A Faster-Than Relation for Semi-Markov Decision Processes

Mathias Ruggaard Pedersen Giorgio Bacci Kim Guldstrand Larsen
Department of Computer Science
Aalborg University
Denmark
{mrp,grbacci,kgl}@cs.aau.dk

When modeling concurrent or cyber-physical systems, non-functional requirements such as time are important to consider. In order to improve the timing aspects of a model, it is necessary to have some notion of what it means for a process to be faster than another, which can guide the stepwise refinement of the model. To this end we study a *faster-than relation* for semi-Markov decision processes and compare it to standard notions for relating systems. We consider the compositional aspects of this relation, and show that the faster-than relation is not a precongruence with respect to parallel composition, hence giving rise to so-called parallel timing anomalies. We take the first steps toward understanding this problem by identifying decidable conditions sufficient to avoid parallel timing anomalies in the absence of non-determinism.

1 Introduction

Timing aspects are important when considering real-time or cyber-physical systems. For example, they are of interest in real-time embedded systems when one wants to verify the worst-case execution time for guaranteeing minimal system performance or in safety-critical systems when one needs to ensure that unavoidable rigid deadlines will always be met [8].

Semi-Markov decision processes are continuous-time Markov decision processes where the residence-time on states is governed by generic distributions on the positive real line. These systems have been extensively used to model real-time cyber-physical systems [12, 20].

For reasoning about timing aspects it is important to understand what it formally means for a real-time or cyber-physical system to operate faster than another. To this end we define the notion of *faster-than relation* for semi-Markov decision processes. The definition of faster-than relation we propose in this paper is a reactive version of an analogous notion of faster-than relation previously introduced in [15] for the case of generative systems. According to our relation, a semi-Markov decision process is faster than another one when it reacts to any sequence of inputs with equal or higher probability than the slower process, within the same time bound.

Often, complex cyber-physical systems are organised as concurrent systems of multiple components running in parallel and interacting with each other. Such systems are better analysed *compositionally*, that is, by breaking them into smaller components that are more easily examined [4]. However, it is not always the case that an analysis on the components carries over to the full composite system. A well known example of this, occurring in real-time systems such as scheduling for processors [3, 9], are *timing anomalies*, that is, when locally faster behaviour leads to a globally slower behaviour [7].

In this paper we study the compositional aspects of the faster-than relation for semi-Markov decision processes. We are interested in the situation where we have a composite system consisting of a context W and a component V, and we want to understand what happens when we replace V with another component U that is faster than V. We consider common notions of parallel composition, and show that timing anomalies can occur using our faster-than relation, even in the absence of non-determinism.

Submitted to:
QAPL 2019

This shows that timing anomalies are not caused by non-determinism, but arise from the linear timing behaviour of processes.

We then take a first step toward recovering compositional reasoning for the faster-than relation, by identifying decidable conditions sufficient for avoiding timing anomalies, which we call *strong monotonicity*.

Related Work. The notion of a faster-than relation has been studied in many different contexts throughout the literature. The work most closely related to ours is that of Pedersen et al. [15], which considers a generative version of the faster-than relation, whereas we study the reactive version. The focus of [15] is on decidability issues, and the faster-than relation is proved undecidable. However, positive results are also given in the form of an approximation algorithm, and a decidability result for unambiguous processes. Baier et al. [1] define, among other relations, a simulation relation for continuous-time Markov chains which can be interpreted as a faster-than relation, and study its logical characterisation. However, none of these works consider compositional aspects.

For process algebras, discrete-time faster-than relations have been defined for variations of Milner's CCS, and shown to be precongruences with respect to parallel composition [5, 10, 13, 17]. Lüttgen and Vogler [11] attempt to unify some of these process algebraic approaches and also consider the issue of parallel timing anomalies. For Petri nets, Vogler [22, 21] considers a testing preorder as a faster-than relation and shows that this is a precongruence with respect to parallel composition.

Work on timing anomalies date back to at least 1969 [6], but the most influential paper in the area is probably that of Lundqvist and Stenström [9], which shows that timing anomalies can occur in dynamically scheduled processors. More recent work has focused on compositional aspects [7] and defining timing anomalies formally, using transition systems as the formalism [3, 16].

2 Notation and Preliminaries

We fix some notation and recall concepts that are used throughout the rest of the paper. Let $\mathbb{N}$ denote the natural numbers and $\mathbb{R}_{\geq 0}$ denote the non-negative real numbers, which we equip with the standard Borel σ-algebra $\mathbb{B}$. For any set X, $\mathscr{D}(X)$ denotes the set of probability measures on X, and $\mathscr{D}_{\leq}(X)$ the set of subprobability measures on X. For an element $x \in X$, we will use δ_x to denote the Dirac measure at x. We fix a non-empty, countable set L of *labels* or *actions* and equip them with the discrete σ-algebra Σ_L.

For a probability measure $\mu \in \mathscr{D}(\mathbb{R}_{\geq 0})$, we denote by F_μ its *cumulative distribution function (CDF)* defined as $F_\mu(t) = \mu([0,t])$, for all $t \in \mathbb{R}_{\geq 0}$. We will denote by $Exp[\theta]$ the CDF of an exponential distribution with rate $\theta > 0$. The *convolution* of two probability measures $\mu, \nu \in \mathscr{D}(\mathbb{R}_{\geq 0})$, written $\mu * \nu$, is the probability measure on $\mathbb{R}_{\geq 0}$ given by $(\mu * \nu)(B) = \int_{-\infty}^{\infty} \nu(B - x)\, \mu(\mathrm{d}x)$, for all $B \in \mathbb{B}$ [2].

3 Semi-Markov Decision Processes

In this section we recall the definition of semi-Markov decision processes.

Definition 3.1. A *semi-Markov decision process (SMDP)* is a tuple $M = (S, \tau, \rho)$ where (1) S is a non-empty, countable set of *states*, (2) $\tau : S \times L \to \mathscr{D}_{\leq}(S)$ is a *transition probability function*, and (3) $\rho : S \to \mathscr{D}(\mathbb{R}_{\geq 0})$ is a *residence-time probability function*. ▲

Notice that Markov decision processes are a special case of SMDPs where for all $s \in S$, $\rho(s) = \delta_0$ (i.e. transitions happen instantaneously), and that continuous-time Markov decision processes are also a special case of SMDPs where, for all states $s \in S$, $F_{\rho(s)} = Exp[\theta_s]$ for some rate $\theta_s \in \mathbb{R}_{\geq 0}$.

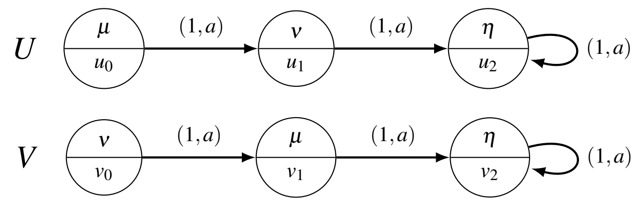

Figure 1: If $F_\mu(t) \geq F_\nu(t)$ for all t, then U is faster than V in the first states, and after that their probabilities are the same, so U is faster than V.

The executions of an SMDP $M = (S,\tau,\rho)$ are infinite timed transition sequences of the form $\pi = (s_1,a_1,t_1)(s_2,a_2,t_2)\cdots \in (S\times L\times \mathbb{R}_{\geq 0})^\omega$, representing the fact that M waited in state s_i for t_i time units after the action a_i was input. We will refer to executions of an SMDP as *paths*. We denote by $(\Pi(M),\Sigma)$ the *measurable space of paths*, where Σ is the smallest σ-algebra generated by the cylinders of the form $S_1\times L_1\times R_1,\dots,S_n\times L_n\times R_n$ for $S_i \in 2^S$, $L_i \in 2^L$, and $R_i \in \mathbb{B}$.

In this paper we assume that external choices are resolved by means of memoryless stochastic schedulers, however all the results we present still hold for schedulers that can remember actions and states.

Definition 3.2. Given an SMDP $M = (S,\tau,\rho)$, a *scheduler* for M is a function $\sigma : S \to \mathscr{D}(L)$ that assigns to each state a probability distribution over action labels. ▲

We will use the notation $\tau^\sigma(s,a)(s')$ as shorthand for $\tau(s,a)(s')\cdot\sigma(s)(a)$ to denote the probability of moving from state s to s' under the stochastic choice of a given by σ. Given an SMDP M and a scheduler σ for it, we define inductively the usual probability on cylinders $\mathbb{P}^\sigma_M(s)(S_1\times L_1\times R_1,\dots,S_n\times L_n\times R_n)$ which can be uniquely extended to the whole measurable space of paths[1]. For example, $\mathbb{P}^\sigma_M(s)(\{s'\}\times\{a\}\times[0,x],\{s''\}\times\{b\}\times\{[0,y]\})$ is the probability of doing an a-transition from s to s' within x time units and then doing a b-transition from s' to s'' within y time units.

4 A Faster-Than Relation

Intuitively, a process U is faster than V if it is able to execute any sequence of actions $a_1,\dots,a_n$ in less time than V. Since our systems are probabilistic, we need to consider the probability of their execution within any time bound. More precisely, for any time t and any scheduler σ for V, we must be able to find a scheduler σ' for U which allows U to execute any sequence of actions within time t with higher or equal probability than V. Hence, the type of events on which we want to focus are the following.

Definition 4.1. For any finite sequence of actions $a_1,\dots,a_n$, and $t\in\mathbb{R}_{\geq 0}$, we say that $\mathfrak{C}(a_1\dots a_n,t) = \{(s_1,b_1,t_1)(s_2,b_2,t_2)\cdots\in\Pi(M) \mid \forall 1\leq i\leq n,\ b_i = a_i \text{ and } \sum_{j=1}^n t_i \leq t\}$ is a *time-bounded cylinder*. The *length* of a time-bounded cylinder is the length of the sequence of actions in the cylinder. ▲

We will use the notation (M,s_0) to indicate that $M = (S,\tau,\rho)$ is an SMDP with initial state $s_0\in S$ and call it *pointed SMDP*. For the rest of the paper, we fix three SMDPs $M = (S,\tau,\rho)$, $U = (S_U,\tau_U,\rho_U)$, and $V = (S_V,\tau_V,\rho_V)$, with initial states $s_0\in S$, $u_0\in S_U$, $v_0\in S_V$, respectively.

Definition 4.2 (Faster-than)**.** We say that U is *faster than* V, written $U \preceq V$, if for all schedulers σ for V, time bounds t, and sequences of actions $a_1\dots a_n$, there exists a scheduler σ' for U such that $\mathbb{P}^{\sigma'}(u_0)(\mathfrak{C}(a_1\dots a_n,t)) \geq \mathbb{P}^\sigma(v_0)(\mathfrak{C}(a_1\dots a_n,t))$. ▲

[1] This is guaranteed by the Hahn-Kolmogorov theorem [19].

Example 4.3. Consider the pointed SMDPs (U,u_0) and (V,v_0) that are depicted in Figure 1. Assuming that $F_\mu(t) \geq F_\nu(t)$ for all t, we now show that $U \preceq V$.

(Case $n=1$**)** In this case we get $\mathbb{P}^\sigma(u_0)(\mathfrak{C}(a,t)) = F_\mu(t)$ and $\mathbb{P}^\sigma(v_0)(\mathfrak{C}(a,t)) = F_\nu(t)$. Since we assumed $F_\mu(t) \geq F_\nu(t)$ for all t, this implies $\mathbb{P}^\sigma(u_0)(\mathfrak{C}(a,t)) \geq \mathbb{P}^\sigma(v_0)(\mathfrak{C}(a,t))$.

(Case $n>1$**)** We have both $\mathbb{P}^\sigma(u_0)(\mathfrak{C}(a^n,t)) = (\mu * \nu * \eta^{*(n-2)})([0,t])$ and $\mathbb{P}^\sigma(v_0)(\mathfrak{C}(a^n,t)) = (\nu * \mu * \eta^{*(n-2)})([0,t])$, where η^{*n} is the n-fold convolution of η, defined inductively by $\eta^{*0} = \delta_0$ and $\eta^{*(n+1)} = \eta * \eta^{*n}$. Since convolution is commutative and associative, and δ_0 is the identity for convolution, we obtain $\mathbb{P}^\sigma(u_0)(\mathfrak{C}(a^n,t)) = \mathbb{P}^\sigma(v_0)(\mathfrak{C}(a^n,t))$. ♦

4.1 Comparison with Simulation and Bisimulation

The standard notions used to compare processes are bisimulation [14] and simulation [1]. Next we recall their definitions, naturally extended to our setting of SMDPs.

Definition 4.4. For an SMDP M, a relation $R \subseteq S \times S$ is a *bisimulation relation* (resp. *simulation relation*) on M if for all $(s_1,s_2) \in R$ we have (1) $F_{\rho(s_1)}(t) = F_{\rho(s_2)}(t)$ (resp. $F_{\rho(s_1)}(t) \leq F_{\rho(s_2)}(t)$) for all $t \in \mathbb{R}_{\geq 0}$ and (2) for all $a \in L$ there exists a weight function $\Delta_a : S \times S \to [0,1]$ 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 $s' \in S$.

If there is a bisimulation relation (resp. simulation relation) R such that $(s_1,s_2) \in R$, then we say that s_1 and s_2 are *bisimilar* (resp. s_2 *simulates* s_1) and write $s_1 \sim s_2$ (resp. $s_1 \precsim s_2$). ▲

We lift bisimulation and simulation relations to two different SMDPs by considering the disjoint union of the two and comparing their initial states. We denote by $\sim$ the largest bisimulation relation and by $\precsim$ the largest simulation relation. Furthermore, we say that U and V are *equally fast* and write $U \equiv V$ if $U \preceq V$ and $V \preceq U$.

Theorem 4.5. *$\precsim$ and $\preceq$ are incomparable, $\sim$ and $\preceq$ are incomparable, and we have $\sim \not\subseteq \equiv$.*

Theorem 4.5 holds both for memoryless and for memoryful schedulers that can remember actions and states, but it is still unclear what happens if schedulers are allowed to remember time.

5 Compositionality

Next we introduce the notion of composition of SMDPs. As argued in [18], the style of synchronous CSP composition is the most natural one to consider for reactive probabilistic systems, so this is the one we will adopt. However, we leave the composition of the residence-times as a parameter, so that we can compare different kinds of composition.

Definition 5.1. A function $\star : \mathscr{D}(\mathbb{R}_{\geq 0}) \times \mathscr{D}(\mathbb{R}_{\geq 0}) \to \mathscr{D}(\mathbb{R}_{\geq 0})$ is called a *residence-time composition function* if it is commutative, i.e. $\star(\mu,\nu) = \star(\nu,\mu)$ for all $\mu,\nu \in \mathscr{D}(\mathbb{R}_{\geq 0})$. ▲

In this paper we consider the three following composition functions:

Minimum composition: $F_{\star(\mu,\nu)}(t) = \min\{F_\mu(t), F_\nu(t)\}$

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

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

Note that product composition only makes sense when we know that the residence-time distributions are exponential, such as in the case of continuous-time Markov decision processes.

Definition 5.2. Let $\star$ be a residence-time composition function. Then the $\star$-*composition* of U and V, denoted by $U \parallel_\star V = (S,\tau,\rho)$, is given by (1) $S = U \times V$, (2) $\tau((u,v),a)((u',v')) = \tau_U(u,a)(u') \cdot \tau_V(v,a)(v')$ for all $a \in L$ and $(u',v') \in S$, and (3) $\rho((u,v)) = \star(\rho_U(u), \rho_V(v))$. ▲

We write $u \parallel_\star v$ to denote the composite state (u,v) of $U \parallel_\star V$ where $u \in S_U$ and $v \in S_V$.

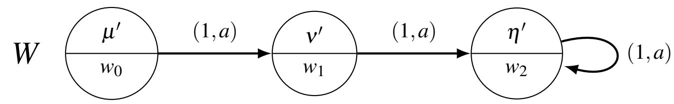

Figure 2: For different instantiations of μ', ν', and η', the context W leads to parallel timing anomalies for product, minimum, and maximum rate composition, respectively.

5.1 Parallel Timing Anomalies

If we have two components U and V, and we know that U is faster than V, then if V is in parallel with some context W, we would expect this composition to become faster when we replace the component V with the component U. However, sometimes this fails to happen, and we will call such an occurrence a *parallel timing anomaly*. Consider the two SMDPs U and V depicted in Figure 1. For the examples in this section, let $F_\mu = Exp[2]$, $F_\nu = Exp[0.5]$, and let $\eta = \eta'$ be arbitrary. From Example 4.3 we know that $U \preceq V$.

Example 5.3 (Product composition)**.** Let $F_{\mu'} = Exp[10]$ and $F_{\nu'} = Exp[0.1]$. In $U \parallel_\star W$, the rates in the first two states will then be 20 and 0.05, and in $V \parallel_\star W$ they will be 5 and 0.2. Consider the time-bounded cylinder $\mathfrak{C}(aa,2)$. Then we see that $\mathbb{P}(u_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.09$ and $\mathbb{P}(v_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.30$, showing that $U \parallel_\star W \not\preceq V \parallel_\star W$. Hence we have a parallel timing anomaly. ♦

Example 5.4 (Minimum composition)**.** Let $F_{\mu'} = Exp[1]$ and $F_{\nu'} = Exp[2]$. The rates of $U \parallel_\star W$ are then 1 and 0.5, whereas they are 0.5 and 2 in $V \parallel_\star W$. Then we get $\mathbb{P}(u_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.40$ and $\mathbb{P}(v_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.51$, so $U \parallel_\star W \not\preceq V \parallel_\star W$. ♦

Example 5.5 (Maximum composition)**.** Let $F_{\mu'} = Exp[2]$ and $F_{\nu'} = Exp[1]$. $U \parallel_\star W$ then has rates 2 and 1, and $V \parallel_\star W$ has rates 2 and 2. Then $\mathbb{P}(u_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.75$ and $\mathbb{P}(v_0 \parallel_\star w_0)(\mathfrak{C}(aa,2)) \approx 0.91$, so $U \parallel_\star W_3 \not\preceq V \parallel_\star W_3$. ♦

5.2 Avoiding Parallel Timing Anomalies

We have seen in the previous section that parallel timing anomalies can occur. We now wish to understand what kind of contexts do not lead to timing anomalies. In this section we assume that the set L of transition labels is a finite set. Also, we fix a residence-time composition function $\star$ and two additional SMDPs $(W, w_0) = (S_W, \tau_W, \rho_W)$ and $(W', w'_0) = (S_{W'}, \tau_{W'}, \rho_{W'})$ which should be thought of as contexts.

We first give conditions that over-approximate the faster-than relation between the composite systems by requiring that when U and W are put in parallel, then the composite system is point-wise faster than U along all paths. Likewise, we require that when V and W are put in parallel, the composite system is point-wise slower than V along all paths. If we already know that U is faster than V, this will imply by transitivity that $U \parallel_\star W$ is faster than $V \parallel_\star W$. We have already seen in Example 4.3 that a process U need not be point-wise faster than V along all paths in order for U to be faster than V. However, by imposing this condition, we do not need to compare convolutions of distributions, but can compare the distributions directly.

We will say that an SMDP M has a *deterministic Markov kernel* if for all states s and labels a, there is at most one state s' such that $\tau(s,a)(s') > 0$. A *state path* of length n in M is a sequence $s_1 \dots s_n$ such that s_1 is the initial state and for all i there exists $a \in L$ such that $\tau(s_i,a)(s_{i+1}) > 0$.

Definition 5.6. We say that $\star$ is *strongly n-monotonic* in U, V, W, and W' and write $(U,W) \leqq^n_\star (V,W')$ if W' has a deterministic Markov kernel and for all state paths $u_1 \dots u_n$, $v_1 \dots v_n$, $w_1 \dots w_n$, and $w'_1 \dots w'_n$,

- $F_{\rho(u_i \|_\star w_i)}(t) \geq F_{\rho_U(u_i)}(t)$ and $F_{\rho_V(v_i)}(t) \geq F_{\rho(v_i \|_\star w'_i)}(t)$ for all $t \in \mathbb{R}_{\geq 0}$ and $1 \leq i \leq n$
- for all schedulers σ_U for U and $\sigma_{U,W}$ for $U \|_\star W$, $\tau^{\sigma_{U,W}}(u_i \|_\star w_i, a)(u_{i+1} \|_\star w_{i+1}) \geq \tau_U^{\sigma_U}(u_i, a)(u_{i+1})$
- for all schedulers $\sigma_{V,W'}$ for $V \|_\star W'$ and σ_V for V, $\tau_V^{\sigma_V}(v_i, a)(v_{i+1}) \geq \tau^{\sigma_{V,W'}}(v_i \|_\star w'_i, a)(v_{i+1} \|_\star w'_{i+1})$

for all $a \in L$ and $1 \leq i < n$. If $(U,W) \leqq_\star^n (V,W')$ for all $n \in \mathbb{N}$, we say that $\star$ is *strongly monotonic* in U, V, W, and W' and write $(U,W) \leqq_\star (V,W')$. ▲

Example 5.7. Let U and V be given by Figure 1 with $F_\mu \geq F_\nu$ as in Example 4.3. Let $\star$ be minimum rate composition and consider the context W from Figure 2, where $\mu' = \mu$, $\nu' = \nu$, and $\eta' = \eta$. There is only one possible scheduler σ, which is the Dirac measure at a, and hence it is clear that the second and third conditions are satisfied. We also find that $F_{\rho(u_i \|_\star w_i)}(t) = F_{\rho_U(u_i)}(t)$ and $F_{\rho_V(v_i)}(t) = F_{\rho(v_i \|_\star w_i)}(t)$, so the first condition is also satisfied. Hence we conclude $(U,W) \leqq_\star (V,W)$. ◆

Notice that the examples in Section 5.1 are not strongly monotonic. The following is the main theorem of our paper, showing that strong monotonicity implies the absence of parallel timing anomalies.

Theorem 5.8. *If $(U,W) \leqq_\star (V,W')$ as well as $U \preceq V$ and $W \preceq W'$, then $U \|_\star W \preceq V \|_\star W'$.*

We now wish to show that it is decidable whether $(U,W) \leqq_\star (V,W')$ for finite SMDPs, thereby giving a decidable condition for avoiding timing anomalies. To do this, we first show that in order to establish strong monotonicity, it is enough to consider paths up to length $m = \max\{|S_U| \cdot |S_W|, |S_V| \cdot |S_{W'}|\} + \max\{|S_U|, |S_V|, |S_W|, |S_{W'}|\} + 1$, due to the fact that they start repeating.

Lemma 5.9. *Let U, V, W, and W' be finite. If $(U,W) \leqq_\star^m (V,W')$, then $(U,W) \leqq_\star (V,W')$.*

Theorem 5.10. *Let U, V, W, and W' be finite. If for all state paths $u_1 \dots u_n$, $v_1 \dots v_n$, $w_1 \dots w_n$, and $w'_1 \dots w'_n$ we have that $\{t \in \mathbb{R}_{\geq 0} \mid F_{\rho(u_i \|_\star w_i)}(t) \geq F_{\rho_U(u_i)}(t)\}$ and $\{t \in \mathbb{R}_{\geq 0} \mid F_{\rho_V(v_i)}(t) \geq F_{\rho(v_i \|_\star w'_i)}(t)\}$ are semialgebraic sets for all $1 \leq i \leq m$, then it is decidable whether $(U,W) \leqq_\star (V,W')$.*

For uniform and exponential distributions with minimum or maximum composition, the corresponding sets are all semialgebraic, and the same is true for exponential distributions with product composition.

It turns out that strong monotonicity implies the absence of non-determinism.

Proposition 5.11. *If $(U,W) \leqq_\star (V,W')$, then L is a singleton set or u_0 is a deadlock state, i.e. the transition probability is zero from u_0 to any other state.*

However, strong monotonicity still makes sense as a condition, since all our examples of timing anomalies in Section 5.1 have no non-determinism.

6 Conclusion

In this paper, we have investigated the notion of a process being faster than another process in the context of semi-Markov decision processes. We have given a trace-based definition of a faster-than relation, and shown how this notion relates to the usual notions of simulation and bisimulation. By considering composition as being parametric in how the residence times of states are combined, we have given examples showing that our faster-than relation gives rise to parallel timing anomalies for many of the popular ways of composing rates. We have therefore given sufficient conditions for how such parallel timing anomalies can be avoided, and we have shown that these conditions are decidable.

While the conditions for strong monotonicity are so strict that they rule out non-determinism, it seems difficult to obtain better results that are still decidable. This is because the faster-than relation itself is undecidable whenever non-determinism is introduced, and even without non-determinism, it has close connections to a long-standing open problem in number theory [15], and its decidability status for a single label is therefore still open. Because of this, an interesting future work direction is to explore the boundaries of decidability for conditions to avoid timing anomalies.

References

[1] Christel Baier, Joost-Pieter Katoen, Holger Hermanns & Verena Wolf (2005): *Comparative branching-time semantics for Markov chains*. *Inf. Comput.* 200(2), pp. 149–214, doi:10.1016/j.ic.2005.03.001.

[2] Patrick Billingsley (1995): *Probability And Measure*, 3rd edition. Wiley-Interscience.

[3] Franck Cassez, René Rydhof Hansen & Mads Chr. Olesen (2012): *What is a Timing Anomaly?* In Tullio Vardanega, editor: *12th International Workshop on Worst-Case Execution Time Analysis, WCET 2012, July 10, 2012, Pisa, Italy*, *OASICS* 23, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 1–12, doi:10.4230/OASIcs.WCET.2012.1.

[4] Edmund M. Clarke, David E. Long & Kenneth L. McMillan (1989): *Compositional Model Checking*. In: *Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS '89), Pacific Grove, California, USA, June 5-8, 1989*, IEEE Computer Society, pp. 353–362, doi:10.1109/LICS.1989.39190.

[5] Flavio Corradini, Roberto Gorrieri & Marco Roccetti (1995): *Performance Preorder: Ordering Processes with Respect to Speed*. In Jirí Wiedermann & Petr Hájek, editors: *Mathematical Foundations of Computer Science 1995, 20th International Symposium, MFCS'95, Prague, Czech Republic, August 28 - September 1, 1995, Proceedings*, *Lecture Notes in Computer Science* 969, Springer, pp. 444–453, doi:10.1007/3-540-60246-1_150.

[6] Ronald L. Graham (1969): *Bounds on Multiprocessing Timing Anomalies*. *SIAM Journal of Applied Mathematics* 17(2), pp. 416–429. Available at `http://www.jstor.org/stable/2099572`.

[7] Raimund Kirner, Albrecht Kadlec & Peter P. Puschner (2009): *Precise Worst-Case Execution Time Analysis for Processors with Timing Anomalies*. In: *21st Euromicro Conference on Real-Time Systems, ECRTS 2009, Dublin, Ireland, July 1-3, 2009*, IEEE Computer Society, pp. 119–128, doi:10.1109/ECRTS.2009.8.

[8] Edward A. Lee (2008): *Cyber Physical Systems: Design Challenges*. In: *11th IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2008), 5-7 May 2008, Orlando, Florida, USA*, IEEE Computer Society, pp. 363–369, doi:10.1109/ISORC.2008.25.

[9] Thomas Lundqvist & Per Stenström (1999): *Timing Anomalies in Dynamically Scheduled Microprocessors*. In: *Proceedings of the 20th IEEE Real-Time Systems Symposium, Phoenix, AZ, USA, December 1-3, 1999*, IEEE Computer Society, pp. 12–21, doi:10.1109/REAL.1999.818824.

[10] Gerald Lüttgen & Walter Vogler (2001): *A Faster-than Relation for Asynchronous Processes*. In Kim Guldstrand Larsen & Mogens Nielsen, editors: *CONCUR 2001 - Concurrency Theory, 12th International Conference, Aalborg, Denmark, August 20-25, 2001, Proceedings*, *Lecture Notes in Computer Science* 2154, Springer, pp. 262–276, doi:10.1007/3-540-44685-0_18.

[11] Gerald Lüttgen & Walter Vogler (2006): *Bisimulation on speed: A unified approach*. *Theor. Comput. Sci.* 360(1-3), pp. 209–227, doi:10.1016/j.tcs.2006.03.004.

[12] Sneana Maovi, Saa Stoi & Rade Hajdin (2015): *Application of Semi-Markov Decision Process in Bridge Management*. *IABSE Symposium Report* 105(28), pp. 1–8, doi:doi:10.2749/222137815818358204.

[13] Faron Moller & Chris M. N. Tofts (1991): *Relating Processes With Respect to Speed*. In Jos C. M. Baeten & Jan Friso Groote, editors: *CONCUR '91, 2nd International Conference on Concurrency Theory, Amsterdam, The Netherlands, August 26-29, 1991, Proceedings*, *Lecture Notes in Computer Science* 527, Springer, pp. 424–438, doi:10.1007/3-540-54430-5_104.

[14] Martin R. Neuhäußer & Joost-Pieter Katoen (2007): *Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes*. In Luís Caires & Vasco Thudichum Vasconcelos, editors: *CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings*, *Lecture Notes in Computer Science* 4703, Springer, pp. 412–427, doi:10.1007/978-3-540-74407-8_28.

[15] Mathias Ruggaard Pedersen, Nathanaël Fijalkow, Giorgio Bacci, Kim G. Larsen & Radu Mardare (2018): *Timed Comparisons of Semi-Markov Processes*. In Shmuel Tomi Klein, Carlos Martín-Vide & Dana Shapira, editors: *Language and Automata Theory and Applications - 12th International Conference, LATA 2018,*

Ramat Gan, Israel, April 9-11, 2018, Proceedings, *Lecture Notes in Computer Science* 10792, Springer, pp. 271–283, doi:10.1007/978-3-319-77313-1_21.

[16] Jan Reineke, Björn Wachter, Stephan Thesing, Reinhard Wilhelm, Ilia Polian, Jochen Eisinger & Bernd Becker (2006): *A Definition and Classification of Timing Anomalies*. In Frank Mueller, editor: *6th Intl. Workshop on Worst-Case Execution Time (WCET) Analysis, July 4, 2006, Dresden, Germany*, *OASICS* 4, Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany. Available at `http://drops.dagstuhl.de/opus/volltexte/2006/671`.

[17] Ichiro Satoh & Mario Tokoro (1994): *A Formalism for Remotely Interacting Processes*. In Takayasu Ito & Akinori Yonezawa, editors: *Theory and Practice of Parallel Programming, International Workshop TPPP'94, Sendai, Japan, November 7-9, 1994, Proceedings*, *Lecture Notes in Computer Science* 907, Springer, pp. 216–228, doi:10.1007/BFb0026571.

[18] Ana Sokolova & Erik P. de Vink (2004): *Probabilistic Automata: System Types, Parallel Composition and Comparison*. In Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, Joost-Pieter Katoen & Markus Siegle, editors: *Validation of Stochastic Systems - A Guide to Current Research*, *Lecture Notes in Computer Science* 2925, Springer, pp. 1–43, doi:10.1007/978-3-540-24611-4_1.

[19] Terence Tao (2013): *An Introduction to Measure Theory*. Graduate studies in mathematics, American Mathematical Society.

[20] O. S. Thomas & J. O. Sobanjo (2016): *Semi-Markov Decision Process: A Decision Tool for Transportation Infrastructure Management Systems*. In: *International Conference on Transportation and Development*, doi:10.1061/9780784479926.036.

[21] Walter Vogler (1995): *Timed Testing of Concurrent Systems*. *Inf. Comput.* 121(2), pp. 149–171, doi:10.1006/inco.1995.1130.

[22] Walter Vogler (2003): *Faster asynchronous systems*. *Inf. Comput.* 184(2), pp. 311–342, doi:10.1016/S0890-5401(03)00065-8.