2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

# Opportunistic Communication in Robot Teams

Daniel Mox, Kashish Garg, Alejandro Ribeiro, and Vijay Kumar

*Abstract*— **In this paper we present a new approach to *Mobile Infrastructure on Demand* (MID) where a dedicated team of robots creates and sustains a wireless network that satisfies the communication requirements of a different team of task-oriented robots seeking to coordinate their actions in the absence of existing communication infrastructure. Different from previous works, our approach forgoes heuristics for network performance such as algebraic-connectivity or network flow optimizations and instead positions communication support robots to directly maximize the probability of packet delivery by the underlying opportunistic routing protocol. Our system is task agnostic and practical to implement and operate on robots equipped with off-the-shelf WiFi radios. We demonstrate this through a set of experiments showing our MID system maintaining the delivery of critical mission data in a situational awareness setting and enabling foraging robots to effectively coordinate their actions during multi-robot exploration.**

## I. INTRODUCTION

The promise of multi-robot systems is that they can complete tasks faster and more efficiently than any single robot. However, these benefits are only realized if the robot team can communicate in order to effectively coordinate their actions. Many robot teams operate in tightly controlled environments with access to existing network infrastructure in order to communicate. With continued improvement in on-board autonomy, robots are increasingly being deployed in challenging environments beyond the reaches of existing communication networks [1]. In these scenarios, robots can utilize their on-board wireless radios to form peer-to-peer or ad-hoc networks in order to exchange vital coordination information. However, actions taken to complete task objectives often take robots out of direct communication range with one another (e.g. in multi-robot exploration, foraging robots disperse to map out a space and avoid revisiting areas already covered by another member of the team) introducing a tension between task fulfillment and maintaining contact with the rest of the team.

Research in the area of communication in robot teams has saught to reconcile this tension between competing task and communication objectives. Early work considered the graph induced by the known positions of each robot in the team coupled with a channel model predicting the condition of communication links between them [2]–[6]. These approaches saught to maintain or improve the algebraic-connectivity of this communication graph in both a centralized and decentralized manner. Algebraic-connectivity is inherently a heuristic for network performance and later work saught to bridge the gap by introducing more sophisticated network performance objectives such as end-to-end bit-error-rate [7] or a packet flow optimization [8]. Other work eschewed communication models entirely, instead estimating the spatial variation of communication signals and moving robots to maximize expected received signal strength [9]. More recent work introduced the most sophisticated system yet with global and local planners for controlling a team of robots to complete a task objective while satisfying a robust packet routing optimization [10], [11].

One critical choice in robotic communication systems is the network performance metric. While algebraic-connectivity affords a convenient mathematical framework for judging the quality of the network it is inherently a heuristic for performance. In practice one would have to operate an ad-hoc routing protocol to handle network traffic and it is unclear if maximizing algebraic-connectivity also maximizes the performance this protocol. More sophisticated models don't necessarily solve this problem. While packet flow algorithms like the ones in [8], [10], [11] are useful for high level network load balancing they are difficult to operate in practice as we found in our previous work [12]. Robot teams often require numerous flows of information causing the dimension of the optimization to increase significantly; furthermore, not all exchanges of information in robot teams can be accurately modeled as constant bandwidth flows (e.g. intermittent exchanges of map data in an exploration setting).

We also desire solutions for communication in robot teams that are suitable for deployment in highly dynamic scenarios to support complex tasks. Ensuring routing optimization programs are satisfied is a computational expensive task and not amenable to on-line deployment. In addition, other works require specific restrictions such as stationary task agents [9] or a single source and destination node [7].

In this paper we seek to address these challenges by introducing a new approach to the problem of *Mobile Infrastructure on Demand* (MID) formalized in [12] and similar to [9], [13]. In MID a set of communication robots that we control act as mobile routers and support the network requirements of a different set of client robots collaborating to accomplish some objective which we observe. This approach

---

The work of Daniel Mox was supported by the National Science Foundation Graduate Research Fellowship under Grant DGE-1845298. This work was supported in part by ARL under Grant DCIST CRA W911NF-17-2-0181, in part by NSF under Grant CNS-1521617, in part by ARO under Grant W911NF-13-1-0350, in part by ONR under Grants N00014-20-1-2822 and N00014-20-S-B001, and in part by Qualcomm Research. (Corresponding author: Daniel Mox.)

Daniel Mox, Kashish Garg, and Vijay Kumar are with the General Robotics, Automation, Sensing and Perception (GRASP) Laboratory at the University of Pennsylvania, Philadelphia, PA 19104 USA (email: mox@seas.upenn.edu, kumar@seas.upenn.edu)

Alejandro Ribeiro is with the Electrical and Systems Engineering Department at the University of Pennsylvania, Philadelphia, PA 19104 USA (email: aribeiro@seas.upenn.edu)

979-8-3503-8457-4/24/$31.00 ©2024 IEEE

decouples task planning from network maintenance so that the task oriented agents can go about their objective without considering the impact their actions have on their ability to communicate. Additionally, the communication or MID team simply requires knowledge of the position of the task agents and their communication requirements so that solutions to the MID problem are task agnostic. Critically, in the place of a heuristic for network performance, we introduce a planner that directly optimizes the performance of the underlying routing protocol handling network traffic. Following our planner, MID agents move to maximize the probability that packets get delivered to their intended destination. We also show the practicality of our method through a set of experiments where our MID system supports task agents performing situational awareness and multi-robot exploration [1].

## II. METHODOLOGY

In this work we seek a solution to the problem of *Mobile Infrastructure on Demand*. Given the positions of a set of task-oriented robots, we seek target positions for a different set of robots (i.e. the MID or network or communication team) that move to ensure flows of information in the task team are preserved over the duration of the mission. As we elaborate on in this section, our controller seeks to satisfy these communication requirements by maximizing the probability that packets in each flow are successfully delivered by the underlying routing protocol handling traffic in the network. In Section II-A we detail this important choice of routing protocol, then in Section II-B we introduce a gradient-based local planner, and finally in Section II-C we discuss how these pieces are integrated into a complete MID system.

### *A. Opportunistic Routing*

Mobile ad-hoc routing is a well studied problem within the wireless systems community and there exists an extensive body of literature dedicated to developing capable solutions [14]. Routing protocols can be categorized in many different ways. One particular axis of comparison involves when the choice of next hop relay is made. Traditional protocols are proactive in nature, seeking to use cached network state information in order to plan the best possible packet path before transmission. Consider the scenario shown in Fig. 1 where node A seeks to transmit to node D. A traditional routing protocol might select node C as the next-hop relay. However, one feature of wireless networks is that they operate via a shared medium. In other words, only one node can transmit at a time and while this transmission occurs any node within range can decode the packet regardless of if they are the intended recipient [2]. While node C may be chosen as the next hop, node B may also receive the packet. In a proactive routing setting, if C does not receive the transmission node A must retry the transmission at a later time even though B received the packet and is closer to the desired destination D than node A. A new class of opportunistic routing protocols seek to capitalize on these fortuitous receptions by deferring the choice of next hop relay till after transmission. The set of nodes that could act as next hop relays are embedded in the packet header and then, after transmission, the candidate relays utilize a time slot mechanism to coordinate which one relays. Returning to the example in Fig. 1, an opportunistic protocol might list both B and C as relays in order of increasing priority; if C receives the packet it will immediately relay it to D. While this is happening B is monitoring the airwaves; if B does not overhear a transmission from C after a specified period of time then it assumes C did not relay the packet and thus did not receive it to begin with. In this case, B takes charge and relays the packet. While opportunistic protocols require this additional coordination step they have been shown to outperform their proactive counterparts in networks with lossy links, such as those often encountered in robotics [15], [16]. In addition, by allowing more than one node to act as the next hop relay opportunistic protocols are more robust to errors in cached network state information and increase the probability that the packet makes progress towards the destination for each transmission. Note that the choice of allowable next hop relays is typically a proper subset of the network, distinguishing opportunistic protocols from flooding protocols. This is illustrated in the example shown in Fig. 1 where nodes B and C were selected to relay but not node E. The set of allowable relays at each node for each flow is carefully selected using collected network state information so that this distributed coordination mechanism is likely to succeed. In this paper we closely follow the protocol SOAR and further details about its operation such as the relay selection algorithm can be found in [16]. In the remainder of this work, we assume that the network state information and routing table constructed by the protocol are available.

### *B. Local Planner*

With knowledge of the relays the protocol is using, we can enumerate the possible paths a packet might take through the network from source to destination. An example of the routing table for the scenario in Fig. 1 is shown below in Table I.

TABLE I: an example routing table for the network shown in Fig. 1 for flows between nodes $A$ and $D$.

| flow | A relays | B relays | C relays | D relays |
|---|---|---|---|---|
| A → D | C, B | D, C | D | - |
| D → A | - | A | A, B | B, C |

Traversing the routing table, we can formulate the set of possible paths a packet might follow for the flow $A \rightarrow D$ as $\mathcal{R} = \{ABD, ACD, ABCD\}$. A route $r = \{ABD\} \in \mathcal{R}$ is composed of a set of links $l_1 = AB$ and $l_2 = BD$ and each link has an associated probability of packet delivery $P_l$. The

---

[1] code and multimedia available at: `www.danmox.com/projects/oc.html`

[2] considering standard 802.11 WiFi operating in ad-hoc/IBSS mode and ignoring spatial reuse and modern amendments like OFDMA

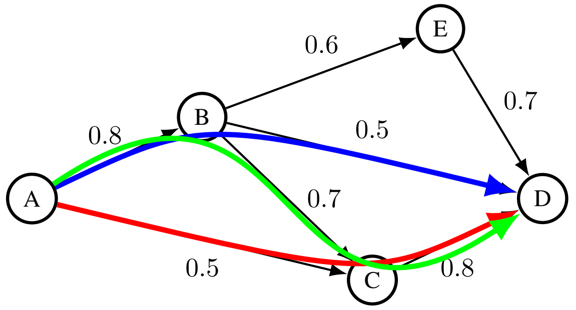

Fig. 1: An example ad-hoc network where node A seeks to transmit a packet to node D. The arrows show the probability of delivery for each link and the possible paths a packet might take through the network if B and C are selected as intermediate relays. A critical decision by the protocol is which nodes can participate in routing; here B and C were selected but E was not.

probability that a packet gets delivered along at least one of the paths in the flow follows from the probability that they all fail:

$$\begin{aligned} P_d &= 1-(1-P_{AB}P_{BD})(1-P_{AC}P_{CD})(1-P_{AB}P_{BC}P_{CD}) \\ &= 1-\prod_{r\in\mathcal{R}}\left(1-\prod_{l\in r}P_l\right). \end{aligned} \quad (1)$$

Supposing we have a channel model that provides the packet delivery probability of a link as a function of the distance between the two nodes, we can compute the gradient of Eq. (1) with respect to the position of a MID agent $x_i$ as:

$$\frac{\partial P_d}{\partial x_i}=\sum_{r\in\mathcal{R}}\frac{\partial P_r}{\partial x_i}\prod_{s\in\mathcal{R}\setminus r}(1-P_s). \quad (2)$$

In multi-robot tasks there are often many flows of information between different robots in the task team. The user can specify a set of source/destination pairs of importance or the system can operate under the assumption that all possible flows between task agents should be maintained. In either case, the route enumeration process described previously is repeated for each source/destination pair and the gradients of each flow are summed:

$$\frac{\partial P}{\partial x_i}=\sum_{k=0}^{F}\sum_{r\in\mathcal{R}_k}\frac{\partial P_r}{\partial x_i}\prod_{s\in\mathcal{R}_k\setminus r}(1-P_s) \quad (3)$$

where there are $F$ total flows and $\mathcal{R}_k$ is the set of possible paths for flow $k\in\{1,2,\cdots,F\}$. Finally, we employ Eq. (3) in a gradient ascent scheme to find the locally optimal goal position that MID node $i$ should move to in order to increase the probability packets get delivered.

In previous works, a great deal of effort was devoted to selecting sufficiently rich functions to capture the inherent uncertainty in wireless communication [7], [10]–[12]. In this work, we take a different approach and use a linear function to model the probability of delivery between two nodes:

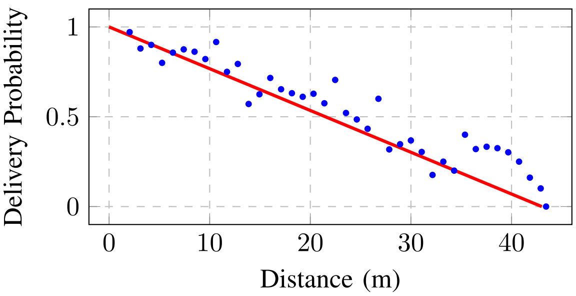

Fig. 2: Packet delivery probability as a function of distance. Roughly 3000 packets transmission attempts were made between two nodes at varying distance in the experimental space of Section III. Transmission attempts were divided into 40 bins based on distance with the blue dots showing the delivery probability of each bin.

$$P(d)=1-d\ /\ d_c \quad (4)$$

where $d_c$ is the cutoff distance after which delivering packets along the link drops to zero. While this approach might seem overly simplistic, we note a few things. First, delivery probability statistics collected in the space we conducted experiments exhibit a roughly linear trend with distance as can be seen in Fig. 2. Second, since the opportunistic protocol on which our local planner is built capitalizes on a diversity of links for forwarding, it makes the choice of channel model used for planning less critical. For these reasons, the linear model in Eq. (4) is sufficient for our MID system. Nevertheless, there is nothing preventing the use of more complicated channel models with our approach provided they are differentiable with respect to node position.

Gradient ascent on Eq. (3) with Eq. (4) as the channel model produces locally optimal configurations. While finding globally optimal solutions is generally infeasible, we can perform exhaustive search on the simple line scenario shown in Fig. 3a with two task agents and two MID agents to gain insight into the kinds of configurations our planner might converge to. When nodes A and D are fixed relatively close together, the MID configuration that maximizes the probability of packet delivery is both MID relays lumped together at the midpoint (Fig. 3b). This configuration also maximizes the number of possible routes from source to destination and minimizes the number of hops in each path. Of course, as $d_t$ grows beyond the range of a single agent, the optimal configuration is a daisy chain evenly dividing the space between A and D (Fig. 3c). Interestingly, Fig. 3a illustrates that our planner prefers path diversity over link quality. In other words, configurations with many packet paths with few hops albeit via poorer links have higher delivery probabilities than configurations with fewer possible packet paths with more hops via higher quality links.

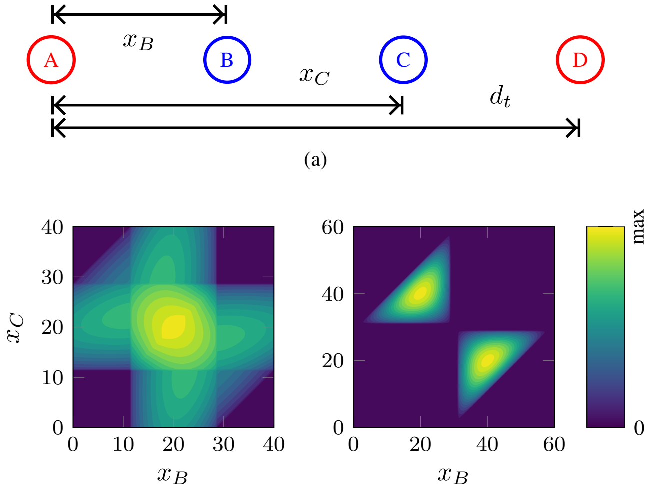

(a)

(b) $d_t = 40$

(c) $d_t = 60$

Fig. 3: (a) A scenario with two task agents (red) supported by two MID agents (blue) and the corresponding delivery probability reward surfaces for different values of the distance between the task agents $d_t$ in (b), (c). In each case, the linear channel model of Eq. (4) with $d_c = 30$ is used.

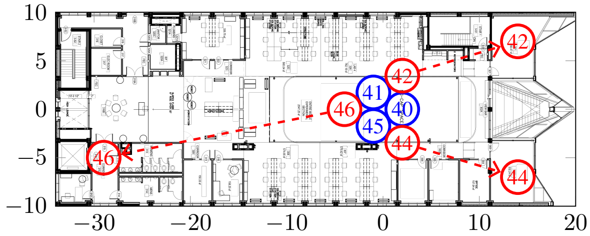

Fig. 4: An illustration of the situational awareness experiment. The red task agents of 42, 44, and 46 start at a central depot and spread out to their indicated destinations. Robot 46 seeks to send 8 Kbps of traffic to both 42 and 44. Robots 40, 41, and 45 are MID agents.

### *C. System Implementation*

Each robot is equipped with an 802.11 WiFi radio configured in ad-hoc/IBSS mode. Without an available implementation of SOAR, we developed our own in C/C++ utilizing the broadcast channel and unix sockets. Each node estimates the state of the network (i.e. link ETX [17]) by periodically transmitting / listening for beacon packets, which we also utilize to disseminate this link state information along with the latest position of each robot. Link state information is fed into the route selection algorithm used to update the routing table applied by the protocol. Robot positions coupled with Eq. (4) are used to generate link state and routing information that acts as a smooth approximation to their beacon based counterparts suitable for control. Our routing protocol is taylored for ROS and acts as a multi-master similar to the popular FKIE Multimaster [18]. Arbitrary ROS traffic that must be sent to other nodes in the network is captured by the protocol and serialized into IP packets transmitted by the physical layer. Because we utilize the broadcast channel, each node receives all nearby network traffic regardless of destination which enables the opportunistic gains discussed earlier. Packets arriving at a node are processed by the protocol and relayed or re-published on the local ROS instance as appropriate. Every node in the network runs an instance of our protocol implementation while only MID robots run the planner. In this way, all traffic exchanged between robots regardless of team affiliation flows through the protocol.

## III. RESULTS

The goal of our proposed MID system is to support the communication requirements of a set of task agents seeking to coordinate their actions. Due to the stochastic nature of wireless networking, the true test of a communication system is its ability to operate in the noisy, lossy, cluttered environments in which robots are often called upon to operate. To this end, we conducted a suite of experiments utilizing the system described in the previous section with the Scarab robots developed at the University of Pennsylvania [19]. Scarabs are differential drive ground robot platforms equipped with a powerful onboard computer capable of running a complete ROS-based autonomy stack in addition to other application related programs. Each is outfitted with a 2D Hokuyo Lidar for navigation and obstacle avoidance, an ASUS Xtion structured light sensor providing limited field of view RGBD images of the environment, and off the shelf WiFi radios. In the following sections we detail experiments where our MID system supports a team of task agents performing situational awareness and multi-robot exploration.

### *A. Situational Awareness*

In the situational awareness task, patrolling agents must remain in contact with a fixed base station [10], [11], [20]. This task poses a challenge as patrolling trajectories stretch the network, requiring the supporting MID system to adapt and reconfigure online. In our case, Scarabs 42, 44, and 46 acting as task agents begin at a central depot and then disperse to predefined target locations; while this is happening, 46 sends critical mission information in the form of 100 Byte packets at a rate of 10 Hz to 42 and 44. An illustration of this scenario is shown in Fig. 4. We conducted the test twice: with and without MID support. A plot of the packet delivery probability over the duration of each test is shown in Fig. 5 and snapshots of the team configurations are shown in Fig. 6.

A clear takeaway from Fig. 5 is the stark difference between the case with MID support and without. This is not particularly surprising. Without MID support, the task agents were forced to route packets among themselves and eventually along low probability links as the team dispersed. While task agents 42 and 44 could act as relays for each other adding some level of receiver redundancy the overall success of broadcasts across the space was low. As is apparent in Fig. 5 after $t$ = 250s, the links between 46 and the other

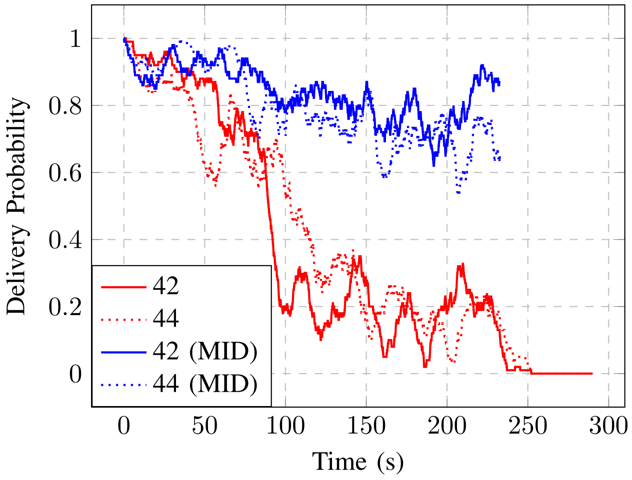

Fig. 5: Delivery probability for packet streams beginning at Scarab 46 destined for Scarabs 42 and 44 for the cases with and without MID support.

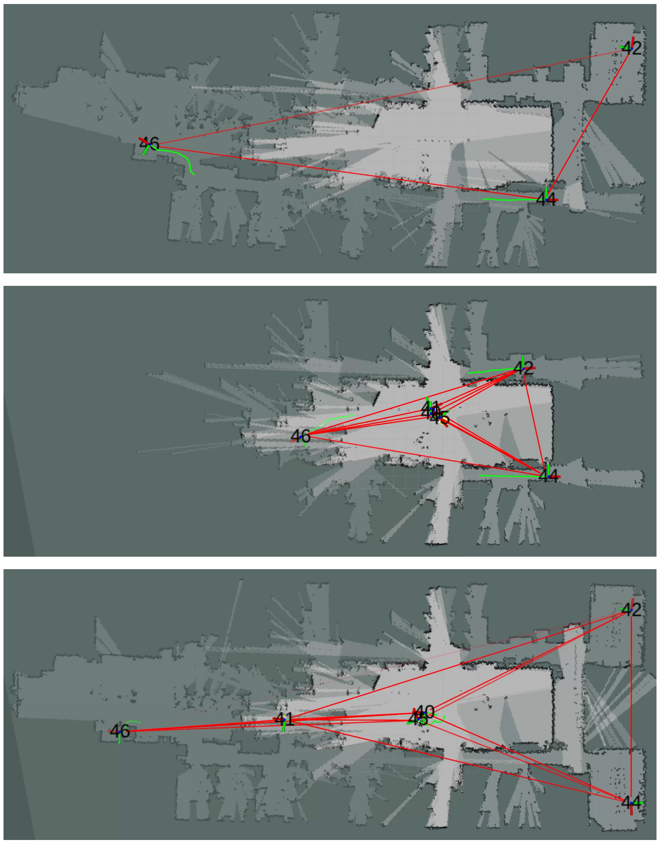

Fig. 6: Snapshots of the situational awareness test without MID at $t = 120$s (top) and with MID at $t = 80$s (middle) and $t = 200$s (bottom). The link probabilities estimated through beaconing are shown as red lines with proportional transparency.

task agents eventually broke and task traffic ceased entirely.

In contrast, with MID support the flows of packets in the task team were maintained with a high degree of reliability throughout the experiment. As the task agents grew further and further apart, the MID agents dynamically positioned themselves in order to maximize the probability of end-to-end packet delivery. Even as the task agents reached the extents of the space, little change can be noticed in the delivery probability in Fig. 5. Initially, the optimal configuration tracked by the MID team was a cluster positioned at the center of convex hull formed by the task agents (see Fig. 6 middle). The position of the MID team is less consequential at these early stages of the experiment as the task agents are able to effectively communicate with one another directly. However, as the experiment progressed and the task agents grow further apart, the cluster of MID agents split up with agent 41 moving towards 46 and agents 40, 45 remaining together closer to 42 and 44 (Fig. 6 bottom). While we cannot make claims about the global optimality of this configuration it is consistent with our intuition from Section II-B that the planner seeks to maximize the probability of packet delivery through path diversity. As the task agents reach their ultimate destinations, traffic in the network is primarily relayed through the MID agents as can be seen in Fig. 6 (bottom).

One might wonder why the delivery probability isn't closer to 100% for the MID case. In both experiment runs, packets were transmitted by the protocol in a best effort manner with acknowledgments / retries disabled so that the results in Fig. 5 are a function of the network configuration produced by our MID planner and not of the capability of the protocol itself. As such, some packets are dropped due to the lossy links in the network. However, as we will see in Section III-B reliable transport with retries / acknowledgments can increase the reliability of packet delivery close to 100%.

### B. *Multi-Robot Exploration*

The second task we support with our MID system is multi-robot exploration (MRE), where team of robots are tasked with autonomously building a map of the environment. Multiple robots can explore a space faster than any single robot can provided they coordinate their actions and choose to visit different regions. A natural consequence of this behavior is that communication links between exploring robots often break as they grow further apart, conspiring against their ability to efficiently complete the task. In this experiment, we demonstrate the use of our MID system to ensure a team of exploring robots can successfully coordinate their goals.

Autonomous exploration with robots is a well studied problem. The fundamental challenge lies in analyzing partial maps of the environment to determine where to travel next to complete the map in a fast, efficient manner. The traditional approach involves repeatedly seeking the border between known and unknown space, known as the frontier, until all free space has been visited [21]. While other more sophisticated approaches exist (e.g. [22], [23]), the focus of our experiment is on demonstrating our MID system. Thus, we choose frontier exploration as it is simple, effective, and lends itself well to multi-robot coordination. Indeed, many MRE approaches build off of frontier exploration and involve sharing frontier locations and computing goals for the team that do not coincide [24]–[26]. Some approaches assume robots can communicate throughout exploration by utilizing existing network infrastructure [24] while others propose coordination strategies that handle intermittent connectivity but cannot guarantee robots won't re-explore visited regions [25], [26]. Enabling systems like these to operate without existing communication networks and without sacrificing performance is a main motivation of our work.

In our experiments, we employ a distributed, asynchronous coordination approach based on sharing known frontiers and current goals. In a nutshell, each exploring robot maintains a synchronized set of frontiers goals that none of them have visited. This is accomplished by sharing candidate frontier goals with the team and notifying other agents when one falls within an already visited space. In addition, each robot periodically shares their current goal and remaining path cost so that no two robots pursue the same one.

We conducted our multi-robot mapping experiment across a 10,000 sq. ft. indoor office space with scarabs $42$, $44$, $46$ dedicated to exploration and scarabs $40$, $41$, $45$ running our MID system. Snapshots of the experiment at key points can be seen in Fig. 7. This time, packet flows between task agents in this experiment were transmitted in reliable mode with retries and acknowledgments enabled. While exploring, agents $42$, $44$, and $46$ exchanged a total of $3,396$ coordination messages among which $46$ were dropped for a packet delivery rate of $98.6\%$. Following our coordination scheme, the exploring agents remained largely separate during the experiment, initially beginning in a cluster towards the center of the space but quickly spread out to visit different areas of the map. Note that the opacity of the underlying occupancy grid indicates the number of robots that have covered the space with faint regions visited by a single robot and fully opaque regions by the entire team. With our MID system for support, the exploring robots could effectively coordinate throughout the mission and after $450$s seconds the entire space was explored with little overlap as seen in highlighted regions in the bottom of Fig. 7.

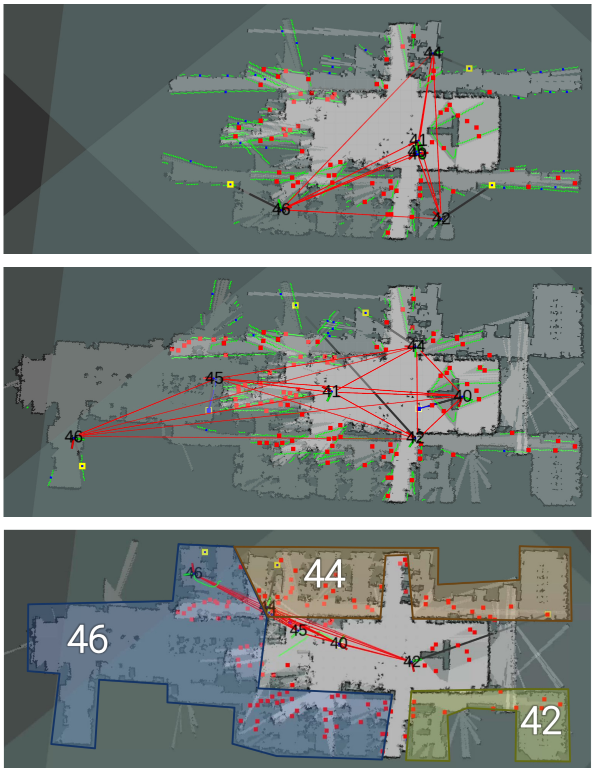

Fig. 7: Snapshots of the multi-robot exploration experiment with MID support at $t = 150$s (top), $t = 290$s (middle), and at the end at $t = 450$s (bottom). Link probabilities are shown as red lines with proportional transparency. Frontier cells are green lines, visited goals red squares, unvisited goals blue squares, and current goals yellow squares.

Supporting multi-robot exploration posed much more of a challenge for our MID planner. Over the course of the task, the exploration team's configuration translated and morphed, requiring MID agents to handle more interference from obstacles as compared with the situational awareness task. While no explicit global planning was performed, each robot was equipped with a navigation stack capable of avoiding static and dynamic obstacles. Utilizing this capability, the MID agents pursued their goals, navigating around obstacles if necessary and seeking to get as close as possible to the target goal if it was unreachable. Because of this, MID agents would occasionally fall out of formation for a short period of time (e.g. agent 40 in Fig. 7 at $t = 290$s) but were able to swiftly recover their target position in the network. It should be noted that while the experiment space is a cluttered, obstacle filled environment, it is largely a simply connected domain in which our local controller with obstacle avoidance functions fairly well. In order to operate in a more complex environment (e.g. with long connecting corridors or larger obstacles) our MID system would require global planning in order to reconfigure the team when the gradient controller becomes trapped in a local minimum. This remains an avenue for future work.

## IV. CONCLUSION

In this work we have presented a new solution to *Mobile Infrastructure on Demand*, enabling teams of task oriented robots to successfully coordinate their actions beyond the reach of existing communication networks. Different from previous work, our approach does not use a heuristic for network performance but instead adjusts network agent positions in order to directly maximize the end-to-end packet delivery probability of the underlying opportunistic routing protocol. Critically, our system is task agnostic and practical to implement and operate on mobile robots equipped with off the shelf WiFi radios. To show this, we conducted a situational awareness task where our MID approach improved the delivery rate of critical mission information and a multi-robot exploration task where it enabled foraging robots to effectively coordinate their actions.

# REFERENCES

[1] Darpa subterranean (subt) challenge. `https://www.darpa.mil/program/darpa-subterranean-challenge`. Accessed: 2023-08-28.

[2] Yoonsoo Kim and Mehran Mesbahi. On maximizing the second smallest eigenvalue of a state-dependent graph laplacian. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 99–103. IEEE, 2005.

[3] Maria Carmela De Gennaro and Ali Jadbabaie. Decentralized control of connectivity for multi-agent systems. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 3628–3633. IEEE, 2006.

[4] Meng Ji and Magnus Egerstedt. Distributed coordination control of multiagent systems while preserving connectedness. *IEEE Transactions on Robotics*, 23(4):693–703, 2007.

[5] Ethan Stump, Ali Jadbabaie, and Vijay Kumar. Connectivity management in mobile robot teams. In *2008 IEEE international conference on robotics and automation*, pages 1525–1530. IEEE, 2008.

[6] Michael M Zavlanos, Magnus B Egerstedt, and George J Pappas. Graph-theoretic connectivity control of mobile robot networks. *Proceedings of the IEEE*, 99(9):1525–1540, 2011.

[7] Yuan Yan and Yasamin Mostofi. Robotic router formation in realistic communication environments. *IEEE Transactions on Robotics*, 28(4):810–827, 2012.

[8] Michael M Zavlanos, Alejandro Ribeiro, and George J Pappas. Network integrity in mobile robotic networks. *IEEE Transactions on Automatic Control*, 58(1):3–18, 2012.

[9] Stephanie Gil, Swarun Kumar, Dina Katabi, and Daniela Rus. Adaptive communication in multi-robot systems using directionality of signal strength. *The International Journal of Robotics Research*, 34(7):946–968, 2015.

[10] Jonathan Fink, Alejandro Ribeiro, and Vijay Kumar. Robust control of mobility and communications in autonomous robot teams. *IEEE Access*, 1:290–309, 2013.

[11] James Stephan, Jonathan Fink, Vijay Kumar, and Alejandro Ribeiro. Concurrent control of mobility and communication in multirobot systems. *IEEE Transactions on Robotics*, 33(5):1248–1254, 2017.

[12] Daniel Mox, Miguel Calvo-Fullana, Mikhail Gerasimenko, Jonathan Fink, Vijay Kumar, and Alejandro Ribeiro. Mobile wireless network infrastructure on demand. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7726–7732. IEEE, 2020.

[13] Onur Tekdas, Wei Yang, and Volkan Isler. Robotic routers: Algorithms and implementation. *The International Journal of Robotics Research*, 29(1):110–126, 2010.

[14] Azzedine Boukerche, Begumhan Turgut, Nevin Aydin, Mohammad Z Ahmad, Ladislau Bölöni, and Damla Turgut. Routing protocols in ad hoc networks: A survey. *Computer networks*, 55(13):3032–3080, 2011.

[15] Sanjit Biswas and Robert Morris. Exor: Opportunistic multi-hop routing for wireless networks. In *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 133–144, 2005.

[16] Eric Rozner, Jayesh Seshadri, Yogita Mehta, and Lili Qiu. Soar: Simple opportunistic adaptive routing protocol for wireless mesh networks. *IEEE transactions on Mobile computing*, 8(12):1622–1635, 2009.

[17] Douglas SJ De Couto, Daniel Aguayo, John Bicket, and Robert Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 134–146, 2003.

[18] Fkie multimaster - ros packages for multimaster support. `https://fkie.github.io/multimaster_fkie/`. Accessed: 2023-09-15.

[19] Nathan Michael, Jonathan Fink, and Vijay Kumar. Experimental testbed for large multirobot teams. *IEEE robotics & automation magazine*, 15(1):53–61, 2008.

[20] Ethan Stump, Nathan Michael, Vijay Kumar, and Volkan Isler. Visibility-based deployment of robot formations for communication maintenance. In *2011 IEEE international conference on robotics and automation*, pages 4498–4505. IEEE, 2011.

[21] Brian Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97.'Towards New Computational Principles for Robotics and Automation'*, pages 146–151. IEEE, 1997.

[22] Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. Information-theoretic mapping using cauchy-schwarz quadratic mutual information. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4791–4798. IEEE, 2015.

[23] Kelsey Saulnier, Nikolay Atanasov, George J Pappas, and Vijay Kumar. Information theoretic active exploration in signed distance fields. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4080–4085. IEEE, 2020.

[24] John G Rogers, Carlos Nieto-Granda, and Henrik I Christensen. Coordination strategies for multi-robot exploration and mapping. In *Experimental Robotics: The 13th International Symposium on Experimental Robotics*, pages 231–243. Springer, 2013.

[25] Wolfram Burgard, Mark Moors, Cyrill Stachniss, and Frank E Schneider. Coordinated multi-robot exploration. *IEEE Transactions on robotics*, 21(3):376–386, 2005.

[26] Robert Zlot, Anthony Stentz, M Bernardine Dias, and Scott Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings 2002 IEEE international conference on robotics and automation (Cat. No. 02CH37292)*, volume 3, pages 3016–3023. IEEE, 2002.