Project no.: **ICT-FP7-STREP-214755**

Project full title: **Quantitative System Properties in Model-Driven Design**

Project Acronym: **QUASIMODO**

Deliverable no.: D2.2

Title of Deliverable: Symbolic data structures and analysis of models with multiple quantitative aspects

Contractual Date of Delivery to the CEC:	**Month 18**
Actual Date of Delivery to the CEC:	**Month 18 (July 1, 2009)**
Organisation name of lead contractor for this deliverable:	**CNRS**
Author(s):	**Nicolas Markey, Jasper Berendsen, Alexandre David, Tingting Han, Holger Hermanns, Kim G. Larsen, Martin Neuhäußer.**
Participant(s):	**P01 AAU, P02 ESI, P03 CNRS, P04 RWTH, P05 SU**
Work package contributing to the deliverable:	**WP 2**
Nature:	**R+P**
Version:	**0.99**
Total number of pages:	**15**
Start date of project:	1 Jan. 2008 **Duration:** 36 month

Project co-funded by the European Commission within the Seventh Framework Programme (2007-2013)

Dissemination Level

PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Abstract:

This deliverable reports on the works carried out inside the QUASIMODO consortium on the study of models involving several quantitative aspects. In most cases, the primary quantitative aspect is time; timed models are then extended with costs (for measuring some quantity such as energy, level of water in a tank, ...), probabilities or games (for coping with uncertainty in the description of the model), or a combination of them.

Keyword list: priced timed automata, Markov chains, Markov decision processes.

Contents

Introduction

Markov chains and timed automata are two fundamental quantitative models that are widely used in the field of formal verification (especially for model checking), as they provide a way of modelling quantitative aspects of models while remaining tractable (at least from a practical point of view).

We study several quantitative extensions of these quantitative models: continuous-time Markov chains, priced timed automata, and probabilistic priced timed automata.

Discrete- and continuous-time Markov chains (DTMCs/CTMCs) and Markov decision processes (DTMDPs/CTMDPs)

Discrete- and continuous-time Markov chains are one of the most important models in performance and dependability analysis. They are exploited in a broad range of applications, and constitute the underlying semantical model of a plethora of modeling formalisms for real-time probabilistic systems.

Discrete- and continuous-time Markov decision processes, also known as *controlled* Markov chains, have been used for, among others, the control of queueing systems, epidemic, and manufacturing processes. The analysis of Markov decision processes is focused on determining optimal schedulers for criteria such as expected total reward and expected (long-run) average reward.

Priced timed automata (PTAs)

Priced timed automata extend timed automata with information on the *price* to pay for delaying in a location or firing a transition. This real-valued variable can be used to represent for instance energy consumption (or harvesting), level of oil in a pump, memory usage, bandwidth consumption, ...

Probabilistic priced timed automata (PPTAs)

While both extensions (with probabilities and with prices) already raise many difficult theoretical as well as practical questions, the Quasimodo consortium has also introduced and studied a combination of both: PPTAs include timing constraints in the behaviour of the model, probabilistic choices on the transitions, and prices for quantitative evaluation of the model.

In the sequel, we present our recent results on all three models, both on the analysis of these models and on the development of symbolic data structurtes to handle them in practice.

1 Model-checking discrete- and continuous-time Markov chains and Markov decision processes

1.1 Probabilistic Reachability for Parametric Markov Models

1.1.1 Participants

- Lijun Zhang, Ernst Moritz Hahn and Holger Hermanns, Saarland University, Germany.

1.1.2 Contributions

Discrete time Markov chains (DTMCs) have been applied successfully to reason about quantitative properties in a large number of areas such as computer science, engineering, mathematics, biological systems. Often, several variants of a probabilistic model are of interest. For example, it would be interesting to evaluate several variants of sensor networks with different reliabilities of the wireless connection, without doing a complete analysis for each instance.

We call a DTMC in which certain probabilities or other properties are not fixed but given as parameters of the model a *parametric* DTMC (PDTMC). An analysis of a PDTMC results in a closed-form solution in form of a function in the parameters. Given such a function f, we could also analyze properties of the function itself. If f represents the probability of a certain set of goal states, we could find the parameter values which maximize f to obtain the optimal parameters, without having to do large numbers of costly analyzes to estimate this point.

The efficient analysis of PDTMCs is involved and different approaches than the well-known ones for the analysis of DTMC have to be taken. Our goal is to nevertheless develop an efficient and effective algorithm for PDTMCs and related models.

We have developed algorithms for PDTMCs [HHZ09], based on a variant of the state elimination algorithm. It computes the parametric unbounded reachability probability from the initial state of the PDTMC to a set of target states. The state elimination algorithms is a standard means to derive a regular expression from a finite automaton, by eliminating its states except the initial and final one, while relabeling its transitions by regular expressions instead of just elements of the alphabet. In our adaption, instead of having transitions labeled with regular expressions, we label them with functions of the model parameters into probabilities. Finally, we can obtain the function we wanted to obtain from the only transition remaining in the PDTMC.

We also have an initial approach for models involving nondeterminism. There, we replace nondeterminism by parametric probabilistic choice. This method works well for special cases, as seen in the paper.

In a further extension, we extend our method to compute the expected parametric reward till a set of target states is reached. Rewards are costs or bonuses (depending on the interpretation) obtained from entering a state of the PDTMC or taking a transition from a state to another state. Such reward properties play a crucial role for the estimation of performance properties of probabilistic systems.

The analysis of PDTMCs is more expensive than the analysis of usual DTMCs. Therefore, we use a precomputation to reduce the number of states. This has a great impact of the overall

performance of the method.

The algorithms described here have been implemented in the tool PARAM. Using a number of case studies, we have shown the feasibility of our approach.

- Crowds Protocol [RR98]: an information exchange protocol with aims at protecting the anonymity of its users. We considered the degree of anonymity guarantees possible to users parametric in the number of attackers.
- Zeroconf [BvdSHV03]: a self-configuring network protocol. We considered a variant parametric in the number of possible network addresses. The property under consideration is the probability of duplicate choice of the same address.
- Cyclic Polling Server [IT90]: This model consists of a number of stations which are handled by a polling server. We considered the probability that a certain station is served first, parametric in the speed with which the server works and the rate with which requests are generated.
- Randomized Mutual Exclusion [PZ86]: a variant of the well-known mutual exclusion protocol where processes decide probabilistically whether they will try to enter the critical section in their next step. We compute the expected number of times the processes try to enter the critical section, parametric in the probability that they try it.
- Bounded Retransmission Protocol [HSV94]: a message transfer protocol to transfer data over unreliable channels. Our variant is parametric in the reliabilities of the channels. The property we consider is the maximal probability that the sender of data in this protocol does not finally finish the transmission.

In all of the above case studies, we were able to effectively analyze the parametric properties we wanted to consider. For most of them, it turned out crucial to use preprocessing for state space reduction. We consider the fact that the models under consideration were taken from diverse areas and are very different from each other an indication of the general applicability of our method.

As future work, we are investigating improvements of the implementation with respect performance, especially for the setting with nondeterminism. Additionally, we plan to look into continuous time models with clocks and rewards. Other possible directions include the use of symbolic model representations, such as advanced representations of state spaces. We also want to explore model checking for interval Markov chains.

References

[BvdSHV03] H. Bohnenkamp, P. van der Stok, H. Hermanns, and F. Vaandrager. Cost-optimization of the IPv4 zeroconf protocol. In *DSN*, pages 531–540, 2003.

[HHZ09] E. Moritz Hahn, Holger Hermanns, and Lijun Zhang. Probabilistic reachability for parametric Markov models. In *SPIN*, volume 5578 of *Lecture Notes in Computer Science*, pages 88–106. Springer, 2009.

[HSV94] L. Helmink, M. Sellink, and F. Vaandrager. Proof-checking a data link protocol. In *TYPES*, volume 806 of *Lecture Notes in Computer Science*, pages 127–165. Springer, 1994.

[IT90] O. Ibe and K. Trivedi. Stochastic Petri net models of polling systems. *IEEE Journal on Selected Areas in Communications*, 8(9):1649–1657, 1990.

[PZ86] A. Pnueli and L. Zuck. Verification of multiprocess probabilistic protocols. *Distrib. Comput.*, 1(1):53–72, 1986.

[RR98] M. K. Reiter and A. D. Rubin. Crowds: anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1):66–92, 1998.

1.2 Quantitative model checking of continuous-time Markov chains against timed automata specifications

1.2.1 Participants

- Taolue Chen, ESI, University of Twente, the Netherlands;
- Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre, RWTH Aachen, Germany.

1.2.2 Contributions

CTMC model-checking has been focused on the *branching-time* temporal logic CSL (Continuous Stochastic Logic) [ASSB00, BHHK03]. One of the key ingredients of model checking CTMC against CSL is that reachability probabilities for time-bounded until-formulae can be approximated arbitrarily closely by a reduction to transient analysis in CTMCs. This results in a polynomial-time algorithm that has been realized in model checkers such as PRISM and MRMC.

Our work concerns the problem of verifying CTMCs versus *linear* real-time specifications, which are based on timed automata. Concretely speaking, we explore the following problem: given a CTMC $\mathcal{C}$, and a linear real-time property provided as a *deterministic timed automaton* [AD94] (DTA) $\mathcal{A}$, what is the probability of the set of paths of $\mathcal{C}$ which are accepted by $\mathcal{A}$ ($\mathcal{C} \models \mathcal{A}$)? We set off to show that this problem is well-defined in the sense that the path set is *measurable*. Computing its probability can then be reduced to computing the reachability probability in a *piecewise deterministic Markov process* (PDP) [Dav93], a model that is used in, e.g., stochastic control theory and financial mathematics. This result relies on a product construction of CTMC $\mathcal{C}$ and DTA $\mathcal{A}$, denoted $\mathcal{C} \otimes \mathcal{A}$, yielding *deterministic Markov timed automata* (DMTA), a variant of DTA in which, besides the usual ingredients of timed automata, like guards and clock resets, the location residence time is exponentially distributed. We show that the probability of

$\mathcal{C} \models \mathcal{A}$ coincides with the reachability probability of accepting paths in $\mathcal{C} \otimes \mathcal{A}$. The underlying PDP of a DMTA is obtained by a slight adaptation of the standard region construction. The desired reachability probability is characterized as the least solution of a system of *integral equations* that is obtained from the PDP. Finally, this probability is shown to be approximated by solving a system of *partial differential equations* (PDEs). For single-clock DTA, we show that the system of integral equations can be transformed into a system of *linear equations*, where the coefficients are solutions of some *ordinary differential equations* (ODEs), which can either have an analytical solution (for small state space) or an arbitrarily closely approximated solution efficiently.

Related work is model checking of asCSL [BCH+07] and CSL^{TA} [DHS09]. asCSL allows to impose a time constraint on action sequences described by regular expressions, while in CSL^{TA}, time constra ints (of until modalities) are specified by *single-clock* DTA. Compared to [DHS09], our approach does not restrict the number of clocks and supports more specifications than CSL^{TA}. For the single-clock case, our approach produces the same result as [DHS09], but yields a conceptually simpler formulation whose correctness can be derived from the simplification of the system of integral equations obtained in the general case. Moreover, measurability has not been addressed in [DHS09].

References

[AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.

[ASSB00] Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert K. Brayton. Model-checking continous-time Markov chains. *ACM Trans. Comput. Log.*, 1(1):162–170, 2000.

[BCH+07] Christel Baier, Lucia Cloth, Boudewijn R. Haverkort, Matthias Kuntz, and Markus Siegle. Model checking Markov chains with actions and state labels. *IEEE Trans. Software Eng.*, 33(4):209–224, 2007.

[BHHK03] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Software Eng.*, 29(6):524–541, 2003.

[Dav93] Mark H. A. Davis. *Markov Models and Optimization*. Chapman & Hall, 1993.

[DHS09] Susanna Donatelli, Serge Haddad, and Jeremy Sproston. Model checking timed and stochastic properties with CSL^{TA}. *IEEE Trans. Softare Eng.*, 35(2):224–240, 2009.

1.3 Delayed Nondeterminism in Continuous-Time Markov Decision Processes

1.3.1 Participants

- Martin R. Neuhäußer, Joost-Pieter Katoen, RWTH Aachen, Germany;
- Mariëlle Stoelinga, ESI, University of Twente, the Netherlands.

1.3.2 Contribution

In general, the delay to jump to a next state in a CTMDP is determined by the action selected by the scheduler on entering the current state. We investigate under which conditions this resolution of nondeterminism may be deferred. Rather than focusing on a specific objective, we consider generic (measurable) properties. The core of our study is a transformation—called local uniformization—on CTMDPs which unifies the speed of outgoing transitions per state. Whereas classical uniformization [Gra91, GM84, Jen53] adds self-loops to achieve this, local uniformization uses auxiliary copy-states. In this way, we enforce that schedulers in the original and uniformized CTMDP have (for important scheduler classes) the same power, whereas classical loop-based uniformization allows a scheduler to change its decision when re-entering a state through the added self-loop. As a result, locally uniform CTMDPs allow to defer the resolution of nondeterminism for important scheduler classes, by dissolving the intrinsic dependency between state residence times and the scheduler. Thus, they can be viewed as MDPs with exponentially distributed state residence times.

More concretely, we show that total time positional (TTP) and time-abstract posi tional (TAP) schedulers allow to delay nondeterminism for all measures. As TTP schedulers are optimal for time-bounded reachability objectives, this shows that local uniformization preserves the probability of such objectives. Further, we prove that by delaying the scheduling decision, we can define delayed variants of these schedulers which in general induce strictly better bounds for the probabilities of quantitative properties.

However, we also identify two classes of schedulers, which do not allow for delaying nondeterminism. Our study results in a hierarchy of time-dependent schedulers which are characterized w.r.t. their ability to preserve probability measures under local uniformization. Moreover, we solve an open problem in [BHKH05] concerning TAP schedulers.

References

[BHKH05] Christel Baier, Holger Hermanns, Joost-Pieter Katoen, and Boudewijn R. Haverkort. Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. *Theoretical Computer Science*, 345(1):2–26, 2005.

[GM84] Donald Gross and Douglas R. Miller. The randomization technique as a modeling tool and solution procedure for transient Markov processes. *Operations Research*, 32(2):343–361, 1984.

[Gra91] W. K. Grassmann. Finding transient solutions in Markovian event systems through randomization. In W. J. Stewart, editor, *Numerical Solutions of Markov Chains*, pages 357–371. 1991.

[Jen53] A. Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 3:87–91, 1953.

2 Model checking priced timed automata and probabilistic timed automata

2.1 Discount optimal runs for priced timed automata

2.1.1 Participants

- Uli Fahrenberg, Kim G. Larsen, Claus Thrane, CISS, Aalborg University, Denmark

2.1.2 Contribution

Priced timed automata are timed automata whose locations and edges are decorated with prices: prices in locations are *prices per time unit spent in the location*, while prices on edges are paid each time the transition is fired. As opposed to hybrid systems, prices cannot serve in guards on transitions, which makes model checking and optimization decidable in priced timed automata. In [Lar09] we highlight recent (un)decidability results related to priced timed automata as well as point to a number of open problems.

One class of results is concerned with the derivation of infinite schedules for timed automata that are in some sense optimal. In previous work [BBL08] we have shown computability of synthesis of infinite runs being optimal in terms of minimal (or maximal) mean-payoff, *i.e.*, limit of the fraction between accumulation of cost and accummulated time. This is done by a reduction of the problem to the determination of optimal mean-cycles in finite graphs with weighted edges. This reduction is obtained by introducing the so-called corner-point abstraction, a powerful abstraction technique of which we showed that it preserves optimal schedules.

A serious drawback of the corner-point abstraction is the immediate explosion in the number of (symbolic) states to be considered—being exponential in both the number of clocks as well as in the size of the constants appearing in guards. As yet no efficient (in practice) zone-based algorithm for computing optimal infinite schedules in the above sense has been designed. Instead an alternative measure of optimality of infinite runs has recently been studied [FL09a] allowing for a fixed-point characterization and hence having the potiential of an efficient implementation. The optimality measure is based on a new discounting semantics for priced timed automata. Discounting provides a way to model optimal-cost problems for infinite traces and has applications in optimal scheduling and other areas. In the discounting semantics, prices decrease exponentially, so that the contribution of a certain part of the behaviour to the overall cost depends on how far into the future this part takes place. We show that the problem of finding an infinite path with minimal dsicounted price is computable, by a reduction to a similar problem on finite weighted graphs.

In [FL09b] we show that when postulating a certain natural additivity property for the discounted weights of runs, there is essentially only one possible way to introduce a discounting semantics. Our proof relies on the fact that a certain functional equation essentially only has one solution, for which we provide an elementary proof.

References

[BBL08] Patricia Bouyer, Ed Brinksma, and Kim G. Larsen. Optimal infinite scheduling for multi-priced timed automata. *Formal Methods in System Design*, 32(1):3–23, 2008.

[FL09a] Uli Fahrenberg and Kim G. Larsen. Discount-optimal infinite runs in priced timed automata. *Electr. Notes Theor. Comput. Sci.*, 239:179–191, 2009.

[FL09b] Uli Fahrenberg and Kim G. Larsen. Discounting in time. *Electr. Notes Theor. Comput. Sci.*, 253(3):25–31, 2009.

[Lar09] Kim G. Larsen. Priced timed automata: Theory and tools. In Ravi Kannan and K Narayan Kumar, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2009)*, volume 4 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 417–425, Dagstuhl, Germany, 2009. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

2.2 Priced timed automata with energy constraints

2.2.1 Participants

- Patricia Bouyer, Nicolas Markey, CNRS/LSV, ENS Cachan, France;
- Uli Fahrenberg, Kim G. Larsen, CISS, Aalborg University, Denmark.

2.2.2 Contribution

The Hydac case study led us to consider an alternative semantics for priced timed automata: in this case study, a pump must be turned on and off regularly in order to maintain the level of oil in a tank between two bounds. This can be modelled by priced timed automata under *energy constraints*: the accumulated price since the beginning of the run must be kept between a lower and an upper bound. The name *energy constraint* comes from the analogy with batteries, where energy can be consummed and regain, with the aim to never run out of energy.

We have defined the general problem in [BFL$^+$08], where we solved the untimed case (with only prices on transitions) and some of the 1-clock problems. In particular, we proved that the game version of the untimed lower-bound problem ("is there a strategy to maintain the accumulated price above a lower bound?") is equivalent to the classical *mean-payoff games*, which has given rise to improved algorithms for this fundamental theoretical problem [DGR09].

Recently, we investigated the case of 1-clock PTAs with lower bound, and proposed an exponential-time algorithm to solve the problem of optimizing the final credit under a lower-bound constraint. We also extended our setting to *exponential prices*: instead of evolving linearly with time, prices now satisfy first-order differential equations in locations, of the form $\frac{dp}{dt} = k \cdot p$, where k is the "rate" of the location and p is the value of the current accumulated price. Hence the accumulate dprice follows an exponential evolution: $p(t) = p(t_0) \cdot \exp(k \cdot (t - t_0))$. For

this setting also (under the restriction that prices on edges must be nonpositive), we designed an exponential-time algorithm for optimizing the final accumulated credit [BFLM09].

References

[BFL+08] Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, Nicolas Markey, and Jiří Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *Lecture Notes in Computer Science*, pages 33–47, Saint-Malo, France, September 2008. Springer.

[BFLM09] Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, and Nicolas Markey. Timed automata with observers under energy constraints. submitted, October 2009.

[DGR09] Laurent Doyen, Rafaella Gentilini, and Jean-Franois Raskin. Faster pseudo-polynomial algorithms for mean-payoff games. Submitted, June 2009.

2.3 Probabilistic timed automata

2.3.1 Participants

- Alexandre David, Arild Haugstad, Kim G. Larsen, CISS, Aalborg University, Denmark.

2.3.2 Contribution

We have developed new algorithms and data-structures for our family of tools, namely our main UPPAAL tool, UPPAAL TIGA, and UPPAAL PRO. The improvements range from added functionalities in the graphical interface to new capabilities in our model-checking engine.

Regarding the model checking of probabilistic timed automata, we have adapted the algorithms of "Minimal State Graph Generation", Bouajjani *et al.*, 1992 and "Minimization of Timed Transition Systems", Alur *et al.*, 1992 to probabilistic UPPAAL models. The novelties implemented in our tool are: 1) the computation is done on-the-fly with a partitioning algorithm and 2) we support all the modeling power offered by the rich language of the tool. The algorithm computes under- and over-approximations and can stop whenever the refinement has enough precision (provided by the user).

The tool is currently under development and a version is available on demand. We are working on how to better use the computed bounds to improve performance.

3 Model checking probabilistic priced timed automata

3.1 Participants

- Jasper Berendsen, David N. Jansen, Frits W. Vaandrager, ESI, University of Nijmegen, the Netherlands;
- Joost-Pieter Katoen, RTWH Aachen, Germany.

3.2 Contribution

In [BJK06], we introduce and study an extension of priced timed automata with probabilities: PPTAs equip timed automata with prices and probabilities on discrete transitions. Price *rates* indicate the increase of cost per time unit, whereas prices on discrete transitions indicate instantaneous costs. PPTAs are the orthogonal extension of both probabilistic timed automata (PTAs) [KNSS02] and priced timed automata [BFH$^+$01, ATP01], as PTAs extend timed automata with probabilities on discrete transitions and priced timed automata extend timed automata with prices. We proposed an algorithm which determines whether the probability to reach a (set of) goal location(s) within a given price bound (and time bound) can exceed a threshold $p \in [0, 1]$. We prove that the algorithm is partially correct and show an example for which termination cannot be guaranteed.

We then developped the tool FORTUNA, the first model checking tool that is able to deal with the combination of probabilities, costs and timing. FORTUNA is able to compute the fundamental problem of cost-bounded maximal probabilistic reachability (CBMR) for PPTA. CBMR determines the maximal probability by which a state can be reached under a certain cost-bound (and time bound).

As PTAs are PPTAs with trivial cost parameters, we were able to compare the performance of FORTUNA with existing approaches for PTAs that compute maximal probabilistic reachability. The comparison is made on a number of existing PTA case studies to the best approaches available: the game-based verification of [KNP09], the backwards reachability approach of [KNSW07], and the mcpta tool [HH09]. Surprisingly, although FORTUNA is more general, it shows better performance, sometimes by several orders of a magnitude.

In addition, FORTUNA is applied to an enhanced model of the IEEE 802.3 CSMA/CD protocol. The CSMA/CD is a protocol to avoid data collision on a single channel. An existing PTA model models the real-time and probabilistic aspects of the protocol. The extension is to model also its energy usage, and to see whether a message can reliably be sent within energy constraints.

FORTUNA employs an optimization of the algorithm described in [BJK06]. The algorithm of [BJK06] performs symbolic backwards exploration, in the spirit of the backwards reachability approach of [KNSW07]. Like that work, FORTUNA only adds intersections of symbolic states to the state space, thereby minimizing the number of stored states. To compute the probability, the explored symbolic state graph is transformed into a Markov decision process that is analysed with existing techniques. For PPTA FORTUNA may not terminate, since the problem is shown to

be undecidable in general [BCJ09]. But, for increasing exploration depth, the produced sequence of probabilities is non-decreasing and converges to the maximum probability [BJK06].

A number of optimizations is applied that increase performance drastically. They are proven to generate abstractions that preserve probabilistic reachability. The proofs are done in a rigorous way by the use of (probabilistic) simulation relations. The last optimization employs Hasse diagram data structures to speed up comparisons between symbolic states.

The main two other approaches for maximal probabilistic reachability on PTA use quite different techniques. Game-based verification [KNP09] uses an abstraction-refinement scheme to iteratively generate tighter lower and upper bounds on the probability that terminates in a finite number of steps. mcpta can be seen as an optimization of the digital clocks approach of [KNPS06], where clock values are discretized.

The FORTUNA tool, source code, and models discussed in this paper are freely available through the URL `http://www.cs.ru.nl/J.Berendsen/fortuna/`.

References

[ATP01] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. In Maria Domenica Di Benedetto and Alberto L. Sangiovanni-Vincentelli, editors, *HSCC*, volume 2034 of *Lecture Notes in Computer Science*, pages 49–62. Springer, 2001.

[BCJ09] Jasper Berendsen, Taolue Chen, and David N. Jansen. Undecidability of cost-bounded reachability in priced probabilistic timed automata. In Jianer Chen and S. Barry Cooper, editors, *TAMC*, volume 5532 of *Lecture Notes in Computer Science*, pages 128–137. Springer, 2009.

[BFH+01] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim Guldstrand Larsen, Paul Pettersson, Judi Romijn, and Frits W. Vaandrager. Minimum-cost reachability for priced timed automata. In Maria Domenica Di Benedetto and Alberto L. Sangiovanni-Vincentelli, editors, *HSCC*, volume 2034 of *Lecture Notes in Computer Science*, pages 147–161. Springer, 2001.

[BJK06] Jasper Berendsen, David N. Jansen, and Joost-Pieter Katoen. Probably on time and within budget: On reachability in priced probabilistic timed automata. In *QEST*, pages 311–322. IEEE Computer Society, 2006.

[HH09] Arnd Hartmanns and Holger Hermanns. A modest approach to checking probabilistic timed automata. In *QEST*, pages 187–196. IEEE Computer Society, 2009.

[KNP09] Marta Z. Kwiatkowska, Gethin Norman, and David Parker. Stochastic games for verification of probabilistic timed automata. In Joël Ouaknine and Frits W. Vaandrager, editors, *FORMATS*, volume 5813 of *Lecture Notes in Computer Science*, pages 212–227. Springer, 2009.

[KNPS06] Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Jeremy Sproston. Performance analysis of probabilistic timed automata using digital clocks. *Formal Methods in System Design*, 29(1):33–78, 2006.

[KNSS02] Marta Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282(1):101–150, 2002.

[KNSW07] Marta Z. Kwiatkowska, Gethin Norman, Jeremy Sproston, and Fuzhi Wang. Symbolic model checking for probabilistic timed automata. *Information and Computation*, 205(7):1027–1077, 2007.