

A Networked Multi-Agent System for Mobile Wireless Infrastructure on Demand

Miguel Calvo-Fullana, Mikhail Gerasimenko, Daniel Mox, Leopoldo Agorio, Mariana del Castillo, Vijay Kumar, Alejandro Ribeiro, and Juan Andrés Bazerque

Abstract—Despite the prevalence of wireless connectivity in urban areas around the globe, there remain numerous and diverse situations where connectivity is insufficient or unavailable. To address this, we introduce *mobile wireless infrastructure on demand*, a system of UAVs that can be rapidly deployed to establish an ad-hoc wireless network. This network has the capability of reconfiguring itself dynamically to satisfy and maintain the required quality of communication. The system optimizes the positions of the UAVs and the routing of data flows throughout the network to achieve this quality of service (QoS). By these means, task agents using the network simply request a desired QoS, and the system adapts accordingly while allowing them to move freely. We have validated this system both in simulation and in real-world experiments. The results demonstrate that our system effectively offers mobile wireless infrastructure on demand, extending the operational range of task agents and supporting complex mobility patterns, all while ensuring connectivity and being resilient to agent failures.

Index Terms—Multi-robot systems, networked robots, cooperating robots, aerial systems: applications.

I. INTRODUCTION

Mobile wireless connectivity has become widespread, with access to either Wi-Fi, LTE or 5G channels appearing to be always at reach in most cities around the globe. However, outside urban environments connectivity is still far from ubiquitous. Unfortunately, this disparity is intrinsic, as expanding network infrastructure to remote areas is often unprofitable. Indeed, the low cost of Wi-Fi is offset by its low range. While cellular technologies can reach further, they require significant investment on structures and equipment for cell sites. Even in urban environments, reliable communications are not guaranteed, with occasional high-demand events resulting in temporary network capacity overflows and outages. Dimensioning a system to prevent outages is impractical, as

This work was supported by the IoT4Ag Engineering Research Center funded by the National Science Foundation under Award EEC-1941529; and the Agencia Estatal de Investigación under grant RYC2021-033549-I.

M. Calvo-Fullana is with Universitat Pompeu Fabra, Barcelona, Spain (e-mail: miguel.calvo@upf.edu).

M. Gerasimenko was with Tampere University, Tampere, Finland. He is now with Nokia, Finland. (e-mail: mikhail.gerasimenko@nokia.com).

D. Mox was with University of Pennsylvania, Philadelphia, USA. He is now with Zoon, California, USA. (e-mail: dan.c.mox@gmail.com)

L. Agorio and J. A. Bazerque are with University of Pittsburgh, Pittsburgh, USA (e-mail: {lca31,juanbazerque}@pitt.edu).

M. del Castillo is with Universidad de la República, Montevideo, Uruguay (e-mail: mdelcastillo@fing.edu.uy).

V. Kumar and A. Ribeiro are with University of Pennsylvania, Philadelphia, USA (e-mail: {kumar,aribeiro}@seas.upenn.edu).



Fig. 1: Mobile wireless infrastructure on demand. A team of UAVs acting as network providers continuously reconfigure their positions and communication routes, provisioning task users with their required quality of service in terms of wireless connectivity.

it would imply extraordinary costs to cover for outliers that occur with low probability.

The concept of mobile infrastructure on demand (MID) has been proposed to address these issues [1]. The idea of MID is to deploy a team of autonomous unmanned aerial vehicles (UAVs) to provide wireless connectivity on demand (Figure 1). Agents of the MID team reconfigure their positions and communication routes to satisfy quality of service (QoS) requirements of a team of task agents (Figure 2). Task agents can be themselves UAVs or they can be any other type of end user. The distinction between task and MID agents is that the movement of task agents is not controlled, only their communication routes are.

In this paper, we build a MID system. This is a cyber-physical control system that dynamically reconfigures (physical) positions of the MID agents and (cyber) routing tables of MID and task agents. We considered a series of tests in which task agents move apart, follow undisclosed trajectories, and hold their positions while a MID agent is replaced. Through the evaluation of these tests in simulated scenarios and their corroboration in real experiments, we were able to demonstrate that MID enables the following features:

- (F1) **Extension of operational range.** In exploration scenarios the use of a MID team extends the operational range by providing packet relaying capabilities.
- (F2) **Support for tasks with complex mobility patterns.** In tasks whose execution requires complex mobility patterns MID offers more reliable communication. It does so by continuously adapting positions and routes in response to the movement of the task agents.
- (F3) **Robust autonomous reconfiguration.** In scenarios in which MID agents fail or exhaust their batteries, other

team members reconfigure their positions and routes to limit the effect on end-to-end QoS for the task team.

Features **(F1)**–**(F3)** hold significant practical importance across many tasks. One reference application where feature **(F1)** proves valuable is agricultural monitoring. In this particular scenario, a group of agents moves through a crop field or an orchard, gathering measurements on soil humidity or assessing harvest readiness. These measurements must be relayed back in real-time to a coordination center. Feature **(F1)** of the MID system allows for larger fields to be covered using the same amount of communication infrastructure, which is now *mobile* infrastructure. An additional example highlighting feature **(F2)** is a search and rescue mission spanning a large area. In such a scenario, the agents may move in unpredictable ways, and it is crucial for them to maintain communication quality at all times. The reconfiguration capabilities provided by feature **(F2)** ensure the needed communication QoS. Feature **(F3)** complements the previous features **(F1)** and **(F2)** by providing operational flexibility during task execution. Indeed, autonomous network reconfiguration guarantees the best possible connectivity under adverse conditions, such as the temporary recall or recharge of a UAV during deployment, and prevents the abortion of a mission when coping with agent failures during task execution.

While these examples serve to highlight the system’s capabilities in specific applications, our main objective is to offer a flexible solution for providing mobile wireless infrastructure on demand in diverse applications. We are particularly motivated in providing ad-hoc solutions for situations that demand fast deployments during the span of temporary missions. We envision the MID system being utilized in a variety of scenarios, including rescue missions, off-path events, agricultural evaluation, and monitoring of livestock. Additionally, MID can provide wireless service at crowded city gatherings such as sports events, marathons, art festivals, and live concerts. These use cases, among others, serve as a non-exhaustive demonstration of the versatility of our system.

In our modeling abstraction, as depicted in Fig. 2, there is a team of agents carrying out a specific task. Depending on the application at hand, these agents would represent a combination of patrolling UAVs, ground robots collaborating on a mission, livestock with wireless trackers, or even humans using their cell phones or radio equipment. All these scenarios require wireless communication in order to be successful. In this context, our objective is to design a system comprising a second team of agents that provide communication support to the primary team. We argue that UAVs are an appropriate platform for wireless networking as they facilitate fast and flexible deployments and they can reach a proper altitude to provide reliable communication links. Hence, a supplementary team of UAVs will form the system that provides a mobile wireless infrastructure on demand, operating as a set of communication relays that configure an ad-hoc wireless network. While UAVs are ideal for this role, other robotic platforms with different motion dynamics could also be utilized within the same mobile infrastructure system, though they might require distinct motion planning approaches.

Compared with standard mobile applications where the




Fig. 2: Task agents (red) demand a level of connectivity to the MID system, provisioned by the network agents (blue).

positions of the communication base stations are fixed, the ad-hoc network considered here is capable of deploying promptly and being adaptive and robust. The network team can move and track the positions of the task agents in order to maintain and maximize the communication rate they offer. For this optimal design to be possible two technical problems must be solved. First, how to route data across the network in order to maximize throughput for a given spatial arrangement of the task and network teams. Secondly, how to move the network agents to increase the throughput and maintain network connectivity.

The framework on which we base our system includes an optimization scheme for robust routing introduced in [1], which takes the positions of the agents in the primary (task) team as parameters. In this paper, we also incorporate Laplacian methods with stochastic communication channels [2]–[4] to optimize the trajectories of the network agents. A distinctive aspect of our approach is that we define the network team as an autonomous separate entity adapting to the task agents, which do not participate in the path planning. With this concept, the task team can move freely, and our system and control strategy generalizes to provide communication services to a range of applications where task agents could model mobile entities as diverse as robots or livestock.

We further proceed to demonstrate the success of our methods in real-world experiments. For this purpose, we developed an experimental UAV platform shown in Fig. 1, with configurable wireless communication interfaces and sufficient computational power to solve the robust routing and Laplacian optimization in real-time. More crucially, this platform allows us to demonstrate our system for mobile infrastructure on demand, and showcase the features **(F1)**–**(F3)** proposed above, namely, extending the operational range, supporting tasks with complex mobility patterns, and providing a robust reconfigurable network.

A. Related work

Extensive research effort has been dedicated to addressing the challenge of maintaining interconnectivity among a team of

robots [5]. Broadly, this research has focused on investigating the integration of communications and control systems in the context of facilitating coordination among autonomous agents. Achieving accurate coordination necessitates careful consideration of both the wireless channel and the communication network. Graph-theoretic approaches have emerged as a common methodology, leveraging graph connectivity as a practical measure to evaluate network strength [6]–[15]. In these approaches, the edges of the graph represent point-to-point wireless links. These wireless channels can be modeled, in increasing complexity, as binary links [14]–[17], distance-dependent links [8]–[10], [18]–[21], statistical models [22]–[24], or even direct high-fidelity system-level simulations of the wireless communication system [25], [26]. Further works consider, not only connectivity, but also the routing of information through the network [1], [3], [20], [27]–[31]. These latter approaches take a more holistic view of the problem. In the case of fixed task agents, [27] proposes an algorithm that addresses both the placement of network agents to ensure connectivity and their information flow decisions. For dynamic agents, [28] considers maintaining communication flows by dynamically constructing connected graphs through the addition and removal of graph edges, following the approach in [32]. Other metrics have also been considered, with [29] focusing on placement and flow decisions aimed at minimizing latency. Additionally, approaching information flows as packet routing [20] adds an additional layer of nuance beyond graph connectivity. For agents attempting to accomplish a task while maintaining communication, [30], [31] develop robust routing schemes. In [1], task accomplishment is decoupled from communication, introducing the mobile infrastructure on demand system, with robust routing communication and a search-based local planner to control spatial deployment.

Still, the majority of these developments lack extensive experimental validation or remain predominantly theoretical. In this work, we leverage the robust routing principles of [1] and expand on it to support features **(F1)**–**(F3)**. This expansion relies in a more intricate graph-based connectivity planner, inspired by the framework developed in [4].

II. SYSTEM AND PROBLEM FORMULATION

Consider Fig. 2, illustrating a system with $L = N + M$ agents, where a team of N task agents is deployed in an area of interest to carry out a specific task. These are interconnected by a supplementary team of M network agents that will operate as a set of communication relays. The goal of this section is to find the optimal trajectories and routing strategies for the network team. A prerequisite to answering these questions is to adopt a model for the wireless communication links between agents.

A. Wireless channel model

In a free space without obstacles, the wireless transmission of power by an antenna reaches the receptor with a path-loss that depends quadratically on the distance between them [33]. This quadratic loss of power is inevitable since the wavefront is an expanding surface in space that distributes the transmitted




Fig. 3: Characterization of the expected rate function $R(x_i, x_j)$ and one standard deviation following model equations (1) and (2), where $d = \|x_i - x_j\|$ is the distance between the agents x_i and x_j , and $P_{L_0} = -53$ dBm, $n = 2.52$, $P_{N_0} = -70$ dBm, $a = 0.2$ and $b = 0.6$.

power in different directions. Antenna designs are successful in shaping this surface to direct power towards the intended receptor, with each maker providing corresponding radiation patterns and antenna gains, but the quadratic loss remains.

In more crowded environments, waves are typically obstructed, reflected, and superimposed on their way to the target by interfering objects such as buildings and cars. This further affects how much power reaches the receiving antenna. The path-loss is still modeled as decaying with distance but with an exponent larger than two, whose value is estimated from experimental data. Furthermore, in many cases, communication systems are deployed in environments where the interfering objects move relative to the transmitter-receiver pair. This would be typically the case in an urban environment with passing vehicles. Or even in the case of static objects, such as buildings or trees, which are perceived as moving from the inertial system of moving antennas, for instance, when they are assembled on moving vehicles such as UAVs. In this case, shadowing effects of obstructions and reflections are too complex for a deterministic model, and therefore it is standard to resort to probabilistic models. When adopting a stochastic model, the Shannon information theory identifies the signal-to-noise ratio as the main parameter determining the transmission rate that can be achieved over a communication link.

These stochastic aspects of mobile wireless links are captured in the channel model [34] between two antennas at locations $x_i, x_j \in \mathbb{R}^3$, which is characterized in Fig. 3. It relates the transmission power P_{L_0} , noise at the receiver P_{N_0} , distance $d = \|x_i - x_j\|$ and path-loss exponent n to the normalized achievable rate $R_{ij}(x_i, x_j) \in [0, 1]$. These rates are modeled as stochastic variables with mean $\bar{R}_{ij}(x_i, x_j)$ and variance $\tilde{R}_{ij}(x_i, x_j)$ given by

$$\bar{R}_{ij}(x_i, x_j) = \text{erf} \left(\sqrt{\frac{P_{L_0}}{P_{N_0}} \|x_i - x_j\|^{-n}} \right) \quad (1)$$

and

$$\tilde{R}_{ij}(x_i, x_j) = \frac{a \|x_i - x_j\|}{b + \|x_i - x_j\|}. \quad (2)$$

where $\text{erf}(x) = \frac{1}{\sqrt{\pi}} \int_{-x}^x e^{-t^2} dt$ is the Gauss error function and a and b are decay uncertainty parameters. This model is adaptable to diverse hardware by adjusting its parameters to fit empirical measurements [31].

B. Optimal probabilistic routing

In practice, the distance between agents is often large enough to cause sufficient path-loss to impede direct communication. Hence, communication occurs over a network of multiple interconnected wireless links distributed across space as in Fig. 2. Information is transmitted by agents, which flows through this network. We consider agents transmitting up to K information flows. These flows consist of data packets, relayed across the network, with different flows representing distinct data with different requirements, e.g., voice and video streaming. The process of relaying data packets is governed by routing variables $\alpha_{ij}^k \in [0, 1]$. These variables specify the probability, or fraction of time, that the i -th agent sends data to the j -th agent for the k -th data flow. Further, let the rate $b_i^k(\boldsymbol{\alpha}, \mathbf{x})$ be the data injected into the network by the i -th agent for the k -th flow, with the vector $\boldsymbol{\alpha} \triangleq \{\alpha_{ij}^k\}$ collecting all the routing variables and $\mathbf{x} \triangleq \{x_i\}$ collecting all the agent positions. Then, the difference between data transmitted and received at a node is given by the flow equation

$$b_i^k(\boldsymbol{\alpha}, \mathbf{x}) = \sum_{j=1}^L \alpha_{ij}^k R_{ij}(x_i, x_j) - \sum_{j=1}^L \alpha_{ji}^k R_{ji}(x_i, x_j). \quad (3)$$

This flow equation (3) indicates that $b_i^k(\boldsymbol{\alpha}, \mathbf{x})$ must equal the aggregated traffic outgoing to the other nodes of the network during the fractions α_{ij}^k , after subtracting the incoming traffic from the rest of the network into the node. If the node is a network agent, then it does not inject or absorb data, so that $b_i^k(\boldsymbol{\alpha}, \mathbf{x}) = 0$, and hence all the incoming data is forwarded according to the flow equation. Given (3), and assuming that the random variables $R_{ij}(x_i, x_j)$ are independent, the mean and variance of the flows in (3) are given by

$$\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) = \sum_{j=1}^L \alpha_{ij}^k \bar{R}_{ij}(x_i, x_j) - \sum_{j=1}^L \alpha_{ji}^k \bar{R}_{ji}(x_i, x_j) \quad (4)$$

$$\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) = \sum_{j=1}^L (\alpha_{ij}^k)^2 \tilde{R}_{ij}(x_i, x_j) + \sum_{j=1}^L (\alpha_{ji}^k)^2 \tilde{R}_{ji}(x_i, x_j). \quad (5)$$

Using (4) and (5), we can establish quality of service (QoS) requirements for the communication network, setting them as constraints of a routing optimization problem. We consider the case in which the system must guarantee minimum rate requirements $b_i^k \geq m_i^k \geq 0$ at each node of the network. For networking agents that are only intended to forward data, we set $m_i^k = 0$. On the other hand, each task agent can request a minimum data rate guarantee $m_i^k > 0$ from the network. Since the variables involved are stochastic, we must admit a certain probability of outage $1 - \epsilon_i^k$. That is, the constraint can only be satisfied probabilistically as follows

$$\mathbb{P} [b_i^k(\boldsymbol{\alpha}, \mathbf{x}) \geq m_i^k] \geq \epsilon_i^k. \quad (6)$$

In this case, $1 - \epsilon_i^k$ is the probability of outage (the requirement not being satisfied). Thus, the quality of service that a task agent can demand from the network team is the pair (m_i^k, ϵ_i^k) . That is, to provide an average data rate of m_i^k with a confidence level ϵ_i^k . In order to expand expression (6), let

us consider achievable rates $R_{ij}(x_i, x_j)$ modeled as Gaussian variables. Then, the following bound holds

$$\frac{\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) - m_i^k}{\sqrt{\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})}} \geq \Phi^{-1}(\epsilon_i^k) \quad (7)$$

with $\Phi^{-1}(\cdot)$ being the inverse cumulative Gaussian distribution function. Additionally, we define

$$q_i^k \triangleq \bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) - \Phi^{-1}(\epsilon_i^k) \sqrt{\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})} \quad (8)$$


as the lowest QoS rate, i.e., the statistical minimum rate that needs to be greater or equal to m_i^k for the QoS requirement (6) to be satisfied. Furthermore, if the following problem is feasible, then the condition (7) can be satisfied with an extra margin $s \geq 0$ for all i and k ,

$$\begin{aligned} & \text{maximize} && s \\ & \boldsymbol{\alpha} \in \mathcal{A}, s \geq 0 \\ & \text{subject to} && \frac{\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) - m_i^k - s}{\sqrt{\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})}} \geq \Phi^{-1}(\epsilon_i^k) \quad \forall i, k \end{aligned} \quad (9)$$

where the set $\mathcal{A} \triangleq \{\boldsymbol{\alpha} : \alpha_{ij}^k \in [0, 1], \sum_{i,k} \alpha_{ij}^k \leq 1, \sum_{j,k} \alpha_{ij}^k \leq 1\}$ is defined so that the constraint $\boldsymbol{\alpha} \in \mathcal{A}$ makes α_{ij}^k a fraction, and that these fractions can be scheduled without overlapping so that the confluent wireless channels are not overbooked. Solving (9) defines the optimal routing scheme $\boldsymbol{\alpha}$ satisfying the QoS requirements. Notice that the positions of the UAVs are parameters for this problem so that they are not optimized by solving (9). Controlling the network agent trajectories is the goal of the ensuing subsection. Problem (9) aims to produce routing tables that maximize the rates across the network. It admits an equivalent reformulation as a convex second-order cone program (SOCP), which ensures that the global optimal solution $\boldsymbol{\alpha}^*$ and s^* can be solved efficiently in polynomial time. Appendix A includes the details on how to construct this SOCP from (9).

Implementing the optimal routing strategy $\boldsymbol{\alpha}^*$ amounts to agent i transmitting data corresponding to the k -th flow to each agent j with probability $(\alpha_{ij}^k)^*$, effectively a fraction of the time. The actual rates achieved when routing according to $\boldsymbol{\alpha}^*$ are $m_i^k + s^*$, which means that this routing strategy guarantees that all rate requirements m_i^k are satisfied and exceeded with probability ϵ_i^k , and that the surplus of network capabilities is shared equally among agents and flows offering each one an excess of rate equal to s^* .

The rationale behind (9) is that the rate surplus s^* is maximized by two complementary means. One possibility is to increase the mean rate $\bar{b}_i^k(\boldsymbol{\alpha}^*, \mathbf{x})$ by using the links with higher channel gains $\bar{R}_{ij}(x_i, x_j)$ more frequently, which amounts to giving higher values to fractions α_{ij}^k corresponding to those links, cf. (4). An alternative strategy to increase the surplus is to reduce the variance $\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})$ by selecting the channels with lower variance $\tilde{R}_{ij}(x_i, x_j)$ and splitting data over multiple nodes, cf. (5). Which of these two options is the best one depends on the QoS requirements in terms of rate m_i^k and confidence ϵ_i^k , and it is resolved numerically by solving the optimization problem (9).



(a) Low rate margin with high confidence ($m_i^k = 0.2, \epsilon_i^k = 0.95$). (b) High rate margin with low confidence ($m_i^k = 0.5, \epsilon_i^k = 0.7$).

Fig. 4: Illustrative example of the robust routing solution obtained at the source by solving (9). A low probability of error requires less variance on the transmission rate, which is achieved by splitting data between the direct link and the relays. On the other hand, a high margin requires all the data to be transmitted to the relays, which present the higher channel gain.

Fig. 4 aims to give intuition on the dependence between QoS parameters and routing options in a simple network with a source node (colored red and marked with an S), which sends data to a destination node (colored red and marked with a D), with the option of using their direct link or routing through two relay nodes (colored blue). When high confidence $\epsilon_i^k = 0.95$ is required, the variance needs to be brought down, so that the optimal strategy is to split data between the two possible routes. Indeed, given its quadratic form in (5), the variance is reduced by dividing the routing variables into smaller fractions. On the other hand, when the confidence margin is relaxed to $\epsilon_i^k = 0.7$, but the rate requirement m_i^k increases, then the optimal strategy is to route all the data through the relays, which are closer to both the source and the destination, and thus present higher channel gains (cf. Fig. 3). These two examples also illustrate the robustness to noise of the optimal routing strategy, which takes into account the variance of the channel gains to ensure that the rate requirements are satisfied within the specified confidence margins. It balances mean and variance so that the channels with higher average gains are sacrificed if they cannot guarantee that the QoS requirements (m_i^k, ϵ_i^k) are satisfied.

C. Optimal spatial deployment

Solving the robust routing problem (9) provides us with optimal routing decisions given the positions of the agents. One exceptional aspect of the ad-hoc network of communication devices that we propose is that they are mobile. Hence, the position of network agents can be actively controlled with the goal of increasing data rates. To this end, we consider a network optimization problem over a communication graph. Consider the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of L nodes in the network and $\mathcal{E} \subseteq \mathcal{L} \times \mathcal{L}$ is the set of communication links. The graph is paired with an adjacency matrix $\mathbf{A} \in \mathbb{R}^{L \times L}$ that indicate how strong each link is, using the rate model previously introduce in (1), we can define the following entries of the adjacency matrix by

$$A_{ij} \triangleq \bar{R}_{ij}(x_i, x_j) = \text{erf} \left(\sqrt{\frac{P_{L0}}{P_{N0}} \|x_i - x_j\|^{-n}} \right). \quad (10)$$

The goal of the following optimal controller is to maximize the strength of all links across the network. To do so, we obtain the Laplacian of the graph, given by

$$\mathbf{L} = \text{diag}(\mathbf{A}\mathbf{1}) - \mathbf{A}. \quad (11)$$

In Appendix B we illustrate how strengthening the connectivity of the graph is equivalent to maximizing the second eigenvalue of the Laplacian [35], also known as *Fiedler value*. Before proceeding to maximize this eigenvalue, we consider a linear Taylor approximation \hat{A}_{ij} of the data rates A_{ij} around the current positions $\hat{\mathbf{x}}$ of the agents, which will induce convexity of the ensuing optimization problem. This Taylor approximation is given by

$$\hat{A}_{ij} = \hat{R}_{ij} + \nabla_{x_i} \hat{R}_{ij}^T (x_i - \hat{x}_i) + \nabla_{x_j} \hat{R}_{ij}^T (x_j - \hat{x}_j) \quad (12)$$

where we have defined $\hat{R}_{ij} \triangleq \bar{R}_{ij}(\hat{x}_i, \hat{x}_j)$, i.e., the evaluation of the mean rate at the current agent positions. Now we are ready to maximize the Fiedler value, denoted by γ .

$$\begin{aligned} & \underset{\mathbf{x}, \gamma}{\text{maximize}} && \gamma \\ & \text{subject to} && \mathbf{P}^T \hat{\mathbf{L}} \mathbf{P} \succeq \mathbf{I} \gamma \\ & && \hat{\mathbf{L}} = \text{diag}(\hat{\mathbf{A}}\mathbf{1}) - \hat{\mathbf{A}} \\ & && \hat{A}_{ij} = \hat{R}_{ij} + \nabla_{x_i} \hat{R}_{ij}^T (x_i - \hat{x}_i) + \nabla_{x_j} \hat{R}_{ij}^T (x_j - \hat{x}_j) \\ & && \|\mathbf{x} - \hat{\mathbf{x}}\|_{\infty} \leq \Delta \end{aligned} \quad (13)$$

The constraint $\mathbf{P}^T \hat{\mathbf{L}} \mathbf{P} \succeq \mathbf{I} \gamma$, ensures that γ is indeed the second eigenvalue of the Laplacian [18]. For this purpose, the columns of the constant matrix \mathbf{P} must be selected to form a basis of $\mathbf{1}^{\perp}$, that is the complement of the vector of all ones. In addition, the constraint $\|\mathbf{x} - \hat{\mathbf{x}}\|_{\infty} \leq \Delta$ is set so that the Taylor approximation remains valid. Appendix B shows that the optimization problem (13) is convex and it can be cast as a semidefinite program (SDP).

All in all, the solution of (13) yields a new set of positions for the network agents that increases the strength of the communication network. An agent designated as the connectivity planner must solve (13) repeatedly, and publish the solution to the network agents, so that they travel across space tracking the sequence of optimal locations.

D. Network and connectivity planner architecture

In order to implement the high-level control of the mobile wireless infrastructure on demand a network agent is designated as the planning agent, wherein two planning functionalities coexist (Fig. 5). The network planner implements the optimal probabilistic routing (Section II-B) and the connectivity planner implements the optimal spatial deployment (Section II-C). These planning functionalities are supported by a rate estimation module (Section II-A).

First, the positions are collected for all network and task agents, i.e. $\{x_i\}_{i=1}^N$ and $\{x_j\}_{i=1}^M$. These are fed to the Rate Estimation Unit, which computes the mean and variance of channel rates between each pair of agents according to the rate model (1) and (2). These means and variances are inputs for the Network Planner. The Network Planner also takes the QoS requirements (m_i^k, ϵ_i^k) as prescribed configurable

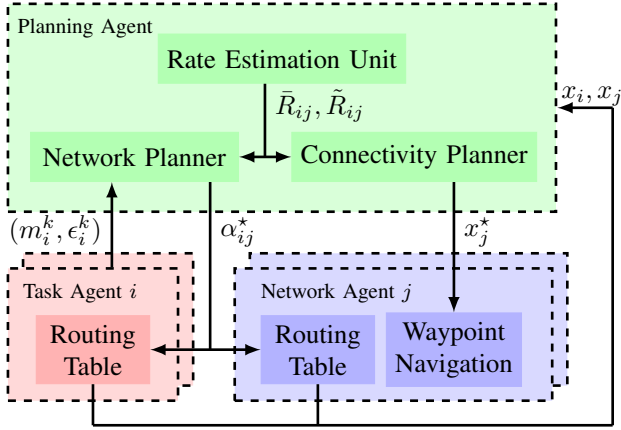


Fig. 5: The high-level control loop uses the current positions of network and task agents to estimate the channel rates. Then the network planner selects the optimal routing strategy according to these rates, and the connectivity planner provides trajectories of the network agents to ensure a cohesive communication network.

parameters, and uses them to specify problem (9) with variables α and s . In particular, α enters the constraint in (9) via (4) and (5). The equivalent convex SOCP reformulation detailed in Appendix A is solved and the resulting α^* are shared with all network and task agents. Then each i -th agent adjusts its routing table to send data to its j -th pair during a fraction α_{ij}^k of the time interval between two consecutive executions of the SOCP solver. Notice that task agents have to be capable of processing these parameters and adapting their routing tables accordingly.

In parallel, the Connectivity Planner determines the next position of each agent in the network team. From the Rate Estimation Unit, it receives the Taylor approximation of the data rates and uses them together with the set of all current agent positions as parameters for solving the equivalent SDP formulation of (13) shown in Appendix B. The solution $\{x_j^*\}_{j=1}^M$ is published to all network agents, which use them as waypoints to which to navigate. Next, the actual positions are measured and communicated to the planning agent starting a new iteration of the loop in Fig 5. A pseudo-code for the overall planning mechanism is given in Algorithm 1.

III. IMPLEMENTATION AND EXPERIMENTAL PLATFORM

A main goal of this work is to develop a *system* for mobile wireless infrastructure on demand. To this end, the functionalities presented in the previous section were implemented in ROS and an experimental UAV platform was built to provide each agent with sufficient onboard computational capabilities to solve the optimal network and connectivity planning problems. As described in the previous section, the agent acting as a planning agent must solve (13) in real time to command the MID team to a new set of points in space. It must also solve (9) in order to choose the routing strategy dynamically. Both (9) and (13) are recast as standard convex programs, and then solved by the convex optimization library `cvxpy` for Python. The code used to solve (9) and (13) using `cvxpy` is wrapped in a planner ROS node that publishes the optimal locations \mathbf{x} and routes α sequentially.

Algorithm 1 Network and Connectivity Planner

```

Task agent demands QoS  $(m_i^k, \epsilon_i^k)$ 
while active do
  Collect  $\{x_i\}_{i=1}^N, \{x_j\}_{j=1}^M$  from task and network agents
  Rate Estimation Unit
  Compute  $\bar{R}_{ij}, \tilde{R}_{ij}, \hat{R}_{ij}, \nabla_{x_i} \hat{R}_{ij}, \nabla_{x_j} \hat{R}_{ij}$ 
  Network Planner
  Solve SOCP reformulation of (9) as Algorithm 2
  Distribute solution  $\alpha^*$  to all agents
  Connectivity Planner
  Solve SDP reformulation of (13) as Algorithm 3
  Distribute solution  $\{x_j^*\}_{j=1}^M$  to all network agents
end while

```

A. Hardware components

For this to be possible, the hardware to be used in experiments must have sufficient compute to be able to support real-time execution of `cvxpy` over ROS. Most commercial UAV platforms only possess an onboard flight controller with insufficient computation power for our needs. In the absence of an off-the-shelf commercial UAV with these capabilities, we built an experimental quadrotor with both a flight controller and a relatively powerful onboard computer. The choice of a quadrotor, as opposed to other forms of UAVs, has been guided by their availability, flexibility, and simplicity. The main hardware components of this prototype are highlighted in Fig. 6.

A small Intel NUC board is used for this purpose. An Intel 2.4 GHz i3-7100U CPU is mounted on a NUC7i3DNBE motherboard, which is then wired via USB to a Pixhawk Pix 2.4.8 flight controller running the PX4 flight stack. The flight controller receives positioning directives from the NUC via the MAVLink protocol over a USB connection. This low-level flight controller is commanded by MAVROS. The Pixhawk is also provided with signals from a u-blox NEO-M8 GPS receiver, which is taken as a global input to define the position of each agent relative to the common reference of the team. In addition, a downward-facing Garmin LIDAR-Lite v3 is used to provide Pixhawk with a complementary measure of altitude, assisting the PX4 controller at take-off and landing. The two processors and peripherals are mounted on a standard F450 quadcopter frame along with four 920KV brushless motors attached to four 1045 propellers, each one connected to the Pixhawk via a dedicated 30A ESC controller. The F450 structure together with these components is lifted with the power of a lithium polymer battery of four cells and 5200 mAh, which feeds power to the rotors directly, and both to the Pixhawk controller and the onboard computer motherboard, through voltage regulators.

While the UAVs are designed to fly autonomously, each unit can be connected wirelessly to a manual remote controller (RC) on the ground for safety purposes. For this purpose, a R9DS receiver is wired to the Pixhawk controller onboard and communicates with the RC using spread-spectrum in the band of 2.4 GHz. Likewise, each UAV has the capability to take flight directives and report flight statistics to a computer on

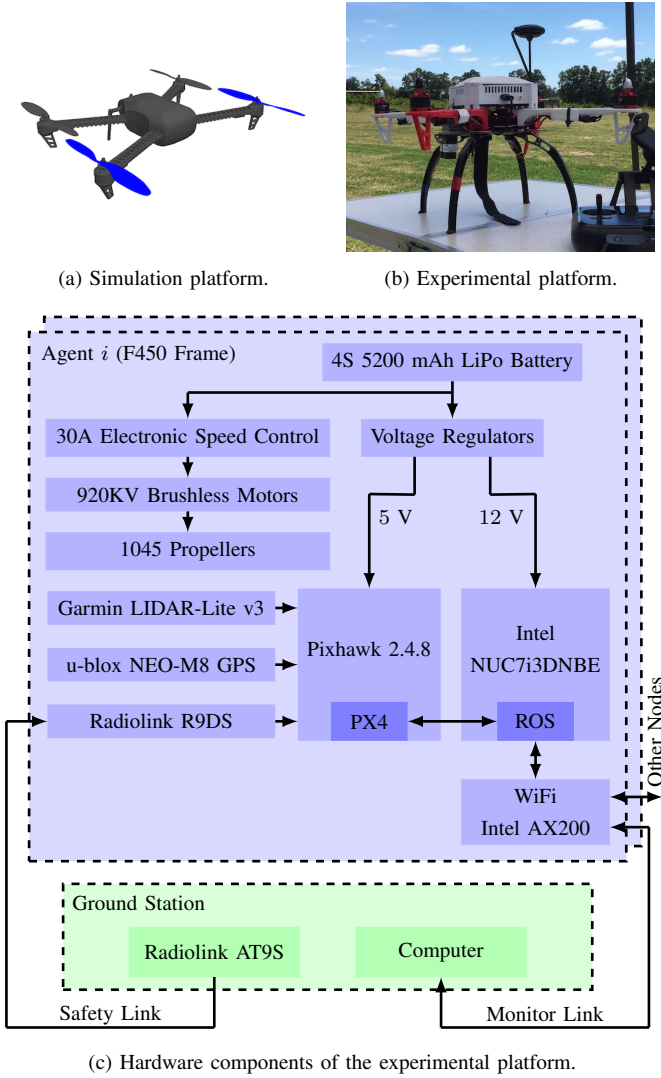


Fig. 6: Platforms used for simulation and for experimental deployments. A UAV model in Gazebo is used in simulation, while a custom-built UAV platform is used in experiments. The experimental UAV has a low-level flight controller and a high-level processor to optimize the trajectories and routing.

the ground via Wi-Fi. This data can be exchanged with the onboard computer as ROS topics or parameters, or bypassed to the Pixhawk controller via the same MAVLink connection described before.

B. Networking, routing and coordination

Each UAV agent is assigned a unique IP address within the communication network. The high-level communication among UAVs in the swarm is achieved via IP over an ad-hoc wireless network, which is composed of links between Intel AX200 transceiver cards. These cards are mounted on each onboard computer motherboard, and implement the IEEE 802.11ax wireless standard, radiating power through a pair of dual-band 2.4 – 5.0 GHz patch Molex 6E antennas. The AX200 cards are configured in IBBS ad-hoc mode operating at 2.4 GHz, which enables a decentralized network as the one described in Section II.

The solution of the robust routing network planner is a probabilistic choice of route. In effect, it indicates the probability

with which an agent should use a specific communication link. Equivalently, this is the fraction of time that a link should be used between planning loops. Effectively, the optimal routing strategy obtained by solving (9) is implemented by subsequent modification of the routing tables in the Ubuntu 18.04 (Bionic Beaver) operating system running onboard the NUC. These routing table modifications are effected by the `ip`, `ifconfig` and `route` commands available in the standard set of network tools available in unix-based kernels. An optimal fraction $\alpha_{ij}^k \leq 1$ resulting from solving the optimization procedure is effected in practice by a sequence of probabilistic updates. Specifically, every half a second a Bernoulli variable with probability α_{ij}^k is flipped for each i , j , and k , and if the result is positive the IP of agent j is set as an entry of the routing table of agent i .

A key step towards a decentralized multi-agent implementation is to have ROS running in multi-master mode. Setting up this mode also prevents ROS from flooding the wireless network by keeping the ROS topics communicated among agents to a minimum. With each UAV implementing a ROS master, only the topics that report positions and routing fractions are shared among UAVs. The complementary ones needed for each master to run are communicated to the internal localhost IP and do not reach the wireless network, leaving more room for user data. This is an issue of practical nature in the ROS framework. In the future, building a MID system on top of the ROS2 framework, with its Data Distribution Service implementation, would resolve this issue.

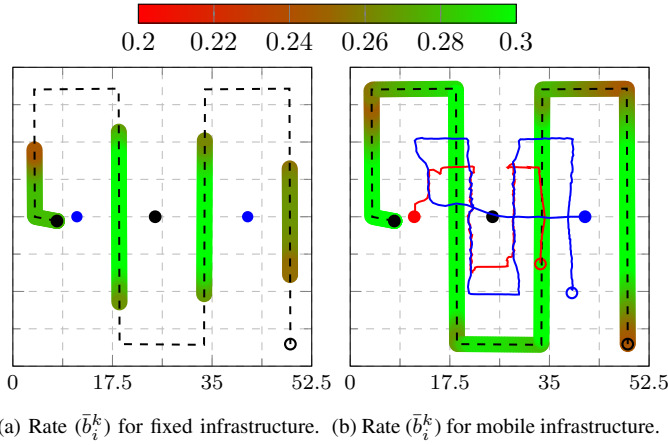
Although in our current configuration, both the network planner and connectivity planner run on a designated UAV, and their outputs are broadcast to the other agents via Wi-Fi, all UAVs are built with the same computational capabilities, so that they could run local versions of these planners, which would produce the same outputs, but in a decentralized manner. In both cases, for these mechanisms to work, agents need to be in a common reference system, this is a practical issue that is addressed in Appendix C.

IV. SIMULATIONS


We proceed now to perform a verification of the proposed mobile infrastructure method. To this end, we conduct numerical simulations of the system described in Section II. The system is implemented in ROS [36] and simulations are performed with the aid of Gazebo [37]. The hardware used for simulations is comparable to the compute assembled on the experimental platforms. Network connectivity is simulated according to the channel model shown in Fig. 3 and described by equations (1)–(2). We consider three different scenarios, each one highlighting a different feature of the mobile wireless infrastructure on demand system.

- (F1) Extension of operational range.
- (F2) Support for tasks with complex mobility patterns.
- (F3) Robust autonomous reconfiguration.

All scenarios consider the following premise, a set of task agents is deployed in the environment, performing a specific task. These agents are supported by the MID system. The mobile infrastructure team executes our proposed algorithms



(a) Rate (\bar{b}_i^k) for fixed infrastructure. (b) Rate (\bar{b}_i^k) for mobile infrastructure.



(c) Lowest QoS rate (q_i^k) with respect to the distance between task agents.

Fig. 7: Extending operational range. A task agent performs a square-wave trajectory covering an area of around $50 \text{ m} \times 50 \text{ m}$. This agent must relay information back to another stationary task agent placed in the center of the map. A bidirectional flow of traffic with QoS of $(m_i^k, \epsilon_i^k) = (0.2, 0.7)$ is requested. Black dots represent task agents, while red and blue dots represent network agents. Starting positions are shown by a solid dot, while hollow dots correspond to end positions. The lines connecting these dots show the trajectories taken by the agents during the simulation. The rate attained by the task agents experiences a sudden drop when they are 20 m apart and using fixed infrastructure. In contrast, there is a smooth decay when using the MID system, in which case the range extends up to 40 m.

and enables the successful accomplishment of the task agents' mission by supporting their network connectivity. For comparison purposes, we also simulate the deployment of static network teams, whose placement is predefined. The fixed system uses the optimal robust routing as described in Section II, but does not reposition itself. This comparison baseline corresponds to the usual solution of spreading relays in the environment, which has to be done beforehand and are not repositioned dynamically.

A. Extending operational range

The first evaluation scenario that we consider is that of operational range extension. The goal here is to demonstrate the capability of the mobile infrastructure to reposition itself and thus cover more area than a predefined fixed deployment. By these means, the MID system can provide a larger operative range with the same number of network agents when compared to a fixed network deployment. Practical applications of this feature match area coverage tasks or data collection situations, in which data must be sent back to a center for post-processing. A clear example would be something akin to agricultural

monitoring in which an agent (or a group of agents) moves in a predetermined path through a field, acquiring measurements that must be relayed back in real-time to a coordination center.

For purposes of this simulation, let us consider the scenario illustrated in Fig. 7. In this example, a task agent performs a square-wave trajectory covering the $50 \text{ m} \times 50 \text{ m}$ area. Another task agent is placed in the center of the environment and is in charge of collecting information. A bidirectional flow of traffic with QoS given by $m_i^k = 0.2$ with confidence $\epsilon_i^k = 0.7$ is requested between them. In the case of the fixed configuration, two network agents are deployed in a reasonably equidistant manner over the x-axis. However, the need of covering a square with only two agents clearly leaves the perpendicular y-axis of the map with lower overall coverage, resulting in severe outages along the edges, see Fig. 7(a). The same deployment is used in the case of the mobile infrastructure. In this case, the mobile system is capable of repositioning itself to support the required communication rate along the entirety of the map, cf. Fig. 7(b).

This behavior is also certified by looking at the supported rate with respect to the distance between the moving task agent and the receiving task agent (node in the center). This is shown in Fig. 7(c). The fixed deployment has a clear cutoff at around 20 m, where the rate decreases significantly and its variance increases. The latter is due to the fact that moving the same distance along the y-axis results in worse rates than moving the same distance along the x-axis, as the relays are spaced on the x-axis. In contrast, the mobile infrastructure deployment shows a smooth decrease in rate with respect to the distance between agents, and it always outperforms the fixed infrastructure at the same distance. Specifically, when the moving task agent relies on our MID system, in contrast to the fixed setup, it can double the travel distance from 20 m to 40 m while keeping communication to the base at a rate over the required minimum of $m_i^k = 0.2$. Hence, the mobile infrastructure effectively extends the functional range of the system when operating with the same number of agents.

B. Supporting complex task behaviors

In most cases, the behavior of task agents is either unpredictable or follows a complex trajectory. Under these circumstances, deploying fixed infrastructure at optimal locations is not a trivial endeavor. This fact either leads to denser infrastructure deployments than needed, or QoS not being satisfied due to a lack of network agents. Next, we study one such scenario. We consider three task agents conducting a patrol following the four-leafed clover trajectory shown in Fig. 8. Two network agents are employed to support a circular flow of traffic among the task agents. In this case, the node placement of the fixed infrastructure is not trivial, as task agents might go in and out of sync of being close or far from the center of the trajectory. Thus, we place the fixed infrastructure across the x-axis. The exact placement of the two network agents for the fixed infrastructure is shown by the solid red and blue dots in Fig. 8(a).

The mean rate \bar{b}_i^k and lower QoS rate q_i^k for agent $i = 1$ and flow $k = 1$ are shown in Fig. 8(d). The red and blue

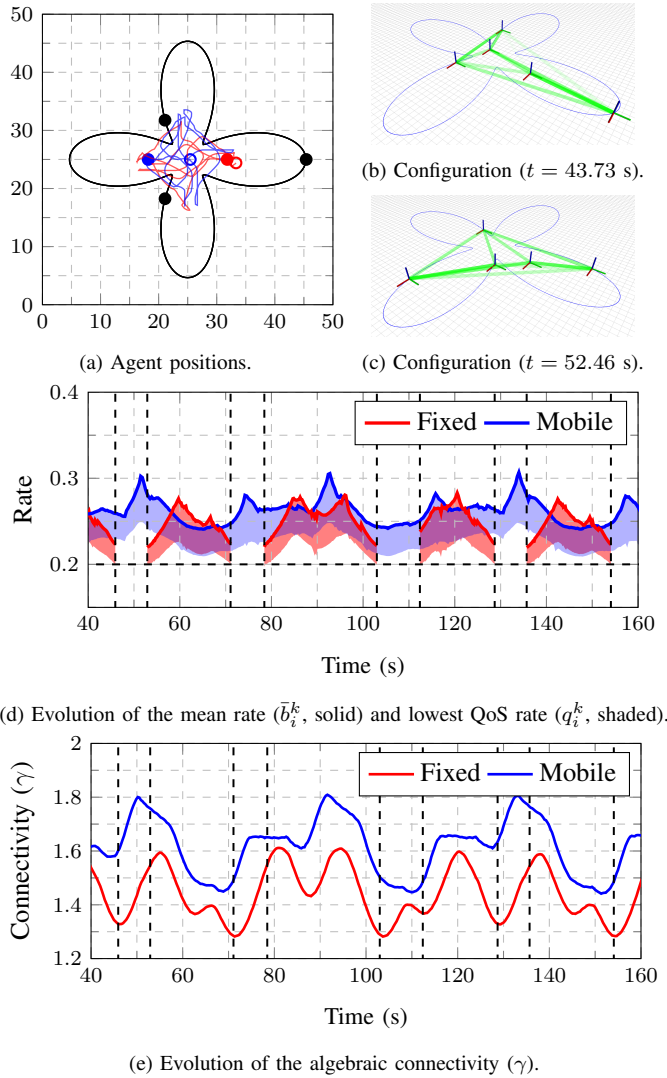


Fig. 8: Supporting complex task behaviors. Three task agents perform the patrol shown by the black line. A QoS of $(m_i^k, e_i^k) = (0.2, 0.9)$ for a circular flow of traffic among task agents is requested. Black dots represent task agents, while red and blue dots represent network agents. Starting positions are shown by a solid dot, while hollow dots correspond to end positions. The lines connecting these dots show the trajectories taken by the agents during the simulation. The probability of routing information among agents is shown by the opacity of the connecting green lines, with more probable routing directions shown as more opaque. In contrast to fixed infrastructure, which suffers from communication outages 28.19% of the time, the MID system guarantees that the QoS requirements are satisfied 100% of the time.

shaded bands correspond to rates obtained using the fixed infrastructure and MID system, respectively. They extend between the mean rate $\bar{b}_i^k(t)$ (solid line) and the lowest QoS rate $q_i^k(t)$ (lowest shaded region). This figure illustrates that while fixed infrastructure suffers from frequent communication outages, mobile infrastructure is capable of reconfiguring dynamically and supporting the required rate in a relatively stable manner. Due to the use of only two network agents, fixed infrastructure is unable to fully cover the trajectory of the task agents. Specifically, the fixed setup cannot hold the required QoS during recurrent time intervals, corresponding to communication outages 28.19% of the total time span of

the simulation. The naive fixed deployment used spreads the network agents over the x-axis, thus when task agents move far in the y-axis, communication outages occur. In contrast, the capability of the mobile infrastructure to dynamically adapt means that two network agents can efficiently cover the whole trajectory. As seen in Fig. 8(a), the overall trajectories of the network agents stay relatively close in range to the naive fixed infrastructure locations. However, they rotate and adapt to track the movement of the task agents.

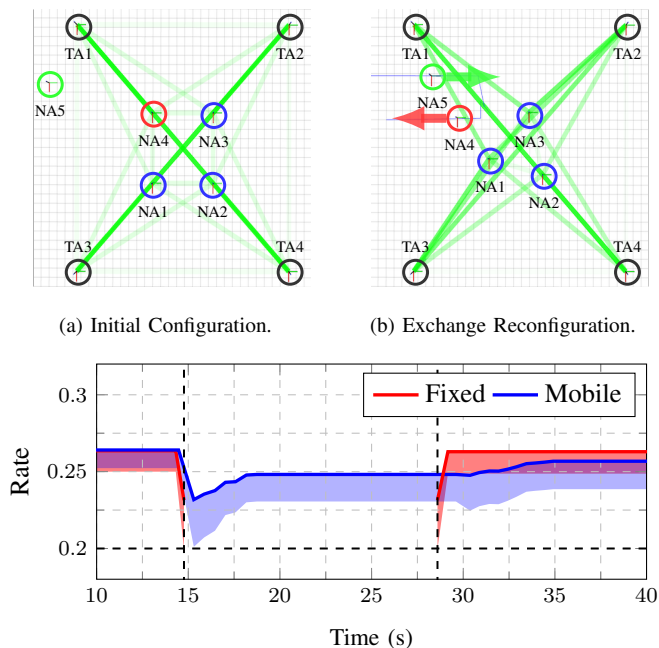
The movement of the network agents is better analyzed by snapshots of their configuration shown in Fig. 8(b)–8(c). The mobile infrastructure places network agents inside the triangle connecting the task agents, attempting to stay relatively close to the center of it. When one of the task agents moves far away in one axis while the other two task agents remain close (e.g., Fig. 8(b)), the network agents elongate their positions maintaining the quality of service required. In this case, fixed infrastructure is incapable of maintaining connectivity and produces a communication outage (see Fig. 8(d), $t = 43.73$ s). Furthermore, routing decisions are also affected by the agents' positions. For instance, at $t = 52.46$ s (Fig. 8(c)), the task agents are at a point in their trajectories such that they are closer to each other, this results in routing decisions that are more spread over with some probability of traffic going directly from task agent to task agent. With the same number of network agents, in contrast to fixed infrastructure which suffers from communication outages 28.19% of the time, the MID system guarantees that the QoS requirements are satisfied 100% of the time. Thus, the MID system adapts to support task agents with complex mobility patterns where fixed infrastructure fails.

Another interesting aspect to study is the relationship between the resulting rate, Fig. 8(d), and the algebraic connectivity, Fig. 8(e). Recall that the measure of true interest is rate and we use algebraic connectivity as a surrogate to design the connectivity planner (13). Fig. 8(e) shows that, in effect, algebraic connectivity serves as a good proxy measure for the resulting rate of the system.

We have further used this simulation scenario to obtain estimates of the time required to execute the network and connectivity planning loop (Algorithm 1). Independent simulations with the same task agent trajectories and QoS demands and network agents ranging from 2 to 14 nodes show average re-planning times of 0.45 s to 5.06 s with standard deviations of 0.019 s and 0.124 s, respectively.

C. Robust autonomous reconfiguration

We now consider another simulation case wherein one network agent must be replaced by another. The motivation underlying this scenario is practical. In many cases, fixed infrastructure is deployed to cover a large dense area. Some of the fixed nodes in this infrastructure might need to be replaced due to malfunction, battery limitations, or simply due to the need for continuous operation or network design. Specifically, for purposes of the simulation, in this scenario, task agents are simply stationary. We set a single network agent to be replaced by a different agent as shown in Fig. 9. This simulation



(c) Evolution of the mean rate (\bar{b}_i^k , solid) and lowest QoS rate (q_i^k , shaded).

Fig. 9: Robust autonomous reconfiguration. Task agents are the outside square of four agents. The network agents are the inner set of four agents plus the substitute. A flow of traffic with QoS of $(m_i^k, c_i^k) = (0.2, 0.7)$ is requested from node 2 to node 4 and the same requirements for a flow from node 3 to node 1. The MID system prevents a connectivity outage of 12 s that would result from using a static setup in the event that a networking node is deactivated and needs to be replaced.

consists of an outer square of task agents, requesting a flow of traffic among some of them. An inner square of network agents is deployed to support these communication needs. Under this initial configuration, each task agent mainly routes information to its nearest network agent, with alternative routes occurring with low probability.

At a specific point during the simulation, the top left network agent depicted in red deactivates itself from the network configuration and needs to be replaced by another network agent. During this replacement interval, the agents in the mobile infrastructure reconfigure themselves to provide the required connectivity while the new agent arrives at the formation. Figures 9(a)–9(b) shows the clear change in the agents' positions. Network agents move slightly to cover the region of the network agent being replaced. The routing decisions are also altered. Information flow is more evenly spread through the network, in contrast to it being routed with high probability to the closest network agent. As the replacement network agent enters the network, the nodes begin to reposition themselves again, reaching the same initial configuration with four agents as shown in Fig. 9(a). The temporary reconfiguration is necessary to maintain the quality of service requested by the task agents. In contrast to this adaptive solution, removing a single agent from the fixed infrastructure might result in the network being unable to support the required rates. This is the case of this simulation, as it is apparent from the mismatching data rates shown in Fig. 9(c).

In the static case, when the network agent deactivates itself to be substituted and goes out of the system, the lowest QoS rate (q_i^k , the lower limit of the shaded region) drops under the required threshold, resulting in a connectivity outage. This outage lasts for the complete duration of the replacement procedure (15 s to 27 s in Fig. 9(c)). In contrast, the mobile infrastructure immediately reconfigures itself, recovering the rate and managing to maintain connectivity during the network agent exchange process. When the replacement agent arrives, the network configuration goes back to the same initial configuration shown in Fig. 9(a).

This illustrates the ability of the mobile infrastructure system to provide more than simply extended coverage. Since a fixed infrastructure design needs to place nodes in specific predetermined positions, a failure of a single node can be catastrophic to overall network connectivity unless we explicitly introduce redundancy into the system, with the corresponding increase in the number of required nodes. In contrast, the mobile infrastructure system is robust to system failure by providing a capable reconfiguration mechanism to maintain the requested level of connectivity.

An important observation common to the three scenarios presented in this section is the fact that even if these are simulated experiments, they run in real-time on similar hardware to the real-life experimental platform. A consequence of this is that there is a small delay between reconfigurations, as the agents need to solve the optimization problems discussed in Section II. Specifically, for the simulation in this section, the replanning process runs at around 2.85 Hz. This publishing rate could be increased with better computing power, but the current case is representative of what will be observed in experimental deployments. Nonetheless, while faster computers (or avoiding real-time operation) would lead to better rate results due to more adaptive replanning, mobile infrastructure outperforms fixed infrastructure even with the current conditions.

In this section, we have presented three illustrative instances that showcase the capabilities of the MID system. Firstly, we have successfully expanded the operational range, feature **(F1)**, effectively doubling it for an area coverage task. Secondly, in a task involving intricate mobility patterns, feature **(F2)**, we have effectively mitigated communication outages from 28.19% to 0%, entirely eliminating disruptions in a complex patrol scenario. Thirdly, our last simulation reveals the inherent robustness of the MID system, feature **(F3)**. In a scenario where an agent failure occurs, the MID system instantaneously reconfigures itself to seamlessly support network requirements, avoiding a 12 second outage, thereby completely avoiding any task disruptions.

V. EXPERIMENTS

Complementarily to the simulations in the previous section, we performed experimental demonstrations. These experiments are conducted using the hardware platforms described in Section III. These UAVs implement the required algorithms and networking capabilities to test our proposed MID system. We configure the UAVs' wireless cards to operate in ad-hoc




Fig. 10: Top-down view of the location used for the experimental sessions, where each square is approximately $50 \text{ m} \times 50 \text{ m}$. A snapshot of the range extension experiment being conducted is shown on the right. Task agent 10.42.0.2 progressively moves away from the stationary task agent 10.42.0.1, relaying information back to it. A mobile infrastructure agent 10.42.0.3 redeploys itself dynamically to support this experiment.


mode. Hence, the nodes themselves create their own mesh network, which they use for communication. To test the network capabilities we use the `iperf3` and `traceroute` tools of the Linux system to measure both the actual throughput and delay. Effectively, these tools do the following: `iperf3` injects TCP/IP or UDP packets into the network interface of choice and with a given destination, allowing for the measurement of communication throughput; on the other hand, `traceroute` provides the round trip time of packets between two nodes, to measure the delay experienced in the network.

The location used for conducting the experiments is on the outskirts of Montevideo, Uruguay. This area, shown in Fig. 10(a), spans several football and agricultural fields. Due to its location, minimal wireless interference exists in this area. Similarly to the simulations in the previous section, we conduct three experiments, each highlighting a feature (F1)–(F3) of the MID system.


A. Extending operational range

In the first experiment, the MID system extends the operational range of the agents, similar to the previously performed simulation in Section IV-A. The node placement for this experimental test is shown in Fig. 10(b). We label nodes by their IP addresses in the ad-hoc network. In this experiment, a stationary task agent (10.42.0.1) is placed at one end of the experiment area. This node takes off but does not move, maintaining a hovering state. Another task agent (10.42.0.2) takes off and moves away from the first agent in a straight line. A network agent is placed in the middle (10.42.0.3) which effectively acts as a relay between the two task agents. We conduct two distinct tests, one with fixed infrastructure placed at 20 m away from the stationary task agent, in which the network agent takes off but does not reposition itself, simply executing the solution to the robust routing problem. Then, we conduct another test using the mobile infrastructure, in which the network agent repositions itself dynamically using the MID system.

The results from this experiment are shown in Fig. 11. We measure the communication rate between the task agents and plot it against their distance from each other (Fig. 11(c)). The capability of the MID system to extend the operational range is apparent from Fig. 11(c), since the rate at any given distance is always lower when using the fixed infrastructure.



(a) Fixed infrastructure ($d = 178.23 \text{ m}$). (b) Mobile infrastructure ($d = 178.17 \text{ m}$).




(c) Rate with respect to the distance between task agents. Individual rate measurements and a moving average over 25 meters are shown.

Fig. 11: Extending operational range. The MID system operates above 2.5 Mbps at distances of up to 220 m, while the fixed infrastructure can only support this data rate for up to 160 m. While the fixed infrastructure is limited by its predetermined placement, the mobile infrastructure places itself approximately in the middle of the two task agents, supporting longer ranges and overall higher rates than the fixed infrastructure.

Due to the limited wireless interference at the location of these experiments, the antenna directivity pattern and the complete line of sight between agents, very long operational ranges can be achieved. Nevertheless, the mobile infrastructure is shown to be capable of operating at around 2.5 Mbps at distances of up to 220 m, while the fixed infrastructure can only support this data rate at distances of up to 160 m.

The optimal placement of the network node at such long distances is intuitive, one should place it in the middle between the task agents. We verify this in Fig. 11(a)–11(b). While the relay is incapable of moving when the infrastructure is fixed, the network agent in the mobile system always keeps itself near the center point between the two task agents. This enables the extended operational range. Since the data rate decreases with respect to distance (See e.g., Fig. 3 and equation (1), for the channel model), this results in overall better data rates.

This experiment verifies that mobile infrastructure is capable of making more efficient use of network agents. Indeed, by dynamically routing and repositioning, further operational ranges can be achieved without the need of additional agents. The MID system also proved itself robust to model inaccuracies. The parameters fitted in Fig. 3 resulted to be too conservative in the experiment site probably due to low noise levels. However, problem (13) remains feasible so that the connectivity planner will always work under model mismatch, and the network planner can protect itself by lowering the

(a) Configuration ($t = 5$ s).(b) Configuration ($t = 145$ s).

(c) Trajectories. Grid space is 10 m.

(d) Executed routing ($t = 145$ s).


Fig. 12: Spatial configurations and routing solutions of the mobile infrastructure during the patrol experiment. Three task agents rotate in a counter-clockwise circle of 40 m diameter. Three network agents are deployed to provide the required connectivity.

parameter m_i^k if needed, as the border case $m_i^k = 0$ ensures the feasibility of (9).


B. Supporting complex task behaviors

We proceed with a second experiment consisting of a circular patrol by three task agents. This experiment emulates the complex behavior simulation discussed in Section IV-B. For practical reasons, in the following two experiments, we opted to position the antennas on the arms of the UAV frame as opposed to the legs. This deliberate placement significantly diminishes the antenna gain within the plane parallel to the ground. While this outcome would typically be considered undesirable, it enables us to conduct multi-agent experiments within a more confined space. In this experiment, three task agents (10.42.0.1, 10.42.0.2, and 10.42.0.3) and three network agents (10.42.0.4, 10.42.0.5, 10.42.0.6) are deployed utilizing the MID system. As the experiment initiates, the task agents move into a circular patrol of 40 m diameter, rotating in a counter-clockwise manner.

The resulting spatial and routing configurations can be seen in Fig. 12. As seen in Fig. 12(a)–12(b), compared to the starting configuration, once the task agents pull away, the routing probabilities favor a flow of information towards the network agents, avoiding direct communication between the task agents. The trajectories taken during the experiment are shown in Fig. 12(c). Overall, the network agents assume a somewhat triangular inner circle formation, tracking the rotation of the task agents. At any given instance, the execution of the probabilistic robust routing solution requires a sampling of probabilities. Though occurring in the same way, this is



(a) Average data rate in the network.



(b) Average delay in the network.

Fig. 13: Patrol experiment. Instantaneous rate and delay measurements are collected every second and a moving average over 20 seconds is shown. The average data rate and delay achieved by the MID system stabilizes after 60 s and is kept around 5 Mbps and below 115 ms.

something that we have not explicitly shown in our previous experiments. Fig. 12(d) shows the exact sampling of the routing configuration Fig. 12(b) at a specific time instance ($t = 145$ s), which is then employed at the communication interface. As previously discussed, the probability of direct communication between task agents is very low, and thus traffic moves through the inner network agents before reaching the destination task agents.

The average data rate across the network observed during the experiment is shown in Fig. 13(a). As the experiment begins, the task agents move away to their specified patrol radius. As they do this, the data rate drops, at around $t = 25$ s, they begin the patrol and at the same time, the mobile infrastructure team begins operating. As the network team repositions, the data rate starts recovering to a baseline level of 5 Mbps at time $t = 60$ s, which is kept stable during the whole duration of the experiment. Similar behavior occurs with the delay experienced in the network, shown in Fig. 13(b). In this case, as the task agents pull away from the center, the delay increases and is brought back to stable levels under 115 ms when the mobile infrastructure team begins operating. Hence, the MID system is adaptive to complex trajectories of multiple task agents, ensuring consistent data rates and delays.

C. Robust autonomous reconfiguration

Finally, we conduct an experimental test to study network robustness and reconfiguration. This experiment is in the same vein as the simulation described in Section IV-C. The experiment setup is as follows. Two stationary task agents




Fig. 14: Spatial configurations and routing solutions of the mobile infrastructure during the reconfiguration experiment. After the replacement agent enters the system, the network nodes go back to the initial configuration. Black circles represent task agents, while red, blue, and green (replacement agent) circles represent network agents. In the trajectories, solid dots represent initial positions, while hollow dots correspond to end positions. The lines connecting these dots show the trajectories taken by the agents during the experiment.

(10.42.0.1 and 10.42.0.2) are placed at a separation of about 30 m. They take off and keep hovering without moving from their initial positions. Two network agents (10.42.0.3 and 10.42.0.4) are then deployed using our MID system to support traffic between the task agents. An additional network agent (10.42.0.5) is placed aside as a backup. At a given point during the experiment, one of the network agents (10.42.0.3) deactivates itself and the backup UAV is activated. The mobile infrastructure reconfigures itself dynamically and without human input to support the communication needs of the task agents.

The trajectories, routing, and spatial configurations during the experiment are shown in Fig. 14. In the initial configuration, the network agents position themselves in between near the middle of the two task agents. The optimal positioning is such as to provide a certain level of redundancy, instead of acting purely as a two-hop relay between the task agents. As such, the optimal routing configuration splits approximately half of the traffic from each task agent between the two network agents. When the network agent 10.42.0.4 drops out and the substitute network agent 10.42.0.5 is activated into the MID system, the network nodes reposition themselves. Agent 10.42.0.3 progressively moves to the position of the network agent that left the network, and the new network agent takes the previous position of agent 10.42.0.4. Once the repositioning stabilizes, the final configuration is similar to the initial configuration.

The data rate attained during the experiment can be seen in Fig. 15. When the starting network agent deactivates itself ($t = 13$ s), the average data rate drops and then progressively increases as the new network agent enters the system and the mobile infrastructure reconfigures itself. This effect is more evident when looking at the instantaneous data rate. When the network agent leaves the system ($t = 13$ s), the data rate immediately drops drastically. As the network reconfigures itself, the data rate progressively recovers, with the average communication rate remaining above 4.5 Mbps during the exchange and stabilizing to an average rate of 8 Mbps, comparable to the pre-exchange level.

Through the execution of this experiment, we have demonstrated how the MID system reacts when faced with the deactivation of a network agent, effectively recovering the network




Fig. 15: Achieved data rate for the reconfiguration experiment. The instantaneous rate collected at intervals of one second and a moving average over 20 seconds are shown. At $t = 13$ s the network agent 10.42.0.4 drops from the system and is subsequently substituted by agent 10.42.0.5. The MID team adapts during the exchange to hold an average communication rate slightly above 4.5 Mbps and stabilizes eventually to an average rate of 8 Mbps, which is comparable to the level before the agent dropout.

from system failure. As in the previous two experiments, our experimental finding not only corroborates the behavior observed in simulations, but also provides evidence for the practical viability of the proposed mobile infrastructure on demand system.

VI. CONCLUSIONS

We have introduced the MID system, utilizing UAVs to deliver wireless connectivity on demand. The UAVs serve as communication relays, forming an ad-hoc network and dynamically adapting their routing decisions and positions to guarantee the required quality of service demanded by the system users. Simulated and experimentally validated, the MID system has demonstrated its capability to serve many scenarios. Distilled into three distinctive features: **(F1)** expanding the operational range, **(F2)** supporting tasks involving complex mobility patterns, and **(F3)** exhibiting robustness via autonomous reconfiguration. Simulations and experiments have shown that MID outperforms fixed infrastructure solutions, offering greater efficiency with fewer infrastructure elements, enhanced area coverage, increased data rates, lower delays, and robustness to node failure. Given these observations, the MID system holds considerable potential in enhancing the quality of service and bridging connectivity gaps in previously underserved areas.

REFERENCES

- [1] D. Mox, M. Calvo-Fullana, M. Gerasimenko, J. Fink, V. Kumar, and A. Ribeiro, "Mobile wireless network infrastructure on demand," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7726–7732.
- [2] J. Fink, A. Ribeiro, and V. Kumar, "Robust control for mobility and wireless communication in cyber-physical systems with application to robot teams," *Proc. of the IEEE*, vol. 100, no. 1, pp. 164–178, 2011.
- [3] J. Stephan, J. Fink, V. Kumar, and A. Ribeiro, "Concurrent control of mobility and communication in multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1248–1254, 2017.
- [4] D. Mox, V. Kumar, and A. Ribeiro, "Learning connectivity-maximizing network configurations," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5552–5559, 2022.
- [5] A. Muralidharan and Y. Mostofi, "Communication-aware robotics: Exploiting motion for communication," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 115–139, 2021.
- [6] D. P. Spanos and R. M. Murray, "Robust connectivity of networked vehicles," in *2004 43rd IEEE Conference on Decision and Control (CDC)*, vol. 3. IEEE, 2004, pp. 2893–2898.
- [7] M. M. Zavlanos and G. J. Pappas, "Controlling connectivity of dynamic graphs," in *Proceedings of the 44th IEEE Conference on Decision and Control*. IEEE, 2005, pp. 6388–6393.
- [8] M. C. De Gennaro and A. Jadbabaie, "Decentralized control of connectivity for multi-agent systems," in *Proceedings of the 45th IEEE Conference on Decision and Control*. IEEE, 2006, pp. 3628–3633.
- [9] M. M. Zavlanos and G. J. Pappas, "Potential fields for maintaining connectivity of mobile networks," *IEEE Transactions on robotics*, vol. 23, no. 4, pp. 812–816, 2007.
- [10] E. Stump, A. Jadbabaie, and V. Kumar, "Connectivity management in mobile robot teams," in *2008 IEEE international conference on robotics and automation*. IEEE, 2008, pp. 1525–1530.
- [11] D. V. Dimarogonas and K. H. Johansson, "Decentralized connectivity maintenance in mobile networks with bounded inputs," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1507–1512.
- [12] 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*, vol. 46, no. 2, pp. 390–396, 2010.
- [13] M. M. Zavlanos, M. B. Egerstedt, and G. J. Pappas, "Graph-theoretic connectivity control of mobile robot networks," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1525–1540, 2011.
- [14] M. Ji and M. Egerstedt, "Distributed coordination control of multi-agent systems while preserving connectedness," *IEEE Transactions on Robotics*, vol. 23, no. 4, pp. 693–703, 2007.
- [15] M. M. Zavlanos and G. J. Pappas, "Distributed connectivity control of mobile networks," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1416–1428, 2008.
- [16] D. P. Spanos and R. M. Murray, "Motion planning with wireless network constraints," in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 87–92.
- [17] K. Savla, G. Notarstefano, and F. Bullo, "Maintaining limited-range connectivity among second-order agents," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 187–205, 2009.
- [18] Y. Kim and M. Mesbahi, "On maximizing the second smallest eigenvalue of a state-dependent graph laplacian," in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 99–103.
- [19] M. Schuresko and J. Cortés, "Distributed motion constraints for algebraic connectivity of robotic networks," *Journal of Intelligent and Robotic Systems*, vol. 56, no. 1, pp. 99–126, 2009.
- [20] M. M. Zavlanos, A. Ribeiro, and G. J. Pappas, "Network integrity in mobile robotic networks," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 3–18, 2012.
- [21] O. Tekdas, W. Yang, and V. Isler, "Robotic routers: Algorithms and implementation," *The International Journal of Robotics Research*, vol. 29, no. 1, pp. 110–126, 2010.
- [22] Y. Mostofi, "Communication-aware motion planning in fading environments," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 3169–3174.
- [23] Y. Mostofi, A. Gonzalez-Ruiz, A. Gaffarkhah, and D. Li, "Characterization and modeling of wireless channels for networked robotic and control systems—a comprehensive overview," in *2009 IEEE/RSJ Inter. Conf. on Intelligent Robots and Systems*. IEEE, 2009, pp. 4849–4854.
- [24] Y. Yan and Y. Mostofi, "Robotic router formation in realistic communication environments," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 810–827, 2012.
- [25] M. Calvo-Fullana, D. Mox, A. Pyattaev, J. Fink, V. Kumar, and A. Ribeiro, "Ros-netsim: A framework for the integration of robotic and network simulators," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1120–1127, 2021.
- [26] M. Calvo-Fullana, A. Pyattaev, D. Mox, S. Andreev, and A. Ribeiro, "Communications and robotics simulation in uavs: A case study on aerial synthetic aperture antennas," *IEEE Communications Magazine*, vol. 59, no. 1, pp. 22–27, 2021.
- [27] R. K. Williams, A. Gasparri, and B. Krishnamachari, "Route swarm: Wireless network optimization through mobility," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3775–3781.
- [28] V. S. Varadharajan, D. St-Onge, B. Adams, and G. Beltrame, "Swarm relays: Distributed self-healing ground-and-air connectivity chains," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5347–5354, 2020.
- [29] M. A. Schack, J. G. Rogers, Q. Han, and N. T. Dantam, "Robot team data collection with anywhere communication," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 705–711.
- [30] J. Fink, A. Ribeiro, and V. Kumar, "Robust control of mobility and communications in autonomous robot teams," *IEEE Access*, vol. 1, pp. 290–309, 2013.
- [31] J. Fink, "Communication for teams of networked robots," Ph.D. dissertation, University of Pennsylvania, 2011.
- [32] P. D. Hung, T. Q. Vinh, and T. D. Ngo, "Hierarchical distributed control for global network integrity preservation in multirobot systems," *IEEE transactions on cybernetics*, vol. 50, no. 3, pp. 1278–1291, 2019.
- [33] A. Goldsmith, *Wireless communications*. Cambridge university press, 2005.
- [34] R. Shorey, A. Ananda, M. C. Chan, and W. T. Ooi, *Mobile, wireless, and sensor networks: technology, applications, and future directions*. John Wiley & Sons, 2006.
- [35] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [36] "Robot Operating System (ROS)," <https://www.ros.org/>.
- [37] "Gazebo Simulator," <http://gazebosim.org/>.
- [38] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

APPENDIX

A. Optimal routing as an SOCP

The following problem is expressed in the standard convex SOCP formulation [38, Chapter 4.4.2] which can be solved using off-the-shelf optimization solvers, including the `cvxpy` library used in our ROS implementation.

$$\begin{aligned} & \underset{\mathbf{z}}{\text{minimize}} && \mathbf{f}^T \mathbf{z} \\ & \text{subject to} && \|\mathbf{C}_p \mathbf{z} + \mathbf{r}_p\| \leq \mathbf{d}_p^T \mathbf{z} + n_p, \quad p = 1, \dots, P \end{aligned} \quad (14)$$

To bring the optimal routing problem (9) into this standard form, let us define $P = KL + 2L + 2KL^2 + 1$, with K being the number of data flows and $L = N + M$ the total number of agents. Also define $\mathbf{z} = (\boldsymbol{\alpha}^T, s)^T$ with the probabilities α_{ij}^k collected in vector $\boldsymbol{\alpha}$. Then, the costs of (14) and (9) become identical by setting $\mathbf{f} = (\mathbf{0}^T, -1)^T$. If \mathbf{C}_p and \mathbf{r}_p are set to zero, then the corresponding constraint in (14) becomes linear in \mathbf{z} , and thus also linear in $\boldsymbol{\alpha}$. Hence, (14) admits the equivalent form

$$\begin{aligned} & \underset{\boldsymbol{\alpha}, s \geq 0}{\text{maximize}} && s \\ & \text{subject to} && \|\mathbf{C}_i^k \boldsymbol{\alpha}\| \leq \boldsymbol{\alpha}^T \mathbf{d}_i^k - s - m_i^k \quad \forall i, k \\ & && \mathbf{M} \boldsymbol{\alpha} \leq \mathbf{1} \\ & && \mathbf{0} \leq \boldsymbol{\alpha} \leq \mathbf{1} \end{aligned} \quad (15)$$

by redefining $\mathbf{d}_p = ((\mathbf{d}_i^k)^T, -1)^T$, $n_p = -m_i^k$, and $\mathbf{C}_i^k \boldsymbol{\alpha} = \mathbf{C}_p \mathbf{z}$ (removing the last column of \mathbf{C}_p) for the first $K(M + N)$ constraints, and then transforming the remaining ones into linear constraints and grouping them as $s \geq 0$, $\mathbf{M}\boldsymbol{\alpha} \leq \mathbf{n}$ and $\mathbf{0} \leq \boldsymbol{\alpha} \leq \mathbf{1}$. With (14) and (15) being equivalent, we consider henceforth problem (15) as our definition of an SOCP.

Next, we show how to set the constraints of (9) in the form of (15). In particular, (9) has the constraint

$$\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) - m_i^k - s \geq \zeta_i^k \sqrt{\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})} \quad (16)$$

with $\zeta_i^k \triangleq \Phi^{-1}(\epsilon_i^k)$ being the percentile of the error. Notice that (16) includes $\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})$ which is linear in $\boldsymbol{\alpha}$, so it can be rewritten as $\bar{b}_i^k(\boldsymbol{\alpha}, \mathbf{x}) = \boldsymbol{\alpha}^T \mathbf{d}_i^k$ by properly defining \mathbf{d}_i^k in terms of the average rate \bar{R}_{ij} (see Algorithm 2). Similarly, $\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})$ in (5) is quadratic in $\boldsymbol{\alpha}$ so that $\zeta_i^k \sqrt{\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})}$ can be rewritten using the ℓ_2 -norm in the form

$$\zeta_i^k \sqrt{\tilde{b}_i^k(\boldsymbol{\alpha}, \mathbf{x})} = \|\mathbf{C}_i^k \boldsymbol{\alpha}\| \quad (17)$$

by properly defining the diagonal matrix \mathbf{C}_i^k in terms of the rate variances \hat{R}_{ij} and the percentile ζ_i^k (also in Algorithm 2). Under these definitions, it is apparent that the constraints in (16) and (15) coincide. The remaining $\boldsymbol{\alpha} \in \mathcal{A}$ and $s \geq 0$ in (14) are linear, and thus they can also be written as in (15). More details on how to select the matrix \mathbf{M} and vector \mathbf{n} for that purpose are given in Algorithm 2, where \mathbf{e}_i are vectors of the canonical basis.

B. Optimal connectivity as an SDP

Let us start by arguing that the connectivity of a graph is strengthened by maximizing the second eigenvalue of the graph Laplacian, and then reformulate this maximization problem to solve it with a standard convex optimization algorithm. By construction, the smallest eigenvalue of the Laplacian is zero, and its corresponding eigenvector is the vector $\mathbf{1}$ which has all elements equal to one. Furthermore, the multiplicity of the null eigenvalue gives the number of disconnected sub-graphs. Indeed, if the graph has R disconnected regions, the Laplacian will be block diagonal with each of the R blocks being the Laplacian of each sub-graph. For instance, if a network with 5 nodes is such that nodes 1, 2, and 3 are pairwise connected but isolated from the pair of nodes 4 and 5, then the eigenvalue corresponding to the eigenvector $\mathbf{v} = (1, 1, 1, 0, 0)$ will also be null. This is because the first three ones will nullify the 3×3 Laplacian of the sub-graph that contains nodes 1, 2, and 3, and will not multiply the other block. Thus, the eigenvalue associated with $\mathbf{v} = (1, 1, 1, 0, 0)$ is also null. In this case, the null eigenvalue has multiplicity $R = 2$ with eigenvalues $\mathbf{v} = (1, 1, 1, 1, 1)$ and $\mathbf{v} = (1, 1, 1, 0, 0)$. For the network to be fully connected, the null eigenvalue must have multiplicity 1. Using a continuity argument, a second smallest eigenvalue being close to zero would unveil a fully connected graph, with some weak links, since the network will disconnect as this eigenvalue tends to zero. Finally, the magnitude of the second eigenvalue indicates the strength of the network.

In the direction of maximizing the second eigenvalue, consider the family of semidefinite programs (SDPs). These

Algorithm 2 Network Planner (SOCP Form)

```

Input  $\bar{R}_{ij}, \hat{R}_{ij}, m_i^k, \epsilon_i^k$ 
for  $k = 1, \dots, K$  do
  for  $i = 1, \dots, L$  do
    Set  $\mathbf{c}_i^k = \mathbf{d}_i^k = \mathbf{0}$ 
    for  $j = 1, \dots, L$  do
      Set  $m = (k-1)L^2 + (i-1)L + j$ 
      Set  $n = (k-1)L^2 + (j-1)L + i$ 
      Set  $p = i$ 
      Set  $q = L + j$ 
       $\mathbf{c}_i^k = \mathbf{c}_i^k + \mathbf{e}_m \bar{R}_{ij} + \mathbf{e}_n \hat{R}_{ji}$ 
       $\mathbf{d}_i^k = \mathbf{d}_i^k + \mathbf{e}_m \bar{R}_{ij} - \mathbf{e}_n \hat{R}_{ji}$ 
       $\mathbf{M} = \mathbf{M} + \mathbf{e}_p \mathbf{e}_m^T + \mathbf{e}_q \mathbf{e}_n^T$ 
    end for
    Set  $\mathbf{C}_i^k = \text{diag}(\mathbf{c}_i^k) / (\Phi^{-1}(\epsilon_i^k))$ 
  end for
end for
Solve SOCP (15) and output  $\boldsymbol{\alpha} = \text{vec}(\alpha_{ij}^k)$ 

```

are convex problems that admit the following standard form [38, Chapter 4.6.2],

$$\begin{aligned} & \underset{\mathbf{z}}{\text{minimize}} && \mathbf{f}^T \mathbf{z} \\ & \text{subject to} && \mathbf{F}_0 + z_1 \mathbf{F}_1 + z_2 \mathbf{F}_2 + \dots + z_P \mathbf{F}_P \preceq \mathbf{0} \quad (18) \\ & && \mathbf{Mz} = \mathbf{r} \end{aligned}$$

Both the semidefinite and the linear constraints can be separated into multiple constraints by selecting matrices \mathbf{F}_p and \mathbf{M} as block diagonal matrices, and the variables can be included, or not, in each of these multiple constraints by setting to zero, or not, the corresponding blocks. Even a nonnegative constraint can be incorporated using a block of dimension one. Hence we are going to adopt the following equivalent problem as our definition of an SDP

$$\begin{aligned} & \underset{\gamma, \mathbf{l}, \mathbf{a}, \mathbf{x}, \mathbf{u} \geq 0, \mathbf{v} \geq 0}{\text{maximize}} && \gamma \\ & \text{subject to} && \gamma \mathbf{I} - \sum_{n=1}^{L^2} l_n \mathbf{F}_n \preceq \mathbf{0} \\ & && \mathbf{x} - \hat{\mathbf{x}} + \mathbf{u} = \Delta \mathbf{1} \\ & && \hat{\mathbf{x}} - \mathbf{x} + \mathbf{v} = \Delta \mathbf{1} \\ & && \mathbf{l} = (\mathbf{C} - \mathbf{I})\mathbf{a} \\ & && \mathbf{a} = \mathbf{Bx} + \mathbf{d} \end{aligned} \quad (19)$$

where we also separated $\mathbf{z} = (\gamma, \mathbf{l}^T, \mathbf{a}^T, \mathbf{x}^T, \mathbf{u}^T, \mathbf{v}^T)^T$ and renamed the blocks of the matrices \mathbf{F}_p and \mathbf{M} and vector \mathbf{r} in a way that will be useful to prove the next equivalence under the connectivity problem (13). Finally, for the costs of (19) and (18) to coincide, we must select $\mathbf{f} = (-1, \mathbf{0}^T)^T$.

To put (13) in the form of (19), define $\mathbf{l} = \text{vec}(\hat{\mathbf{L}})$ and $\mathbf{a} = \text{vec}(\hat{\mathbf{A}})$, where the $\mathbf{y} = \text{vec}(\mathbf{Y})$ operator returns a vector \mathbf{y} formed by the concatenation of the columns of its matrix argument \mathbf{Y} . The cost of (13) is already equal to that of (19). The constraint $\mathbf{P}^T \mathbf{L} \mathbf{P} \succeq \gamma \mathbf{I}$ is equivalent to $\gamma \mathbf{I} - \mathbf{P}^T \mathbf{L} \mathbf{P} \preceq \mathbf{0}$ that can be rewritten as in (19) by properly defining the matrices \mathbf{F}_n in terms of the columns of \mathbf{P} as it is detailed in

Algorithm 3 *Connectivity Planner (SDP Form)*

Input \hat{R}_{ij} , $\nabla_{x_i} \hat{R}_{ij}$, $\nabla_{x_j} \hat{R}_{ij}$, $\hat{\mathbf{x}}$, and \mathbf{P} with rows \mathbf{p}_i
 Set $\mathbf{B} = \mathbf{0}$, $\mathbf{C} = \mathbf{0}$, $\mathbf{d} = \mathbf{0}$
for $i = 1, \dots, L$ **do**
 $m = (i - 1)L + i$
 for $j = 1, \dots, L$ **do**
 $n = (i - 1)L + j$
 $\mathbf{F}_n = \mathbf{p}_i \mathbf{p}_j^T$
 $\mathbf{C} = \mathbf{C} + \mathbf{e}_m \mathbf{e}_n^T$
 Set $\mathbf{g}_n = \nabla_{x_i} \hat{R}_{ij}$ and $\mathbf{h}_n = \nabla_{x_j} \hat{R}_{ij}$
 $\mathbf{B} = \mathbf{B} + \mathbf{e}_n ((\mathbf{e}_i \otimes \mathbf{g}_n)^T + (\mathbf{e}_j \otimes \mathbf{h}_n)^T)$
 $\mathbf{d} = \mathbf{d} + \mathbf{e}_n (\hat{R}_{ij} - \mathbf{g}_n^T \hat{\mathbf{x}}_i - \mathbf{h}_n^T \hat{\mathbf{x}}_j)$
 end for
end for
 Solve SDP (19) and output \mathbf{x}

Algorithm 3. The constraint $\|\mathbf{x} - \hat{\mathbf{x}}\|_\infty \leq \Delta$ can be put as the system of linear equations in (19) by adding the nonnegative auxiliary variables \mathbf{u} and \mathbf{v} . The remaining constraints in (13) are linear in \mathbf{a} , \mathbf{l} , and \mathbf{x} , so that they can be rewritten as $\mathbf{l} = (\mathbf{C} - \mathbf{I})\mathbf{a}$ and $\mathbf{a} = \mathbf{B}\mathbf{x} + \mathbf{d}$ by properly defining \mathbf{B} , \mathbf{C} , and \mathbf{d} . More details, including the construction of \mathbf{B} , \mathbf{C} , and \mathbf{d} can be found in Algorithm 3, where the symbol \otimes stands for the Kronecker product.

C. Mapping to a common reference

By default, the position x_i of a UAV is referenced to its own local reference system s_i where they were initialized. In our experiments, we define a nearby common reference s_r to be shared by all UAVs. We select this point as a clear visual reference within the zone where we are going to carry out the experiments. In order to obtain a transformation H_i^r that maps the position x_i to this common reference system s_r , the mapping node follows the following procedure. It takes the GPS coordinates of s_r as input parameters and receives from PX4 the GPS coordinates of the current position s_p . With these parameters, it computes the transformation H_g^r between the global GPS reference s_g and the common reference s_r . It also computes the transformation H_g^p between the global GPS reference s_g and the current GPS coordinates s_p . Then it obtains from PX4 the transformation H_i^p from the initialization to the current position. These three transformations give the sufficient information to compute $H_r^i = (H_g^r)^{-1} H_g^p (H_i^p)^{-1}$. Hence, using H_r^i it is possible to map the position x_i to the common reference s_r , transforming the position of the UAV, from its internal reference system s_i , to the common reference system s_r . At a high level, all variables are expressed in these common coordinates when processed by the network and connectivity planners in ROS, and the mapping node is called to communicate to and from the lower-level flight controller.