

Deciding Probabilistic Bisimilarity over Infinite-State Probabilistic Systems^{*}

Tomáš Brázdil, Antonín Kučera, and Oldřich Stražovský

Faculty of Informatics, Masaryk University,
Botanická 68a, 60200 Brno, Czech Republic,
{brazdil, tony, strazovsky}@fi.muni.cz

Abstract. We prove that probabilistic bisimilarity is decidable over probabilistic extensions of BPA and BPP processes. For normed subclasses of probabilistic BPA and BPP processes we obtain polynomial-time algorithms. Further, we show that probabilistic bisimilarity between probabilistic pushdown automata and finite-state systems is decidable in exponential time. If the number of control states in PDA is bounded by a fixed constant, then the algorithm needs only polynomial time.

1 Introduction

Theory of probabilistic systems is a formal basis for modeling and verification of systems that exhibit some kind of uncertainty [20, 18]. For example, this uncertainty can be caused by unpredictable errors (such as message loss in unreliable channels), randomization (as in randomized algorithms), or simply underspecification in some of the system components. The semantics of probabilistic systems is usually defined in terms of homogeneous Markov chains or Markov decision processes. The former model allows to specify just probabilistic behavioural aspects, while the latter one combines the paradigms of nondeterministic and probabilistic choice. In this paper we consider a generalized model of [23] which subsumes both of the aforementioned formalisms and also “ordinary” non-probabilistic systems. As we shall see, this means that the majority of our results generalize the ones which were previously established for non-probabilistic infinite-state systems.

Methods for formal verification of probabilistic systems follow the two standard approaches of *model-checking* and *equivalence-checking*. In the model-checking approach, desired properties of the system are specified as a formula of a suitable probabilistic temporal logic (such as PCTL or PCTL* [7]), and then it is shown that the system satisfies the formula. In the equivalence-checking approach, one proves that the verified system is semantically equivalent to its *specification*, which is another probabilistic system. Here the notion of semantic equivalence can be formally captured in many ways. Most of the existing equivalences are probabilistic extensions of their non-probabilistic counterparts. One consequence of this is that various variants of *probabilistic bisimilarity* [21] play a very important role in this setting.

^{*} The work has been supported by the Grant Agency of the Czech Republic, grant No. 201/03/1161.

The state of the art: Algorithmic support for formal verification of probabilistic systems has so far been limited to finite-state systems [11, 15, 3, 12, 16, 6, 20, 13, 10]. Only recently, model-checking algorithms for infinite-state models of fully probabilistic lossy channel systems [17, 5, 1, 2, 22] and fully probabilistic pushdown automata [14] appeared. However, the authors are not aware of any results about equivalence-checking with probabilistic infinite-state systems.

Our Contribution: In the first part of our work we consider probabilistic extensions of the well-known families of BPA and BPP processes, which are denoted pBPA and pBPP, respectively. We have chosen a general extension based on the idea that process constants have finitely many basic transitions of the form $X \rightarrow \mu$ where μ is a probability distribution over pairs of the form (a, α) , where a is an action and α a sequence of BPA/BPP constants (in the case of BPP, sequences of constants are considered modulo commutativity and thus the concatenation operator models a simple form of parallel composition without synchronization). Basic transitions then define transitions performable from sequences of constants by adjusting the target distributions accordingly. Hence, our model subsumes the original (non-probabilistic) BPA and BPP, which can be understood as those subclasses of pBPA and pBPP where all distributions used in basic transitions are Dirac. Moreover, pBPA also subsumes a fully probabilistic extension of BPA. We prove that probabilistic bisimilarity (both in its combined and non-combined variant) is decidable for pBPA and pBPP processes. Moreover, for normed subclasses of pBPA and pBPP we have polynomial-time algorithms. Our results generalize the ones for non-probabilistic BPA and BPP by extending and adapting the original notions and proofs. Intuitively, such an extension is possible because probabilistic bisimilarity has similar algebraic and transfer properties as “ordinary” non-probabilistic bisimilarity. These properties can be reformulated and reproved in the probabilistic setting by incorporating some ideas for finite-state systems (e.g., the use of geometrical algorithms for finitely-generated convex spaces in the style of [10]), and there are also new techniques for handling problems which are specific to infinite-state probabilistic systems. After reestablishing these crucial properties, we can basically follow the original proofs because they mostly rely just on algebraic arguments. This can be seen as a nice evidence of the robustness of the original ideas.

In Section 4 we concentrate on checking probabilistic bisimilarity between processes of probabilistic pushdown automata (pPDA) and probabilistic finite-state automata. Our results are based on a generic method for checking semantic equivalences between PDA and finite-state processes proposed in [19]. This method clearly separates generic arguments (applicable to every behavioral equivalence which is a right PDA congruence in the sense of Definition 31) from the equivalence-specific parts that must be supplied for each behavioral equivalence individually. This method works also in the probabilistic setting, but the application part would be unnecessarily long and complicated if we used the original scheme of [19]. Therefore, the generic part of the method is first adjusted into a more “algebraic” form which simplifies some of the crucial steps. The method is then used to prove that probabilistic bisimilarity is decidable between pPDA and finite-state processes in exponential time. Actually, this algorithm is *polynomial* if the number of pPDA control states is bounded by a fixed constant (in particular, this holds for pBPA).

In all sections we tried to avoid repeating of the known things as much as possible; unfortunately, this inevitably means that the material is not completely self-contained. We did our best to provide enough information and intuition so that our presentation is understandable even for a reader who is not familiar with “classical” results on BPA and BPP presented in [9], and who does not know anything about the recent results of [19]. We always clearly mark the results which are not to be considered as a part of this work.

The results presented in this paper generate many questions. Some of them are summarized in Section 5. Due to space constraints, most proofs had to be omitted. These can be found in a full version of this paper [8].

2 Basic Definitions

We start by recalling basic notions of probability theory. A *discrete probability measure* (or *distribution*) over a set X is a function $\mu : 2^X \rightarrow \mathbb{R}^{\geq 0}$ such that, for each countable collection $\{X_i\}_{i \in I}$ of pairwise disjoint subsets of X , $\mu(\bigcup_{i \in I} X_i) = \sum_{i \in I} \mu(X_i)$, and moreover $\mu(X) = 1$. The set of all distributions over a set X is denoted $Disc(X)$. A *Dirac* distribution is a distribution which assigns 1 to exactly one object. A *rational* distribution is a distribution which assigns a rational number to each object. For every $\mu \in Disc(X)$ we define its *support*, denoted $supp(\mu)$, as the set $\{x \in X \mid \mu(x) > 0\}$. A *discrete probability space* is a pair (X, μ) where X is a set called *sample space* and μ a distribution over X .

The underlying semantics of probabilistic systems is usually defined in terms of labelled Markov chains or labelled Markov decision processes, depending mainly on whether the considered system is sequential or parallel. Since some of our results are applicable to both sequential and parallel probabilistic systems, we use a more general formalism of [23] which subsumes the aforementioned models.

Definition 1. An action-labelled probabilistic transition system (or *just* transition system) is a triple $\mathcal{S} = (S, Act, D)$ where S is a finite or countably infinite set of states, $Act \neq \emptyset$ is a set of actions, and $D \subseteq S \times Disc(Act \times S)$ is a finite or countably infinite transition relation. An element $(s, \mu) \in D$ is called a transition and alternatively denoted by $s \rightarrow \mu$. A (probabilistic) process is a state of some transition system.

For the rest of this section, let us fix a probabilistic transition system $\mathcal{S} = (S, Act, D)$.

We say that $t \in S$ is *reachable* from $s \in S$ under a word $w = a_1 \cdots a_k \in Act^*$, written $s \xrightarrow{w} t$ (or simply $s \rightarrow^* t$ if w is irrelevant), if there is a finite sequence $s = s_0, s_1, \dots, s_k = t$ of states such that $(s_i, \mu_i) \in D$ and $\mu_i(a_{i+1}, s_{i+1}) > 0$ for each $0 \leq i < k$. For each transition $s \rightarrow \mu$ we define the set of μ -successors of s by $succ(s, \mu) = \{t \in S \mid \mu(a, t) > 0 \text{ for some } a \in Act\}$. For each state s we define the set of successors by $succ(s) = \bigcup_{s \rightarrow \mu} succ(s, \mu)$.

For every $s \in S$, let $D(s) = \{(s, \mu) \in D\}$ be the set of transitions that leave from s . Every distribution $\sigma \in Disc(D(s))$ determines a unique distribution $\mu_\sigma \in Disc(Act \times S)$ defined for each $(a, t) \in Act \times S$ as $\mu_\sigma(a, t) = \sum_{(s, \mu) \in D(s)} \sigma(s, \mu) \mu(a, t)$. Note that the sum $\sum_{(s, \mu) \in D(s)} \sigma(s, \mu) \mu(a, t)$ exists because the set $D(s)$ is finite or countably infinite. A *combined transition relation* $D_C \subseteq S \times Disc(Act \times S)$ is defined by

$D_C = \{(s, \mu_\sigma) \mid s \in S, \sigma \in \text{Disc}(D(s))\}$. We write $s \rightarrow_C \mu$ instead of $(s, \mu) \in D_C$. Obviously, introducing combined transitions does not influence the reachability relation. However, a single state can have uncountably many outgoing combined transitions. Therefore, the triple (S, Act, D_C) cannot be generally seen as a transition system in the sense of Definition 1.

Semantic equivalence of probabilistic processes can be formally captured in many ways. Existing approaches extend the ideas originally developed for non-probabilistic processes, and the resulting notions have similar properties as their non-probabilistic counterparts. One consequence of this is that probabilistic extensions of *bisimulation-like equivalences* play a very important role in this setting. First we introduce some useful notions and notation. For the rest of this section, let us fix a transition system $S = (S, \text{Act}, D)$. Let $E \subseteq S \times S$ be an equivalence relation. We say that two distributions $\mu, \nu \in \text{Disc}(\text{Act} \times S)$ are *equivalent according to E* , denoted $\mu E \nu$, iff for each $a \in \text{Act}$ and each equivalence class $C \in S/E$ we have that $\mu(a, C) = \nu(a, C)$, where $\mu(a, C) = \sum_{s \in C} \mu(a, s)$. In other words, the equivalence E (defined on states) determines a unique equivalence on distributions that is also denoted by E .

Definition 2. Let E be an equivalence on S , and let $(s, t) \in S \times S$. We say that (s, t) expands in E iff

- for each $s \rightarrow \mu$ there is $t \rightarrow \nu$ such that $\mu E \nu$;
- for each $t \rightarrow \mu$ there is $s \rightarrow \nu$ such that $\mu E \nu$.

A relation $R \subseteq S \times S$ expands in E iff each $(s, t) \in R$ expands in E . An equivalence E on S is a probabilistic bisimulation iff E expands in E . We say that $s, t \in S$ are bisimilar, written $s \sim t$, iff they are related by some probabilistic bisimulation.

The notions of combined expansion, combined bisimulation, and combined bisimilarity (denoted \sim_C), are defined in the same way as above, using \rightarrow_C instead of \rightarrow .

It can be shown that probabilistic bisimilarity is a proper refinement of combined probabilistic bisimilarity (we refer to [23] for a more detailed comparison of the two equivalences). Since most of our results are valid for both of these equivalences, we usually refer just to “bisimilarity” and use the \rightarrow and \simeq symbols to indicate that a given construction works both for \rightarrow and \sim , and for \rightarrow_C and \sim_C , respectively. The word “expansion” is also overloaded in the rest of this paper.

Bisimilarity can also be used to relate processes of different transition systems by considering bisimulations on the disjoint union of the two systems.

Given a binary relation R over a set X , the symbol \equiv_R denotes the least equivalence on X subsuming R . We start with a sequence of basic observations.

Lemma 3. Let R_1, R_2 be binary relations on S such that $R_1 \subseteq R_2$. Then for all $\mu, \nu \in \text{Disc}(\text{Act} \times S)$ we have that if $\mu \equiv_{R_1} \nu$, then also $\mu \equiv_{R_2} \nu$.

Lemma 4. Let R be a relation on S and E be an equivalence on S . If R expands in E , then \equiv_R expands in E .

An immediate corollary to the previous lemmas is the following:

Corollary 5. \simeq is a bisimulation.

Proof. \simeq expands in \equiv_{\simeq} by Lemma 3, hence \equiv_{\simeq} expands in \equiv_{\simeq} by Lemma 4. Therefore, \equiv_{\simeq} is a bisimulation and $\equiv_{\simeq} \subseteq \simeq$. \square

Lemma 6. *Suppose that $(s, t) \in E$ where E is a bisimulation on S . If $s \xrightarrow{w} s'$ for some $w \in Act^*$, then there is $t \xrightarrow{w} t'$ such that $(s', t') \in E$.*

2.1 Approximating bisimilarity

Bisimilarity can be approximated by a family of equivalences \simeq_i , $i \in \mathbb{N}_0$, defined inductively as follows:

- $\simeq_0 = S \times S$;
- \simeq_{i+1} consists of those $(s, t) \in \simeq_i$ which expand in \simeq_i .

Clearly $\simeq \subseteq \bigcap_{i=0}^{\infty} \simeq_i$, and the other inclusion holds if each process $s \in S$ is *finitely branching*, i.e., the set $\{\mu \mid s \rightarrow \mu\}$ is finite. It is worth mentioning that this observation can be further generalized.

Lemma 7. *Let $s, t \in S$, and let us assume that each state t' reachable from t is finitely branching (i.e., s can still be infinitely-branching). Then $s \simeq t$ iff $s \simeq_i t$ for each $i \in \mathbb{N}_0$.*

Lemma 7 can be seen as a generalization of a similar result for non-probabilistic processes and strong bisimilarity presented in [4]. Also note that Lemma 7 does not impose any restrictions on distributions which can have an infinite support.

Definition 8. *We say that a process $s \in S$ is well-defined if s is finitely branching and for each transition $s \rightarrow \mu$ we have that μ is a rational distribution with a finite support.*

For example, pBPA, pBPP, and pPDA processes which are introduced in next sections are well-defined.

Lemma 9. *Let us assume that Act is finite, and let $s, t \in S$ be well-defined states. Let E be an equivalence over $\text{succ}(s) \cup \text{succ}(t)$ (represented as a finite set of its elements). The problem if (s, t) expands in E^1 is decidable in time polynomial in $|D(s)| + |D(t)|$. Here*

$$|D(s)| = \sum_{s \rightarrow \mu} \sum_{\substack{(a,t) \in Act \times S \\ \mu(a,t) > 0}} |(\mu(a,t), a, t)|$$

where $|(\mu(a,t), a, t)|$ is the length of the corresponding binary encoding of the triple $(\mu(a,t), a, t)$ (note that $\mu(a,t)$ is a rational number).

A direct corollary to Lemma 7 and Lemma 9 is the following:

Corollary 10. *Let us assume that Act is finite and each $s \in S$ is well-defined. Then $\not\sim$ over $S \times S$ is semidecidable.*

¹ Strictly speaking, we consider expansion in $E \cup \{(s, s) \mid s \in S\}$ because E is not an equivalence over S (which is required by Definition 2).

3 Deciding Bisimilarity over pBPA and pBPP Processes

In this section we show that bisimilarity is decidable over pBPA and pBPP processes, which are probabilistic extensions of the well-known process classes BPA and BPP [9]. Moreover, we also show that bisimilarity over normed subclasses of pBPA and pBPP is decidable in polynomial time.

Let $\mathcal{S} = (S, Act, D)$ be a transition system, and let “ \cdot ” be a binary operator on S . For every $R \subseteq S \times S$, the symbol $\overset{R}{\equiv}$ denotes the least congruence over S wrt. “ \cdot ” subsuming R .

Lemma 11. *Let $R \subseteq S \times S$, and let $Pre(R)$ be the least set such that $R \subseteq Pre(R)$, and if $(s, t) \in Pre(R)$ then also $(su, tu), (us, ut) \in Pre(R)$ for every $u \in S$. Then $\overset{Pre(R)}{\equiv} = \overset{R}{\equiv}$.*

Now we formulate three abstract conditions which guarantee the semidecidability of \simeq over $S \times S$. As we shall see, pBPA and pBPP classes satisfy these conditions.

1. For every finite relation $R \subseteq S \times S$ we have that if R expands in $\overset{R}{\equiv}$, then $\overset{R}{\equiv} \subseteq \simeq$.
2. There is a finite relation $\mathcal{B} \subseteq S \times S$ such that $\overset{\mathcal{B}}{\equiv} = \simeq$ over $S \times S$ (\mathcal{B} is called a *bisimulation base*).
3. The definition of \mathcal{S} is effective in the following sense: the set of states S is recursively enumerable, each state $s \in S$ is well-defined, and the problem if $s = t \cdot u$ for given $s, t, u \in S$ is semidecidable.

Lemma 12. *If the three conditions above are satisfied, then \simeq over $S \times S$ is semidecidable (and thus decidable by applying Corollary 10).*

Now we formally introduce pBPA and pBPP processes. Let $N = \{X, Y, \dots\}$ be a countably infinite set of *constants* and $Act = \{a, b, \dots\}$ a countably infinite set of actions. The elements of N^* are denoted α, β, \dots , and the empty word by ε .

Let $\mu \in Disc(Act \times N^*)$ be a distribution. For each $\alpha \in N^*$, the symbol $\mu\alpha$ denotes the distribution such that $(\mu\alpha)(a, \beta\alpha) = \mu(a, \beta)$, and $(\mu\alpha)(a, \gamma) = 0$ if α is not a suffix of γ .

Definition 13. *A pBPA (pBPP) system Δ is a finite set of rules of the form $X \rightarrow \mu$ where $\mu \in Disc(Act \times N^*)$ is a rational distribution with a finite support.*

The sets of all constants and actions occurring in Δ are denoted $N(\Delta)$ and $Act(\Delta)$, respectively. We require that for each $X \in N(\Delta)$ there is at least one rule of the form $X \rightarrow \mu$ in Δ .

To Δ we associate the transition system $\mathcal{S}_\Delta = (N(\Delta)^*, Act(\Delta), D)$ where the transitions of D are determined as follows:

$$\frac{X \rightarrow \nu \in \Delta}{X\alpha \rightarrow \nu\alpha} \alpha \in N(\Delta)^*$$

The elements of $N(\Delta)^*$ are called pBPA processes (of Δ).

pBPP systems and processes are defined in the same way, but the elements of $N(\Delta)^*$ are understood modulo commutativity (intuitively, this corresponds to an unsynchronized parallel composition of constants).

Observe that “ordinary”, i.e., non-probabilistic BPA and BPP systems can be understood as those pBPA and pBPP where all distributions used in basic transitions are Dirac (see Section 2). Moreover, to every pBPA/pBPP system Δ we associate its *underlying* non-probabilistic BPA/BPP system Δ^u defined as follows: for every rule $X \rightarrow \mu \in \Delta$ we add to Δ^u the rules $X \xrightarrow{a} \alpha$ for each $(a, \alpha) \in \text{supp}(\mu)$. If we consider \simeq as a relation on the states of \mathcal{S}_{Δ^u} , we can readily confirm that \simeq is a (non-probabilistic) strong bisimulation; this follows immediately from Lemma 6. However, \simeq is generally *finer* than strong bisimilarity over the states of \mathcal{S}_{Δ^u} .

Definition 14. *Let Δ be a pBPA or pBPP system. A given $X \in N(\Delta)$ is normed if there is some $w \in \text{Act}(\Delta)^*$ such that $X \xrightarrow{w} \varepsilon$. The norm of X , denoted $n(X)$, is the length of the shortest such w . If $X \in N(\Delta)$ is not normed, we put $n(X) = \infty$. We say that Δ is normed if every $X \in N(\Delta)$ is normed.*

Note that $n(\varepsilon) = 0$, and if we adopt the usual conventions for ∞ , then $n(\alpha\beta) = n(\alpha) + n(\beta)$. Also note that bisimilar processes must have the same norm. Transition systems generated by pBPA and pBPP systems are clearly effective in the sense of condition 3 above. Now we check that conditions 1 and 2 are also satisfied. This is where new problems (which are specific to the probabilistic setting) arise.

Lemma 15 (condition 1). *Let Δ be a pBPA or a pBPP system. Let R be a binary relation over $N(\Delta)^*$, and let E be a congruence over $N(\Delta)^*$ where $R \subseteq E$. If R expands in E , then $\stackrel{R}{\equiv} \stackrel{R}{\equiv} \stackrel{R}{\equiv}$ expands in E .*

It follows from Lemma 15 that $\stackrel{R}{\equiv} \subseteq \simeq$ whenever R expands in $\stackrel{R}{\equiv}$.

Corollary 16. *\simeq is a congruence over processes of a given pBPA or pBPP system.*

Proof. \simeq expands in $\stackrel{\simeq}{\equiv}$, hence $\stackrel{\simeq}{\equiv} \subseteq \simeq$ by Lemma 15. □

It remains to check that bisimilarity over pBPA and pBPP processes can be represented by a finite base (condition 2 above).

Lemma 17 (condition 2 for pBPP). *Let Δ be a pBPP system. There is a finite relation $\mathcal{B} \subseteq N(\Delta)^* \times N(\Delta)^*$ such that $\stackrel{\mathcal{B}}{\equiv} = \simeq$ over $N(\Delta)^* \times N(\Delta)^*$.*

Proof. The proof in [9] for (non-probabilistic) BPP relies just on the fact that (non-probabilistic) bisimilarity is a congruence. Due to Corollary 16, we can use the same proof also for pBPP. □

In the case of pBPA, the situation is more complicated. Let $N_n \subseteq N(\Delta)$ be the set of all normed variables, and $N_u = N(\Delta) \setminus N_n$ the set of all unnormed ones.

Lemma 18. *Let $X \in N(\Delta)$ and $\alpha \in N(\Delta)^*$. If $n(X) = \infty$, then $X \simeq X\alpha$.*

Note that due to Lemma 18 we need only ever consider states $\alpha \in N_n^* \cup (N_n^* \times N_u)$, the others being immediately transformed into such a bisimilar state by erasing all symbols following the first infinite-norm variable.

A careful inspection of the construction for non-probabilistic BPA (as presented in [9]) reveals the following:

Proposition 19 (see [9]). *Let Δ be a (non-probabilistic) BPA system. Let $\stackrel{\circ}{=} \subseteq N(\Delta)^* \times N(\Delta)^*$ be an equivalence satisfying the following properties:*

1. *if $\alpha \stackrel{\circ}{=} \beta$ and $\alpha \xrightarrow{w} \alpha'$, then there is $\beta \xrightarrow{w} \beta'$ such that $\alpha' \stackrel{\circ}{=} \beta'$ (note that it implies that $n(\alpha) = n(\beta)$);*
2. *$\stackrel{\circ}{=}$ is a congruence;*
3. *if $\alpha\gamma \stackrel{\circ}{=} \beta\gamma$ for infinitely many pairwise non-equivalent γ 's, then $\alpha \stackrel{\circ}{=} \beta$;*

Then there is a finite base \mathcal{B} such that $\stackrel{\mathcal{B}}{=} = \stackrel{\circ}{=}$ over $N_n^ \cup (N_n^* \times N_u)$.*

So, it suffices to prove that \simeq (when considered as an equivalence over the states of the underlying BPA system Δ^u) satisfies the conditions 1–3 of Proposition 19. The first condition follows immediately from Lemma 6, and the second condition follows from Corollary 16. Condition 3 is proven below, together with one auxiliary result.

Lemma 20. *Let α, β be processes of a pBPA system. If $\alpha \simeq \gamma\alpha$ and $\beta \simeq \gamma\beta$ for some $\gamma \neq \varepsilon$, then $\alpha \simeq \beta$.*

Lemma 21. *Let α, β be processes of a pBPA system. If $\alpha\gamma \simeq \beta\gamma$ for infinitely many pairwise non-bisimilar γ 's, then $\alpha \simeq \beta$.*

An immediate consequence of Proposition 19, Lemma 6, Corollary 16, and Lemma 21, is the following:

Lemma 22 (condition 2 for pBPA). *Let Δ be a pBPA system. There is a finite relation $\mathcal{B} \subseteq N(\Delta)^* \times N(\Delta)^*$ such that $\stackrel{\mathcal{B}}{=} = \simeq$ over $N_n^* \cup (N_n^* \times N_u)$.*

Now we can formulate the first theorem of our paper:

Theorem 23. *Bisimilarity for pBPA and pBPP processes is decidable.*

3.1 Polynomial-time algorithms for normed pBPA and normed pBPP

In this subsection we show that the polynomial-time algorithms deciding (non-probabilistic) bisimilarity over the normed subclasses of BPA and BPP processes (see [9]) can also be adapted to the probabilistic case. We concentrate just on crucial observations which underpin the functionality of these algorithms, and show that they can be reformulated and reproved in the probabilistic setting. We refer to [9] for the omitted parts.

In the probabilistic setting, the polynomial-time algorithms deciding non-probabilistic bisimilarity over normed BPA and normed BPP processes are modified as follows: Given a normed pBPA or normed pBPP system Δ , we run the non-probabilistic

algorithm on the underlying system Δ^u , where the only modification is that *the expansion is considered in the probabilistic transition system \mathcal{S}_Δ (instead of \mathcal{S}_{Δ^u})*. To see that the modified algorithm is again polynomial-time, we need to realize that the problem if a given pair of pBPA or pBPP processes expands in a polynomially computable equivalence is decidable in polynomial time. However, it is a simple consequence of Lemma 9.

Lemma 24. *Let Δ be a pBPA or pBPP system, and E a polynomially computable equivalence over $N(\Delta)^*$. Let α, β be processes of Δ . It is decidable in polynomial time whether (α, β) expands in E .*

The authors have carefully verified that bisimilarity has all the properties which imply the correctness of these (modified) algorithms. Some of the most important observations are listed below; roughly speaking, the original non-probabilistic algorithms are based mainly on the unique decomposition property, which must be reestablished in the probabilistic setting.

A pBPA or pBPP process α is a *prime* iff whenever $\alpha \simeq \beta\gamma$, then either $\beta = \varepsilon$ or $\gamma = \varepsilon$ (note that $\alpha \in N$).

Lemma 25. *Let α, β, γ be processes of a normed pBPA system. Then $\alpha\gamma \simeq \beta\gamma$ implies $\alpha \simeq \beta$.*

Theorem 26. *Every normed pBPA process α decomposes uniquely (up to bisimilarity) into prime components.*

Proof. We can use the same proof as in [9]. It relies on Lemma 25, Corollary 16, and Lemma 6. □

Theorem 27. *Every normed pBPP process decomposes uniquely (up to bisimilarity) into prime components.*

Proof. As in [9]. It relies on Lemma 6. □

Now we have all the “tools” required for adapting the observations about non-probabilistic normed BPA/BPP to the probabilistic setting which altogether imply the following:

Theorem 28. *Bisimilarity is decidable for normed pBPA and normed pBPP processes in polynomial time.*

4 Deciding Bisimilarity between pPDA and pFS Processes

Definition 29. A probabilistic pushdown automaton (pPDA) is a tuple $\Delta = (Q, \Gamma, Act, \delta)$ where Q is a finite set of control states, Γ is a finite stack alphabet, Act is a finite set of actions, and $\delta : (Q \times \Gamma) \rightarrow 2^{Disc(Act \times (Q \times \Gamma^*))}$ is a transition function such that the set $\delta(p, X)$ is finite and each $\mu \in \delta(p, X)$ is a rational distribution with a finite support for all $p \in Q$ and $X \in \Gamma$.

We write $p\alpha$ instead of (p, α) and $pA \rightarrow \mu$ instead of $\mu \in \delta(p, A)$. Let $\nu \in \text{Disc}(\text{Act} \times (Q \times \Gamma^*))$ be a distribution. For each $\beta \in \Gamma^*$, the symbol $\nu\beta$ denotes the distribution such that $(\nu\beta)(a, p\alpha\beta) = \nu(a, p\alpha)$, and $(\nu\beta)(a, p\gamma) = 0$ if β is not a suffix of γ . Each pPDA Δ induces a unique transition system \mathcal{S}_Δ where $Q \times \Gamma^*$ is the set of states, Act is the set of actions, and transitions are given by the following rule:

$$\frac{pX \rightarrow \nu \in \delta}{pX\beta \rightarrow \nu\beta} \beta \in \Gamma^*$$

The states of \mathcal{S}_Δ are called pPDA processes of Δ , or just pPDA processes if Δ is not significant.

Our aim is to show that \simeq between pPDA processes and finite-state processes is decidable in exponential time. For this purpose we adapt the results of [19], where a generic framework for deciding various behavioral equivalences between PDA and finite-state processes is developed. In this framework, the generic part of the problem (applicable to every behavioral equivalence which is a right PDA congruence in the sense of Definition 31) is clearly separated from the equivalence-specific part that must be supplied for each behavioral equivalence individually. The method works also in the probabilistic setting, but the application part would be unnecessarily complicated if we used the original scheme proposed in [19]. Therefore, we first develop the generic part of the method into a more “algebraic” form, and then apply the new variant to probabilistic bisimilarity. The introduced modification is generic and works also for other (non-probabilistic) behavioral equivalences.

For the rest of this section, we fix a pPDA $\Delta = (Q, \Gamma, \text{Act}, \delta)$ of size m and a finite-state system $\mathcal{S} = (F, \text{Act}, D)$ of size n (the size of a given $\mu \in \text{Disc}(\text{Act} \times (Q \times \Gamma^*))$ is defined similarly as in Lemma 9). In our complexity estimations we also use the parameter $z = |F|^{|Q|}$.

We start by recalling some notions and results of [19]. To simplify our notation, we introduce all notions directly in the probabilistic setting. We denote $F_\perp = F \cup \{\perp\}$, where $\perp \notin F$ stands for “undefined”.

Definition 30. For every process $p\alpha$ of Δ we define the set $M_{p\alpha} = \{q \in Q \mid p\alpha \rightarrow^* q\varepsilon\}$. A function $\mathcal{F} : Q \rightarrow F_\perp$ is compatible with $p\alpha$ iff $\mathcal{F}(q) \neq \perp$ for every $q \in M_{p\alpha}$. The class of all functions that are compatible with $p\alpha$ is denoted $\text{Comp}(p\alpha)$.

For every process $p\alpha$ of Δ and every $\mathcal{F} \in \text{Comp}(p\alpha)$ we define the process $p\alpha\mathcal{F}$ whose transitions are determined by the following rules:

$$\frac{p\alpha \rightarrow \mu}{p\alpha\mathcal{F} \rightarrow \mu\mathcal{F}} \mathcal{F} \in \text{Comp}(p\alpha) \qquad \frac{\mathcal{F}(p) \rightarrow \mu}{p\mathcal{F} \rightarrow \mu\mathcal{F}} \mathcal{F} \in \text{Comp}(p\varepsilon)$$

Here $\mu\mathcal{F}$ is a distribution which returns a non-zero value only for pairs of the form $(a, q\beta\mathcal{F})$, where $(\mu\mathcal{F})(a, q\beta\mathcal{F}) = \mu(a, q\beta)$, and $\mu_{\mathcal{F}}$ is a distribution which returns a non-zero value only for pairs of the form $(a, p\mathcal{F}[s/p])$, where $\mu(a, p\mathcal{F}[s/p]) = \mu(a, s)$. Here $\mathcal{F}[s/p] : Q \rightarrow F_\perp$ is the function which returns the same result as \mathcal{F} for every argument except for p where $\mathcal{F}[s/p](p) = s$. In other words, $p\alpha\mathcal{F}$ behaves like $p\alpha$ until the point when the stack is emptied and a configuration of the form $q\varepsilon$ is entered; from that point on, $p\alpha\mathcal{F}$ behaves like $\mathcal{F}(q)$. Note that if $\mathcal{F} \in \text{Comp}(p\alpha)$ and $p\alpha \rightarrow^* q\beta$,

then $\mathcal{F} \in \text{Comp}(q\beta)$. We also put $\text{Stack}(\Delta, F) = \Gamma^* \cup \{\alpha\mathcal{F} \mid \alpha \in \Gamma^*, \mathcal{F} \in (F_\perp)^Q\}$, and $\mathcal{P}(\Delta, F) = \{p\alpha \mid p \in Q, \alpha \in \Gamma^*\} \cup \{p\alpha\mathcal{F} \mid p \in Q, \alpha \in \Gamma^*, \mathcal{F} \in \text{Comp}(p\alpha)\}$.

Definition 31. We say that an equivalence E over $\mathcal{P}(\Delta, F) \cup F$ is a right pPDA congruence (for Δ and S) iff the following conditions are satisfied:

- For every process $p\alpha$ of Δ and all $\varphi, \psi \in \text{Stack}(\Delta, F)$ we have that if $(q\varphi, q\psi) \in E$ for each $q \in M_{p\alpha}$, then also $(p\alpha\varphi, p\alpha\psi) \in E$.
- $(p\mathcal{F}, \mathcal{F}(p)) \in E$ for every $p\mathcal{F} \in \mathcal{P}(\Delta, F)$.

Let R be a binary relation over $\mathcal{P}(\Delta, F) \cup F$. The least right pPDA congruence over $\mathcal{P}(\Delta, F) \cup F$ subsuming R is denoted $\overset{R}{\equiv}_r$. Further, $Rpre(R)$ denotes the least relation over $\mathcal{P}(\Delta, F) \cup F$ subsuming R satisfying the following condition: For every process $p\alpha$ of Δ and all $\varphi, \psi \in \text{Stack}(\Delta, F)$ we have that if $(q\varphi, q\psi) \in Rpre(R)$ for each $q \in M_{p\alpha}$, then also $(p\alpha\varphi, p\alpha\psi) \in Rpre(R)$. In general, $\equiv_{Rpre(R)}$ is a proper subset of $\overset{R}{\equiv}_r$; the relationship between $Rpre(R)$ and $\overset{R}{\equiv}_r$ is revealed in the following lemma:

Lemma 32. Let R be a binary relation over $\mathcal{P}(\Delta, F) \cup F$. For every $i \in \mathbb{N}_0$ we define a binary relation R^i over $\mathcal{P}(\Delta, F) \cup F$ inductively as follows: $R^0 = R$, and $R^{i+1} = \equiv_{Rpre(R^i)}$. Then $\overset{R}{\equiv}_r = \bigcup_{i \in \mathbb{N}_0} R^i$.

For the rest of this section, let us fix a right pPDA congruence $\overset{\circ}{\equiv}$ over $\mathcal{P}(\Delta, F) \cup F$ which is decidable for finite-state processes and satisfies the following transfer property: if $s \overset{\circ}{\equiv} t$ and $s \rightarrow^* s'$, then there exists t' such that $t \rightarrow^* t'$ and $s' \overset{\circ}{\equiv} t'$. The following definitions are also borrowed from [19].

Definition 33. Let $\varphi \in \text{Stack}(\Delta, F)$ and $\mathcal{F} : Q \rightarrow F_\perp$. We write $\varphi \overset{\circ}{\equiv} \mathcal{F}$ iff for all $p \in Q$ we have that if $\mathcal{F}(p) \neq \perp$, then $p\varphi \overset{\circ}{\equiv} \mathcal{F}(p)$.

Further, for every relation $K \subseteq \text{Stack}(\Delta, F) \times (F_\perp)^Q$ we define the set $I(K)$ of K -instances as follows: $I(K) = \{(p\varphi, \mathcal{F}(p)) \mid (\varphi, \mathcal{F}) \in K, \mathcal{F}(p) \neq \perp\}$.

Definition 34. Let $K = \{(\varepsilon, \mathcal{F}) \mid \varepsilon \overset{\circ}{\equiv} \mathcal{F}\} \cup \{(\mathcal{G}, \mathcal{F}) \mid \mathcal{G} \overset{\circ}{\equiv} \mathcal{F}\} \cup K'$ where $K' \subseteq \Gamma \times (F_\perp)^Q \cup ((\Gamma \times (F_\perp)^Q) \times (F_\perp)^Q)$. (That is, K' consists of (some) pairs of the form (X, \mathcal{F}) and $(X\mathcal{G}, \mathcal{F})$.) We say that K is well-formed iff K satisfies the following conditions:

- if $(X\mathcal{G}, \mathcal{F}) \in K$ and $\mathcal{F}(p) \neq \perp$, then $\mathcal{G} \in \text{Comp}(pX)$;
- if $(X, \mathcal{F}) \in K$ (or $(X\mathcal{G}, \mathcal{F}) \in K$) and $(\mathcal{F}, \mathcal{H}) \in K$, then also $(X, \mathcal{H}) \in K$ (or $(X\mathcal{G}, \mathcal{H}) \in K$, resp.).

It is clear that there are only finitely many well-formed sets, and that there exists the greatest well-formed set G whose size is $\mathcal{O}(|\Gamma| \cdot |F|^{2 \cdot |Q|})$. Observe that G is effectively constructible because $\overset{\circ}{\equiv}$ is decidable for finite-state processes.

Intuitively, well-formed sets are finite representations of certain infinite relations between processes of $\mathcal{P}(\Delta, F)$ and F , which are “generated” from well-formed sets using the rules introduced in our next definition:

Definition 35. Let K be a well-formed set. The closure of K , denoted $Cl(K)$, is the least set L satisfying the following conditions:

- (1) $K \subseteq L$;
- (2) if $(\alpha\mathcal{G}, \mathcal{F}) \in L$, $(\varepsilon, \mathcal{G}) \in K$, and $\alpha \neq \varepsilon$, then $(\alpha, \mathcal{F}) \in L$;
- (3) if $(\alpha\mathcal{G}, \mathcal{F}) \in L$, $(\mathcal{H}, \mathcal{G}) \in K$, and $\alpha \neq \varepsilon$, then $(\alpha\mathcal{H}, \mathcal{F}) \in L$;
- (4) if $(\alpha\mathcal{G}, \mathcal{F}) \in L$, $(X, \mathcal{G}) \in K$, and $\alpha \neq \varepsilon$, then $(\alpha X, \mathcal{F}) \in L$;
- (5) if $(\alpha\mathcal{G}, \mathcal{F}) \in L$, $(X\mathcal{H}, \mathcal{G}) \in K$, and $\alpha \neq \varepsilon$, then $(\alpha X\mathcal{H}, \mathcal{F}) \in L$.

Further, we define $Gen(K) = I(Cl(K))$.

Observe that Cl and Gen are monotonic and that $Gen(K) \subseteq \mathcal{P}(\Delta, F) \times F$ for every well-formed set K .

An important property of Gen is that it generates only “congruent pairs” as stated in the following lemma.

Lemma 36. *Let K be a well-formed set. Then $Gen(K) \subseteq \overset{I(K)}{\cong_r}$.*

The following well-formed set is especially important.

Definition 37. *The base \mathcal{B} is defined as follows: $\mathcal{B} = \{(\varepsilon, \mathcal{F}) \mid \varepsilon \doteq \mathcal{F}\} \cup \{(\mathcal{G}, \mathcal{F}) \mid \mathcal{G} \doteq \mathcal{F}\} \cup \{(X, \mathcal{F}) \mid X \doteq \mathcal{F}\} \cup \{(X\mathcal{G}, \mathcal{F}) \mid X\mathcal{G} \doteq \mathcal{F}\}$.*

The importance of \mathcal{B} is clarified in the next lemma.

Lemma 38 (see [19]). *$Gen(\mathcal{B})$ coincides with \doteq over $\mathcal{P}(\Delta, F) \times F$.*

Let (\mathcal{W}, \subseteq) be the complete lattice of all well-formed sets, and let $Exp : \mathcal{W} \rightarrow \mathcal{W}$ be a function satisfying the four conditions listed below:

1. $Exp(\mathcal{B}) = \mathcal{B}$.
2. Exp is monotonic, i.e. $K \subseteq L$ implies $Exp(K) \subseteq Exp(L)$.
3. If $K = Exp(K)$, then $Gen(K) \subseteq \doteq$.
4. The membership to $Exp(K)$ is decidable.

According to condition 1, the base \mathcal{B} is a fixed-point of Exp . We prove that \mathcal{B} is the greatest fixed-point of Exp . Suppose that $K = Exp(K)$ for some well-formed set K . By definition of $Gen(K)$ and condition 3 we have that $I(K) \subseteq I(Cl(K)) = Gen(K) \subseteq \doteq$. Since for each $(\varphi, \mathcal{F}) \in K$ we have that $\mathcal{F}(p) \neq \perp$ implies $p\varphi \doteq \mathcal{F}(p)$, we can conclude that $(\varphi, \mathcal{F}) \in \mathcal{B}$.

Hence, \mathcal{B} can be computed by a simple algorithm which iterates Exp on G until a fixed-point is found. These conditions are formulated in the same way as in [19] except for condition 3 which is slightly different. As we shall see, with the help of the new “algebraic” observations presented above, condition 3 can be checked in a relatively simple way. This is the main difference from the original method presented in [19].

Similarly as in [19], we use finite multi-automata to represent certain infinite subsets of $\mathcal{P}(\Delta, F)$.

Definition 39. *A multi-automaton is a tuple $\mathcal{M} = (S, \Sigma, \gamma, Acc)$ where*

- S is a finite set of states such that $Q \subseteq S$ (i.e. the control states of Δ are among the states of \mathcal{M});
- $\Sigma = \Gamma \cup \{\mathcal{F} \mid \mathcal{F} : Q \rightarrow F_\perp\}$ is the input alphabet (the alphabet has a special symbol for each $\mathcal{F} : Q \rightarrow F_\perp$);

- $\gamma \subseteq S \times \Sigma \times S$ is a transition relation;
- $Acc \subseteq S$ is a set of accepting states.

Every multi-automaton \mathcal{M} determines a unique set

$$\mathcal{L}(\mathcal{M}) = \{pw \mid p \in Q, w \in \Sigma^*, \gamma(p, w) \cap Acc \neq \emptyset\}$$

The following tool will be useful for deciding the membership to $Exp(K)$.

Lemma 40. *Let K be a well-formed set. The relation $R = (\equiv_{Gen(K)} \cap (F \times F))$ is computable in time polynomial in m, n, z . Moreover, for each equivalence class $C \in F/R$ there is a multiautomaton $\mathcal{M}_{K,C}$ accepting the set $C' \subseteq \mathcal{P}(\Delta, F)$ where $C \cup C' \in (\mathcal{P}(\Delta, F) \cup F) / \equiv_{Gen(K)}$. The automaton $\mathcal{M}_{K,C}$ is constructible in time polynomial in m, n, z .*

4.1 Deciding \simeq between pPDA and finite-state processes

We apply the abstract framework presented in the previous section. That is, we show that \simeq is a right pPDA congruence and define an appropriate function Exp satisfying the four conditions given earlier. We start with an auxiliary result.

Lemma 41. *Let R be a binary relation over $\mathcal{P}(\Delta, F) \cup F$. If R expands in $\overset{R}{\equiv}_r$ then $\overset{R}{\equiv}_r \subseteq \simeq$.*

The next lemma follows immediately from Lemma 41.

Lemma 42. *\simeq is a right pPDA congruence.*

Definition 43. *Given a well-formed set K , the set $Exp(K)$ consists of all pairs $(\varphi, \mathcal{F}) \in K$ such that for each $p \in Q$ we have that if $\mathcal{F}(p) \neq \perp$, then $(p\varphi, \mathcal{F}(p))$ expands in $\equiv_{Gen(K)}$.*

Now we verify the four conditions that must be satisfied by Exp . The first condition follows easily from the fact that $Gen(\mathcal{B})$ coincides with \simeq over $\mathcal{P}(\Delta, F) \times F$, because if $(p\varphi, \mathcal{F}(p)) \in I(\mathcal{B})$, then $\simeq = \equiv_{Gen(\mathcal{B})}$ over $succ(p\varphi) \cup succ(\mathcal{F}(p))$. The second condition is obvious.

Lemma 44. *$Exp(K) = K$ implies $\equiv_{Gen(K)} \subseteq \simeq$.*

Proof. $Exp(K) = K$ implies that each pair of $I(K)$ expands in $\equiv_{Gen(K)}$. But then each pair of $I(K)$ expands in $\overset{I(K)}{\equiv}_r$ by Lemma 3 and Lemma 36. Thus, $\equiv_{Gen(K)} \subseteq \overset{I(K)}{\equiv}_r \subseteq \simeq$ by Lemma 41. \square

Lemma 45. *$Exp(K)$ is computable in time polynomial in m, n, z .*

Proof. Let $(p\alpha, \mathcal{F}(p)) \in I(K)$ and $U = succ(p\alpha) \cup succ(\mathcal{F}(p))$. It follows immediately from Lemma 40 that the equivalence relation $\equiv_{Gen(K)} \cap (U \times U)$ can be computed in time polynomial in m, n, z . The claim then follows from Lemma 9. \square

Now we can formulate our next theorem.

Theorem 46. *Probabilistic bisimilarity between pPDA and finite-state processes is decidable in time which is polynomial in m, n, z . That is, the problem is decidable in exponential time for general pPDA, and in polynomial time for every subclass of pPDA where the number of control states is bounded by some constant (in particular, this applies to pBPA).*

Proof. Let $p\alpha$ be a pPDA process and f a finite-state process. We can assume (w.l.o.g.) that $\alpha = X$ for some $X \in \Gamma$. The algorithm computes the base \mathcal{B} by first computing the greatest well-formed relation G and then iterating Exp until a fixed-point is found. Then, it suffices to find out if there is a pair $(X, \mathcal{F}) \in \mathcal{B}$ such that $\mathcal{F}(p) = f$. Note that this takes time polynomial in m, n, z , because

- G is computable in time polynomial in m, n, z . This is because the size of G is $\mathcal{O}(|\Gamma| \cdot |F|^{2 \cdot |Q|})$ and \simeq over finite-state systems is decidable in polynomial time [10].
- Exp is computable in time polynomial in m, n, z due to Lemma 45.
- The algorithm needs at most $|G|$, i.e., $\mathcal{O}(|\Gamma| \cdot |F|^{2 \cdot |Q|})$ iterations to reach a fixed-point. \square

5 Conclusions

The results presented in this paper show that various forms of probabilistic bisimilarity are decidable over certain classes of infinite-state systems. In particular, this paper advocates the use of algebraic methods which were originally developed for non-probabilistic systems. These methods turn out to be surprisingly robust and can be applied also in the probabilistic setting.

An obvious question is whether the decidability/tractability results for other non-probabilistic infinite-state models can be extended to the probabilistic case. We conjecture that the answer is positive in many cases, and we hope that the results presented in this paper provide some hints and guidelines on how to achieve that. Another interesting question is whether we could do better than in the non-probabilistic case. In particular, undecidability results and lower complexity bounds do not carry over to fully probabilistic variants of infinite-state models (fully probabilistic systems are probabilistic systems where each state s has at most one out-going transition $s \rightarrow \mu$). It is still possible that methods specifically tailored to fully probabilistic models might produce better results than their non-probabilistic counterparts. This also applies to probabilistic variants of other behavioural equivalences, such as trace or simulation equivalence.

References

- [1] P.A. Abdulla, C. Baier, S.P. Iyer, and B. Jonsson. Reasoning about probabilistic channel systems. In *Proceedings of CONCUR 2000*, volume 1877 of *LNCS*, pages 320–330. Springer, 2000.

- [2] P.A. Abdulla and A. Rabinovich. Verification of probabilistic systems with faulty communication. In *Proceedings of FoSSaCS 2003*, volume 2620 of *LNCS*, pages 39–53. Springer, 2003.
- [3] A. Aziz, V. Singhal, F. Balarin, R. Brayton, and A. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In *Proceedings of CAV'95*, volume 939 of *LNCS*, pages 155–165. Springer, 1995.
- [4] J.C.M. Baeten, J.A. Bergstra, and J.W. Klop. On the consistency of Koomen's fair abstraction rule. *TCS*, 51(1):129–176, 1987.
- [5] C. Baier and B. Engelen. Establishing qualitative properties for probabilistic lossy channel systems: an algorithmic approach. In *Proceedings of 5th International AMAST Workshop on Real-Time and Probabilistic Systems (ARTS'99)*, volume 1601 of *LNCS*, pages 34–52. Springer, 1999.
- [6] C. Baier, H. Hermanns, and J. Katoen. Probabilistic weak simulation is decidable in polynomial time. *Information Processing Letters*, 89(3):123–130, 2004.
- [7] A. Bianco and L. de Alfaro. Model checking of probabalistic and nondeterministic systems. In *Proceedings of FST&TCS'95*, volume 1026 of *LNCS*, pages 499–513. Springer, 1995.
- [8] T. Brázdil, A. Kučera, and O. Stražovský. Deciding probabilistic bisimilarity over infinite-state probabilistic systems. Technical report FIMU-RS-2004-06, Faculty of Informatics, Masaryk University, 2004.
- [9] O. Burkart, D. Caucal, F. Moller, and B. Steffen. Verification on infinite structures. *Handbook of Process Algebra*, pages 545–623, 1999.
- [10] S. Cattani and R. Segala. Decision algorithms for probabilistic bisimulation. In *Proceedings of CONCUR 2002*, volume 2421 of *LNCS*, pages 371–385. Springer, 2002.
- [11] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *Proceedings of 29th Annual Symposium on Foundations of Computer Science*, pages 338–345. IEEE, 1988.
- [12] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *JACM*, 42(4):857–907, 1995.
- [13] L. de Alfaro, M.Z. Kwiatkowska, G. Norman, D. Parker, and R. Segala. Symbolic model checking of probabilistic processes using MTBDDs and the Kronecker representation. In *Proceedings of TACAS 2000*, volume 1785 of *LNCS*, pages 395–410. Springer, 2000.
- [14] J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*. IEEE, 2004. To appear.
- [15] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.
- [16] M. Huth and M.Z. Kwiatkowska. Quantitative analysis and model checking. In *Proceedings of LICS'97*, pages 111–122. IEEE, 1997.
- [17] S.P. Iyer and M. Narasimha. Probabilistic lossy channel systems. In *Proceedings of TAPSOFT'97*, volume 1214 of *LNCS*, pages 667–681. Springer, 1997.
- [18] B. Jonsson, W. Yi, and K.G. Larsen. Probabilistic extensions of process algebras. *Handbook of Process Algebra*, pages 685–710, 1999.
- [19] A. Kučera and R. Mayr. A generic framework for checking semantic equivalences between pushdown automata and finite-state automata. In *Proceedings of IFIP TCS'2004*. Kluwer, 2004. To appear.
- [20] M.Z. Kwiatkowska. Model checking for probability and time: from theory to practice. In *Proceedings of LICS 2003*, pages 351–360. IEEE, 2003.
- [21] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *I&C*, 94(1):1–28, 1991.
- [22] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proceedings of ICALP 2003*, volume 2719 of *LNCS*, pages 1008–1021. Springer, 2003.
- [23] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *NJC*, 2(2):250–273, 1995.