Noname manuscript No.
(will be inserted by the editor)

Bisimulation and Simulation Algorithms on Probabilistic Transition Systems by Abstract Interpretation

Silvia Crafa · Francesco Ranzato

Abstract We show how bisimulation equivalence and simulation preorder on probabilistic LTSs (PLTSs), namely the main behavioural relations on probabilistic nondeterministic processes, can be characterized by abstract interpretation. Both bisimulation and simulation can be obtained as completions of partitions and preorders, viewed as abstract domains, w.r.t. a pair of concrete functions that encode a PLTS. This approach provides a general framework for designing algorithms that compute bisimulation and simulation on PLTSs. Notably, (i) we show that the standard bisimulation algorithm by Baier et al. [2000] can be viewed as an instance of such a framework and (ii) we design a new efficient simulation algorithm that improves the state of the art.

1 Introduction

Randomization phenomena in concurrent systems have been widely studied in probabilistic extensions of process algebras like Markov chains and probabilistic labeled transition systems (PLTSs). Most standard tools for studying nonprobabilistic processes, like behavioural equivalences, temporal logics and model checking, have been investigated for these probabilistic models. In particular, bisimulation equivalence and simulation preorder relations, namely the main behavioural relations between concurrent systems, have been extended and studied in a probabilistic setting [6,9,11,17].

Abstract interpretation [2,3] is a well-known general theory for specifying the approximation of formal semantics. Abstract domains play an essential role in any abstract interpretation design, since they encode in an ordered structure how concrete semantic properties are approximated. A number of behavioural relations, including bisimulation, stuttering bisimulation and simulation, have been characterized in abstract interpretation as complete refinements, so-called forward complete shells, of abstract domains w.r.t. logical/temporal operators of suitable modal logics [15]. One notable benefit of this approach is that it provides a general framework for designing basic algorithms that compute behavioral relations as forward complete shells of abstract domains. As a remarkable example, this abstract interpretation-based approach led to an efficient algorithm for computing the simulation preorder for nonprobabilistic processes [14,16] that features the best time complexity among the simulation algorithms.

Dipartimento di Matematica, University of Padova, Italy. E-mail: {crafa,ranzato}@math.unipd.it

The goal of this paper is to investigate whether the abstract interpretation approach can be applied to probabilistic LTSs in order (i) to characterize bisimulation equivalence and simulation preorder as logical completions of abstract domains and (ii) to design bisimulation and simulation algorithms.

Main Results. We consider probabilistic processes specified as PLTSs, a general model that exhibits both non-deterministic choice (as in LTSs) and probabilistic choice (as in Markov chains). In [15], bisimulation in LTSs has been characterized in terms of forward complete shells of partitions w.r.t. the predecessor operator of LTSs. We show that this same idea scales to the case of PLTSs by considering the probabilistic predecessor operator that defines the transitions of a PLTS together with a probabilistic function that encodes the distributions in the PLTS (this latter operator is somehow reminiscent of a probabilistic connective in Parma and Segala's [13] modal logics for probabilistic bisimulation and simulation). Bisimulation equivalence in PLTSs is thus characterized as a domain refinement through a complete shell w.r.t. the above two operators. On the other hand, the simulation preorder in PLTSs turns out to be the same complete shell of abstract domains w.r.t. the same two operators, but using different underlying abstract domains: for bisimulation, the complete shell is computed in a space of abstractions that are state and distribution partitions, while for simulation the same complete shell is instead computed over abstractions that are preorders on states and distributions.

Complete shells of abstract domains may in general be obtained through a simple fixpoint computation. We show how such a basic procedure can be instantiated to obtain two algorithms that iteratively compute bisimulation and simulation on PLTSs. Interestingly, the standard procedure for computing bisimulations in PLTSs, namely Baier-Engelen-Majster's algorithm [1], can be actually viewed as an implementation of our complete shell procedure that characterizes bisimulation. On the other hand, we show that the corresponding complete shell for computing the simulation preorder yields a new efficient probabilistic simulation algorithm that advances the state-of-the-art: in fact, its time and space complexity bounds improve on the best known simulation algorithm for PLTSs by Zhang et al. [18].

This is an extended and revised version of the conference paper [4] that includes full proofs.

2 Bisimulation and Simulation in PLTSs

Given a set X, $\mathrm{Distr}(X)$ denotes the set of (stochastic) distributions on X, i.e., functions $d{:}X \to [0,1]$ such that $\sum_{x\in X} d(x) = 1$. The support of a distribution d is defined by $\mathrm{supp}(d) \triangleq \{x \in X \mid d(x) > 0\}$; also, if $S \subseteq X$ then $d(S) \triangleq \sum_{s\in S} d(s)$. 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 $\mathcal{S} = \langle \Sigma, \mathit{Act}, \to\rangle$ where Σ is a set of states, Act is a set of actions and $\to \subseteq \Sigma \times \mathit{Act} \times \mathrm{Distr}(\Sigma)$ is a transition relation, where $(s, a, d) \in \to$ is also denoted by $s \xrightarrow{a} d$. We denote by $\mathrm{Distr} \triangleq \{d \in \mathrm{Distr}(\Sigma) \mid \exists s \in \Sigma. \exists a \in \mathit{Act} \,.\, s \xrightarrow{a} d\}$ the set of target distributions in $\mathcal{S}$. Given $D \subseteq \mathrm{Distr}$, we write $s \xrightarrow{a} D$ when there exists $d \in D$ such that $s \xrightarrow{a} d$. For any $a \in \mathit{Act}$, the predecessor and successor operators $\mathrm{pre}_a : \wp(\mathrm{Distr}) \to \wp(\Sigma)$ and $\mathrm{post}_a : \wp(\Sigma) \to \wp(\mathrm{Distr})$ are defined as usual by $\mathrm{pre}_a(D) \triangleq \{s \in \Sigma \mid s \xrightarrow{a} D\}$ and $\mathrm{post}_a(S) \triangleq \{d \in \mathrm{Distr} \mid \exists s \in S.\, s \xrightarrow{a} d\}$. For any $d \in \mathrm{Distr}$ and $s \in \Sigma$, we define $\mathrm{in}(d) \triangleq \{a \in \mathit{Act} \mid \mathrm{pre}_a(d) \neq \varnothing\}$ and $\mathrm{out}(s) \triangleq \{a \in \mathit{Act} \mid \mathrm{post}_a(s) \neq \varnothing\}$.

Bisimulation. A partition of a set X is a set $P \subseteq \wp(X)$ of nonempty subsets of X (called blocks) that are pairwise disjoints and whose union gives X. Let $\mathrm{Part}(X)$ denote the set of partitions of X. If $P \in \mathrm{Part}(X)$ and $x \in X$ then $P(x)$ denotes the unique block of P that contains x. A partition P will be also viewed as a mapping $P : \wp(X) \to \wp(X)$ defined by $P(Y) \triangleq \cup_{y \in Y} P(y)$. Any partition $P \in \mathrm{Part}(X)$ induces an equivalence relation (which can be equivalently given as a partition) over distributions $\equiv_P \in \mathrm{Part}(\mathrm{Distr}(X))$ which is defined as follows: for any $d, e \in \mathrm{Distr}(X)$, $d \equiv_P e$ if for any $B \in P$, $d(B) = e(B)$. In words, two distributions are $\equiv_P$-equivalent whenever they give the same probability to the blocks of P.

Given a PLTS $\mathcal{S} = \langle \Sigma, Act, \to \rangle$, a partition $P \in \mathrm{Part}(\Sigma)$ is a bisimulation on $\mathcal{S}$ when for all $s, t \in \Sigma$ and $d \in \mathrm{Distr}$, if $P(s) = P(t)$ and $s \xrightarrow{a} d$ then there exists $e \in \mathrm{Distr}$ such that $t \xrightarrow{a} e$ and $d \equiv_P e$. Bisimilarity $P_{\mathrm{bis}} \in \mathrm{Part}(\Sigma)$ is defined as follows: for any $s \in \Sigma$, $P_{\mathrm{bis}}(s) \triangleq \cup\{P(s) \mid P \text{ is a bisimulation on } \mathcal{S}\}$. P_{bis} turns out to be the greatest bisimulation on $\mathcal{S}$ which is also called the bisimulation partition on $\mathcal{S}$.

Simulation. A preorder on a set X is a reflexive and transitive relation $R \subseteq X \times X$. Let $\mathrm{PreOrd}(X)$ denote the set of preorders on X. If $R \in \mathrm{PreOrd}(X)$ and $S \subseteq X$ then $R(S) \triangleq \{x \in X \mid \exists s \in S. (s, x) \in R\}$ denotes the image of S for R. Similarly to the case of partitions, any preorder $R \in \mathrm{PreOrd}(X)$ induces a preorder $\leq_R$ on $\mathrm{Distr}(X)$ which is defined as follows: for any $d, e \in \mathrm{Distr}(X)$, $d \leq_R e$ if for any $S \subseteq X$, $d(S) \leq e(R(S))$. Such a definition of $\leq_R$ can be equivalently stated in terms of so-called weight functions between distributions and of maximum flows between networks. We briefly recall its equivalent formulation based on the maximum flow problem since our simulation algorithm, as well as the simulation algorithms by Baier et al. [1] and Zhang et al. [18], are based on this notion.

Let $\overline{X} \triangleq \{\overline{x} \mid x \in X\}$, where $\overline{x}$ are pairwise distinct new elements; $\bot$ (the source) and $\top$ (the sink) are a pair of new distinct elements not contained in $X \cup \overline{X}$. Let $R \subseteq X \times X$ and $d, e \in \mathrm{Distr}(X)$. The network $\mathcal{N}(d, e, R)$ is defined as follows: the set of nodes is $V \triangleq \mathrm{supp}(d) \cup \overline{\mathrm{supp}(e)} \cup \{\bot, \top\}$ while the set of edges $E \subseteq V \times V$ is defined by

$$E \triangleq \{(x, \overline{y}) \mid (x, y) \in R\} \cup \{(\bot, x) \mid x \in \mathrm{supp}(d)\} \cup \{(\overline{y}, \bot) \mid y \in \mathrm{supp}(e)\}.$$

The capacity function $c : V \times V \to [0, 1]$ is defined as follows: for all $x \in \mathrm{supp}(d)$, $c(\bot, x) \triangleq d(x)$; for all $y \in \mathrm{supp}(e)$, $c(\overline{y}, \top) \triangleq e(y)$; for all the remaining edges $(x, \overline{y}) \in E$, $c(x, \overline{y}) \triangleq 1$. It turns out that $d \leq_R e$ if and only if the maximum flow of the network $\mathcal{N}(d, e, R)$ is 1 (see [1,5,18,20]).

Given a PLTS $\mathcal{S} = \langle \Sigma, Act, \to \rangle$, a preorder $R \in \mathrm{PreOrd}(\Sigma)$ is a simulation on $\mathcal{S}$ when for all $s, t \in \Sigma$ and $d \in \mathrm{Distr}$, if $t \in R(s)$ and $s \xrightarrow{a} d$ then there exists $e \in \mathrm{Distr}$ such that $t \xrightarrow{a} e$ and $d \leq_R e$. The simulation preorder $R_{\mathrm{sim}} \in \mathrm{PreOrd}(\Sigma)$ on $\mathcal{S}$ is defined as follows: for all $s \in \Sigma$, $R_{\mathrm{sim}}(s) \triangleq \cup\{R(s) \mid R \text{ is a simulation on } \mathcal{S}\}$. It turns out that R_{sim} is the greatest simulation preorder on $\mathcal{S}$. The simulation partition P_{psim} on $\mathcal{S}$ is the kernel of the simulation preorder, i.e., for all $s, t \in \Sigma$, $P_{\mathrm{psim}}(s) = P_{\mathrm{psim}}(t)$ iff $s \in R_{\mathrm{sim}}(t)$ and $t \in R_{\mathrm{sim}}(s)$.

Example 2.1 Consider the PLTS depicted in Figure 1, where $\Sigma = \{s_1, s_2, s_3, x_1, \ldots, x_6, t, u, v, w\}$, $Act = \{a, b, c, d\}$ and $\mathrm{Distr} = \{d_1, d_2, d_3, \delta_t, \delta_u, \delta_v, \delta_w\}$. We have that s_2 simulates s_1 while s_1 does not simulate s_2 since starting from s_2 a d-transition can be fired, whereas starting from s_1 this is not possible. Moreover, even if x_3 simulates both x_5 and x_6, we have that $d_3 \not\leq_{R_{\mathrm{psim}}} d_2$ since $1 = d_3(\{x_5, x_6\}) \not\leq d_2(R_{\mathrm{psim}}(\{x_5, x_6\})) = 0.5$ since $x_4 \notin R_{\mathrm{psim}}(\{x_5, x_6\})$. Therefore, s_2 does not simulate s_3. $\square$

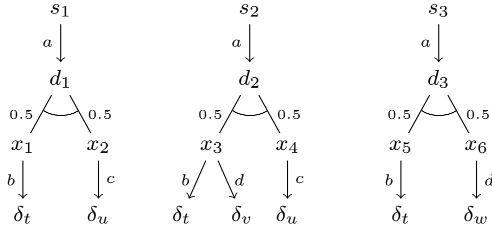

Fig. 1 A PLTS.

3 Shells

3.1 Forward Completeness

In standard abstract interpretation [2,3], approximations of a concrete semantic domain are encoded by abstract domains (or abstractions), that are specified by Galois insertions (GIs for short) or, equivalently, by adjunctions. Concrete and abstract domains are defined as complete lattices $\langle C, \leq_C\rangle$ and $\langle A, \leq_A\rangle$ where $x \leq y$ means that y approximates x both concretely and abstractly. A GI of A into C 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). Recall that GIs of a common concrete domain C are preordered w.r.t. their relative precision: $\mathcal{G}_1 = (\alpha_1, C, A_1, \gamma_1) \trianglelefteq \mathcal{G}_2 = (\alpha_2, C, A_2, \gamma_2)$, i.e. A_1/A_2 is a refinement/simplification of A_2/A_1, if and only if for all $c \in C$, $\gamma_1(\alpha_1(c)) \leq_C \gamma_2(\alpha_2(c))$. Moreover, $\mathcal{G}_1$ and $\mathcal{G}_2$ are equivalent when $\mathcal{G}_1 \trianglelefteq \mathcal{G}_2$ and $\mathcal{G}_2 \trianglelefteq \mathcal{G}_1$. We denote by $\mathrm{Abs}(C)$ the family of abstract domains of C up to the above equivalence. It is well known that $\langle \mathrm{Abs}(C), \trianglelefteq\rangle$ is a complete lattice. Given a family of abstract domains $\mathcal{X} \subseteq \mathrm{Abs}(C)$, their lub $\sqcup\mathcal{X}$ is therefore the most precise domain in $\mathrm{Abs}(C)$ which is a simplification of any domain in $\mathcal{X}$.

Let $f : C \to D$ be some concrete semantic function defined on a pair of concrete domains C and D, and let $A \in \mathrm{Abs}(C)$ and $B \in \mathrm{Abs}(D)$ be a pair of abstractions. In the following, we will denote by $\sqsubseteq_{X \to Y}$ the pointwise ordering relation between functions in $X \to Y$. Given an abstract function $f^\sharp : A \to B$, we have that $\langle A, B, f^\sharp\rangle$ is a sound abstract interpretation of f when $f \circ \gamma_{A,C} \sqsubseteq_{A \to D} \gamma_{B,D} \circ f^\sharp$. Forward completeness [3,7] corresponds to the following strengthening of soundness: $f \circ \gamma_{A,C} = \gamma_{B,D} \circ f^\sharp$, meaning that the abstract function $f^\sharp$ is able to replicate the behaviour of f on the abstract domains A and B with no loss of precision. If an abstract interpretation $\langle A, B, f^\sharp\rangle$ is forward complete then it turns out that the abstract function $f^\sharp$ indeed coincides with $\alpha_{D,B} \circ f \circ \gamma_{A,C}$, which is the best correct approximation of the concrete function f on the pair of abstractions $\langle A, B\rangle$. Hence, the notion of forward completeness of an abstract interpretation $\langle A, B, f^\sharp\rangle$ does not depend on the choice of the abstract function $f^\sharp$ but only depends on the chosen abstract domains A and B. Accordingly, a pair of abstract domains $\langle A, B\rangle \in \mathrm{Abs}(C) \times \mathrm{Abs}(D)$ is called forward complete for f (or simply f-complete) iff $f \circ \gamma_{A,C} = \gamma_{B,D} \circ (\alpha_{D,B} \circ f \circ \gamma_{A,C})$. Equivalently, $\langle A, B\rangle$ is f-complete iff the image of f in D is contained in $\gamma_{B,D}(B)$, namely, $f(\gamma_{A,C}(A)) \subseteq \gamma_{B,D}(B)$. If $\mathcal{F} \subseteq C \to D$ is a set of concrete functions then $\langle A, B\rangle$ is $\mathcal{F}$-complete when $\langle A, B\rangle$ is f-complete for all $f \in \mathcal{F}$.

```
ShellAlgo(𝓕, 𝓖, A, B) {
  Initialize();
  while ¬(𝓕-Stable ∧ 𝓖-Stable) do
  | if ¬𝓕-Stable then 𝓖-Stable := Stabilize(𝓕, A, B); 𝓕-Stable := true;
  | if ¬𝓖-Stable then 𝓕-Stable := Stabilize(𝓖, B, A); 𝓖-Stable := true;
}

Initialize() {
  𝓕-Stable := CheckStability(𝓕, A, B); 𝓖-Stable := CheckStability(𝓖, B, A);
}

bool Stabilize(𝓗, X, Y) {
  Y_old := Y;
  Y := ⊔{Y' ∈ Abs | Y' ⊴ Y, ⟨X, Y'⟩ is 𝓗-complete};
  return (Y = Y_old);
}
```

Fig. 2 Basic Shell Algorithm.

3.2 Shells of Abstract Domains

Given a set of semantic functions $\mathcal{F} \subseteq C \to D$ and a pair of abstractions $\langle A, B\rangle \in \mathrm{Abs}(C) \times \mathrm{Abs}(D)$, the notion of forward complete shell [7] formalizes the problem of finding the most abstract pair $\langle A', B'\rangle$ such that $A' \trianglelefteq A$, $B' \trianglelefteq B$ and $\langle A', B'\rangle$ is $\mathcal{F}$-complete, which is a particular case of abstraction refinement [8]. It turns out (see [7]) that any pair $\langle A, B\rangle$ can be minimally refined to its forward $\mathcal{F}$-complete shell:

$$\mathrm{Shell}_{\mathcal{F}}(\langle A, B\rangle) \triangleq \sqcup \{\langle A', B'\rangle \in \mathrm{Abs}(C) \times \mathrm{Abs}(D) \mid \langle A', B'\rangle \trianglelefteq \langle A, B\rangle,\ \langle A', B'\rangle \text{ is } \mathcal{F}\text{-complete}\}.$$

Thus, $\mathrm{Shell}_{\mathcal{F}}(\langle A, B\rangle)$ encodes the least refinement of a pair of abstractions $\langle A, B\rangle$ which is needed in order to achieve forward completeness for $\mathcal{F}$.

Let us now consider a further set of concrete semantic functions $\mathcal{G} \subseteq D \to C$ that operate in the opposite direction w.r.t. $\mathcal{F}$, i.e., from D to C. Given $A \in \mathrm{Abs}(C)$ and $B \in \mathrm{Abs}(D)$, it makes sense to consider both forward $\mathcal{F}$-completeness of $\langle A, B\rangle$ and forward $\mathcal{G}$-completeness of the reversed pair $\langle B, A\rangle$. Thus, $\langle A, B\rangle$ is defined to be $\langle \mathcal{F}, \mathcal{G}\rangle$-complete when $\langle A, B\rangle$ is $\mathcal{F}$-complete and $\langle B, A\rangle$ is $\mathcal{G}$-complete. Here again, any pair $\langle A, B\rangle$ can be minimally refined to its $\langle \mathcal{F}, \mathcal{G}\rangle$-complete shell:

$$\mathrm{Shell}_{\langle \mathcal{F}, \mathcal{G}\rangle}(\langle A, B\rangle) \triangleq \sqcup \{\langle A', B'\rangle \in \mathrm{Abs}(C) \times \mathrm{Abs}(D) \mid \langle A', B'\rangle \trianglelefteq \langle A, B\rangle,\ \langle A', B'\rangle \text{ is } \langle \mathcal{F}, \mathcal{G}\rangle\text{-complete}\}.$$

Such a combined shell $\mathrm{Shell}_{\langle \mathcal{F}, \mathcal{G}\rangle}(\langle A, B\rangle)$ can be obtained through the ShellAlgo() procedure described in Figure 2. This procedure works by iteratively refining the abstractions A and B separately until both $\langle A, B\rangle$ becomes $\mathcal{F}$-complete and $\langle B, A\rangle$ becomes $\mathcal{G}$-complete. The ShellAlgo() procedure crucially relies on the Stabilize() function. In general, given a set of functions $\mathcal{H}$ and a pair of abstractions $\langle X, Y\rangle$, we have that Stabilize($\mathcal{H}, X, Y$) refines the abstraction Y to $Y_{\mathrm{stable}} \triangleq \sqcup\{Y' \mid Y' \trianglelefteq Y,\ \langle X, Y'\rangle \text{ is } \mathcal{H}\text{-complete}\}$, so that $\langle X, Y_{\mathrm{stable}}\rangle$ becomes $\mathcal{H}$-stable (i.e., $\mathcal{H}$-complete). Hence, Stabilize($\mathcal{F}, A, B$) minimally refines B to B' so that $\langle A, B'\rangle$ is $\mathcal{F}$-complete. Hence, while the abstraction B is refined, the abstraction A is left unchanged. Moreover, if B is actually refined into $B' \triangleleft B$, then the $\mathcal{G}$-Stable

flag is set to false so that ShellAlgo() proceeds by $\mathcal{G}$-stabilizing $\langle B', A\rangle$, i.e., by calling Stabilize($\mathcal{G}, B, A$). Thus, ShellAlgo($\mathcal{F}, \mathcal{G}, A, B$) begins by first checking whether $\langle A, B\rangle$ is $\mathcal{F}$-complete and $\langle B, A\rangle$ is $\mathcal{G}$-complete, and then proceeds by iteratively refining the abstractions A and B separately, namely it refines B w.r.t. $\mathcal{F}$ while A is kept fixed and then it refines A w.r.t. $\mathcal{G}$ while B is kept fixed. It turns out that the ShellAlgo() procedure is correct.

Theorem 3.1 ShellAlgo($\mathcal{F}, \mathcal{G}, A, B$) $= \mathrm{Shell}_{\langle \mathcal{F},\mathcal{G}\rangle}(\langle A, B\rangle)$.

Proof Firstly, we observe in general that if $Y_{\mathrm{stable}} \triangleq \sqcup\{Y' \in \mathrm{Abs} \mid Y' \trianglelefteq Y,\ \langle X, Y'\rangle$ is $\mathcal{H}$-complete$\}$ then $\langle X, Y_{\mathrm{stable}}\rangle$ is forward $\mathcal{H}$-complete. In fact, let $\langle X, Y_i\rangle_{i\in I}$ be a family of $\mathcal{H}$-complete pairs with $Y_i \trianglelefteq Y$, then by definition of forward $\mathcal{H}$-completeness we have that for all $i \in I$ and for all $h \in \mathcal{H}$, $h(\gamma(X)) \subseteq \gamma(Y_i)$, hence $h(\gamma(X)) \subseteq \cap_{i\in I}\gamma(Y_i)$, that is $h(\gamma(X)) \subseteq \gamma(\sqcup_{i\in I} Y_i)$; thus, $\langle X, \sqcup_{i\in I} Y_i\rangle$ is $\mathcal{H}$-complete. Secondly, the procedure ShellAlgo() always terminates because if $\langle A, B\rangle$ and $\langle A', B'\rangle$ are, respectively, the current abstraction pairs at the beginning and at the end of the while-loop, then either $A' \triangleleft A$ or $B' \triangleleft B$. Let $\langle A_{\mathrm{ref}}, B_{\mathrm{ref}}\rangle$ be the output abstraction pair of ShellAlgo($\mathcal{F}, \mathcal{G}, A, B$) and let $\langle A^s, B^s\rangle \triangleq \mathrm{Shell}_{\langle \mathcal{F},\mathcal{G}\rangle}(\langle A, B\rangle)$. Then, by the observation above, $\langle A_{\mathrm{ref}}, B_{\mathrm{ref}}\rangle$ is $\mathcal{F}$-complete and $\langle B_{\mathrm{ref}}, A_{\mathrm{ref}}\rangle$ is $\mathcal{G}$-complete, so that $\langle A_{\mathrm{ref}}, B_{\mathrm{ref}}\rangle \trianglelefteq \langle A^s, B^s\rangle$. Conversely, since $\langle A^s, B^s\rangle$ is $\mathcal{F}$-complete and $\langle B^s, A^s\rangle$ is $\mathcal{G}$-complete then it turns out that ShellAlgo($\mathcal{F}, \mathcal{G}, A^s, B^s$) $=$ $\langle A^s, B^s\rangle$. Moreover, observe that ShellAlgo() is monotone, meaning that if $\langle A_1, B_1\rangle \trianglelefteq \langle A_2, B_2\rangle$ then ShellAlgo($\mathcal{F}, \mathcal{G}, A_1, B_1$) $\trianglelefteq$ ShellAlgo($\mathcal{F}, \mathcal{G}, A_2, B_2$). Thus, from $\langle A^s, B^s\rangle \trianglelefteq \langle A, B\rangle$, we derive that $\langle A^s, B^s\rangle =$ ShellAlgo($\mathcal{F}, \mathcal{G}, A^s, B^s$) $\trianglelefteq$ ShellAlgo($\mathcal{F}, \mathcal{G}, A, B$) $=$ $\langle A_{\mathrm{ref}}, B_{\mathrm{ref}}\rangle$. $\square$

4 Bisimulation as a Shell

Bisimulation is commonly computed by coarsest partition refinement algorithms [1, 12] that iteratively refine a current partition until it becomes the bisimulation partition. Coarsest partition refinements can be formalized as shells of partitions: given a property of partitions $\mathbb{P} \subseteq \mathrm{Part}(X)$, the $\mathbb{P}$-shell of $Q \in \mathrm{Part}(X)$ corresponds to the coarsest partition refinement of Q that satisfies $\mathbb{P}$, when this exists. In this section we show how bisimulation in PLTSs can be equivalently stated in terms of forward complete shells of partitions w.r.t. suitable concrete semantic functions. We also show how the above basic shell algorithm ShellAlgo() can be instantiated to compute bisimulations on PLTSs.

4.1 Shells of Partitions

Let us first recall that, given a finite set X, $\langle \mathrm{Part}(X), \preceq, \curlyvee, \curlywedge\rangle$ is a (finite) lattice where $P_1 \preceq P_2$ (i.e., P_2 is coarser than P_1 or P_1 refines P_2) iff $\forall x \in X. P_1(x) \subseteq P_2(x)$, and its top element is $\top_{\mathrm{Part}(X)} = \{X\}$. By following the approach in [15], any partition $P \in \mathrm{Part}(X)$ can be viewed as an abstraction of $\langle \wp(X), \subseteq\rangle$, where any set $S \subseteq X$ is approximated through its minimal cover in the partition P. This is formalized by viewing P as the abstract domain $\mathrm{closed}(P) \triangleq \{S \subseteq X \mid P(S) = S\}$ so that $S \in \mathrm{closed}(P)$ iff $S = \cup_{i\in I} B_i$ for some blocks $\{B_i\}_{i\in I} \subseteq P$. Note that $\varnothing, X \in \mathrm{closed}(P)$ and that $\langle \mathrm{closed}(P), \subseteq, \cup, \cap\rangle$ is a lattice. It turns out that $\langle \mathrm{closed}(P), \subseteq\rangle$ is an abstraction in $\mathrm{Abs}(\wp(X)_{\subseteq})$, where any set $S \subseteq X$ is approximated through the blocks in P covering S, namely by $\cup\{B \in P \mid B \cap S \neq \varnothing\} \in \mathrm{closed}(P)$.

The above embedding of partitions as abstract domains allows us to define a notion of forward completeness for partitions. Let $f : \wp(X) \to \wp(Y)$ be a concrete semantic function that transforms subsets. Then, a pair of partitions $\langle P, Q\rangle \in \mathrm{Part}(X) \times \mathrm{Part}(Y)$ is (forward) f-complete when the pair of abstract domains $\langle \mathrm{closed}(P), \mathrm{closed}(Q)\rangle$ is forward complete for f, that is, $f(\mathrm{closed}(P)) \subseteq \mathrm{closed}(Q)$. In other terms, $\langle P, Q\rangle$ is f-complete when for any union $U \in \mathrm{closed}(P)$ of blocks of P, $f(U)$ still is a union of blocks of Q, namely $f(U) \in \mathrm{closed}(Q)$. Also, if we additionally consider $g : \wp(Y) \to \wp(X)$ then $\langle P, Q\rangle$ is $\langle f, g\rangle$-complete when $\langle P, Q\rangle$ is f-complete and $\langle Q, P\rangle$ is g-complete. Analogously to forward complete shells of generic abstract domains in Section 3, it is easy to see that forward complete shells of partitions exist. Given $\mathcal{F} \subseteq \wp(X) \to \wp(Y)$ and $\mathcal{G} \subseteq \wp(Y) \to \wp(X)$, $\mathrm{Shell}_{\langle \mathcal{F},\mathcal{G}\rangle}(\langle P, Q\rangle)$ is the coarsest pair of partitions that (component-wise) refines the pair $\langle P, Q\rangle$ and is $\langle \mathcal{F}, \mathcal{G}\rangle$-complete, namely

$$\mathrm{Shell}_{\langle \mathcal{F},\mathcal{G}\rangle}(\langle P, Q\rangle) \triangleq$$
$$\curlyvee \{\langle P', Q'\rangle \in \mathrm{Part}(X) \times \mathrm{Part}(Y) \mid \langle P', Q'\rangle \preceq \langle P, Q\rangle,\ \langle P', Q'\rangle \text{ is } \langle \mathcal{F}, \mathcal{G}\rangle\text{-complete}\}.$$

4.2 Bisimulation on PLTSs

Ranzato and Tapparo [15] have shown that bisimulation on a LTS $\mathcal{L}$ can be equivalently defined in terms of forward complete shells of partitions w.r.t. the predecessor operator of $\mathcal{L}$. This same idea scales to the case of PLTSs taking into account that: (i) in a PLTS the target of the transition relation is a set of distributions rather than a set of states, and (ii) bisimulation on the set of states of a PLTS induces an equivalence over distributions that depends on the probabilities that distributions assign to blocks of bisimilar states. Let $\mathcal{S} = \langle \Sigma, Act, \rightarrow\rangle$ be a PLTS and consider the following two functions, where $a \in Act$ and $p \in [0, 1]$:

$$\mathrm{pre}_a \;: \wp(\mathrm{Distr}) \to \wp(\Sigma),\ \mathrm{pre}_a(D) \triangleq \{s \in \Sigma \mid s \xrightarrow{a} D\}$$
$$\mathrm{prob}_p : \wp(\Sigma) \to \wp(\mathrm{Distr}),\ \mathrm{prob}_p(S) \triangleq \{d \in \mathrm{Distr} \mid d(S) \geq p\}$$

$\mathrm{pre}_a(D)$ is the a-predecessor function in the PLTS $\mathcal{S}$ for a set of distributions D while $\mathrm{prob}_p(S)$ returns the distributions whose probability on the set S is greater than or equal to p. Let us define $\mathrm{pre} \triangleq \{\mathrm{pre}_a\}_{a \in Act}$ and $\mathrm{prob} \triangleq \{\mathrm{prob}_p\}_{p\in[0,1]}$. It is worth noticing that this pair of sets of functions provides an encoding of the PLTS $\mathcal{S}$: pre encodes the transition relation $\rightarrow$ of $\mathcal{S}$, while any distribution d in $\mathcal{S}$ can be retrieved through functions in prob. For instance, the support of a distribution $d \in \mathrm{Distr}$ is given by the minimal set of states S such that $d \in \mathrm{prob}_1(S)$, while, for any $s \in \Sigma$, $d(s) = \sup \{p \in [0, 1] \mid d \in \mathrm{prob}_p(\{s\})\}$.

In the following, the bisimulation problem is stated in terms of forward completeness of pairs of abstract domains $\langle P, \mathcal{P}\rangle \in \mathrm{Part}(\Sigma) \times \mathrm{Part}(\mathrm{Distr})$ w.r.t. the concrete semantic functions $\mathrm{prob} \subseteq \wp(\Sigma) \to \wp(\mathrm{Distr})$ and $\mathrm{pre} \subseteq \wp(\mathrm{Distr}) \to \wp(\Sigma)$.

Lemma 4.1 *Consider a pair of partitions $\langle P, \mathcal{P}\rangle \in \mathrm{Part}(\Sigma) \times \mathrm{Part}(\mathrm{Distr})$.*

(i) $\langle P, \mathcal{P}\rangle$ *is* prob*-complete if and only if $e \in \mathcal{P}(d)$ then $d \equiv_P e$, i.e., if and only if for any block $B \in P$ and any distribution $d \in \mathrm{Distr}$, the set $\{e \in \mathrm{Distr} \mid e(B) = d(B)\}$ is a union of blocks of $\mathcal{P}$;*

(ii) $\langle \mathcal{P}, P\rangle$ *is* pre*-complete if and only if for any $a \in Act$, $s \xrightarrow{a} d$ and $t \in P(s)$ imply $t \xrightarrow{a} \mathcal{P}(d)$, i.e., if and only if for any block $C \in \mathcal{P}$ and for any incoming label $a \in \mathrm{in}(C)$, $\mathrm{pre}_a(C)$ is a union of blocks of P.*

Proof Let us prove (i). Assume that $\langle P, \mathcal{P}\rangle$ is prob-complete. Hence, for any $p \in [0,1]$ and for any block $B \in P$, if $d \in \mathrm{prob}_p(B)$ and $\mathcal{P}(d) = \mathcal{P}(e)$ then $e \in \mathrm{prob}_p(B)$. Consider therefore $d, e \in \mathrm{Distr}$ such that $\mathcal{P}(d) = \mathcal{P}(e)$. For any $B \in P$, let us define $p_B \triangleq d(B)$. Since $d \in \mathrm{prob}_{p_B}(B)$, we have that $e \in \mathrm{prob}_{p_B}(B)$, hence $e(B) \geq p_B = d(B)$. Thus, for any $B \in P$, $e(B) \geq d(B)$. Thus, $d(B) = 1 - \sum_{C\in P, C\neq B} d(C) = \sum_{C\in P} e(C) - \sum_{C\in P, C\neq B} d(C) \geq \sum_{C\in P} e(C) - \sum_{C\in P, C\neq B} e(C) = e(B)$, hence $d(B) = e(B)$ for any $B \in P$, namely $d \equiv_P e$.
For the opposite direction, assume that if $\mathcal{P}(d) = \mathcal{P}(e)$ then $d \equiv_P e$ and let us show that $\langle P, \mathcal{P}\rangle$ is prob-complete. Consider $X \subseteq \Sigma$, $p \in [0,1]$, $d \in \mathrm{prob}_p(P(X))$, $e \in \mathcal{P}(d)$ and let us show that $e \in \mathrm{prob}_p(P(X))$. We have that $P(X) = \cup_{i\in I} B_i$ for some set of blocks $\{B_i\}_{i\in I} \subseteq P$. Hence, from $d \equiv_P e$, we obtain that $d(\cup_i B_i) = e(\cup_i B_i)$, so that from $d(\cup_i B_i) \geq p$ we get $e(\cup_i B_i) \geq p$, namely $e \in \mathrm{prob}_p(P(X))$.
It remains to prove that $\langle P, \mathcal{P}\rangle$ is prob-complete iff for any block $B \in P$ and any distribution $d \in \mathrm{Distr}$, $\{e \in \mathrm{Distr} \mid e(B) = d(B)\}$ is a union of blocks of $\mathcal{P}$.
($\Rightarrow$) Consider $B \in P$ and $d, e, f \in \mathrm{Distr}$ such that $d(B) = f(B)$ and $\mathcal{P}(d) = \mathcal{P}(e)$. Let us show that $e(B) = f(B)$. Let $p \triangleq f(B)$, so that $d \in \mathrm{prob}_p(P(B))$. Hence, by hypothesis, $\mathcal{P}(d) \subseteq \mathrm{prob}_p(P(B))$, so that $e \in \mathrm{prob}_p(P(B))$, namely $e(P(B)) = e(B) \geq p$. Let $Y \triangleq \Sigma \smallsetminus B$ so that $Y = \cup_{C\in P, C\neq B} C$. We have that $f(Y) = f(P(Y)) = 1 - p = d(Y) = d(P(Y))$, so that $d \in \mathrm{prob}_{1-p}(P(Y))$. Hence, by hypothesis, $\mathcal{P}(d) \subseteq \mathrm{prob}_{1-p}(P(Y))$, so that $e \in \mathrm{prob}_{1-p}(P(Y))$, namely $e(P(Y)) = e(Y) \geq 1 - p$, from which $e(B) = 1 - e(Y) \leq p$. Hence, $e(B) = p = f(B)$.
($\Leftarrow$) Consider $S \subseteq \Sigma$ and $p \in [0,1]$ and let us show that $\mathrm{prob}_p(P(S))$ is a union of blocks of $\mathcal{P}$. Hence, consider $d, e \in \mathrm{Distr}$ such that $d(P(S)) \geq p$ and $\mathcal{P}(d) = \mathcal{P}(e)$ and let us show that $e(P(S)) \geq p$. Let $P(S) = \cup_{i\in I} B_i$ for some set of blocks $\{B_i\}_{i\in I} \subseteq P$. By hypothesis, for any $i \in I$, $\{h \in \mathrm{Distr} \mid h(B_i) = d(B_i)\}$ is a union of blocks of $\mathcal{P}$, so that $e(B_i) = d(B_i)$ for any $i \in I$, so that $e(P(S)) = \sum_{i\in I} e(B_i) = \sum_{i\in I} d(B_i) = d(P(S)) \geq p$.
Let us consider item (ii). Note that since any function pre_a is additive, $\langle \mathcal{P}, P\rangle$ is pre_a-complete if and only if for any $d \in \mathrm{Distr}$, $\mathrm{pre}_a(\mathcal{P}(d))$ is a union of blocks in P, i.e., if and only if for any $d \in \mathrm{Distr}$, if $s \xrightarrow{a} \mathcal{P}(d)$ and $P(s) = P(t)$ then $t \xrightarrow{a} \mathcal{P}(d)$, i.e., if and only if for any block $C \in \mathcal{P}$ and for any incoming label $a \in \mathrm{in}(C)$, we have that $\mathrm{pre}_a(C)$ is a union of blocks of P. $\square$

As a direct consequence of the above characterization, we have that a partition $P \in \mathrm{Part}(\Sigma)$ is a bisimulation on $\mathcal{S}$ if and only if the pair of partitions $\langle P, \equiv_P\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete. In turn, the coarsest bisimulation P_{bis} on $\mathcal{S}$ can be obtained as a forward complete shell of partitions.

Corollary 4.2 *Let $\mathcal{S} = \langle \Sigma, Act, \to\rangle$ be a PLTS.*

(i) $P \in \mathrm{Part}(\Sigma)$ *is a bisimulation on* $\mathcal{S}$ *if and only if* $\langle P, \equiv_P\rangle$ *is* $\langle \mathrm{prob}, \mathrm{pre}\rangle$*-complete.*
(ii) *Let* $\langle P, \mathcal{P}\rangle \in \mathrm{Part}(\Sigma) \times \mathrm{Part}(\mathrm{Distr})$. *If* $\langle P, \mathcal{P}\rangle$ *is* $\langle \mathrm{prob}, \mathrm{pre}\rangle$*-complete then* P *is a bisimulation on* $\mathcal{S}$ *and* $\mathcal{P} \subseteq\, \equiv_P$.

Theorem 4.3 $\langle P_{\mathrm{bis}}, \equiv_{P_{\mathrm{bis}}}\rangle = \mathrm{Shell}_{\langle \mathrm{prob}, \mathrm{pre}\rangle}(\top_{\mathrm{Part}(\Sigma)}, \top_{\mathrm{Part}(\mathrm{Distr})})$.

Proof Let $\langle P^*, \mathcal{P}^*\rangle = \mathrm{Shell}_{\langle \mathrm{prob}, \mathrm{pre}\rangle}(\top_{\mathrm{Part}(\Sigma)}, \top_{\mathrm{Part}(\mathrm{Distr})})$. Since we have that $\langle P^*, \mathcal{P}^*\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete, by Corollary 4.2 (ii), we have that $P^* \preceq P_{\mathrm{pbis}}$ and $\mathcal{P}^* \subseteq\, \equiv_{P^*}$. Hence, we also have that $\mathcal{P}^* \subseteq\, \equiv_{P_{\mathrm{pbis}}}$. For the opposite direction, by Corollary 4.2 (i), we have that $\langle P_{\mathrm{pbis}}, \equiv_{P_{\mathrm{pbis}}}\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete, so that, by definition of shell, it turns out that $\langle P_{\mathrm{pbis}}, \equiv_{P_{\mathrm{pbis}}}\rangle \preceq \langle P^*, \mathcal{P}^*\rangle$, namely $P_{\mathrm{pbis}} \preceq P^*$ and $\equiv_{P_{\mathrm{pbis}}} \subseteq \mathcal{P}^*$. $\square$

```
PBis(𝒮) {
   Initialize();
   while ¬ (probStable ∧ preStable) do
      if ¬ probStable then preStable := probStabilize(⟨P, 𝒫⟩); probStable := true;
      if ¬ preStable then probStable := preStabilize(⟨𝒫, P⟩); preStable := true;
}

Initialize() {
   forall s ∈ Σ do P(s) := Σ;
   forall d ∈ Distr do 𝒫(d) := Distr;
   preStabilize(⟨𝒫, P⟩); preStable := probStabilize(⟨P, 𝒫⟩); probStable := true;
}

bool preStabilize(⟨𝒫, P⟩) {
   P_old := P;
   forall C ∈ 𝒫 do
      forall a ∈ in(C) do P := Split(P, pre_a(C));
   return (P ≠ P_old)
}

bool probStabilize(⟨P, 𝒫⟩) {
   𝒫_old := 𝒫;
   forall B ∈ P do
      forall d ∈ Distr do 𝒫 := Split(𝒫, {e ∈ Distr | e(B) = d(B)});
   return (𝒫 ≠ 𝒫_old)
}
```

Fig. 3 Bisimulation Algorithm PBis.

4.3 Bisimulation Algorithm

By Theorem 4.3, P_{bis} can be computed as a partition shell by instantiating the basic shell algorithm in Figure 2 to $\mathcal{F} = \{\mathrm{prob}_p\}_{p\in[0,1]}$ and $\mathcal{G} = \{\mathrm{pre}_a\}_{a\in Act}$, and by viewing partitions in $\mathrm{Part}(\Sigma) \times \mathrm{Part}(\mathrm{Distr})$ as abstract domains. This leads to a bisimulation algorithm as described in Figure 3, called PBis, that takes a PLTS $\mathcal{S}$ as input and initializes and stabilizes a pair of state and distribution partitions $\langle P, \mathcal{P}\rangle$ until it becomes $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete.

Stabilization is obtained by means of two auxiliary functions preStablize() and probStabilize(), that implement Lemma 4.1. In particular, the function call preStabilize($\langle \mathcal{P}, P\rangle$) refines the state partition P into P' so that $\langle \mathcal{P}, P'\rangle$ is pre-complete. According to Lemma 4.1, in order to get pre-completeness it is sufficient to minimally refine P so that for any block of distributions $C \in \mathcal{P}$, and for any incoming label $a \in \mathrm{in}(C)$, $\mathrm{pre}_a(C)$ is a union of blocks of P. If $\mathrm{pre}_a(C)$ is not a union of blocks of P then $\mathrm{pre}_a(C) \subseteq \Sigma$ is called a splitter of P, and we denote by $\mathrm{Split}(P, \mathrm{pre}_a(C))$ the partition obtained from P by replacing each block $B \in P$ with the nonempty sets $B \cap \mathrm{pre}_a(C)$ and $B \smallsetminus \mathrm{pre}_a(C)$. Notice that when some $\mathrm{pre}_a(C)$ is already a union of blocks of P we have that $\mathrm{Split}(P, \mathrm{pre}_a(C)) = P$, i.e., we also allow no splitting. Hence, preStabilize() refines P by iteratively splitting P w.r.t. $\mathrm{pre}_a(C)$, for all $C \in \mathcal{P}$ and $a \in \mathrm{in}(C)$. On the other hand, the function call probStabilize($\langle P, \mathcal{P}\rangle$) refines the current distribution partition $\mathcal{P}$ into $\mathcal{P}'$ so that $\langle P, \mathcal{P}'\rangle$ is prob-complete. According to Lemma 4.1, $\langle P, \mathcal{P}\rangle$ is prob-complete when for any block $B \in P$ and any distribution $d \in \mathrm{Distr}$, $\{e \in \mathrm{Distr} \mid e(B) = d(B)\}$ is a union of blocks of $\mathcal{P}$. Thus, probStabilize() iteratively splits the distribution partition $\mathcal{P}$ w.r.t. $\{e \in \mathrm{Distr} \mid e(B) = d(B)\}$, for all $B \in P$ and $d \in \mathrm{Distr}$.

The initialization phase is carried out by the Initialize() function. The two current partitions P and $\mathcal{P}$ are initialized with the top elements $\top_{\mathrm{Part}(\Sigma)}$ and $\top_{\mathrm{Part}(\mathrm{Distr})}$, i.e., $\{\Sigma\}$ and $\{\mathrm{Distr}\}$. Moreover, in order to inizialize the two boolean flags probStable and preStable, Initialize() first calls preStabilize() and then calls probStabilize(). Therefore, the initial value of probStable is true, while that of preStable is either true or false depending on whether the prob-stabilization has invalidated the previous pre-stabilization or not.

Theorem 4.4 *For a finite PLTS $\mathcal{S}$, $\mathrm{PBis}(\mathcal{S})$ terminates and is correct, i.e., if $\langle P, \mathcal{P}\rangle$ is the output of $\mathrm{PBis}(\mathcal{S})$ then $P = P_{\mathrm{bis}}$ and $\mathcal{P} = \equiv_{P_{\mathrm{bis}}}$.*

Proof A consequence of Theorems 3.1 and 4.3. □

Implementation. Baier-Engelen-Majster's two-phased partitioning algorithm [1] is the standard procedure for computing the bisimulation P_{bis}. This bisimulation algorithm can be essentially viewed as an implementation of the above PBis algorithm, since the two phases of Baier et al.'s algorithm (see [1, Figure 9]) coincide with our preStabilize() and probStabilize() functions. The only remarkable difference is that instead of using a single partition over all the distributions in Distr, Baier et al.'s algorithm maintains a so-called step partition, namely, a family of partitions $\{M_a\}_{a\in Act}$ such that, for any $a \in Act$, M_a is a partition of the distributions in $\mathrm{post}_a(\Sigma)$, i.e., the distributions that have an incoming edge labeled with a. As a consequence, in the phase that corresponds to probStabilize(), any partition M_a is split w.r.t. all the splitters $\{e \in \mathrm{post}_a(\Sigma) \mid e(B) = d(B)\}$, where $B \in P$ and $d \in \mathrm{post}_a(\Sigma)$. Baier et al.'s algorithm is implemented by exploiting Hopcroft's "process the smaller half" principle when splitting the state partition w.r.t. a splitter $\mathrm{pre}_a(C)$ and this yields a procedure that computes bisimulation on PLTSs in $O(|{\to}||\Sigma|(\log|{\to}| + \log|\Sigma|))$ time and $O(|{\to}||\Sigma|)$ space.

5 Simulation as a Shell

Let us focus on simulation preorders in PLTSs. We show that the simulation preorder is a complete shell of abstract domains w.r.t. the same operators prob and pre considered above for bisimulation equivalence, whereas the key difference lies in the underlying abstract domains that in this case are preorders, rather than partitions, viewed as abstractions.

5.1 Shells of Preorders

Recall that, given any finite set X, $\langle \mathrm{PreOrd}(X), \subseteq, \cup^t, \cap\rangle$ is a lattice, where $R_1 \cup^t R_2$ is the transitive closure of $R_1 \cup R_2$ and the top element is $\top_{\mathrm{PreOrd}(X)} \triangleq X \times X$. Analogously to partitions, any preorder $R \in \mathrm{PreOrd}(X)$ can be viewed as an abstraction of $\langle \wp(X), \subseteq\rangle$, where any set $S \subseteq X$ is approximated by its R-closure $R(S)$. Formally, a preorder $R \in \mathrm{PreOrd}(X)$ can be viewed as the abstract domain $\mathrm{closed}(R) \triangleq \{S \subseteq X \mid R(S) = S\}$. Observe that $S \in \mathrm{closed}(R)$ iff $S = \cup_{i\in I} R(x_i)$ for some set $\{x_i\}_{i\in I} \subseteq X$ and that $\langle \mathrm{closed}(R), \subseteq, \cup, \cap\rangle$ is a lattice (note that $\varnothing, X \in \mathrm{closed}(R)$). It turns out that $\mathrm{closed}(R) \in \mathrm{Abs}(\wp(X)_\subseteq)$: this means that any set $S \subseteq X$ is approximated by its R-closure, namely by $R(S) \in \mathrm{closed}(R)$.

Given a pair of sets of functions $\langle \mathcal{F}, \mathcal{G}\rangle \subseteq (\wp(X) \to \wp(Y)) \times (\wp(Y) \to \wp(X))$, a pair of preorders $\langle R, S\rangle \in \mathrm{PreOrd}(X) \times \mathrm{PreOrd}(Y)$ is (forward) $\langle \mathcal{F}, \mathcal{G}\rangle$-complete when

for any $f \in \mathcal{F}$ and $g \in \mathcal{G}$, if $\langle U, V\rangle \in \mathrm{closed}(R) \times \mathrm{closed}(S)$ then $\langle f(U), g(V)\rangle \in \mathrm{closed}(S) \times \mathrm{closed}(R)$. Forward complete shells of preorders are therefore defined as follows: $\mathrm{Shell}_{\langle \mathcal{F},\mathcal{G}\rangle}(\langle R, S\rangle)$ is the largest pair of preorders $\langle R', S'\rangle \subseteq \langle R, S\rangle$ which is $\langle \mathcal{F}, \mathcal{G}\rangle$-complete.

5.2 Simulation on PLTSs

Similarly to the case of bisimulation, simulation can be equivalently expressed in terms of forward completeness w.r.t. $\mathrm{prob} = \{\mathrm{prob}_p\}_{p\in[0,1]}$ and $\mathrm{pre} = \{\mathrm{pre}_a\}_{a\in Act}$.

Lemma 5.1 *Consider a pair of preorders $\langle R, \mathcal{R}\rangle \in \mathrm{PreOrd}(\Sigma) \times \mathrm{PreOrd}(\mathrm{Distr})$.*

(i) $\langle R, \mathcal{R}\rangle$ *is* prob-*complete if and only if $e \in \mathcal{R}(d)$ implies $d \leq_R e$;*
(ii) $\langle \mathcal{R}, R\rangle$ *is* pre-*complete if and only if for any $a \in Act$, $t \in R(s)$ and $s \xrightarrow{a} d$ imply that there exists e such that $t \xrightarrow{a} e$ and $e \in \mathcal{R}(d)$.*

Proof Let us prove (i), that is, $\langle R, \mathcal{R}\rangle$ is prob-complete iff $\mathcal{R} \subseteq \leq_R$. Recall (see [20]) that $d \leq_R e$ iff for any $Z \subseteq \Sigma$, $d(R(Z)) \leq e(R(Z))$.
($\Rightarrow$) Assume that $e \in \mathcal{R}(d)$. For any $Z \subseteq \Sigma$, let us define $p_Z \triangleq d(R(Z))$ so that $d \in \mathrm{prob}_{p_Z}(R(Z))$. Thus, by prob-completeness, $e \in \mathrm{prob}_{p_Z}(R(Z))$, namely $e(R(Z)) \geq d(R(Z))$. Since this holds for any $Z \subseteq \Sigma$, we have that $d \leq_R e$.
($\Leftarrow$) Consider $Z \subseteq \Sigma$, $p \in [0, 1]$, $d \in \mathrm{prob}_p(R(Z))$, $e \in \mathcal{R}(d)$ and let us show that $e \in \mathrm{prob}_p(R(Z))$. By hypothesis, $d \leq_R e$, so that $e(R(Z)) \geq d(R(Z)) \geq p$, that is, $e \in \mathrm{prob}_p(R(Z))$.
Let us turn to (ii). Note that since any function pre_a is additive, $\langle \mathcal{R}, R\rangle$ is pre_a-complete if and only if for any $d \in \mathrm{Distr}$, $\mathrm{pre}_a(\mathcal{R}(d)) \in \mathrm{closed}(R)$, that is, if and only if for any $d \in \mathrm{Distr}$, if $s \xrightarrow{a} \mathcal{R}(d)$ and $t \in R(s)$ then $t \in \mathrm{pre}_a(\mathcal{R}(d))$, i.e., there exists e such that $t \xrightarrow{a} e$ with $e \in \mathcal{R}(d)$. □

Thus, a preorder $R \in \mathrm{PreOrd}(\Sigma)$ is a simulation on $\mathcal{S}$ if and only if the pair $\langle R, \leq_R\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete. In turn, the greatest simulation preorder R_{sim} can be obtained as a preorder shell.

Corollary 5.2 *Let $\mathcal{S} = \langle \Sigma, Act, \rightarrow\rangle$ be a PLTS.*

(i) *$R \in \mathrm{PreOrd}(\Sigma)$ is a simulation on $\mathcal{S}$ if and only if $\langle R, \leq_R\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete.*
(ii) *Let $\langle R, \mathcal{R}\rangle \in \mathrm{PreOrd}(\Sigma) \times \mathrm{PreOrd}(\mathrm{Distr})$. If $\langle R, \mathcal{R}\rangle$ is $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete, then R is a simulation on $\mathcal{S}$ and $\mathcal{R} \subseteq \leq_R$.*

Theorem 5.3 $\langle R_{\mathrm{sim}}, \equiv_{R_{\mathrm{sim}}}\rangle = \mathrm{Shell}_{\langle \mathrm{prob},\mathrm{pre}\rangle}(\top_{\mathrm{PreOrd}(\Sigma)}, \top_{\mathrm{PreOrd}(\mathrm{Distr})})$.

Proof Analogous to the proof of Theorem 4.3. □

6 A New Efficient Simulation Algorithm on PLTSs

We show how a new efficient algorithm for computing simulations in PLTSs, called PSim, can be derived by instantiating the basic shell algorithm to $\mathcal{F} = \{\mathrm{prob}_p\}_{p\in[0,1]}$ and $\mathcal{G} = \{\mathrm{pre}_a\}_{a\in Act}$, and by viewing preorders in $\mathrm{PreOrd}(\Sigma)$ and $\mathrm{PreOrd}(\mathrm{Distr})$ as abstract domains. Similarly to the case of bisimulation, PSim, which is described in Figure 4, takes

```
PSim(𝒮) {
  Initialize();
  while ¬ (probStable ∧ preStable) do
      if ¬ probStable then  preStable := probStabilize(⟨R, ℛ⟩);  probStable := true;
      if ¬ preStable then  probStable := preStabilize(⟨ℛ, R⟩);  preStable := true;
}
```

Fig. 4 PSim Algorithm.

```
 1 Initialize() {
      // Initialize R and ℛ
 2    forall s ∈ Σ do R(s) := {t ∈ Σ | out(s) ⊆ out(t)};
 3    forall d ∈ Distr do ℛ(d) := {e ∈ Distr | Init_SMF(d, e, R) = true};
      // Initialize in
 4    forall d ∈ Distr do in(d) := {a ∈ Act | pre_a(d) ≠ ∅};
      // Initialize Count
 5    forall e ∈ Distr do
 6        forall a ∈ in(e) do
 7            forall x ∈ pre_a(Distr) do
 8                Count(x, a, e) := |{d ∈ Distr | x -a-> d, d ∈ ℛ(e), a ∈ in(e)}|;
      // Initialize Remove
 9    forall d ∈ Distr do
10        forall a ∈ in(d) do
11            Remove_a(d) := {s ∈ Σ | a ∈ out(s), s -a-/-> ℛ(d)};
      // Initialize Stability Flags
12    probStable := true;
13    if ∃e ∈ Distr, a ∈ in(e) such that Remove_a(e) ≠ ∅ then preStable := false;
14    else preStable := true;
      // Initialize Listener
15    forall x, y ∈ Σ do Listener(x, y) := {(d, e) | x ∈ supp(d), e ∈ supp(e)};
      // Initialize Deleted Arcs
16    Deleted := ∅;
17 }
```

Fig. 5 High-level definition of Initialize() function.

a PLTS $\mathcal{S}$ as input and initializes and stabilizes a pair of state and distribution preorders $\langle R, \mathcal{R}\rangle \in \mathrm{PreOrd}(\Sigma) \times \mathrm{PreOrd}(\mathrm{Distr})$ until it becomes $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete.

The stabilization functions, which are given in Figure 6, refine the preorders according to Lemma 5.1, that is

- the function preStabilize() makes the pair $\langle \mathcal{R}, R\rangle$ pre-complete by refining the state preorder R until there exists a transition $s \xrightarrow{a} d$ such that and $R(s) \not\subseteq \mathrm{pre}_a(\mathcal{R}(d))$;
- the function probStabilize() makes the pair $\langle \mathcal{R}, R\rangle$ prob-complete by refining the distribution preorder $\mathcal{R}$ by iteratively refining it until there exist $e, d \in \mathrm{Distr}$ such that $e \in \mathcal{R}(d)$ and $d \not\leq_R e$.

The design of these functions allows us to refine the preorders R and $\mathcal{R}$ by following an efficient incremental approach. In particular, preStabilize() refines the preorder R by mimicking the incremental approach of Henzinger et al. [10] simulation algorithm for non-probabilistic LTSs. On the other hand, the function probStabilize() resorts to the incremental

```
1  bool preStabilize(⟨ℛ, R⟩) {
2    Deleted := ∅;
3    while ∃ Remove_a(e) ≠ ∅ do
4      Remove := Remove_a(e); Remove_a(e) := ∅;
5      forall t -a-> e do
6        forall w ∈ Remove do
7          if w ∈ R(t) then Deleted := Deleted ∪ {(t, w)}; R(t) := R(t) ∖ {w};
8    if (Deleted ≠ ∅) then probStable := false;
9    return probStable;
10 }

1  bool probStabilize(⟨R, ℛ⟩) {
2    forall (t, w) ∈ Deleted do
3      forall (d, e) ∈ Listener(t, w) such that e ∈ ℛ(d) do
4        if SMF(d, e, (t, w)) = false then
5          ℛ(d) := ℛ(d) ∖ {e};
6          forall b ∈ in(e) ∩ in(d) do
7            forall s -b-> e do
8              Count(s, b, d)--;
9              if Count(s, b, d) = 0 then
10               Remove_b(d) := Remove_b(d) ∪ {s};  preStable := false;
11   return preStable;
12 }
```

Fig. 6 Stabilization functions.

approach of Zhang et al. [18] simulation algorithm, and stabilizes the distribution preorder $\mathcal{R}$ by computing sequences of maximum flow problems. More precisely, given a pair of distributions (d, e), successive calls to probStabilize() might repeatedly check whether $d \leq_R e$ where R is the current (new) state preorder. This amounts to repeatedly check whether the maximum flow over the net $\mathcal{N}(d, e, R)$ remains 1 with the current (new) preorder R. Zhang et al. [18] observe that the networks for a given pair (d, e) across successive iterations of their algorithm are very similar, since they differ only by deletion of some edges due to the refinement of R. Therefore, in order to incrementally deal with this sequence of tests, Zhang et al.'s algorithm stores after each iteration the current network $\mathcal{N}(d, e, R)$ together with its maximum flow information, so that at the next iteration, instead of computing the maximum flow of the full new network, one can exploit a so-called preflow algorithm which is initialized with the previous maximum flow function. We do not discuss the details of the preflow algorithm by Zhang et al. [18], since it can be used here as a black box that incrementally solves the sequence of maximum flow problems that arise for a same network.

PSim relies on a number of data structures, whose initialization is provided by the Initialize() function, which is described in Figure 5 at high-level and fully implemented in Figure 7. First, the two preorders $R \subseteq \Sigma \times \Sigma$ and $\mathcal{R} \subseteq \mathrm{Distr} \times \mathrm{Distr}$ are stored as boolean matrices and are initialized in such a way that they are coarser than, respectively, R_{sim} and $\leq_{R_{\mathrm{sim}}}$. In particular, the initial preorder R is coarser than R_{sim} since if $s \xrightarrow{a} d$ and $t \not\xrightarrow{a}$ then $t \notin R_{\mathrm{sim}}(s)$. Moreover, line 3 of Figure 5 initializes $\mathcal{R}$ so that $\mathcal{R} = \leq_R$: this is done by calling the function Init_SMF(d, e, R) which in turn calls the preflow algorithm to check whether $d \leq_R e$, and in case this is true, stores the network $\mathcal{N}(d, e, R)$ in order to reuse it in later calls to probStabilize(). The additional data structures used by PSim come

```
1  Initialize() {
     // Initialize In
2    forall d ∈ Distr do
3      in(d) = ∅;
4      forall a ∈ Act such that pre_a(d) ≠ ∅ do in(d) := in(d) ∪ {a};
     // Initialize R
5    forall s, t ∈ Σ do R(s, t) :=true;
6    forall a ∈ Act do
7      forall d ∈ Distr do
8        forall x ∈ pre_a(d) do  mark(x)=true;
9      forall y ∈ Σ such that mark(y)=false do
10       forall d ∈ Distr do
11         forall x ∈ pre_a(d) do R(x, y) :=false;
12     forall x ∈ Σ do mark(x)=false;
     // Initialize ℛ
13   forall d, e ∈ Distr do  ℛ(d, e) :=Init_SMF(d, e, R)
     // Initialize Count
14   forall a ∈ Act do
15     forall d ∈ Distr do
16       forall x ∈ pre_a(d) do
17         forall e ∈ Distr such that pre_a(e) ≠ ∅ do Count(x, a, e) = 0;
18   forall d ∈ Distr do
19     forall a ∈ in(d) do
20       forall x ∈ pre_a(d) do
21         forall e ∈ Distr such that ℛ(e, d) ∧ pre_a(e) ≠ ∅ do Count(x, a, e)++;
     // Initialize Remove
22   forall d ∈ Distr do
23     forall a ∈ in(d) do
24       Remove_a(d) := ∅;
25       forall x ∈ Σ do
26         if Count(x, a, d) = 0 then Remove_a(d) := Remove_a(d) ∪ {x};
     // Initialize Stability Flags
27   probStable:= true;
28   if ∃e ∈ Distr, a ∈ in(e) such that Remove_a(e) ≠ ∅ then preStable := false;
29   else preStable := true;
     // Initialize Listener
30   forall x, y ∈ Σ do Listener(x, y) := ∅;
31   forall d, e ∈ Distr do
32     forall x∈ supp(d), y∈ supp(e) do Listener(x, y):= Listener(x, y)∪{(d, e)};
     // Initialize Deleted Arcs
33   Deleted := ∅;
34 }
```

Fig. 7 Implementation of Initialize() function.

from the incremental refinement methods used in [10] and [18]. Actually, as in [10], for any distribution e and for any incoming action $a \in \mathrm{in}(e)$, we store and maintain a set

$$\mathrm{Remove}_a(e) \triangleq \{s \in \Sigma \mid s \xrightarrow{a},\ s \stackrel{a}{\not\rightarrow} \mathcal{R}(e)\}$$

that is used to prune the relation R to get pre-completeness (lines 5-7 of preStabilize()). The $\{\mathrm{Count}(s, a, e)\}_{e \in \mathrm{Distr}, a \in \mathrm{in}(e), s \in \mathrm{pre}_a(\mathrm{Distr})}$ table records in any entry $\mathrm{Count}(s, a, e)$

the number of a-transitions from state s to a distribution in $\mathcal{R}(e)$, so that it can be used to efficiently refill the remove sets (line 10 of probStabilize()), since it allows to test whether $s \not\xrightarrow{a} \mathcal{R}(e)$ in $O(1)$ by checking whether $\mathrm{Count}(s, a, e)$ is equal to 0. Moreover, in order to get an efficient refinement also for the distribution preorder $\mathcal{R}$, likewise to Zhang et al.'s algorithm, for any pair of states (x, y) we compute and store a set

$$\mathrm{Listener}(x, y) \triangleq \{(d, e) \in \mathrm{Distr} \times \mathrm{Distr} \mid x \in \mathrm{supp}(d),\, y \in \mathrm{supp}(e)\}$$

that contains all the pairs of distributions (d, e) such that the network $\mathcal{N}(d, e, R)$ could contain the edge (x, y), i.e., $\mathrm{Listener}(x, y)$ records the networks that are affected when the pair of states (x, y) is removed from the preorder R. Indeed, these sets are used in probStabilize() to recognize the pairs (d, e) that have been affected by the refinement of R due to the previous call of preStabilize() (lines 2-3 of probStabilize()).

At the end of initialization, the probStable flag is set to true (due to the initialization of $\mathcal{R}$ as $\leq_R$), whereas the preStable flag is set to false if at least a nonempty remove set exists. The main loop of PSim then repeatedly calls the stabilization functions until the pair $\langle R, \mathcal{R}\rangle$ becomes $\langle \mathrm{prob}, \mathrm{pre}\rangle$-complete. More precisely, a call to preStabilize(): (i) refines the relation R in such a way that if $t \xrightarrow{a} e$ then $R(t)$ is pruned to $R(t) \smallsetminus \mathrm{Remove}_a(e)$, (ii) empties all the Remove sets and collects all the pairs removed from R into the set $\mathrm{Deleted}$, and (iii) sets the probStable flag to false when R has changed. On the other hand, a call to probStabilize() exploits the sets $\mathrm{Deleted}$ and $\mathrm{Listener}$ to determine all the networks $\mathcal{N}(d, e, R)$ that have been affected by the refinement of R due to preStabilize(). For any pair (t, w) that has been removed from R, the call SMF$(d, e, (t, w))$ at line 4 removes the edge (t, w) from the network for (d, e) and calls the preflow algorithm to check whether it still has a maximum flow equals to 1. Then, if this is not the case, e is removed from $\mathcal{R}(d)$. Notice that such a pruning may induce an update of some $\mathrm{Remove}_b(d)$ set, which in turn triggers a further call of preStabilize() by setting the preStable flag to false.

Example 6.1 Let us illustrate how the algorithm PSim works on the PLTS in Figure 1, where $\Sigma = \{s_1, s_2, s_3, x_1, \ldots, .x_6, t, u, v, w\}$ and $\mathrm{Distr} = \{d_1, d_2, d_3, \delta_t, \delta_u, \delta_v, \delta_w\}$. The call to Initialize() yields the following preorders and remove sets:

$R(x_1) = R(x_5) = \{x_1, x_3, x_5\}$ $\quad$ $R(s_1) = R(s_2) = R(s_3) = \{s_1, s_2, s_3\}$
$R(x_2) = R(x_4) = \{x_2, x_4\}$ $\quad$ $R(t) = R(u) = R(v) = R(w) = \Sigma$
$R(x_3) = \{x_3\}$ $\quad$ $R(x_6) = \{x_6, x_3\}$

$\mathcal{R}(d_1) = \{d_1, d_2\}$ $\quad$ $\mathcal{R}(d_2) = \{d_2\}$
$\mathcal{R}(d_3) = \{d_3\}$ $\quad$ $\mathcal{R}(\delta_t) = \mathcal{R}(\delta_u) = \mathcal{R}(\delta_v) = \mathcal{R}(\delta_w) = \mathrm{Distr}$

$\mathrm{Remove}_a(d_1) = \{s_3\}$ $\quad$ $\mathrm{Remove}_a(d_2) = \{s_1, s_3\}$
$\mathrm{Remove}_a(d_3) = \{s_1, s_2\}$ $\quad$ $\mathrm{Remove}_b(\delta_t) = \mathrm{Remove}_c(\delta_u) = \varnothing$
$\mathrm{Remove}_d(\delta_v) = \mathrm{Remove}_d(\delta_w) = \varnothing$

The main loop of PSim begins with both preStable and probStable set to true, and a call to preStabilize() refines R of s_1, s_2 and s_3 to: $R(s_1) = \{s_1, s_2\}, R(s_2) = \{s_2\}, R(s_3) = \{s_3\}$. A final vacuous call to probStabilize() terminates the computation. Hence, R_{psim} is as follows:

$R_{\mathrm{psim}}(s_1) = \{s_1, s_2\}$ $\quad$ $R_{\mathrm{psim}}(s_2) = \{s_2\}$
$R_{\mathrm{psim}}(s_3) = \{s_3\}$ $\quad$ $R_{\mathrm{psim}}(x_1) = R_{\mathrm{psim}}(x_5) = \{x_1, x_3, x_5\}$
$R_{\mathrm{psim}}(x_3) = \{x_3\}$ $\quad$ $R_{\mathrm{psim}}(x_2) = R_{\mathrm{psim}}(x_4) = \{x_2, x_4\}$
$R_{\mathrm{psim}}(x_6) = \{x_6, x_3\}$ $\quad$ $R_{\mathrm{psim}}(t) = R_{\mathrm{psim}}(u) = R_{\mathrm{psim}}(v) = R_{\mathrm{psim}}(w) = \Sigma$

In particular, we have that s_2 simulates s_1 while s_1 does not simulate s_2, and s_2 does not simulate s_3. $\square$

6.1 Correctness and Complexity

The correctness of PSim comes as a consequence of Theorems 3.1 and 5.3 and the fact that the procedures in Figure 6 correctly stabilize the preorders R and $\mathcal{R}$.

Lemma 6.2 (Correctness of preStabilize()) *Let $\langle R, \mathcal{R}\rangle \in \mathrm{PreOrd}(\Sigma) \times \mathrm{PreOrd}(\mathrm{Distr})$ and $\langle R', \mathcal{R}'\rangle$ be the pair of preorders at the exit of a call to* preStabilize($\langle \mathcal{R}, R\rangle$). *Then, $\mathcal{R}' = \mathcal{R}$ and $R' \subseteq R$ is such that for any $s, t \in \Sigma$, $d \in \mathrm{Distr}$, $a \in Act$, if $s \xrightarrow{a} d$ and $t \in R'(s)$ then $t \xrightarrow{a} \mathcal{R}(d)$, i.e., $\langle \mathcal{R}', R'\rangle$ is* pre-*complete.*

Proof preStabilize($\langle \mathcal{R}, R\rangle$) does not modify the distribution preorder $\mathcal{R}$, hence $\mathcal{R}' = \mathcal{R}$. Consider $s, t \in \Sigma$, $d \in \mathrm{Distr}$ and $a \in Act$ with $s \xrightarrow{a} d$ and $t \in R'(s)$. Since R' is a refinement of the initial state preorder R, we have that $a \in \mathrm{out}(t)$, so that there exists a distribution e such that $t \xrightarrow{a} e$. Moreover, $e \in \mathcal{R}(d)$, otherwise t would belong to $\mathrm{Remove}_a(d)$, which instead is empty, because at the exit of preStabilize($\langle \mathcal{R}, R\rangle$) every remove set is empty. By Lemma 5.1, $\langle \mathcal{R}', R'\rangle$ is therefore pre-complete. $\square$

Lemma 6.3 (Correctness of probStabilize()) *Let $\langle R, \mathcal{R}\rangle \in \mathrm{PreOrd}(\Sigma) \times \mathrm{PreOrd}(\mathrm{Distr})$ and $\langle R', \mathcal{R}'\rangle$ be the pair of preorders at the exit of a call to* probStabilize($\langle R, \mathcal{R}\rangle$)*, where $\mathcal{R} = \leq_{R \cup \mathrm{Deleted}}$. Then $R' = R$ and $\mathcal{R}' \subseteq \mathcal{R}$ is such that for any $d, e \in \mathrm{Distr}$, if $e \in \mathcal{R}'(d)$, then $d \leq_R e$, i.e., $\langle R', \mathcal{R}'\rangle$ is* prob-*complete.*

Proof probStabilize($\langle R, \mathcal{R}\rangle$) does not modify the state preorder R, hence $R' = R$. Consider $e \in \mathcal{R}'(d)$, so that $e \in \mathcal{R}(d)$. By hypothesis, we have that the maximal flow in the network $\mathcal{N}(d, e, R \cup \mathrm{Deleted})$ is 1. If $\mathrm{Deleted} \cap (\mathrm{supp}(d) \times \mathrm{supp}(e)) = \varnothing$, then $\mathcal{N}(d, e, R \cup \mathrm{Deleted}) = \mathcal{N}(d, e, R)$, and therefore $d \leq_R e$. Otherwise, the fact that $e \in \mathcal{R}'(d)$ means that probStabilizes() has (repeatedly) checked (at line 4) that $\mathcal{N}(d, e, R)$ remained 1, so that $d \leq_R e$. Thus, by Lemma 5.1, $\langle R', \mathcal{R}'\rangle$ is prob-complete. $\square$

Theorem 6.4 (Correctness of PSim) *Let $\mathcal{S}$ be a finite PLTS. Then, $\mathrm{PSim}(\mathcal{S})$ always terminates with output $\langle R_{\mathrm{psim}}, \leq_{R_{\mathrm{psim}}}\rangle$.*

Proof The fact that PSim terminates on a finite input PLTS depends on the fact that at each iteration probStabilize() and/or preStabilize() refine $\mathcal{R}$ and/or R. The correctness of the output comes from Theorems 3.1 and 5.3 together with the above Lemmata 6.2 and 6.3. $\square$

Given a PLTS $\mathcal{S} = \langle \Sigma, Act, \rightarrow\rangle$, the complexity bounds of $\mathrm{PSim}(\mathcal{S})$ are given in terms of the following sizes:

- $|\mathrm{Distr}| = |\bigcup_{a \in Act} \mathrm{post}_a(\Sigma)|$ is the number of distributions appearing as target of some transition in $\mathcal{S}$. Also, the number of edges in $\mathcal{S}$ is $|{\rightarrow}| \leq |\Sigma||\mathrm{Distr}|$. Let us notice that $|\mathrm{Distr}| \leq |{\rightarrow}|$.
- Let us define

$$p \triangleq \sum_{d \in \mathrm{Distr}} |\,\mathrm{supp}(d)| \;\text{ and }\; m \triangleq \sum_{a \in Act} \sum_{s \in \Sigma} \sum_{d \in \mathrm{post}_a(s)} (|\,\mathrm{supp}(d)| + 1).$$

  Thus, p represents the full size of $\mathrm{Distr} = \mathrm{post}(\Sigma)$, being the number of states that appear in the support of some distribution in Distr. On the other hand, m represents

the number of transitions in the PLTS $\mathcal{S}$, that is, the number of transitions "from states to states", where a "state transition" (s,t) is taken into account when $s \xrightarrow{a} d$ and $t \in \mathrm{supp}(d)$. Notice that both $p, m \leq |\Sigma||\,\mathrm{Distr}\,|$. The key point to remark here is that $p \leq m$, since the "states" of $\mathcal{S}$ are always less than or equal the "state transitions" in $\mathcal{S}$.

Lemma 6.5 (Complexity of SMF) *All the calls to* SMF$(d, e, ...)$ *relative to a given pair of distributions* d, e*, including the first call* Init_SMF(d, e, R)*, overall take* $O(|\,\mathrm{supp}(d)||\,\mathrm{supp}(e)|^2)$ *time.*

Proof See [18, Lemma 4.4]. □

Theorem 6.6 (Complexity of PSim**)** *Let* $\mathcal{S}$ *be a finite PLTS.* $\mathrm{PSim}(\mathcal{S})$ *runs in* $O(|\Sigma|(p^2 + |{\rightarrow}|))$*-time and* $O(p^2 + (|\Sigma| + |\mathrm{Distr}|)|{\rightarrow}|)$*-space.*

Proof Let us first discuss how we represent the input PLTS $\mathcal{S}$. Let $s_1, \ldots, s_n$ be a fixed enumeration of the states in Σ, let $a_1, \ldots, a_k$ be a fixed enumeration of the labels in Act and let $d_1, \ldots, d_m$ be a fixed enumeration of the distributions in $\mathrm{Distr} = \bigcup_{a \in Act} \mathrm{post}_a(\Sigma)$. We assume that any distribution $d \in \mathrm{Distr}$ is represented as a record with two components:

- an array $[\mathrm{pre}_{a_1}, \ldots, \mathrm{pre}_{a_k}]$, where the i-th entry pre_{a_i} is a pointer to the list of states s such that $s \xrightarrow{a_i} d$.
- a list of pairs $(s_1, d(s_1)), \ldots, (s_r, d(s_r))$ that enumerates the states in the support of d together with their (non-zero) probability.

We observe the following useful properties:

- $|\mathrm{Distr}| \leq p$ and $|{\rightarrow}| \leq m$, so that $|\mathrm{Distr}| \leq m$ and $|\mathrm{Distr}||{\rightarrow}| \leq m^2$.
- $|\Sigma| \leq |{\rightarrow}|$ because any transition has some source state. Moreover, $|\Sigma| \leq m$, so that $|\Sigma||{\rightarrow}| \leq m^2$.

Time Complexity. The time complexities of the functions called by PSim are as follows, where the cost of Initialize() refers to the implementation given in Figure 7.
Initialize() takes $O(p^2 + (|\Sigma| + |\mathrm{Distr}|)|{\rightarrow}|)$ time, as detailed below:

- The initialization of the $\mathrm{in}(\cdot)$ sets takes $O(|\mathrm{Distr}||Act|)$ time since, given a distribution d, the test $\mathrm{pre}_a(d) \neq \varnothing$ is done in $O(1)$ since it is sufficient to test if the a-component of the array stored in the distribution d is non-null.
- Initialization of R and $\mathcal{R}$: lines 5-12 takes $O(|{\rightarrow}| + |\Sigma||{\rightarrow}|)$ time[1], while the cost of the calls to Init_SMF() will be considered together with the global cost of all the calls to the function SMF() in probStabilize().
- Initialization of the Count table: the cost of lines 14-21 is $O(|\mathrm{Distr}||{\rightarrow}|)$ since, as discussed above, the test $\mathrm{pre}_a(e) \neq \varnothing$ takes constant time.
- Initializing the Remove sets (lines 22-26) takes $O(|\Sigma||{\rightarrow}|)$ time.
- Initializing the stability flags (lines 27-29) takes $O(|{\rightarrow}|)$ time.
- Initialization of the Listener sets: line 30 takes $O(|\Sigma|^2)$, while the cost of lines 31-32 is $\sum_{d \in \mathrm{Distr}} \sum_{e \in \mathrm{Distr}} |\,\mathrm{supp}(d)||\,\mathrm{supp}(e)| \leq p^2$.

All the calls to preStabilize() globally cost $O(|\Sigma||{\rightarrow}|)$ time, similarly to what happens in Henzinger et al.'s simulation algorithm [10]:

[1] A more efficient procedure can be obtained by resorting to a kind of remove sets to avoid the reassignment of false to some entry $R(x, y)$.

- line 4: given a pair (e, a), the overall cost of this line is in $O(|\mathrm{pre}_a(\mathrm{Distr})|)$ since $\mathrm{Remove}_a(e) \subseteq \mathrm{pre}_a(\mathrm{Distr})$ and when the same pair appears as pivot of different iterations, say i and j, the sets $\mathrm{Remove}_a^i(e)$ and $\mathrm{Remove}_a^j(e)$ relative to the different iterations i and i are disjoint. Hence, the overall cost of line 4 is $\sum_{e\in\mathrm{Distr}} \sum_{a\in\mathrm{in}(e)} |\mathrm{pre}_a(\mathrm{Distr})|$, which is in $O(|\Sigma||{\rightarrow}|)$.
- Since the sets $\mathrm{Remove}_a^i(e)$ and $\mathrm{Remove}_a^j(e)$ relative to two different iterations i and j are disjoint, we have that any transition $t \xrightarrow{a} e$ in line 5 is traversed at most $| \cup_i \{w \mid w \in \mathrm{Remove}_a^i(e)\}|$ times, namely, at most $|\{w \mid a \in \mathrm{out}(w)\}|$ times. Hence, the overall cost of lines 5-6-and the test at line 7 is $\sum_{w\in\Sigma} \sum_{a\in\mathrm{out}(w)} |\xrightarrow{a}|$ and therefore is in $O(|\Sigma||{\rightarrow}|)$.
- The body of the if statement at line 7 globally takes $|\Sigma|^2$ time since the pairs (t, w) such that the if-test is positive are globally pairwise disjoint.

The cost of all the calls to probStabilize() is in $O(p^2|\Sigma|)$ time, as detailed below:

- Let us first consider the cost of all the calls to SMF() and Init_SMF(). By Lemma 6.5, for all the pairs (d, e), with $d, e \in \mathrm{Distr}$, the cost of all the calls to SMF() and Init_SMF() is $\sum_{d\in\mathrm{Distr}} \sum_{e\in\mathrm{Distr}} |\mathrm{supp}(d)||\mathrm{supp}(e)|^2$, which is in $O(p^2|\Sigma|)$ since for any distribution e it holds that $|\mathrm{supp}(e)| \leq |\Sigma|$.
- The same pair (d, e) may appear in at most $|\mathrm{supp}(d)||\mathrm{supp}(e)|$ different Listener sets. The set of pairs $(t, w) \in \mathrm{Deleted}$ are pairwise disjoint throughout all the calls to probStabilize(), thus the overall cost of lines 2-3 is $\sum_{d\in\mathrm{Distr}} \sum_{e\in\mathrm{Distr}} |\mathrm{supp}(d)||\mathrm{supp}(e)| = p^2$.
- To estimate the overall cost of lines 5-10 of probStabilize(), observe that the test at line 4 is positive at most once for every pair d, e, because after a positive test e is removed from $\mathcal{R}(d)$ and never put back. Hence, the overall cost of lines 5-10 is $\sum_{d\in\mathrm{Distr}} \sum_{e\in\mathrm{Distr}} (1 + \sum_{b\in\mathrm{in}(e)\cap\mathrm{in}(d)} | \mathrm{pre}_b(e)|)$, which is in $O(|\mathrm{Distr}||{\rightarrow}|)$.
- Summing up, the overall cost of all the calls to probStabilize() is $O(p^2|\Sigma| + |\mathrm{Distr}||{\rightarrow}|)$ time. This bound is in $O(p^2|\Sigma|)$ since $|\mathrm{Distr}| \leq p$ and $|{\rightarrow}| \leq |\Sigma||\mathrm{Distr}|$, so that $|\mathrm{Distr}||{\rightarrow}| \leq p^2|\Sigma|$.

Finally, notice that the test of the main while loop of PSim is performed $|\Sigma|^2$ times since the relation R can be refined at most $|\Sigma|^2$ times. Thus, summing up, it turns out that the time complexity of PSim is in $O(|\Sigma|(|{\rightarrow}| + p^2))$.

Space Complexity. PSim relies on the following data structures:

- the in lists of labels, that take $O(|\mathrm{Distr}||Act|)$ space;
- the networks $\mathcal{N}(d, e, R)$ that are updated at each iteration. According to [18], the space needed to store these networks is $\sum_{d\in\mathrm{Distr}} \sum_{e\in\mathrm{Distr}} | \mathrm{supp}(d)|| \mathrm{supp}(e)| \leq p^2$;
- the two boolean matrixes $\{R(s, t)\}_{s,t\in\Sigma}$ and $\{\mathcal{R}(d, e)\}_{d,e\in\mathrm{Distr}}$ that take, respectively, $O(|\Sigma|^2)$ and $O(|\mathrm{Distr}|^2)$ space.
- the integer table $\{\mathrm{Count}(s, a, e)\}_{e\in\mathrm{Distr},\ a\in\mathrm{in}(e),\ s\in\mathrm{pre}_a(\mathrm{Distr})}$ and the lists of states $\{\mathrm{Remove}_a(e)\}_{e\in\mathrm{Distr},a\in\mathrm{in}(e)}$ take, respectively, $O(|\mathrm{Distr}||{\rightarrow}|)$ and $O(|\Sigma||{\rightarrow}|)$ space;
- the sets $\{\mathrm{Listener}(x, y)\}_{x,y\in\Sigma}$ take p^2 space because, as above, the same pair (d, e) may appear in at most $| \mathrm{supp}(d)|| \mathrm{supp}(e)|$ different Listener sets;
- the set of deleted arcs in Deleted takes $O(|\Sigma|^2)$ space.

Thus, the overall space complexity of PSim is in $O(p^2 + (|\Sigma| + |\mathrm{Distr}|)|{\rightarrow}|)$. □

It is easily seen that $(|\Sigma| + |\mathrm{Distr}|)|{\rightarrow}| \leq m^2$, so that PSim turns out to be more efficient than the most efficient probabilistic simulation algorithm in literature, that is Zhang

et al.'s algorithm [18], that runs in $O(|\Sigma|m^2)$-time and $O(m^2)$-space. Our reduction from the size m to p, that is from the size of the "state transitions" to the size of the "state" space, basically depends on the fact that in Zhang et al.'s algorithm the same test $d \not\leq_R e$ is repeated for every pair of states (s_i, t_i) such that $s_i \in \mathrm{pre}_a(d), t_i \in \mathrm{pre}_a(e)$, whereas in PSim once the test $d \not\leq_R e$ has been performed, every state t_i is removed from $R(s_i)$. Such a difference becomes evident when the input PLTS $\mathcal{S}$ degenerates to a LTS. In this case a call to the function SMF() can be executed in constant time, so that the time complexity of Zhang et al.'s algorithm boils down to $O(|{\rightarrow}|^2)$, whereas in this case PSim runs in $O(|\Sigma||{\rightarrow}|)$-time, essentially reducing to Henzinger et al.'s nonprobabilistic simulation algorithm [10]. As a further key difference, it is worth observing that Zhang et al.'s algorithm relies on a positive strategy that at each iteration i computes the pairs (s_i, t_i) such that $t_i \in R_i(s_i)$, whereas PSim follows a dual, negative, strategy that removes from R_i the pairs (s_i, t_i) such that $t_i \notin R_i(s_i)$.

7 Future Work

We have shown how abstract interpretation can be fruitfully applied in the context of behavioral relations between probabilistic processes. We focused here on bisimulation/simulation relations on PLTSs and we proved how efficient algorithms that compute these behavioral relations can be systematically derived. As future work, we plan to investigate how this abstract interpretation approach can be adapted to characterize the weak variants of bisimulation/simulation and the so-called probabilistic bisimulations/simulations on PLTSs [17]. We also plan to apply a coarsest partition refinement approach to design a "symbolic" version of our PSim simulation algorithm. Analogously to the symbolic algorithm by Ranzato and Tapparo [14, 16] for nonprobabilistic simulation, the basic idea is to symbolically represent the relations R on states and $\mathcal{R}$ on distributions through partitions (of states and distributions) and corresponding relations between blocks of these relations. It is worth noting that this partition refinement approach has been already applied by Zhang [19] to design a space-efficient simulation algorithm for PLTSs. Finally, we envisage to study how the abstract interpretation approach can be related to the logical characterizations of behavioral relations of probabilistic processes studied e.g. in [13].

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

References

1. C. Baier, B. Engelen and M. Majster-Cederbaum. Deciding bisimilarity and similarity for probabilistic processes. *J. Comp. Syst. Sci.*, 60:187-231, 2000.
2. 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*, pp. 238–252, 1977.
3. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Proc. 6th ACM POPL*, pp. 269–282, 1979.
4. S. Crafa and F. Ranzato. Probabilistic bisimulation and simulation algorithms by abstract interpretation. In *Proc. ICALP'11*, LNCS 6756, pp. 295-306, Springer, 2011.
5. J. Desharnais. *Labelled Markov Processes*. PhD thesis, McGill Univ., 1999.
6. J. Desharnais, A. Edalat and P. Panangaden. Bisimulation for labelled Markov processes. *Information and Computation*, 179:163-193, 2002.
7. R. Giacobazzi and E. Quintarelli. Incompleteness, counterexamples and refinements in abstract model checking. In *Proc. 8th SAS*, LNCS 2126, pp. 356-373, 2001.

8. R. Giacobazzi and F. Ranzato. Refining and compressing abstract domains. In *Proc. 24th ICALP*, LNCS 1256, pp. 771-781, Springer, 1997.
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. M.R. Henzinger, T.A. Henzinger and P.W. Kopke. Computing simulations on finite and infinite graphs. In *Proc. 36th FOCS*, pp. 453–462, 1995.
11. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1-28, 1991.
12. R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM J. Comput.*, 16(6):973-989, 1987
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. F. Ranzato and F. Tapparo. A new efficient simulation equivalence algorirthm. In *Proc. IEEE LICS'07*, pp. 171-180, 2007.
15. F. Ranzato and F. Tapparo. Generalized strong preservation by abstract interpretation. *J. Logic and Computation*, 17(1):157-197, 2007.
16. F. Ranzato and F. Tapparo. An efficient simulation algorithm based on abstract interpretation. *Information and Computation*, 208(1):1-22, 2010.
17. R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *Nordic J. of Computing*, 2(2):250-273, 1995.
18. 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.
19. L. Zhang. A space-efficient probabilistic simulation algorithm. In *Proc. CONCUR'08*, LNCS 5201, pp. 248-263, 2008.
20. L. Zhang. *Decision Algorithms for Probabilistic Simulations*. PhD thesis, Univ. des Saarlandes, 2009.