A Spectrum of Behavioral Relations over LTSs on Probability Distributions

SILVIA CRAFA FRANCESCO RANZATO

University of Padova, Italy

Abstract. Probabilistic nondeterministic processes are commonly modeled as probabilistic LTSs (PLTSs, a.k.a. probabilistic automata). A number of logical characterizations of the main behavioral relations on PLTSs have been studied. In particular, Parma and Segala [2007] define a probabilistic Hennessy-Milner logic interpreted over distributions, whose logical equivalence/preorder when restricted to Dirac distributions coincide with standard bisimulation/simulation between the states of a PLTS. This result is here extended by studying the full logical equivalence/preorder between distributions in terms of a notion of bisimulation/simulation defined on a LTS of probability distributions (DLTS). We show that the standard spectrum of behavioral relations on nonprobabilistic LTSs as well as its logical characterization in terms of Hennessy-Milner logic scales to the probabilistic setting when considering DLTSs.

1 Introduction

Formal methods for concurrent and distributed system specification and verification have been extended to encompass randomized phenomena exhibited by the behavior of probabilistic systems. In a standard nonprobabilistic setting, systems are commonly modeled as labeled transition systems (LTSs) and model checking techniques are based on two major tools: temporal logics and behavioral relations. Logics are used to specify the properties that systems have to satisfy, while behavioral equivalence and preorder relations are used as appropriate abstractions that reduce the state space. Precise relationships have been established between these two approaches: van Glabbeek [8] shows how a wide spectrum of observational equivalences for concurrent processes is logically characterized in terms of Hennessy-Milner-like modal logics (HML).

A number of probabilistic behavioral relations and probabilistic temporal logics have been proposed (see e.g. [4,9,10,11,12,13,15]). Probabilistic LTSs (PLTSs, a.k.a. probabilistic automata) are a prominent model for formalizing probabilistic systems since they allow to model both probabilistic and nondeterministic behaviors. In PLTSs, a state s evolves through a labeled transition to a state distribution d that defines the probabilities of reaching the possible successor states of s. Accordingly, the standard probabilistic extension [15] of the simulation relation requires that if a state s progresses to a distribution d, then a simulating state s' needs to mimic such a transition by moving to a distribution d' that is related to d through a so-called weight function. This definition is a conservative extension of the simulation relation on LTSs since a LTS can be viewed as a particular PLTS where the target of transitions are Dirac distributions, i.e., distributions δ_s such that $\delta_s(s) = 1$ and $\delta_s(t) = 0$ for any $t \neq s$.

A number of modal logics have been proposed in order to provide a logical characterization of probabilistic simulation and bisimulation. Larsen and Skou [12]'s logic as well as Hansson and Jonsson [10]'s PCTL logic are interpreted over states of probabilistic systems, such as reactive models and discrete time Markov chains, that do not express nondeterminism. On the other hand, Parma and Segala [13] show that richer probabilistic models that encode pure nondeterminism (besides probabilistic choice), such as PLTSs, call for a richer logic. They propose a probabilistic extension of HML whose formulae are interpreted over distributions rather than states, and they show that two states s and t are similar (the same holds for bisimilarity) if and only if their corresponding Dirac distributions δ_s and δ_t satisfy the same set of formulae. However, nothing is stated about logically equivalent distributions that are not Dirac distributions.

In this paper we study the full logical equivalence between (possibly non-Dirac) distributions that is induced by Parma and Segala [13]'s logic. We show that this logic actually characterizes a novel and natural notion of simulation (bisimulation) between distributions of a PLTS, so that the standard state simulation (bisimulation) on PLTSs can be indeed retrieved by a suitable restriction to Dirac distributions. Furthermore, the transition relation of a PLTS is lifted to a transition relation between distributions that gives rise to a corresponding LTS on distributions (called DLTS). This allows us to lift behavioral relations on PLTSs to corresponding behavioral relations on DLTSs. Such a move from PLTSs to DLTSs yields a number of byproducts:

- Parma and Segala [13]'s logic turns out to be equivalent to a logic $\mathcal{L}$ whose diamond operator is interpreted on the DLTS in accordance with its standard interpretation on LTSs. Hence, this logic best suits as probabilistic extension of HML.
- This logic $\mathcal{L}$ characterizes a (bi)simulation relation between distributions which is equivalent to that characterized by Parma and Segala's logic, but that naturally admits a game-theoretic characterization.
- A spectrum of behavioral relations can be defined on DLTSs along the lines of the standard approach on LTSs [9]. These preorders and equivalences between distributions can be then projected back to states, thus providing a spectrum of (probabilistic) preorders and equivalences between states of PLTSs.
- Shifting the problem from PLTSs to DLTSs opens the way to the reuse of efficient model checking techniques available for LTSs.

This approach is studied on a number of well known probabilistic relations appearing in literature, namely simulation, probabilistic simulation, failure simulation, and their corresponding bisimulations. A discussion about related approaches is the subject of the final section, that also hints at future work.

2 Probabilistic Simulation and Bisimulation

Given a set X and a relation $R \subseteq X \times X$, we write xRy for $(x, y) \in R$; if $x \in X$ and $Y \subseteq X$ then $R(x) \triangleq \{y \in X \mid xRy\}$ and $R(Y) \triangleq \cup_{x \in Y} R(x)$.

$\mathrm{Distr}(X)$ denotes the set of (stochastic) distributions on a set X, i.e., the set of functions $d : X \to [0, 1]$ such that $\sum_{x \in X} d(x) = 1$. The support of a distribution d is defined by $\mathrm{supp}(d) \triangleq \{x \in X \mid d(x) > 0\}$; moreover, if $Y \subseteq X$, then $d(Y) \triangleq$

$\sum_{y \in Y} d(y)$. The Dirac distribution on $x \in X$, denoted by δ_x, is the distribution that assigns probability 1 to x (and 0 otherwise).

A probabilistic LTS (PLTS) is a tuple $\mathcal{M} = \langle \Sigma, Act, \rightarrow \rangle$ where Σ is a (denumerable) set of states, Act is a (denumerable) set of actions, and $\rightarrow \subseteq \Sigma \times Act \times \mathrm{Distr}(\Sigma)$ is a transition relation, where $(s, a, d) \in \rightarrow$ is tipically denoted by $s \xrightarrow{a} d$. For any $a \in Act$, the predecessor operator $\mathrm{pre}_a : \wp(\mathrm{Distr}(\Sigma)) \rightarrow \wp(\Sigma)$ is defined by $\mathrm{pre}_a(D) \triangleq \{s \in \Sigma \mid \exists d \in D.\ s \xrightarrow{a} d\}$.

The definitions of probabilistic behavioral relations often rely on so-called weight functions [14], that are used to lift a relation between states to a relation between distributions. We do not recall here the definition of weight functions, as we will use the following equivalent characterizations (see [6,11,17]).

Definition 2.1 (Lifting). Let $R \subseteq X \times X$ be any relation. Then, the *lifting* of R to distributions is the relation $\sqsubseteq_R \subseteq \mathrm{Distr}(X) \times \mathrm{Distr}(X)$ that can be equivalently defined in one of the following ways:

- $d \sqsubseteq_R e$ iff there exists a weight function for (d, e) w.r.t. R;
- $d \sqsubseteq_R e$ iff $d(U) \leq e(R(U))$ for any set $U \subseteq \mathrm{supp}(d)$;
- when R is an equivalence on X, then $d \sqsubseteq_R e$ iff $d(B) = e(B)$ for any equivalence class B of R. □

It is easy to see that if $R \subseteq R'$ then $\sqsubseteq_R \subseteq \sqsubseteq_{R'}$; moreover, if R is symmetric then $\sqsubseteq_R$ is also a symmetric relation, that we denote with $\equiv_R$.

Definition 2.2 (Simulation). Given a PLTS $\mathcal{M}$, a relation $R \subseteq \Sigma \times \Sigma$ is a simulation on $\mathcal{M}$ if for all $s, t \in \Sigma$ such that sRt,

- if $s \xrightarrow{a} d$ then there exists $e \in \mathrm{Distr}(\Sigma)$ such that $t \xrightarrow{a} e$ and $d \sqsubseteq_R e$. □

Let $R_{\mathrm{sim}} \triangleq \cup\{R \subseteq \Sigma \times \Sigma \mid R \text{ is a simulation on } \mathcal{M}\}$. Then, R_{sim} turns out to be a preorder relation which is the greatest simulation on $\mathcal{M}$ and is called simulation preorder on $\mathcal{M}$. Simulation equivalence P_{sim} on $\mathcal{M}$ is defined as the kernel of the simulation preorder, i.e., $P_{\mathrm{sim}} \triangleq R_{\mathrm{sim}} \cap R_{\mathrm{sim}}^{-1}$.

Definition 2.3 (Bisimulation). A symmetric relation $S \subseteq \Sigma \times \Sigma$ is a bisimulation on $\mathcal{M}$ if for all $s, t \in \Sigma$ such that sSt,

- if $s \xrightarrow{a} d$ then there exists $e \in \mathrm{Distr}(\Sigma)$ such that $t \xrightarrow{a} e$ and $d \equiv_S e$. □

Let $P_{\mathrm{bis}} \triangleq \cup\{S \subseteq \Sigma \times \Sigma \mid S \text{ is a bisimulation on } \mathcal{M}\}$. Then, P_{bis} turns out to be an equivalence relation which is the greatest bisimulation on $\mathcal{M}$ and is called bisimilarity on $\mathcal{M}$.

3 An Operational View of Probabilistic HML

In order to logically characterize behavioral relations on probabilistic models that encode pure nondeterminism, such as PLTSs, Parma and Segala [13] put forward an extension of Hennessy-Milner logic whose formulae are interpreted over distributions on

the states of a PLTS. They show that two states are bisimilar if and only if their corresponding Dirac distributions satisfy the same set of formulae. However, nothing is stated about logically equivalent distributions that are not Dirac distributions. In the following, we give a novel notion of simulation (and correspondingly bisimulation) between distributions which (i) characterizes the full logical equivalence of Parma and Segala's logic and (ii) boils down to standard simulation (and bisimulation) between the states of a PLTS when restricted to Dirac distributions.

Parma and Segala's logic [13] is syntactically defined as follows:

$$\phi ::= \top \mid \bigwedge_{i \in I} \phi_i \mid \neg\phi \mid \langle\!\rangle_a \phi \mid [\phi]_p$$

where I is a possibly infinite (denumerable) set of indices, $a \in Act$ and p is a rational number in $[0, 1]$. Given a PLTS $\langle \Sigma, Act, \rightarrow \rangle$, the semantics of the formulae is inductively defined as follows: for any distribution $d \in \mathrm{Distr}(\Sigma)$,

$$\begin{array}{l} d \models \top \\ d \models \bigwedge_I \phi_i \ \text{ iff for any } \ i \in I,\ d \models \phi_i \\ d \models \neg\phi \ \text{ iff } \ d \not\models \phi \\ d \models \langle\!\rangle_a \phi \ \text{ iff } \ \forall x \in \mathrm{supp}(d).\, \exists e \in \mathrm{Distr}(\Sigma).\ x \xrightarrow{a} e \text{ and } e \models \phi \\ d \models [\phi]_p \ \text{ iff } \ d(\{s \in \Sigma \mid \delta_s \models \phi\}) \geq p \end{array}$$

The first three clauses are standard. The modal connective $\langle\!\rangle_a$ is a probabilistic counterpart of HML's diamond operator. $\langle\!\rangle_a \phi$ is satisfied by a distribution $d \in \mathrm{Distr}(\Sigma)$ whenever *any state* $x \in \mathrm{supp}(d)$ reaches through an a-labeled transition a distribution e that satisfies the formula ϕ. As the formulae $\langle\!\rangle_a \phi$ only deal with transitions of the PLTS, a further modal operator $[\cdot]_p$ needs to take into account the probabilities that distributions assign to sets of related states. More precisely, a distribution d satisfies a formula $[\phi]_p$ when d assigns a probability at least p to the set of states whose Dirac distributions satisfy the formula ϕ. This logic is here referred to as $\mathcal{L}_\forall$ in order to stress the *universal* nature of its diamond operator $\langle\!\rangle_a$.

Definition 3.1 (Logical equivalence and preorder). Two distributions $d, e \in \mathrm{Distr}(\Sigma)$ are logically equivalent for $\mathcal{L}_\forall$, written $d \equiv_{\mathcal{L}_\forall} e$, when, for any $\phi \in \mathcal{L}_\forall$, $d \models \phi$ iff $e \models \phi$. We write $d \leq_{\mathcal{L}_\forall} e$ for the corresponding logical preorder, i.e., when for any $\phi \in \mathcal{L}_\forall$, $d \models \phi$ implies $e \models \phi$. □

Let $\mathcal{L}^+_\forall$ be the negation-free and finitely disjunctive fragment of $\mathcal{L}_\forall$, that is:

$$\phi ::= \top \mid \bigwedge_{i \in I} \phi_i \mid \phi \vee \phi \mid \langle\!\rangle_a \phi \mid [\phi]_p$$

The following result by Parma and Segala [13] (see also [11]) shows that the logical equivalence induced by $\mathcal{L}_\forall$ and the logical preorder induced by $\mathcal{L}^+_\forall$, when restricted to Dirac distributions, correspond, respectively, to bisimulation and simulation. Notice that the simulation preorder is logically characterized by negation-free formulae, reflecting the fact that simulation, differently from bisimulation, is not a symmetric relation. However, the logic for simulation requires finite disjunction to characterize probabilistic choice.

Theorem 3.2 ([13]). *Consider R_{sim} and P_{bis} on a given PLTS. Then, for all $s, t \in \Sigma$,*

- *$s \, R_{\mathrm{sim}} \, t$ if and only if $\delta_s \leq_{\mathcal{L}^+_\forall} \delta_t$;*
- *$s \, P_{\mathrm{bis}} \, t$ if and only if $\delta_s \equiv_{\mathcal{L}_\forall} \delta_t$.*

Our main goal is to define a notion of simulation and bisimulation between distributions that represents the operational match of the full logical preorder $\leq_{\mathcal{L}^+_\forall}$ and equivalence $\equiv_{\mathcal{L}_\forall}$ between distributions. Firstly, notice that any relation on distributions $\mathcal{R} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ embeds a corresponding relation on states that can be obtained by restricting $\mathcal{R}$ to Dirac distributions. This is formalized by a mapping $\Delta : \wp(\mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)) \to \wp(\Sigma \times \Sigma)$ defined as follows:

$$\Delta(\mathcal{R}) \triangleq \{(s,t) \in \Sigma \times \Sigma \mid (\delta_s, \delta_t) \in \mathcal{R}\}.$$

Note that if $\mathcal{R}$ is a symmetric/preorder/equivalence relation then $\Delta(\mathcal{R})$ is correspondingly a symmetric/preorder/equivalence relation on Σ.

Our definition of (bi)simulation between distributions (called d-(bi)simulation) is directly inspired by the logic $\mathcal{L}_\forall$. In particular, the two distinctive modal operators of $\mathcal{L}_\forall$ are mirrored in two defining conditions of (bi)simulation between distributions. More precisely, the semantics of the diamond operator suggests a kind of transfer property that (bi)similar distributions should respect (cf. condition (1)). On the other hand, a second condition, peculiar of the probabilistic setting, deals with the probabilities assigned by (bi)similar distributions to sets of related states (cf. condition (2)).

Definition 3.3 (∀d-simulation). A relation $\mathcal{R} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ is a *∀d-simulation* on a PLTS if for all $d, e \in \mathrm{Distr}(\Sigma)$, if $d \, \mathcal{R} \, e$ then:

(1) for all $D \subseteq \mathrm{Distr}(\Sigma)$, if $\mathrm{supp}(d) \subseteq \mathrm{pre}_a(D)$ then $\mathrm{supp}(e) \subseteq \mathrm{pre}_a(\mathcal{R}(D))$;
(2) $d \sqsubseteq_{\Delta(\mathcal{R})} e$. □

Definition 3.4 (∀d-bisimulation). A symmetric relation $\mathcal{S} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ is a *∀d-bisimulation* on a PLTS if for all $d, e \in \mathrm{Distr}(\Sigma)$, if $d \, \mathcal{S} \, e$ then:

(1) for all $D \subseteq \mathrm{Distr}(\Sigma)$, if $\mathrm{supp}(d) \subseteq \mathrm{pre}_a(D)$ then $\mathrm{supp}(e) \subseteq \mathrm{pre}_a(\mathcal{S}(D))$;
(2) $d \equiv_{\Delta(\mathcal{S})} e$. □

Given a PLTS $\mathcal{M}$, let $\mathcal{R}^\forall_{\mathrm{sim}} \triangleq \cup\{\mathcal{R} \mid \mathcal{R} \text{ is a } \forall\text{d-simulation on } \mathcal{M}\}$. Then, it turns out that $\mathcal{R}^\forall_{\mathrm{sim}}$ is the greatest ∀d-simulation on $\mathcal{M}$ and is a preorder, called the ∀d-simulation preorder on $\mathcal{M}$. Analogously, let $\mathcal{P}^\forall_{\mathrm{bis}} \triangleq \cup\{\mathcal{S} \mid \mathcal{S} \text{ is a } \forall\text{d-bisimulation on } \mathcal{M}\}$, so that $\mathcal{P}^\forall_{\mathrm{bis}}$ turns out to be the greatest ∀d-bisimulation on $\mathcal{M}$ and an equivalence relation, called the ∀d-bisimilarity on $\mathcal{M}$

It turns out that ∀d-simulation preorder fully captures the logical preorder induced by $\mathcal{L}^+_\forall$ while ∀d-bisimilarity fully captures the logical equivalence induced by $\mathcal{L}_\forall$.

Theorem 3.5. *For any $d, e \in \mathrm{Distr}(\Sigma)$,*

- *$d \, \mathcal{R}^\forall_{\mathrm{sim}} \, e$ if and only $d \leq_{\mathcal{L}^+_\forall} e$;*
- *$d \, \mathcal{P}^\forall_{\mathrm{bis}} \, e$ if and only $d \equiv_{\mathcal{L}_\forall} e$.*

A closer look at the semantics of the diamond operator of $\mathcal{L}_\forall$ points out a key difference with the semantics of the standard diamond operator in HML. In the case of LTSs, the diamond operator of HML induces the predecessor operator of the LTS. Similarly, the semantic definition of the diamond operator of $\mathcal{L}_\forall$ induces the following operator $\mathrm{ppre}_a^\forall$, that we call probabilistic predecessor operator:

$$\mathrm{ppre}_a^\forall : \wp(\mathrm{Distr}(\Sigma)) \to \wp(\mathrm{Distr}(\Sigma))$$
$$\mathrm{ppre}_a^\forall(D) \triangleq \{d \in \mathrm{Distr}(\Sigma) \mid \mathrm{supp}(d) \subseteq \mathrm{pre}_a(D)\}$$

where $\mathrm{pre}_a : \wp(\mathrm{Distr}(\Sigma)) \to \wp(\Sigma)$ is the PLTS predecessor operator. However, differently from the predecessor operators of LTSs and PLTSs, this probabilistic predecessor $\mathrm{ppre}_a^\forall$ *does not preserve set unions*, i.e., it is not true in general that, for any $D_1, D_2 \subseteq \mathrm{Distr}(\Sigma)$, $\mathrm{ppre}_a^\forall(D_1 \cup D_2) = \mathrm{ppre}_a^\forall(D_1) \cup \mathrm{ppre}_a^\forall(D_2)$. In fact, $\mathrm{supp}(d) \subseteq \mathrm{pre}_a(D_1 \cup D_2)$ does not imply $\mathrm{supp}(d) \subseteq \mathrm{pre}_a(D_1)$ nor $\mathrm{supp}(d) \subseteq \mathrm{pre}_a(D_2)$. It is here worth noting that, in general, an operator $f : \wp(X) \to \wp(X)$ defined on a powerset $\wp(X)$ preserves set unions if and only if there exists a relation $R \subseteq X \times X$ whose corresponding predecessor operator $\mathrm{pre}_R = \lambda Y.\{x \in X \mid \exists y \in Y. xRy\}$ coincides with f. As a consequence, *one cannot define* a transition relation between distributions whose corresponding predecessor operator coincides with $\mathrm{ppre}_a^\forall$. The lack of a transition relation between distributions is particularly troublesome when defining coinductive behavioral relations between distributions. Consider the transfer property of $\forall$d-simulations, namely condition (1) of Definition 3.3: this can be equivalently stated as

$$\text{if } d \in \mathrm{ppre}_a^\forall(D) \text{ then } e \in \mathrm{ppre}_a^\forall(\mathcal{R}(D)) \tag{1}$$

Since $\mathrm{ppre}_a^\forall$ does not preserve set unions, the statement $d \in \mathrm{ppre}_a^\forall(D)$ is not equivalent to $\exists f \in D.\, d \in \mathrm{ppre}_a^\forall(f)$, so that the above condition (1) does not scale to the standard transfer property of (bi)simulations on LTSs that naturally admits a game characterization. It is therefore interesting to ask whether a suitable definition of an additive (i.e., union-preserving) probabilistic predecessor operator between distributions can be found. In the following, this question will be positively answered.

3.1 LTS on Distributions

Let us consider the following alternative definition of probabilistic predecessor operator:

$$\mathrm{ppre}_a : \wp(\mathrm{Distr}(\Sigma)) \to \wp(\mathrm{Distr}(\Sigma))$$
$$\mathrm{ppre}_a(D) \triangleq \{d \in \mathrm{Distr}(\Sigma) \mid \mathrm{supp}(d) \cap \mathrm{pre}_a(D) \neq \varnothing\}$$

This definition is much less restrictive than that of the above $\mathrm{ppre}_a^\forall$ operator: in order for a distribution d to be a probabilistic predecessor of a distribution e it is now sufficient that the support of d contains some state that reaches e. In this sense, ppre_a has an *existential* flavour as opposed to the *universal* flavour of $\mathrm{ppre}_a^\forall$. In the following, this observation will be also formalized by means of abstract interpretation [1,2].

Since the ppre_a operator actually preserves set unions, a corresponding transition relation between distributions can be defined as follows: $d \xrightarrow{a} e$ iff $d \in \mathrm{ppre}_a(\{e\})$, namely,

$$d \xrightarrow{a} e \text{ iff } \exists s \in \mathrm{supp}(d).\, s \xrightarrow{a} e \tag{$*$}$$

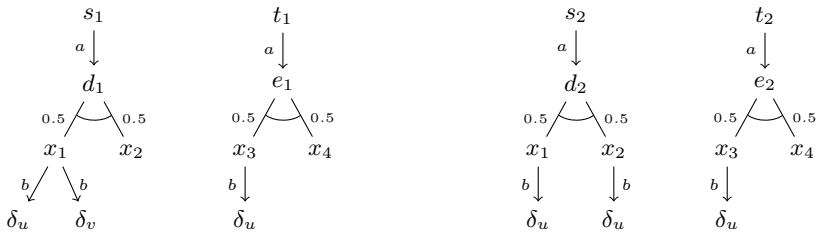

Fig. 1. A pair of PLTSs.

This allows us to lift a PLTS to an LTS of distributions, that we call DLTS. Hence, the following notions of simulation/bisimulation based on the standard transfer property naturally arise.

Definition 3.6 (d-simulation). Given a PLTS $\mathcal{M}$, a relation $\mathcal{R} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ is a *d-simulation* on $\mathcal{M}$ if for all $d, e \in \mathrm{Distr}(\Sigma)$, if $d\,\mathcal{R}\,e$ then:

(1) if $d \xrightarrow{a} f$ then there exists $g \in \mathrm{Distr}(\Sigma)$ such that $e \xrightarrow{a} g$ and $f\,\mathcal{R}\,g$;
(2) $d \sqsubseteq_{\Delta(\mathcal{R})} e$. □

Definition 3.7 (d-bisimulation). Given a PLTS $\mathcal{M}$, a symmetric relation $\mathcal{S} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ is a *d-bisimulation* on $\mathcal{M}$ if for all $d, e \in \mathrm{Distr}(\Sigma)$, if $d\,\mathcal{S}\,e$ then:

(1) if $d \xrightarrow{a} f$ then there exists $g \in \mathrm{Distr}(\Sigma)$ such that $e \xrightarrow{a} g$ and $f\,\mathcal{S}\,g$;
(2) $d \equiv_{\Delta(\mathcal{S})} e$. □

Given a PLTS $\mathcal{M}$, let $\mathcal{R}_{\mathrm{sim}} \triangleq \cup\{\mathcal{R} \mid \mathcal{R} \text{ is a d-simulation on } \mathcal{M}\}$ and $\mathcal{P}_{\mathrm{bis}} \triangleq \cup\{\mathcal{S} \mid \mathcal{S} \text{ is a d-bisimulation on } \mathcal{M}\}$. Then, $\mathcal{R}_{\mathrm{sim}}$ turns out to be the greatest d-simulation on $\mathcal{M}$ and a preorder, called the d-simulation preorder on $\mathcal{M}$. Likewise, it turns out that $\mathcal{P}_{\mathrm{bis}}$ is the greatest d-bisimulation on $\mathcal{M}$ and an equivalence, called d-bisimilarity on $\mathcal{M}$

Interestingly, d-simulations (and analogously for d-bisimulations) enjoy a neat correspondence with state simulations in a PLTS. More precisely, the state simulation preorder R_{sim} can be recovered from the d-simulation preorder $\mathcal{R}_{\mathrm{sim}}$ by restricting $\mathcal{R}_{\mathrm{sim}}$ to Dirac distributions. On the other hand, the d-simulation preorder coincides with the lifting to distributions of the state simulation preorder.

Theorem 3.8.

- $\Delta(\mathcal{R}_{\mathrm{sim}}) = R_{\mathrm{sim}}$ *and* $\mathcal{R}_{\mathrm{sim}} = \sqsubseteq_{R_{\mathrm{sim}}}$.
- $\Delta(\mathcal{P}_{\mathrm{bis}}) = P_{\mathrm{bis}}$ *and* $\mathcal{P}_{\mathrm{bis}} = \equiv_{P_{\mathrm{bis}}}$.

It is worth noting that this result opens the way to define new model checking tools that compute (bi)simulations on PLTSs by adapting to DLTSs the standard (bi)simulation techniques/algorithms designed in the nonprobabilistic framework.

Example 3.9. Consider the leftmost PLTS depicted in Figure 1. The relation $\mathcal{R}_1 = \{(\delta_{s_1}, \delta_{t_1}), (d_1, e_1)\} \cup \{(d, d) \mid d \in \mathrm{Distr}(\Sigma)\}$ is not a d-simulation since $d_1 \xrightarrow{b} \delta_v$ but $e_1 \xrightarrow{b} \delta_u$ and $\delta_u \notin \mathcal{R}_1(\delta_v)$. Moreover, even if $\mathcal{R}_1$ respects the transfer property of $\forall$d-simulations (since there is no set $D \subseteq \mathrm{Distr}(\Sigma)$ such that $\mathrm{supp}(d_1) \subseteq \mathrm{pre}_a(D)$), $\mathcal{R}_1$ is not even a $\forall$d-simulation because $d_1 \not\sqsubseteq_{\Delta(\mathcal{R}_1)} e_1$, since, for instance, $0.5 = d(\{x_2\}) \not\leq e(\Delta(\mathcal{R}_1)(\{x_2\})) = e(\{x_2\}) = 0$. Nevertheless, s_1 and t_1 are bisimilar states since there exists a ($\forall$)d-bisimulation containing the pair $(\delta_{s_1}, \delta_{t_1})$. Let $\mathcal{R}$ be the equivalence relation corresponding to the partition $\{\{\delta_{s_1}, \delta_{t_1}\}, \{d_1, e_1\}, \{\delta_{x_1}, \delta_{x_3}\}, \{\delta_{x_2}, \delta_{x_4}, \delta_u, \delta_v\}\}$. It is not difficult to check that $\mathcal{R}$ is a ($\forall$)d-bisimulation: every pair in $\mathcal{R}$ respects the transfer property and is $\equiv_{\Delta(\mathcal{R})}$-equivalent, where $\Delta(\mathcal{R}) = \{\{s_1, t_1\}, \{x_1, x_3\}, \{u, v, x_2, x_4\}\}$.

As a further example, consider the rightmost PLTS in Figure 1. We have that s_2 simulates t_2 but t_2 does not simulate s_2. In fact, consider the relation

$$\mathcal{R}_2 = \{(\delta_{t_2}, \delta_{s_2}), (e_2, d_2), (\delta_u, \delta_u)\} \cup \{(\delta_{x_4}, \delta_{x_i})\}_{i=1,\ldots,4} \cup \{(\delta_{x_3}, \delta_{x_i})\}_{i=1,2,3}.$$

Then, $\mathcal{R}_2$ is a ($\forall$)d-simulation since every pair respects the transfer property and belongs to $\sqsubseteq_{\Delta(\mathcal{R}_2)}$. For instance, let us check that $e_2 \sqsubseteq_{\Delta(\mathcal{R}_2)} d_2$: by Definition 2.1, it is enough to check that for all $U \subseteq \mathrm{supp}(e_2)$, $e_2(U) \leq d_2(\Delta(\mathcal{R}_2)(U))$. The nonempty subsets of $\mathrm{supp}(e_2)$ are: $U_1 = \{x_3\}$, $U_2 = \{x_4\}$ and $U_3 = \{x_3, x_4\}$, so that we have

$$\begin{aligned}
0.5 = e_2(\{x_3\}) &\leq d_2(\Delta(\mathcal{R}_2)(\{x_3\})) = d_2(\{x_1, x_2, x_3\}) = 1 \\
0.5 = e_2(\{x_4\}) &\leq d_2(\Delta(\mathcal{R}_2)(\{x_4\})) = d_2(\{x_1, x_2, x_3, x_4\}) = 1 \\
1 = e_2(\{x_3, x_4\}) &\leq d_2(\Delta(\mathcal{R}_2)(\{x_3, x_4\})) = d_2(\{x_1, x_2, x_3, x_4\}) = 1
\end{aligned}$$

The fact that t_2 does not simulate s_2 depends on the fact that e_2 does not simulate d_2 since this would imply that there exists a ($\forall$)d-simulation $\mathcal{R}$ such that $d_2 \sqsubseteq_{\Delta(\mathcal{R})} e_2$. However, the latter statement implies that $1 = d_2(\{x_1, x_2\}) \leq e_2(\Delta(\mathcal{R})(\{x_1, x_2\}))$, which is true only if $\mathrm{supp}(e_2) = \{x_3, x_4\} \subseteq \Delta(\mathcal{R})(\{x_1, x_2\})$; hence, in particular, we would obtain $\delta_{x_4} \in \mathcal{R}(\{\delta_{x_1}, \delta_{x_2}\})$, which is a contradiction since x_4 cannot simulate a b-transition. □

Besides the above notions of d-simulation/d-bisimulation, the operator ppre_a allows us to provide a corresponding new interpretation for the diamond connective. Let $\mathcal{L}$ denote the logic whose syntax coincides with $\mathcal{L}_\forall$ and whose semantics is identical to that of $\mathcal{L}_\forall$ but for the diamond connective, which is interpreted as follows:

$$d \models \Diamond_a \phi \text{ iff } \exists e.\ d \xrightarrow{a} e \text{ and } e \models \phi$$

This is therefore the *standard interpretation* of the diamond operator on a DLTS, namely a LTS whose "states" are distributions and whose transitions are defined by $(*)$. In the following, we will argue that $\mathcal{L}$ is best suited as probabilistic extension of Hennessy-Milner logic. As a first result, it turns out that the preorder $\leq_{\mathcal{L}^+}$ and the equivalence $\equiv_{\mathcal{L}}$ logically characterize, respectively, d-simulations and d-bisimulations.

Theorem 3.10. *For any $d, e \in \mathrm{Distr}(\Sigma)$,*

- $d\, \mathcal{R}_{\mathrm{sim}}\, e$ *if and only* $d \leq_{\mathcal{L}^+} e$*;*
- $d\, \mathcal{P}_{\mathrm{bis}}\, e$ *if and only* $d \equiv_{\mathcal{L}} e$*.*

3.2 Comparing $\mathcal{L}_\forall$ and $\mathcal{L}$

It turns out that $\forall$d-(bi)simulations and d-(bi)simulations are equivalent notions. In spite of the fact they rely on quite different transfer properties (cf. condition (1) of Definitions 3.3 and 3.6), their second defining condition, peculiar to the probabilistic setting, is powerful enough to bridge their gap. More precisely, this depends on the following key property: if $d \sqsubseteq_R e$ then for any state s in the support of d there exists a state t in the support of e such that $t \in R(s)$, and viceversa, for any state t in the support of e there exists a state s in the support of d such that $t \in R(s)$.

Lemma 3.11. *Consider a PLTS $\mathcal{M}$ and a relation $\mathcal{R} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$. Then, $\mathcal{R}$ is a $\forall$d-(bi)simulation on $\mathcal{M}$ iff $\mathcal{R}$ is a d-(bi)simulation on $\mathcal{M}$.*

As a consequence, we have that $\mathcal{R}^\forall_{\mathrm{sim}} = \mathcal{R}_{\mathrm{sim}}$ and $\mathcal{P}^\forall_{\mathrm{bis}} = \mathcal{P}_{\mathrm{bis}}$, so that, by Theorems 3.5 and 3.10, the two modal logics $\mathcal{L}_\forall$ and $\mathcal{L}$ induce the same equivalence on distributions while $\mathcal{L}^+_\forall$ and $\mathcal{L}^+$ induce the same preorder on distributions. As far as their relative expressive powers are concerned, we have that $\mathcal{L}_\forall$ and $\mathcal{L}$ are equivalent, while this is not the case for their negation-free fragments.

Theorem 3.12.

- *$\mathcal{L}_\forall$ and $\mathcal{L}$ have the same expressive power (and therefore $\equiv_{\mathcal{L}_\forall} = \equiv_{\mathcal{L}}$).*
- *$\mathcal{L}^+_\forall$ is strictly less expressive than $\mathcal{L}^+$, although $\leq_{\mathcal{L}^+_\forall} = \leq_{\mathcal{L}^+}$.*

Let us observe that the equivalence between $\mathcal{L}_\forall$ and $\mathcal{L}$ depends on the fact that the semantics of the diamond operator of $\mathcal{L}_\forall$ can be encoded in $\mathcal{L}$ and viceversa. In particular, the semantics of the $\mathcal{L}_\forall$-formula $\Diamond_a\phi$, i.e. $[\![\Diamond_a\phi]\!]_{\mathcal{L}_\forall} = \{d \mid \mathrm{supp}(d) \subseteq \mathrm{pre}_a(\{e \mid e \models_{\mathcal{L}_\forall} \phi\})\}$, can be expressed in $\mathcal{L}$ by the formula $[\Diamond_a\phi]_1$, whose semantics is indeed $[\![[\Diamond_a\phi]_1]\!]_{\mathcal{L}} = \{d \mid d(\{x \mid \delta_x \models_{\mathcal{L}} \Diamond_a\phi\}) = 1\} = \{d \mid \mathrm{supp}(d) \subseteq \{x \mid \delta_x \models_{\mathcal{L}} \Diamond_a\phi\}\}$. On the other hand, the encoding of $\mathcal{L}$'s diamond as a $\mathcal{L}_\forall$-formula is more tricky. The semantics of $\Diamond_a\phi$ viewed as a $\mathcal{L}$-formula is given by all the distributions whose support contains at least a state that moves to a distribution that satisfies ϕ, i.e., $[\![\Diamond_a\phi]\!]_{\mathcal{L}} = \{d \mid d(\{x \mid \exists e.\, x \xrightarrow{a} e,\ e \models_{\mathcal{L}} \phi\}) > 0\}$. This semantics can be therefore expressed in $\mathcal{L}_\forall$ by requiring that $d \models_{\mathcal{L}_\forall} [\Diamond_a\phi]_p$ for some $p > 0$. However, in general the existence of a rational number $p > 0$ can be expressed as a logical formula only by means of an infinite (countable) disjunction, hence it is expressible in the full $\mathcal{L}_\forall$ logic, but not in its negation-free and finitely disjunctive fragment $\mathcal{L}^+_\forall$.

Let us describe an example showing that the logic $\mathcal{L}^+_\forall$ is strictly less expressive than $\mathcal{L}^+$. Consider a PLTS $\mathcal{M} = \langle\{x_1, x_2\}, \{a\}, \{x_1 \xrightarrow{a} d = (x_1/0.5, x_2/0.5)\}\rangle$ that contains two states x_1, x_2 and a single transition from x_1 to the distribution $d = (x_1/0.5, x_2/0.5)$. In the logic $\mathcal{L}$, we have that $[\![\Diamond_a\top]\!]_{\mathcal{L}} = \mathrm{Distr}(\Sigma) \setminus \{\delta_{x_2}\}$, since any distribution different from δ_{x_2} contains x_1 in its support, and therefore has an outgoing a-transition. Let us show that there is no formula in $\mathcal{L}_\forall$ whose semantics is $\mathrm{Distr}(\Sigma) \setminus \{\delta_{x_2}\}$. Consider the $\mathcal{L}_\forall$-formulae $\top$, $\Diamond_a\top$ and $[\Diamond_a\top]_p$, with $p > 0$, whose semantics are as follows: $[\![\top]\!]_{\mathcal{L}_\forall} = \mathrm{Distr}(\Sigma)$, $[\![\Diamond_a\top]\!]_{\mathcal{L}_\forall} = \{\delta_{x_1}\}$, $[\![[\Diamond_a\top]_p]\!]_{\mathcal{L}_\forall} = \{d \mid d(\{x_1\}) \geq p\}$. It is easily seen that the semantics of any other formula in $\mathcal{L}_\forall$ is in the set $\mathit{Sem}_{\mathcal{L}_\forall} = \{[\![\top]\!]_{\mathcal{L}_\forall}, [\![\Diamond_a\top]\!]_{\mathcal{L}_\forall}\} \cup \{[\![[\Diamond_a\top]_p]\!]_{\mathcal{L}_\forall} \mid p > 0\}$, which is indeed closed under infinite intersections, *finite* unions, probabilistic predecessor and the semantics of

the operator $[\cdot]_p$. It is thus enough to observe that $\mathrm{Distr}(\Sigma) \setminus \{\delta_{x_2}\} \notin \mathit{Sem}_{\mathcal{L}_\forall}$: actually, $\mathrm{Distr}(\Sigma) \setminus \{\delta_{x_2}\}$ can only be expressed as the *infinite* union $\cup_{p>0}[\![[\Diamond_a \top]_p]\!]_{\mathcal{L}_\forall}$. □

Example 3.13. Consider again the rightmost PLTS in Figure 1. We have already observed that s_2 simulates t_2 whilst t_2 does not simulate s_2. The fact that t_2 does not simulate s_2 can be easily proved by exhibiting a formula that is satisfied by δ_{s_2} but not by δ_{t_2}. We provide both a formula in $\mathcal{L}_\forall$ and an equivalent formula in $\mathcal{L}$:

(1) let $\phi \triangleq \Diamond_a \Diamond_b \top \in \mathcal{L}_\forall$; then $\delta_{s_2} \models_{\mathcal{L}_\forall} \phi$ and $\delta_{t_2} \not\models_{\mathcal{L}_\forall} \phi$

(2) let $\phi' \triangleq \Diamond_a [\Diamond_b \top]_1 \in \mathcal{L}$; then $\delta_{s_2} \models_{\mathcal{L}} \phi'$ and $\delta_{t_2} \not\models_{\mathcal{L}} \phi'$

To see (1), observe that $\delta_{s_2} \models_{\mathcal{L}_\forall} \phi$ since $\mathrm{supp}(\delta_{s_2}) \subseteq \mathrm{pre}_a(d_2)$ and $\mathrm{supp}(d_2) \subseteq \mathrm{pre}_b(\delta_u)$ with $\delta_u \models_{\mathcal{L}_\forall} \top$. On the other hand, $\mathrm{supp}(\delta_{t_2}) \subseteq \mathrm{pre}_a(e_2)$ but $\mathrm{supp}(e_2) \not\subseteq \mathrm{pre}_b(f)$ for some distribution f such that $f \models_{\mathcal{L}_\forall} \top$. To show (2), notice that $\delta_{s_2} \models_{\mathcal{L}} \phi'$ since $\delta_{s_2} \xrightarrow{a} d_2$, and $d_2(\{x \mid \delta_x \models_{\mathcal{L}} \Diamond_b \top\}) = 1$ since for any $x \in \mathrm{supp}(d_2)$ it holds $\delta_x \xrightarrow{b} \delta_u$ with $\delta_u \models_{\mathcal{L}} \top$. On the other hand, $\delta_{t_2} \not\models_{\mathcal{L}} \phi'$ since $\delta_{t_2} \xrightarrow{a} e_2$ and $e_2(\{x \mid \delta_x \models_{\mathcal{L}} \Diamond_b \top\}) = 0.5 \not\geq 1$ since for $x_4 \in \mathrm{supp}(e_2)$ it holds $\delta_{x_4} \not\models_{\mathcal{L}} \Diamond_b \top$. □

4 States as Abstract Interpretation of Distributions

Differently from LTSs and their behavioral relations, whose definitions rely on a single notion of system state, PLTSs as well as their corresponding spectra of behavioral relations in some sense involve two notions of system state, namely a bare state and a probabilistic state modeled as a state distribution. We have shown above how PLTSs can be embedded into DLTSs, that is, LTSs of probabilistic states that involve a single (but richer) notion of system state, i.e. state distributions. We show in this section how to formalize a systematic embedding of states into distributions by viewing states as abstract interpretation of distributions.

Intuitively, Dirac distributions allow us to view states as an abstraction of distributions, namely the map $\delta : \Sigma \to \mathrm{Distr}(\Sigma)$ such that $\delta(x) \triangleq \delta_x$ may be viewed as a function that embeds states into distributions. The other way round, the support map $\mathrm{supp} : \mathrm{Distr}(\Sigma) \to \wp(\Sigma)$ can be viewed as a function that abstracts a distribution d as the set of states in its support.

Let us recall that in standard abstract interpretation [1,2], approximations of a concrete semantic domain are encoded by abstract domains that are specified by Galois insertions (GIs for short) or, equivalently, by adjunctions. Approximation on a concrete/abstract domain is encoded by a partial order where traditionally $x \leq y$ means that y is a concrete/abstract approximation of x. Concrete and abstract approximation orders, denoted by $\leq_C$ and $\leq_A$, must be related by a GI. Recall that a GI of an abstract domain $\langle A, \leq_A \rangle$ into a concrete domain $\langle C, \leq_C \rangle$ is determined by a surjective abstraction map $\alpha : C \to A$ and a 1-1 concretization map $\gamma : A \to C$ such that $\alpha(c) \leq_A a \Leftrightarrow c \leq_C \gamma(a)$ and is denoted by (α, C, A, γ). In a GI, intuitively $\alpha(c)$ provides the best approximation in A of a concrete value c while $\gamma(a)$ is the concrete value that a abstractly represents.

In our case, in order to cast δ as a concretization map in abstract interpretation, we need to lift its definition from sets of states to sets of distributions, namely we

need to provide its so-called "collecting" version [1,2]. Observe that $\{\delta(x)\} = \{d \in \mathrm{Distr}(\Sigma) \mid \mathrm{supp}(d) \subseteq \{x\}\}$. This leads us to define the following concretization function $\gamma^{\forall} : \wp(\Sigma) \to \wp(\mathrm{Distr}(\Sigma))$:

$$\gamma^{\forall}(S) \triangleq \{d \in \mathrm{Distr}(\Sigma) \mid \mathrm{supp}(d) \subseteq S\}.$$

This is a *universal* concretization function, meaning that $d \in \gamma^{\forall}(S)$ iff all the states in $\mathrm{supp}(d)$ are contained into S. Hence, one can dually define an *existential* concretization map $\gamma^{\exists} : \wp(\Sigma) \to \wp(\mathrm{Distr}(\Sigma))$ as

$$\gamma^{\exists}(S) \triangleq \{d \in \mathrm{Distr}(\Sigma) \mid \mathrm{supp}(d) \cap S \neq \varnothing\},$$

where $d \in \gamma^{\exists}(S)$ if there exists some state in the support of d which is contained into S. Actually, these two mappings give rise to a pair of GIs (i.e., approximations in abstract interpretation) where $\wp(\mathrm{Distr}(\Sigma))$ and $\wp(\Sigma)$ play, respectively, the role of concrete and abstract domains. The approximation order is encoded by the subset relation (i.e., logical implication) in the case of $\gamma^{\forall}$ and by the superset relation (i.e., logical co-implication) in the case of $\gamma^{\exists}$. The dual maps, systematically obtained by adjunction from $\gamma^{\forall}$ and $\gamma^{\exists}$, are $\alpha^{\forall}, \alpha^{\exists} : \wp(\mathrm{Distr}(\Sigma)) \to \wp(\Sigma)$ defined as follows:

$$\begin{aligned}
\alpha^{\forall}(X) &\triangleq \{s \in \Sigma \mid \exists d \in X.\ s \in \mathrm{supp}(d)\} \\
\alpha^{\exists}(X) &\triangleq \{s \in \Sigma \mid \forall d \in \mathrm{Distr}(\Sigma).\ s \in \mathrm{supp}(d) \Rightarrow d \in X\}
\end{aligned}$$

Lemma 4.1. $(\alpha^{\forall}, \wp(\mathrm{Distr}(\Sigma))_{\subseteq}, \wp(\Sigma))_{\subseteq}, \gamma^{\forall})$ *and* $(\alpha^{\exists}, \wp(\mathrm{Distr}(\Sigma))_{\supseteq}, \wp(\Sigma))_{\supseteq}, \gamma^{\exists})$ *are GIs.*

Observe that $\alpha^{\forall}/\gamma^{\forall}$ and $\alpha^{\exists}/\gamma^{\exists}$ are dual abstractions, i.e.,

$$\alpha^{\exists} = \neg\, \alpha^{\forall} \neg \quad \text{and} \quad \gamma^{\exists} = \neg\, \gamma^{\forall} \neg$$

where $\neg\, \alpha^{\forall} \neg(X) = \Sigma \setminus \alpha^{\forall}(\mathrm{Distr}(\Sigma) \setminus X)$ and $\neg\, \gamma^{\forall} \neg(S) = \mathrm{Distr}(\Sigma) \setminus \gamma^{\forall}(\Sigma \setminus S)$. Moreover, it is not hard to see that $\alpha^{\forall}$ is the additive extension of the supp function, while $\alpha^{\exists}$ is its co-additive extension, i.e.,

$$\alpha^{\forall}(X) = \cup_{d \in X} \mathrm{supp}(d) \quad \text{and} \quad \alpha^{\exists}(\mathrm{Distr} \setminus X) = \cap_{d \in X} \Sigma \setminus \mathrm{supp}(d).$$

These two abstract domains thus provide dual universal/existential ways for logically approximating sets of distributions into sets of states. The interesting point in these formal abstractions lies in the fact that they allow us to systematically obtain the above probabilistic predecessor operators $\mathrm{ppre}_a^{\forall}$ and ppre_a in a DLTS from the predecessor operator pre_a of the corresponding PLTS. Recall that in a PLTS the predecessor operator $\mathrm{pre}_a : \wp(\mathrm{Distr}(\Sigma)) \to \wp(\Sigma)$ maps a set of distributions into a set of states. Here, $\wp(\Sigma)$ can therefore be viewed as a universal/existential abstraction of $\wp(\mathrm{Distr}(\Sigma))$, so that, correspondingly, pre_a can be viewed as an *abstract predecessor function*, since its co-domain actually is an abstract domain. Consequently, the output of this abstract function can be projected back to distributions using the corresponding concretization map. Interestingly, it turns out that the corresponding *concrete predecessor functions*, obtained by composing the operator pre_a with either $\gamma^{\forall}$ or $\gamma^{\exists}$, exactly coincide with the two probabilistic predecessors $\mathrm{ppre}_a^{\forall}$ and ppre_a.

Lemma 4.2. $\mathrm{ppre}_a^{\forall} = \gamma^{\forall} \circ \mathrm{pre}_a$ *and* $\mathrm{ppre}_a = \gamma^{\exists} \circ \mathrm{pre}_a$.

Thus, in equivalent terms, the predecessor operator pre_a is the best correct universal/existential approximation of the operators $\mathrm{ppre}_a^{\forall}/\mathrm{ppre}_a$, for the universal/existential abstractions $\alpha^{\forall}/\gamma^{\forall}$ and $\alpha^{\exists}/\gamma^{\exists}$.

5 A Spectrum of Probabilistic Relations over DLTSs

The approach developed above suggests a general methodology for defining behavioral relations between the states of a PLTS: first define a "lifted" behavioral relation between distributions of the corresponding DLTS and then restrict this definition to Dirac distributions. As discussed above, this approach works satisfactorily for simulation and bisimulation on PLTSs. In what follows, we show that this technique is indeed more general since it can be applied to a number of known probabilistic behavioral relations.

5.1 Probabilistic Simulation

Segala and Lynch [15] put forward a variant of simulation where a state transition $s \xrightarrow{a} d$ can be matched by a so-called combined transition from a state t, namely a convex combination of distributions reachable from t. We show that this same idea can be lifted to transitions in DLTSs.

Let $\mathcal{M} = \langle \Sigma, Act, \rightarrow \rangle$ be a PLTS, let $\{s \xrightarrow{a} d_i\}_{i \in I}$ be a (denumerable) family of transitions of $\mathcal{M}$ and let $\{p_i\}_{i \in I}$ be a corresponding family of probabilities in $[0, 1]$ such that $\sum_{i \in I} p_i = 1$. Let $d \in \mathrm{Distr}(\Sigma)$ be the convex combination $d = \sum_{i \in I} p_i d_i$. Then, $\{s, a, \sum_{i \in I} p_i d_i\}$, denoted by $s \overset{a}{\rightsquigarrow} d$, is called a *combined transition* in $\mathcal{M}$. This notion of combined transition can be lifted to distributions as follows.

Definition 5.1 (Combined d-transitions and Hyper transitions).

- Let $d, e \in \mathrm{Distr}(\Sigma)$. Then, $d \overset{a}{\rightsquigarrow} e$ if there exists $s \in \mathrm{supp}(d)$ such that $s \overset{a}{\rightsquigarrow} e$. $d \overset{a}{\rightsquigarrow} e$ is called a *combined d-transition*.
- Let $\{d \xrightarrow{a} d_i\}_{i \in I}$ be a family of transitions in a DLTS, and let $\{p_i\}_{i \in I}$ be a corresponding family of probabilities such that $\sum_{i \in I} p_i = 1$. Let $d = \sum_{i \in I} p_i d_i$. Then, the triple $\{d, a, \sum_{i \in I} p_i d_i\}$, compactly denoted by $d \overset{a}{\rightrightarrows} e$, is called a *hyper transition*. □

It is worth noting that the notion of hyper transition is "stronger" than that of combined d-transition, in that $d \overset{a}{\rightsquigarrow} e$ implies $d \overset{a}{\rightrightarrows} e$ but not viceversa. Moreover, our definition of hyper transition can be compared with analogous notions of hyper transition defined in [16] and [5]. In particular, it can be shown that a hyper transition in the sense of both Stoelinga [16] and Deng et al. [5] is a hyper transition in our sense, but not vice versa. Anyhow, the notion of combined d-transition is sufficient to lift probabilistic (bi)simulations of [15] to distributions.

In what follows, we focus on simulation relations only, since the same results scale to bisimulations. A probabilistic simulation is defined as a simulation in a PLTS apart

from using combined transitions rather than standard transitions of a PLTS. Correspondingly, probabilistic d-simulation is defined as in Definition 3.6, but using combined d-transitions rather than transitions in a DLTS.

Let R_{psim} ($\mathcal{R}_{\mathrm{psim}}$) be the union of all the probabilistic (d-)simulations on $\mathcal{M}$. Then, all the results obtained in Section 3 also hold for probabilistic simulation, and they are collected in the following theorem. In particular, as before, the probabilistic simulation preorder between states can be recovered from the probabilistic d-simulation preorder by restricting to Dirac distributions. Dually, the probabilistic d-simulation preorder coincides with the lifting of the probabilistic simulation preorder. As far as the logic is concerned, Parma and Segala [13] show that the probabilistic relations between the states of a PLTS are logically characterized by the logical equivalence/preorder — restricted to Dirac distributions — of a modal logic that has the same syntax of $\mathcal{L}_\forall$ but whose diamond operator is defined in terms of combined transitions on the PLTS. Let $\mathcal{L}_P$ denote the logic $\mathcal{L}$ (which is equivalent to $\mathcal{L}_\forall$) where the semantics of the diamond operator is defined in terms of combined d-transitions. Then, as in Section 3, the result in [13] can be extended by showing that the full logical preorder of $\mathcal{L}_P^+$ coincides with the probabilistic d-simulation preorder.

Theorem 5.2.

- $\Delta(\mathcal{R}_{\mathrm{psim}}) = R_{\mathrm{psim}}$ *and* $\sqsubseteq_{R_{\mathrm{psim}}} = \mathcal{R}_{\mathrm{psim}}$.
- $\mathcal{R}_{\mathrm{psim}} = \leq_{\mathcal{L}_P^+}$.

5.2 Failure Simulation

One nice consequence of defining DLTSs as LTSs of distributions lies in the fact that the standard van Glabbeek's spectrum [8] of behavioral relations on LTSs can be reformulated in terms of transitions between distributions of a DLTS. This leads to a spectrum of d-relations between distributions of a DLTS, that can be projected into a spectrum of relations between states of a PLTS by restricting the d-relations to Dirac distributions. As an example we show how this approach works on failure simulation [8]. A formalization and generalization of such a "lifting schema" in a suitable framework like abstract interpretation or coalgebras is left as future work.

Definition 5.3 (Failure Simulation). A relation $R \subseteq \Sigma \times \Sigma$ is a *failure simulation* on a PLTS when for any $s, t \in \Sigma$, if sRt then:

- if $s \xrightarrow{a} d$ then there exists $e \in \mathrm{Distr}(\Sigma)$ such that $t \xrightarrow{a} e$ and $d \sqsubseteq_R e$;
- if $s \xrightarrow{A}$ then $t \xrightarrow{A}$ for any $A \subseteq Act$. □

Definition 5.4 (Failure d-Simulation). A relation $\mathcal{R} \subseteq \mathrm{Distr}(\Sigma) \times \mathrm{Distr}(\Sigma)$ is a *failure d-simulation* on a PLTS when for all $d, e \in \mathrm{Distr}(\Sigma)$, if $d \,\mathcal{R}\, e$ then:

(1) if $d \xrightarrow{a} f$ then there exists $g \in \mathrm{Distr}(\Sigma)$ such that $e \xrightarrow{a} g$ and $f \,\mathcal{R}\, g$;
(2) if $d \xrightarrow{A}$ then $e \xrightarrow{A}$ for any $A \subseteq Act$;
(3) $d \sqsubseteq_{\Delta(\mathcal{R})} e$. □

The lifting of a relation between states of a PLTS to a relation between distributions of the corresponding DLTS is obtained by resorting to the standard transfer property and by adding the condition (i.e., condition (3) in Definition 5.4) that deals with probabilities assigned to sets of related states. Let R_{fail} and $\mathcal{R}_{\mathrm{fail}}$ be, respectively, the failure simulation and d-simulation preorders on a PLTS $\mathcal{M}$. According to the LTS spectrum, failure simulation can be logically characterized through a modality that characterizes which transitions cannot be fired. We follow this same approch and we denote by $\mathcal{L}_F^+$ the logic obtained from $\mathcal{L}^+$ by adding a modality $\mathsf{ref}\langle A\rangle$, where $A \subseteq Act$, and whose semantics is defined as follows: for any $d \in \mathrm{Distr}(\Sigma)$, $d \models_{\mathcal{L}_F^+} \mathsf{ref}\langle A\rangle$ iff $d \stackrel{A}{\not\rightarrow}$.

Theorem 5.5.

- $\Delta(\mathcal{R}_{\mathrm{fail}}) = R_{\mathrm{fail}}$ *and* $\sqsubseteq_{R_{\mathrm{fail}}} = \mathcal{R}_{\mathrm{fail}}$*;*
- $\mathcal{R}_{\mathrm{fail}} = \leq_{\mathcal{L}_F^+}$.

6 Related and Future work

Simulation and bisimulation relations on PLTSs have been introduced by Segala and Lynch [15] as two equivalences that preserve significant classes of temporal properties in the probabilistic logic PCTL [10]. Since then a number of works put forward probabilistic extensions of Hennessy-Milner logic in order to logically characterize these equivalences. Larsen and Skou [12] and Desharnais et al. [7] investigated a probabilistic diamond operator that enhances the diamond operator of HML with the probability bounds of transitions. However, these logics are adequate just for reactive and alternating systems, which are probabilistic models that are strictly less expressive than PLTSs. Two further probabilistic variants of HML are available [13,5]. The first one is that of Parma and Segala [13] (see also [11]), whose formulae are interpreted on sets of probability distributions over the states of a PLTS. One distinctive operator of this logic is a modal operator $[\phi]_p$, whose semantics is the set of distributions that assigns at least probability p to the set of states whose Dirac distributions satisfy ϕ. This paper has shown that such a logic admits an equivalent formulation that retains the probabilistic operator $[\phi]_p$ and retrieves the diamond operator of HML by lifting it to distributions. Deng et al. [5] follow a different approach. They propose a probabilistic variant of HML that is interpreted on sets of processes of the pCSP process calculus. In their logic the semantics of the diamond operator is defined in terms of hyper transitions between distributions: this notion of hyper transition is more complex than ours and has been compared with our notion of hyper transitions in Section 5. Moreover, Deng et al.'s logic features a probabilistic operator $\bigoplus_{i\in I} p_i\phi_i$ that is satisfied by processes that correspond to distributions that can be decomposed into convex combinations of distributions that satisfy ϕ_i. Besides (bi)simulation and probabilistic (bi)simulation, this logic is able to characterize two notions of failure and forward simulation that have been proved to agree with the testing preorders on pCSP processes (see [5]).

Deng et al. [5]'s definition of failure simulation is quite different from ours, that we directly derived from the standard LTS spectrum. One major difference is that we define a relation between states of a PLTS which is then lifted to a relation between distributions, whereas Deng et al. consider a relation between states and distributions.

A precise comparison between the spectrum of behavioral relations on DLTSs and the behavioral relations defined by Deng et al. [5] is left as subject for future work. We also plan to investigate weak transitions in DLTSs that abstract from internal, invisible, actions. Weak variants of simulation, probabilistic simulation, forward and failure simulation have been studied both in [5] and [13].

As a further avenue of future work we plan to study whether and how behavioral relations on PLTSs can be computed by resorting to standard algorithms for LTSs that compute the corresponding lifted relations on a DLTS. A first step in this direction has been taken in [3], where efficient algorithms to compute simulation and bisimulation on PLTSs have been derived by resorting to abstract interpretation techniques.

Acknowledgements. This work was partially supported by University of Padova under the projects "AVIAMO" and "BECOM".

References

1. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proc. 4th ACM POPL*, 1977.
2. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Proc. 6th ACM POPL*, pp. 269–282, 1979.
3. S. Crafa and F. Ranzato Probabilistic bisimulation and simulation algorithms by abstract interpretation. In *Proc. ICALP'11*, LNCS 6756, 2011.
4. Y. Deng and R. van Glabbeek. Characterising probabilistic processes logically. In *Proc. LPAR'10*, LNCS 6397, pp. 278-293, 2010.
5. Y. Deng, R. van Glabbeek, M. Hennessy and C. Morgan. Characterising testing preorders for finite probabilistic processes. *Logical Methods in Computer Science*, 4(4), 2008.
6. J. Desharnais. *Labelled Markov Processes*. PhD thesis, McGill Univ., 1999.
7. J. Desharnais, V. Gupta, R. Jagadeesan and P. Panangaden. Weak bisimulation is sound and complete for PCTL. In *Proc. CONCUR'02*, LNCS 2421, pp. 355-370, 2002.
8. R.J. van Glabbeek. The linear time – branching time spectrum I; the semantics of concrete, sequential processes. In *Handbook of Process Algebra*, chapter 1, pp. 3-99, Elsevier, 2001.
9. R.J. van Glabbeek, S. Smolka, B. Steffen and C. Tofts. Reactive, generative and stratified models for probabilistic processes. In *Proc. IEEE LICS'90*, pp. 130-141, 1990.
10. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512-535, 1994.
11. H. Hermanns, A. Parma, R. Segala, B. Wachter and L. Zhang. Probabilistic logical characterization. *Information and Computation*, 209(2):154-172, 2011.
12. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1-28, 1991.
13. A. Parma and R. Segala. Logical characterizations of bisimulations for discrete probabilistic systems. In *Proc. FOSSACS'07*, LNCS 4423, p. 287-301, 2007.
14. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
15. R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *Nordic J. of Computing*, 2(2):250-273, 1995.
16. M. Stoelinga. *Alea Jacta Est: Verification of Probabilistic, Real-Time and Parametric Systems*. PhD thesis, University of Nijmegen, The Netherlands, 2002.
17. L. Zhang, H. Hermanns, F. Eisenbrand and D.N. Jansen. Flow faster: efficient decision algorithms for probabilistic simulations. *Logical Methods in Computer Science*, 4(4), 2008.