

Timed Comparisons of Semi-Markov Processes^{*}

Mathias R. Pedersen¹, Nathanaël Fijalkow², Giorgio Bacci¹, Kim G. Larsen¹,
and Radu Mardare¹

¹ Department of Computer Science, Aalborg University, Selma Lagerlöfsvej 300,
Denmark {mrp,grbacci,kgl,mardare}@cs.aau.dk

² The Alan Turing Institute, 96 Euston Road, London NW1 2DB, United Kingdom,
nfijalkow@turing.ac.uk

Abstract. Semi-Markov processes are Markovian processes in which the firing time of transitions is modelled by probabilistic distributions over positive reals interpreted as the probability of firing a transition at a certain moment in time.

In this paper we consider the trace-based semantics of semi-Markov processes, and investigate the question of how to compare two semi-Markov processes with respect to their time-dependent behaviour. To this end, we introduce the relation of being “faster than” between processes and study its algorithmic complexity. Through a connection to probabilistic automata we obtain hardness results showing in particular that this relation is undecidable. However, we present an additive approximation algorithm for a time-bounded variant of the faster-than problem over semi-Markov processes with slow residence-time functions, and a **coNP** algorithm for the exact faster-than problem over unambiguous semi-Markov processes.

Keywords: Automata for system analysis and programme verification, real-time systems, semi-Markov processes, probabilistic automata.

1 Introduction

Semi-Markov processes are Markovian stochastic systems that model the firing time of transitions as probabilistic distribution over positive reals; thus, one can encode the probability of firing a certain transition within a certain time interval. For example, continuous-time Markov processes are particular case of semi-Markov processes where the timing distributions are always exponential.

Semi-Markov processes have been used extensively to model real-time systems such as power plants [15] and power supply units [16]. For such real-time systems, non-functional requirements are becoming increasingly important.

^{*} This work was supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1, the Danish FTP project ASAP, the ERC Advanced Grant LASSO, and the Sino-Danish Basic Research Center IDEA4CPS. This is a post-peer-review, pre-copyedit version of an article published in Lecture Notes in Computer Science. The final authenticated version is available online at: http://dx.doi.org/10.1007/978-3-319-77313-1_21

Many of these requirements, such as response time and throughput, depend heavily on the timing behaviour of the system in question. It is therefore natural to understand and be able to compare the timing behaviour of different systems.

Moller and Tofts [11] proposed the notion of a *faster-than* relation for systems with discrete-time in the context of process algebras. Their goal was to be able to compare processes that are functionally behaviourally equivalent, except that one process may execute actions faster than the other. This line of study was continued by Lüttgen and Vogler [10], who moreover considered upper bounds on time, in order to allow for reasoning about worst-case timing behaviours. For timed automata, Guha et al. [9] introduced a bisimulation-like faster-than relation and studied its compositional properties. For continuous-time probabilistic systems, Baier et al. [3] considered a simulation relation where the timing distribution on each state is required to stochastically dominate the other. They introduced both a weak and a strong version of their simulation relation, and gave a logical characterization of these in terms of the logic CSL.

In the literature, less attention has been drawn to trace-based notions of faster-than relations although trace equivalence and inclusion are important concepts when considering linear-time properties such as liveness or safety [2]. In this paper we propose a simple and intuitive notion of trace inclusion for semi-Markov processes, which we call *faster-than* relation, that compares the relative speed of processes with respect to the execution of arbitrary sequences of actions.

Differently from trace inclusion, our relation does not make a step-wise comparison of the timing delays for each individual action in a sequence, but over the overall execution time of the sequence. As an example, consider the semi-Markov process in Fig. 1. The states s and s' , although performing the same sequences of actions, are not related by trace inclusion because the first two actions in any sequence are individually executed at opposite order of speeds (here governed by exponential-time distributions). Instead, according to our relation, s is faster-than s' (but not vice versa) because it executes single-action sequences at a faster rate than s' , and action sequences of length greater than one at the same speed — this is due to the fact that the execution time of each action is governed by random variables that are independent of each other and the sum of independent random variables is commutative.




Fig. 1. A semi-Markov process where s is faster than s' . The states of the process are annotated with their timing distributions and each action-labelled transition is decorated with its probability to be executed.

In this paper we investigate the algorithmic complexity of various problems regarding the faster-than relation, emphasising their connection with classical algorithmic problems over Rabin's probabilistic automata. In particular, we prove

that the faster-than problem over generic semi-Markov processes is undecidable and that it is Positivity-hard when restricted to processes with only one action label. The reduction from the Positivity problem is important because it relates the faster-than problem to the Skolem problem, an important problem in number theory, whose decidability status has been an open problem for at least 80 years [12, 1].

We show that undecidability for the faster-than problem can not be tackled even by approximation techniques: via the same connection with probabilistic automata we are able to prove that the faster-than problem can not be approximated up to a multiplicative constant. However, as a positive result, we show that a time-bounded variant of the faster-than problem, which compares processes up to a given finite time bound, although still undecidable, admits approximated solutions up to an *additive* constant over semi-Markov processes with slow residence-time distributions. These include the important cases of uniform and exponential distributions.

Finally, we present a **coNP** algorithm for solving the faster-than problem exactly over unambiguous semi-Markov processes, where a process is unambiguous if every transition to a next state is unambiguously determined by the label that it outputs.

A full version of the paper with proofs and additional material can be found in [14].

2 Semi-Markov Processes and Faster-than Relation

For a finite set S we let $\mathcal{D}(S)$ denote the set of subdistributions over S , i.e. functions $\delta : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \delta(s) \leq 1$. The subset of total distributions is $\mathcal{D}_{=1}(S)$. We let \mathbb{N} denote the natural numbers and $\mathbb{R}_{\geq 0}$ denote the non-negative real numbers. We equip $\mathbb{R}_{\geq 0}$ with the Borel σ -algebra \mathcal{B} , so that $(\mathbb{R}_{\geq 0}, \mathcal{B})$ is a measurable space. Let $\mathcal{D}(\mathbb{R}_{\geq 0})$ denote the set of (sub)distributions over $(\mathbb{R}_{\geq 0}, \mathcal{B})$, i.e. measures $\mu : \mathcal{B} \rightarrow [0, 1]$ such that $\mu(\mathbb{R}_{\geq 0}) \leq 1$. Throughout the paper we will write $\mu(t)$ for $\mu([0, t])$. To avoid confusion we will refer to μ in $\mathcal{D}(\mathbb{R}_{\geq 0})$ as timing distributions, and to δ in $\mathcal{D}(S)$ as distributions.

Definition 1 (Semi-Markov process). *A semi-Markov process is a tuple $\mathcal{M} = (S, \text{Out}, \Delta, \rho)$ where*

- S is a (finite) set of states,
- Out is a (finite) set of output labels,
- $\Delta : S \rightarrow \mathcal{D}(S \times \text{Out})$ is a transition function,
- $\rho : S \rightarrow \mathcal{D}(\mathbb{R}_{\geq 0})$ is a residence-time function.

The operational behaviour of a semi-Markov process can be described as follows. In a given state $s \in S$, the process fires a transition within time t with probability $\rho(s)(t)$, leading to the state $s' \in S$ while outputting the label $a \in \text{Out}$ with probability $\Delta(s)(s', a)$.

We aim at defining $\mathbb{P}_{\mathcal{M}}(s, w, t)$, the probability that from the state s , the output of the semi-Markov process \mathcal{M} within time t starts with the word w . It

is important to note here that time is accumulated: we sum together the time spent in all states along the way, and ask that this total time is less than the specified bound t .

In order to account for the accumulated time in the probability, we need the notion of convolution. The convolution of two timing distributions μ and ν is $\mu * \nu$ defined, for any Borel set $E \subseteq \mathbb{R}_{\geq 0}$, as follows

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

Convolution is both associative and commutative. Let X and Y be two independent random variables with timing distributions μ and ν , i.e. $\mathbb{P}(X \in E) = \mu(E)$ and $\mathbb{P}(Y \in E) = \nu(E)$, then $\mathbb{P}(X + Y \in E) = (\mu * \nu)(E)$.

Definition 2 (Probability). Consider a semi-Markov process \mathcal{M} . We define the timing distribution $\mathbb{P}_{\mathcal{M}}(s, w)$ inductively on w , as follows, for any word $w \in \text{Out}^*$, label $a \in \text{Out}$, and time $t \in \mathbb{R}_{\geq 0}$

$$\begin{aligned} \mathbb{P}_{\mathcal{M}}(s, \varepsilon)(t) &= 1 \\ \mathbb{P}_{\mathcal{M}}(s, aw)(t) &= \sum_{s' \in S} \Delta(s)(s', a) \cdot (\rho(s) * \mathbb{P}_{\mathcal{M}}(s', w))(t) . \end{aligned}$$

We will then write $\mathbb{P}_{\mathcal{M}}(s, w, t)$ to mean $\mathbb{P}_{\mathcal{M}}(s, w)(t)$.

Timed Comparisons

We introduce the following relation which will be the focus of our paper.

Definition 3 (Faster-than relation). Consider a semi-Markov process \mathcal{M} and two states s and s' . We say that s is faster than s' , denoted $s \preceq s'$, if for all words $w \in \text{Out}^*$, for all time $t \in \mathbb{R}_{\geq 0}$,

$$\mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t) .$$

The algorithmic problem we consider in this paper is the *faster-than problem*: given a semi-Markov process and two states s and s' , determine whether $s \preceq s'$.

Algorithmic Considerations

The definition we use for semi-Markov processes is very general, because we allow for any residence-time function. The aim of the paper is to give generic algorithmic results which apply to *effective* classes of timing distributions, a notion we define now. Recall that a residence-time function associates with each state a timing distribution. We first give some examples of classical timing distributions.

- The prime example is exponential distributions, defined by the timing distribution $\mu(t) = 1 - e^{-\lambda t}$ for some parameter $\lambda > 0$ usually called the rate.

- Another interesting example is piecewise polynomial distributions. Consider finitely many polynomials P_1, \dots, P_n and a finite set of pairwise disjoint intervals $I_1 \cup I_2 \cup \dots \cup I_n$ covering $[0, \infty)$ such that for every k , P_k is non-negative over I_k and $\sum_k \int_{I_k} P_k = 1$. This induces the timing distribution

$$\mu(t) = \sum_k \int_{I_k \cap [0, t]} P_k(t) .$$

- Another important special case of piecewise polynomial distributions are the uniform distributions with parameters $0 \leq a < b$.
- The simplest example is given by Dirac distributions defined for the parameter a by $\mu(E) = 1$ if a is in E , and 0 otherwise.

The following definition captures these examples, and more. For a class \mathcal{C} of timing distributions, we let $\text{Convex}(\mathcal{C})$ be the smallest class of timing distributions containing \mathcal{C} and closed under convex combinations, and similarly $\text{Conv}(\mathcal{C})$ adding closure under convolutions.

Lemma 4. *Let \mathcal{C} be a class of timing distributions. Consider a semi-Markov process \mathcal{M} whose residence-time function uses timing distributions from \mathcal{C} . Then, for any state $s \in \mathcal{M}$ and word $w \in \text{Out}^*$, $\mathbb{P}_{\mathcal{M}}(s, w) \in \text{Conv}(\mathcal{C})$.*

In the rest of the paper we will consider only distributions that are suitable for algorithmic manipulation. Clearly, we must be able to give them as input to a computational device, hence we assume they can be described by finitely many rational parameters. Moreover, we require that testing inequalities between them is decidable, since this is essential for determining the faster-than relation. The next definition formalises this intuition.

Definition 5 (Effective timing distributions). *A class \mathcal{C} of timing distributions is effective if, for any $\varepsilon \geq 0$, $b \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, and $\mu_1, \mu_2 \in \text{Conv}(\mathcal{C})$, it is decidable whether $\mu_1(t) \geq \mu_2(t) - \varepsilon$, for all $t \leq b$.*

Proposition 6. *The following classes of timing distributions are effective: exponential, piecewise polynomial, uniform, and Dirac distributions.*

We do not provide in the conference version a full proof of Proposition 6, as it is mostly folklore but rather tedious. In particular, for exponential and piecewise polynomial distributions one relies on decidability results for the existential theory of the reals [17], implying that the most demanding operations above can be performed in polynomial space [4].

Although in this paper we give algorithmic results for generic effective classes of timing distributions, the semi-Markov processes we will focus on have only finitely many states, and hence can only use finitely many timing distributions from the same class. For our decidability results we will therefore focus on finite classes of timing distributions.

Moreover, in our complexity analyses, we will always assume that the operations on the timing distributions have a unit cost.

3 Hardness Results

We start the technical part of this article by presenting a series of hardness results for semi-Markov processes inherited from Markov processes.

A Markov process is a tuple $\mathcal{M} = (S, \text{Out}, \Delta)$ consisting of a (finite) set of states S , a (finite) set of labels Out , and a transition function $\Delta: S \rightarrow \mathcal{D}(S \times \text{Out})$. For a Markov process \mathcal{M} we define the probability $\mathbb{P}_{\mathcal{M}}(s)$ on Out^* inductively, for $a \in \text{Out}$ and $w \in \text{Out}^*$, as follows

$$\mathbb{P}_{\mathcal{M}}(s, \varepsilon) = 1 \quad \text{and} \quad \mathbb{P}_{\mathcal{M}}(s, aw) = \sum_{s' \in S} \Delta(s)(s', a) \cdot \mathbb{P}_{\mathcal{M}}(s')(w) .$$

The faster-than relation \preceq for Markov processes is defined similarly to the case of semi-Markov processes: $s \preceq s'$ if, for all words w , $\mathbb{P}_{\mathcal{M}}(s, w) \geq \mathbb{P}_{\mathcal{M}}(s', w)$.

We show that the faster-than problem for Markov processes, and hence also for semi-Markov processes, is (i) undecidable, (ii) can not be multiplicatively approximated, and (iii) is Positivity-hard even over the restricted case of Markov processes with one single output label. These limitations shape and motivate our positive results, which will be the topic of the remaining sections.

We first explain how hardness results for Markov processes directly imply hardness results for semi-Markov processes. The following lemma formalises the two ways semi-Markov processes subsume Markov processes.

Lemma 7. *Consider a semi-Markov process $\mathcal{M} = (S, \text{Out}, \Delta, \rho)$ and its induced Markov process $\mathcal{M}' = (S, \text{Out}, \Delta)$.*

- *If ρ is constant, i.e. for all s, s' we have $\rho(s) = \rho(s')$, then for all w , for all t , we have $\mathbb{P}_{\mathcal{M}}(s, w, t) = \mathbb{P}_{\mathcal{M}'}(s, w) \cdot \underbrace{(\rho(s) * \dots * \rho(s))}_{|w| \text{ times}}(t)$.*
- *If for all s , $\rho(s)$ is the Dirac distribution for 0, then for all w , for all t , we have $\mathbb{P}_{\mathcal{M}}(s, w, t) = \mathbb{P}_{\mathcal{M}'}(s, w)$.*

In particular in both cases, the following holds: for s, s' two states, we have $s \preceq s'$ in \mathcal{M} if, and only if, $s \preceq s'$ in \mathcal{M}' .

The hardness results of this section will be based on a connection to Rabin's probabilistic automata. A probabilistic automaton is given by

$$\mathcal{A} = (Q, A, q_0, \Delta : Q \times A \rightarrow \mathcal{D}_{=1}(Q), F) ,$$

where Q is the set of states, A is the alphabet, q_0 is an initial state, Δ is the transition function, and F is a set of final or accepting states. Any probabilistic automaton \mathcal{A} induces the probability $\mathbb{P}_{\mathcal{A}}(w)$ that a run over $w \in A^*$ is accepting, i.e. starts in q_0 and ends in F .

The key property of probabilistic automata that we will exploit is the undecidability of the universality problem, which was proved in [13], see also [8]. The universality problem is as follows: given a probabilistic automaton \mathcal{A} , determine whether $\mathbb{P}_{\mathcal{A}}(w) \geq \frac{1}{2}$, for all nonempty words $w \in A^+$.

Given a probabilistic automaton \mathcal{A} we define the *derived Markov process* $\mathcal{M}(\mathcal{A})$ as follows. The set of states of $\mathcal{M}(\mathcal{A})$ is $Q \times \{\ell, r\} \cup \{\top\}$, where \top is a new state; the set of output labels is A , and the transition function Δ' is defined as follows, for $p, q \in Q$ and $a \in \text{Out}$:

$$\begin{aligned} \Delta'(p, \ell)((q, \ell), a) &= \frac{1}{2|A|} \Delta(p, a)(q) & \Delta'(p, \ell)(\top, a) &= \frac{1}{2} \text{ if } p \in F \\ \Delta'(p, r)((q, r), a) &= \frac{1}{2|A|} \Delta(p, a)(q) & \Delta'(p, r)(\top, a) &= \frac{1}{2} . \end{aligned}$$

Let $s = (q_0, \ell)$ and $s' = (q_0, r)$, where q_0 is the initial state of \mathcal{A} . Then, for the Markov process $\mathcal{M}(\mathcal{A})$, we can then verify the following equalities:

$$\mathbb{P}_{\mathcal{M}(\mathcal{A})}(s, wa) = \frac{1}{(2|A|)^{|w|}} \mathbb{P}_{\mathcal{A}}(w) \quad \text{and} \quad \mathbb{P}_{\mathcal{M}(\mathcal{A})}(s', wa) = \frac{1}{(2|A|)^{|w|}} \frac{1}{2} .$$

Theorem 8. *The faster-than problem is undecidable for Markov processes.*

We discuss three approaches to recover decidability. A first approach is to look for *structural restrictions* on the underlying graph. However, the undecidability result above for probabilistic automata is quite robust in this respect, as it already applies when the underlying graph is acyclic, meaning that the only loops are self-loops. In spite of this, we present in Sect. 5 an algorithm to solve the faster-than problem for *unambiguous* semi-Markov processes.

A second approach is to restrict the *observations*. Interestingly, specialising the construction above to only one output letter yields a reduction from the Positivity problem. Formally, the Positivity problem reads: given a linear recurrence sequence, are all terms of the sequence non-negative? It has been shown that the universality problem for probabilistic automata with one letter alphabet is equivalent to the Positivity problem [1]. Thus, using again the derived Markov process $\mathcal{M}(\mathcal{A})$ for a probabilistic automaton \mathcal{A} with only one label, we obtain the following result.

Theorem 9. *The faster-than problem is Positivity-hard over Markov processes with one output label.*

A third approach is *approximations*. However, we can exploit further the connection we made with probabilistic automata, obtaining an impossibility result for *multiplicative approximation*. We rely on the following celebrated theorem for probabilistic automata due to Condon and Lipton [5]. The following formulation of their theorem is described in detail in [6].

Theorem 10 ([5]). *Let $0 < \alpha < \beta < 1$ be two constants. There is no algorithm which, given a probabilistic automaton \mathcal{A} ,*

- *if for all w we have $\mathbb{P}_{\mathcal{A}}(w) \geq \beta$, returns YES,*
- *if there exists w such that $\mathbb{P}_{\mathcal{A}}(w) \leq \alpha$, returns NO.*

Theorem 11. *Let $0 < \varepsilon < \frac{1}{3}$ be a constant. There is no algorithm which, given a Markov process \mathcal{M} and two states s, s' ,*

- *if for all w we have $\mathbb{P}_{\mathcal{M}}(s, w) \geq \mathbb{P}_{\mathcal{M}}(s', w)$, returns YES,*
- *if there exists w such that $\mathbb{P}_{\mathcal{M}}(s, w) \leq \mathbb{P}_{\mathcal{M}}(s', w) \cdot (1 - \varepsilon)$, returns NO.*

From these hardness results for Markov processes together with Lemma 7, we get the following hardness results for semi-Markov processes.

Corollary 12. *The following holds for semi-Markov processes for any class of timing distributions.*

- *The faster-than problem is undecidable.*
- *The faster-than problem with only one output label is Positivity-hard.*
- *The faster-than problem can not be multiplicatively approximated.*

4 Time-Bounded Additive Approximation

Instead of considering multiplicative approximation, we can also consider additive approximation, meaning that we want to decide whether for all w and t we have $\mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t) - \varepsilon$ for some constant $\varepsilon > 0$. In this section, we present an algorithm to solve the problem of approximating additively the faster-than relation with two assumptions:

- *time-bounded:* we only look at the behaviours up to a given bound b in $\mathbb{R}_{\geq 0}$,
- *slow residence-time functions:* each transition takes *some* time to fire.

As we will show, the combination of these two assumptions imply that the relevant words have bounded length. This is in contrast to the impossibility of approximating the faster-than relation multiplicatively showed in Sect. 3.

More precisely, we consider the *time-bounded* variant of the faster-than problem: given a time bound $b \in \mathbb{R}_{\geq 0}$, and two states s and s' in \mathcal{M} determine whether for all $t \leq b$ and w it holds that $\mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t)$.

We first observe that this restriction of the faster-than problem does not make any of the problems in Sect. 3 easier for semi-Markov processes. Indeed, if the residence-time functions are all Dirac distributions at 0, then all transitions are fired instantaneously, and the time-bounded restriction is immaterial. Thus we focus on distributions that do not fire instantaneously, as made precise by the following definition.

Definition 13 (Slow distributions). *We say that a class \mathcal{C} of timing distributions is slow if for all finite subset \mathcal{C}_0 of \mathcal{C} , there exists a computable function $\varepsilon: \mathbb{N} \times \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ such that for all n, t , and $\mu_1, \dots, \mu_n \in \text{Convex}(\mathcal{C}_0)$ we have $(\mu_1 * \dots * \mu_n)(t) \leq \varepsilon(n, t)$ and $\lim_{n \rightarrow \infty} \varepsilon(n, t) = 0$.*

Given a slow and effective class \mathcal{C} of timing distributions, we can do additive approximation of the time-bounded faster-than problem in the following way.

We introduce the following notation. Fix a semi-Markov process \mathcal{M} . Let $\mathcal{C}_{\mathcal{M}} = \text{Convex}(\{\rho(s) \mid s \in S\})$, and $n \in \mathbb{N}$. We define the timing distribution $F_{\mathcal{M},n}$ by $F_{\mathcal{M},n}(t) = 1$ if $n = 0$ and otherwise

$$F_{\mathcal{M},n}(t) = \sup \{(\mu_1 * \dots * \mu_n)(t) \mid \mu_1, \dots, \mu_n \in \mathcal{C}_{\mathcal{M}}\} .$$

Lemma 14. *For all s and all w , we have $\mathbb{P}_{\mathcal{M}}(s, w) \leq F_{\mathcal{M},|w|}$.*

Theorem 15. *For a constant $\varepsilon > 0$, there exists an algorithm which, given a semi-Markov process \mathcal{M} with slow and effective timing distributions, two states s, s' , and a bound $b \in \mathbb{R}_{\geq 0}$, determines whether*

$$\forall w, \forall t \leq b, \mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t) - \varepsilon .$$

Proof. Let $\mathcal{C}_{\mathcal{M}} = \text{Convex}(\{\rho(s) \mid s \in S\})$, since S is finite there exists a computable function $\varepsilon: \mathbb{N} \times \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ such that for all n, t , and $\mu_1, \dots, \mu_n \in \mathcal{C}_{\mathcal{M}}$ we have $(\mu_1 * \dots * \mu_n)(t) \leq \varepsilon(n, t)$ and $\lim_{n \rightarrow \infty} \varepsilon(n, t) = 0$. Given $\varepsilon > 0$, there exists N such that $\varepsilon(N, b) < \varepsilon$. For $n \geq N$. By assumption $(\mu_1 * \dots * \mu_n)(b) \leq \varepsilon(n, b) \leq \varepsilon(N, b) < \varepsilon$ for all $\mu_1, \dots, \mu_n \in \mathcal{C}_{\mathcal{M}}$. Taking the supremum over μ_1, \dots, μ_n , we then get $F_{\mathcal{M},n}(b) < \varepsilon$, and by Lemma 14, this means that for all w of length at least N , we have $\mathbb{P}_{\mathcal{M}}(s', w, b) < \varepsilon$. Hence it holds trivially that for all $t \leq b$ and w of length at least N , we have $\mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t) - \varepsilon$.

Thus the algorithm checks whether for all words of length less than N , for all $t \leq b$, we have $\mathbb{P}_{\mathcal{M}}(s, w, t) \geq \mathbb{P}_{\mathcal{M}}(s', w, t) - \varepsilon$, which is decidable thanks to the effectiveness of \mathcal{C} . \square

Next we show that there are interesting classes of timing distributions that are indeed slow. For this we introduce a class of timing distributions that are not just slow, but furthermore are guaranteed to converge to zero rapidly. We say that a timing distribution μ is *very slow* if there exists a computable function $\varepsilon: \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ such that $\lim_{t \rightarrow 0} \frac{\varepsilon(t)}{t} = 0$ and for all t , we have $\mu(t) \leq \varepsilon(t)$.

Theorem 16. *The following classes of timing distributions are slow: very slow, uniform, and exponential distributions.*

The proof of Theorem 16 depends on closed forms for the n -fold convolution of exponential distributions and uniform distributions, both of which converge to 0 as n goes to infinity. For exponential distributions, this closed form is the well-known Gamma distribution.

5 Unambiguous Semi-Markov Processes

In order to regain decidability of the faster-than relation, we can look at structurally simpler special cases of semi-Markov processes. Here we will focus on semi-Markov processes such that each output word induces at most one trace of states. More precisely, we will say that a semi-Markov process is *unambiguous* if for every state s and output label $a \in \text{Out}$, there exists at most one state s' such




Fig. 2. An example of an unambiguous semi-Markov process.

that $\Delta(s)(s', a) \neq 0$. A related notion of bounded ambiguity has been utilised to obtain decidability results in the context of probabilistic automata [7]. We introduce the following notation for unambiguous semi-Markov processes: $T(s, w)$ is the state reached after emitting w from s .

Example 17. Figure 2 gives an example of an unambiguous semi-Markov process. For each of the three states, there is at most one state that can be reached by a given output label. However, there need not be a transition for each output label from every state. In this example, the state s_2 has no b -transition, so for instance $T(s_1, ab) = s_2$, but $T(s_1, abb)$ is undefined.

Theorem 18. *The faster-than problem is decidable in **coNP** over unambiguous semi-Markov processes for all effective classes of timing distributions.*

Theorem 18 follows from the next proposition.

Proposition 19. *Consider an unambiguous semi-Markov process \mathcal{M} and two states s, s' . Let $L(s, s')$ be the set of loops reachable from (s, s') :*

$$\left\{ (p, p', v) \in S^2 \times \text{Out}^{\leq S^2} \mid \exists w \in \text{Out}^{\leq S^2}, \begin{array}{l} T(s, w) = p, T(s', w) = p', \\ T(p, v) = p, T(p', v) = p' \end{array} \right\}.$$

We have $s \preceq s'$ if, and only if

- for all $w \in \text{Out}^{\leq S^2}$, we have $\mathbb{P}_{\mathcal{M}}(s, w) \geq \mathbb{P}_{\mathcal{M}}(s', w)$, and
- for all $(p, p', v) \in L(s, s')$, we have $\mathbb{P}_{\mathcal{M}}(p, v) \geq \mathbb{P}_{\mathcal{M}}(p', v)$.

Before going into the proof, we explain how to use Proposition 19 to construct an algorithm solving the faster-than problem over unambiguous semi-Markov processes.

1. The first step is to compute $L(s, s')$, which can be done in polynomial time using a simple graph analysis,
2. The second step is to check the two properties, which both can be reduced to exponentially many queries of the form: $\mu_1 \geq \mu_2$ for μ_1, μ_2 in $\text{Conv}(\mathcal{C})$.

To obtain a **coNP** algorithm, in the second step we guess which of the two properties is not satisfied and a witness of polynomial length, which is either a word of quadratic length for the first property, or two states and a word of quadratic length for the second property.

Proof (of Proposition 19). (\implies) Assume that s is faster than s' and let (p, p') be in $L(s, s')$. There exist $w, v \in \text{Out}^*$ such that $T(s, w) = p$, $T(s', w) = p'$, $T(p, v) = p$, $T(p', v) = p'$. Let $n \in \mathbb{N}$. Since s is faster than s' , we have $\mathbb{P}_{\mathcal{M}}(s, wv^n) \geq \mathbb{P}_{\mathcal{M}}(s', wv^n)$. We have

$$\begin{aligned} \mathbb{P}_{\mathcal{M}}(s, wv^n) &= \mathbb{P}_{\mathcal{M}}(s, w) * \underbrace{\mathbb{P}_{\mathcal{M}}(p, v) * \cdots * \mathbb{P}_{\mathcal{M}}(p, v)}_{n \text{ times}} \\ \mathbb{P}_{\mathcal{M}}(s', wv^n) &= \mathbb{P}_{\mathcal{M}}(s', w) * \underbrace{\mathbb{P}_{\mathcal{M}}(p', v) * \cdots * \mathbb{P}_{\mathcal{M}}(p', v)}_{n \text{ times}} . \end{aligned}$$

Let $X_{s,w}$ be the random variable measuring the time elapsed from s emitting w . Similarly, we define $X_{p,v}, Y_{s',w}$ and $Y_{p',v}$. We have: for all $n \in \mathbb{N}$, for all t ,

$$\mathbb{P}_{\mathcal{M}}(X_{s,w} + nX_{p,v} \leq t) \geq \mathbb{P}_{\mathcal{M}}(Y_{s',w} + nY_{p',v} \leq t) ,$$

Dividing both sides by n yields

$$\mathbb{P}_{\mathcal{M}}\left(\frac{X_{s,w}}{n} + X_{p,v} \leq \frac{t}{n}\right) \geq \mathbb{P}_{\mathcal{M}}\left(\frac{Y_{s',w}}{n} + Y_{p',v} \leq \frac{t}{n}\right) .$$

We make the change of variables $x = \frac{t}{n}$: for all $n \in \mathbb{N}$, for all x we have

$$\mathbb{P}_{\mathcal{M}}\left(\frac{X_{s,w}}{n} + X_{p,v} \leq x\right) \geq \mathbb{P}_{\mathcal{M}}\left(\frac{Y_{s',w}}{n} + Y_{p',v} \leq x\right) .$$

Letting $n \rightarrow \infty$, we then obtain, for all x $\mathbb{P}_{\mathcal{M}}(X_{p,v} \leq x) \geq \mathbb{P}_{\mathcal{M}}(Y_{p',v} \leq x)$, which is equivalent to $\mathbb{P}_{\mathcal{M}}(p, v) \geq \mathbb{P}_{\mathcal{M}}(p', v)$.

(\Leftarrow) We prove that for all w , we have $\mathbb{P}_{\mathcal{M}}(s, w) \geq \mathbb{P}_{\mathcal{M}}(s', w)$ by induction on the length of w . For w of length at most S^2 , this is ensured by the first assumption. Let w be a word longer than S^2 . There exist two states p, p' such that p is reached by s and p' by s' after emitting i letters of w and again after emitting j letters of w , with j at most S^2 . Let $w = w_1 v w_2$ where v starts at position i and ends at position j . By construction (p, p', v) is in $L(s, s')$. We have

$$\begin{aligned} \mathbb{P}_{\mathcal{M}}(s, w) &= \mathbb{P}_{\mathcal{M}}(s, w_1) * \mathbb{P}_{\mathcal{M}}(p, v) * \mathbb{P}_{\mathcal{M}}(p, w_2) \\ &= \mathbb{P}_{\mathcal{M}}(s, w_1 w_2) * \mathbb{P}_{\mathcal{M}}(p, v) \\ &\geq \mathbb{P}_{\mathcal{M}}(s', w_1 w_2) * \mathbb{P}_{\mathcal{M}}(p', v) && \text{(inductive hypothesis)} \\ &= \mathbb{P}_{\mathcal{M}}(s', w) . \end{aligned}$$

□

6 Conclusion and Open Problems

We have introduced a trace-based relation on semi-Markov processes called the faster-than relation which asks that for any time bound, the probability of outputting any word within the time bound is higher in the faster process than in the

slower process. We have shown through a connection to probabilistic automata that the faster-than relation is highly undecidable. It is undecidable in general, and remains Positivity-hard even restricting to processes with one output label. Furthermore, approximating the faster-than relation up to a multiplicative constant is shown to be impossible.

However, we constructed algorithms for special cases of the faster-than problem. We have shown that if one considers approximating up to an additive constant rather than a multiplicative constant, and if one gives a bound on the time up to which one is interested in comparing the two processes, then approximation can be done for timing distributions in which we are sure to spend some amount of time to take a transition. In addition, we have shown that the faster-than relation over unambiguous processes is decidable and in **coNP**.

In this paper, we have focused on the generative model, where the labels are treated as outputs. An alternative viewpoint would be to consider reactive models, where the labels are instead treated as inputs [18]. While all the undecidability and hardness results we have shown can also easily be shown to hold for reactive Markov processes, the same is not true for the algorithms we have constructed. It is non-trivial to extend these algorithms for the case of reactive semi-Markov processes: the main obstacle is that for reactive systems, one has to also handle schedulers, often uncountably many. It is therefore still an open question whether our decidability results carry over to reactive models.

References

1. S. Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for Markov chains. *Inf. Process. Lett.*, 115(2):155–158, 2015.
2. Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
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. John F. Canny. Some Algebraic and Geometric Computations in PSPACE. In *STOC*, pages 460–467. ACM, 1988.
5. Anne Condon and Richard J. Lipton. On the complexity of space bounded interactive proofs (extended abstract). In *FOCS*, 1989.
6. Nathanaël Fijalkow. Undecidability results for probabilistic automata. *SIGLOG News*, 4(4), 2017.
7. Nathanaël Fijalkow, Cristian Riveros, and James Worrell. Probabilistic automata of bounded ambiguity. In *CONCUR*, pages 19:1–19:14, 2017.
8. Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP*, 2010.
9. Shibashis Guha, Chinmay Narayan, and S. Arun-Kumar. On decidability of pre-bisimulation for timed automata. In *CAV*, pages 444–461, 2012.
10. Gerald Lüttgen and Walter Vogler. A faster-than relation for asynchronous processes. In *CONCUR*, pages 262–276, 2001.
11. Faron Moller and Chris M. N. Tofts. Relating processes with respect to speed. In *CONCUR*, pages 424–438, 1991.

12. Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *SODA*, 2014.
13. Azaria Paz. *Introduction to Probabilistic Automata*. Academic Press, 1971.
14. Mathias Ruggaard Pedersen, Nathanaël Fijalkow, Giorgio Bacci, Kim Guldstrand Larsen, and Radu Mardare. Timed comparisons of semi-Markov processes. *CoRR*, 2017. <http://arxiv.org/abs/1711.10216>.
15. 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.
16. 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.
17. Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California press, 1951.
18. Rob J. van Glabbeek, Scott A. Smolka, and Bernhard Steffen. Reactive, generative and stratified models of probabilistic processes. *Inf. Comput.*, 121(1):59–80, 1995.