

A Hemimetric Extension of Simulation for Semi-Markov Decision Processes ^{*}

Mathias R. Pedersen¹, Giorgio Bacci¹, Kim G. Larsen¹, and Radu Mardare¹

Department of Computer Science, Aalborg University, Denmark

Abstract. Semi-Markov decision processes (SMDPs) are continuous-time Markov decision processes where the residence-time on states is governed by generic distributions on the positive real line.

In this paper we consider the problem of comparing two SMDPs with respect to their time-dependent behaviour. We propose a hemimetric between processes, which we call *simulation distance*, measuring the least acceleration factor by which a process needs to speed up its actions in order to behave at least as fast as another process. We show that this distance can be computed in time $\mathcal{O}(n^2(f(l) + k) + mn^7)$, where n is the number of states, m the number of actions, k the number of atomic propositions, and $f(l)$ the complexity of comparing the residence-time between states. The theoretical relevance and applicability of this distance is further argued by showing that (i) it is suitable for compositional reasoning with respect to CSP-like parallel composition and (ii) has a logical characterisation in terms of a simple Markovian logic.

1 Introduction

Semi-Markov decision processes (SMDPs) are Markovian stochastic decision processes modelling the firing time of transitions via real-valued random variables describing the residence-time on states. Semi-Markov decision processes provide a more permissive model than continuous-time Markov decision processes, since they allow as residence-time distributions any generic distribution on the positive real line, rather than only exponential ones. The generality offered by SMDPs has been found useful in modelling several real-case scenarios. Successful examples include power plants [19] and power supply units [20], to name a few.

When considering such real-time stochastic processes, non-functional requirements are important, particularly requirements like response time and throughput, which depend on the timing behaviour of the process. We therefore wish to understand and be able to compare the timing behaviour of different processes.

To cope with the need for comparing the timing behaviour of different systems, in this paper we propose and study a quantitative extension of the simulation relation by Baier et al. [2], called ε -simulation, which puts the focus on the timing aspect of processes. The intuition is that a process s_2 ε -*simulates* another

^{*} The final authenticated version is available online at https://doi.org/10.1007/978-3-319-99154-2_21

process s_1 if after accelerating the actions of s_2 by a factor $\varepsilon > 0$ it reacts to the inputs from the external environment as s_1 with at least the same speed.

This type of quantitative reasoning is not new in the literature, and it dates back to the seminal work of Jou and Smolka [14, 9], who proposed the concept of probabilistic ε -bisimulation. This line of work has led to much work on probabilistic bisimulation distances [5, 7, 8]. While our work is conceptually similar to the bisimulation distances, it is technically very different. This is because bisimulation distances are constructed from a coalgebraic view as fixed points of operators. However, for the kind of timed systems that we are investigating, the coalgebraic perspective is much less understood. Moreover, since our distance generalises a preorder relation and not a congruence as the other distances do, it is not symmetric, which brings in new technical challenges

Following the work of Jou and Smolka, our notion of ε -simulation naturally induces a distance between processes: For any pair of states s_1 and s_2 , we define their *simulation distance* as the least acceleration factor needed by s_2 to speed up its actions in order to behave at least as fast as s_1 . This definition does not provide a distance in the usual sense, but rather a *multiplicative hemimetric*, i.e. an asymmetric notion of distance satisfying a multiplicative version of the triangle inequality. Such a notion is not new, as it is extensively applied in the context of differential privacy to measure information leakage of systems (see e.g. [1, 4]).

The theoretical relevance and applicability of the simulation distance is argued by means of the following results, which are the main technical contributions of the paper:

1. We provide an algorithm for computing the simulation distance between arbitrary states of an SMDP running in time $\mathcal{O}(n^2(f(l) + k) + mn^7)$, where n is the number of states, m the number of actions, k the number of atomic propositions, and $f(l)$ the complexity of comparing the residence time distributions on states.
2. We show that under some mild conditions on how residence-time distributions are combined in the parallel composition of two states, CSP-like parallel composition of SMDPs is non-expansive with respect to our hemimetric. This shows that the simulation distance is suitable for compositional reasoning.
3. We provide a logical characterisation of the distance in terms of a simple Markovian logic, stating that the distance from s_1 to s_2 is less than or equal to ε if and only if s_2 satisfies the ε -perturbation of any logical property that s_1 satisfies. Moreover, we prove that ε -simulation preserves the ε -perturbation of time-bounded reachability properties.

A full version of the paper containing all the technical proofs can be found at [18].

Notation and Preliminaries. Let \mathbb{N} denote the natural numbers, $\mathbb{Q}_{\geq 0}$ the non-negative rational numbers, $\mathbb{R}_{\geq 0}$ the non-negative real numbers, and $\bar{\mathbb{R}}_{> 0}$ the strictly positive ones. Given a set X , we will denote by $\mathcal{D}(X)$ the set of all probability measures on X . If $\mu \in \mathcal{D}(\mathbb{R}_{\geq 0})$, then the *cumulative distribution*

function (CDF) will be denoted by F_μ and is given by $F_\mu(t) = \mu([0, t])$. For $x \in \mathbb{R}_{\geq 0}$, we will write δ_x for the Dirac measure at x , which is defined by $\delta_x(E) = 1$ if $x \in E$ and $\delta_x(E) = 0$ otherwise. For any $\theta \in \mathbb{R}_{> 0}$, we will write $Exp[\theta]$ for the CDF of an exponential distribution with rate θ , and for $a, b \in \mathbb{R}_{\geq 0}$ such that $a < b$, we will write $Unif[a, b]$ for the CDF of a uniform distribution.

We will use the convention that $\infty + x = \infty$ for $x \in \mathbb{R}_{\geq 0}$ and $\infty \cdot y = \infty$ for $y \in \mathbb{R}_{> 0}$. A function $d: X \times X \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is called a *hemimetric* if it satisfies $d(x, x) = 0$ and the triangle inequality $d(x, z) \leq d(x, y) + d(y, z)$. It is called a *pseudometric* if it is also symmetric, i.e. $d(x, y) = d(y, x)$, and it is called a *metric* if it is symmetric and furthermore $d(x, y) = 0$ if and only if $x = y$.

2 Semi-Markov Decision Processes

In this section, we introduce semi-Markov decision processes, which are continuous-time reactive probabilistic systems. A semi-Markov decision process has residence time on states governed by generic distributions on the positive real line and reacts to inputs from an external environment by making a probabilistic transition to a next state.

Hereafter, we consider a non-empty finite set of *input actions* A , and a non-empty, finite set of *atomic propositions* AP .

Definition 1. A semi-Markov decision process (SMDP) is given by a tuple $M = (S, \tau, \rho, L)$ where

- S is a non-empty, countable set of states,
- $\tau: S \times A \rightarrow \mathcal{D}(S)$ is the transition function,
- $\rho: S \rightarrow \mathcal{D}(\mathbb{R}_{\geq 0})$ is the residence-time function, and
- $L: S \rightarrow 2^{AP}$ is the labelling function.

The operational behaviour of an SMDP is as follows. The SMDP at a given state $s \in S$, after receiving an input $a \in A$, goes to state $s' \in S$ within time t with probability $\tau(s, a)(s') \cdot \rho(s)([0, t])$. An SMDP is said to be finite if it has a finite set of states. For $s \in S$, we will write F_s for the CDF of $\rho(s)$, i.e. $F_s(t) = \rho(s)([0, t])$.

Continuous-time Markov decision processes are a special case of SMDPs in which all residence-time functions are exponentially distributed, and discrete-time Markov decision processes are a special case of SMDPs where the residence-time distribution in each state is the Dirac measure at 0, representing the fact that transitions are taken instantaneously.

In defining simulation and bisimulation for SMDPs, we will use ingredients from the definition of simulation and bisimulation for Markov decision processes [21] and simulation and bisimulation for continuous-time Markov chains [3]. However, since we are also generalising to arbitrary distributions on time rather than just exponential distributions, the condition on rates for exponential distributions must be replaced with a more general condition on the distributions. There is a rich literature on so-called stochastic orders [22], which impose an ordering

on random variables. We will consider here the most commonly used of these, known as the *usual stochastic order*.

Definition 2. For an SMDP $M = (S, \tau, \rho, L)$, a relation $R \subseteq S \times S$ is a simulation (resp. bisimulation) on M if for all $(s_1, s_2) \in R$ we have

1. $L(s_1) = L(s_2)$,
2. $F_{s_2}(t) \geq F_{s_1}(t)$ (resp. $F_{s_2}(t) = F_{s_1}(t)$) for all $t \in \mathbb{R}_{\geq 0}$, and
3. for all $a \in A$ there exists a weight function or coupling $\Delta_a : S \times S \rightarrow [0, 1]$ between $\tau(s_1, a)$ and $\tau(s_2, a)$ such that
 - (a) $\Delta_a(s, s') > 0$ implies $(s, s') \in R$,
 - (b) $\tau(s_1, a)(s) = \sum_{s' \in S} \Delta_a(s, s')$ for all $s \in S$, and
 - (c) $\tau(s_2, a)(s') = \sum_{s \in S} \Delta_a(s, s')$ for all $s' \in S$.

We say that s_2 simulates (resp. is bisimilar to) s_1 , written $s_1 \preceq s_2$ (resp. $s_1 \sim s_2$), if there is a simulation (resp. bisimulation) relation containing (s_1, s_2) .

It is easy to show that the *similarity* relation \preceq is the largest simulation relation, and analogously that the *bisimilarity* relation \sim is the largest bisimulation relation. The coupling ensures that the simulation relation is preserved by successor states. Intuitively, s_1 simulates s_2 if the CDF of $\rho(s_2)$ is pointwise greater than or equal to the CDF of $\rho(s_1)$, and the transition probability distribution of s_1 can be matched by the transition probability function s_2 by means of a coupling, in such a way that if two successor states s'_1 and s'_2 have a non-zero coupling, then s'_1 again simulates s'_2 . For bisimulation, we instead require that the CDFs behave exactly the same in each step.

Given a set $C \subseteq S$ and a relation $R \subseteq S \times S$, let

$$R(C) = \{s' \in S \mid (s, s') \in R \text{ for some } s \in C\}$$

be the R -closure of C . If R is a preorder, $R(C)$ is the upward closure of C .

The following result, which is a trivial generalisation of [25, Lemma 4.2.4], gives a different but equivalent definition of simulation which is sometimes useful.

Proposition 1. For finite S , $R \subseteq S \times S$ is a simulation relation if and only if for any $(s_1, s_2) \in R$ the first two conditions for simulation are satisfied and

$$\tau(s_1, a)(C) \leq \tau(s_2, a)(R(C)), \text{ for all } C \subseteq S .$$

The following generalises [3, Proposition 25(3)] to the case of SMDPs.

Proposition 2. $\preceq \cap \preceq^{-1} = \sim$.

The above is analogous to a result stating that bisimulation and simulation equivalence coincide for deterministic labelled transition systems [2]. In our case, Proposition 2 holds because reactive systems are inherently deterministic.

3 Comparing the Speed of Residence-Time Distributions

For comparing the random variables describing the residence time on states, the similarity relation uses the usual stochastic order: if $s_1 \preceq s_2$ then, for all $t \in \mathbb{R}_{\geq 0}$, $F_{s_1}(t) \leq F_{s_2}(t)$. In words, if s_2 simulates s_1 , it is more likely that s_2 will take a transition before s_1 , that is, s_2 is stochastically faster than s_1 in reacting to an input.

In this section, we propose a different way of comparing residence-time distributions. The idea is to get quantitative information about how much a distribution should be accelerated to become at least as fast as another distribution.

Definition 3. *Let F and G be CDFs and $\varepsilon \in \mathbb{R}_{>0}$. We say that F is ε -faster than G , written $F \sqsubseteq_\varepsilon G$, if for all t we have $F(\varepsilon \cdot t) \geq G(t)$.*

Consider two states s_1 and s_2 , having residence time governed by the distributions F_{s_1} and F_{s_2} , respectively, and assume that $F_{s_1} \sqsubseteq_\varepsilon F_{s_2}$ holds. If $0 < \varepsilon \leq 1$, then this means that s_1 is stochastically faster than s_2 , even if the residence time in s_1 is slowed down by a factor ε . If instead we have $\varepsilon > 1$, then s_1 is stochastically slower than s_2 , but if we accelerate its residence-time distribution by a factor ε , then it becomes stochastically faster than s_2 .

In the rest of the section we will argue that \sqsubseteq_ε is a good notion for gathering quantitative information about the speed of residence-time distributions on states. We will do this by comparing the most common distributions used in the literature for modelling residence time on states on stochastic systems: Dirac distributions, exponential distributions, and uniform distributions.

The Dirac measure at zero is the fastest measure, in the following sense.

Proposition 3. *Let F be any CDF. The following holds for any $\varepsilon \in \mathbb{R}_{>0}$.*

1. $\text{Dirac}[0] \sqsubseteq_\varepsilon F$.
2. If $F \neq \text{Dirac}[0]$, then $F \not\sqsubseteq_\varepsilon \text{Dirac}[0]$.

For comparing exponential distributions, it is simple to show that it is enough to accelerate by the ratio between the two rates. The same is true for uniform distributions, except we also need to consider whether the uniform distributions start at 0, since if a uniform distribution starts at 0, then we can only hope to make another uniform distribution faster than it if this other uniform distribution also starts at 0.

Proposition 4.

1. $\text{Exp}[\theta_1] \sqsubseteq_\varepsilon \text{Exp}[\theta_2]$, where $\varepsilon = \frac{\theta_2}{\theta_1}$.
2. If $c = 0$ and $a > 0$, then $\text{Unif}[a, b] \not\sqsubseteq_\varepsilon \text{Unif}[c, d]$ for any $\varepsilon \in \mathbb{R}_{>0}$.
3. If $c = 0$ and $a = 0$, then $\text{Unif}[a, b] \sqsubseteq_\varepsilon \text{Unif}[c, d]$, where $\varepsilon = \frac{b}{d}$.
4. If $c > 0$, then $\text{Unif}[a, b] \sqsubseteq_\varepsilon \text{Unif}[c, d]$, where $\varepsilon = \max \left\{ \frac{a}{c}, \frac{b}{d} \right\}$.

In all cases, the given ε is the least such that the ε -faster than relation holds.

Moreover, an exponential distribution can never be made faster than a uniform distribution, since uniform distributions become 1 eventually, whereas exponential distributions tend asymptotically to 1 but never reach it. Furthermore, whether or not a uniform distribution can be made faster than an exponential distribution depends on its value at 0.

Proposition 5.

1. $Exp[\theta] \not\sqsubseteq_{\varepsilon} Unif[a, b]$ for all $\varepsilon \in \mathbb{R}_{>0}$.
2. If $a > 0$, then $Unif[a, b] \not\sqsubseteq_{\varepsilon} Exp[\theta]$ for all $\varepsilon \in \mathbb{R}_{>0}$.
3. If $a = 0$, then $Unif[a, b] \sqsubseteq_{\varepsilon} Exp[\theta]$, where $\varepsilon = \theta \cdot b$. Furthermore, this is the least ε such that the ε -faster-than relation holds.

The ε -faster-than relation enjoys a kind of monotonicity property, which is simply a consequence of the fact that CDFs are increasing.

Lemma 1. *Let $\varepsilon \leq \varepsilon'$ and assume that $F \sqsubseteq_{\varepsilon} G$. Then $F \sqsubseteq_{\varepsilon'} G$.*

The probability distribution of the sum of two independent random variables is the *convolution* of their individual distributions. The general formula for the convolution of two measures μ and ν on the real line is given by

$$(\mu * \nu)(E) = \int_0^{\infty} \nu(E - x) \mu(dx) .$$

Notably, the ε -faster-than relation is a congruence with respect to convolution of measures.

Proposition 6. *If $F_{\mu_1} \sqsubseteq_{\varepsilon} F_{\mu_2}$ and $F_{\nu_1} \sqsubseteq_{\varepsilon} F_{\nu_2}$, then $F_{(\mu_1 * \nu_1)} \sqsubseteq_{\varepsilon} F_{(\mu_2 * \nu_2)}$.*

In Sect. 7.1 we will see that the above property is essential for the preservation of reachability properties. Intuitively, this is because convolution corresponds to sequential composition of the residence-time behaviour.

There are other possible ways to compare the relative speed of residence-time distributions quantitatively. In the following we explore some alternative definitions of the notion of the ε -faster-than relation, and argue that none of them are preferable to the one given in Definition 3. Given two CDFs F and G , we consider the following three alternative definitions of $F \sqsubseteq_{\varepsilon} G$:

1. for all t , $F(t) \cdot \varepsilon \geq G(t)$,
2. for all t , $F(t) + \varepsilon \geq G(t)$, and
3. for all t , $F(\varepsilon + t) \geq G(t)$.

If $\sqsubseteq_{\varepsilon}$ is defined as in (1), then we see that $Unif[a, b] \not\sqsubseteq_{\varepsilon} Unif[c, d]$, for any $\varepsilon \in \mathbb{R}_{>0}$ whenever $c < a$. This is because $Unif[a, b](a) \cdot \varepsilon = 0 < Unif[c, d](a)$. Hence we lose the properties of Proposition 4.

If $\sqsubseteq_{\varepsilon}$ is defined as in (2), we trivially get that whenever $\varepsilon \geq 1$, $F \sqsubseteq_{\varepsilon} G$, for any two CDFs F and G . Hence (2) is only interesting for $0 \leq \varepsilon < 1$. However, even in this case we would still lose the properties of Proposition 4. Indeed,

whenever $a \geq d$, $Unif[a, b] \not\sqsubseteq_\varepsilon Unif[c, d]$, for any $0 \leq \varepsilon < 1$. This follows because $Unif[a, b](a) + \varepsilon = \varepsilon < 1 = Unif[c, d](a)$.

Lastly, if \sqsubseteq_ε is defined as in (3), then it would not be a congruence with respect to convolution of distributions, i.e., Proposition 6 would not hold. For a counterexample, take $F_{\mu_1} = Unif[2, 4]$, $F_{\mu_2} = Unif[1, 3]$, $F_{\nu_1} = Unif[3, 4]$, and $F_{\nu_2} = Unif[2, 4]$. Then $F_{\mu_1} \sqsubseteq_1 F_{\mu_2}$ and $F_{\nu_1} \sqsubseteq_1 F_{\nu_2}$, but $F_{(\mu_1 * \mu_2)} \not\sqsubseteq_1 F_{(\nu_1 * \nu_2)}$.

4 A Hemimetric for Semi-Markov Decision Processes

In this section, we are going to extend the definition of simulation relation between SMDPs to the quantitative setting. We will see that this relation naturally induces a notion of distance between SMDPs, describing the least acceleration factor required globally on the residence-time distributions to make a given SMDP as fast as another one.

Definition 4. Let $\varepsilon \in \mathbb{R}_{>0}$. For an SMDP $M = (S, \tau, \rho, L)$, a relation $R \subseteq S \times S$ is a ε -simulation relation on M if for all $(s_1, s_2) \in R$ we have that the first and third condition for simulation are satisfied, and $F_{s_2} \sqsubseteq_\varepsilon F_{s_1}$. We say that s_2 ε -simulates s_1 , written $s_1 \preceq_\varepsilon s_2$, if there is a ε -simulation relation R such that $(s_1, s_2) \in R$.

Example 1. Let $A = \{a\}$ and consider the SMDP $M = (S, \tau, \rho, L)$ given by $S = \{s_1, s_2\}$, $\tau(s_1, a)(s_1) = 1 = \tau(s_2, a)(s_2)$, $F_{s_1} = Exp[4]$, $F_{s_2} = Exp[2]$, and $L(s_1) = L(s_2)$. By Proposition 4 we see that $s_1 \preceq_2 s_2$ and $s_2 \preceq_{\frac{1}{2}} s_1$.

It is easy to show that the ε -similarity relation \preceq_ε is the largest simulation relation, and with the previous section in mind, one immediately sees that \preceq_1 coincides with \preceq . Moreover, the following holds.

Proposition 7. For any $\varepsilon \leq 1$, if $s_1 \preceq_\varepsilon s_2$, then $s_1 \preceq s_2$.

If $\varepsilon > 1$, the above implication does not hold. For an easy counterexample consider s_1 and s_2 from Example 1 where $s_1 \preceq_2 s_2$ but $s_1 \not\preceq s_2$.

For $\varepsilon > 1$, we can obtain a result similar to Proposition 7 only if we “accelerate” the overall behaviour of s_2 . Formally, for a given SMDP $M = (S, \tau, \rho, L)$, we define the SMDP $M_\varepsilon = (S_\varepsilon, \tau_\varepsilon, \rho_\varepsilon, L_\varepsilon)$ as follows:

$$\begin{aligned} S_\varepsilon &= S \cup \{(s)_\varepsilon \mid s \in S\} , & \tau_\varepsilon(s, a)(s') &= \tau(s, a)(s') , \\ L_\varepsilon(s) &= L(s) , & \tau_\varepsilon(s, a)((s')_\varepsilon) &= 0 , \\ L_\varepsilon((s)_\varepsilon) &= L(s) , & \tau_\varepsilon((s)_\varepsilon, a)(s') &= 0 , \\ \rho_\varepsilon(s)([0, t]) &= \rho(s)([0, t]) , & \tau_\varepsilon((s)_\varepsilon, a)((s')_\varepsilon) &= \tau(s, a)(s') . \\ \rho_\varepsilon((s)_\varepsilon)([0, t]) &= \rho(s)([0, \varepsilon \cdot t]) , & & \end{aligned}$$

Intuitively, the states $s \in S$ in M_ε are identical copies of those in M , whereas the states $(s)_\varepsilon$ react to each input $a \in A$ functionally identically to s but faster,

since the residence-time on the states are all equally accelerated by a factor ε , thus $(s)_\varepsilon \preceq_\varepsilon s$. For this reason $(s)_\varepsilon$ is called the ε -acceleration of s .

Given the definition of accelerated state, Proposition 7 can be generalised to arbitrary values of $\varepsilon \in \mathbb{R}_{>0}$ in the following way.

Proposition 8. *For any $\varepsilon \in \mathbb{R}_{>0}$, $s_1 \preceq_\varepsilon s_2$ if and only if $s_1 \preceq (s_2)_\varepsilon$.*

The relevance of the above statement is twofold: it clarifies the relation between \preceq_ε and \preceq , and also provides a way to modify the behaviour of a state s_2 of an SMDP in order to simulate a state s_1 whenever $s_1 \preceq_\varepsilon s_2$ holds.

Having this characterisation of similarity in terms of acceleration of processes one can think about the following problem: given two states, s_1 and s_2 such that $s_1 \not\preceq s_2$, what is the least $\varepsilon \geq 1$ (if it exists) such that $s_1 \preceq (s_2)_\varepsilon$ holds? We can answer this question by means of the following distance.

Definition 5. *The simulation distance $d: S \times S \rightarrow [1, \infty]$ between two states s_1 and s_2 is given by*

$$d(s_1, s_2) = \inf\{\varepsilon \geq 1 \mid s_1 \preceq_\varepsilon s_2\} .$$

As usual, if there is no $\varepsilon \geq 1$ such that $s_1 \preceq_\varepsilon s_2$, then $d(s_1, s_2) = \infty$, because $\inf \emptyset = \infty$. It is clear from the definition that $s_1 \preceq s_2$ if and only if $d(s_1, s_2) = 1$.

Note that the definition above does not give a distance in the usual sense, for two reasons: d is not symmetric and it does not satisfy the triangle inequality. One can show instead that d satisfies a multiplicative version of the triangle inequality, namely, that for all $s_1, s_2, s_3 \in S$, $d(s_1, s_3) \leq d(s_1, s_2) \cdot d(s_2, s_3)$. This is a direct consequence of the following properties of \preceq_ε . The first property states that \preceq_ε is monotonic with respect to increasing values of ε .

Lemma 2. *If $s_1 \preceq_\varepsilon s_2$ and $\varepsilon \leq \varepsilon'$, then $s_1 \preceq_{\varepsilon'} s_2$.*

The second property is a quantitative generalisation of transitivity from which the multiplicative inequality discussed above follows.

Lemma 3. *If $s_1 \preceq_\varepsilon s_2$ and $s_2 \preceq_{\varepsilon'} s_3$, then $s_1 \preceq_{\varepsilon \cdot \varepsilon'} s_3$.*

Typically, one still uses the term distance for such multiplicative distances, because by applying the logarithm one does obtain a hemimetric.

Theorem 1. *$\log d(s_1, s_2)$ is a hemimetric.*

Example 2. Consider again the SMDP from Example 1. We can now see that $d(s_1, s_2) = 2$ and $d(s_2, s_1) = \frac{1}{2}$. This also shows that our distance is not symmetric, and hence not a pseudometric.

5 Computing the Simulation Distance

In this section we provide an algorithm to compute the simulation distance given in Definition 5 for finite SMDPs. The algorithm is shown to run in polynomial time for the distributions we have considered so far.

The following technical lemma will provide a sound basis for the correctness of the algorithm. Given two CDFs F and G , let

$$c(F, G) = \inf\{\varepsilon \geq 1 \mid F \sqsubseteq_\varepsilon G\}$$

denote the least acceleration factor needed by F to be faster than G .

Lemma 4. *For an SMDP M , define the set $\mathcal{C}(M) = \{c(F_{s'}, F_s) \mid s, s' \in S\}$. If $d(s_1, s_2) \neq \infty$, then*

- $s_1 \preceq_c s_2$, for some $c \in \mathcal{C}(M)$ and
- $d(s_1, s_2) = \min\{c \in \mathcal{C}(M) \mid s_1 \preceq_c s_2\}$.

Lemma 4 provides a strategy for computing the simulation distance between any two states s_1 and s_2 of a given SMDP M as follows. First, one constructs the set $\mathcal{C}(M)$. If $s_1 \preceq_c s_2$ does not hold for any $c \in \mathcal{C}(M)$, then the distance must be infinite; otherwise, it is the smallest $c \in \mathcal{C}(M)$ for which $s_1 \preceq_c s_2$ holds.

In order for this strategy to work, we need two ingredients: first, we should be able to compute the set $\mathcal{C}(M)$ and second, for any $c \in \mathcal{C}(M)$, we need an algorithm for checking whether $s_1 \preceq_c s_2$.

Recall that SMDPs allow for *arbitrary* residence-time distributions on states. Therefore, it is not guaranteed that for any SMDP M the set $\mathcal{C}(M)$ can be computed. With the following definition we identify the class of SMDPs for which this can be done.

Definition 6. *A class \mathcal{C} of CDFs is effective if for any $F, G \in \mathcal{C}$, $c(F, G)$ is computable. An SMDP M is effective if $\{F_s \mid s \in S\}$ is an effective class.*

In particular, for any pair of states s, s' of an effective SMDP M , we can decide whether $F_{s'} \sqsubseteq_\varepsilon F_s$ by simply checking whether $\varepsilon \geq c(F_{s'}, F_s)$. We will denote by $f(l)$ the complexity of computing $c(F_{s'}, F_s)$ for two arbitrary $s, s' \in S$ as a function of the length l of the representation of the residence-time distributions of M .

Let \mathcal{C}_A denote the class consisting of the Dirac distribution at 0 as well as uniform and exponential distributions with rational parameters. By Propositions 3–5 we immediately see that \mathcal{C}_A is an effective class, and in fact it can be computed using only simple operations such as multiplication, division, and taking maximum. Hence $f(l)$ has constant complexity¹ whenever M takes residence-time distributions from \mathcal{C}_A .

Next we consider how to decide $s_1 \preceq_\varepsilon s_2$ for a given rational $\varepsilon \geq 1$. A decision procedure can be obtained by adapting to our setting the algorithm by

¹ As is standard, we consider numbers to be represented as floating points of bounded size in their binary representation.

Baier et al. [2] for deciding the simulation preorder between probabilistic labelled transition systems. The algorithm from [2] uses a partition refinement approach to compute the largest simulation relation and runs in time $\mathcal{O}(mn^7/\log n)$ for reactive systems, where $m = |A|$ is the number of actions, and $n = |S|$ is the number of states. Given $\varepsilon \geq 1$, we can proceed correspondingly to compute ε -similarity: we start from the relation $R = S \times S$ and update it by removing all the pairs (s, s') of states not satisfying the conditions of Definition 4. This process is repeated on the resulting updated relation until no more pairs of states are removed. The resulting relation is the largest ε -simulation. Hence, checking $s_1 \preceq_\varepsilon s_2$ corresponds to determining whether the pair (s_1, s_2) is contained in the relation returned by the above algorithm.

Theorem 2. *Let M be a finite and effective SMDP. Given $s_1, s_2 \in S$ and $\varepsilon \geq 1$, deciding whether $s_1 \preceq_\varepsilon s_2$ can be done in time $\mathcal{O}(n^2(f(l) + k) + (mn^7)/\log n)$, where $k = |AP|$ is the number of atomic propositions.*

The algorithm for computing the simulation distance is given in Algorithm 1. The algorithm starts by ordering the elements of $\mathcal{C}(M)$ as c_1, \dots, c_n while removing ∞ from the list. Then it searches for the smallest c_i such that $s_1 \preceq_{c_i} s_2$ holds. This is done by means of a bisection method. If $s_1 \preceq_{c_1} s_2$ holds, then c_1 is the smallest element such that this holds, so we return it. If $s_1 \preceq_{c_n} s_2$ does not hold, then, by Lemma 2, $s_1 \preceq_{c_i} s_2$ does not hold for any $1 \leq i \leq n$, so we return ∞ . If none of the above apply, at this point of the algorithm (line 4) we have that $s_1 \not\preceq_{c_1} s_2$ and $s_1 \preceq_{c_n} s_2$.

We use the variables i and j , respectively, as the left and right endpoints of the bisection interval. The bisection interval keeps track of those elements c_n for which we still do not know whether $s_1 \preceq_{c_n} s_2$. At the beginning, $i = 1$ and $j = n$. At line 7, $h = \lceil \frac{j-i}{2} \rceil$ is used as the decrement factor for the length of the bisection interval at each step. Since $h > 0$, the bisection interval decreases in size for each iteration. If $s_1 \preceq_{c_{j-h}} s_2$ holds, then $j - h$ is the current smallest element in $\mathcal{C}(M)$ for which this holds, hence $j - h$ will become the new right endpoint of the interval; otherwise $i + h$ is the new left endpoint. The bisection method stops when the endpoints meet or cross each other, at which point we know that $s_1 \not\preceq_{c_n} s_2$ for all $n < j$ and $s_1 \preceq_{c_n} s_2$ for all $n \geq j$, and hence we return c_j .

Computing the set $\mathcal{C}(M)$ at line 1 has complexity $n^2 f(l)$, and sorting it can be done in time $\mathcal{O}(n \log n)$ using mergesort. By Theorem 2, and since we have already computed $\mathcal{C}(M)$, each of the ε -simulation checks in lines 2, 3, and 8 can be done in time $\mathcal{O}(n^2 k + (mn^7)/\log n)$, but the complexity $n^2 k$ from comparing labels only needs to be computed once. Since the bisection interval is halved each time, the while-loop is taken at most $\log n$ times. We therefore get an overall time complexity of $\mathcal{O}(n^2(f(l) + k) + mn^7)$.

Theorem 3. *Let M be a finite and effective SMDP. The simulation distance between any two states can be computed in time $\mathcal{O}(n^2(f(l) + k) + mn^7)$.*

```

1 Order the elements of  $\mathcal{C}(M) \setminus \{\infty\}$  such that  $c_1 < c_2 < \dots < c_n$ ;
2 if  $s_1 \preceq_{c_1} s_2$  then return  $c_1$  ;
3 else if  $s_1 \preceq_{c_n} s_2$  then return  $\infty$  ;
4 else
5    $i \leftarrow 1, j \leftarrow n$ ;
6   while  $i < j$  do
7      $h \leftarrow \lceil \frac{j-i}{2} \rceil$ ;
8     if  $s_1 \preceq_{c_{j-h}} s_2$  then  $j \leftarrow j - h$  ;
9     else  $i \leftarrow i + h$  ;
10  end
11  return  $c_j$ ;
12 end

```

Algorithm 1: Computing the simulation distance between s_1 and s_2 .

6 Compositional Properties of the Simulation Distance

In this section we will prove that some natural notions of parallel composition on SMDPs are non-expansive with respect to the simulation distance.

First we define what it means to compose two SMDPs in parallel. As argued in [23], the style of synchronous CSP is the one that is most suitable for SMDPs, so this is the one we will adopt here.

Definition 7. A function $\star : \mathcal{D}(\mathbb{R}_{\geq 0}) \times \mathcal{D}(\mathbb{R}_{\geq 0}) \rightarrow \mathcal{D}(\mathbb{R}_{\geq 0})$ is a residence-time composition function if it is commutative.

Definition 8. Let \star be a residence-time composition function. Then the \star -composition of $M_1 = (S_1, \tau_1, \rho_1, L_1)$ and $M_2 = (S_2, \tau_2, \rho_2, L_2)$, denoted $M_1 \parallel_\star M_2 = (S, \tau, \rho, L)$, is given as follows, for arbitrary $s_1, s'_1 \in S_1$, $s_2, s'_2 \in S_2$, and $a \in A$.

1. $S = S_1 \times S_2$,
2. $\tau((s_1, s_2), a)((s'_1, s'_2)) = \tau_1(s_1, a)(s'_1) \cdot \tau_2(s_2, a)(s'_2)$,
3. $\rho((s_1, s_2)) = \star(\rho_1(s_1), \rho_2(s_2))$, and
4. $L(s_1 \parallel_\star s_2) = L(s_1) \cup L(s_2)$.

Given a composite system $M_1 \parallel_\star M_2 = (S, \tau, \rho, L)$, we write $s_1 \parallel_\star s_2$ to mean $(s_1, s_2) \in S$. The residence-time composition function \star allows us to accommodate many different ways of combining timing behaviour, including those found in the literature on process algebras. We recall here some of these.

Maximum composition: $F_{\star(\mu, \nu)}(t) = \max(F_\mu(t), F_\nu(t))$.

For exponential distributions, $F_\mu = \text{Exp}[\theta]$ and $F_\nu = \text{Exp}[\theta']$, the following alternatives can be found.

Product rate composition: $F_{\star(\mu, \nu)} = \text{Exp}[\theta \cdot \theta']$.

Minimum rate composition: $F_{\star(\mu, \nu)} = \text{Exp}[\min\{\theta, \theta'\}]$.

Maximum rate composition: $F_{\star(\mu, \nu)} = \text{Exp}[\max\{\theta, \theta'\}]$.

Maximum composition is used for interactive Markov chains [11], product rate composition is used in SPA [12], minimum rate composition is used in PEPA [13], and maximum rate composition is used in TIPP [10].

In order to have non-expansiveness for \star -composition of SMDPs, we will need to restrict to residence-time composition functions \star that are monotonic.

Definition 9. *A residence-time composition function \star is monotonic if for all $\varepsilon \geq 1$ and $\mu, \nu, \eta \in \mathcal{D}(\mathbb{R}_{\geq 0})$, it holds that $F_\mu \sqsubseteq_\varepsilon F_\nu$ implies $F_{\star(\mu, \eta)} \sqsubseteq_\varepsilon F_{\star(\nu, \eta)}$.*

Requiring monotonicity is not a significant restriction, as many of the composition functions that are found in the literature are indeed monotonic.

Lemma 5. *Maximum composition as well as product, minimum, and maximum rate composition are all monotonic.*

Now we can prove that the \star -composition of finite SMDPs is non-expansive with respect to the simulation distance, provided that \star is monotonic.

Theorem 4. *For finite SMDPs and monotonic \star ,*

$$d(s_1, s_2) \leq \varepsilon \quad \text{implies} \quad d(s_1 \parallel_\star s_3, s_2 \parallel_\star s_3) \leq \varepsilon.$$

We conclude this section by exploring the computational aspects of composition of SMDPs. In particular, we would like to be able to also compute the distance between composite systems.

By Lemma 4, we know that computing the simulation distance amounts to being able to compute the constants $c(F_s, F_{s'})$, for each pair of states s, s' of the SMDP. Hence we would like that, whenever two distributions μ and ν have effective CDFs then also their composition $\star(\mu, \nu)$ has an effective CDF. By Proposition 4, it is easy to see that this holds for product, minimum, and maximum rate composition, since these compositions are still exponential distributions.

For maximum composition, the class \mathcal{C}_A is unfortunately not closed under composition. However, the following result holds.

Proposition 9. *Let \star be maximum composition. For any $\mu, \nu, \eta \in \mathcal{C}_A$, the constants $c(F_\mu, F_{\star(\nu, \eta)})$ and $c(F_{\star(\mu, \eta)}, F_\nu)$ are computable.*

The above results tells us that if we are interested in computing the distance $d(s_1, s_2 \parallel_\star s_3)$ or $d(s_1 \parallel_\star s_2, s_3)$, when \star is maximum composition, then we can indeed compute the constants c that are needed for Algorithm 1 to work.

7 Logical Properties of the Simulation Distance

If the distance between two processes is small, then we would also expect that they satisfy almost the same properties. In order to make this idea precise, in this section we introduce and study a slight extension of Markovian logic [15], which we will call *timed Markovian logic* (TML). The syntax of TML is given by the following grammar, where $\alpha \in AP$, $p \in \mathbb{Q}_{\geq 0} \cap [0, 1]$, $t \in \mathbb{Q}_{\geq 0}$, and $a \in A$.

$$\text{TML} : \quad \varphi ::= \alpha \mid \neg\alpha \mid \ell_p t \mid m_p t \mid L_p^a \varphi \mid M_p^a \varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi'$$

The semantics of TML is given by

$$\begin{array}{lll} s \models \alpha & \text{iff } \alpha \in L(s) & s \models \ell_p t \quad \text{iff } F_s(t) \geq p \\ s \models \neg\alpha & \text{iff } \alpha \notin L(s) & s \models m_p t \quad \text{iff } F_s(t) \leq p \\ s \models \varphi \wedge \varphi' & \text{iff } s \models \varphi \text{ and } s \models \varphi' & s \models L_p^a \varphi \quad \text{iff } \tau(s, a)(\llbracket \varphi \rrbracket) \geq p \\ s \models \varphi \vee \varphi' & \text{iff } s \models \varphi \text{ or } s \models \varphi' & s \models M_p^a \varphi \quad \text{iff } \tau(s, a)(\llbracket \varphi \rrbracket) \leq p \end{array}$$

where $\llbracket \varphi \rrbracket = \{s \in S \mid s \models \varphi\}$ is the set of states satisfying φ .

We also isolate the following two fragments of TML.

$$\text{TML}^{\geq} : \quad \varphi ::= \alpha \mid \neg\alpha \mid \ell_p t \mid L_p^a \varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi'$$

$$\text{TML}^{\leq} : \quad \varphi ::= \alpha \mid \neg\alpha \mid m_p t \mid M_p^a \varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi'$$

Intuitively, the formula $L_p^a \varphi$ says that the probability of taking an a -transition to where φ holds is *at least* p , and $M_p^a \varphi$ says the probability is *at most* p . $\ell_p t$ and $m_p t$ are similar in spirit, but talk about the probability of firing a transition instead. Thus, $\ell_p t$ says that the probability of firing a transition before time t is *at least* p , whereas $m_p t$ says that the probability is *at most* p .

For any $\varphi \in \text{TML}$ and $\varepsilon \geq 1$ we denote the ε -perturbation of φ by $(\varphi)_\varepsilon$ and define it inductively as

$$\begin{array}{llll} (\alpha)_\varepsilon = \alpha & (\ell_p t)_\varepsilon = \ell_p \varepsilon \cdot t & (L_p^a \varphi)_\varepsilon = L_p^a (\varphi)_\varepsilon & (\varphi \wedge \varphi')_\varepsilon = (\varphi)_\varepsilon \wedge (\varphi')_\varepsilon \\ (\neg\alpha)_\varepsilon = \neg\alpha & (m_p t)_\varepsilon = m_p \varepsilon \cdot t & (M_p^a \varphi)_\varepsilon = M_p^a (\varphi)_\varepsilon & (\varphi \vee \varphi')_\varepsilon = (\varphi)_\varepsilon \vee (\varphi')_\varepsilon \end{array}$$

By making use of the alternative characterisation for simulation given in Proposition 1 and drawing upon ideas from [6], we can now prove the following logical characterisation of the ε -simulation relation.

Theorem 5. *Let $\varepsilon \in \mathbb{Q}_{\geq 0}$ with $\varepsilon \geq 1$. For any finite SMDP, the following holds.*

- $s_1 \preceq_\varepsilon s_2$ if and only if $\forall \varphi \in \text{TML}^{\geq}. s_1 \models \varphi \implies s_2 \models (\varphi)_\varepsilon$.
- $s_1 \succeq_\varepsilon s_2$ if and only if $\forall \varphi \in \text{TML}^{\leq}. s_2 \models \varphi \implies s_1 \models (\varphi)_\varepsilon$.

As a special case of Theorem 5, we have also shown that TML^{\geq} and TML^{\leq} characterise simulation for SMDPs. Conceptually, Theorem 5 says that if s_1 ε -simulates s_2 , then s_2 satisfies the ε -perturbation of any property that s_1 satisfies for the TML^{\geq} fragment of TML, and vice versa for the TML^{\leq} fragment.

By Lemma 4 and Theorem 5, we get the following corollary, connecting our simulation distance with the properties expressible in the logic TML.

Corollary 1. *Let $\varepsilon \in \mathbb{Q}_{\geq 0}$ with $\varepsilon \geq 1$. For finite SMDPs the following holds.*

- $d(s_1, s_2) \leq \varepsilon$ if and only if $\forall \varphi \in \text{TML}^{\geq}. s_1 \models \varphi \implies s_2 \models (\varphi)_\varepsilon$.
- $d(s_1, s_2) \leq \varepsilon$ if and only if $\forall \varphi \in \text{TML}^{\leq}. s_2 \models \varphi \implies s_1 \models (\varphi)_\varepsilon$.

By Proposition 2, we also get a logical characterisation of bisimulation for SMDPs in terms of TML, which is simpler than the one given in [17, 24].

Theorem 6. *For any finite SMDP, it holds that*

$$s_1 \sim s_2 \quad \text{if and only if} \quad \forall \varphi \in \text{TML}. s_1 \models \varphi \iff s_2 \models \varphi .$$

7.1 Reachability Properties

We will now argue that the simulation distance behaves nicely also with respect to linear-time properties, by proving preservation of reachability properties up to perturbations.

The probability of reaching a given set of states in an SMDP depends on the choice of actions in each state. The non-determinism introduced by this choice is typically resolved by means of *schedulers*. Here we consider probabilistic schedulers σ of type $S^* \rightarrow \mathcal{D}(A)$, telling us what the probability is of selecting an action $a \in A$ depending on the history of the states visited so far. Given a scheduler σ , we denote by $\mathbb{P}_s^\sigma(\diamond^t X)$ the probability, under the scheduler σ , that starting from the state s the SMDP will eventually reach a state in $X \subseteq S$ within time $t \geq 0$ (for a rigorous definition of this probability see e.g. [17]).

Given our notion of ε -simulation, we can prove the following result.

Theorem 7. *Let β be a Boolean combination of atomic propositions. If we have $s_1 \preceq_\varepsilon s_2$, then for any scheduler σ there exists a scheduler σ' such that*

$$\mathbb{P}_{s_1}^\sigma(\diamond^t \llbracket \beta \rrbracket) \leq \mathbb{P}_{s_2}^{\sigma'}(\diamond^{\varepsilon \cdot t} \llbracket \beta \rrbracket) \quad (\text{or equivalently, } \mathbb{P}_{s_1}^\sigma(\neg \diamond^t \llbracket \beta \rrbracket) \geq \mathbb{P}_{s_2}^{\sigma'}(\neg \diamond^{\varepsilon \cdot t} \llbracket \beta \rrbracket)) .$$

Note that the above result might find useful applications for speeding up the computation time required by model checking tools to disprove certain types of reachability properties. For example, consider the atomic proposition `bad`, identifying all the states considered “not safe” in the SMDP. Usually, given a process s , one wants to verify that, under all possible schedulers σ , the probability $\mathbb{P}_s^\sigma(\neg \diamond^t \llbracket \text{bad} \rrbracket)$ is above a certain threshold value $\delta \leq 1$, meaning that the SMDP is unlikely to end up in an unsafe configuration within a time horizon bounded by t . Then, to disprove this property one only needs to provide a scheduler σ' and a process s' such that $s' \preceq_\varepsilon s$ and $\mathbb{P}_{s'}^{\sigma'}(\neg \diamond^{\frac{t}{\varepsilon}} \llbracket \text{bad} \rrbracket) < \delta$. Indeed, given that $s' \preceq_\varepsilon s$, by Theorem 7

$$\mathbb{P}_{s'}^{\sigma'}(\neg \diamond^{\frac{t}{\varepsilon}} \llbracket \text{bad} \rrbracket) < \delta \xrightarrow{\text{Th. 7}} \exists \sigma. \mathbb{P}_s^\sigma(\neg \diamond^t \llbracket \text{bad} \rrbracket) < \delta .$$

Since s simulates s' , s' can be thought of as a simplified abstraction of s , which is usually a smaller process. Hence, finding a scheduler σ' for s' which gives a counterexample may be much simpler than finding one for s . Moreover, the above technique is robust to perturbations of ε .

8 Conclusions and Open Problems

We have proposed a quantitative extension of the notion of simulation relation on SMDPs, called ε -simulation, comparing the relative speed of different processes. This quantitative notion of simulation relation induces a multiplicative hemimetric, which we call simulation distance, measuring the least acceleration factor needed by a process to speed up its actions in order to behave at least as fast as another process.

We have given an efficient algorithm to compute the simulation distance and identified a class of distributions for which the algorithm works on finite SMDPs. Furthermore, we have shown that, under mild conditions on the composition of residence-time distributions on states, a generalised version of CSP-like parallel composition on SMDPs is non-expansive with respect to this distance, showing that our distance is suitable for compositional reasoning. Lastly, we have shown the connection between our distance and properties expressible in a timed extension of Markovian logic. Namely, we have shown that if the simulation distance between s_1 and s_2 is at most ε , then s_1 satisfies the ε -perturbation of any property that s_2 satisfies. This result also gives a novel logical characterisation of simulation and bisimulation for semi-Markov decision processes.

Instead of using the usual stochastic order to relate the timing behaviour of states as we have done, one could also consider many other kinds of stochastic orders, for example ones that compare the expected value of the distributions. This may be more natural for applications where one wants to consider an exponential distribution with a high enough rate to be faster than a uniform distribution.

We have shown that the timing distributions that are obtained when composing systems are compatible with the algorithm for computing the distance only in the case when composing systems either on the left or on the right. A more general result showing that this also happens when composing on both sides an arbitrary number of components seems difficult. Nonetheless, we are confident that such a result can be obtained for any concrete case involving common types of distributions used in the literature.

Since we have both a distance and a logical characterisation, it makes sense to ask whether the set of states satisfying a formula is a closed or an open set in the topology induced by the distance. If such a set is indeed closed, this means that the approximate reasoning at the limit is sound: if every state in a sequence satisfies a formula, then the limit of that sequence also satisfies the formula. Such an investigation has already been done for Markovian logic [16], and some of the ideas from there may carry over to our setting.

Acknowledgments. We thank the anonymous reviewers for helpful suggestions as well as Robert Furber and Giovanni Bacci for insightful discussions. This research was supported by the Danish FTP project ASAP, the ERC Advanced Grant LASSO, and the Sino-Danish Basic Research Center IDEA4CPS.

References

- [1] Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. Additive and multiplicative notions of leakage, and their capacities. In *CSF*, pages 308–322, 2014.
- [2] Christel Baier, Bettina Engelen, and Mila E. Majster-Cederbaum. Deciding bisimilarity and similarity for probabilistic processes. *J. Comput. Syst. Sci.*, 60(1):187–231, 2000.

- [3] Christel Baier, Joost-Pieter Katoen, Holger Hermanns, and Verena Wolf. Comparative branching-time semantics for Markov chains. *Inf. Comput.*, 200(2):149–214, 2005.
- [4] Konstantinos Chatzikokolakis, Miguel E. Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *PETS*, pages 82–102, 2013.
- [5] Di Chen, Franck van Breugel, and James Worrell. On the complexity of computing probabilistic bisimilarity. In *FoSSaCS*, pages 437–451, 2012.
- [6] Josée Desharnais. Logical characterization of simulation for labelled Markov chains. In *PROBMIV*, pages 33–48. University of Birmingham, Technical Report, CS-99-8, August 1999.
- [7] Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled Markov processes. *Theor. Comput. Sci.*, 318(3):323–354, 2004.
- [8] Norm Ferns, Prakash Panangaden, and Doina Precup. Bisimulation metrics for continuous Markov decision processes. *SIAM J. Comput.*, 40(6):1662–1714, 2011.
- [9] Alessandro Giacalone, Chi-Chang Jou, and Scott A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proc. IFIP TC2 Working Conference on Programming Concepts and Methods*, pages 443–458. North-Holland, 1990.
- [10] Norbert Götz, Ulrich Herzog, and Michael Rettelbach. Multiprocessor and distributed system design: The integration of functional specification and performance analysis using stochastic process algebras. In Lorenzo Donatiello and Randolph D. Nelson, editors, *Performance Evaluation of Computer and Communication Systems*, volume 729 of *Lecture Notes in Computer Science*, pages 121–146. Springer, 1993.
- [11] Holger Hermanns. *Interactive Markov Chains: The Quest for Quantified Quality*, volume 2428 of *Lecture Notes in Computer Science*. Springer, 2002.
- [12] Holger Hermanns, Ulrich Herzog, and Vassilis Mertsiotakis. Stochastic process algebras - between LOTOS and Markov chains. *Computer Networks*, 30(9-10):901–924, 1998.
- [13] Jane Hillston. *A Compositional Approach to Performance Modelling*. Distinguished Dissertations in Computer Science. Cambridge University Press, 2005.
- [14] Chi-Chang Jou and Scott A. Smolka. Equivalences, congruences, and complete axiomatizations for probabilistic processes. In *CONCUR*, pages 367–383, 1990.
- [15] Dexter Kozen, Radu Mardare, and Prakash Panangaden. Strong completeness for Markovian logics. In *MFCS*, pages 655–666, 2013.
- [16] Kim Guldstrand Larsen, Radu Mardare, and Prakash Panangaden. Taking it to the limit: Approximate reasoning for Markov processes. In *MFCS*, pages 681–692, 2012.
- [17] Martin R. Neuhäüßer and Joost-Pieter Katoen. Bisimulation and logical preservation for continuous-time Markov decision processes. In *CONCUR*, pages 412–427, 2007.
- [18] Mathias R. Pedersen, Giorgio Bacci, Kim G. Larsen, and Radu Mardare. A hemimetric extension of simulation for semi-Markov decision processes. Technical report, Aalborg University, Department of Computer Science, 2018. <http://people.cs.aau.dk/mrp/pubs/simuldist.pdf>.
- [19] Mihael Perman, Andrej Senegacnik, and Matija Tuma. Semi-Markov models with an application to power-plant reliability analysis. *IEEE Transactions on Reliability*, 46(4):526–532, Dec 1997.
- [20] Antonio Pievatolo, Enrico Tironi, and Ivan Valade. Semi-Markov processes for power system reliability assessment with application to uninterruptible power supply. *IEEE Transactions on Power Systems*, 19(3):1326–1333, Aug 2004.

- [21] Roberto Segala and Nancy A. Lynch. Probabilistic simulations for probabilistic processes. *Nord. J. Comput.*, 2(2):250–273, 1995.
- [22] Moshe Shaked and George Shanthikumar. *Stochastic Orders*. Springer Series in Statistics. Springer, 2007.
- [23] Ana Sokolova and Erik P. de Vink. Probabilistic automata: System types, parallel composition and comparison. In Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, Joost-Pieter Katoen, and Markus Siegle, editors, *Validation of Stochastic Systems - A Guide to Current Research*, volume 2925 of *Lecture Notes in Computer Science*, pages 1–43. Springer, 2004.
- [24] Lei Song, Lijun Zhang, and Jens Chr. Godskesen. Bisimulations and logical characterizations on continuous-time Markov decision processes. In *VMCAI*, pages 98–117, 2014.
- [25] Lijun Zhang. *Decision algorithms for probabilistic simulations*. PhD thesis, Saarland University, Saarbrücken, Germany, 2009.