Distributed Adaptive-Neighborhood Control for Stochastic Reachability in Multi-Agent Systems

Anna Lukina
Technische Universität Wien, Austria

Scott A. Smolka
Stony Brook University, USA

Ashish Tiwari
SRI International, USA

Radu Grosu
Technische Universität Wien, Austria

ABSTRACT

We present DAMPC, a distributed, adaptive-horizon and adaptive-neighborhood algorithm for solving the stochastic reachability problem in multi-agent systems, in particular, flocking modeled as a Markov decision process. At each time step, every agent first calls a centralized, adaptive-horizon model-predictive control (AMPC) algorithm to obtain an optimal solution for its local neighborhood. Second, the agents derive the flock-wide optimal solution through a sequence of consensus rounds. Third, the neighborhood is adaptively resized using a flock-wide cost-based Lyapunov function. This way DAMPC improves efficiency without compromising convergence. We evaluate DAMPC's performance using statistical model checking. Our results demonstrate that, compared to AMPC, DAMPC achieves considerable speed-up (two-fold in some cases) with only a slightly lower rate of convergence. The smaller average neighborhood size and lookahead horizon demonstrate the benefits of the DAMPC approach for stochastic reachability problems involving any controllable multi-agent system that possesses a cost function.

KEYWORDS

adaptive, predictive control, multi-agent, stochastic reachability

ACM Reference Format:

Anna Lukina, Ashish Tiwari, Scott A. Smolka, and Radu Grosu. 2019. Distributed Adaptive-Neighborhood Control, for Stochastic Reachability in Multi-Agent Systems. In *The 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19)*, April 8–12, 2019, Limassol, Cyprus. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3297280.3297370>

1 INTRODUCTION

V-formation in a flock of birds is a quintessential example of emergent behavior in a stochastic multi-agent system. V-formation 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. In addition, it offers each bird a clear frontal view, unobstructed by any flockmate. It is therefore not surprising that interest in V-formation is on the rise in the aircraft industry [1].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '19, April 8–12, 2019, Limassol, Cyprus

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5933-7/19/04...\$15.00

<https://doi.org/10.1145/3297280.3297370>

Figure 1: 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.

Recent work on V-formation has shown that the problem can be viewed as one of optimal control, model-predictive control (MPC) [3], in particular. In [16], authors 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. 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, with global knowledge of the state of the flock.

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 sub-flock. 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 such as those offered by its centralized counterpart AMPC.

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 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 attempted to monotonically decrease k , but this time starting from a lower value of J , until V-formation is reached. A smoother convergence scenario is shown in Fig. 3. In this case, the efficiency gains of adaptive neighborhood resizing are more evident, as the cost function J follows an almost purely monotonically decreasing trajectory. A formal proof of global convergence of DAMPC with high probability is given as one of paper’s main results.

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. A main benefit of the adaptive scheme, apart from efficiency, is that it gives a path towards completeness. 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).

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 describes the flocking model and the cost function. Section 4 defines the stochastic reachability problem. Section 5 introduces DAMPC, our main contribution. Section 6 provides proofs of our theoretical results. Section 7 presents our statistical evaluation of the algorithm. Section 8 offers concluding remarks and indicates directions for future research.

2 RELATED WORK

In [12], the ARES algorithm generates plans incrementally, segment by segment, using adaptive horizons (AH) to find the next best step towards the global optimum. In particular, ARES calls a collection of particle swarm optimizers (PSO), each with its own AH. PSOs with the best results are cloned, while the others are restarted from the current level on the way to the goal (similar to importance splitting [10]). When the V-formation is achieved, the complete plan is composed of the best segments. The AHs are chosen such that the best PSOs can succeed to decrease the objective cost by at least a pre-defined value, implicitly defining a Lyapunov function that guarantees global convergence.

In [19], the problem of taking an arbitrary initial configuration of n agents to a final configuration, where every pair of stationary “neighbors” is a fixed distance d apart, is considered. The authors present centralized and distributed algorithms for this problem, both of which use MPC to determine the next action. The problem in [19] is related to our work. However, we consider nonconvex and nonlinear cost functions, which require overcoming local minima to ensure convergence. In contrast, [19] deals with convex functions, which do not suffer from problems introduced by the presence of multiple local minima. Furthermore, in the distributed control procedure of [19], each agent publishes the control value it locally computed, which is then used by other agents to calculate their own. A quadratic number of such steps is performed before each agent fixes its control input for the next time step. In our work, we limit this number to linear.

Other related work, including [5, 7, 18], focuses on distributed controllers for flight formation that operate in an environment where the multi-agent system is already in the desired formation and the distributed controller’s objective is to maintain formation in the presence of disturbances. A distinguishing feature of these approaches is the particular formation they are seeking to maintain, including a half-vee [7], a ring and a torus [5], and a leader-follower formation [18]. These works are specialized for capturing the dynamics of moving-wing aircraft. In contrast, we use DAMPC with dynamic neighborhood resizing to bring a flock from a random initial configuration to a stable V-formation.

It is worth noting that, although DAMPC uses global consensus, our main focus is on adaptive neighborhood resizing and global convergence, and not on fault tolerance [6, 15]. In [14], the authors consider such MPC-inspired approaches to system self-adaptation as PLA [13]. The common ground between these approaches and DAMPC is that the future behavior of the system is predicted based on a model, in the case of PLA, a Markov decision process, and a sequence of actions is computed from the current state for the length of prediction horizon. PLA resembles our approach also in the way it synthesizes action plans at each discrete time step while only applying the first of them to the plant. Unlike PLA, however, our approach adapts the prediction horizon and neighborhood size based on the value of the cost function. Moreover, DAMPC is a distributed procedure guaranteed to converge to the goal state.

There has been a body of work done on decentralized multi-robot planning and coordination using partially observable Markov decision processes (POMDPs). In [4], the authors tackle a problem of planning under uncertainty. Their approach achieves better

scalability than previous POMDP-based algorithms. The authors, however, admit that they do so on the expense of optimality of their solutions. In contrast, our algorithm is focused on providing statistical guarantees of reaching the optimal desired state without any role assignment as in [4]. Another related work making use of multi-agent POMDPs ([2]) presents an efficient application of factored statistics and factored trees for improving scalability of on-line planning. The authors achieve significant improvement for the problems with 10 agents. Their approach is built upon coordination graphs, which we would like to consider in our future research. For this paper, we aim to avoid imposing any topological structure on the flock to allow for emergent behavior.

3 BACKGROUND ON V-FORMATION

Dynamical model. In the flocking model used, the state of each bird is given by four variables: a 2-dimensional vector $\mathbf{x}$ denoting the position of the bird in 2D continuous space, and a 2-dimensional vector $\mathbf{v}$ denoting the velocity of the bird. We use $\mathbf{s} = \{\mathbf{x}_i, \mathbf{v}_i\}_{i=1}^B$ to denote a state of a flock with B birds. The control actions of each bird are 2-dimensional accelerations $\mathbf{a}$.

Let $\mathbf{x}_i(t)$, $\mathbf{v}_i(t)$, and $\mathbf{a}_i(t)$ denote the position, velocity, and acceleration, of i -th bird at time t , $i \in \{1, \dots, B\}$, respectively. Given an initial configuration $\mathbf{x}_i(0) = \mathbf{x}_i^0$, $\mathbf{v}_i(0) = \mathbf{v}_i^0$ inside a bounding box of a given size, the discrete-time behavior of each bird is in Eq. 1:

$$\begin{aligned}\mathbf{v}_i(t+1) &= \mathbf{v}_i(t) + \mathbf{a}_i(t), \\ \mathbf{x}_i(t+1) &= \mathbf{x}_i(t) + \mathbf{v}_i(t).\end{aligned}\quad (1)$$

In the simulations described in Section 7, the time interval between two successive updates of the positions and velocities of all birds in the flock equals one. The initial state is generated uniformly at random inside a bounding box. The accelerations are the output of the particle swarm optimization (PSO) algorithm [11], which samples particles uniformly at random subject to the following constraints on the maximum velocities and accelerations: $\forall i \in \{1, \dots, B\} \|\mathbf{v}_i(t)\| \leq \mathbf{v}_{max}$, $\|\mathbf{a}_i(t)\| \leq \rho \|\mathbf{v}_i(t)\|$, where $\mathbf{v}_{max}$ is a constant and $\rho \in (0, 1)$.

We chose this model for its simplicity, as the cost function described in Section 3, is nonlinear, nonconvex, nondifferentiable, and therefore sufficient enough to make reachability analysis extremely challenging. We propose an approximate algorithm to tackle reachability. In principle, more sophisticated models of flocking dynamics can be considered, but we leave those for future work, and focus on the simplest one.

The problem of bringing a flock from an arbitrary configuration to a V-formation can be posed as a reachability question, where the goal is the set of states representing a V-formation. A key assumption is that the reachability goal can be specified as $J(\mathbf{s}) \leq \varphi$, where J is a cost function that assigns a nonnegative real value to each state $\mathbf{s}$, and φ is a small positive constant.

Cost Function. To define the cost function, we take the metrics determining the cost of a state as described in [17]. *Clear View:* $CV(\mathbf{s})$ is defined by accumulating the percentage of a cone with angle θ , blocked by other birds. The minimum value is $CV^* = 0$ and attained in a perfect V-formation where all birds have an unobstructed view. *Velocity Matching:* $VM(\mathbf{s})$ for flock state $\mathbf{s}$ is defined as the difference between the velocity of a given bird and all other

birds, summed up over all birds in the flock. The minimum value is $VM^* = 0$ and attained in a perfect V-formation where all birds have the same velocity. *Upwash Benefit:* the trailing upwash is generated near the wingtips of a bird, while downwash is in the center of a bird. An upwash measure is defined on the 2D space using a Gaussian-like model that peaks at the appropriate upwash and downwash regions. $UB(\mathbf{s})$ for flock state $\mathbf{s}$ is the sum of the upwash benefit each bird can obtain UB_i for $1 \leq i \leq B$. The upwash benefit $UB(\mathbf{s})$ in V-formation is $UB^* = 1$, as all birds, except for the leader, have minimum upwash-benefit metric ($UB_i = 0$), while the leader has an upwash-benefit metric of 1 ($UB_i = 1$).

Given the above metrics, the overall objective function J is defined as a sum-of-squares of VM , CV , and UB , as follows, where state $\mathbf{s}^*$ is considered to be a V-formation if $J(\mathbf{s}^*) \leq \varphi$, for $0 < \varphi \ll 1$:

$$J(\mathbf{s}) = (CV(\mathbf{s}) - CV^*)^2 + (VM(\mathbf{s}) - VM^*)^2 + (UB(\mathbf{s}) - UB^*)^2. \quad (2)$$

4 STOCHASTIC REACHABILITY PROBLEM

Definition 4.1. A **Markov decision process (MDP)** is a 5-tuple $\mathcal{M} = (S, A, T, J, I)$ consisting of a set of states S , a set of actions A , a transition function $T : S \times A \times S \mapsto [0, 1]$, where $T(\mathbf{s}, a, \mathbf{s}')$ is the probability of transitioning from state $\mathbf{s}$ to state $\mathbf{s}'$ under action a , a cost function $J : S \mapsto \mathbb{R}$, where $J(\mathbf{s})$ is the cost associated with state $\mathbf{s}$, and an initial state distribution $I : S \mapsto [0, 1]$.

Given the stochasticity introduced by PSO, the V-formation problem can be formulated in terms of a reachability problem for a Markov chain, induced by the composition of a Markov decision process (MDP) and a controller. The MDP $\mathcal{M}$ modeling a flock of B birds is defined as follows. The set of states S is $S = \mathbb{R}^{4B}$, as each bird has a 2D position and a 2D velocity vector, and the flock contains B birds. The set of actions A is $A = \mathbb{R}^{2B}$, as each bird takes a 2D acceleration action and there are B birds. The cost function J is defined by Eq. 2. The transition function T is defined by Eq. 1. As the acceleration vector $\mathbf{a}_i(t)$ for bird i at time t is a random variable, the state vector $\mathbf{s}_i = \{\mathbf{x}_i(t+1), \mathbf{v}_i(t+1)\}$ is also a random variable. The initial state distribution I is a uniform distribution from a region of state space where all birds have positions and velocities in a range defined by fixed lower and upper bounds.

Before we can define traces, or executions, of $\mathcal{M}$, we need to fix a controller, or strategy, that determines which action from A to use at any given state of the system. We focus on randomized strategies. A *randomized strategy (controller)* σ over $\mathcal{M}$ is a function of the form $\sigma : S \mapsto PD(A)$, where $PD(A)$ is the set of probability distributions over A . That is, σ takes a state $\mathbf{s}$ and returns an action consistent with the probability distribution $\sigma(\mathbf{s})$. Applying a policy σ to the MDP $\mathcal{M}$ defines the Markov chain. $\mathcal{M}_\sigma$. We use the terms strategy and controller interchangeably.

In the bird-flocking problem, a controller would be a function that determines the accelerations for all the birds given their current positions and velocities. Once we fix a controller, we can iteratively use it to (probabilistically) select a sequence of flock accelerations. The goal is to generate a sequence of actions that takes an MDP from an initial state $\mathbf{s}$ to a state $\mathbf{s}^*$ with $J(\mathbf{s}^*) \leq \varphi$.

We approach the stochastic reachability problem by designing a controller and quantifying its probability of success in reaching the goal states. In [12], a stochastic reachability problem was solved by appropriately designing centralized controllers σ . In this paper,

we design a distributed procedure with an adaptive horizon and adaptive neighborhood resizing and evaluate its performance.

Definition 4.2. Let $\mathcal{M} = (S, A, T, J, I)$ be an MDP, and let $G \subseteq S$ be the set of goal states $G = \{s | J(s) \leq \varphi\}$ of $\mathcal{M}$. The **stochastic reachability problem** is to design a controller $\sigma : S \mapsto PD(A)$ for $\mathcal{M}$ such that for a given δ , the probability of the underlying Markov chain $\mathcal{M}_\sigma$ to reach a state in G in m steps (for a given m) starting from an initial state, is at least $1 - \delta$.

5 ADAPTIVE-NEIGHBORHOOD DISTRIBUTED CONTROL

In contrast to [12, 16], we consider a distributed setting with the following assumptions about the system model. (1) Each bird is equipped with the means for communication without delays. The communication radius of each bird i changes its size adaptively. The measure of the radius is the number of birds covered and we refer to it as the bird's local neighborhood N_i , including the bird itself. (2) All birds use the same algorithm to satisfy their local reachability goals, i.e. to bring the local cost $J(s_{N_i})$, $i \in \{1, \dots, B\}$, below the given threshold φ . (3) The birds move in continuous space and change accelerations synchronously at discrete time points. (4) After executing its local algorithms, each bird broadcasts the obtained solutions to its neighbors. This way every bird receives solution proposals, which differ due to the fact that each bird has its own local neighborhood. To find consensus, each bird takes as its best action the one with the minimal cost among the received proposals. The solutions for the birds in the considered neighborhood are then fixed. The consensus rounds repeat until all birds have fixed solutions. (5) Every time step the value of the cost function $J(s)$ is obtained globally for all birds in the flock and checked for improvement. The neighborhood for each bird is then resized based on this global check. (6) The upwash modeled in Section 3 maintains connectivity of the flock along the computations, while our algorithm manages collision avoidance.

The central result presented in this paper is a *distributed adaptive-neighborhood and adaptive-horizon model-predictive control algorithm* we call DAMPC. At each time step, each bird runs AMPC to determine the best acceleration for itself and its neighbors (while ignoring the birds outside its neighborhood). The birds then exchange the computed accelerations with their neighbors, and the whole flock arrives at a consensus that assigns each bird to a unique (fixed) acceleration value. Before reaching consensus, it may be the case that some of i 's neighbors already have fixed solutions (accelerations) – these accelerations are not updated when i runs AMPC. A key idea of our algorithm is to adaptively resize the extent of a bird's neighborhood.

5.1 The Distributed AMPC Algorithm

DAMPC (see Alg. 1) takes as input an MDP $\mathcal{M}$, a threshold φ defining the goal states G , the maximum horizon length h_{max} , the maximum number of time steps m , the number of birds B , and a scaling factor β . It outputs a state s_0 in I and a sequence of actions $\mathbf{a}^{1:m}$ taking $\mathcal{M}$ from s_0 to a state in G .

The initialization step (Line 1) chooses an initial state s_0 from I , fixes an initial level ℓ_0 as the cost of s_0 , sets the initial time t and number of birds to process k . The outer while-loop (Lines 2-22) is

Table 1: Table of Notation

H, h_i	$\triangleq$	Maximum and current local horizon lengths
N_i	$\triangleq$	neighborhood of the i 's bird
k	$\triangleq$	the number of birds in the neighborhood ($ N_i $)
m	$\triangleq$	number of time-steps allowed by the property φ
$\mathbf{a}^{1:m}$	$\triangleq$	sequence of synthesized acceleration for all birds for each time-step
$?$	$\triangleq$	acceleration that has not yet been fixed
$1, !$	$\triangleq$	superscript for the first and last, respectively, elements in the horizon sequence
$\mathbf{a}_{N_i}^{1:t}, \mathbf{s}_{N_i}^{1:t}$	$\triangleq$	sequence of accelerations and corresponding states of the horizon length reached at time-step t by bird i locally in its neighborhood N_i
Δ_i	$\triangleq$	dynamical threshold defined based on the last achieved local cost $J(\mathbf{s}_{N_i}^1)$ in the neighborhood N_i
$\mathbf{a}^1, \mathbf{s}^1$	$\triangleq$	accelerations and corresponding states for all birds achieved globally as unions of the last elements in the best horizon sequences reached locally in each neighborhood
$\mathbf{a}^1(t), \mathbf{s}^1$	$\triangleq$	accelerations and states for all birds achieved globally as unions of the first elements in the best horizon sequences reached locally in each neighborhood
ℓ_t	$=$	$J(\mathbf{s}^1)$ – level achieved globally at time-step t after applying $\mathbf{a}^{1:t}$ to the current state
Δ	$\triangleq$	dynamical threshold defined based on the last achieved global level

Algorithm 1: DAMPC

```

Input :  $\mathcal{M} = (S, A, T, J, I)$ ,  $\varphi$ ,  $h_{max}$ ,  $m$ ,  $B$ ,  $\beta$ 
Output:  $s_0, \mathbf{a}^{1:m} = [\mathbf{a}(t)]_{1 \leq t \leq m}$ 

1  $s_0 \leftarrow \text{sample}(I)$ ;  $s \leftarrow s_0$ ;  $\ell_0 \leftarrow J(s)$ ;  $t \leftarrow 1$ ;  $k \leftarrow B$ ;  $H \leftarrow h_{max}$ ;
2 while  $(\ell_{t-1} > \varphi) \wedge (t < m)$  do
3    $\forall i: \mathbf{a}_i^{1:t}(t) \leftarrow ?$ ; // No bird has a fixed solution yet
4   while  $(R \leftarrow \{j \mid \mathbf{a}_j(t) = ?\}) \neq \emptyset$  do
5     for  $i \in R$  do in parallel
6        $N_i \leftarrow \text{Neighbors}(i, k)$ ; //  $k$  neighbors of  $i$ 
7        $\Delta_i \leftarrow J(\mathbf{s}_{N_i}^1) / (m - t)$ ;
8        $(\mathbf{s}_{N_i}^{1:t}, \mathbf{a}_{N_i}^{1:t}) \leftarrow \text{LocalAMPC}(\mathcal{M}, \mathbf{s}_{N_i}^{1:t}, \mathbf{a}_{N_i}^{1:t}, \Delta_i, H, \beta)$ ;
9     end
10     $i^* \leftarrow \arg \min_{j \in R} J(\mathbf{s}_{N_j}^1)$ ; // Best solution in  $R$ 
11    // Fix  $i^*$ 's neighbors solutions
12    for  $i \in \text{Neighbors}(i^*, k)$  do
13       $\mathbf{a}_i^{1:t}(t) \leftarrow \mathbf{a}_{N_{i^*}}^{1:t}[i]$ ; // The solution for bird  $i$ 
14    end
15  end
16  // First action and next state
17   $\mathbf{a}(t) \leftarrow \mathbf{a}^1(t)$ ;  $\mathbf{s}^1 \leftarrow \bigcup_i \mathbf{s}_{N_i}^1$ ;  $\mathbf{s}^1 \leftarrow \bigcup_i \mathbf{s}_{N_i}^1$ ;  $s \leftarrow \mathbf{s}^1$ ;
18  if  $\ell_{t-1} - J(\mathbf{s}^1) > \Delta$  then
19     $\ell_t \leftarrow J(\mathbf{s}^1)$ ;  $t \leftarrow t + 1$ ; // Proceed to the next level
20  end
21   $k \leftarrow \text{NeighSize}(J(\mathbf{s}^1), k)$ ; // Adjust neighborhood size
22 end

```

Figure 2: (a) First round of synchronization for neighborhood size four where Bird 2 runs Local AMPC taking as an input for PSO accelerations of Birds 1, 3, and 4 with ? value. (b) Second synchronization round where Bird 5 takes as an input for PSO fixed accelerations of Birds 3 and 4, and value ? for acceleration of Bird 6. (c) Third synchronization round during the same time step where Bird 7 is the only one whose acceleration has not been fixed yet and it simply has to compute the solution for its neighborhood given fixed accelerations of Birds 4, 5, and 6.

active as long as $\mathcal{M}$ has not reached G and time has not expired. In each time step, DAMPC first sets the sequences of accelerations $\mathbf{a}_i^{1:!(t)}$ for all i to ? (not yet fixed), and then iterates lines 4-15 until all birds fix their accelerations through global consensus (Line 10). This happens as follows. First, all birds determine their neighborhood (subflock) N_i and the cost decrement Δ_i that will bring them to the next level (Lines 6-7). Second, they call LocalAMPC (see Section 5.2), which takes sequences of states and actions fixed so far and extends them such that (line 8) the returned sequence of actions $\mathbf{a}_{N_i}^{1:!}$ and corresponding sequence of states $\mathbf{s}_{N_i}^{1:!}$ decrease the cost of the subflock by Δ_i . Here notation $1:!$ means the whole sequence including the last element ! (some number, the farthest point in the future where the state of the subflock is fixed), which can differ from one neighborhood to another depending on the length of used horizon. Note that an action sequence passed to LocalAMPC as input $\mathbf{a}_{N_i}^{1:!}$ contains ? and the goal is to fill in the gaps in solution sequence by means of this iterative process. In Line 10 we use the value of the cost function in the last resulting state $J(\mathbf{s}_{N_i}^!)$ as a criterion for choosing the best action sequence proposed among neighbors $j \in R$. Then the acceleration sequences of all birds in this subflock are fixed (Lines 12-14).

After all accelerations sequences are fixed, that is all ? are eliminated, the first accelerations in this sequence are selected for the output (Line 17). The next state $\mathbf{s}^1$ is set to the union of $\mathbf{s}_{N_i}^1$ for all neighbors $i = 1:B$, the state of the flock after executing $\mathbf{a}(t)$ is set to the union of $\mathbf{s}_{N_i}^1$. If we found a path that eventually decreases the cost by Δ , we reached the next level, and advance time (Lines 18-20) In that case, we optionally decrease the neighborhood, and increase it otherwise (Line 21).

The algorithm is distributed and with a dynamically changing topology. Lines 4, 10, and 18 require synchronization which can be achieved by broadcasting corresponding information to a central hub of the network. This can be a different bird or a different base station at each time step.

Fig. 2 illustrates DAMPC for synchronization rounds within two consecutive time steps including neighborhood resizing. Bigger yellow circles represent birds that are running LocalAMPC. Smaller

blue circles represent birds whose acceleration sequences are not completely fixed yet. Black squares mark birds with already fixed accelerations. Connecting lines are neighborhood relationship.

Working with a real CPS flock requires careful consideration of energy consumption. Our algorithm accounts for this by using the smallest neighborhood necessary during next control input computations. Regarding deployment, we see the following approach. Alg. 2 can be implemented as a local controller on each drone and communication will require broadcasting positions and output of the algorithm to other drones in the neighborhood through a shared memory. In this case, according to Alg. 1, a central agent will be needed to periodically compute the global cost and resize the neighborhood.

5.2 The Local AMPC Algorithm

LocalAMPC is a modified version of the AMPC algorithm [16], as shown in Alg. 2. Its input is an MDP $\mathcal{M}$, the current state $\mathbf{s}_{N_i}^{1:!}$ of a subflock N_i , a vector of acceleration sequences $\mathbf{a}_{N_i}^{1:!}$, one sequence for each bird in the subflock, a cost decrement Δ_i to be achieved, a maximum horizon H and a scaling factor β . In $\mathbf{a}_{N_i}^{1:!}$ some accelerations may not be fixed yet, that is, they have value ?.

Its output is a vector of acceleration sequences $\mathbf{a}_{N_i}^{1:!}$, one for each bird, that decreased the cost of the flock at most, the state $\mathbf{s}_{N_i}^{1:!}$ of the subflock after executing all actions.

Algorithm 2: LocalAMPC

Input : $\mathcal{M} = (S, A, T, J, I)$, $\mathbf{s}_{N_i}^{1:!}$, $\mathbf{a}_{N_i}^{1:!}$, Δ_i , H , β
Output: $\mathbf{s}_{N_i}^{1:!}$, $\mathbf{a}_{N_i}^{1:!}$

```

1  $p \leftarrow 2 \cdot \beta \cdot B$ ; // Initial swarm size
2  $h_i \leftarrow 1$ ; // Initial horizon  $\forall j \in N_i : \mathbf{a}_j^1 = ?$ 
3 repeat
4   // Run PSO with local information  $\mathbf{s}_{N_i}^{1:!}$  and  $\mathbf{a}_{N_i}^{1:!}$ 
5    $(\mathbf{ts}_{N_i}^{1:!}, \mathbf{ta}_{N_i}^{1:!}) \leftarrow \text{PSO}(\mathcal{M}, \mathbf{s}_{N_i}^{1:!}, \mathbf{a}_{N_i}^{1:!}, p, h_i)$ ;
6    $h_i \leftarrow h_i + 1$ ;  $p \leftarrow 2 \cdot \beta \cdot h_i \cdot B$ ; // increase horizon, swarm size
7 until  $(J(\mathbf{ts}_{N_i}^!) - \ell_{t-1} < \Delta_i) \wedge (h_i \leq H)$ 
8  $\mathbf{s}_{N_i}^{1:!} \leftarrow \mathbf{ts}_{N_i}^{1:!}$ ;  $\mathbf{a}_{N_i}^{1:!} \leftarrow \mathbf{ta}_{N_i}^{1:!}$ ; // Return temporary sequences
```

LocalAMPC first initializes (Line 1) the number of particles p to be used by PSO, proportionally to the input horizon h_i , to the number of birds B , and the scaling factor β . It then tries to decrement the cost of the subflock by at least Δ_i , as long as the maximum horizon H is not reached (Lines 3-7).

For this purpose it calls PSO (Line 5) with an increasingly longer horizon, and an increasingly larger number of particles. The idea is that the flock might have to first overcome a cost bump, before it gets to a state where the cost decreases by at least Δ_i . PSO extends the input sequences of fixed actions to the desired horizon with new actions that are most successful in decreasing the cost of the flock, and it computes from scratch the sequence of actions, for the ? entries. The result is returned in $\mathbf{a}_{N_i}^{1:!}$. PSO also returns the states $\mathbf{s}_{N_i}^{1:!}$ of the flock after applying the whole sequence of actions. Using this information, it computes the actual cost achieved.

Figure 3: Left: Blue bars are the values of the cost function in every time step. Red dashed line in the value of the Lyapunov function serving as a threshold for the algorithm. Black solid line is resizing of the neighborhood for the next step given the current cost. Right: Step-by-step evolution of the flock from an arbitrary initial configuration in the left lower corner towards a V-formation in the right upper corner of the plot.

LEMMA 5.1 (LOCAL CONVERGENCE). Given $\mathcal{M} = (S, A, T, J, I)$, an MDP with cost function cost , and a nonempty set of target states $G \subset S$ with $G = \{s \mid J(s) \leq \varphi\}$. If the transition relation T is controllable with actions in A for every (local) subset of agents, then there exists a finite (maximum) horizon $h_{\max}$ such that LocalAMPC is able to find the best actions $\mathbf{a}_{N_i}^{1:h}$ that decreases the cost of a neighborhood of agents in the states $s_{N_i}^{1:h}$ by at least a given Δ .

PROOF. In the input to LocalAMPC, the accelerations of some birds in N_i may be fixed (for some horizon). As a consequence, the MDP $\mathcal{M}$ may not be fully controllable within this horizon. Beyond this horizon, however, PSO is allowed to freely choose the accelerations, that is, the MDP $\mathcal{M}$ is fully controllable again. The result now follows from convergence of AMPC (Theorem 1 from [16]). $\square$

5.3 Dynamic Neighborhood Resizing

The key feature of DAMPC is that it *adaptively resizes neighborhoods*. This is based on the following observation: *as the agents are gradually converging towards a global optimal state, they can explore smaller neighborhoods* when computing actions that will improve upon the current configuration.

Adaptation works on lookahead cost, which is the cost that is reachable in some future time. Line 19 of DAMPC is reached (and the level ℓ_t is incremented) whenever we are able to decrease this look-ahead cost. If level ℓ_t is incremented, neighborhood size $k \in [k_{\min}, k_{\max}]$ is decremented, and incremented otherwise, as follows: $\text{NeighSize}(J, k) =$

$$\begin{cases} \min \left(\max \left(k - \left\lceil \left(1 - \frac{J(s^*)}{k} \right) \right\rceil, k_{\min} \right), k_{\max} \right), & \text{the next level} \\ \min(k + 1, k_{\max}), & \text{otherwise.} \end{cases} \quad (3)$$

In Fig. 3 we depict a simulation-trace example, demonstrating how levels and neighborhood size are adapting to the current value of the cost function.

6 CONVERGENCE AND STABILITY

Since we are solving a nonlinear nonconvex optimization problem, the cost J itself may not decrease monotonically. However, the look-ahead cost – the cost of some future reachable state – monotonically decreases. These costs are stored in level variables ℓ_t in Algorithm DAMPC and they define a Lyapunov function V .

$$V(t) = \ell_t \quad \text{for levels } t = 0, 1, 2, \dots \quad (4)$$

where the levels decrease by at least a minimum dynamically defined threshold: $V(t+1) < V(t) - \Delta$.

LEMMA 6.1. $V(t) : \mathbb{Z} \rightarrow \mathbb{R}$ defined by (4) is a valid Lyapunov function, i.e., it is positive-definite and monotonically decreases until the system reaches its goal state.

PROOF. Note that the cost function $J(s)$ is positive by definition, and since ℓ_t equals $J(s)$ for some state s , V is nonnegative. Line 18 of Algorithm DAMPC guarantees that V is monotonically decreasing by at least Δ . $\square$

LEMMA 6.2 (GLOBAL CONSENSUS). Given Assumptions 1-7 in Section 5, all agents in the system will fix their actions in a finite number of consensus rounds.

PROOF. During the first consensus round, each agent i in the system runs LocalAMPC for its own neighborhood N_i of the current size k . Due to Lemma 5.1, $\exists \hat{h}$ such that a solution, i.e. a set of action (acceleration) sequences of length $\hat{h}$, will be found for all agents in the considered neighborhood N_i . Consequently, at the end of the round the solutions for at least all the agents in N_{i^*} , where i^* is the agent which proposed the globally best solution, will be fixed. During the next rounds the procedure recurses. Hence, the set R of all agents with nfy values is monotonically decreasing with every consensus round. $\square$

Global consensus is reached by the system during communication rounds. However, to achieve the global optimization goal we prove that the consensus value converges to the desired property.

Definition 6.3. Let $\{s(t) : t = 1, 2, \dots\}$ be a sequence of random vector-variables and s^* be a random or non-random. Then $s(t)$ **converges with probability one** to s^* if

$$\mathbb{P} \left[\bigcup_{\varepsilon > 0} \bigcap_{N < \infty} \bigcup_{n \geq N} |s(t) - s^*| \geq \varepsilon \right] = 0.$$

LEMMA 6.4 (MAX-NEIGHBORHOOD CONVERGENCE). If DAMPC is run with constant neighborhood size B , then it behaves identically to centralized AMPC.

PROOF. If DAMPC uses neighborhood B , then it behaves like the centralized AMPC, because the accelerations of all birds are fixed in the first consensus round. $\square$

THEOREM 6.5 (GLOBAL CONVERGENCE). Let $\mathcal{M} = (S, A, T, J, I)$ be an MDP with a positive and continuous cost function J and a nonempty set of target states $G \subset S$, with $G = \{s \mid J(s) \leq \varphi\}$. If there exists a finite horizon $h_{\max}$ and a finite number of execution steps m , such that centralized AMPC is able to find a sequence of actions $\{a(t) : t = 1, \dots, m\}$ that brings $\mathcal{M}$ from a state in I to a state in G , then DAMPC is also able to do so, with probability one.

PROOF. We illustrate the proof by our example of flocking. Note that the theorem is valid in the general formulation above for the fact that as global Lyapunov function approaches zero, the local dynamical thresholds will not allow neighborhood solutions to significantly diverge from reaching the state obtained as a result of repeated consensus rounds. Owing to Lemma 5.1, after the first consensus round, Alg. 2 finds a sequence of best accelerations of length h_{i^*} , for birds in subflock N_{i^*} , decreasing their cost by Δ_{i^*} . In the next consensus round, birds j outside N_{i^*} have to adjust the accelerations for their subflock N_j , while keeping the accelerations of the neighbors in $N_{i^*} \cap N_j$ to the already fixed solutions. If bird j fails to decrease the cost of its subflock N_j with at least Δ_j within prediction horizon h_{i^*} , then it can explore a longer horizon h_j up to h_{max} . This allows PSO to compute accelerations for the birds in $N_{i^*} \cap N_j$ in horizon interval $h_j < h \leq h_{i^*}$, decreasing the cost of N_j by Δ_j . Hence, the entire flock decreases its cost by Δ (this defines Lyapunov function V in Eq. 4) ensuring convergence to a global optimum. If h_{max} is reached before the cost of the flock was decreased by Δ , the size of the neighborhood will be increased by one, and eventually it would reach B . Consequently, using Theorem 1 in [16], there exists a horizon h_{max} that ensures global convergence. For this choice of h_{max} and for maximum neighborhood size, the cost is guaranteed to decrease by Δ , and we are bound to proceed to the next level in DAMPC. The Lyapunov function on levels guarantees that we have no indefinite switching between “decreasing neighborhood size” and “increasing neighborhood size” phases, and we converge (see Fig. 1). $\square$

Fig. 1 illustrates the proof of global convergence of our algorithm, where we overcome a local minimum by gradually adapting the neighborhood size to proceed to the next level defined by the Lyapunov function. In the plot on the right, we see 7 birds starting from an arbitrary initial state near the origin $(x, y) = (0, 0)$, and eventually reaching V-formation at position $(x, y) \approx (300, 100)$. However, around $x \approx 50$, the flock starts to drift away from a V-formation, but our algorithm is able to bring it back to a V-formation. Let us see how this is reflected in terms of changing cost and neighborhood sizes. In the plot on the left, we see the cost starting very high (blue lines), but mostly decreasing with time steps initially. When we see an unexpected rise in cost value at time steps in the range 11–13 (corresponding to the divergence at $x \approx 50$), our algorithm adaptively increases the horizon h first, and eventually the neighborhood size, which eventually increases back to 7, to overcome the divergence from V-formation, and maintain the Lyapunov property of the red function. Note that the neighborhood size eventually decreases to three, the minimum for maintaining a V-formation.

The result presented in [16] applied to our distributed approach, together with Theorem 6.5, ensure the following corollary.

COROLLARY 6.6 (GLOBAL STABILITY). *Assume the set of target states $G \in S$ has been reached and one of the following perturbations of the system dynamics has been applied: a) the best next action is chosen with probability zero (crash failure); b) an agent is displaced (sensor noise); c) an action of a player with opposing objective is performed. Then applying Algorithm 1 the system converges with probability one from a disturbed state to a state in G .*

7 EXPERIMENTAL RESULTS

We comprehensively evaluated DAMPC to compute statistical estimates of the success rate of reaching a V-formation from an arbitrary initial state in a finite number of steps m . We considered flocks of size $B = \{5, 7, 9\}$ birds. The specific reachability problem we addressed is as follows. Given a flock MDP $\mathcal{M}$ with B birds and the randomized strategy $\sigma : S \mapsto PD(A)$ of Alg. 1, estimate the probability of reaching a state s where the cost function $J(s) \leq \varphi$, starting from an initial state in the underlying Markov chain $\mathcal{M}_\sigma$ induced by σ on $\mathcal{M}$.

Since the exact solution to this stochastic reachability problem is intractable (infinite/continuous state and action spaces), we solve it approximately using statistical model checking (SMC). In particular, as the probability estimate of reaching a V-formation under our algorithm is relatively high, we can safely employ the *additive error* (ϵ, δ) -Monte-Carlo-approximation scheme [8]. This requires L i.i.d. executions (up to a maximum time horizon), determining in Z_l if execution l reaches a V-formation, and returning the mean of the random variables $Z_1, \dots, Z_L$. We compute $\bar{\mu}_Z = \sum_{l=1}^L Z_l / L$ by using Bernstein’s inequality to fix $L \propto \ln(1/\delta)/\epsilon^2$ and obtain $\mathbb{P}[\mu_Z - \epsilon \leq \bar{\mu}_Z \leq \mu_Z + \epsilon] \geq 1 - \delta$, where $\bar{\mu}_Z$ approximates μ_Z with additive error ϵ and probability $1 - \delta$. In particular, we are interested in a Bernoulli random variable Z returning 1 if the cost $J(s)$ is less than φ and 0 otherwise. In this case, we can use the Chernoff-Hoeffding instantiation of the Bernstein’s inequality, and further fix the proportionality constant to $N = 4 \ln(2/\delta)/\epsilon$ [9]. Executing the algorithm 10^3 times for each flock size gives us a confidence ratio $\delta = 0.05$ and an additive error of $\epsilon = 10^{-2}$.

We used the following parameters: number of birds $B \in \{5, 7, 9\}$, cost threshold $\varphi = 10^{-1}$, maximum horizon $h_{max} = 3$, number of particles in PSO $p = 200 \cdot h \cdot B$. DAMPC is allowed to run for a maximum of $m = 60$ steps. The initial configurations are generated independently, uniformly at random, subject to the following constraints on the initial positions and velocities: $\forall i \in \{1, \dots, B\} \mathbf{x}_i(0) \in [0, 3] \times [0, 3]$ and $\mathbf{v}_i(0) \in [0.25, 0.75] \times [0.25, 0.75]$. To perform the SMC evaluation of DAMPC, and to compare it with the centralized AMPC from [16], we designed the above experiments for both algorithms in C, and ran them on the 2x Intel Xeon E5-2660 Okto-Core, 2.2 GHz, 64 GB platform.

Our experimental results are given in Table 2. We used three different ways of computing the average number of neighbors for successful runs. Assuming a successful run converges after m' steps, we (1) compute the average over the first m' steps, reported as “for good runs until convergence”; (2) extend the partial m' -step run into a full m -step run and compute the average over all m steps, reported as “for good runs over m steps”; or (3) take an average across $> m$ steps, reported as “for good runs after convergence”, to illustrate global stability.

We obtain a high success rate for 5 and 7 birds, which does not drop significantly for 9 birds. The average convergence duration, horizon, and neighbors, respectively, increase monotonically when we consider more birds, as one would expect. The average neighborhood size is smaller than the number of birds, indicating that we improve over AMPC [16] where all birds need to be considered for synthesizing the next action. We also observe that the average number of neighbors for good runs until convergence is larger than

Table 2: Comparison of DAMPC and AMPC [16] on 10^3 runs.

NUMBER OF BIRDS	DAMPC			AMPC		
	5	7	9	5	7	9
Success rate, $\tilde{\mu}_Z$	0.98	0.92	0.80	0.99	0.95	0.88
Avg. convergence duration, m	7.40	10.15	15.65	9.01	12.39	17.29
Avg. horizon, h	1.35	1.36	1.53	1.29	1.55	1.79
Avg. execution time in sec.	295s	974s	$\propto 10^3s$	644s	3120s	$\propto 10^4s$
Avg. neighborhood size, k						
for good runs until convergence	3.69	5.32	6.35	5.00	7.00	9.00
for good runs over m steps	3.35	4.86	5.58	5.00	7.00	9.00
for good runs after convergence	4.06	5.79	6.75	5.00	7.00	9.00
for bad runs	4.74	6.43	6.99	5.00	7.00	9.00

the one for bad runs, except for 5 birds. The reason is that in some bad runs the cost drops quickly to a small value resulting in a small neighborhood size, but gets stuck in a local minimum (e.g., the flock separates into two groups) due to the limitations imposed by fixing the parameters h_{max} , p , and m . The neighborhood size remains small for the rest of the run leading to a smaller average.

Finally, compared to the centralized AMPC [16], DAMPC is faster (e.g., two times faster for 5 birds). Our algorithm takes fewer steps to converge. The average horizon of DAMPC is smaller. The smaller horizon and neighborhood sizes, respectively, allow PSO to speed up its computation.

8 CONCLUSIONS

We introduced DAMPC, a distributed adaptive-neighborhood and adaptive-horizon model-predictive control algorithm, that synthesizes actions for a controllable Markov decision process (MDP), such that the MDP eventually reaches a state with cost close to zero, provided that the MDP has such a state.

The main contribution of DAMPC is that it adaptively resizes an agent's local neighborhood, while still managing to converge to a goal state with high probability. Initially, when the cost value is large, the neighborhood of an agent is the entire multi-agent system. As the cost decreases, however, the neighborhood is resized to smaller values. Eventually, when the system reaches a goal state, the neighborhood size remains around a pre-defined minimal value.

This is a remarkable result showing that the local information needed to converge is strongly related to a cost-based Lyapunov function evaluated over a global system state. While our experiments were restricted to V-formation in bird flocks, our approach applies to reachability problems for any collection of entities that seek convergence from an arbitrary initial state to a desired goal state, where a notion of distance to it can be suitably defined.

One of the main goals for this work was to evaluate the dynamic neighborhood resizing idea and determine if it can result in a relatively small average neighborhood size. Our evaluation shows that this is indeed the case. We plan to focus on exploring alternative stochastic dynamic models. Another direction for improvement is designing a version of DAMPC where each time step starts with all birds running AMPC in parallel. Local solutions are then subsequently combined through a series of information-exchange rounds into a consistent global solution. Finally, we would like to consider an application of DAMPC to control of drone teams, including the investigation of possible communication faults.

Acknowledgments. The first author would like to thank Josef Widder for very valuable feedback. This work was partially supported by the Doctoral Program Logical Methods in Computer Science and the Austrian National Research Network RiSE/SHiNE (S11405-N23 and S11412-N23) project funded by the Austrian Science Fund (FWF) project W1255-N23, and National Science Foundation grant CCF 1423296.

REFERENCES

- [1] [n. d.]. Boeing Copies Flying Geese to Save Fuel. <https://www.bloomberg.com/news/articles/2017-08-08/boeing-nasa-look-to-flying-geese-in-chase-for-jet-fuel-savings>.
- [2] Christopher Amato, Frans A Oliehoek, et al. 2015. Scalable Planning and Learning for Multiagent POMDPs. In *AAAI* 1995–2002.
- [3] Eduardo F Camacho and Carlos Bordons. 2004. Model predictive control. Advanced textbooks in control and signal processing. Springer-Verlag, London (2004).
- [4] Jesus Capitan, Matthijs TJ Spaan, Luis Merino, and Anibal Ollero. 2013. Decentralized multi-robot cooperation with auctioned POMDPs. *The International Journal of Robotics Research* 32, 6 (2013), 650–671.
- [5] R. D'Andrea and G. E. Dullerud. 2003. Distributed Control Design for Spatially Interconnected Systems. *IEEE Trans. Automat. Control* 48, 9 (2003).
- [6] Danny Dolev, Nancy A Lynch, Shlomit S Pinter, Eugene W Stark, and William E Weihl. 1986. Reaching approximate agreement in the presence of faults. *Journal of the ACM (JACM)* 33, 3 (1986), 499–516.
- [7] J. M. Fowler and R. D'Andrea. 2002. Distributed Control of Close Formation Flight. In *Proc. of 41st IEEE Conference on Decision and Control*.
- [8] R. Grosu, D. Peled, C. R. Ramakrishnan, S. A. Smolka, S. D. Stoller, and J. Yang. 2014. Using Statistical Model Checking for Measuring Systems. In *Proc. of the International Symposium Leveraging Applications of Formal Methods, Verification and Validation (LNCS)*, Vol. 8803. Springer, 223–238. https://doi.org/10.1007/978-3-662-45231-8_16
- [9] Thomas Héroult, Richard Lassaigne, Frédéric Magniette, and Sylvain Peyronnet. 2004. Approximate probabilistic model checking. In *International Workshop on Verification, Model Checking, and Abstract Interpretation*. Springer, 73–84.
- [10] Herman Kahn and Theodore E Harris. 1951. Estimation of particle transmission by random sampling. *National Bureau of Standards applied mathematics series* 12 (1951), 27–30.
- [11] J. Kennedy and R. Eberhart. 1995. Particle Swarm Optimization. In *Proc. of 1995 IEEE International Conference on Neural Networks*. 1942–1948.
- [12] Anna Lukina, Lukas Esterle, Christian Hirsch, Ezio Bartocci, Junxing Yang, Ashish Tiwari, Scott A. Smolka, and Radu Grosu. 2017. 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. 286–302. https://doi.org/10.1007/978-3-662-54580-5_17
- [13] Gabriel A Moreno, Javier Cámara, David Garlan, and Bradley Schmerl. 2015. Proactive self-adaptation under uncertainty: a probabilistic model checking approach. In *Proc. of the 2015 10th joint meeting on foundations of software engineering*. ACM, 1–12.
- [14] Gabriel A. Moreno, Alessandro V. Papadopoulos, Konstantinos Angelopoulos, Javier Cámara, and Bradley Schmerl. 2017. Comparing Model-based Predictive Approaches to Self-adaptation: CobRA and PLA. In *Proc. of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '17)*. IEEE Press, Piscataway, NJ, USA, 42–53. <https://doi.org/10.1109/SEAMS.2017.2>
- [15] Lili Su and Nitin H Vaidya. 2016. Fault-tolerant multi-agent optimization: optimal iterative distributed algorithms. In *Proc. of the 2016 ACM Symposium on Principles of Distributed Computing*. ACM, 425–434.
- [16] Ashish Tiwari, Scott A. Smolka, Lukas Esterle, Anna Lukina, Junxing Yang, and Radu Grosu. 2017. 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. Springer, 446–462. https://doi.org/10.1007/978-3-319-68167-2_29
- [17] J. Yang, R. Grosu, S. A. Smolka, and A. Tiwari. 2016. Love Thy Neighbor: V-Formation as a Problem of Model Predictive Control (Extended Abstract). In *Proc. of CONCUR 2016, 27th International Conference on Concurrency Theory*.
- [18] D. Ye, J. Zhang, and Z. Sun. 2017. Extended State Observer-Based Finite-Time Controller Design for Coupled Spacecraft Formation with Actuator Saturation. *Advances in Mechanical Engineering* 9, 4 (2017), 1–13.
- [19] Jingyuan Zhan and Xiang Li. 2013. Flocking of Multi-Agent Systems Via Model Predictive Control Based on Position-Only Measurements. *IEEE Trans. Industrial Informatics* 9, 1 (2013), 377–385. <https://doi.org/10.1109/TII.2012.2216536>