Resilient Control and Safety for Multi-Agent Cyber-Physical Systems

Anna Lukina
Technische Universität Wien, Vienna, Austria
anna.lukina@tuwien.ac.at

Abstract

I develop novel intelligent approximation algorithms for solving modern problems of Cyber-Physical Systems (CPS), such as control and verification, by combining advanced statistical methods. it is important for the control algorithms underlying the class of multi-agent CPS to be resilient to various kinds of attacks. I designed a very general adaptive receding-horizon synthesis approach to planning and control that can be applied to controllable stochastic dynamical systems. Apart from being fast and efficient, it provides statistical guarantees of convergence. The optimization technique based on the best features of Model Predictive Control and Particle Swarm Optimization proves to be robust in finding a winning strategy in the stochastic non-cooperative games against a malicious attacker. The technique can further benefit probabilistic model checkers and real-world CPS.

1 Motivation

The nature of many Cyber-Physical Systems (CPS) is highly distributed, representing a multitude of agents computing their actions, collectively contributing to the emergent behavior. As an appealing example of such a distributed CPS, the drone swarms have been increasingly applied in battlefield surveillance and reconnaissance [Condliffe, 2017]. The emergent behavior they exhibit is that of flight formation.

The physical component of CPS introduces a wealth of uncertainties which raises the problems of verification and control. Statistical Model-Checking (SMC) has recently grown in popularity as a way to avoid state-space explosion. SMC, however, might fall a victim of rare events, e.g., diverging from the target, crashing. After estimating such events with high confidence, it is essential to develop a control policy guaranteeing avoidance of rare events on the way to the goal. Computing an appropriate action response to the environmental changes includes probabilistic estimation of the current state, as well as prediction of the optimal action by simulating the future.

The controller synthesis problem has been widely studied [Kwiatkowska, 2016]. Dynamic Programming (DP) is one of the most popular approaches [Bellman, 1957]. Given a predefined asymptotic error, it converges to the optimal value of the functional by improving it at each iteration. In contrast with DP, which might suffer from state-space explosion when considering all possible states of the system with respect to environmental uncertainties, approximate algorithms [Mannor *et al.*, 2003] take into account only the paths leading to the desired goal. One of the most efficient such techniques is Particle Swarm Optimization (PSO) [Kennedy and Eberhart, 1995]. PSO is a very powerful optimization tool but in the case of multi-agent CPS, it is too time-consuming for PSO to converge. In [Yang *et al.*, 2016] PSO is adapted for finding the next best step of Model-Predictive Control (MPC) [Garca *et al.*, 1989]. MPC determines the value of a control variable at the current time by looking several steps ahead and finding the best horizon-length control sequence that can bring the system from its current state to a new one with the improvement of the objective function. However, MPC provides no convergence guarantees for stochastic systems. Sequential Monte-Carlo methods proved to be efficient in tackling the question of control for linear stochastic systems.

As a result, in my thesis, I develop a novel intelligent approximation algorithm for solving modern problems of CPS, such as control and verification, by combining advanced statistical methods. It is important for the control algorithms underlying the class of multi-agent CPS to be resilient to various kinds of cyber-attacks.

2 Contributions

To tackle the problem of rare events, a novel framework of SMC as a feedback control for CPS was proposed in [Kalajdzic *et al.*, 2016]. In this work, I investigated a combination of two sequential Monte-Carlo methods, importance sampling (ISam) and importance splitting (ISpl), for steering a stochastic system towards a rare property and estimating its true probability. Using ISam to estimate the current state and ISpl to iteratively simulate the CPS for bounded time, I computed the probabilities of given rare properties for two running examples of stochastic models, which illustrated good performance of the algorithm.

Considering various forms of flight formations, V-formation is of a particular interest for attacking the control question, especially for long-distance travels where energy conservation is the key. The V-formation is characteristic for migratory birds, e.g., Canada geese, when a bird can take ad-

vantage of the upwash region of the bird flying in front of it and save energy this way. This type of formation also provides clear view and matching velocities for each bird, i.e., none of them is obstructed by the others or needs to accelerate to maintain the formation of the flock.

Working on the problem above, I was interested to see if flying in V-formation is a policy optimizing energy conservation. As a consequence, I developed ARES, an efficient approximation algorithm for synthesizing optimal plans that take a Markov decision process (MDP) from its initial state to a state whose cost is below a specified (convergence) threshold. My algorithm employs PSO, with adaptive values for both receding horizon (look-ahead) and the swarm size. Inspired by ISpl, the length of the horizon and the number of particles in the swarm are chosen such that at least one particle reaches the next level, that is, a state where the objective function (the cost) is decreased by a required delta compared to the previous level. The level relation on the states and the resulting plans implicitly define a Lyapunov function and an optimal policy, respectively. They could be further explicitly generated by applying ARES to all states of the MDP, up to some topological equivalence relation. To assess the effectiveness of ARES, I statistically evaluated its success rate in generating optimal plans. For flocks with 7 birds, ARES can generate a plan that leads to a V-formation in 95% of the 8,000 random initial configurations within 63 seconds, on average. To the best of our knowledge, this adaptive-sizing approach is the first to provide convergence guarantees in receding-horizon techniques [Lukina *et al.*, 2017].

Later the adaptive level-based approach used in ARES proved to be resilient to various types of attacks. ARES was customized into adaptive MPC (AMPC). The concept of V-formation game was introduced as a stochastic two-player game between an adaptive controller and an attacker. In this framework, I conducted experiments with various attack models: removing or displacing a bird, as well as an AMPC-enabled attacker. After reaching a V-formation, the goal of the games was to reach it again in bounded time under attack. Algorithm performance evaluation demonstrated that AMPC-controller almost always wins the games (submitted to ATVA 2017 [Smolka *et al.*, 2017]).

3 Current and Anticipated Progress

Inspired by the success of ARES, I further improved the execution time by parallelizing the algorithm and proceeded to adapt it to distributed multi-agent CPS.

Resilient Distributed Control of Autonomous Agents. One of my greatest aspirations is to build a distributed controller for autonomous systems that would be resilient to the physical and cyber attacks. Apart from the challenge of developing a robust consensus algorithm, this task also implies dealing with local policy computation under imperfect information. The authors of [Saulnier *et al.*, 2017] propose a resilient controller for a networked team of mobile robots. They, however, use a global observer for synchronization.

Model-checking for MDPs. After testing ARES on benchmark objective function with multiple local optima, I would like to compare my approach against popular model-checkers, such as PRISM, to improve their performance.

Deep Reinforcement Learning. Currently, when synthesizing an optimal plan the objective function and its optimal value are known in advance. I find deep reinforcement learning very promising approach to learning the cost functions for a given multi-agent CPS, such as a drone swarm, inside the optimization procedure of ARES. This could lead to an algorithmic solution applicable for real-world multi-agents.

References

[Bellman, 1957] Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[Condliffe, 2017] Jamie Condliffe. A 100-Drone Swarm, Dropped from Jets, Plans Its Own Moves. *MIT Technology Review*, January 2017.

[Garca *et al.*, 1989] Carlos E. Garca, David M. Prett, and Manfred Morari. Model Predictive Control: Theory and Practice – A Survey. *Automatica*, 25(3):335–348, 1989.

[Kalajdzic *et al.*, 2016] Kenan Kalajdzic, Cyrille Jégourel, Anna Lukina, Ezio Bartocci, Axel Legay, Scott A. Smolka, and Radu Grosu. Feedback Control for Statistical Model Checking of Cyber-Physical Systems. In *ISoLA 2016, October 10-14, 2016, Proceedings, Part I*, volume 9952 of *LNCS*, pages 46–61, 2016.

[Kennedy and Eberhart, 1995] James Kennedy and Russell Eberhart. Particle Swarm Optimization. In *Proceedings of 1995 IEEE International Conference on Neural Networks*, pages 1942–1948, 1995.

[Kwiatkowska, 2016] Marta Kwiatkowska. Advances and Challenges of Quantitative Verification and Synthesis for Cyber-Physical Systems. In *Cyber-Physical Systems Workshop (SOSCYPS), Science of Security for*, pages 1–5. IEEE, 2016.

[Lukina *et al.*, 2017] Anna Lukina, Lukas Esterle, Christian Hirsch, Ezio Bartocci, Junxing Yang, Ashish Tiwari, Scott A. Smolka, and Radu Grosu. ARES: adaptive receding-horizon synthesis of optimal plans. In *TACAS 2017, April 22-29, 2017, Proceedings, Part II*, volume 10206 of *LNCS*, pages 286–302, 2017.

[Mannor *et al.*, 2003] S. Mannor, R. Y. Rubinstein, and Y. Gat. The cross entropy method for fast policy search. In *ICML*, pages 512–519, 2003.

[Saulnier *et al.*, 2017] Kelsey Saulnier, David Saldana, Amanda Prorok, George J. Pappas, and Vijay Kumar. Resilient Flocking for Mobile Robot Teams. *IEEE Robotics and Automation Letters*, 2(2):1039–1046, 2017.

[Smolka *et al.*, 2017] Scott A Smolka, Ashish Tiwari, Lukas Esterle, Anna Lukina, Junxing Yang, and Radu Grosu. Attacking the V: On the Resiliency of Adaptive-Horizon MPC. *arXiv preprint arXiv:1702.00290*, 2017.

[Yang *et al.*, 2016] Junxing Yang, Radu Grosu, Scott A. Smolka, and Ashish Tiwari. Love Thy Neighbor: V-Formation as a Problem of Model Predictive Control. In *CONCUR 2016, August 23-26, 2016*, volume 59 of *LIPIcs*, pages 4:1–4:5. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.