

Parameter synthesis in Markov chains: algorithms, complexity and applications [★]

Christel Baier

Faculty of Computer Science
Technische Universität Dresden
Dresden, Germany

Christel.Baier@tu-dresden.de

Results of the quantitative analysis of systems modeled as Markov chains crucially depend on the concrete transition probabilities. Even small perturbations of the probability values can affect the analysis results. This can be problematic in cases where only estimates of the transition probabilities are available. This, for instance, applies to cases where the values of transition probabilities are derived using statistical or learning methods.

This motivated the introduction of Markov chains where intervals are attached to the transitions rather than concrete transition probabilities [24, 30, 9, 14, 7, 5]. Two semantics have been introduced for interval-valued Markov chains. The *uncertain semantics* treats interval-valued Markov chains as families of Markov chains with the same state space and transition probabilities in the intervals. The *nondeterministic semantics* considers interval-valued Markov chains as MDPs where the concrete probability values are chosen nondeterministically.

Parametric Markov chains can be seen as a generalization of interval-valued models with the uncertain semantics. Instead of intervals of potential probability values, they use polynomial functions over a fixed set of parameters to specify the transition probabilities [12, 27, 19]. Such parametric models can be seen to define a family of concrete probabilistic models that arise by plugging in concrete values for the parameters. Formally, a parametric Markov chain with parameters x_1, \dots, x_k is a tuple $\mathcal{M}[x_1, \dots, x_k] = (S, E, \Upsilon)$ where (S, E) is a finite directed graph, i.e., S is a finite state space and $E \subseteq S \times S$ specifies the transitions, $\Upsilon: E \rightarrow \mathbb{Q}[x_1, \dots, x_k]$ a function that assigns to each transition $s \rightarrow s'$ a polynomial over x_1, \dots, x_k with rational coefficients, and Γ a finite set of polynomial constraints for x_1, \dots, x_k . More precisely, Γ is a finite set of constraints of the form $f \bowtie c$ where $f \in \mathbb{Q}[x_1, \dots, x_k]$, \bowtie a comparison operator in $\{=, \neq, <, \leq, >, \geq\}$ and $c \in \mathbb{Q}$. A parameter valuation $\xi: \{x_1, \dots, x_k\} \rightarrow \mathbb{R}$ assigning real numbers to the parameters is said to be *admissible* if (1) $\sum_{s' \in E(s)} \Upsilon(s, s')(\xi_1, \dots, \xi_k) = 1$ for each non-trap state $s \in S$ where $E(s) = \{s' \in S : (s, s') \in E\}$ and (2) all constraints in Γ are satisfied, i.e., $f(\xi_1, \dots, \xi_k) \bowtie c$ for all $f \bowtie c$ in Γ . The (concrete) Markov chain given by an admissible parameter valuation ξ is $\mathcal{M}_\xi = (S, P_\xi)$ where $P_\xi(s, s') = \Upsilon(s, s')(\xi_1, \dots, \xi_k)$ if $(s, s') \in E$ and $P_\xi(s, s') = 0$ if $(s, s') \notin E$. Then, the semantics of $\mathcal{M}[x_1, \dots, x_k]$ is the family of (concrete) Markov chains induced by the admissible parameter valuations.

[★] The author is supported by the DFG through the collaborative research centres CRC TRR 248 (Foundations of Perspicuous Software Systems) and CRC 912 (HAEC) and the Excellence Strategy by the German Federal and State Governments (cluster of excellence CeTI) and the projects DFG BA-1679/11-1 and BA-1679/12-1.

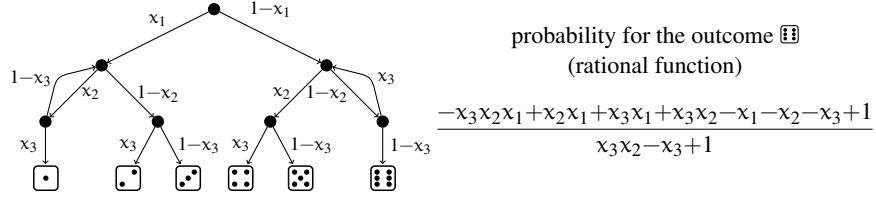


Fig. 1. Parametric Markov chain for the simulation of a dice by flipping biased coins

On the left of Figure 1 a parametric Markov chain is depicted, describing a parametric variant of Knuth and Yao’s protocol for simulating a six-sided dice by coin-flipping.

Probabilities for reachability conditions (and even for all ω -regular properties) can be represented as the solution of a linear equation system $A \cdot p = b$ where A is a matrix with coefficients represented as polynomials derived from Υ , p is the probability vector, and b a column vector. Thus, we can view A as a matrix over the field $\mathbb{Q}(x_1, \dots, x_k)$ of rational functions with parameters x_1, \dots, x_k and then apply Gaussian elimination (or the related state-elimination approach known for generating regular expressions for finite automata) to compute the solution vector p . This approach has been first proposed in [12, 27] and later refined using computer-algebra tools to simplify the rational functions obtained in intermediate steps [19] or using decompositions into strongly connected components and factorization techniques for polynomials [23]. These techniques have been implemented in the tools PARAM [18], its reimplement in PRISM [26], and in STORM [13]. More recently, it has been observed that the expensive computations of greatest common divisors of polynomials can be avoided using one-step fraction-free Gaussian elimination [22]. The latter yields that the rational functions for reachability probabilities in parametric Markov chains are computable in time $\mathcal{O}(\text{poly}(n, d)^k)$ where n is the number of states, d the maximal degree of the polynomials in Υ and k the number of parameters. In particular, in the univariate case ($k=1$) or for any fixed number k of parameters, the rational functions for reachability probabilities are computable in polynomial time. Even stronger, the problem to decide the existence of a parameter valuation satisfying a Boolean combination of threshold conditions for reachability probabilities is in P. The multivariate case is, however, challenging and threshold problems, e.g., “does there exists a parameter valuation such that the probability for reaching a given set of states exceeds a given threshold?” lie between NP and PSPACE [30, 9, 22].

Analogous results hold for expected accumulated weights or the expected mean payoff. If, however, the transition probabilities are non-parametric and the weights attached to states are parametric then parametric expressions for the expected accumulated weight or the expected mean payoff are computable in polynomial time [22].

The applications of parametric Markov chains and the algorithms to compute rational functions for reachability probabilities or expected values are manifold. These include the analysis of systems with unknown transition probabilities where the task is a family-based analysis to provide guarantees for all admissible parameter valuations,

but also the treatment of systems with tiny transition probabilities, which, e.g., applies often to models that capture the effect of very exceptional errors and where standard analysis techniques for the analysis fail due to numerical problems.

The parameter-synthesis problem goes one step further and asks to find parameter valuations where a set of probability or expectation constraints holds, possibly in combination with other side constraints. For a simple example, we regard the Markov chain shown in Figure 1. While for a fair coin, $x_1 = x_2 = x_3$, the result of the outcome “six” is $1/6$, one might ask how to manipulate the coins such that the probability for “six” is $1/2$. Assuming the same coin is used in all rounds (i.e., $x_1 = x_2 = x_3$), the task is to find a value p such that $(-p^3 + 3p^2 - 3p + 1)/(p^2 - p + 1) = 1/2$. With Newton’s method we obtain $p \approx 0.26102$. While this is a toy example, several important applications of the parameter-synthesis problem have been studied in the literature. Let us mention a few of them.

The use of parametric model checking for supporting decision making at run-time in adaptive software has been investigated by several authors, see, e.g., [8, 15, 16, 32]. The idea is to precompute the rational functions for a set of quantitative properties, which can be efficiently evaluated at run-time for guiding software adaptations when the environment changes. An iterative decision-making scheme for MDPs with interval estimates for the transition probabilities has been proposed in [31]. This approach serves to tackle the trade-off between accuracy, data usage and computational overhead. It successively computes strategies that are optimal for a cost-bounded reachability condition and checks their confidence optimality. If not, the iteration returns to data sampling.

Besides supporting decisions at run-time, the parametric approach can also be very useful to find (nearly) optimal system configurations at design-time. For example, [1] employs parametric model-checking techniques to find probability distributions for a probabilistic self-stabilizing protocol that achieves minimum average recovery time. [28] uses the parametric approach to determine (an approximation of) the optimal frequency of periodic reboots for an inter-process communication protocol of a space probe where optimality is understood with respect to the long-run availability. Another application of parametric Markov chains in the context of fault-tolerant systems is reported in [21] where the task is addressed to set the number of instructions per transaction and the maximal number of repeated executions of transactions before aborting that guarantee correct termination with probability at least 0.999 and minimize the overhead. In [4] the model-repair problem is considered where some transition probabilities of a Markov chain are declared to be controllable and the task is to modify the controllable probabilities such that a PCTL property holds for the modified Markov chain and the costs for deriving the new Markov chain are minimal. The latter is formalized using a nonlinear cost function, which leads to a nonlinear programming problem and is shown to be related to the optimal-controller synthesis problem for discrete linear dynamical systems. An alternative model-repair approach using parametric Markov chains and greedy methods has been proposed in [29].

Beyond parametric Markov chains, several authors have addressed the parametric setting to different and more expressive stochastic models. This includes methods to synthesize parametric rate values in continuous-time Markov chains that ensure the validity of bounded reachability properties [20]. Timeout-synthesis problems for

continuous-time stochastic models have been investigated, e.g., in [6, 2]. The parameter-synthesis problem for interval Markov chains with parametric lower and upper endpoints of intervals has been addressed in [3] where methods to find concrete values for the parameters such that the resulting interval-valued Markov chain maximizes the probability for a reachability condition have been proposed. Parametric approaches for MDPs have been studied in [17, 10, 11]. Another recent approach to deal with unknown transition probabilities uses online learning techniques, e.g., applied on MDP models in [25], where only the support of distributions is known in advance and where the task is to synthesize a strategy that almost surely satisfies a parity condition and is nearly optimal w.r.t. a mean-payoff objective.

References

1. Saba Aflaki, Matthias Volk, Borzoo Bonakdarpour, Joost-Pieter Katoen, and Arne Storjohann. Automated fine tuning of probabilistic self-stabilizing algorithms. In *36th IEEE Symposium on Reliable Distributed Systems (SRDS)*, pages 94–103. IEEE Computer Society, 2017.
2. Christel Baier, Clemens Dubslaff, Ľuboš Korenčiak, Antonín Kučera, and Vojtěch Řehák. Mean-payoff optimization in continuous-time Markov chains with parametric alarms. In *14th International Conference on Quantitative Evaluation of Systems (QEST)*, volume 10503 of *Lecture Notes in Computer Science*, pages 190–206. Springer, 2017.
3. Anicet Bart, Benoît Delahaye, Paulin Fournier, Didier Lime, Eric Monfroy, and Charlotte Truchet. Reachability in parametric interval Markov chains using constraints. *Theoretical Computer Science*, 747:48–74, 2018.
4. Ezio Bartocci, Radu Grosu, Panagiotis Katsaros, C. R. Ramakrishnan, and Scott A. Smolka. Model repair for probabilistic systems. In *17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 6605 of *Lecture Notes in Computer Science*, pages 326–340. Springer, 2011.
5. Michael Benedikt, Rastislav Lenhardt, and James Worrell. LTL model checking of interval Markov chains. In *19th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 7795 of *LNCS*, pages 32–46. Springer, 2013.
6. Tomáš Brázdil, Lubos Korenciak, Jan Krcál, Petr Novotný, and Vojtech Rehák. Optimizing performance of continuous-time stochastic systems using timeout synthesis. In *12th International Conference on Quantitative Evaluation of Systems (QEST)*, volume 9259 of *Lecture Notes in Computer Science*, pages 141–159. Springer, 2015.
7. Benoît Caillaud, Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wasowski. Constraint markov chains. *Theoretical Computer Science*, 412(34):4373–4404, 2011.
8. Radu Calinescu, Carlo Ghezzi, Marta Z. Kwiatkowska, and Raffaella Mirandola. Self-adaptive software needs quantitative verification at runtime. *Communications of the ACM*, 55(9):69–77, 2012.
9. Krishnendu Chatterjee, Koushik Sen, and Thomas A. Henzinger. Model-checking omega-regular properties of interval Markov chains. In *11th Int. Conf. on Foundations of Software Science and Computational Structures (FoSSaCS)*, volume 4962 of *LNCS*, pages 302–317. Springer, 2008.
10. Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, Ivan Papusha, Hasan A. Poonawala, and Ufuk Topcu. Sequential convex programming for the efficient verification of parametric MDPs. In *TACAS (2)*, volume 10206 of *LNCS*, pages 133–150, 2017.

11. Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, and Ufuk Topcu. Synthesis in pMDPs: A tale of 1001 parameters. In *16th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, volume 11138 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2018.
12. Conrado Daws. Symbolic and parametric model checking of discrete-time Markov chains. In *1st Int. Colloquium on Theoretical Aspects of Computing (ICTAC)*, volume 3407 of *LNCS*, pages 280–294. Springer, 2005.
13. Christian Dehnert, Sebastian Junges, Joost-Pieter Katoen, and Matthias Volk. A Storm is coming: A modern probabilistic model checker. In *29th Int. Conf. on Computer Aided Verification (CAV)*, volume 10427 of *LNCS*, pages 592–600. Springer, 2017.
14. Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wasowski. Decision problems for interval markov chains. In *5th International Conference on Language and Automata Theory and Applications (LATA)*, volume 6638 of *Lecture Notes in Computer Science*, pages 274–285. Springer, 2011.
15. Antonio Filieri, Giordano Tamburrelli, and Carlo Ghezzi. Supporting self-adaptation via quantitative verification and sensitivity analysis at run time. *IEEE Transactions on Software Engineering*, 42(1):75–99, 2016.
16. Simos Gerasimou. *Runtime quantitative verification of self-adaptive systems*. PhD thesis, University of York, UK, 2016.
17. Ernst Moritz Hahn, Tingting Han, and Lijun Zhang. Synthesis for PCTL in parametric markov decision processes. In *NASA Formal Methods - Third International Symposium (NFM)*, volume 6617 of *Lecture Notes in Computer Science*, pages 146–161. Springer, 2011.
18. Ernst Moritz Hahn, Holger Hermanns, Björn Wachter, and Lijun Zhang. PARAM: A model checker for parametric Markov models. In *22nd Int. Conference on Computer Aided Verification (CAV)*, volume 6174 of *LNCS*, pages 660–664. Springer, 2010.
19. Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. Probabilistic reachability for parametric Markov models. *Int. Journal on Software Tools for Technology Transfer*, 13(1):3–19, 2011.
20. Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. Approximate parameter synthesis for probabilistic time-bounded reachability. In *29th IEEE Real-Time Systems Symposium (RTSS)*, pages 173–182. IEEE Computer Society, 2008.
21. Linda Herrmann, Christel Baier, Christof Fetzer, Sascha Klüppelholz, and Markus Napierkowski. Formal parameter synthesis for energy-utility-optimal fault tolerance. In *15th European Workshop on Computer Performance Engineering (EPEW)*, volume 11178 of *Lecture Notes in Computer Science*, pages 78–93. Springer, 2018.
22. Lisa Hutschenreiter, Christel Baier, and Joachim Klein. Parametric Markov chains: PCTL complexity and fraction-free Gaussian elimination. In *8th Int. Symp. on Games, Automata, Logics and Formal Verification (GandALF)*, volume 256 of *EPTCS*, pages 16–30, 2017.
23. Nils Jansen, Florian Corzilius, Matthias Volk, Ralf Wimmer, Erika Ábrahám, Joost-Pieter Katoen, and Bernd Becker. Accelerating parametric probabilistic verification. In *11th Conference on Quantitative Evaluation of Systems (QEST)*, volume 8657 of *LNCS*, pages 404–420. Springer, 2014.
24. Bengt Jonsson and Kim Guldstrand Larsen. Specification and refinement of probabilistic processes. In *Sixth Annual Symposium on Logic in Computer Science (LICS)*, pages 266–277. IEEE Computer Society, 1991.
25. Jan Kretínský, Guillermo A. Pérez, and Jean-François Raskin. Learning-based mean-payoff optimization in an unknown MDP under omega-regular constraints. In *29th International Conference on Concurrency Theory (CONCUR)*, volume 118 of *LIPICs*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.

26. Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *23rd International Conference on Computer Aided Verification (CAV)*, volume 6806 of *Lecture Notes in Computer Science*, pages 585–591, 2011.
27. Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina. Parametric probabilistic transition systems for system design and analysis. *Formal Aspects of Computing*, 19(1):93–109, 2007.
28. Linda Leuschner, Martin Küttler, Tobias Stumpf, Christel Baier, Hermann Härtig, and Sascha Klüppelholz. Towards automated configuration of systems with non-functional constraints. In *16th Workshop on Hot Topics in Operating Systems (HotOS)*, pages 111–117. ACM, 2017.
29. Shashank Pathak, Erika Ábrahám, Nils Jansen, Armando Tacchella, and Joost-Pieter Katoen. A greedy approach for the efficient repair of stochastic models. In *7th International Symposium on NASA Formal Methods (NFM)*, volume 9058 of *Lecture Notes in Computer Science*, pages 295–309. Springer, 2015.
30. Koushik Sen, Mahesh Viswanathan, and Gul Agha. Model-checking Markov chains in the presence of uncertainties. In *12th Int. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 3920 of *LNCS*, pages 394–410. Springer, 2006.
31. Guoxin Su, Taolue Chen, Yuan Feng, David S. Rosenblum, and P. S. Thiagarajan. An iterative decision-making scheme for Markov decision processes and its application to self-adaptive systems. In *19th International Conference on Fundamental Approaches to Software Engineering (FASE)*, volume 9633 of *Lecture Notes in Computer Science*, pages 269–286. Springer, 2016.
32. Guoxin Su, David S. Rosenblum, and Giordano Tamburrelli. Reliability of run-time quality-of-service evaluation using parametric model checking. In *38th International Conference on Software Engineering (ICSE)*, pages 73–84. ACM, 2016.