Approximation methods for the soundness of control laws derived by machine learning

The design of controllers for large cyber physical systems (CPS, i.e. systems driven both by physical equations and digital controllers) is challenged today by machine learning approaches, and specifically reinforcement learning. The latter however still fail to provide guarantees on the behavior of the controllers it provides. The objective of this thesis is to explore a range of techniques that would make control design for CPS or any other large-scale complex system sound and scalable. The focus will be on quantitative methods, that provide performance guarantees, for example PAC bounds (probably approximately correct).

Several research directions are envisioned. The main one concerns model approximation. For a given dynamic system with discrete state, like a stochastic automaton, this may mean reducing the size of the state space while preserving as much as possible the distribution over generated runs, which requires computing or estimating distances between models. Starting from a CPS with continuous state variables, this means finding the best discretization with bounded state size. One may as well take as starting point a (possibly infinite) collection of representative runs of that system, or a black box trace generator, and be interested in learning a model from these traces (system identification) in order to capture the most characteristic features of their dynamics. For all these directions, one will be interested both in designing approximation algorithms, in characterizing their convergence properties, and in providing bounds for their accuracy.

A second research direction concerns approximation techniques in view of control design. There, the model (a Markov Decision Process for example) comes as the support to design an efficient control policy, toward some quantitative objective. Optimal control laws generally derive from iterative methods that do not scale up with model dimension, in particular if the latter come from discretization of continuous variables. The objective will be to explore various approximation techniques that would improve scalability, convergence speeds and provide both performance bounds and readability of the control laws. Model approximations are one possible way, but also controller regularization (for example through state aggregation), or approximations in the iterative procedure that yield optimal laws, or even control objective relaxations.

As a possible use-case for the above techniques, we aim at designing distributed controllers for large CPS, for example a fleet of trains on a subway line. The objective will be both to design multi-agent control strategies, to estimate their performances and to verify safety properties like maintaining minimal headways. Applications to other complex mechanical devices are also envisioned, like those of the OpenAi Gym.

This PhD will take place in the SUMO Team at INRIA Rennes (Brittany, France), under the joint supervision of Loïc Hélouët and Eric Fabre. Funding is secured for this PhD, as a 3 years contract. This research will be connected to the Maveriq ANR project.

Candidates for this PhD should have a strong interest in formal methods. Former experience in probabilistic or quantitative model checking, or in learning techniques will be appreciated.
For more information, please contact:
eric.fabre@inria.fr
loic.helouet@inria.fr

Applicants must send their resume (CV), master marks, and a list of references.

Bibliography :

- “A canonical form for weighted automata and applications to approximate minimization,” Balle, Panangaden, Precup, LICS 2015
- “Adaptive state space partitioning for reinforcement learning,” Lee, Lau, Engineering Applications of Artificial Intelligence, 2004
- “Reinforcement learning with soft state aggregation,” Singh, Jaakkola, Jordan, NeurIPS 1994
- “On Satisficing in Quantitative Games,” Bansal, Chatterjee, Vardy, TACAS 2021
- “On Time with Minimal Expected Cost !” David, Jensen, Larsen et al., ATVA 2014