Software Modeling and Verification Chair | RWTH AACHEN UNIVERSITY

— Bachelor's or Master's Thesis —

Theoretical Aspects of Parametric Markov Decision Processes

What is it all about?

Markov decision processes (MDPs) are *the* model to reason about sequential processes under (stochastic) uncertainty and non-determinism. Often, probability distributions in these models are difficult to assess precisely during design time of a system. There exist modeling formalisms that allow for interval-labeled transitions; however, this typically leads to too pessimistic analysis results: The actual probabilities on the transitions are considered to be non-deterministically and locally chosen. Parametric MDPs (pMDPs) [LMT07, HHZ11] overcome this limitation by adding dependencies (or couplings) between various transitions — they add global restrictions to the selection of the probability distributions. Intuitively, the probability of flipping heads with a biased coin can be arbitrary, but should be independent of the system's state, e.g. an agent's local memory.

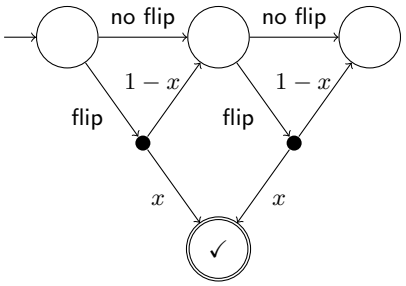

For example, suppose that an agent may flip the same biased coin at most two times in order to achieve the outcome "heads" at least once. The depicted pMDP models this situation, denoting the probability for heads with x. A prototypical question is then "Does there exist a probability x such that for one/all behaviors of the agent the probability to see heads once is at least/at most 0.5"?

What is to be done? The aim of this project is to investigate theoretical questions about pMDPs, in particular questions related to (computational) complexity. As a starting point, various problems left open in [WJPK19] may be addressed:

1. An *minimal optimal strategy set* (MOSS) contains exactly one optimal strategy (=resolution of the non-determinism) for each point of the parameter space. What is the size of a MOSS in general? What are the implications for computational complexity?
2. Once a strategy is fixed, the reachability probabilities can be described as *rational functions*. How can the class of rational functions that may occur in a pMDP be characterized?
3. What are the complexity bounds on approximative sampling based techniques that are being used in practice?
4. Investigate the relationship to other models such as stochastic games more in detail.

This list is of course non-exhaustive!

What you can expect:

- Get a chance to work on relevant open problems of theoretical nature.
- Work closely together with us – you can work in our student's room at the chair whenever you like!

What we expect:

- Solid background in theoretical computer science and maths – ideally you have already taken theoretical CS and/or maths electives.
- Passion and endurance for solving difficult (mathmatical) problems.

Contact

- Tobias Winkler, tobias.winkler@cs.rwth-aachen.de

References

[HHZ11] Ernst Moritz Hahn, Tingting Han, and Lijun Zhang, Synthesis for PCTL in Parametric Markov Decision Processes, NASA Formal Methods (Mihaela Bobaru, Klaus Havelund, Gerard J. Holzmann, and Rajeev Joshi, eds.), vol. 6617, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 146–161 (en).

[LMT07] Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina, Parametric probabilistic transition systems for system design and analysis, Formal Aspects of Computing **19** (2007), no. 1, 93–109 (en).

[WJPK19] Tobias Winkler, Sebastian Junges, Guillermo A. Pérez, and Joost-Pieter Katoen, On the complexity of reachability in parametric markov decision processes, 30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands, 2019, pp. 14:1–14:17.