

ACHORD: Communication-Aware Multi-Robot Coordination with Intermittent Connectivity

Maira Saboia^{1,2*}, Lillian Clark^{3*}, Vivek Thangavelu^{4*}, Jeffrey A. Edlund¹, Kyohei Otsu¹, Gustavo J. Correa⁵, Vivek Shankar Varadharajan⁶, Angel Santamaria-Navarro⁷, Thomas Touma⁸, Amanda Bouman⁸, Hovhannes Melikyan¹, Torkom Pailevanian¹, Sung-Kyun Kim¹, Avak Archanian¹, Tiago Stegun Vaquero¹, Giovanni Beltrame⁶, Nils Napp⁴, Gustavo Pessin², Ali-akbar Agha-mohammadi¹

Abstract—Communication is an important capability for multi-robot exploration because (1) inter-robot communication (comms) improves coverage efficiency and (2) robot-to-base comms improves situational awareness. Exploring comms-restricted (e.g., subterranean) environments requires a multi-robot system to tolerate and anticipate intermittent connectivity, and to carefully consider comms requirements, otherwise mission-critical data may be lost. In this paper, we describe and analyze ACHORD (Autonomous & Collaborative High-Bandwidth Operations with Radio Droppables), a multi-layer networking solution which tightly co-designs the network architecture and high-level decision-making for improved comms. ACHORD provides bandwidth prioritization and timely and reliable data transfer despite intermittent connectivity. Furthermore, it exposes low-layer networking metrics to the application layer to enable robots to autonomously monitor, map, and extend the network via droppable radios, as well as restore connectivity to improve collaborative exploration. We evaluate our solution with respect to the comms performance in several challenging underground environments including the DARPA SubT Finals competition environment. Our findings support the use of data stratification and flow control to improve bandwidth-usage.

Index Terms—Networked Robots; Multi-Robot Systems; Co-operating Robots; Distributed Robot Systems; Field Robots.

I. INTRODUCTION

MULTI-robot systems can accomplish tasks which are unrealistic for a single robot, especially when it comes to spatially distributed objectives, and allow for redundancy and robustness to individual robot failures [1], [2]. Particularly in harsh environments, supervised autonomy minimizes risks to the individual robots and to the overall mission, but relies on comms with the remote supervisor. To mitigate stress on other parts of the solution, it is advantageous to design a multi-robot autonomy solution that maximizes comms capability.

We consider the setting where exploring robots must (1) navigate in an environment where comms might be lost

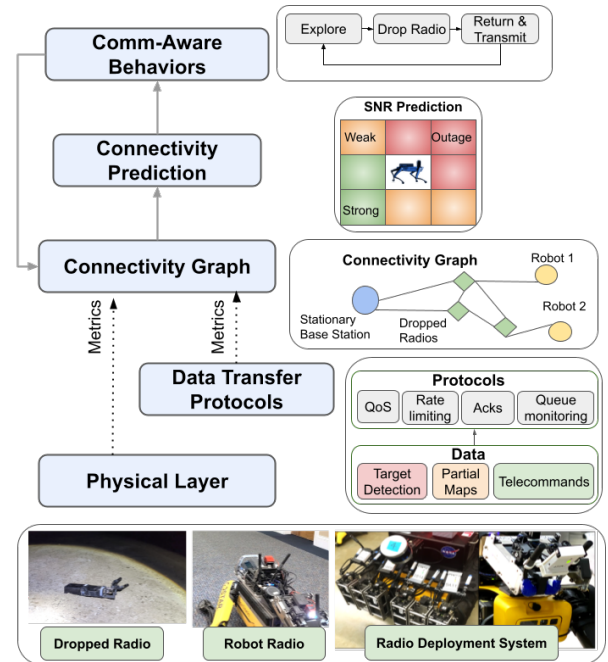


Fig. 1. ACHORD: At the physical layer, robots use wireless radios (comms nodes) to form a mesh network with droppable radios they deploy. Metrics from the lower-layers of the networking stack (e.g., SNR and queue size) inform high-level connectivity monitoring, prediction, and comms-aware behaviors.

temporarily, (2) communicate findings to a stationary base which requires situational awareness, and (3) share data with nearby robots. Maintaining comms links is challenging during autonomous exploration of large-scale environments, especially those with winding passages or obstacles which prevent line-of-sight comms. Maintaining links at all times constrains the maximum coverage a multi-robot team can achieve. Even taking approaches to expand the network through relay nodes (static nodes or robots), the effective comms range of the stationary base is limited by the number of nodes. Furthermore, the available bandwidth decreases with each additional network node. An efficient usage of this bandwidth requires careful management and an understanding of the network health (congestion, available bandwidth, interference, etc) considering the large amounts of data typically generated by robotic exploration (e.g., 3D maps of large environments). Transferring pointclouds, maps, or streaming video in real-time requires high-bandwidth and stable comms links.

Manuscript received February 24, 2022; revised July 12, 2022; accepted June 27, 2022.

This paper was recommended for publication by M. Ani Hsieh upon evaluation of the Associate Editor and Reviewers' comments.

* denotes equal contribution.

¹Jet Propulsion Laboratory, California Institute of Technology, CA, USA. ²Instituto Tecnológico Vale, MG, Brazil. ³University of Southern California, CA, USA. ⁴Cornell University, NY, USA. ⁵University of California, Riverside, CA, USA. ⁶Polytechnique Montréal, QC, CA. ⁷Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Spain. ⁸California Institute of Technology, Pasadena, CA, USA.

©2022. All rights reserved.

Digital Object Identifier (DOI): see top of this page.

In this paper we address two challenges. Firstly, we present a comms architecture which meets the needs of exploration: dedicated protocols for distinct classes of data and exposure of statistics on link-quality and queue sizes to upper layers of the networking stack for planning. Secondly, we present the implementation of behaviors during exploration which meet the needs of comms: robots autonomously determine when to drop additional comms nodes and when to prioritize timely data transfer.

A. Related Work

Comms-aware exploration: The communication challenges posed by multi-robot exploration have drawn increased attention [3]–[5]. In particular, ensuring the availability of comms links between all robots [6]–[8] or via dedicated relay robots [9]–[11] are well-researched objectives. More recently, *intermittent* connectivity has been considered as a more flexible objective, and several works present path-planning methods which ensure intermittent connectivity by requiring robots meet infinitely often [12]–[14] in a known environment. During exploration, maintaining connectivity (even intermittently) poses unique advantages as large-scale exploration and connectivity are opposing objectives [3], [15]. Designing comms-aware exploration strategies requires modeling connectivity, and existing works model link qualities as a function of distance [16] or other factors like shadowing and multipath components [17], [18]. However, decision-making based solely on connectivity fails to consider whether robots have new information to transfer, and realistic memory constraints [4]. In this work, we consider these practical needs and incorporate information about the network state including the size of data buffers at each robot. Recent works have verified this concept theoretically and in simulation [5], [15], [19], [20], while we consider its implementation in practice as one component of our multi-layer architecture.

Comms protocols: The most common protocols for data transfer in multi-robot systems are Data Distribution Service (DDS), Robot Operating System (ROS), and Message Queuing Telemetry Transport (MQTT). DDS prioritizes performance [21], while MQTT focuses on a lightweight solution for the Internet of Things and ROS focuses on enabling modular development. Despite its popularity, previous work has shown that ROS is not well-suited for networks of robots subject to intermittent connectivity [22]. With this in mind, most field-hardened networking approaches rely on custom solutions built directly on UDP/TCP which act as a bridge between ROS-enabled robots [23]–[26]. ROS2 is better-suited for multi-robot systems than ROS and provides significant improvements by leveraging DDS and configurable Quality of Service (QoS) parameters, which allow differentiating between classes of data. However, ROS2 is still missing some key features needed for multi-robot autonomy with intermittent and variable bandwidth connectivity.

CHORD: In our previous work [27], we introduced CHORD (Collaborative High-Bandwidth Operations with Radio Droppables). CHORD is a hybrid ROS1/2 data transfer solution and demonstrates the advantages of using QoS for

radio traffic while using ROS TCP connections for intra-robot communication. However, we observed some issues with network congestion when using ROS2’s *reliable* QoS to resend messages after a transmitter rejoins the network. Furthermore, CHORD lacked the necessary bookkeeping to provide queue size monitoring at the application layer. This bookkeeping, introduced in this work in Sec. III-B, enables high-level autonomy (described in Sec. IV) which explicitly considers these metrics. Our cross-layer design philosophy allows ACHORD to jointly consider low-level networking and high-level decision-making, while the latter was out of the scope of our previous work.

B. Contributions

We present an overview of ACHORD, as shown in Fig. 1, a multi-layer networking solution which focuses on scalability and bandwidth-usage in the joint design of low-level networking and high-level decision-making. We analyze its performance as part of the larger NeBula framework [28] on a network of up to six robots and up to 13 relay nodes in varied environments. Contributions unique to ACHORD include:

- **Bandwidth-aware prioritization:** To efficiently use the available bandwidth, we introduce a novel classification of data and leverage dedicated protocols to meet QoS needs.
- **Network state representations:** To enable comms-aware exploration, we propose a rich representation of the network which considers the radio propagation environment, network congestion, and data queue sizes.
- **Predictive signal modeling:** To adapt to changes in the dynamic network, we introduce the use of radio propagation models to predict link quality.
- **Comms-aware coordination:** We propose and implement behaviors such that robots can autonomously prioritize timely data transfer without sacrificing exploration.

II. PROBLEM FORMULATION

We consider a team of heterogeneous robots which explores a large-scale environment in the absence of existing comms infrastructure (e.g., wifi access points) and must communicate findings with a stationary, remote base station. This operational scenario poses a number of requirements on the system.

Limited range: First, the range at which robots can communicate directly with the base station is severely limited by the scale of the environment and obstacles or winding passageways which prevent line-of-sight comms. This imposes the need for additional comms infrastructure to extend the effective range of the base, for example relay network nodes placed throughout the environment. The complexities of the environment also render simple connectivity models insufficient, and more careful monitoring of the network is needed.

Intermittent connectivity: Second, the exploration objective is to reduce uncertainty by visiting and sensing unexplored areas. Thus, even with the use of relay nodes, it is expected that the robots will explore beyond the effective comms range of the base. This imposes the need for a comms architecture which gracefully handles intermittent connectivity. Data

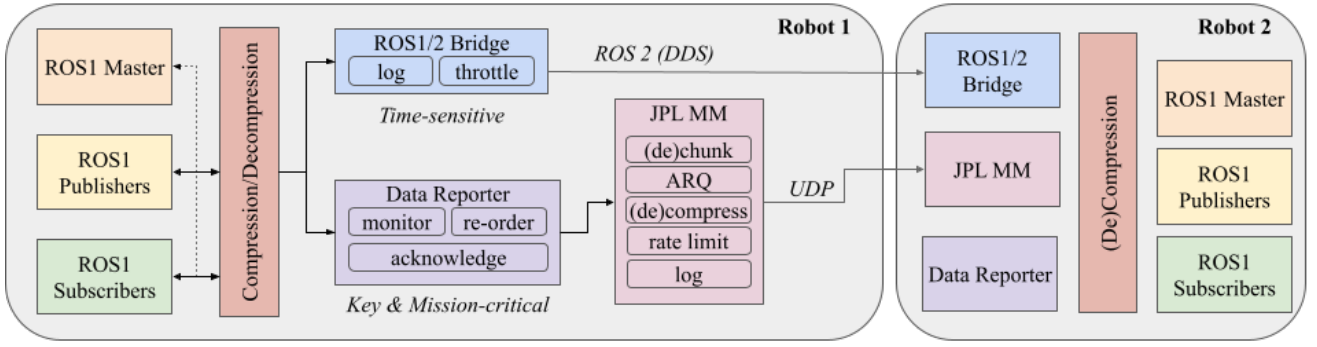


Fig. 2. Diagram of software architecture for inter-robot comms. Comms to and from the base is similar.

transfer protocols should be disruption-tolerant, and reliably transmit all data which is critical to the mission. Further, data transfer protocols should achieve low-latency during periods when connectivity is available, to enable up-to-date situational awareness or teleoperation by human supervisors.

Limited bandwidth: Third, we assume the multi-robot team generates megabytes of data per minute. This imposes the need for efficient use of the limited bandwidth, especially as distance and non-line-of-sight stress link quality. After periods without connectivity, robots accumulate large quantities of data and the comms architecture needs to avoid the risk of this flooding the network.

Network layer metrics: Finally, although the dynamics of the network can be seen as a challenge, the opportunity for controlled mobility presents advantages. Leveraging the ability of robots to navigate to areas with connectivity or areas which would benefit from the presence of a dropped relay node imposes the need to circumvent isolation between layers of the network stack. For a cross-layer solution, the comms architecture must expose metrics from the lower layers to the application layer for comms-aware operations.

Given these design requirements, we decompose the high-level problem of ensuring robots can communicate all mission-critical data to the base in a timely manner into layers. The following sections describe the details of ACHORD.

III. HIGH-BANDWIDTH MULTI-ROBOT NETWORKING

This section describes the lower layers of our architecture. We discuss the wireless mesh network and present protocols specific to multiple robots with intermittent connectivity. To cover large-scale, comms-limited environments, we create a mesh network using commercially available wireless layer 2 devices (e.g., radios from Rajant or Silvus).

A. Mesh Network Deployment

To extend the effective comms range and establish a backbone network, robots are equipped with the NeBula Communication Deployment System (NCDS) [28], [29], shown in Fig. 1. The NCDS has a modular design suited for wheeled and legged robots and can carry up to six droppable radios. The NCDS can deploy radios upon request from the human supervisor or autonomously via a *scheduler*. A serial-ROS connection provides real-time sensory feedback and detailed state logging on which slots have radios loaded and which have

deployed. An important component of the NCDS is automatic jam detection, which allows the scheduler to report a drop failure and initiate a new drop. This improves the resiliency of the system to operational challenges.

Each robot monitors the Signal-to-Noise Ratio (SNR) [30] of links between itself and all other radios and uses this information to autonomously deploy the droppable radios to prevent loss of signal. SNR can be measured passively for any received message, and thus does not introduce comms overhead. The Shannon capacity, or theoretical max data capacity for any channel, is given by $C = B \log_2 (1 + \text{SNR})$ where C is the capacity and B is the bandwidth available [30]. This means for a defined bandwidth, increasing SNR increases the amount of data which a link can support.

The goal is to keep a lower bound on the *bottleneck SNR* between the robot and base station. The bottleneck SNR of a multi-hop route is the SNR of the weakest link along the route, and we assume data flows along the route with the highest bottleneck SNR. When this value falls below the desired threshold, the NCDS scheduler is triggered. This typically occurs as the robot gets farther from closest dropped radio, as the backbone network links are above the lower bound by design. If the bottleneck SNR falls below the lower bound, the robot is able to backtrack along its path before committing to the drop location. Our solution then evaluates the traversability risk [31], the terrain inclination, and the environment geometry [32] to locally select the exact drop position, favoring junctions, dry, and flat surfaces.

B. Data Transfer Protocols

This section describes our inter-robot comms software architecture which meets the requirements enumerated in Sec. II, as shown in Fig. 2. For intra-robot comms, we leverage ROS for ease of development. Using ROS for intra-robot comms also allows us to use high bandwidth links on the robot without the limitations of ROS2's throughput controller and isolates the radio traffic to only topics explicitly shared with ROS2 or JPL MM.

Data classes: For inter-robot comms, we consider three types of data to transfer: (i) key; (ii) mission-critical; and (iii) time-sensitive. (i) Key data refers to information that needs to be shared periodically and in-order, which is required for the nominal multi-robot mission control, but has no timing restrictions. Examples of key data are the telemetry of the robots

or the incremental maps they share. (ii) Mission-critical data includes crucial asynchronous information. While we desire low-latency, the correct transfer of this information is of higher priority than its transmission time. An example of mission-critical data is the detection and localization of a target object. (iii) Time-sensitive data is selected considering potential harms or vehicle integrity risks, thus an example includes sharing relative positioning between neighboring robots in a collision trajectory. All data which is transmitted over the wireless network is compressed into a generic, compressed data blob using bzip2¹.

ROS2/DDS: For neighbor discovery and transmitting time-sensitive data, we use ROS2 DDS (eProsima FastRTPS). As described in [27], each robot has a ROS1/2 bridge, based on the `ros1_bridge` package². While ROS2/DDS offers many improvements specific to multi-robot networking, it has two shortcomings that we did not address in CHORD [27]: (1) it lacks the option to resend only certain messages after reconnecting with the network, limiting the developer to select a static number of messages to resend (if too large, this will flood the network), and (2) it does not offer the application-layer a way to monitor the amount of data waiting to transmit, which is desirable in our case for high-level decision-making. For disruption-tolerant networks, a resend policy at the application layer which also offers buffer size monitoring is key.

Data Reporter: To address this, we introduce the *Data Reporter*, which monitors reliable data transfer. For each reliable (key or mission-critical) message sent over the network, an acknowledgement message (ACK) is sent from the receiver back to the sender. On the sender side, messages from data publishers wait in per-topic queues (buffers) to be transmitted and are removed from the queue when an ACK is received. On the receiver side, per-topic queues hold received messages which are then published via ROS to data subscribers. The receive side queue implements re-ordering for key data. Because the data reporter is implemented in ROS, metrics on per-topic queue sizes can easily propagate up to higher levels of ACHORD as shown in Fig. 1.

JPL MM: Key and mission-critical topics require a solution which is highly configurable and provides guarantees of message delivery. JPL multi-master (JPL MM) [25] is a module which provides inter-robot comms that is compatible with ROS and built on top of the User Datagram Protocol (UDP). JPL MM allows specifying configurations for each ROS topic and network port, and provides many customizable features. The primary responsibility of JPL MM on the sender side is to chunk the data into UDP datagrams and transmit them. Then on the receiver side data is reassembled into messages. JPL MM provides Selective Repeat ARQ (automatic repeat-request) using datagram-ACKs to ensure reliability at the transport layer. It also supports additional compression as needed. One of the main advantages of JPL MM for this application is that it provides token bucket rate limiting. A fixed number of tokens is allocated representing the maximum bandwidth available, and these tokens are allocated to each




Fig. 3. Information RoadMap (IRM) constructed collaboratively during exploration of the DARPA SubT Finals competition environment. The environment state is indicated by green, orange, and red comms checkpoints which represent strong, weak, and no comms, respectively.

key/mission-critical ROS topic. This lets us prevent certain types of messages from overwhelming the network.

IV. COMMS-AWARE MULTI-ROBOT EXPLORATION

This section describes the higher layers of our architecture. We discuss modeling and prioritizing connectivity while maximizing coverage of unknown environments.

A. State Representation

We introduce two representations of state used for comms-aware decision-making: the spatial environment state and network performance state.

Spatial environment state: We maintain a semantic, spatial representation of the environment called an Information RoadMap (IRM) [33]. As illustrated in Fig. 3, an IRM is a generic graph that captures the environment via nodes and edges, where nodes represent locations and edges connect nodes if a robot can travel between them. We categorize IRM nodes into four types: i) frontier nodes in unexplored space, ii) breadcrumb nodes in previously visited locations, iii) comms checkpoints, which are nodes associated with a signal strength, and iv) dropped radio nodes. Thus, the backbone network topology (as illustrated in Fig. 1) is captured by the IRM. It is an incrementally built, shared structure such that when robots meet or return to the base they merge their IRMs, favoring more recent data.

Each comms checkpoint stores the bottleneck SNR value a robot would experience at that location. Comms checkpoints with $\text{SNR} \geq T_C = 20\text{dB}$ are considered strong (green tag markers in Fig. 3), while checkpoints with $0 \geq \text{SNR} < T_C$ are considered weak (orange tag markers). When $\text{SNR} = 0$

¹<https://www.sourceware.org/bzip2/>

²https://github.com/ros2/ros1_bridge




Fig. 4. Connectivity map of the DARPA SubT Finals competition environment based on predicted signal strength, which can inform autonomy and augment situational awareness for the human supervisor.

(red tag markers), we nominally do not have comms, although the network may transmit some packets sporadically.

Network performance state: The spatial comms representation is unaware of the network usage; a location with a high SNR may exhibit slow data transfer rates due to congestion or interference. For effective comms-aware operations, the robots also maintain statistics on the reliable data that needs to be transferred to the base or other robots: (1) buffer size: the amount of data (in bytes) that needs to be transferred; (2) measured data rate: the amount of data transferred per unit time; and (3) estimated transfer time: the amount of time required to empty the buffer. These properties together determine the network state and help autonomy to prioritize exploration over connectivity or vice versa.

B. Predictive Comm Modeling

Our network topology is dynamic by design, given the mobility of the robots and the deployment of relay nodes. This introduces a challenge for relying on the environment state, as comms checkpoints can become outdated. To overcome this, our solution uses a signal propagation model to update the spatial environment representation with predicted SNR values. We assume the noise level σ_{dB} at a comms checkpoint location is invariant and focus on predicting signal strength.

Signal propagation model: Radio signal propagation is a multi-scale process, but first-order models typically calculate the free space path loss, which quantifies the expected attenuation in an obstacle-free environment [30]. Path loss is modeled as a logarithmic function of distance d given by

$$PL_{dB}(d) = PL(d_0)_{dB} + \eta 10 \log_{10}(d/d_0). \quad (1)$$

$PL(d_0)_{dB}$ is the reference path loss in dB at a known distance d_0 , and η is the path loss exponent which captures how quickly the signal falls off. Values $\eta = 3.83$ and $PL(d_0 = 1m)_{dB} = 34$ were selected after fitting this model to experimental data

via linear regression, giving the predicted SNR received by j from transmitter i :

$$SNR_{dB}(i, j) = Tx_{dB}(i) - PL_{dB}(d(i, j)) - \sigma_{dB}(j) \quad (2)$$

$Tx_{dB}(i)$ is the transmit power and $d(i, j)$ is the distance between i and j . In a previous work, we have shown that more accurate prediction can be achieved by modeling second and third order effects of the environment, at the cost of increased complexity [18].

Connectivity maps: Beyond updating checkpoints, having a predictive comms model allows estimating the connectivity offered at arbitrary locations based on the position estimates of the relay nodes, as illustrated in Fig. 4. As depicted in Fig. 1, the information in this predicted map flows up to the comms-aware behaviors discussed in the next section.

C. Comms-Aware Coordination

Equipped with the metrics exposed by our low-layer protocols and the state representation described in Sec. IV-A, robots autonomously make comms-aware decisions. Here we describe several elements of coordinated comms-aware operations.

Return to Comms: When the buffer size exceeds an upper bound ($T_B^u = 300KB$), the robot will sacrifice nominal exploration and instead move towards an area with strong comms. Our solution selects the closest comms checkpoint from the IRM with $SNR \geq T_C$, or a frontier neighboring this strong comms checkpoint, and moves towards it. If the buffer size drops below a desired threshold ($T_B^l = 200KB$) before reaching the target checkpoint, nominal exploration continues immediately. Otherwise, the robot will wait at the target checkpoint until the buffer size drops below T_B^l . If the buffer size does not decrease within a timeout (60 secs), which could indicate that the network is congested, the robot will choose a strong comms checkpoint closer to the base station as the new target. It will continue in this manner until the buffer size drops.

Radio deployment coordination: Coordination between robots is required during radio deployment to prevent redundant deployments in the same area. To enable this coordination, robots communicate their intentions via the shared IRM. When a robot reaches the triggering condition to deploy a radio, as described in Sec. III-A, the robot includes a dedicated node in the IRM that represents the intention to deploy a radio at that position. Before deploying an additional radio, other robots will consider whether the expected radio coverage of the two dropped radios would have significant overlap, and skip redundant deployments.

V. EXPERIMENTAL RESULTS

This section evaluates ACHORD in several comms-restricted environments and discusses lessons learned.

Our robot team included three Boston Dynamics Spot robots and three Clearpath Robotics Husky robots. The field test took place in a limestone mine (KY Underground) depicted in Fig. 5. The mine spanned hundreds of meters, with thick (20m) columns preventing line-of-sight comms. The other environment we tested in was constructed for the DARPA




Fig. 5. Snapshots of the network during exploring of KY Underground. Colors indicate the signal strength of links (yellow is strong, purple is weak). Timestamps (sec) are given in the lower right corner of each snapshot.




Fig. 6. Snapshots of the network during exploration of the DARPA SubT Finals competition environment, day 2.

Subterranean Challenge depicted in Fig. 6. The DARPA SubT Finals course was characterized by narrow, winding passageways and had three distinct subsections: an urban environment similar to a subway station, a mine-like environment, and a subterranean cave-like environment.

Given our priority of high-bandwidth comms, ACHORD leverages commercial off-the-shelf MIMO radios from Silvus Technologies (Streamcaster 4240 for the robots and 4400 for the base station) which are designed for mobile ad-hoc networks. The Silvus radios offer multi-hop routing and layer 2 protocols, and equipping the robots with these mesh nodes enables them to act as relays or data mules as needed. Silvus also provides an API for collecting link quality metrics like SNR, loss rate, noise level, etc. This allows us to propagate these metrics from the physical and link layer up to the application layer.

Table I summarizes results on the overall performance of our system with respect to its ability to enable high-bandwidth comms and give the human supervisor a thorough understanding of the environment. The number of dropped radios indicates how much network infrastructure the robots were able to autonomously deploy, which also depends on the scale of the environment. The much larger field test environment required more deployed nodes. The maximum delay indicates

TABLE I
SUMMARY OF RESULTS

	Field Test	Day 1	Day 2	Day 3
Time exploring (mins)	60	30	30	60
Exploring robots (#)	6	6	4	6
Deployed radios (#)	13	1	6	7
Maximum delay (sec)	813	49	1058	98
Effective comm range (m)	173	70	86	68
Up time (mins)	16	14	20	47
Peak data rate Base → Robots (Mbps)	22.61	19.60	34.58	19.39
Peak data rate Robots → Base (Mbps)	12.43	14.41	17.14	20.30




Fig. 7. KY Underground data rates (kbps) received by the base station (red) and transmitted by the base station (blue). Robots actively explore from 1019 to 4609 sec. After exploration ended, as the robots returned to the base, the data rates increased.

the greatest period of time for which any robot was not able to transfer data to the base station. Note that exploration beyond the range of the deployed infrastructure, which is desirable, requires outages. The effective comms range is a measure of the longest distance from the base any robot was able to travel without losing connectivity through one or more hops. Up time captures the maximum percentage of time an exploring robot is able to maintain connectivity with the base through one or more hops, as indicated by low data reporter buffer sizes. The data rates at the base station, for both incoming and outgoing traffic, are an indicator of overall bandwidth the network is able to support (see Fig. 7. Traffic from the base includes primarily mission status updates and the aggregated mapping data. The key takeaways are that ACHORD enables comms over more than 150 meters of exploration and our network can support up to 20Mbps of data from six exploring robots.

Comparison with CHORD: While the many novel aspects of ACHORD make a direct comparison to CHORD [27] challenging, we can highlight the performance improvement offered by certain features. For example, Fig. 8 shows the sizes of data reporter queues at each robot during exploration for two settings of ACHORD. In the first setting (top graph), we required ordered receipt for mission-critical topics, as in CHORD. We observed that the robots built up large queues, and even after restoring connectivity (e.g., spot4 at 1400s), we observed congestion likely due to unnecessary retransmissions which strained the available bandwidth. In the second setting (bottom graph), we introduced the additional data stratification presented in Sec. III-B and observed significantly less build up.

Lessons learned: We observed a scenario in which a robot

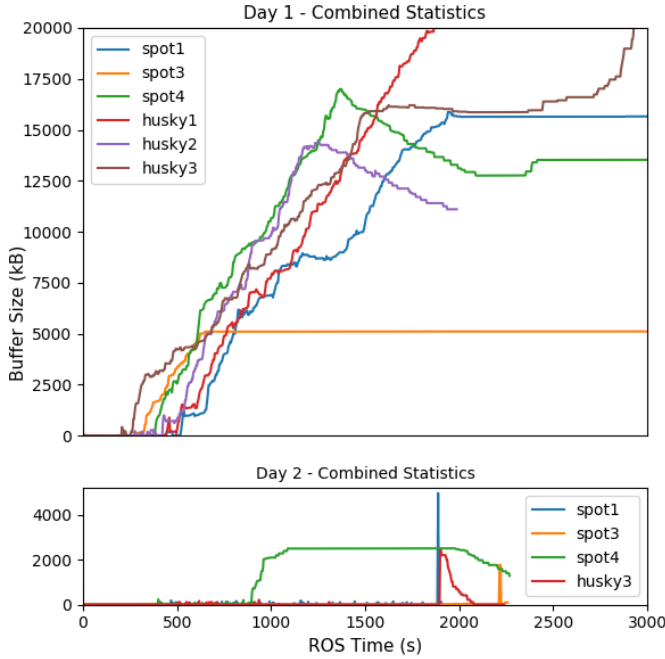


Fig. 8. Size of the data reporter buffer (summed over all key and mission-critical topics). The top graph depicts results from Day 1 with the data classification from CHORD and the bottom graph depicts results from Day 2 with the new data classification.

got stuck while out of comms range and another robot acted as a data relay, recovering data which would have otherwise been lost. This capability proved a significant advantage, and leads us to conclude that equipping high-level decision-making with inter-robot comms models is key. With this in mind, high-fidelity simulation of the wireless network would have aided in development and testing of comms-aware autonomy. In smaller environments, we found it was sufficient to drop nodes generously. The criteria for autonomous node deployment needs careful consideration in expansive environments where droppable nodes are a relatively scarce resource. If the network infrastructure covers enough of the environment, we found it was sufficient to return to comms sparingly. With sophisticated autonomy that can reliably determine when to return and transfer data, a small effective comms range and long outages are permissible. On the other hand, with a large effective comms range, a human supervisor has better opportunities to control and intervene, and less sophisticated autonomy is permissible. To design a resilient multi-robot system, we found it is not enough to focus on only one of these objectives, and considering both connectivity and autonomy is key.

VI. CONCLUSION

In this work, we present an overview of ACHORD, our multi-layer networking solution which provides timely and reliable data transfer for intermittently connected multi-robot teams and leverages droppable radios and comms-aware operations to improve connectivity. ACHORD is field-hardened through experiments in several underground environments with teams of up to six robots. Our finding indicate that

taking the radio environment and network state into account is advantageous for multi-robot exploration when a remote base needs to be kept situationally aware. Autonomous relay node deployment extends the effective comms range of the system, improves signal quality, and reduces delays and connectivity outages. Data prioritization and efficient bandwidth usage are key to enabling exploration of large-scale environments with multiple robots. Better handling of data prioritization at the semantic level and improved strategies for controlling access to the shared wireless medium are two directions for further study.

ACKNOWLEDGEMENTS

We gratefully acknowledge all members of Team CoSTAR. We also thank Belal Wang and Silvus Technologies for their hardware and technical expertise, and Kentucky Underground Storage for access to their facilities. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was supported in part by FUNDEP, by NASA Space Technology Research Fellowship Grant No. 80NSSC19K1189, and by NSF Grant No. #1846340. Further, it was partially supported by the Spanish Ministry of Science and Innovation under the project EB-CON (PID2020-119244GB-I00, MCIN/AEI/10.13039/501100011033) and by the European Commission under the project CANOPIES (H2020-ICT-2020-2-101016906).

REFERENCES

- [1] F. Rossi, S. Bandyopadhyay, M. Wolf, and M. Pavone, "Review of multi-agent algorithms for collective behavior: a structural taxonomy," *IFAC-PapersOnLine*, vol. 51, no. 12, pp. 112–117, 2018.
- [2] A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," *Ieee Access*, vol. 6, pp. 28 573–28 593, 2018.
- [3] F. Amigoni, J. Banfi, and N. Basilico, "https://ieeexplore.ieee.org/abstract/document/8267592Multirobot exploration of communication-restricted environments: A survey," *IEEE Intelligent Systems*, vol. 32, no. 6, pp. 48–57, 2017.
- [4] M. Guo and M. M. Zavlanos, "Multirobot data gathering under buffer constraints and intermittent communication," *IEEE transactions on robotics*, vol. 34, no. 4, pp. 1082–1097, 2018.
- [5] F. Klaesson, P. Nilsson, T. S. Vaquero, S. Tepsuporn, A. D. Ames, and R. M. Murray, "Planning and optimization for multi-robot planetary cave exploration under intermittent connectivity constraints," 2020.
- [6] E. Stump, A. Jadbabaie, and V. Kumar, "https://ieeexplore.ieee.org/abstract/document/4543418Connectivity management in mobile robot teams," in *IEEE Int. Conf. Robot. Autom. (ICRA)*. IEEE, 2008, pp. 1525–1530.
- [7] P. Robuffo Giordano, A. Franchi, C. Secchi, and H. H. Bühlhoff, "https://journals.sagepub.com/doi/abs/10.1177/0278364912469671A passivity-based decentralized strategy for generalized connectivity maintenance," *Int. Journal Robotics Research*, vol. 32, no. 3, pp. 299–323, 2013.
- [8] "Multi-robot exploration under the constraints of wireless networking."
- [9] C. Dixon and E. W. Frew, "Maintaining optimal communication chains in robotic sensor networks using mobility control," *Mobile Networks and Applications*, vol. 14, no. 3, pp. 281–291, 2009.
- [10] Y. Yan and Y. Mostofi, "Robotic router formation in realistic communication environments," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 810–827, 2012.
- [11] O. Tekdas, W. Yang, and V. Isler, "Robotic routers: Algorithms and implementation," *The Int. Journal of Robotics Research*, vol. 29, no. 1, pp. 110–126, 2010.

- [12] H. Rovina, T. Salam, Y. Kantaros, and M. A. Hsieh, "Asynchronous adaptive sampling and reduced-order modeling of dynamic processes by robot teams via intermittently connected networks," in *2020 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4798–4805.
- [13] Y. Kantaros, M. Guo, and M. M. Zavlanos, "Temporal logic task planning and intermittent connectivity control of mobile robot networks," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4105–4120, 2019.
- [14] R. Aragues, D. V. Dimarogonas, P. Guallar, and C. Sagues, "Intermittent connectivity maintenance with heterogeneous robots," *IEEE Transactions on Robotics*, vol. 37, no. 1, pp. 225–245, 2020.
- [15] L. Clark, J. Galante, B. Krishnamachari, and K. Psounis, "https://ieeexplore.ieee.org/abstract/document/9361138A Queue-Stabilizing Framework for Networked Multi-Robot Exploration," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2091–2098, 2021.
- [16] Y. Pei, M. W. Mutka, and N. Xi, "https://onlinelibrary.wiley.com/doi/abs/10.1002/wcm.1145Connectivity and bandwidth-aware real-time exploration in mobile robot networks," *Wireless Communications and Mobile Computing*, vol. 13, no. 9, pp. 847–863, 2013.
- [17] Y. Mostofi, M. Malmirchegini, and A. Ghaffarkhah, "Estimation of communication signal strength in robotic networks," in *2010 IEEE Int. Conf. on Robotics and Automation*. IEEE, 2010, pp. 1946–1951.
- [18] L. Clark, J. A. Edlund, M. S. Net, T. S. Vaquero, and A.-a. Agha-Mohammadi, "Propem-l: Radio propagation environment modeling and learning for communication-aware multi-robot exploration," in *Robotics: Science and Systems (RSS)*, 2022.
- [19] V. Spirin, S. Cameron, and J. De Hoog, "Time preference for information in multi-agent exploration with limited communication," in *Conf. Towards Autonomous Robotic Systems*. Springer, 2013, pp. 34–45.
- [20] A. Gasparri and B. Krishnamachari, "Throughput-optimal robotic message ferrying for wireless networks using backpressure control," in *2014 IEEE 11th Int. Conf. on Mobile Ad Hoc and Sensor Systems*. IEEE, 2014, pp. 488–496.
- [21] S. Profanter, A. Tekat, K. Dorofeev, M. Rickert, and A. Knoll, "Opca versus ros, dds, and mqtt: Performance evaluation of industry 4.0 protocols," in *2019 IEEE Int. Conf. on Industrial Technology (ICIT)*, 2019, pp. 955–962.
- [22] K. S. Sikand, L. Zartman, S. Rabiee, and J. Biswas, "Robofleet: Secure open source communication and management for fleets of autonomous robots," *arXiv preprint arXiv:2103.06993*, 2021.
- [23] N. Hudson, F. Talbot, M. Cox, J. Williams, T. Hines, A. Pitt, B. Wood, D. Frousheger, K. L. Surdo, T. Molnar *et al.*, "Heterogeneous ground and air platforms, homogeneous sensing: Team csiro data61's approach to the darpa subterranean challenge," *arXiv preprint arXiv:2104.09053*, 2021.
- [24] M. T. Ohradzansky, E. R. Rush, D. G. Riley, A. B. Mills, S. Ahmad, S. McGuire, H. Biggie, K. Harlow, M. J. Miles, E. W. Frew *et al.*, "Multi-agent autonomy: Advancements and challenges in subterranean exploration," *arXiv preprint arXiv:2110.04390*, 2021.
- [25] K. Otsu, S. Tepsuporn, R. Thakker, T. S. Vaquero, J. A. Edlund, W. Walsh, G. Miles, T. Heywood, M. T. Wolf, and A.-A. Agha-Mohammadi, "Supervised autonomy for communication-degraded subterranean exploration by a robot team," in *2020 IEEE Aerospace Conf.* IEEE, 2020.
- [26] M. Tranzatto, F. Mascarich, L. Bernreiter, C. Godinho, M. Camurri, S. M. K. Khattak, T. Dang, V. Reijgwart, J. Loeje, D. Wisth *et al.*, "Cerberus: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the darpa subterranean challenge," *Journal of Field Robotics*, 2021.
- [27] M. F. Ginting, K. Otsu, J. A. Edlund, J. Gao, and A.-A. Agha-Mohammadi, "Chord: Distributed data-sharing via hybrid ros 1 and 2 for multi-robot exploration of large-scale complex environments," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5064–5071, 2021.
- [28] A. Agha, K. Otsu, B. Morrell, D. D. Fan, R. Thakker, A. Santamaria-Navarro, S.-K. Kim, A. Bouman, X. Lei, J. Edlund *et al.*, "Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge," *arXiv preprint arXiv:2103.11470*, 2021.
- [29] N. Funabiki, B. Morrell, J. Nash, and A.-a. Agha-mohammadi, "Range-aided pose-graph-based slam: Applications of deployable ranging beacons for unknown environment exploration," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 48–55, 2020.
- [30] T. Schwengler, "Wireless & cellular communications, version 3.9," *Telecommunication Systems Laboratory, Colorado, USA.*, vol. 2, 2016.
- [31] D. D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick, and A.-A. Agha-Mohammadi, "Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation," *arXiv preprint arXiv:2103.02828*, 2021.
- [32] "Traversability-aware signal coverage planning for communication node deployment in planetary cave exploration."
- [33] S.-K. Kim*, A. Bouman*, G. Salhotra, D. D. Fan, K. Otsu, J. Burdick, and A.-a. Agha-mohammadi, "PLGRIM: Hierarchical value learning for large-scale exploration in unknown environments," in *Proceedings of the Int. Conf. on Automated Planning and Scheduling*, vol. 31, 2021, pp. 652–662.