Adaptive Neighborhood Resizing for Stochastic Reachability in Multi-Agent Systems

Anna Lukina[1], Ashish Tiwari[3], Scott. A. Smolka[2], and Radu Grosu[1,2]

[1] Cyber-Physical Systems Group, Technische Universität Wien, Vienna, Austria
[2] Department of Computer Science, Stony Brook University, New York, USA
[3] Microsoft, USA
anna.lukina@tuwien.ac.at

V-formation in a flock of birds is a quintessential example of emergent behavior in a stochastic multi-agent system. It brings numerous benefits to the flock. It is primarily known for being energy-efficient due to the upwash benefit a bird in the flock enjoys from its frontal neighbor. It is therefore not surprising that interest in V-formation is on the rise in the aircraft industry [1]. Recent work on V-formation has shown that the problem can be viewed as one of optimal control, model-predictive control (MPC) in particular. In [3], we introduced adaptive-horizon MPC (`AMPC`), a highly effective control algorithm for multi-agent cyber-physical systems (CPS) modeled as a Markov decision process (MDP). Traditional MPC uses a fixed prediction horizon, i.e. number of steps to compute ahead, to determine the optimal, cost-minimizing control action. The downside of the fixed look-ahead is that the algorithm may get stuck in a local minimum. For a controllable MDP, `AMPC` chooses its prediction horizon dynamically, extending it out into the future until the cost function (shown in blue in Fig. 1) decreases sufficiently [2]. This implicitly endows `AMPC` with a Lyapunov function (shown in red in Fig. 1), providing statistical guarantees of convergence to a goal state such as V-formation, even in the presence of adversarial agents. It should be noted that `AMPC` works in a centralized manner.

This paper introduces `DAMPC`, a distributed version of `AMPC` that extends it along several dimensions. First, at every time step, `DAMPC` runs a *distributed consensus algorithm* to determine the optimal action (acceleration) for every agent in the flock. In particular, each agent i starts by computing the optimal actions for its local subflock. The subflocks then communicate in a sequence of consensus rounds to determine the optimal actions for the entire flock. Secondly, `DAMPC` features *adaptive neighborhood resizing* (black line in Fig. 1) in an effort to further improve the algorithm's efficiency. In a similar way as for the prediction horizon in `AMPC`, neighborhood resizing utilizes the implicit Lyapunov function to guarantee eventual convergence to a minimum neighborhood size. `DAMPC` thus treats the neighborhood size as another controllable variable that can be dynamically adjusted for efficiency purposes. This leads to reduced communication and computation compared to the centralized solution, without sacrificing statistical guarantees of convergence. The proof of statistical global convergence is intricate. For example, consider the scenario shown in Fig. 1. `DAMPC` is decreasing the neighborhood size k for all agents, as the system-wide cost function J follows a decreasing trajectory. Suddenly and without warning, the flock begins

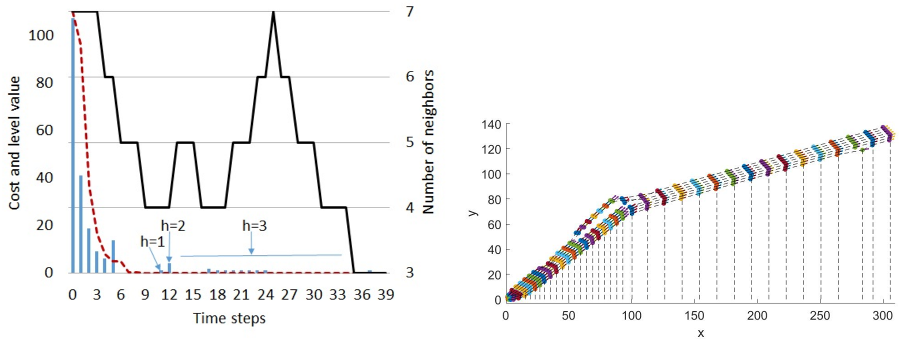

Fig. 1. Left: Blue bars are the values of the cost function in every time step. Red dashed line is the cost-based Lyapunov function used for horizon and neighborhood adaptation. Black solid line is neighborhood resizing for the next step given the current cost. Right: Step-by-step evolution of the flock of seven birds bringing two separate formations together. Each color-slice is a configuration of the birds at a particular time step.

to split into two, undoubtedly owing to an unsuitably low value of k, leading to an abrupt upward turn in J. `DAMPC` reacts accordingly and promptly, increasing its prediction horizon first and then k, until system stability is restored. The ability for `DAMPC` to do this is guaranteed, for in the worst case k will be increased to B, the total number of birds in the flock. It can then again attempt to monotonically decrease k, but this time starting from a lower value of J, until V-formation is reached.

Apart from the novel adaptive-horizon adaptive-neighborhood distributed algorithm to synthesize a controller, and its verification using statistical model checking, we believe the work here is significant in a deeper way. The problem of synthesizing a sequence of control actions to drive a system to a desired state can be also viewed as a falsification problem, where one tries to find values for (adversarial) inputs that steer the system to a bad state. These problems can be cast as constraint satisfaction problems, or as optimization problems. As in case of V-formation, one has to deal with non-convexity, and popular techniques, such as convex optimization, will not work. Our approach can be seen as a tool for solving such highly nonlinear optimization problems that encode systems with notions of time steps and spatially distributed agents. Our work demonstrates that a solution can be found efficiently by adaptively varying the time horizon and the spatial neighborhood. By allowing adaptation to consider longer time horizons, and larger neighborhoods (possibly the entire flock), one can provide convergence guarantees that would be otherwise impossible (say, in a fixed-horizon MPC).

References

1. Johnsson, J.: Boeing copies flying geese to save fuel (2017), https://www.bloomberg.com/news/articles/2017-08-08/boeing-nasa-look-to-flying-geese-in-chase-for-jet-fuel-savings

2. Lukina, A., Esterle, L., Hirsch, C., Bartocci, E., Yang, J., Tiwari, A., Smolka, S.A., Grosu, R.: ARES: adaptive receding-horizon synthesis of optimal plans. In: Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017. LNCS, vol. 10206, pp. 286–302 (2017)
3. Tiwari, A., Smolka, S.A., Esterle, L., Lukina, A., Yang, J., Grosu, R.: Attacking the V: on the resiliency of adaptive-horizon MPC. In: Automated Technology for Verification and Analysis - 15th International Symposium, ATVA 2017. LNCS, vol. 10482, pp. 446–462. Springer (2017)