)), \mathcal{D}(S), \emptyset, \subseteq, \cup, \cap)$. Intuitively, we identify a formula with the set of distributions that make it true. For example, `true` holds for all distributions and dually `false`

$$
\begin{array}{rcl}
[\![\bigwedge_{i\in I}\varphi_i]\!]_\rho = \bigcap_{i\in I}[\![\varphi_i]\!]_\rho & \text{so} & [\![\mathtt{true}]\!]_\rho = \mathcal{D}(S) \\
[\![\bigvee_{i\in I}\varphi_i]\!]_\rho = \bigcup_{i\in I}[\![\varphi_i]\!]_\rho & \text{so} & [\![\mathtt{false}]\!]_\rho = \emptyset \\
[\![\neg\varphi]\!]_\rho = \mathcal{D}(S)\setminus[\![\varphi]\!]_\rho & & \\
[\![\langle a\rangle\varphi]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \exists\Delta' : \Delta\xrightarrow{a}\Delta' \;\wedge\; \Delta'\in[\![\varphi]\!]_\rho\,\} & & \\
[\![[a]\varphi]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \forall\Delta' : \Delta\xrightarrow{a}\Delta' \;\Rightarrow\; \Delta'\in[\![\varphi]\!]_\rho\,\} & & \\
[\![\bigoplus_{i\in I}\varphi_i]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \Delta=\sum_{i\in I}p_i\cdot\Delta_i \text{ for some } p_i \text{ with } \sum_{I\in I}p_i=1 & & \\
\qquad\wedge\ \forall i\in I:\ \Delta_i\in[\![\varphi_i]\!]_\rho\,\} & & \\
[\![\bigoplus_{i\in I}p_i\cdot\varphi_i]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \Delta=\sum_{i\in I}p_i\cdot\Delta_i \;\wedge\; \forall i\in I:\ \Delta_i\in[\![\varphi_i]\!]_\rho\,\} & & \\
[\![{\downarrow}\varphi]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \forall s\in\lceil\Delta\rceil : \overline{s}\in[\![\varphi]\!]_\rho\,\} & & \\
[\![X]\!]_\rho = \rho(X) & & \\
[\![\mu X.\varphi]\!]_\rho = \bigcap\{\,V\subseteq\mathcal{D}(S) \mid [\![\varphi]\!]_{\rho[X\mapsto V]}\subseteq V\,\} & & \\
[\![\nu X.\varphi]\!]_\rho = \bigcup\{\,V\subseteq\mathcal{D}(S) \mid [\![\varphi]\!]_{\rho[X\mapsto V]}\supseteq V\,\} & & \\
\hline
[\![\langle a\rangle\varphi]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \exists\Delta' : \Delta\stackrel{\hat{a}}{\Longrightarrow}\Delta' \;\wedge\; \Delta'\in[\![\varphi]\!]_\rho\,\} & & \\
[\![[a]\varphi]\!]_\rho = \{\,\Delta\in\mathcal{D}(S) \mid \forall\Delta' : \Delta\stackrel{\hat{a}}{\Longrightarrow}\Delta' \;\Rightarrow\; \Delta'\in[\![\varphi]\!]_\rho\,\} & &
\end{array}
$$

Table 1. Strong and weak semantics of the probabilistic modal mu-calculus

holds for no distribution. Conjunction and disjunction are interpreted by intersection and union of sets, and negation by complement. The formula $\langle a\rangle\varphi$ holds for a distribution Δ if there is a distribution Δ' that can be reached after an a-transition and that satisfies φ. Dually, $[a]\varphi$ holds for Δ if all distributions reachable from Δ by an a-transition satisfy φ. The formulas $\bigoplus_{i\in I}\varphi_i$ and $\bigoplus_{i\in I}p_i\cdot\varphi_i$ hold for Δ if the distribution can be decomposed into a convex combination of some distributions Δ_i and each of them satisfies the corresponding sub-formula φ_i; the first of these modalities allows *any* convex combination, whereas the second one specifies a particular one. The formula ${\downarrow}\varphi$ holds for Δ if all states in its support satisfy φ. The characterisation of the *least fixpoint formula* $\mu X.\varphi$ and the *greatest fixpoint formula* $\nu X.\varphi$ follows from the well-known Knaster-Tarski fixpoint theorem [33].

The weak semantics reflects the unobservable nature of internal actions; it differs from the strong semantics only in the use of the relations $\stackrel{\hat{a}}{\Longrightarrow}$ instead of $\xrightarrow{a}$ in the interpretation of the modalities $\langle a\rangle$ and $[a]$.

Note that there is some redundancy in the syntax of pMu: each of the constructs $\bigwedge_{i\in I}$, $\langle a\rangle$ and μ can be expressed in terms of its dual $\bigvee_{i\in I}$, $[a]$ and ν with the aid of negation. However, negation may not be redundant, as the dual of $\bigoplus_{i\in I}p_i\cdot\varphi_i$ does not appear to be expressible without using negation; moreover this dual lacks the intuitive appeal for introducing it as a new primitive.

We shall consider (closed) *equation systems* of formulae of the form

$$
\begin{aligned}
E: X_1 &= \varphi_1 \\
&\ \vdots \\
X_n &= \varphi_n
\end{aligned}
$$

where $X_1,...,X_n$ are mutually distinct variables and $\varphi_1,...,\varphi_n$ are formulae having at most $X_1,...,X_n$ as free variables. Here E can be viewed as a function

- Rule 1: $E \to F$
- Rule 2: $E \to G$
- Rule 3: $E \to H$ if $X_n \notin fv(\varphi_1, ..., \varphi_n)$

E:	F:	G:	H:
$X_1 = \varphi_1$	$X_1 = \varphi_1$	$X_1 = \varphi_1[\varphi_n/X_n]$	$X_1 = \varphi_1$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$X_{n-1} = \varphi_{n-1}$	$X_{n-1} = \varphi_{n-1}$	$X_{n-1} = \varphi_{n-1}[\varphi_n/X_n]$	$X_{n-1} = \varphi_{n-1}$
$X_n = \varphi_n$	$X_n = \nu X_n.\varphi_n$	$X_n = \varphi_n$	

Table 2. Transformation rules

$E : \mathsf{Var} \to \mathcal{L}$ defined by $E(X_i) = \varphi_i$ for $i = 1, ..., n$ and $E(Y) = Y$ for other variables $Y \in \mathsf{Var}$.

An environment ρ is a *solution* of an equation system E if its assignment to X_i coincides with the interpretation of φ_i in the environment, that is,

$$\forall i : \rho(X_i) = [\![\varphi_i]\!]_\rho.$$

The existence of solutions for an equation system can be seen from the following arguments. The set Env, which includes all candidates for solutions, together with the partial order $\sqsubseteq$ defined by

$$\rho \sqsubseteq \rho' \text{ iff } \forall X \in \mathsf{Var} : \rho(X) \subseteq \rho'(X)$$

forms a complete lattice. The *equation functional* $\mathcal{F}_E : \mathsf{Env} \to \mathsf{Env}$ given in the notation of the λ-calculus by

$$\mathcal{F}_E := \lambda\rho.\lambda X.[\![E(X)]\!]_\rho$$

is monotonic, which can be shown by induction on the structure of $E(X)$. Thus, the Knaster-Tarski fixpoint theorem guarantees existence of solutions, and the greatest solution

$$\nu_E := \bigsqcup \{ \rho \mid \rho \sqsubseteq \mathcal{F}_E(\rho) \} \tag{1}$$

is the supremum of the set of all post-fixpoints of $\mathcal{F}_E$.

An expression $\nu_E(X)$, with X one of the variables used in E, denotes a set of distributions. Below we will use such expressions as if they were valid syntax in our probabilistic mu-calculus, with $[\![\nu_E(X)]\!]_\rho := \nu_E(X)$. This amounts to extending the greatest fixpoint operator ν to apply to finite sets of fixpoint equations, instead of single equations; the expression $\nu X.\varphi$ amounts to the special case $\nu_E(X)$ in which E consists of the single equation $X = \varphi$.

The use of expressions $\nu_E(X)$ is justified because they can be seen as syntactic sugar for authentic pMu expressions. As explained in [26], the three transformation rules in Table 2 can be used to obtain from an equation system E a pMu formula whose interpretation coincides with the interpretation of X_1 in the greatest solution of E.

Theorem 1. *Given a finite equation system E that uses the variable X, there is a pMu formula φ such that $\nu_E(X) = [\![\varphi]\!]$.* □

4 Characteristic equation systems

Following [32], the behaviour of a finite-state process can be characterised by an equation system of modal formulae. In the current section we show that this idea also applies in the probabilistic setting. For each behavioural relation $\mathcal{R}$ over a finite state space, ranging over the various simulation preorders and bisimulation equivalences reviewed in Section 2, we establish an equation system E of modal formulae in pMu.

$$\begin{aligned} E : X_{s_1} &= \varphi_{s_1} \\ &\vdots \\ X_{s_n} &= \varphi_{s_n} \end{aligned}$$

There is exactly one such equation for each state s_i, and the formulae φ_{s_i} do not contain fixpoint operators. This equation system is guaranteed to have a greatest solution ν_E which has the nice property that, for any states s, t in the state space in question, s is related to t via $\mathcal{R}$ if and only if the point distribution $\overline{t}$ belongs to the set of distributions assigned to the variable X_s by ν_E. Thus $\nu_E(X_s)$ is a *characteristic formula* for s w.r.t. $\mathcal{R}$ in the sense that $s\ \mathcal{R}\ t$ iff $\overline{t}$ satisfies $\nu_E(X_s)$.

Strong probabilistic bisimulation The key ingredient for the modal characterisation of strong probabilistic bisimulation is to construct an equation system that captures all the transitions of a pLTS. For each state s we build an equation $X_s = \varphi_s$, where X_s is a variable and φ_s is of the form $\varphi'_s \wedge \varphi''_s$ with φ'_s a formula describing the actions enabled by s and φ''_s a formula describing the consequences of performing these actions. Intuitively, if state s is related to state t in a bisimulation game, then φ'_s expresses the transitions that should be matched up by t and φ''_s expresses the capability of s to match up the transitions initiated by t. More specifically, the equation system is given by the following definition.

Definition 6. Given a pLTS, its *characteristic equation system* for strong probabilistic bisimulation consists of one equation $X_s = \varphi_s$ for each state $s \in S$, where

$$\varphi_s := (\bigwedge_{s \xrightarrow{a} \Delta} \langle a \rangle X_\Delta) \wedge (\bigwedge_{a \in \mathsf{Act}_\tau} [a] \bigoplus_{s \xrightarrow{a} \Delta} X_\Delta)^{1} \qquad (2)$$

with $X_\Delta := \bigoplus_{s \in \lceil \Delta \rceil} \Delta(s) \cdot {\downarrow} X_s$.

The equation system thus constructed, interpreted according to the strong semantics of pMu, has the required property, as stated by the theorem below.

[1] The subformula $\bigoplus_{s \xrightarrow{a} \Delta} X_\Delta$ is equivalent to $\bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta$, and this is the form that we use to prove Theorem 2. If the given pLTS has nondeterministic choices among different transitions labelled with the same action, this disjunction is infinite. For example, if $s \xrightarrow{a} \overline{s_i}$ for $i = 1, 2$, then $\overline{s} \xrightarrow{a} \Delta_p$, where $\Delta_p = p \cdot \overline{s_1} + (1-p) \cdot \overline{s_2}$, for any $p \in [0, 1]$. The set $\{\Delta_p \mid p \in [0, 1]\}$ is uncountable, though it is finitely generable, as the convex closure of the two-element set $\{\Delta_0, \Delta_1\}$. The formula $\bigoplus_{s \xrightarrow{a} \Delta} X_\Delta$ exploits that fact to bypass the infinite disjunction; this formula is finite if the underlying pLTS is finitary.

Theorem 2. *Let E be the characteristic equation system for strong probabilistic bisimulation on a given pLTS. Then, for all states s and t,*

1. *$s\,\mathcal{R}\,t$ for some strong probabilistic bisimulation $\mathcal{R}$ if and only if $\overline{t} \in \rho(X_s)$ for some post-fixpoint ρ of $\mathcal{F}_E$.*
2. *In particular, $s \sim t$ if and only if $\overline{t} \in [\![\nu_E(X_s)]\!]$, i.e., $\nu_E(X_s)$ is a characteristic formula for s w.r.t. strong probabilistic bisimilarity.*

Proof. Let E be the characteristic equation system for strong probabilistic bisimulation on a given pLTS. We only consider the first statement, from which the second statement follow immediately.

($\Leftarrow$) For this direction, assuming a post-fixpoint ρ of $\mathcal{F}_E$, we construct a probabilistic bisimulation relation that includes all state pairs (s,t) satisfying $\overline{t} \in \rho(X_s)$. Let $\mathcal{R}= \{\,(s,t) \mid \overline{t} \in \rho(X_s)\,\}$. We first show that

$$\Theta \in [\![X_\Delta]\!]_\rho \text{ implies } \Delta\,\mathcal{R}^\dagger\,\Theta. \tag{3}$$

Let $X_\Delta = \bigoplus_{i\in I} p_i \cdot {\downarrow} X_{s_i}$, so that $\Delta = \sum_{i\in I} p_i \cdot \overline{s_i}$. Suppose $\Theta \in [\![X_\Delta]\!]_\rho$. We have that $\Theta = \sum_{i\in I} p_i \cdot \Theta_i$ and, for all $i \in I$ and all $t \in \lceil \Theta_i \rceil$, that $\overline{t} \in [\![X_{s_i}]\!]_\rho$, i.e. $s_i\,\mathcal{R}\,t$. It follows that $\overline{s_i}\,\mathcal{R}^\dagger\,\Theta_i$ and thus $\Delta\,\mathcal{R}^\dagger\,\Theta$, using Proposition 2(1).

Now we show that $\mathcal{R}$ is a probabilistic bisimulation.

1. Suppose $s\,\mathcal{R}\,t$ and $s \xrightarrow{a} \Delta$. Then $\overline{t} \in \rho(X_s) \subseteq [\![\varphi_s]\!]_\rho$. It follows from (2) that $\overline{t} \in [\![\langle a \rangle X_\Delta]\!]_\rho$. So there exists some Θ such that $\overline{t} \xrightarrow{a} \Theta$ and $\Theta \in [\![X_\Delta]\!]_\rho$. Now we apply (3).
2. Suppose $s\,\mathcal{R}\,t$ and $t \xrightarrow{a} \Theta$. Then $\overline{t} \in \rho(X_s) \subseteq [\![\varphi_s]\!]_\rho$. It follows from (2) that $\overline{t} \in [\![[a] \bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta]\!]$. Notice that it must be the case that $\overline{s} \xrightarrow{a}$, otherwise, $\overline{t} \in [\![[a]\mathtt{false}]\!]_\rho$ and thus $t \not\xrightarrow{a}$, in contradiction with the assumption $t \xrightarrow{a} \Theta$. Therefore, $\Theta \in [\![\bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta]\!]_\rho$, which implies $\Theta \in [\![X_\Delta]\!]_\rho$ for some Δ with $\overline{s} \xrightarrow{a} \Delta$. Now we apply (3).

($\Rightarrow$) Given a strong probabilistic bisimulation $\mathcal{R}$, we construct a post-fixpoint of $\mathcal{F}_E$ such that whenever $s\,\mathcal{R}\,t$ then $\overline{t}$ falls into the set of distributions assigned to X_s by that post-fixpoint. We define the environment $\rho_\mathcal{R}$ by

$$\rho_\mathcal{R}(X_s) := \{\,\overline{t} \mid s\,\mathcal{R}\,t\,\}$$

and show that $\rho_\mathcal{R}$ is a post-fixpoint of $\mathcal{F}_E$, i.e.

$$\rho_\mathcal{R} \sqsubseteq \mathcal{F}_E(\rho_\mathcal{R}). \tag{4}$$

We first show that

$$\Delta\,\mathcal{R}^\dagger\,\Theta \text{ implies } \Theta \in [\![X_\Delta]\!]_{\rho_\mathcal{R}}. \tag{5}$$

Suppose $\Delta\,\mathcal{R}^\dagger\,\Theta$, we have that (i) $\Delta = \sum_{i\in I} p_i \cdot \overline{s_i}$, (ii) $\Theta = \sum_{i\in I} p_i \cdot \overline{t_i}$, (iii) $s_i\,\mathcal{R}\,t_i$ for all $i \in I$. We know from (iii) that $\overline{t_i} \in [\![X_{s_i}]\!]_{\rho_\mathcal{R}}$ and thus $\overline{t_i} \in [\![{\downarrow} X_{s_i}]\!]_{\rho_\mathcal{R}}$. Using (ii) we have that $\Theta \in [\![\bigoplus_{i\in I} p_i \cdot {\downarrow} X_{s_i}]\!]_{\rho_\mathcal{R}}$. Using (i) we obtain $\Theta \in [\![X_\Delta]\!]_{\rho_\mathcal{R}}$.

Now we are in a position to show (4). Suppose $\overline{t} \in \rho_{\mathcal{R}}(X_s)$. We must prove that $\overline{t} \in [\![\varphi_s]\!]_{\rho_{\mathcal{R}}}$, i.e.

$$\overline{t} \in \big(\bigcap_{s \xrightarrow{a} \Delta} [\![\langle a \rangle X_\Delta]\!]_{\rho_{\mathcal{R}}} \big) \cap \big(\bigcap_{a \in \mathsf{Act}_\tau} [\![[a] \bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta]\!]_{\rho_{\mathcal{R}}} \big)$$

by (2). This can be done by showing that $\overline{t}$ belongs to each of the two parts of the outermost intersection.

1. Assume that $s \xrightarrow{a} \Delta$ for some $a \in \mathsf{Act}_\tau$ and $\Delta \in \mathcal{D}(S)$. Since $s \mathrel{\mathcal{R}} t$, there exists some Θ such that $\overline{t} \xrightarrow{a} \Theta$ and $\Delta \mathrel{\mathcal{R}^\dagger} \Theta$. By (5), we get $\Theta \in [\![X_\Delta]\!]_{\rho_{\mathcal{R}}}$. It follows that $\overline{t} \in [\![\langle a \rangle X_\Delta]\!]_{\rho_{\mathcal{R}}}$.
2. Let $a \in Act_\tau$. Whenever $\overline{t} \xrightarrow{a} \Theta$, then by $s \mathrel{\mathcal{R}} t$ there must be some Δ such that $\overline{s} \xrightarrow{a} \Delta$ and $\Delta \mathrel{\mathcal{R}^\dagger} \Theta$. By (5), we get $\Theta \in [\![X_\Delta]\!]_{\rho_{\mathcal{R}}}$ and thus $\Theta \in [\![\bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta]\!]_{\rho_{\mathcal{R}}}$. As a consequence, $\overline{t} \in [\![[a] \bigvee_{\overline{s} \xrightarrow{a} \Delta} X_\Delta]\!]_{\rho_{\mathcal{R}}}$. □

Strong probabilistic simulation In a simulation game, if state s is related to state t, we only need to check that all transitions initiated by s should be matched up by transitions from t, and we do not care about the inverse direction: the capability of s to simulate t. Therefore, it is not surprising that characteristic equation systems for strong probabilistic simulation are defined as in Definition 6 except that we drop the second part of the conjunction in (2), so φ_s takes the form

$$\varphi_s := \bigwedge_{s \xrightarrow{a} \Delta} \langle a \rangle X_\Delta \tag{6}$$

With this modification, we have the expected property for strong probabilistic simulation, which can be shown by using the ideas in the proof of Theorem 2, but with fewer cases to analyse.

Weak probabilistic bisimulation Characteristic equation systems for weak probabilistic bisimulation are defined as in Definition 6 except that the weak semantics of pMu is employed and φ_s takes the form

$$\varphi_s := \big(\bigwedge_{s \xrightarrow{a} \Delta} \langle a \rangle X_\Delta \big) \wedge \big(\bigwedge_{a \in \mathsf{Act}_\tau} [a] \bigvee_{\overline{s} \xRightarrow{\hat{a}} \Delta} X_\Delta \big)^{\,2} \tag{7}$$

Weak probabilistic simulation Characteristic equation systems for weak probabilistic simulation are in exactly the same form as characteristic equation systems for strong probabilistic simulation (cf. (6)), but using the weak semantics of pMu.

Forward simulation Characteristic equation systems for forward simulation are in the same form as characteristic equation systems for weak probabilistic simulation, but with $X_\Delta := \bigoplus_{s \in \lceil \Delta \rceil} \Delta(s) \cdot X_s$, i.e. dropping the $\downarrow$.

[2] Using results from Markov Decision Processes [28], in a finitary pLTS also this infinite disjunction can be expressed as finite convex combination; however, we will not elaborate this here.

Failure simulation To give a modal characterisations for failure simulation we need to add modal formulae of the form $\mathbf{ref}(A)$ with $A \subseteq \mathsf{Act}$, first introduced in [7], to pMu, with the meaning given by

$$[\![\mathbf{ref}(A)]\!]_\rho = \{\, \Delta \in \mathcal{D}(S) \mid \exists \Delta' : \Delta \stackrel{\hat{\tau}}{\Longrightarrow} \Delta' \ \wedge\ \Delta' \stackrel{A}{\not\longrightarrow} \}$$

The formula $\mathbf{ref}(A)$ holds for Δ if by doing internal actions only Δ can evolve into a distribution such that no state in its support can perform an action from $A \cup \{\tau\}$. This time φ_s takes the form

$$\varphi_s := \begin{cases} \bigwedge_{s \stackrel{a}{\longrightarrow} \Delta} \langle a \rangle X_\Delta & \text{if } s \stackrel{\tau}{\longrightarrow} \\ (\bigwedge_{s \stackrel{a}{\longrightarrow} \Delta} \langle a \rangle X_\Delta) \wedge \mathbf{ref}(\{\, a \mid s \stackrel{a}{\not\longrightarrow} \}) & \text{otherwise} \end{cases} \tag{8}$$

with $X_\Delta := \bigoplus_{s \in \lceil \Delta \rceil} \Delta(s) \cdot X_s$. Inspired by [7], here we distinguish two cases, depending on the possibility of making an internal transition from s.

With the above modifications, we have the counterpart of Theorem 2, with a similar proof.

Theorem 3. *Let $E_\prec$ be the characteristic equation system for strong probabilistic simulation on a given pLTS. Let $E_\approx$ ($E_{\precapprox}, E_{\lhd_S}, E_{\lhd_{FS}}$, respectively) be the characteristic equation system for weak probabilistic bisimulation (weak probabilistic simulation, forward simulation, failure simulation, respectively) on a given divergence-free pLTS. Then, for all states s, t and distributions Θ,*

1. *$s\ \mathcal{R}\ t$ for some strong probabilistic simulation (weak probabilistic bisimulation, weak probabilistic simulation, respectively) $\mathcal{R}$ if and only if $\overline{t} \in \rho(X_s)$ for some post-fixpoint ρ of $\mathcal{F}_{E_\prec}$ ($\mathcal{F}_{E_\approx}$, $\mathcal{F}_{E_{\precapprox}}$, respectively).*
2. *$s\ \mathcal{R}\ \Theta$ for some forward simulation (failure simulation) $\mathcal{R}$ if and only if $\Theta \in \rho(X_s)$ for some post-fixpoint ρ of $\mathcal{F}_{E_{\lhd_S}}$ ($\mathcal{F}_{E_{\lhd_{FS}}}$).*
3. *In particular,*
   (a) $s \prec t$ if and only if $\overline{t} \in [\![\nu_{E_\prec}(X_s)]\!]$.
   (b) $s \approx t$ if and only if $\overline{t} \in [\![\nu_{E_\approx}(X_s)]\!]$.
   (c) $s \precapprox t$ if and only if $\overline{t} \in [\![\nu_{E_{\precapprox}}(X_s)]\!]$.
   (d) $s \lhd_{FS} \Theta$ if and only if $\Theta \in [\![\nu_{E_{\lhd_S}}(X_s)]\!]$.
   (e) $s \lhd_{FS} \Theta$ if and only if $\Theta \in [\![\nu_{E_{\lhd_{FS}}}(X_s)]\!]$. □

We can also consider the strong case for $\lhd_S$ and $\lhd_{FS}$ by treating τ as an external action, and give characteristic equation systems. In the strong case for $\lhd_{FS}$ only the "otherwise" in (8) applies, with $\mathbf{ref}(A)$ represented as $\bigwedge_{a \in A} [a]\mathtt{false}$.

5 Modal characterisations

In the previous sections we have pursued logical characterisations for various behavioural relations by characteristic formulae. A weaker form of characterisation, which is commonly called a modal characterisation of a behavioural relation, consists of isolating a class of formulae with the property that two states are equivalent if and only if they satisfy the same formulae from that class.

Definition 7. Let $\mathcal{L}^\mu_\sim$ be simply the class $\mathcal{L}$ of modal formulae defined in Section 3, equipped with the strong semantics of Table 1. With $\mathcal{L}^\mu_\prec$ we denote the fragment of this class obtained by skipping the modalities $\neg$ and $[a]$. The classes $\mathcal{L}^\mu_\approx$ and $\mathcal{L}^\mu_{\precsim}$ are defined likewise, but equipped with the weak semantics. Moreover, $\mathcal{L}^\mu_{\lhd_S}$ is the fragment of $\mathcal{L}^\mu_{\precsim}$ obtained by skipping $\downarrow$, and $\mathcal{L}^\mu_{\lhd_{FS}}$ is obtained from $\mathcal{L}^\mu_{\lhd_S}$ by addition of the modality $\mathbf{ref}(A)$.

In all cases, dropping the superscript μ denotes the subclass obtained by dropping the variables and fixpoint operators.

For $\mathcal{R} \in \{\sim, \prec, \approx, \precsim, \lhd_S, \lhd_{FS}\}$ we write $\Delta \sqsubseteq^\mu_\mathcal{R} \Theta$ just when $\Delta \in [\![\varphi]\!] \Rightarrow \Theta \in [\![\varphi]\!]$ for all closed $\varphi \in \mathcal{L}^\mu_\mathcal{R}$, and $\Delta \sqsubseteq_\mathcal{R} \Theta$ just when $\Delta \in [\![\varphi]\!] \Rightarrow \Theta \in [\![\varphi]\!]$ for all $\varphi \in \mathcal{L}_\mathcal{R}$.

Note that the relations $\sqsubseteq^\mu_\sim$, $\sqsubseteq^\mu_\approx$, $\sqsubseteq_\sim$ and $\sqsubseteq_\approx$ are symmetric. For this reason we will employ the symbol $\equiv$ instead of $\sqsubseteq$ when referring to them.

We have the following modal characterisation for strong probabilistic bisimilarity, strong probabilistic similarity, weak probabilistic bisimilarity, weak probabilistic similarity, forward similarity, and failure similarity.

Theorem 4 (Modal characterisation).
Let s and t be states in a divergence-free pLTS.

1. $s \sim t$ iff $\overline{s} \equiv^\mu_\sim \overline{t}$ iff $\overline{s} \equiv_\sim \overline{t}$.
2. $s \prec t$ iff $\overline{s} \sqsubseteq^\mu_\prec \overline{t}$ iff $\overline{s} \sqsubseteq_\prec \overline{t}$.
3. $s \approx t$ iff $\overline{s} \equiv^\mu_\approx \overline{t}$ iff $\overline{s} \equiv_\approx \overline{t}$.
4. $s \precsim t$ iff $\overline{s} \sqsubseteq^\mu_\precsim \overline{t}$ iff $\overline{s} \sqsubseteq_\precsim \overline{t}$.
5. $s \lhd_S \Theta$ iff $\overline{s} \sqsubseteq^\mu_{\lhd_S} \Theta$ iff $\overline{s} \sqsubseteq_{\lhd_S} \Theta$.
6. $s \lhd_{FS} \Theta$ iff $\overline{s} \sqsubseteq^\mu_{\lhd_{FS}} \Theta$ iff $\overline{s} \sqsubseteq_{\lhd_{FS}} \Theta$.

Note that $\overline{s} \equiv^\mu_\sim \overline{t} \Rightarrow s \sim t$ is an immediate consequence of Theorem 2: From $s \sim s$ we obtain $\overline{s} \in [\![\nu_E(X_s)]\!]$. Together with $\overline{s} \equiv^\mu_\sim \overline{t}$ this yields $\overline{t} \in [\![\nu_E(X_s)]\!]$, hence $s \sim t$.

Proof. We only prove the first statement; the others can be shown analogously. In fact we establish the more general result that

$$\Delta \sim^\dagger \Theta \qquad \Leftrightarrow \qquad \Delta \equiv^\mu_\sim \Theta \qquad \Leftrightarrow \qquad \Delta \equiv_\sim \Theta$$

from which statement 1 of Theorem 4 follows immediately. The implication $\Delta \sim^\dagger \Theta \Rightarrow \Delta \equiv^\mu_\sim \Theta$ expresses the *soundness* of the logic $\mathcal{L}^\mu_\sim$ w.r.t. the relation $\sim^\dagger$, whereas the implication $\Delta \equiv_\sim \Theta \Rightarrow \Delta \sim^\dagger \Theta$ expresses the *completeness* of $\mathcal{L}_\sim$ w.r.t. $\sim^\dagger$. The implication $\Delta \equiv^\mu_\sim \Theta \Rightarrow \Delta \equiv_\sim \Theta$ is trivial.

(Soundness) An environment $\rho : \mathsf{Var} \to \mathcal{P}(\mathcal{D}(S))$ is called *compatible with* $\sim^\dagger$ if for all $X \in \mathsf{Var}$ we have that

$$\Delta \sim^\dagger \Theta \Rightarrow (\Delta \in \rho(X) \Rightarrow \Theta \in \rho(X)).$$

We will show by structural induction on φ that

$$\Delta \sim^\dagger \Theta \Rightarrow (\Delta \in [\![\varphi]\!]_\rho \Rightarrow \Theta \in [\![\varphi]\!]_\rho)$$

for any environment ρ that is compatible with $\sim^\dagger$. By restricting attention to closed φ this implies the soundness of $\mathcal{L}^\mu_\sim$ w.r.t. $\sim^\dagger$. We consider a few interesting cases.

- Let $\Delta \sim^\dagger \Theta$ and $\Delta \in [\![\langle a\rangle\varphi]\!]_\rho$. Then $\Delta \xrightarrow{a} \Delta'$ and $\Delta' \in [\![\varphi]\!]_\rho$ for some Δ'. It follows that there is some Θ' with $\Theta \xrightarrow{a} \Theta'$ and $\Delta' \sim^\dagger \Theta'$. By induction we have $\Theta' \in [\![\varphi]\!]_\rho$, thus $\Theta \models \langle a\rangle\varphi$.
- Let $\Delta \sim^\dagger \Theta$ and $\Delta \in [\![[a]\varphi]\!]_\rho$. Suppose $\Theta \xrightarrow{a} \Theta'$. It can be seen that there is a Δ' with $\Delta \xrightarrow{a} \Delta'$ and $\Delta' \sim^\dagger \Theta'$. As $\Delta \in [\![[a]\varphi]\!]_\rho$ it must be that $\Delta' \in [\![\varphi]\!]_\rho$, and by induction we have $\Theta' \in [\![\varphi]\!]_\rho$. Thus $\Theta \in [\![[a]\varphi]\!]_\rho$.
- Let $\Delta \sim^\dagger \Theta$ and $\Delta \in [\![\bigoplus_{i\in I} p_i \cdot \varphi_i]\!]_\rho$. So $\Delta = \sum_{i\in i} p_i \cdot \Delta_i$ and for all $i \in I$ we have $\Delta_i \in [\![\varphi_i]\!]_\rho$. Since $\Delta \sim^\dagger \Theta$, by Proposition 2(2) we have $\Theta = \sum_{i\in I} p_i \cdot \Theta_i$ and $\Delta_i \sim^\dagger \Theta_i$. So by induction we have $\Theta_i \in [\![\varphi_i]\!]_\rho$ for all $i \in I$. Therefore, $\Theta \in [\![\bigoplus_{i\in I} p_i \cdot \varphi_i]\!]_\rho$. The case $\Delta \in [\![\bigoplus_{i\in I} \varphi_i]\!]_\rho$ goes likewise.
- Let $\Delta \sim^\dagger \Theta$ and $\Delta \in [\![{\downarrow}\varphi]\!]_\rho$. So for all $s \in \lceil\Delta\rceil$ we have $\overline{s} \in [\![\varphi]\!]_\rho$. From $\Delta \sim^\dagger \Theta$ it follows that for each $t \in \lceil\Theta\rceil$ there is an $s \in \lceil\Delta\rceil$ with $s \sim t$, thus $\overline{s} \sim^\dagger \overline{t}$. So by induction we have $\overline{t} \in [\![\varphi]\!]_\rho$ for all $t \in \lceil\Theta\rceil$. Therefore, $\Theta \in [\![{\downarrow}\varphi]\!]_\rho$.
- Suppose $\Delta \sim^\dagger \Theta$ and $\Theta \notin [\![\mu X.\varphi]\!]_\rho$. Then $\exists V \subseteq \mathcal{D}(S)$ with $\Theta \notin V$ and $[\![\varphi]\!]_{\rho[X\mapsto V]} \subseteq V$. Let $V' := \{\Delta' \mid \forall\Theta'.(\Delta' \sim^\dagger \Theta' \Rightarrow \Theta' \in V)\}$. Then $\Delta \notin V'$. It remains to show that $[\![\varphi]\!]_{\rho[X\mapsto V']} \subseteq V'$, because this implies $\Delta \notin [\![\mu X.\varphi]\!]_\rho$, which has to be shown.
  So let $\Delta' \in [\![\varphi]\!]_{\rho[X\mapsto V']}$. Take any Θ' with $\Delta' \sim^\dagger \Theta'$. By construction of V', the environment $\rho[X \mapsto V']$ is compatible with $\sim^\dagger$. Therefore, the induction hypothesis yields $\Theta' \in [\![\varphi]\!]_{\rho[X\mapsto V']}$. We have $V' \subseteq V$, and as $[\![\ ]\!]$ is monotonic we obtain $\Theta' \in [\![\varphi]\!]_{\rho[X\mapsto V']} \subseteq [\![\varphi]\!]_{\rho[X\mapsto V]} \subseteq V$. It follows that $\Delta' \in V'$.
- Suppose $\Delta \sim^\dagger \Theta$ and $\Delta \in [\![\nu X.\varphi]\!]_\rho$. Then $\exists V \subseteq \mathcal{D}(S)$ with $\Delta \in V$ and $[\![\varphi]\!]_{\rho[X\mapsto V]} \supseteq V$. Let $V' := \{\Theta' \mid \exists\Delta' \in V.\ \Delta' \sim^\dagger \Theta'\}$. Then $\Theta \in V'$. It remains to show that $[\![\varphi]\!]_{\rho[X\mapsto V']} \supseteq V'$, because this implies $\Theta \in [\![\nu X.\varphi]\!]_\rho$, which has to be shown.
  So let $\Theta' \notin [\![\varphi]\!]_{\rho[X\mapsto V']}$. Take any Δ' with $\Delta' \sim^\dagger \Theta'$. By construction of V', the environment $\rho[X \mapsto V']$ is compatible with $\sim^\dagger$. Therefore, the induction hypothesis yields $\Delta' \notin [\![\varphi]\!]_{\rho[X\mapsto V']}$. We have $V' \supseteq V$, and as $[\![\ ]\!]$ is monotonic we obtain $\Delta' \notin [\![\varphi]\!]_{\rho[X\mapsto V']} \supseteq [\![\varphi]\!]_{\rho[X\mapsto V]} \supseteq V$. It follows that $\Theta' \notin V'$.

(Completeness) Let $\mathcal{R} = \{(s,t) \mid \overline{s} \equiv_\sim \overline{t}\}$. We show that $\mathcal{R}$ is a strong probabilistic bisimulation. Suppose $s\ \mathcal{R}\ t$ and $s \xrightarrow{a} \Delta$. We have to show that there is some Θ with $\overline{t} \xrightarrow{a} \Theta$ and $\Delta\ \mathcal{R}^\dagger\ \Theta$. Consider the set

$$T := \{\Theta \mid \overline{t} \xrightarrow{a} \Theta \wedge \Theta = \sum_{s'\in\lceil\Delta\rceil} \Delta(s') \cdot \Theta_{s'} \wedge \exists s' \in \lceil\Delta\rceil, \exists t' \in \lceil\Theta_{s'}\rceil : \overline{s'} \not\equiv_\sim \overline{t'}\}$$

For each $\Theta \in T$ there must be some $s'_\Theta \in \lceil\Delta\rceil$ and $t'_\Theta \in \lceil\Theta_{s'_\Theta}\rceil$ and a formula φ_Θ with $\overline{s'_\Theta} \models \varphi_\Theta$ but $\overline{t'_\Theta} \not\models \varphi_\Theta$. So $\overline{s'} \models \bigwedge_{\{\Theta\in T \mid s'_\Theta = s'\}} \varphi_\Theta$ for each $s' \in \lceil\Delta\rceil$, and for each $\Theta \in T$ with $s'_\Theta = s'$ there is some $t'_\Theta \in \lceil\Theta_{s'}\rceil$ with $\overline{t'_\Theta} \not\models \bigwedge_{\{\Theta\in T \mid s'_\Theta = s'\}} \varphi_\Theta$. Let

$$\varphi := \langle a\rangle \bigoplus_{s'\in\lceil\Delta\rceil} \Delta(s') \cdot {\downarrow} \bigwedge_{\{\Theta\in T \mid s'_\Theta = s'\}} \varphi_\Theta.$$

It is clear that $\overline{s} \models \varphi$, hence $\overline{t} \models \varphi$ by $s\mathcal{R}t$. It follows that there must be a Θ^* with $\overline{t} \xrightarrow{a} \Theta^*$, $\Theta^* = \sum_{s'\in\lceil\Delta\rceil} \Delta(s') \cdot \Theta^*_{s'}$ and for each $s' \in \lceil\Delta\rceil$, $t' \in \lceil\Theta^*_{s'}\rceil$

we have $\overline{t'} \models \bigwedge_{\{\Theta \in T \mid s'_\Theta = s'\}} \varphi_\Theta$. This means that $\Theta^* \notin T$ and hence for each $s' \in \lceil \Delta \rceil$, $t' \in \lceil \Theta^*_{s'} \rceil$ we have $\overline{s'} \equiv_\sim \overline{t'}$, i.e. $s' \mathcal{R} t'$. Consequently, we obtain $\Delta \, \mathcal{R}^\dagger \, \Theta^*$. By symmetry all transitions of t can be matched up by transitions of s. □

Modal characterisation of strong and weak probabilistic bisimulation has been studied in [27]. It is also based on a probabilistic extension of the Hennessy-Milner logic. Instead of our modalities $\bigoplus$ and $\downarrow$ they use a modality $[\cdot]_p$. Intuitively, a distribution Δ satisfies the formula $[\varphi]_p$ when the set of states satisfying φ is measured by Δ with probability at least p. So the formula $[\varphi]_p$ can be expressed by our logics in terms of the probabilistic choice $\bigoplus_{i\in I} p_i \cdot \varphi_i$ by setting $I = \{1,2\}$, $p_1 = p$, $p_2 = 1-p$, $\varphi_1 = \downarrow\varphi$, and $\varphi_2 = \mathtt{true}$. Furthermore, instead of our modality $\langle a \rangle$, they use a modality $\Diamond a$ that can be expressed in our logic by $\Diamond a\varphi = \langle a \rangle \downarrow \varphi$. We conjecture that our modalities $\langle a \rangle$ and $\bigoplus$ cannot be expressed in terms of the logic of [27], and that a logic of that type is unsuitable for characterising forward simulation or failure simulation.

When restricted to deterministic pLTSs (i.e., for each state and for each action, there exists at most one outgoing transition), probabilistic bisimulations can be characterised by simpler forms of logics, as observed in [20,11,27].

References

1. E. Bandini & R. Segala (2001): *Axiomatizations for Probabilistic Bisimulation.* In: *Proc. ICALP'01*, LNCS 2076, Springer, pp. 370–381.
2. I. Christoff (1990): *Testing equivalences and fully abstract models for probabilistic processes.* In: *Proc. CONCUR'90*, LNCS 458, Springer, pp. 126–140.
3. R. Cleaveland, S. Purushothaman Iyer & M. Narasimha (2005): *Probabilistic temporal logics via the modal mu-calculus. Theoretical Computer Science* 342(2-3), pp. 316–350.
4. P.R. D'Argenio, N. Wolovick, P.S. Terraf & P. Celayes (2009): *Nondeterministic Labeled Markov Processes: Bisimulations and Logical Characterization.* In: *Proc. QEST'09*, IEEE Computer Society, pp. 11–20.
5. Y. Deng & W. Du (2009): *A Local Algorithm for Checking Probabilistic Bisimilarity.* In: *Proc. FCST'09*, IEEE Computer Society, pp. 401–407.
6. Y. Deng & R.J. van Glabbeek. *Characterising Probabilistic Processes Logically.* Available at `http://arxiv.org/abs/1007.5188`. Full version of this paper.
7. Y. Deng, R.J. van Glabbeek, M. Hennessy & C.C. Morgan (2008): *Characterising Testing Preorders for Finite Probabilistic Processes. Logical Methods in Computer Science* 4(4):4.
8. Y. Deng, R.J. van Glabbeek, M. Hennessy & C.C. Morgan (2009): *Testing Finitary Probabilistic Processes.* In: *Proc. CONCUR'09*, LNCS, Springer, pp. 274–288.
9. Y. Deng, R.J. van Glabbeek, M. Hennessy, C.C. Morgan & C. Zhang (2007): *Remarks on Testing Probabilistic Processes. ENTCS* 172, pp. 359–397.
10. Y. Deng & C. Palamidessi (2007): *Axiomatizations for probabilistic finite-state behaviors. Theoretical Computer Science* 373(1-2), pp. 92–114.
11. J. Desharnais, A. Edalat & P. Panangaden (1998): *A logical characterization of bisimulation for labelled Markov processes.* In: *Proc. LICS'98*, IEEE Computer Society, pp. 478–489.

12. A. Giacalone, C.-C. Jou & S.A. Smolka (1990): *Algebraic reasoning for probabilistic concurrent systems.* In: *Proc. IFIP TC 2 Working Conference on Programming Concepts and Methods*, pp. 443–458.
13. H. Hansson & B. Jonsson (1990): *A Calculus for Communicating Systems with Time and Probabilities.* In: *Proc. RTSS'90*, IEEE Computer Society, pp. 278–287.
14. H. Hansson & B. Jonsson (1994): *A Logic for Reasoning about Time and Reliability. Formal Aspects of Computing* 6(5), pp. 512–535.
15. M. Hennessy & R. Milner (1985): *Algebraic laws for nondeterminism and concurrency. Journal of the ACM* 32(1), pp. 137–161.
16. M. Huth & M. Kwiatkowska (1997): *Quantitative analysis and model checking.* In: *Proc. LICS'97*, IEEE Computer Society, pp. 111–122.
17. B. Jonsson & K.G. Larsen (1991): *Specification and refinement of probabilistic processes.* In: *Proc. LICS'91*, Computer Society Press, pp. 266–277.
18. B. Jonsson & Wang Yi (2002): *Testing preorders for probabilistic processes can be characterized by simulations. Theoretical Computer Science* 282(1), pp. 33–51.
19. D. Kozen (1983): *Results on the propositional mu-calculus. Theoretical Computer Science* 27, pp. 333–354.
20. K.G. Larsen & A. Skou (1991): *Bisimulation through probabilistic testing. Information and Computation* 94(1), pp. 1–28.
21. G. Lowe (1995): *Probabilistic and Prioritized Models of Timed CSP. Theoretical Computer Science* 138, pp. 315–352.
22. A.K. McIver & C.C. Morgan (1997): *An expectation-based model for probabilistic temporal logic.* Technical Report PRG-TR-13-97, Oxford University Computing Laboratory.
23. A.K. McIver & C.C. Morgan (2007): *Results on the Quantitative Mu-Calculus. ACM Transactions on Computational Logic* 8(1).
24. R. Milner (1989): *Communication and Concurrency.* Prentice Hall.
25. M.W. Mislove, J. Ouaknine & J. Worrell (2004): *Axioms for Probability and Nondeterminism. ENTCS* 96, pp. 7–28.
26. M. Müller-Olm (1998): *Derivation of Characteristic Formulae. ENTCS* 18, pp. 159–170.
27. A. Parma & R. Segala (2007): *Logical Characterizations of Bisimulations for Discrete Probabilistic Systems.* In: *Proc. FOSSACS'07*, LNCS 4423, Springer, pp. 287–301.
28. M.L. Puterman (1994): *Markov Decision Processes.* Wiley.
29. Y.S. Ramakrishna & S.A. Smolka (1997): *Partial-order reduction in the weak modal mu-calculus.* In: A. Mazurkiewicz & J. Winkowski, editors: *Proc. CONCUR'97*, LNCS 1234, Springer, pp. 5–24.
30. R. Segala (1995): *Modeling and Verification of Randomized Distributed Real-Time Systems.* Technical Report MIT/LCS/TR-676, PhD thesis, MIT, Dept. of EECS.
31. R. Segala & N.A. Lynch (1994): *Probabilistic Simulations for Probabilistic Processes.* In: *Proc. CONCUR'94*, LNCS 836, Springer, pp. 481–496.
32. B. Steffen & A. Ingólfsdóttir (1994): *Characteristic Formulae for Processes with Divergence. Information and Computation* 110, pp. 149–163.
33. A. Tarski (1955): *A lattice-theoretical fixpoint theorem and its application. Pacific Journal of Mathematics* 5(2), pp. 285–309.
34. R. Tix, K. Keimel & G.D. Plotkin (2005): *Semantic Domains for Combining Probability and Non-Determinism. ENTCS* 129, pp. 1–104.
35. Wang Yi & K.G. Larsen (1992): *Testing Probabilistic and Nondeterministic Processes.* In: *Proc. IFIP TC6/WG6.1 PSTV'92*, C-8, North-Holland, pp. 47–61.