Stop, Think, and Roll: Online Gain Optimization for Resilient Multi-robot Topologies

Marco Minelli[1], Marcel Kaufmann[2], Jacopo Panerati[2], Cinara Ghedini[3], Giovanni Beltrame[2], and Lorenzo Sabattini[1]

[1]Department of Sciences and Methods for Engineering,
Università degli Studi di Modena e Reggio Emilia

[2]Department of Software and Computer Engineering,
Polytechnique Montréal

[3]Departamento de Computação Científica,
Instituto Tecnológico de Aeronáutica

Abstract. Efficient networking of many-robot systems is considered one of the grand challenges of robotics. In this article, we address the problem of achieving resilient, dynamic interconnection topologies in multi-robot systems. In scenarios in which the overall network topology is constantly changing, we aim at avoiding the onset of single points of failure, particularly situations in which the failure of a single robot causes the loss of connectivity for the overall network. We propose a method based on the combination of multiple control objectives and we introduce an online distributed optimization strategy that computes the optimal choice of control parameters for each robot. This ensures that the connectivity of the multi-robot system is not only preserved but also made more resilient to failures, as the network topology evolves. We provide simulation results, as well as experiments with real robots to validate theoretical findings and demonstrate the portability to robotic hardware.

Keywords: resilience; multi-robot systems; connectivity; graph theory; control; online optimization; robotic hardware; experimental validation.

1 Introduction

Robot swarms are defined as one of the grand challenges of Science Robotics [21], and resilient networking in particular is a key technology for their successful implementation and deployment. Yang et al. [21] define resilience as a "property that is about systems that can bend without breaking. Resilient systems are self-aware and self-regulating, and can recover from large-scale disruptions". In this paper, we study the problem of optimizing the resilience of the connectivity of a dynamic multi-robot network, with respect to robotic failures.

When looking at multi-robot and swarm systems, we can observe that they are inherently redundant. Being composed of multiple entities, failure of a single robot does not, in general, cause the complete failure of the multi-robot system. It might still be able to complete its objectives, possibly with a decreased performance by exhibiting graceful degradation.

To cooperate and achieve shared objectives, robots need to exchange information among each other. When considering groups of mobile robots with limited-range communication capabilities, the topology of the robot network changes as the robots move. It is then imperative to impose constraints on the robots' motion, so that connectivity is preserved. Several strategies can be found in the literature to address the connectivity preservation problem [3, 8, 15, 17, 19, 22].

It is worth noting that, when considering real robotic systems, connectivity preservation is not always a sufficient safeguard for network resilience. In general, robots can be prone to failures and in many interconnection topologies, individual failures can generate the immediate loss of connectivity of the overall network, possibly preventing the multi-robot system from achieving the desired cooperative results. The inherent reliability of multi-robot systems (due to redundancy) is then heavily reduced by the presence of single points of failure.

To address this issue, Ghedini et al. propose a control strategy [4] to ensure that a multi-robot system preserves a high level of connectivity, while performing its assigned task. The proposed methodology is based on the combination of multiple control laws aiming at adjusting the interconnection topology when potentially vulnerable topological configurations are identified. In [11], this method was implemented on a real multi-robot system, controlled to perform area coverage, and its performance was evaluated against the shortcomings of the real-world, such as imperfect communication. Again, this is achieved through the combination of different control laws, each associated to a gain. These gain combinations determine the overall performance of the system. To automate the choice of the overall parameter set (i.e., the control law gains), we introduce an offline optimization method [11].

The main drawback of this approach is that the chosen gains are optimal values obtained by averaging metrics over a set of different topologies, therefore losing the advantage of optimizing for the current topology of the network. This is particularly significant for a mobile robot network, where the topology changes dynamically as the robots move.

Consequently, in this paper, we propose an online optimization strategy, that allows the multi-robot system to estimate, at run time, the optimal (or close to optimal) set of control law gains that optimizes the overall performance of the system.

The rest of this paper is organized as follows. The necessary background on network properties is presented in Section 2. The system model and the problem addressed here are discussed in Section 3. Section 4 outlines the control architecture. The optimized control strategy is described in details in Section 5, and the experimental validation on real set-up in presented in Section 6. Finally, Section 7 discusses the future directions and concludes the document.

2   Preliminaries: network properties

Consider an undirected graph $\mathcal{G}$, where $\mathcal{V}(\mathcal{G})$ and $\mathcal{E}(\mathcal{G}) \subset \mathcal{V}(\mathcal{G}) \times \mathcal{V}(\mathcal{G})$ are the vertex set and the edge set, respectively. Moreover, let $W \in \mathbb{R}^{N \times N}$ be the weight

matrix: each element w_{ij} is a positive number if an edge exists between nodes i and j, zero otherwise. Since $\mathcal{G}$ is undirected, then $w_{ij} = w_{ji}$.

Thus, let $\mathcal{L} \in \mathbb{R}^{N \times N}$ be the Laplacian matrix of graph $\mathcal{G}$ and $D = \text{diag}\,(\{k_i\})$ be the degree matrix, where k_i is the degree of the i-th node of the graph, i.e. $k_i = \sum_{j=1}^{N} w_{ij}$. The (weighted) Laplacian matrix of the graph is then defined as $L = D - W$. As it is well known from algebraic graph theory, the Laplacian matrix of an undirected graph $\mathcal{G}$ exhibits some remarkable properties regarding its connectivity [7]. Let $\lambda_i,\ i = 1, \ldots, N$ be the eigenvalues of the Laplacian matrix, then:

- The eigenvalues are real, and can be ordered such that $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_N$
- Define now $\lambda = \lambda_2$. Then, $\lambda > 0$ if and only if the graph is connected. Therefore, λ is defined as the **algebraic connectivity** of the graph: in a weighted graph, λ is a non-decreasing function of each edge weight.

Connected graphs may get disconnected in case of failure of one or more nodes. Different nodes have different roles in maintaining the overall network connectivity. The concept of *centrality* is usually exploited to identify the most important nodes within a graph [9]. In particular, referring to connectivity maintenance, we will consider the concept of *Betweenness Centrality* (*BC*) [20], which establishes higher scores for nodes that are contained in most of the shortest paths between every pair of nodes in the network. According to this definition of centrality, removing the most central nodes might quickly lead to network fragmentation. We, therefore, introduce the following definition of the *Robustness level*.

Definition 1 (Robustness level [6]) *Consider a graph $\mathcal{G}$ with N nodes. Let $[v_1, \ldots, v_N]$ be the list of nodes ordered by descending value of BC. Let $\varphi < N$ be the minimum index $i \in [1, \ldots, N]$ such that, removing nodes $[v_1, \ldots, v_i]$ leads to disconnecting the graph, that is, the graph including only nodes $[v_{\varphi+1}, \ldots, v_N]$ is disconnected. Then, the network robustness level of $\mathcal{G}$ is defined as:*

$$\Theta(\mathcal{G}) = \frac{\varphi}{N} \tag{1}$$

The robustness level defines the fraction of central nodes that need to be removed from the network to obtain a disconnected network. Small values of $\Theta(\mathcal{G})$ imply that a small fraction of node failures may fragment the network. Therefore, increasing this value means increasing the network resilience, that is, its robustness to failures. We observe that $\Theta(\mathcal{G})$ is only an estimate of how far the network is from getting disconnected w.r.t. fraction of nodes removed. In fact, it might be the case that different orderings of nodes with the same BC produce different values of $\Theta(\mathcal{G})$.

While the robustness level refers to the overall state of the network, from a local perspective, a heuristic for estimating the vulnerability level of a node by means of the information acquired from its 1-hop and 2-hops neighbors was

proposed in [6]. We summarize it as follows: let $d(v, u)$ be the shortest path between nodes v and u, i.e., the minimum number of edges that connect nodes v and u. Subsequently, define $\Pi(v)$ as the set of nodes from which v can acquire information:

$$\Pi(v) = \{u \in V(G) : d(v, u) \leq 2\} \tag{2}$$

Moreover, let $|\Pi(v)|$ be the number of elements of $\Pi(v)$. In addition, define $\Pi_2(v) \subseteq \Pi(v)$ as the set of the 2-hop neighbors of v, that comprises only nodes whose shortest path from v is exactly equal to 2 hops, namely:

$$\Pi_2(v) = \{u \in V(G) : d(v, u) = 2\} \tag{3}$$

Now define $L(v, u)$ as the *number of paths* between nodes v and u, and let $Path_\beta(v) \subseteq \Pi_2(v)$ be the set of v's 2-hop neighbors that are reachable through at most β paths, namely:

$$Path_\beta(v) = \{u \in \Pi_2(v) : L(v, u) \leq \beta\} \tag{4}$$

Thus, β defines the threshold for the maximal number of paths between a node v and each of its u neighbors that are necessary to include u in $Path_\beta(v)$. Therefore, using a low value for β, it is possible to identify the most weakly connected 2-hop neighbors. Hence, the value of $|Path_\beta(v)|$ is an indicator of the magnitude of node fragility w.r.t. connectivity, and **the vulnerability level of a node regarding failures** is given by $P_\theta(v) \in (0, 1)$:

$$P_\theta(v) = \frac{|Path_\beta(v)|}{|\Pi(v)|} \tag{5}$$

We will hereafter use $\beta = 1$, in order to identify 2-hop neighbors that are connected by a single path, which can represent a critical situation for network connectivity.

3 System model and problem formulation

We consider a multi-robot system composed of N mobile robots and we assume that each robot is able to communicate with other robots within a communication radius R. The resulting communication topology is represented by an undirected graph $\mathcal{G}$.

Let the state of each robot be its position $p_i \in \mathbb{R}^m$, and let $p = \begin{bmatrix} p_1^T \dots p_N^T \end{bmatrix}^T \in \mathbb{R}^{Nm}$ be the state vector of the multi-robot system. Let each robot be modeled as a single integrator system, whose velocity can be directly controlled:

$$\dot{p}_i = u_i \tag{6}$$

where $u_i \in \mathbb{R}^m$ is a control input. For each robot, the control input has to be defined so that a global objective can be achieved. As an example of a commonly implemented application, in the rest of the paper, we will refer to a scenario in which the robots are controlled to spread in a given area while avoiding collisions.

However, the proposed methodology can be easily extended to other coordinated control objectives.

It is worth noting that coordinated objectives can be achieved only if information can be exchanged among the robots, that is if the communication graph is connected, and the robots keep this property as the system evolves. However, when considering real robotic systems, failures can not be neglected: robots can stop working unexpectedly and become unable to collaborate.

In this paper we combine different control laws, guaranteeing the achievement of a common objective (area coverage, in our case) while ensuring the preservation of the connectivity for the communication graph, even in the presence of failures. The combination of the different control laws aims at maximizing a global performance index. This index defines a trade-off between the area actually covered by the robots, and the level of connectivity of the communication network.

4   Overview of the control architecture

Referring to the kinematic model in Equation (6), in the following, we consider each robot to be controlled by means of a control input defined as the superposition of three different terms, that is:

$$u_i = \sigma u_i^c + \psi u_i^r + \zeta u_i^d \tag{7}$$

The contributions that constitute the control inputs are defined as follows:

- The term $u_i^c \in \mathbb{R}^m$ represents the connectivity preservation control input. The role of this control input is to enforce that, if the communication graph is initially connected, then it will remain connected as the system evolves.
- The term $u_i^r \in \mathbb{R}^m$ represents the topology resilience improvement control input. This term aims at improving the robustness of the topology against failures. In other words, its purpose is to minimize the presence of single points of failure that could induce a disconnection in the communication graph in case of failure of one or more robots.
- The term $u_i^d \in \mathbb{R}^m$ represents the desired control action. This encodes the coordinated objective that the multi-robot system needs to achieve. As a representative example, in this paper, we consider the objective to be the uniform coverage of a given area.
- The terms $\sigma, \psi, \zeta \geq 0$ represent linear combination gains. They define the relative importance of the separate control laws.

It is worth noting that the overall behavior of the multi-robot system is defined by the way in which each individual control action is defined, and by how they are combined. Indeed, a different choice of the linear combination gains leads to a different behavior of the multi-robot system.

In the following subsections we will introduce representative examples of the individual control actions, that we will consider for implementation in the rest of the paper.

4.1   Connectivity preservation

We consider the connectivity preservation control term u_i^c to be designed, as in [18], to ensure that the value of the algebraic connectivity λ never goes below a given threshold $\epsilon > 0$. As in [18], the following *energy function* can be used for generating the decentralized connectivity maintenance control strategy:

$$V\left(\lambda\right) = \begin{cases} \coth\left(\lambda - \epsilon\right) & \text{if } \lambda > \epsilon \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

The control law is then designed to drive the robots to perform a gradient descent of $V\left(\cdot\right)$, which ensures preservation of the graph connectivity. Considering the robot model introduced in (6), the control law is defined as follows:

$$u_i = u_i^c = -\frac{\partial V\left(\lambda\right)}{\partial p_i} = -\frac{\partial V\left(\lambda\right)}{\partial \lambda}\frac{\partial \lambda}{\partial p_i}. \tag{9}$$

We observe that the connectivity preservation framework can be enhanced to consider also additional objectives. In particular, as shown in [16], the concept of *generalized connectivity* can be utilized to simultaneously guarantee **connectivity maintenance and collision avoidance** with environmental obstacles and among the robots.

4.2   Topology resilience improvement

We consider the topology resilience improvement control term u_i^r to be designed—in accordance with the methodology defined in [4,5]—to drive the robots toward an improved resilience of the interconnection topology.

Based on the concept of vulnerability level introduced in (5), this control strategy aims at increasing the number of links of a potentially vulnerable node i towards its 2-hop neighbors that are in $Path_\beta(i)$, for a given value of β. Let $x_\beta^i \in \mathbb{R}^m$ be the barycenter of the positions of the robots in $Path_\beta(i)$. Considering the robot model introduced in (6), the control law is defined as follows:

$$u_i^r = \xi_i \frac{x_\beta^i - p_i}{\left\|x_\beta^i - p_i\right\|} \alpha\left(t\right), \tag{10}$$

where $\alpha\left(t\right) \in \mathbb{R}$ is the linear velocity of the robots[1].

Parameter ξ_i takes into account the vulnerability state of a node i, i.e., $\xi_i = 1$ if node i identifies itself as vulnerable or $\xi_i = 0$ otherwise. As in [4,5], we set as vulnerable those robots i exhibiting high values for $P_\theta(i)$: then, ξ_i is defined as follows

$$\xi_i = \begin{cases} 1 \text{ if } P_\theta(i) > r \\ 0 \text{ otherwise,} \end{cases} \tag{11}$$

---

[1] Pathological situations may exist in which (10) is not well defined, namely when $p_i = x_\beta^i$. However, this corresponds to the case where the i-th robot is exactly in the barycenter of its weakly connected 2-hop neighbors: in practice, this never happens when a robot detects itself as vulnerable.

where $r \in (0,1)$ is a random number drawn from a uniform distribution, i.e., if $P_\theta(i) > r$, then the i-th robot considers itself as vulnerable. It is worth remarking that, according to (5), each robot can evaluate its vulnerability level in a decentralized manner.

To summarize, this control law drives the vulnerable robots towards the barycenter of the robots in their $Path_\beta$, thus decreasing their distance to them, thus eventually creating new edges in the communication graph.

4.3   Area coverage and collision avoidance

To control the robot to evenly spread over a given area while avoiding collisions, we propose to use the well-known control strategy based on the Lennard-Jones potential [2].

At distance x from its origin, the potential and the desired control action equations are:

$$P_{LJ} = \iota \left(\left(\frac{\delta}{x} \right)^a - 2 \cdot \left(\frac{\delta}{x} \right)^b \right) ; \qquad u_i^d = -\iota \left(\left(\frac{a \cdot \delta^a}{x^{a+1}} \right)^a - 2 \cdot \left(\frac{b \cdot \delta}{x^{b+1}} \right)^b \right) \quad (12)$$

Parameters ι and δ represent the depth and distance from the origin of the potential's minimum, respectively. Exponents a and b are set to 4 and 2. For the sake of collision avoidance, we set δ to be larger than the communication range of the robots.

5   Optimized control strategy

In this section, we introduce a new methodology to achieve online optimization of the linear combination gains σ, ψ, ζ introduced in (7). The objective is to make the robots able to identify the *best* set of parameters during the evolution of the system.

The *best* solution is defined starting from the system-level objective we are considering, that is, achieving area coverage while keeping a sufficiently high level of connectivity. For this purpose, we define the following objective function:

$$f_{obj}(t) = \lambda(t)\mathcal{A}(t) \quad (13)$$

where $\lambda(t)$ is the algebraic connectivity of the communication graph at time t, and $\mathcal{A}(t)$ is the value of the covered area at time t.

The choice of the objective function is motivated by the fact that this is a fundamentally multi-objective problem of two inversely proportional functions at different scales. A common way of avoiding an adaptive normalization scheme for the two functions is to consider their product [10]. This approach favors solutions that lead to a trade-off between maximizing the algebraic connectivity and maximizing the covered area.

5.1   Proposed optimization method

The proposed optimization method aims at finding the optimal combination of the gains σ, ψ, ζ such that the objective function introduced in (13) is maximized at the end of each iteration. It is worth noting that the objective function is a product of nonlinear contributes, being computed using the algebraic connectivity (which, being an eigenvalue of the Laplacian matrix, is inherently nonlinear) and the covered area (which is calculated considering circular overlapping areas covered by each robot and therefore, it is also nonlinear).

Consequently, we consider the use of optimization methods that are well suited for nonlinear problems. In particular, we adopted the following methods [1]:

1. The **default search** optimization provides a uniform and default screening of the variable domain space. The main advantage of this method is the accuracy of the solution, which can be freely defined if the system performance is independent from the computation time.
2. **Random search**. This stochastic method does not require the gradient of the objective function and can consider non continuous or non differentiable objective functions. The optimal set of parameters is found randomly probing the domain space and selecting the set which returns the highest objective function value. Random search algorithms are typically used to achieve high computational speed at the cost of losing formal guarantees of optimality.
3. The **augmented Lagrangian** optimization algorithm method. Suited for constrained optimization problems, it requires to first penalize the objective function, translate the constrained optimization problem to a series of unconstrained problems, and then adds a term designed to mimic a Lagrange multiplier and improve precision and convergence speed. This algorithm exploits the gradient of the objective function to get the optimal set of parameters. Since the gradient is hard to compute for nonlinear functions such as (13), numerical differentiation is exploited.

5.2   Implementation and evaluation

We compared the optimization methods in terms of quality of the solution and required computation resources. To achieve this, we implemented the following procedure:

1. At each time step, the positions of all the robots are shared with all the other robots, according to the protocol described in [11]. To achieve this in a reasonable time without a fully connected network we use a consensus implementation—*virtual stigmergy* [13]—that was shown to scale with thousands of robots.
2. Based on the shared positions, each robot computes the output of its own control law (7), and its local estimate of each other robot's control law.
3. Every O_p time steps, each robot runs the optimization process to find an updated set of gains to be used in (7) according to the following sub-steps:

(a) Each robot transforms the positions and the control laws in its local reference frame.
(b) A set of G_p different values of the gains is generated (according to the considered optimization method described in Section 5.1).
(c) For each value of the gains, the robots compute the predicted position of all the robots at the subsequent time step, and evaluate the objective function introduced in (13).
(d) The gains that provide the highest value of the objective function are selected.

For evaluating the different optimization methods, we implemented the control law using the Buzz scripting language [12], and run a set of simulations using the multi-physics environment of ARGoS [14]. The optimization framework was written as a C++ module which communicates, by means of a socket, with the simulation set-up.

We evaluate the performance of each optimization method considering a network of 8 robots and compare against a system constant gains, considering all the possible combinations of the following sets:

$$\psi = \{0, 1, 2\} \quad \sigma = \{0, 1, 2\} \quad \zeta = \{0, 1\} \tag{14}$$

The configurations assessed start from the same (randomly selected) initial topology and are compared using the aforementioned three optimization algorithms.

The results of the simulations are summarized in Fig. 1, which depicts the value of the objective function (13) achieved with each optimization algorithm. In particular, the green line represents the objective function obtained with the optimization algorithm, while the red line with the corresponding shadow represents the average value and standard deviation of the objective function obtained with constant gains.

While in general the value of the objective function is typically greater when using the optimization method (with respect to constant gains), results show that *random search* performs significantly better than other methods. Furthermore, the computational requirements are generally smaller for *random search*, in particular if compared with *default search*, whose convergence time is approximately ten times larger.

According to these results, we choose *random search* as the preferred optimization algorithm. Subsequently, we perform a set of simulations to investigate how the choice of the parameters influences the results. In particular, we run simulations for different numbers of generated points $G_p = \{400, 2200, 4000\}$ and for different values of the optimization period $O_p = \{1, 10, 50\}$. The results we obtain from these simulations are summarized in Fig. 2. The sub-plots show the value of the objective function achieved in each configuration. It is possible to observe that different parameter choices provide similar results, in terms of the objective function. Therefore, we conclude that the optimization algorithm can be effectively run choosing the lowest values of generated points (i.e. $G_p = 400$), and with the largest value of the optimization period (i.e. $O_p = 50$): these choices reduce the computational requirements without causing a significant decrease in the quality of the achieved solution.
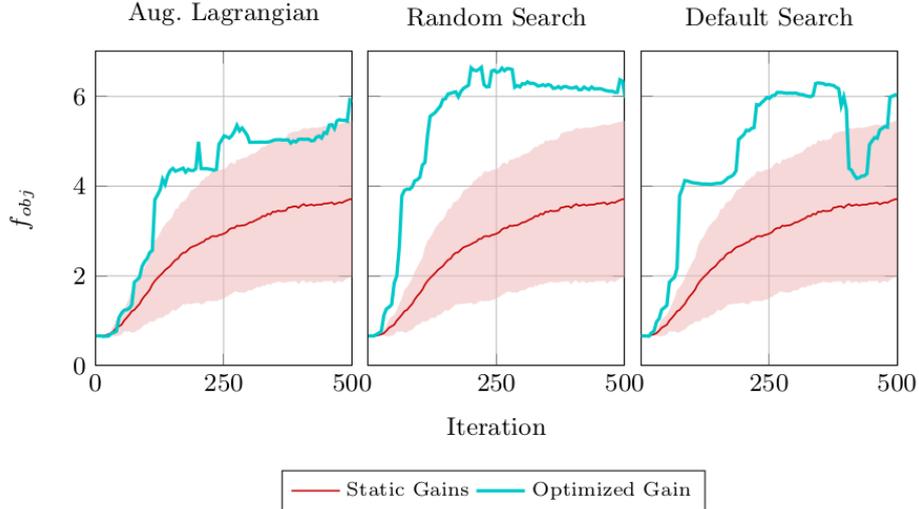
Fig. 1. Objective function evolution comparison. Static gains simulations vs. optimized gains simulations for augmented Lagrangian, random, and default searches.

6   Experimental validation

Segueing from simulation to real robots can be challenging and results in performance degradation, especially with resource constraint hardware [11]. For demonstrating the portability of the proposed online optimization, and to assess how hardware limitations affect the choice of the optimization parameters (i.e., the generated points G_p and optimization period O_p), we transferred our methodology onto a distributed multi-robot system. The robot team consists of eight two-wheeled K-Team Khepera IV (KH4) depicted in Figure 3. Each robot comes with an 800MHz ARM Cortex-A8 and the Yocto operating system[2].

A camera-based tracking system (OptiTrack[3]'s Prime13, see Figure 3), and the `blabbermouth`[4] communication infrastructure are combined to emulate range and bearing sensors for each robot. This also enables point-to-point communication with a limited communication range R between the robots (a similar setup was used in [11]).

The optimization procedure described in Section 5.2 is embedded into the executable `bzzkh4`[5] that runs the Buzz byte code of each robotic controller. Starting from the in simulation investigated parameters, we determine the optimization times Δ_t for processing on the KH4 multi-robot system by varying G_p. A set of optimizations is performed with the initial topology configuration

---

[2] `https://www.k-team.com/mobile-robotics-products/khepera-iv`

[3] `https://optitrack.com/products/prime-13/specs.html`

[4] `https://github.com/MISTLab/blabbermouth`

[5] `https://github.com/MISTLab/BuzzKH4`
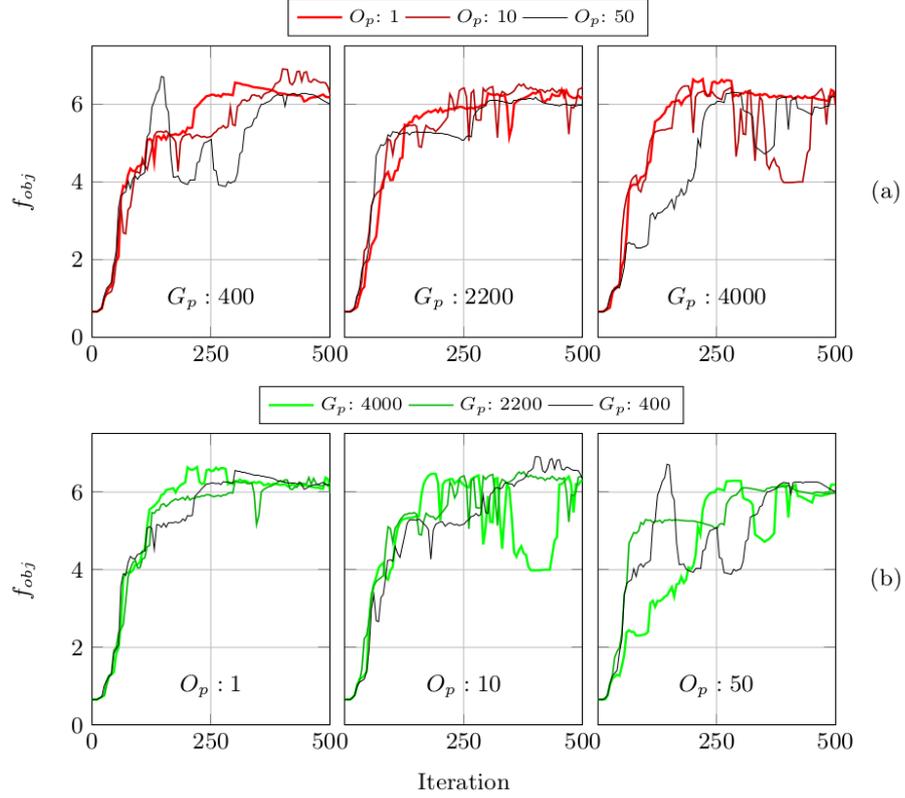
Fig. 2. Random search simulation objective function evolution: (a) comparison at constant generated points, (b) comparison at constant optimization period.

Fig. 3. One of four OptiTrack Prime 13 cameras and one of eight K-Team's Khepera IV robots ($\varnothing= 14.0\,cm, h = 6.0\,cm$) used for the experimental setup in Section 6.

(introduced in Section 5.2). We obtain Δ_t's of $8'41''$, $46'47''$ and $84'23''$ as run-times for 400, 2200 and 4000 generated points G_p, respectively. With increasing G_p, Δ_t increases linearly and ranges from minutes to hours. Considering these computational demands, it is sensible to run the online optimization on the KH4 every $O_p = 50$ steps with $G_p = 400$ points.

Simulations and experimental validation iterate over a fix number of control steps. After 500 such iterations, we consider an experiment to be finished. Due to different processing times on each member of the robot team, the KH4's operate asynchronously. Some reach the end of the experiment earlier and will thus stop communicating and end their operations. Figure 4 shows the evolution of three functions: objective function f_{obj}, the coverage C_a and the connectivity λ_2. The objective function increases, which corresponds to the simulated behavior (compare to Figure 2). The initial decrease in coverage, as well as the increasing of connectivity have been observed in the previous experiments [11].
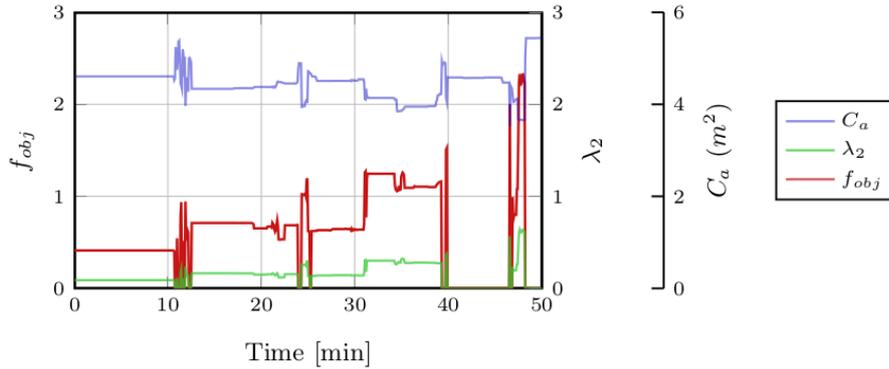
Fig. 4. Objective function f_{obj} evolution achieved in real robot experiments using random search optimization algorithm set with 400 G_p and 50 O_p.

7   Conclusions

In this paper, we propose online optimization to automatically tune the gains of a control law for resilient multi-robot networks. Our starting point is a control law [4] that was proven to increase the robustness of an initially connected multi-robot topology. Here, we extend that work with the following contributions: (i) we implement an online framework to predict and optimize the multi-robot system performance; (ii) within it, we compare three different optimization algorithms; and finally, (iii) we assess the feasibility of implementing this framework on a real robotic setup comprising of eight 2-wheel robots. Simulations demonstrate the effectiveness of the proposed approach as well as its low sensitivity to parameterization. We also demonstrate that the methodology can be executed on robots with limited computational capabilities. Future developments of this work will include the validation of our methodology using ROS-based flying robots and its study with desired control actions other than coverage.

References

1. M. Avriel. *Nonlinear programming: analysis and methods*. Courier Corporation, 2003.
2. M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, Mar 2013.
3. A. Gasparri, L. Sabattini, and G. Ulivi. Bounded control law for global connectivity maintenance in cooperative multi-robot systems. *IEEE Transactions on Robotics*, 33(3):700–717, June 2017.
4. C. Ghedini, C. Ribeiro, and L. Sabattini. Toward fault-tolerant multi-robot networks. *Networks*, 70(4):388–400, 2017.
5. C. Ghedini, C. H. C. Ribeiro, and L. Sabattini. Toward efficient adaptive ad-hoc multi-robot network topologies. *Ad Hoc Networks*, 74:57 – 70, 2018.
6. C. Ghedini, C. Secchi, C. H. C. Ribeiro, and L. Sabattini. Improving robustness in multi-robot networks. In *Proceedings of the IFAC Symposium on Robot Control (SYROCO)*, Salvador, Brazil, aug. 2015.
7. C. Godsil and G. Royle. *Algebraic Graph Theory*. Springer, 2001.
8. M. Ji and M. Egerstedt. Distributed coordination control of multiagent systems while preserving connectedness. *IEEE Transactions on Robotics*, 2007.
9. D. Koschützki, K. A. Lehmann, L. Peeters, S. Richter, D. Tenfelde-Podehl, and O. Zlotowski. Centrality indices. In *Network analysis*, pages 16–61. Springer, 2005.
10. J. Panerati and G. Beltrame. A comparative evaluation of multi-objective exploration algorithms for high-level design. *ACM Trans. Des. Autom. Electron. Syst.*, 19(2):15:1–15:22, Mar. 2014.
11. J. Panerati, M. Minelli, C. Ghedini, L. Meyer, M. Kaufmann, G. Beltrame, and L. Sabattini. Robust connectivity maintenance for fallible robots. *Submitted to Autonomous Robots*, 2018.
12. C. Pinciroli and G. Beltrame. Swarm-oriented programming of distributed robot networks. *Computer*, 49(12):32–41, Dec 2016.
13. C. Pinciroli, A. Lee-Brown, and G. Beltrame. A tuple space for data sharing in robot swarms. In *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (Formerly BIONETICS)*, BICT'15, pages 287–294, ICST, Brussels, Belgium, Belgium, 2016. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
14. C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. Argos: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, Dec 2012.
15. H. A. Poonawala and M. W. Spong. Decentralized estimation of the algebraic connectivity for strongly connected networks. In *American Control Conference (ACC)*, pages 4068–4073. IEEE, 2015.
16. P. Robuffo Giordano, A. Franchi, C. Secchi, and H. H. Bülthoff. A passivity-based decentralized strategy for generalized connectivity maintenance. *The International Journal of Robotics Research*, 32(3):299–323, 2013.
17. L. Sabattini, N. Chopra, and C. Secchi. Decentralized connectivity maintenance for cooperative control of mobile robotic systems. *I. J. Robotic Res.*, 32(12):1411–1423, 2013.
18. L. Sabattini, N. Chopra, and C. Secchi. Decentralized connectivity maintenance for cooperative control of mobile robotic systems. *The International Journal of Robotics Research (SAGE)*, 32(12):1411–1423, October 2013.

19. L. Sabattini, C. Secchi, N. Chopra, and A. Gasparri. Distributed control of multi-robot systems with global connectivity maintenance. *IEEE Transactions on Robotics*, 29(5):1326–1332, October 2013.

20. S. Wasserman, K. Faust, and D. Iacobucci. *Social Network Analysis : Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press, November 1994.

21. G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, et al. The grand challenges of science robotics. *Science Robotics*, 3(14):eaar7650, 2018.

22. P. Yang, , R. A. Freeman, G. J. Gordon, K. M. Lynch, S. S. Srinivasa, and R. Sukthankar. Decentralized estimation and control of graph connectivity for mobile sensor networks. *Automatica*, 46(2):390–396, Feb. 2010.