Analysing decisive stochastic processes*

Nathalie Bertrand[1], Patricia Bouyer[2], Thomas Brihaye[3], and Pierre Carlier[2,3]

1 **Inria Rennes Bretagne Atlantique, France**
2 **LSV, CNRS & ENS Cachan, France**
3 **Université de Mons, Belgium**

Abstract

In 2007, Abdulla *et al.* introduced the elegant concept of decisive Markov chain. Intuitively, decisiveness allows one to lift the good properties of finite Markov chains to infinite Markov chains. For instance, the approximate quantitative reachability problem can be solved for decisive Markov chains (enjoying reasonable effectiveness assumptions) including probabilistic lossy channel systems and probabilistic vector addition systems with states. In this paper, we extend the concept of decisiveness to more general stochastic processes. This extension is non trivial as we consider stochastic processes with a potentially continuous set of states and uncountable branching (common features of real-time stochastic processes). This allows us to obtain decidability results for both qualitative and quantitative verification problems on some classes of real-time stochastic processes, including generalized semi-Markov processes and stochastic timed automata.

1998 ACM Subject Classification D.2.4 Software/Program Verification, F.3.1 Specifying and Verifying and Reasoning about Programs, G.3 Probabilities and Statistics

Keywords and phrases Real-time stochastic processes, Decisiveness, Approximation Scheme

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.XXX

1 Introduction

Given its success for finite-state systems, the model checking approach to verification has been extended to various models based on automata, and including features such as time, probability and infinite data structures. Such models allow one to represent software systems more faithfully, and at the same time, they offer the possibility to consider *quantitative* verification questions. Such problems become particularly hard to solve for infinite-state systems, often requiring the development of dedicated techniques for each class of systems.

A decade ago, Abdulla *et al.* introduced the concept of decisiveness for denumerable Markov chains [2]. A Markov chain is decisive w.r.t. a set of states F if runs almost-surely reach F or a state from which F can no longer be reached. The concept of decisiveness rules out some weird behaviours in denumerable Markov chains, and lifts most good properties of finite Markov chains to infinite Markov chains. In particular, it enables the quantitative model checking of (repeated) reachability properties, by providing an approximation scheme, which is guaranteed to terminate for decisive Markov chains. Decisiveness also elegantly subsumes other concepts such as the existence of finite attractors, or coarseness [2].

Dense time required for representing real-time systems, is a potential source of infinity. However, stochastic real-time systems cannot be handled by the theory of decisive Markov

* The second and the fourth authors are supported by ERC project EQualIS. The third author is partly supported by FP7-EU project Cassting.

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).
Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi; Article No. XXX; pp. XXX:1–XXX:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

chains, as both the state space and the branching are in general non-denumerable. The general philosophy for models with dense time is to design an abstraction that preserves some properties of the original model, and is amenable to efficient model checking techniques. A prominent example of such abstractions is the region graph for timed automata [4]. However, abstractions often do not preserve quantitative properties, and they may be too coarse already for the evaluation of the probability of properties as simple as reachability properties.

In this paper, we generalize the concept of decisiveness to arbitrary stochastic systems, thus including the ones generated by real-time stochastic systems. While stochastic systems are often viewed operational in the model checking community (that is, one considers executions of a system), we take here a more abstract point-of-view, and consider the general mathematical model of stochastic processes.

Our first contribution is to define a notion of decisiveness for stochastic processes, generalizing the concept introduced by Abdulla *et al.* for denumerable Markov chains. This generalization is non trivial as we consider stochastic processes with a potentially continuous state space and uncountable branching, both being common features for modelling real-time stochastic processes. Moreover, in order to discriminate which verification techniques are sound, we refine the notion of decisiveness in three variants.

Our second contribution concerns the qualitative model checking of reachability and repeated reachability properties. We show that, under some decisiveness assumption, the almost-sure model checking of (repeated) reachability properties reduces to a simpler problem, namely to a reachability problem with probability 0. We advocate that this reduction simplifies the problem: in countable models, the 0-reachability amounts to the non existence of a path, in the underlying non-probabilistic system; beyond countable models, checking that a reachability property is satisfied with 0 probability amounts to exhibiting a somehow regular set of executions with positive measure.

A third contribution concerns quantitative model checking, here again for (repeated) reachability properties. Under some further decisiveness assumption, we prove that an approximation scheme, inspired from the path enumeration algorithm [17], is guaranteed to terminate. One can thus approximate, up to a desired precision, the probability of (repeated) reachability properties.

We then realize that non-Zeno real-time stochastic processes have good decisiveness properties when focusing on time-bounded reachability properties, which enables the evaluation of such properties within arbitrary precision.

Last, but not least, we introduce a generic notion of abstraction and explain how to derive decisiveness of the concrete model, using similar properties on the abstraction. We instantiate our framework with generalized semi-Markov processes (GSMP) and stochastic timed automata (STA), two models combining dense-time and probabilities. While the decidability of the qualitative model-checking was already known for STA [9], the current approach yields general approximation results for the quantitative model-checking, which were not known before.

2 Preliminaries

2.1 Stochastic processes

Let $(\Omega, \Sigma, \mathbb{P})$ be a probabilistic space, that is, Ω is a set called the *universe*, Σ is a σ-algebra, and $\mathbb{P}$ is a probability measure over (Ω, Σ). Let (S, Σ') be a measurable space. A *stochastic process* over $(\Omega, \Sigma, \mathbb{P})$ and (S, Σ') is a sequence $X = (X_i)_{i \geq 0}$ of random variables, where

$X_i \colon \Omega \to S$ is a measurable function. Note that we do not assume stochastic processes are homogeneous or discrete.

Let $X = (X_i)_{i\geq 0}$ be a stochastic process. Given B a measurable set of S (that is, $B \in \Sigma'$), we will abuse notation and write B for the *uniform* sequence $(B_i)_{i\in\mathbb{N}}$ such that $B_i = B$ for every $i \geq 0$. Given B a measurable set of S, we will sometimes write $X_i \in B$ for the event $X_i^{-1}(B)$, and when B is a singleton $\{q\}$, we might even write $X_i = q$.

We fix for the rest of the paper a universe $(\Omega, \Sigma, \mathbb{P})$, and a measurable space (S, Σ').

▶ **Example 1.** Let us give an example of a discrete stochastic process representing a random walk over the state space $S = \{q_n \mid n \in \{-1\} \cup \mathbb{N}\}$. The stochastic process X is defined by

- $\mathbb{P}(X_0^{-1}(q_0)) = 1$; $\quad \forall i \geq 0\ \mathbb{P}(X_{i+1}^{-1}(q_{-1}) \mid X_i^{-1}(q_{-1})) = 1$;
- $\forall i \geq 0,\ \forall n \in \mathbb{N}\ \mathbb{P}(X_{i+1}^{-1}(q_{n+1}) \mid X_i^{-1}(q_n))) = \frac{3}{4}$ and $\mathbb{P}(X_{i+1}^{-1}(q_{n-1}) \mid X_i^{-1}(q_n))) = \frac{1}{4}$.

It should be noted that we do not mention the universe Ω. It is always possible to construct a universe that has such a probability measure $\mathbb{P}$, so that it is irrelevant to introduce it (see [14]). This remark holds true in each example of the paper.

▶ **Example 2.** Another example, is the following non-Markovian stochastic process X over the state space $S = \{q_0, q_0', q_1, q_1'\}$:

- $\mathbb{P}(X_0 = q_0) = \mathbb{P}(X_0 = q_0') = \frac{1}{2}$; $\quad \forall i \geq 1,\ \mathbb{P}(X_i = q_1' \mid X_0 = q_0') = 1$;
- $\mathbb{P}(X_1 = q_1 \mid X_0 = q_0) = \lambda_1$ and $\mathbb{P}(X_1 = q_1' \mid X_0 = q_0) = 1 - \lambda_1$;
- $\forall i \geq 1,\ \mathbb{P}(X_{i+1} = q_1 \mid X_i = q_1, X_0 = q_0) = \mathbb{P}(X_{i+1} = q_1 \mid X_i = q_1', X_0 = q_0) = \lambda_i$;
- $\forall i \geq 1,\ \mathbb{P}(X_{i+1} = q_1' \mid X_i = q_1, X_0 = q_0) = \mathbb{P}(X_{i+1} = q_1' \mid X_i = q_1', X_0 = q_0) = 1 - \lambda_i$;

where $(\lambda_i)_{i\in\mathbb{N}}$ is a sequence of reals in $[0, 1]$. Note that X could be made Markovian by changing the state space, with one bit of memory to remember the initial state.

Real-time stochastic processes. A particular class of stochastic processes will be of interest to us, namely real-time stochastic processes, in which the time evolution is important. We define (S_t, Σ_t) as the measurable space defined by $S_t = S \times \mathbb{R}_+$, and where Σ_t is the σ-algebra generated by Σ' and the Borel sets of $\mathbb{R}_+$ (denoted $\mathcal{B}(\mathbb{R}_+)$). A *real-time stochastic process* over (S, Σ) is a stochastic process $Z = (Z_i)_{i\geq 0}$ over (S_t, Σ_t) such that:

- for every $i \geq 0$, $Z_i = (X_i, \tau_i)$, where $X_i \colon \Omega \to S$ and $\tau_i \colon \Omega \to \mathbb{R}_+$ are random variables;
- for each $i \geq 0$, $\mathbb{P}\big(\{\omega \in \Omega \mid \tau_i(\omega) < \tau_{i+1}(\omega)\}\big) = 1$.

The process $X = (X_i)_{i\geq 0}$ somehow represents the spatial behaviour of the system, while the process $\tau = (\tau_i)_{i\geq 0}$ gives the time evolution of the system. The second condition ensures that time almost-surely progresses. We will say that Z is *almost-surely non-Zeno* whenever $\mathbb{P}\big(\{\omega \in \Omega \mid (\tau_i(\omega))_{i\geq 0} \text{ is bounded}\}\big) = 0$.

In Section 5.3, we will see two classes of models that naturally fit into the framework of real-time stochastic processes. We can already mention here continuous-time Markov chains (we can find many examples of applications in [15]), or queuing systems (see below).

▶ **Example 3.** We consider a G/G/1-queue (of infinite capacity). A state of such a queue consists in the number of tasks waiting in the queue, the time delay since the last arrival in the queue (t_a) and the time delay since the last execution (t_e). Task arrivals follow a probability measure F_a, and task services are performed according to probability measure F_e. If at some point, the process is in state (n, t_a, t_e), the next arrival time in the queue is chosen according to $F_{a|t_a}$ and the next execution time is chosen according to $F_{e|t_e}$ where $F_{a|t_a}$ (resp. $F_{e|t_e}$) corresponds to the probability F_a (resp. F_e) given that at least t_a (resp. t_e) has elapsed. This induces a stochastic process $X = (X_i)_{i\in\mathbb{N}}$ where X_i is the state of the queue after i steps. To turn it into a real-time stochastic process, one simply adds global time τ giving at step i the amount of time spent since the beginning.

▶ Remark. Real-time stochastic processes as defined above are *discrete-time* stochastic processes (if we follow standard vocabulary), since the random variables are indexed by $\mathbb{N}$. However they abstract real-time continuous behaviours by giving relevant snapshots of the system (at all times given by the τ_i's). Such abstractions are used for instance in [18, Theorem 1] for abstracting continuous-space pure jump Markov processes while keeping relevant information on the process. We will see in Section 5.3 that these processes capture behaviours of intrinsically time continuous systems.

Events. A stochastic process X over $(\Omega, \Sigma, \mathbb{P})$ and (S, Σ') allows one to define various events expressed using LTL-like notations. Let $\mathcal{L}_{S,\Sigma'}$ be the set of formulas defined by the grammar: $\varphi ::= B\,\mathbf{U}_{\bowtie n}\,B' \mid \mathbf{G}\,\mathbf{F}\,B \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi$, where $B = (B_i)_{i\geq 0}$ and $B' = (B'_i)_{i\geq 0}$ are sequences of measurable subsets of S, $\bowtie \in \{\geq, \leq, =\}$ is a comparison operator and $n \in \mathbb{N}$ is an integer. The semantics of formulas in $\mathcal{L}_{S,\Sigma'}$ in terms of events is defined inductively:

$$\mathsf{Ev}_X(B\,\mathbf{U}_{\bowtie n}B') = \bigcup_{i \bowtie n}\big(X_i^{-1}(B'_i) \cap \bigcap_{0\leq j<i} X_j^{-1}(B_j)\big) \ ; \ \mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\,B) = \bigcap_{i\geq 0}\bigcup_{j\geq i} X_j^{-1}(B_j) \ ;$$
$$\mathsf{Ev}_X(\varphi_1 \vee \varphi_2) = \mathsf{Ev}_X(\varphi_1) \cup \mathsf{Ev}_X(\varphi_2) \ ; \ \mathsf{Ev}_X(\varphi_1 \wedge \varphi_2) = \mathsf{Ev}_X(\varphi_1) \cap \mathsf{Ev}_X(\varphi_2) \ ;$$
$$\mathsf{Ev}_X(\neg\varphi) = \Omega \setminus \mathsf{Ev}_X(\varphi).$$

Note that all these events are measurable in Ω. Following the intuition behind the LTL notations, event $\mathsf{Ev}_X(B\,\mathbf{U}_{\bowtie n}B')$ means that the stochastic process X will eventually satisfy B' (within step constraint $\bowtie n$), and only visit B beforehand. Also, the intuition of $\mathbf{G}\,\mathbf{F}\,B$ is that B should be visited infinitely often. We use classical shorthands: $\top = (S)_{i\geq 0}$; $\bot = (\emptyset)_{i\geq 0}$; $B\,\mathbf{U}\,B' = B\,\mathbf{U}_{\geq 0}B'$; $\mathbf{F}\,B = \top\,\mathbf{U}\,B$; $\mathbf{F}_{\bowtie n}B = \top\,\mathbf{U}_{\bowtie n}B$; $\mathbf{G}\,B = \neg\mathbf{F}\,(\neg B)$, where $\neg B = (S \setminus B_i)_{i\geq 0}$.

2.2 Decisiveness

Abdulla *et al.* originally defined a denumerable Markov chain to be decisive w.r.t. a set of states F if its runs almost-surely reach F or a state from which F can no longer be reached [2]. In order to extend the concept of *decisiveness* to general stochastic processes, we first provide an analogue to the set of states from which F is not reachable.

▶ **Definition 4.** Let B, B' be sequences of measurable sets of S. B' is a *B-avoidance sequence* for the stochastic process X if it satisfies

$$\forall n \geq 0,\ \mathbb{P}\big(\mathsf{Ev}_X(\mathbf{F}_{=n}\,B' \wedge \mathbf{F}_{\geq n}\,B)\big) = 0. \tag{1}$$

Intuitively, B' corresponds to 'states' from which B is almost-surely avoided (due to non-homogeneity of X, it needs to be defined as a sequence by slices).

▶ Remark. For every sequence B, $(\emptyset)_{i\geq 0}$ is a B-avoidance sequence for X. One can also check that B-avoidance sequences are closed under denumerable unions and intersections.

▶ **Example 5.** Let us illustrate the notion of avoidance sequences on the stochastic processes from Examples 1 and 2. In Example 1, we consider the uniform sequence $B = \{q_5\}$. It can be shown that the set of B-avoidance sequences corresponds to all sequences B' with $B'_i \subseteq \{q_{-1}\}$. In Example 2, the following sequence defines a B-avoidance set for $B = \{q_1\}$: $B'_0 = \{q'_0\}$ and for every $n \geq 1$, $B'_n = \emptyset$.

For the rest of the section, we fix a stochastic process $X = (X_i)_{i\geq 0}$ with $X_i \colon \Omega \to S$. Several notions of decisiveness were proposed for discrete-time and implicitly denumerable

Markov chains [2]. In this paper, we define three notions of decisiveness, adapting and refining the ones of [2] to general stochastic processes.

▶ **Definition 6.** Let B be a sequence of measurable sets in S and let B' be a B-avoidance sequence for X. We say that the stochastic process X is

- *initially decisive* (ID) w.r.t. B with *witness* B' if $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B \vee \mathbf{F}\, B')) = 1$
- *initially strongly decisive* (ISD) w.r.t. B with *witness* B' if $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\, B \vee \mathbf{F}\, B')) = 1$
- *persistently decisive* (PD) w.r.t. B with *witness* B' if $\forall n \geq 0$, $\mathbb{P}\big(\mathsf{Ev}_X(\mathbf{F}_{\geq n}\, B \vee \mathbf{F}_{\geq n}\, B')\big) = 1$.

We will then say that X is ID (resp. ISD, PD) w.r.t. B whenever there is some B-avoidance sequence B' such that X is ID (resp. ISD, PD) w.r.t. B with witness B'.

▶ Remark. Note that X might be ID (resp. ISD, PD) w.r.t. B for some witness B', but not for some other B-avoidance sequence B''. However if $B' \subseteq B''$ and X is ID (resp. ISD, PD) w.r.t. B with witness B', then it is also decisive w.r.t. B with witness B'': the larger (for the inclusion) is the B-avoidance sequence, the better it is for decisiveness properties.

We can establish a relationship between the three decisiveness notions.

▶ **Lemma 7.** *Let B be a sequence of measurable sets in S and let B' be a B-avoidance sequence for X. X is* PD *w.r.t. B with witness B' implies that X is* ISD *w.r.t. B with witness B', which, in turns, implies that X is* ID *w.r.t. B with witness B'. Moreover, the converse implications do not hold.*

Example 1 shows that initial decisiveness and initial strong decisiveness are not equivalent (take $B = \{q_5\}$), and Example 2 shows the non-equivalence of initial strong decisiveness and persistent decisiveness (take $\lambda_i = \frac{1}{2}$ for every $i \in \mathbb{N}$, and $B = \{q_1\}$).

In the sequel, we will write B' for an arbitrary B-avoidance sequence for X. However, whenever X is ID (resp. ISD, PD) w.r.t. B with some witness B', we will choose an arbitrary witness and write it $\mathsf{Av}_{\mathsf{dec}}(B)$ (resp. $\mathsf{Av}_{\mathsf{str}}(B)$, $\mathsf{Av}(B)$). When they exist, it is then possible to recursively define B'- (resp. $\mathsf{Av}_{\mathsf{dec}}(B)$-, $\mathsf{Av}_{\mathsf{str}}(B)$- and $\mathsf{Av}(B)$-)avoidance sequences for X: those are then order-two avoidance sequences for B, which record states from which one avoids states, from which states in B are avoided! The previous notations extend in the same way for these order-two avoidance sequences.

▶ **Example 8.** Back to Example 2, for $B = \{q_1\}$, we saw that B' defined by $B'_0 = \{q'_0\}$ and $B'_i = \emptyset$ for each $i \geq 1$, is a B-avoidance sequence for X. Then, we can define a B'-avoidance sequence for X as follows: $B''_0 = \{q_0, q_1, q'_1\}$ and for each $i \geq 1$, $B''_i = \{q_0, q_1, q'_0, q'_1\}$. In fact, we can show that all B'-avoidance sequences are the sequences included in B''.

3 Analysis of decisive stochastic processes

In this section we show how decisiveness properties can help analysing stochastic processes. In the first part, we focus on qualitative (that is, probability 0 or 1) reachability and repeated reachability properties, and we reduce all the corresponding model-checking questions to checking that some reachability property has probability 0. While this could be reduced to graph properties in [2], this is not the case here, since our models might have infinite non-denumerable branching. When we will apply these results in Subsection 5.3, models will have good properties allowing to solve the 0-probability properties of reachability properties. In the two next parts, we will use decisiveness properties to draw general procedures for computing (arbitrary) approximations of the probability of a (repeated) reachability property. Effectiveness of these procedures will of course rely on good effectiveness properties of the models that we want to analyze.

3.1 Qualitative reachability and repeated reachability

We aim at describing a procedure for checking the almost-sure satisfiability of a reachability (resp. a repeated reachability) property, that is, an event of the form $\mathbf{F}\, B$ (resp. $\mathbf{G}\,\mathbf{F}\, B$), where B is a sequence of measurable sets. We fix B' a B-avoidance sequence, and we recall the notations $\mathsf{Av}_{\mathsf{dec}}(B)$, $\mathsf{Av}_{\mathsf{str}}(B)$ and $\mathsf{Av}(B)$ for such sequences when X is ID, resp. ISD, resp. PD w.r.t. B.

▶ **Proposition 9.**
- *If* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B)) = 1$*, then* $\mathbb{P}(\mathsf{Ev}_X(\neg B\,\mathbf{U}\, B')) = 0$*.*
- *If X is* ID *w.r.t. B and* $\mathbb{P}(\mathsf{Ev}_X(\neg B\,\mathbf{U}\,\mathsf{Av}_{\mathsf{dec}}(B))) = 0$*, then* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B)) = 1$*.*

Under an initial decisiveness assumption, this reduces the almost-sure model-checking of reachability properties to checking that some kind of (constrained) reachability property is satisfied with probability 0. Note that, contrary to the case of discrete-time denumerable Markov chains, we cannot reduce to graph properties, yet we advocate that for reachability properties, checking whether the probability is 0, is simpler than checking whether it is 1. We will see in Subsection 5.3 how this can be exploited on specific examples.

Turning to almost-sure repeated reachability, one can show the following proposition:

▶ **Proposition 10.**
- *If* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\, B)) = 1$*, then* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B')) = 0$*;*
- *If X is* ISD *w.r.t. B and* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\,\mathsf{Av}_{\mathsf{str}}(B))) = 0$*, then* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\, B)) = 1$*.*

Under an initial strong decisiveness assumption, this reduces the almost-sure model-checking of a repeated reachability property to the 0-model-checking of some reachability property.

Concerning the positive model-checking of repeated reachability properties, one can show:

▶ **Proposition 11.**
- *If X is* PD *w.r.t. B and* ID *w.r.t.* $\mathsf{Av}(B)$*, and if* $\mathbb{P}\big(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\, B)\big) > 0$*, then* $\mathbb{P}\big(\mathsf{Ev}_X(\mathbf{F}\,\mathsf{Av}_{\mathsf{dec}}(\mathsf{Av}(B)))\big) > 0$*;*
- *If X is* PD *w.r.t B and* $\mathbb{P}\big(\mathsf{Ev}_X(\mathbf{F}\,\mathsf{Av}(B)')\big) > 0$*, then* $\mathbb{P}\big(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\, B)\big) > 0$*.*

Note that the existence of a witness B' such that X is PD w.r.t. B does not imply the existence of a witness such that X is ID w.r.t. B'.

3.2 Quantitative reachability

We assume B is a sequence of measurable sets of S and B' is a B-avoidance sequence. We define the two following sequences ($n \in \mathbb{N}$):

$$\begin{cases} p_n^{\mathsf{Yes}} & = \mathbb{P}\big(\mathsf{Ev}_X(\mathbf{F}_{\leq n}\, B)\big) \\ p_n^{\mathsf{No}} & = \mathbb{P}\big(\mathsf{Ev}_X(\neg B\,\mathbf{U}_{\leq n}\, B')\big) \end{cases}$$

The next proposition gives straightforward properties of these two sequences.

▶ **Proposition 12.** *The sequences* $(p_n^{\mathsf{Yes}})_{n\geq 0}$ *and* $(p_n^{\mathsf{No}})_{n\geq 0}$ *are non-decreasing and converge respectively to* $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B))$ *and* $\mathbb{P}(\mathsf{Ev}_X(\neg B\,\mathbf{U}\, B'))$*.*

▶ **Corollary 13.** *If X is* ID *w.r.t. B and* $B' = \mathsf{Av}_{\mathsf{dec}}(B)$*, then* $\lim_{n\to\infty} p_n^{\mathsf{Yes}} + p_n^{\mathsf{No}} = 1$*.*

Corollary 13 can be used to derive an approximation scheme to evaluate the probability of reachability properties in ID stochastic processes. Indeed, given a fixed error bound $\varepsilon > 0$, in order to compute $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B))$ up to ε, one only needs to iteratively compute the values p_n^{Yes} and p_n^{No} until $1 - p_n^{\mathsf{Yes}} - p_n^{\mathsf{No}} \leq \varepsilon$ to deduce that $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B)) - p_n^{\mathsf{Yes}} \leq \varepsilon$. In case p_n^{Yes} and p_n^{No} cannot be computed exactly, but can only be approximated up to any desired error bound, this scheme can be refined to obtain a 2ε-approximation for $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\, B))$.

▶ Remark. Note that the quality of the above approximation scheme depends on the choice of the sequence $\mathsf{Av}_{\mathsf{dec}}(B)$: the larger $\mathsf{Av}_{\mathsf{dec}}(B)$, the faster convergence. Intuitively, $\mathsf{Av}_{\mathsf{dec}}(B)$ permits to stop the exploration when the reachability goal can no longer be satisfied, hence the sooner the better.

3.3 Quantitative repeated reachability

We assume B is a sequence of measurable sets of S, B' is a B-avoidance sequence, and B'' is a B'-avoidance sequence. We define the two following sequences ($n \in \mathbb{N}$):

$$\begin{cases} q_n^{\mathsf{Yes}} = \mathbb{P}\big(\mathsf{Ev}_X(\neg B' \,\mathbf{U}_{\leq n}\, B'')\big) \\ q_n^{\mathsf{No}} = \mathbb{P}\big(\mathsf{Ev}_X(\neg B'' \,\mathbf{U}_{\leq n}\, B')\big) \end{cases}$$

▶ **Proposition 14.** *The sequences $(q_n^{\mathsf{Yes}})_{n\geq 0}$ and $(q_n^{\mathsf{No}})_{n\geq 0}$ are non-decreasing and converge respectively to $\mathbb{P}\big(\mathsf{Ev}_X(\neg B' \,\mathbf{U}\, B'')\big)$ and $\mathbb{P}\big(\mathsf{Ev}_X(\neg B'' \,\mathbf{U}\, B')\big)$.*

▶ **Proposition 15.** *If X is* PD *w.r.t. B (with witness $B' = \mathsf{Av}(B)$) and* ID *w.r.t. $\mathsf{Av}(B)$ (with witness $B'' = \mathsf{Av}_{\mathsf{dec}}(\mathsf{Av}(B))$), then the two sequences $(q_n^{\mathsf{Yes}})_{n\geq 0}$ and $(1 - q_n^{\mathsf{No}})_{n\geq 0}$ are adjacent and converge to $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\,B))$.*

Here again, the convergence of the two adjacent sequences can be used to derive an approximation scheme for $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\,B))$ in PD stochastic processes.

Note that the persistent decisiveness property is required for the approximation scheme to be correct: consider again Example 2 and assume the sequence $(\lambda_i)_{i\in\mathbb{N}}$ satisfies $\prod_{i\in\mathbb{N}}(1 - \lambda_i) > 0$. Under that hypothesis, one can show that $\mathbb{P}(\mathsf{Ev}_X(\mathbf{G}\,\mathbf{F}\,B)) < \frac{1}{2}$. On the other hand, whatever the choice of the avoidance sequences, we never get that the two sequences $(q_n^{\mathsf{Yes}})_{n\geq 0}$ and $(1 - q_n^{\mathsf{No}})_{n\geq 0}$ converge to that value.

4 Time-bounded reachability in real-time stochastic processes

In this section, we explain how to use decisiveness towards the quantitative analysis of time-bounded reachability (or safety) properties for real-time stochastic processes.

We fix (S_t, Σ_t) the measurable space for real-time stochastic processes we will consider. For $\Delta \in \mathbb{R}_+$ a time bound and B a sequence of measurable sets of S_t, we define the sequence $B \cap (t \leq \Delta)$ by: $(B \cap (t \leq \Delta))_i = \{(s,\tau) \in B_i \mid \tau \leq \Delta\}$. $B \cap (t \leq \Delta)$ is thus the restriction of B in which the time component is bounded by Δ.

First, decisiveness w.r.t. a sequence B propagates to its time-bounded restriction:

▶ **Proposition 16.** *Let $Z = (X_i, \tau_i)_{i\geq 0}$ be a real-time stochastic process, B be a sequence of measurable sets of S_t, and Δ be a time bound. If Z is* ID*, resp.* ISD*, resp.* PD *w.r.t. B, then Z is* ID*, resp.* ISD*, resp.* PD *w.r.t. $B \cap (t \leq \Delta)$.*

More importantly, non-Zeno real-time stochastic processes are PD w.r.t. time-bounded sequences:

▶ **Theorem 17.** *Let $Z = (X_i, \tau_i)_{i\geq 0}$ be a real-time stochastic process, let B be a sequence of measurable sets, and let $\Delta \in \mathbb{R}_+$ be a time bound. If Z is almost-surely non-Zeno, then Z is* PD *w.r.t. any time-bounded sequence $B \cap (t \leq \Delta)$.*

The main argument to establish this theorem is that, assuming almost-sure non-Zenoness, the sequence defined by $t > \Delta$ is a $(B \cap (t \leq \Delta))$-avoidance sequence for every B.

The almost-sure non-Zenoness hypothesis is standard and a desirable property of a system, and it expresses that the system should not have infinitely many discrete changes in a bounded amount of time. This assumption is satisfied by continuous-time Markov chains [6], and is easily enforced in many other models, such as continuous-time Markov processes with bounded transition rates [12], or continuous-space pure jump Markov processes (cPJMPs) assuming a non-explosive property [18].

This result then implies that for all these systems, if B is simple enough (like a uniform sequence of sets), provided one can compute (or approximate) the probability in n steps to reach some set of states, one can approximate the probability of satisfying a time-bounded reachability or safety property. This allows us to partly recover the result for cPJMPs [18, Theorem 3] which was established in a more analytical way.

In its generality, our approximation scheme does not provide any convergence rate, but for stochastic processes for which we can have an upper bound on the probability of completing at least n discrete changes within Δ time units, we will be able to compute a convergence rate for the various schemes. For instance for continuous-time (denumerable) Markov chains whose transition rates are upper-bounded by Λ, that probability can be bounded using a Poisson process of rate Λ.

5 Effectivity through abstraction

Proving decisiveness of general stochastic processes can be a hard task, in particular when their state-space is continuous. Decidability results in this context are often obtained through discrete abstractions. Therefore, in order to analyze the decisiveness of such stochastic processes, we propose to rely on an abstraction. More precisely, we give in this section sufficient conditions on an abstraction to ensure decisiveness of the original stochastic process. The qualitative verification algorithms and quantitative approximation schemes can then be applied to the concrete model. Note that, in general, the abstractions we propose only preserve qualitative properties, so that approximation schemes should be applied to the concrete model, not to the abstraction.

We then explain how this methodology can be applied to two classes of real-time stochastic processes, namely the ones generated by generalized semi-Markov processes and by stochastic timed automata. These models can be abstracted into discrete-time Markov chains while preserving the almost-sure satisfaction of reachability properties; this allows us to derive good decisiveness properties of the original models, and thus to infer approximation schemes.

5.1 Decisiveness for homogeneous denumerable Markov chains

Let us recall some basics of Markov chains. A denumerable Markov chain (MC, for short) is a stochastic process $Y = (Y_i)_{i\geq 0}$ with a denumerable state space T and which has the Markov property: for every $n \geq 0$, for all $t, t_0, t_1, \ldots, t_n \in T$, as soon as $\mathbb{P}(\bigwedge_{i=0}^{n} Y_i = t_i) > 0$, then $\mathbb{P}(Y_{n+1} = t \mid \bigwedge_{i=0}^{n} Y_i = t_i) = \mathbb{P}(Y_{n+1} = t \mid Y_n = t_n)$. The MC is *homogeneous* if for every n and for all $t, t' \in T$, $\mathbb{P}(Y_{n+1} = t' \mid Y_n = t) = \mathbb{P}(Y_n = t' \mid Y_{n-1} = t)$. In that case, the Markov chain is generated by a transition matrix $p_Y : T \times T \to [0,1]$ such that for every $t \in T$, $\sum_{t' \in T} p_Y(t,t') = 1$ and for every $n \geq 0$, $p_Y(t,t') = \mathbb{P}(Y_{n+1} = t' \mid Y_n = t)$. Under the condition that Y is homogeneous, one can simply define runs generated by Y. Precisely, a sequence $t_0, t_1, \cdots, t_n$ is a run, denoted $t_0 \to t_1 \cdots \to t_n$ if $p_Y(t_i, t_{i+1}) > 0$ for every $0 \leq i < n$; n is then the *length* of the run, and we write $t \to^* S$ as soon as there exists a state $t' \in S \subseteq T$ and a run $t = t_0 \to t_1 \cdots \to t_k = t'$.

Let us first explain how to characterize our various notions of decisiveness in the case of homogeneous MCs, and how they compare to the decisiveness of [2]. For every subset of states $C \subseteq T$, borrowing notations from [2], we let $\widetilde{C} = \{t \in T \mid t \not\to^* C\}$. State $t_0 \in T$ is an initial state of Y if $\mathbb{P}_Y(Y_0 = t_0) > 0$, and Y is said initialized at t_0 whenever t_0 is the unique initial state, that is $\mathbb{P}_Y(Y_0 = t_0) = 1$. If Y is a homogeneous MC, we write $Y[t]$ for the MC initialized at t, with transition matrix p_Y.

▶ **Lemma 18.** *Let Y be a homogeneous MC. For every C, $\widetilde{C}$ is a C-avoidance uniform sequence for Y. Moreover, it is maximal for the inclusion.*

The maximality property stated above allows to check decisiveness properties only with the witness $\widetilde{C}$: if Y is decisive with witness B', since $B' \subseteq \widetilde{C}$, it will also be decisive with witness $\widetilde{C}$. Recovering partly the original definitions of [2], we obtain the following characterization of our three notions of decisiveness:

▶ **Corollary 19.** *Let Y be a homogeneous MC, and C a set of states. Then:*

1. *Y is* ID *w.r.t. C iff $\mathbb{P}_Y\big(\mathsf{Ev}_Y(\mathbf{F}\, C \vee \mathbf{F}\, \widetilde{C})\big) = 1$;*
2. *Y is* ISD *w.r.t. C iff $\mathbb{P}_Y\big(\mathsf{Ev}_Y(\mathbf{G}\,\mathbf{F}\, C \vee \mathbf{F}\, \widetilde{C})\big) = 1$;*
3. *Y is* PD *w.r.t. C iff for every $p \geq 0$, $\mathbb{P}_Y\big(\mathsf{Ev}_Y(\mathbf{F}_{\geq p}\, C \vee \mathbf{F}_{\geq p}\, \widetilde{C})\big) = 1$ iff for every state t reachable from an initial state, $Y[t]$ is* ID *w.r.t. C.*

The third characterization implies that the decisiveness notion of [2] corresponds to our persistent decisiveness notion, in the case of homogeneous MCs.

Contrary to the case of general stochastic processes, initial strong decisiveness and persistent decisiveness coincide for homogeneous MCs (we recover here [2, Lemma 3.2]).

▶ **Lemma 20.** *Let Y be a homogeneous MC, and C a set of states. Then, Y* ISD *w.r.t. C iff Y is* PD *w.r.t. C.*

Note though that, even in this restricted context, initial decisiveness is not equivalent to initial strong decisiveness (recall Example 5). Finally, as already noticed in [2]:

▶ **Lemma 21.** *Let Y be a finite homogeneous MC, and C a set of states. Then, Y is* PD *w.r.t. C.*

5.2 Sound abstraction for decisiveness

Let us define a suitable notion of abstraction relating an arbitrary stochastic process X and a homogeneous MC Y such that decisiveness of Y implies decisiveness for X.

▶ **Definition 22.** Let $X = (X_i)_{i\geq 0}$ be a stochastic process, $Y = (Y_i)_{i\geq 0}$ be a homogeneous MC with denumerable state-space T equipped with the discrete σ-algebra $\Theta = 2^T$, and $\alpha : (S, \Sigma') \to (T, \Theta)$ be a mapping such that α and α^{-1} are measurable. The MC Y is an *α-abstraction* of X if for every sequence $A = (A_n)_{n\geq 0}$ of sets in Θ and for every $n \geq 0$

$$\mathbb{P}_Y\big(Y_n^{-1}(A_n) \mid \bigcap_{i<n} Y_i^{-1}(A_i)\big) > 0 \iff \mathbb{P}\big(X_n^{-1}(\alpha^{-1}(A_n)) \mid \bigcap_{i<n} X_i^{-1}(\alpha^{-1}(A_i))\big) > 0 \ .$$

Intuitively, Y is an α-abstraction of X if through the mapping α it preserves the events that may happen with positive probability.

In order to lift avoidance sequence from the abstraction to the concrete stochastic process, we rely on α-closed sets. A set $B \in \Sigma'$ is *α-closed* if $b \in B$ and $\alpha(b) = \alpha(b')$ implies $b' \in B$. Given $B \in \Sigma'$, by Lemma 18, $\widetilde{\alpha(B)}$ is an $\alpha(B)$-avoidance sequence for Y. Assuming B is α-closed, we obtain a B-avoidance sequence for X.

▶ **Lemma 23.** *Let Y be an α-abstraction of X and $B \in \Sigma'$ be an α-closed set. Then $(\alpha^{-1}(\widetilde{\alpha(B)}))$ is a B-avoidance sequence for X.*

However, α-abstractions do not necessarily preserve decisiveness properties, yet these can be ensured thanks to the following soundness notions.

▶ **Definition 24.** Let Y be an α-abstraction of X.

- Y is *sound* if for every α-closed set B, $\mathbb{P}_Y(\mathsf{Ev}_Y(\mathbf{F}\,\alpha(B))) = 1$ implies $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}\,B)) = 1$.
- Y is *persistently sound* if for every α-closed set B and every $p \geq 0$, $\mathbb{P}_Y(\mathsf{Ev}_Y(\mathbf{F}_{\geq p}\,\alpha(B))) = 1$ implies $\mathbb{P}(\mathsf{Ev}_X(\mathbf{F}_{\geq p}\,B)) = 1$.

Roughly said, sound abstractions preserve almost-sure satisfaction of reachability properties. Moreover, they allow to transfer decisiveness properties to the original stochastic process.

▶ **Proposition 25.** *Let Y be an α-abstraction of X and B an α-closed set.*

- *If Y is sound and* ID *w.r.t. $\alpha(B)$, then X is* ID *w.r.t. B with witness $(\alpha^{-1}(\widetilde{\alpha(B)}))$.*
- *If Y is persistently sound and* PD *w.r.t. $\alpha(B)$, then X is* PD *w.r.t. B with witness $(\alpha^{-1}(\widetilde{\alpha(B)}))$.*

▶ **Example 26.** Back to the queue of Example 3, we assume it is M/M/1, that is, F_a and F_e are exponential distributions of parameters λ and μ, respectively. Assuming $\lambda < \mu$, we can exhibit a persistently sound abstraction. Indeed, consider the random walk over $\mathbb{N}$ defined by $p(0,1) = 1$, and if $i \geq 1$, $p(i, i+1) = \frac{\lambda}{\lambda+\mu}$ and $p(i, i-1) = \frac{\mu}{\lambda+\mu}$. Since $\lambda < \mu$, this MC is PD w.r.t. each set of states and thus, the queue is PD w.r.t. each set of states that is closed under the abstraction.

5.3 Applications

We apply the previous study to two classes of systems.

5.3.1 Generalized semi-Markov processes

A generalized semi-Markov process [10, 13] is a stochastic process built on a finite set of events. Each event is equipped with a random variable representing its duration: either a variable-delay defined by a density function or a fixed-delay modelled by a Dirac distribution. A transition is characterized by a set of events which expire, and schedules a set of new events. This model is known to generalize continuous-time Markov chains.

The semantics of a GSMP $\mathcal{M}$ is given as a general state-space Markov chain (GSSMC), defined by a set of configurations and a transition kernel. Configurations of a GSMP are pairs consisting of a state and a valuation assigning a time value to each *scheduled* event. Such a value represents the time elapsed since the event was scheduled. Transitions between configurations combine a time-elapse and the occurrence of some scheduled events and/or the scheduling of new events. The set of configurations can be equipped with a natural σ-algebra $\mathcal{G}$, and the transition system induced by $\mathcal{M}$ is equipped with a transition kernel: for a configuration γ and a set $A \in \mathcal{G}$, $P_{\mathcal{M}}(\gamma, A)$ is the probability to move in one step from configuration γ to some configuration in A. This probability expresses a race between all enabled events, taking into account their residual density functions. The set of runs, *i.e.* infinite sequences of configurations, can then be equipped with a probability measure $\mathbb{P}_{\mathcal{M}}$. The GSSMC associated with a GSMP $\mathcal{M}$ can thus naturally be viewed as a (real-time) stochastic process $X^{\mathcal{M}}$.

Decisiveness. In general, GSMPs do not enjoy any decisiveness property. Indeed, [10, Section 3] presents an example that is not ID w.r.t. any region-closed set. Still, Brázdil *et al.* identified a sufficient condition, that ensures some kind of fairness: GSMP should be *single-ticking* (GSMP with some restriction on fixed-delay events, see [10] for the definition of this condition). Under that condition and using results of [10], one can show that the standard region abstraction for GSMPs is a persistently sound abstraction. More precisely, we consider as an abstraction the (finite-state) Markov chain $Y^{\mathcal{M}}$, whose states are regions, and such that there is a transition between region r and region r' as soon as there is a configuration $\gamma \in r$ from which the probability to reach r' in one step in $X^{\mathcal{M}}$ is positive. Probabilities are assumed to be uniform in $Y^{\mathcal{M}}$.

▶ **Theorem 27.** *Let $\mathcal{M}$ be a single-ticking GSMP with stochastic process $X^{\mathcal{M}}$. Then the region Markov chain $Y^{\mathcal{M}}$ is an persistently sound abstraction of $X^{\mathcal{M}}$.*

Proposition 25 then suffices to derive the decisiveness of the original stochastic process $X^{\mathcal{M}}$:

▶ **Corollary 28.** *Let $\mathcal{M}$ be a single-ticking GSMP with stochastic process $X^{\mathcal{M}}$. Then for every region-closed set B, the stochastic process $X^{\mathcal{M}}$ is* PD *w.r.t. B.*

As a consequence, we can apply all results of Section 3 to single-ticking GSMPs. We remark here that checking whether a reachability property has probability 0 can easily be done using the region graph abstraction: it amounts to checking in the (finite) region abstraction that there is no path from the initial state to the target [10]. Hence all qualitative questions related to region-based (repeated) reachability properties can be solved. For what concerns quantitative verification, assuming the distributions equipping the GSMPs can be handled numerically, this allows one to approximate the probability of reachability or repeated reachability properties, as well as all time-bounded reachability properties. We believe our approach gives new hints into the approximate model-checking problem for GSMPs, for which, up to our knowledge, only few results are known. For instance in [3, 7], the authors approximate the probability of until formulas of the form "the system reaches a target before time T within k discrete events, while staying within a set of safe states" (resp. "the system reaches a target while staying within a set of safe states") for GSMPs (resp. a restricted class of GSMPs which can be proved to be PD), and study numerical aspects. Our result permits to do the same with any reachability (resp. time-bounded reachability) property on the whole class of single-ticking GSMPs (resp. which are a.s. non-Zeno). The numerical aspects in our computations can be dealt with as in [3, 7].

5.3.2 Stochastic timed automata

Stochastic timed automata [9] are stochastic processes derived from timed automata [4] by randomizing both the delays and the edge choices. One can naturally associate a (real-time) stochastic process $X^{\mathcal{A}}$ with an STA $\mathcal{A}$.

Decisiveness. Similarly to GSMP, STA are not decisive in general. Adapting an example from [9], one can indeed exhibit an STA which is not ID w.r.t. a given region-closed set. Still, under some fairness property, one can build an abstraction of the STA, that is sound for the almost-sure model checking of LTL properties and the-like [9]. It turns out that this fairness assumption ensures that the same abstraction is also sound for decisiveness. As for GSMPs, we consider the natural region abstraction, and we define a finite-state Markov chain $Y^{\mathcal{A}}$, in which the states are regions, and there is a transition from one region r to another r' as soon as there exists a configuration γ in r from which the probability to reach r' in one step in $X^{\mathcal{A}}$

is positive. As mentioned earlier, as such, the abstraction $Y^{\mathcal{A}}$ is not sound in general. Yet, it preserves almost-sure satisfaction of LTL properties when the stochastic timed automaton is *almost-surely fair* [9]. Here fairness refers as the following property, which depends on $Y^{\mathcal{A}}$: every edge of $Y^{\mathcal{A}}$ which is enabled infinitely often along a run should be taken infinitely often. Denoting fair this property, the assumption $\mathbb{P}(\mathsf{Ev}_{X^{\mathcal{A}}}(\mathsf{fair})) = 1$ suffices to prove that $Y^{\mathcal{A}}$ is a sound abstraction for the almost-sure model checking for LTL properties [9].

▶ **Theorem 29.** *Let $\mathcal{A}$ be an STA with associated stochastic process $X^{\mathcal{A}}$, and let $Y^{\mathcal{A}}$ be its region abstraction. If $\mathbb{P}(\mathsf{Ev}_{X^{\mathcal{A}}}(\mathsf{fair})) = 1$ then $Y^{\mathcal{A}}$ is a persistently sound abstraction for $X^{\mathcal{A}}$.*

As a consequence of Proposition 25 and Theorem 29, we derive the decisiveness of the original stochastic process $X^{\mathcal{A}}$:

▶ **Corollary 30.** *Let $\mathcal{A}$ be an STA with stochastic process $X^{\mathcal{A}}$ and $Y^{\mathcal{A}}$ its region abstraction. If $\mathbb{P}(\mathsf{Ev}_{X^{\mathcal{A}}}(\mathsf{fair})) = 1$ then for every region-closed set B, $X^{\mathcal{A}}$ is PD w.r.t. B.*

As a consequence, we can apply all results of Section 3 to almost-surely fair stochastic timed automata. While the decidability of qualitative model-checking questions that we can infer from Section 3.1 were already known [9] and can be solved on the region abstraction, the approximation schemes that we can derive from Sections 3.2 and 3.3 are new. Obtaining decidability results even for the qualitative model-checking of large classes of STA required quite some effort (now combined in [9]). Here, we show that our earlier approach importantly implied decisiveness properties for STA. Moreover, the approximation schemes of Sections 3.2 and 3.3 can now be effectively applied to STA, as soon as distributions in the model have good numerical properties. Notice that a first decidability result was obtained in [8] for the quantitative model-checking of a restricted class of single-clock STA: while the current approximation schemes apply to all single-clock STA, the closed-form expression obtained in [8], although more precise requires a condition on the cycles of the automaton.

6 Conclusion and future work

In this paper, we introduced and studied decisiveness for general stochastic processes, setting sufficient conditions for the decidability of qualitative model-checking of (repeated) reachability properties, and more importantly for the approximability of the quantitative evaluation of such properties. We then showed that non-Zeno real-time stochastic processes have good decisiveness properties, allowing one to approximate the probability of all time-bounded properties. Finally we described a framework to obtain decisiveness properties through abstractions, and demonstrated its applicability to generalized semi-Markov processes and stochastic timed automata, thus yielding new approximability results for the quantitative model-checking of stochastic timed automata.

As further work, we would like to extend the applicability of our approach to other classes of stochastic timed systems, like probabilistic extensions of timed lossy channel systems [1] or communicating timed systems [11]. Also, the approximation scheme for reachability properties can be adapted to evaluate an expected accumulated reward, provided the reward evolves linearly in the model, as in Markov reward models [5, 16]. Finally, extending the approximation schemes to Muller conditions would enable the quantitative analysis of properties given as LTL formulas.

References

1 Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Jonathan Cederberg. Timed lossy channel systems. In *Proc. 31sth Conference on Foundations of Software Technology and Theoretical*

Computer Science (FSTTCS'12), volume 18 of *LIPIcs*, pages 374–386. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

2 Parosh Aziz Abdulla, Noomene Ben Henda, and Richard Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4), 2007.

3 Rajeev Alur and Mikhail Bernadsky. Bounded model checking for GSMP models of stochastic real-time systems. In *Proc. 9th International Workshop on Hybrid Systems: Computation and Control (HSCC'06)*, volume 3927 of *Lecture Notes in Computer Science*, pages 19–33. Springer, 2006.

4 Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

5 Christel Baier. Reasoning about cost-utility constraints in probabilistic models. In *Proc. 9th Workshop on Reachability Problems in Computational Models (RP'15)*, volume 9328 of *Lecture Notes in Computer Science*, pages 1–6. Springer, 2015.

6 Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(7):524–541, 2003.

7 Mikhail Bernadsky and Rajeev Alur. Symbolic analysis for GSMP models with one stateful clock. In *Proc. 10th International Workshop on Hybrid Systems: Computation and Control (HSCC'07)*, volume 4416 of *Lecture Notes in Computer Science*, pages 90–103. Springer, 2007.

8 Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *Proc. 5th International Conference on Quantitative Evaluation of Systems (QEST'08)*. IEEE Computer Society Press, 2008.

9 Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, Quentin Menet, Christel Baier, Marcus Größer, and Marcin Jurdziński. Stochastic timed automata. *Logical Methods in Computer Science*, 10(4):1–73, 2014.

10 Tomáš Brázdil, Jan Krčál, Jan Křetínský, and Vojtěch Řehák. Fixed-delay events in generalized semi-Markov processes revisited. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 2011.

11 Lorenzo Clemente, Frédéric Herbreteau, Amélie Stainer, and Grégoire Sutre. Reachability of communicating timed processes. In *Proc. 16th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'13)*, volume 7794 of *Lecture Notes in Computer Science*, pages 81–96. Springer, 2013.

12 Josée Desharnais and Prakash Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming*, 56:99–115, 2003.

13 Peter W. Glynn. A GSMP formalism for discrete event systems. *Proceedings of the IEEE*, 77(1):14–23, 1989.

14 Geoffrey R. Grimmett and David R. Stirzaker. *Probability and Random Processes.* Oxford University Press, 1992.

15 Ernst Moritz Hahn, Holger Hermanns, Björn Wachter, and Lijun Zhang. Time-bounded model checking of infinite-state continuous-time Markov chains. *Fundamenta Informaticae*, 95(1):129–155, 2009.

16 Ronald A. Howard. *Dynamic Probabilistic Systems, Volume II: Semi-Markov and Decision Processes.* John Wiley & Sons, 1971.

17 S. Purushothaman Iyer and Murali Narasimha. Probabilistic lossy channel systems. In *Proc. 7th International Joint Conference on Theory and Practice of Software Develop-*

ment (TAPSOFT'97), volume 1214 of *Lecture Notes in Computer Science*, pages 667–681. Springer, 1997.

18 Sadegh Soudjani, Rupak Majumdar, and Alessandro Abate. Safety verification of continuous-space pure jump Markov processes. In *Proc. 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'16)*, volume 9636 of *Lecture Notes in Computer Science*, pages 147–163. Springer, 2016.