Comparing the speed of probabilistic processes

Mathias R. Pedersen, Nathanaël Fijalkow, Giorgio Bacci,
Kim G. Larsen, and Radu Mardare

1 Introduction

In distributed, embedded, or real-time systems, non-functional requirements are becoming increasingly important. Many of these requirements, such as response time and throughput, depend heavily on the timing behaviour of the system in question. It is therefore crucial to understand and be able to compare the timing behaviour of different systems. To that end, simulations and bisimulations and their relations to logic for continuous-time Markov chains have been thoroughly studied [2], and bisimulation has been extended to the metric setting to allow for more quantitative analysis of systems [1]. Some work has also been done on extending these results to the setting of continuous-time Markov decision processes [6], which also include non-determinism.

On the other hand, the notions of trace equivalence and inclusion for these systems have received less attention. Trace equivalences have been studied for continuous-time Markov chains using the notion of a trace machine [7] which can observe the system in specific ways by pressing buttons. This approach was also extended to continuous-time Markov decision processes [5].

We will study a simple and intuitive notion of trace equivalence and inclusion, which we will call the equally-fast relation and the faster-than relation, respectively. When considering timed traces of system executions, comparing each individual time delay may be too coarse, and in many applications, what one is really interested in is the accumulated amount of time that some sequence of action takes. Hence the idea behind our relations will be to compare the probability of doing some sequence of actions within some given time bound.

2 The equally-fast and faster-than relations

We will consider generative semi-Markov processes where the time it takes before a transition is taken from a state is given by an arbitrary distribution function. If we take as distribution functions exponential distributions, then we of course obtain continuous-time Markov chains, but many other kinds of timing behaviour can be specified by using other distribution functions. If we denote by $\mathbb{P}_s(w,t)$ the probability of doing the sequence of actions given by the word w within time t starting in the state s, then we will say that two states s_1 and s_2 are equally fast if $\mathbb{P}_{s_1}(w,t) = \mathbb{P}_{s_2}(w,t)$ for all words w and all points in time t, and likewise we will say that s_1 is faster than s_2 if $\mathbb{P}_{s_1}(w,t) \geq \mathbb{P}_{s_2}(w,t)$ for all w and t. This definition can easily be extended to decision processes by requiring that for any scheduler for one component, there exists a scheduler for the other component such that the relationship holds.

One important thing to note is that the relations do not require that the relationship also holds for each step in executing the sequence of actions. In particular, if s_1 is faster than s_2, then it is possible that for the successor states s_1' and s_2', s_1' is slower than s_2', as long as s_1 is so much faster than s_2 that it makes up for s_1' being slower than s_2'.

3 Undecidability results

Although both the faster-than and equally-fast relations are intuitive and simple to state, the faster-than relation turns out to be highly undecidable, which can be shown by a reduction from the universality problem for probabilistic automata.

Theorem 1 *The faster-than relation is undecidable.*

Using a result of Condon and Lipton [3] for probabilistic automata, we can even show that is impossible to approximately decide the faster-than relation.

Theorem 2 *Approximating the faster-than relation is impossible.*

Even if we only allow a single output action, the problem of deciding the faster-than relation remains at least as hard as that of the positivity problem for linear recurrence sequences, which has been open for at least 30 years [4].

Theorem 3 *When only a single output action is available, deciding the faster-than relation is positivity-hard.*

4 Decidability results

Not all is bleak, however, as we can also give some decidability results if we restrict ourselves to certain classes of timing behaviour. The following results therefore hold when the timing behaviour of the process is given either by piecewise polynomials or by exponential distributions.

As a contrast to the undecidability of the faster-than relation, the equally-fast relation is in fact decidable, using facts about linear algebra.

Theorem 4 *The equally-fast relation is decidable.*

Although the faster-than relation is undecidable, there are some restrictions that allow us to recover decidability. The first one is when for every state, every output action leads to a unique state. We call such processes unambiguous.

Theorem 5 *For unambiguous semi-Markov processes, the faster-than relation is decidable.*

The second one is to again consider approximations, but this time only compare the two probabilities for all $t \leq b$, where b is some given time bound.

Theorem 6 *Approximating the faster-than relation up to a given time bound is possible.*

5 Open problems

The results we have shown are for generative semi-Markov processes. All of the undecidability results hold also for the reactive case, but the decidability results do not carry over easily. It is therefore still unclear whether the problems that are decidable for the generative case are also decidable for the reactive case, or whether they remain undecidable.

The question of decidability is only a first step towards understanding the faster-than and equally-fast relations. The precise complexity of the decidable problem is still unknown. Furthermore, we still do not understand the logical aspects of these relations well, in particular a logical characterisation seems difficult. We hope to be able to show a logical preservation result for an LTL-like logic, and later extend this to a full logical characterisation. Another key aspect is compositionality. We have shown that when using compositional techniques from process algebra, the faster-than relation is not a precongruence for continuous-time Markov chains. However, we are still trying to understand more precisely under what circumstances we can get such a precongruence.

REFERENCES

[1] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. On-the-Fly Computation of Bisimilarity Distances. *LMCS*, Volume 13, Issue 2, June 2017.

[2] Christel Baier, Joost-Pieter Katoen, Holger Hermanns, and Verena Wolf. Comparative branching-time semantics for markov chains. *Inf. Comput.*, 200(2):149–214, 2005.

[3] Anne Condon and Richard J. Lipton. On the complexity of space bounded interactive proofs (extended abstract). In *FOCS*, 1989.

[4] Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *SODA*, 2014.

[5] Arpit Sharma. Trace relations and logical preservation for continuous-time markov decision processes. *ICTAC 2017, Hanoi, Vietnam*, pages 192–209, 2017.

[6] Lei Song, Lijun Zhang, and Jens Chr. Godskesen. Bisimulations and logical characterizations on continuous-time markov decision processes. *VMCAI 2014, San Diego*, pages 98–117, 2014.

[7] Verena Wolf, Christel Baier, and Mila E. Majster-Cederbaum. Trace machines for observing continuous-time markov chains. *Electr. Notes Theor. Comput. Sci.*, 153(2):259–277, 2006.