

Safe Coverage for Heterogeneous Systems With Limited Connectivity

Annalisa T. Taylor , Thomas A. Berrueta , *Member, IEEE*, Allison Pinosky ,
and Todd D. Murphey , *Senior Member, IEEE*

Abstract—Operating multi-robot teams of diverse agents is an ongoing challenge for emergency deployments, where inter-agent connectivity is rare and environments are unpredictable. Heterogeneous systems must be capable of adapting autonomously while maintaining safety. Here, we develop an algorithm for heterogeneous decentralized multi-robot systems to independently manage safety constraints with provable guarantees for safety and communication in a coverage task. We demonstrate this algorithm in settings where up to 100 agents navigate a simulated cluttered environment with safety constraints that change as agents observe hazards. Further, we show that the performance of a system with a largely disconnected network is equivalent to a fully connected communication network, suggesting that treating connectivity as a constraint may be unnecessary with an appropriate control strategy.

Index Terms—Distributed robot systems, robot safety.

I. INTRODUCTION

IN HIGH-RISK settings like disaster response, heterogeneous robotic systems must autonomously manage basic functions while ensuring the safety of both themselves and bystanders [1]. Inter-agent communication can be sporadic in these environments and maintaining a connected communication network is often infeasible [2], [3]. Therefore, control algorithms for heterogeneous multi-robot systems in emergency applications must guarantee safety even in the case of predominantly disconnected communication networks.

We develop an algorithm for heterogeneous, decentralized multi-robot systems with limited connectivity to manage safety constraints in coverage tasks, which are critical in disaster response efforts such as search and rescue. To demonstrate this method, we simulate a heterogeneous robotic system with up to 100 agents that must conform to safety constraints that change

due to the observations of other agents. We provide guarantees that agents using this controller stay safe with both static obstacles and safety constraints that evolve over time. Additionally, we explore the effects of both communication network topology and degradation on system performance and safety when connections are sparse and infrequent. To compensate for limited communication networks, agents can plan actions to achieve successful environmental coverage in a decentralized manner. We prove that agents using this ergodic control algorithm are guaranteed to communicate, even with different coverage objectives. Through an empirical study, we find, surprisingly, that communication network connectivity does not determine task success. We show that, even with small communication radii, the performance of a system with a generally disconnected network is equivalent to a fully connected communication network. Our results indicate that intermittent and rarely connected networks are sufficient for multi-robot systems to collaborate, demonstrating that highly capable communication networks may not be necessary for the success of multi-robot systems. The contributions of this work are summarized below.

- 1) We develop an algorithm for decentralized safe ergodic coverage for heterogeneous robotic collectives that enables environmental coverage under stringent communication constraints. We simulate our algorithm for heterogeneous systems of up to 100 agents.
- 2) We provide a guarantee that agents using our algorithm will be safe and ergodic with respect to a coverage task and prove that agents will communicate.
- 3) We model communication between heterogeneous agents with the modified two-ray path loss model, which includes ground reflection and transmission altitude, to demonstrate effective system coverage performance in real-world environments.

II. RELATED WORK

A. Multi-Robot Coverage and Task Specifications

This work addresses the challenge of heterogeneous multi-agent spatial coverage in safety-critical environments with unreliable communication, such as disaster response scenarios [4]. Coverage specifications such as temporal logic formulas [5], hierarchical mixed-initiative specifications [6], and compositional motion primitives [7] have been used for multi-agent control. However, many of these methods require prespecifying agent

Received 10 April 2024; accepted 17 August 2024. Date of publication 6 September 2024; date of current version 12 September 2024. This article was recommended for publication by Associate Editor L. Liu and Editor M. A. Hsieh upon evaluation of the reviewers' comments. This work was supported in part by US Army Research Office MURI under Grant W911NF-19-1-0233, in part by the National Science Foundation under Grant CNS-2229170, and in part by the US Office of Naval Research under Grant N00014-21-1-2706. (*Corresponding authors: Annalisa T. Taylor; Todd D. Murphey.*)

The authors are with the Center for Robotics and Biosystems, Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208 USA (e-mail: annalisa.taylor@u.northwestern.edu; t-murphey@northwestern.edu).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2024.3455804>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2024.3455804

trajectories, making them inadequate for dynamic environments. Common coverage solutions, like cellular decomposition algorithms, struggle with adaptation to changing conditions and risk coverage gaps due to agent dropout [8], [9], [10]. Moreover, they implicitly impose strict communication demands as agents must communicate between cells for collaboration to occur.

Our work is motivated by applications where agents must cover unknown environments with sporadic communication. Thus, we select a coverage algorithm that is robust to agent dropout and does not rely on inter-agent connectivity. For this reason, we use the ergodic control algorithm which generates actions so that time spent in an area of a distribution is proportional to its value in that area [11]. This allows the ergodic controller to cover a given task distribution without an additional metric to segment the coverage domain. Previously, we used an ergodic control algorithm with a multi-robot system during experimental field tests where agents shared their trajectory history to collaborate on coverage tasks [12], [13]. We extend prior work with decentralized ergodic control to dynamic networks with sparse connectivity. This improves the robustness of the control architecture for communication constraints typical of real-world conditions [12], [14].

B. Safety Specifications

Reachability analysis and reactive synthesis provide formal guarantees for ensuring safety but face scalability issues with increasing system states [15], [16]. Work in [17], [18] used repulsive vector fields for safety with ergodic control, but relied on a centralized controller and did not ensure agents' safety against various obstacle types. Here we employ control barrier functions (CBFs) to express safety constraints, a formalism compatible with optimal control frameworks. CBFs enforce the forward invariance of sets in the evolution of agent dynamics [19], confining the agent to a safe subset of its operation domain. CBFs have been used for safety in diverse tasks such as constraining the location of a robotic arm end-effector [20] and avoiding collisions with experimentally-deployed agents [21]. Recent work explored the use of the ergodic control algorithm with CBFs [22], [23]. We generalize this work to include heterogeneous multi-robot systems, and contribute theoretical guarantees on the safety and communication of the agents in the system. Instead of assuming a connected communication network, we examine the role of communication topology and its effect on safety and performance. This is an aspect of deploying multi-robot systems in safety critical conditions that has yet to be explored for this class of coverage algorithms.

III. PRELIMINARIES

A. Decentralized Ergodic Control

Next, we discuss the decentralized ergodic control algorithm for coverage of spatial distributions for multi-robot systems [14]. To develop an optimal controller, we assume the system dynamics of N agents are in control-affine form, controllable and

independent:

$$\begin{aligned} \dot{x} &= f(x) + g(x)u \\ &= \begin{bmatrix} f_1(x_1) \\ \vdots \\ f_N(x_N) \end{bmatrix} + \begin{bmatrix} g_1(x_1) & \dots & 0 \\ \vdots & \ddots & \\ 0 & & g_N(x_N) \end{bmatrix} u \end{aligned} \quad (1)$$

where the state of a robotic agent at every point in time $t \in \mathbb{R}^+$ is $x(t) \in \mathbb{R}^n$ and the control is $u(t) \in [u_{min}, u_{max}] \subset \mathbb{R}^m$. Here, $f(x) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g(x) : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are vector fields for the system's free dynamics and response to control inputs. The trajectories of the multi-agent system take the form $x(t) = [x_1(t), \dots, x_N(t)] \in \mathbb{R}^{nN}$. The controller is a receding-horizon optimal control algorithm that calculates actions in time for agents over closed subsets of \mathbb{R}^+ ; we have $t \in [t_i, t_i + \tau]$ for some initial time t_i corresponding to the i^{th} sampling time and time horizon $\tau > 0$. We describe an agent's dynamics as *ergodic* if they satisfy Birkhoff's point-wise ergodic theorem in the sense of [24].

Definition 1: For a collective of N agents with system dynamics given by (1), we say their trajectories $x_i(\cdot) : [0, t] \rightarrow \mathbb{R}^n$, $\forall i \in \{1, \dots, N\}$ are **ergodic** with respect to a distribution p over a task set \mathcal{T} if

$$\lim_{t \rightarrow \infty} d^t(x) = \bar{p}(x) \quad (2)$$

for all $x \in \mathcal{T}$, where $\bar{p}(x)$ is the measure of a spherical set $\mathcal{B}(x, \epsilon)$ centered at x for some $\epsilon > 0$ such that $\bar{p}(x) = \int_{\mathcal{T}} p(y) \mathbf{1}_{\mathcal{B}(x, \epsilon)}(y) dy$, and

$$d^t(x) = \frac{1}{Nt} \sum_{i=1}^N \int_0^t \mathbf{1}_{\mathcal{B}(x, \epsilon)}(x_i(\tau)) d\tau \quad (3)$$

is the fraction of time spent in $\mathcal{B}(x, \epsilon)$.

The task set \mathcal{T} is the set of all states in the coverage domain. The "task distribution" is a density over the states in \mathcal{T} , which encodes the agents' coverage objective (e.g., explore near occluded areas in the domain). For an agent to be ergodic with respect to a task distribution, the amount of time an agent spends in a given region of space must be proportional to the measure of that region. For clarity, sets are notated with calligraphic letters in this work.

Next, we use a metric on ergodicity to compare the statistics of agent trajectories with a task distribution over the coverage space [24]. This metric is evaluated in the space of spatial Fourier coefficients. The Fourier coefficients for the statistics of the trajectory $x(t)$ are $c_k = \frac{1}{\tau} \int_{t_i}^{t_i + \tau} \tilde{F}_k(x(t)) dt$, where $\tilde{F}_k(x(t)) = \frac{1}{N} \sum_j F_k(x_j(t))$ and $F_k(x)$ are cosine basis functions [14]. We also require a Fourier representation of the task distribution given by $\phi_k = \int_{\mathcal{T}} \phi(x) F_k(x) dx$ where $\phi(x)$ is the measure of the agents' task set $\mathcal{T} \subset \mathbb{R}^n$ at x . With these representations, we write an ergodic metric:

$$E(x(t)) = \sum_{k \in \mathbb{N}^n} \Lambda_k (c_k - \phi_k)^2 \quad (4)$$

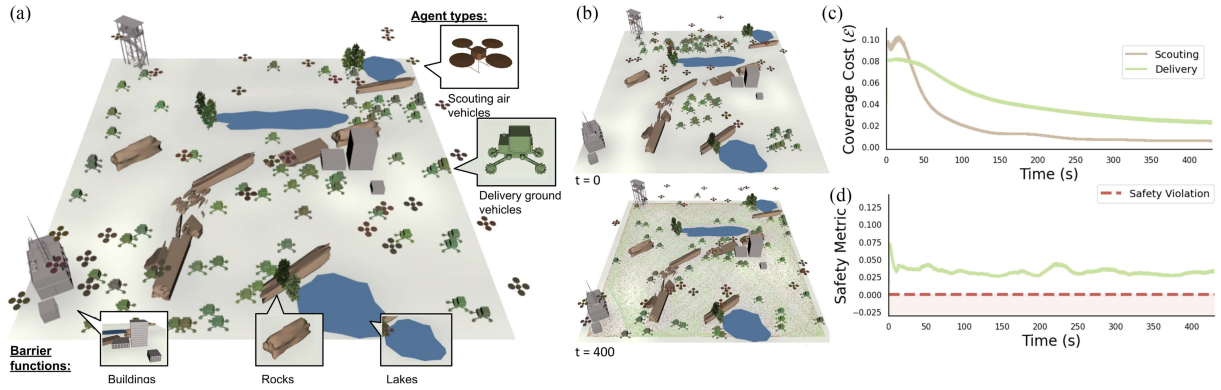


Fig. 1. Results with 100 agents. (a) The decentralized heterogeneous multi-robot system covers an environment with obstacles. (b) Two snapshots of the robots’ trajectories at the beginning and end of a trial. Air vehicles (scouting agents), cover the environment unrestricted while ground vehicles (delivery agents) must stay in the area that scouting agents have observed. Delivery agents must avoid obstacles as they cover their expanding safe area. (c) The ergodic metric (coverage task cost) decreases over the trial length. Delivery agents’ cost decreases slower and remains higher as they are restricted to the safe parts of the domain, but both minimize the ergodic metric. (d) The delivery agents’ safety metric stays above zero (does not violate safety constraints). The safety metric is related to the distance from intersecting with the closest unsafe area.

where Λ_k is a normalization coefficient ensuring the well-posedness of (4). With a method for provably generating ergodic coverage, we turn to safety.

B. Control Barrier Functions

The definition of safety in this work is determined by the properties of CBFs, which confine a dynamical system to the interior of a “safe set” through the property of forward invariance: any system starting in a safe set cannot reach the exterior of that set [19]. We define the safe set \mathcal{S} as follows:

Definition 2: Given a set $\mathcal{G} \subset \mathbb{R}^n$ and a continuously differentiable function $h : \mathcal{G} \rightarrow \mathbb{R}$, the **safe set** \mathcal{S} is a superlevel set such that

$$\begin{aligned} \mathcal{S} &= \{x \in \mathcal{G} : h(x) \geq 0\} \\ \delta\mathcal{S} &= \{x \in \mathcal{G} : h(x) = 0\} \\ \text{Int}(\mathcal{S}) &= \{x \in \mathcal{G} : h(x) > 0\} \end{aligned} \quad (5)$$

where $\delta\mathcal{S}$ is the boundary of set \mathcal{S} and $\text{Int}(\mathcal{S})$ is the interior of set \mathcal{S} . When \mathcal{S} is compact and forward invariant for the dynamics in (1), we say the system is safe with respect to \mathcal{S} . We now define the CBF $h(x)$ in terms of the forward invariance of \mathcal{S} [19].

Definition 3: If $\mathcal{S} \subset \mathcal{A} \subset \mathbb{R}^n$ is the superlevel set of a continuously differentiable function $h : \mathcal{A} \rightarrow \mathbb{R}$, then h is a **control barrier function** for the dynamics in 1 if there exists a class \mathcal{K}_∞ function $\alpha(\cdot)$ such that

$$L_f h(x) + L_g h(x)u \geq -\alpha(h(x)) \quad (6)$$

for all $x \in \mathcal{A}$. L_f and L_g are the Lie derivatives of $h(x)$ with respect to f and g from (1).

C. Problem Setup

For the remainder of the letter, we consider a heterogeneous robot system with two distinct groups of agents. Each uses a decentralized ergodic controller with CBFs for safety. Consistent with the structure of disaster relief operations, we use air and ground vehicles as our two agent types (Fig. 1). Air vehicles are designed to scout and map affected areas while ground vehicles

deliver heavy supplies. Therefore, we refer to the air and ground vehicles as “scouting” and “delivery” agents respectively. Due to their different capabilities and sensors, each agent type has distinct safe sets, with scouting agents being more capable and fewer in number. As a result, we assume that the scouting agents’ safe area encompasses the delivery agents’ safe area. This allows the scouting agents to observe which areas of the environment are safe for delivery agents. Lastly, scouting agents broadcast the coordinates of environmental hazards in a radius around their position, enabling delivery agents to update their safe sets when they come close enough to receive messages. Delivery agents consider areas unsafe by default until they receive updates from scouting agents which gradually expands their safe sets, though static obstacles are always inaccessible.

IV. GUARANTEES ON SAFETY AND COMMUNICATION

In this section, we provide formal guarantees on the safety of the system described in Section III-C and prove that our controller ensures inter-agent communication. Our proofs address three distinct sources of heterogeneity. First, the two groups of agents may have different dynamics (e.g., ground vs. air vehicles). Second, scouting agents influence delivery agents: scouting agents send information about safe and unsafe regions of the task domain to delivery agents, but not vice versa. Third, agents can have different task distributions (i.e., different coverage objectives $\phi(x)$). To keep track of these differences, in what follows we let \mathcal{A}_S represent a set of scouting agents and \mathcal{A}_D a set of delivery agents. While both groups of agents share the same task set \mathcal{T} , their different safety requirements will result in different safe task sets. We will refer to the intersection of an agent’s task set with its safe set as its “effective” task set.

A. Safe Coverage

With ergodic control algorithms, safety violations are inevitable if any unsafe region of the agents’ task set has a non-zero measure. Therefore, we need a different method to construct safety guarantees. Here, we show that a system in the form of (1) will be safe and ergodic using our algorithm. When the task set is connected and contained in the interior of the safe set,

agents using an ergodic controller are guaranteed to be safe and ergodic if the system is initialized in a non-zero measure set of the task set. However, the safe area is often unknown *a priori* or changes over time and the task may not be contained in the safe set. The forward invariance of the safe set prevents a safe ergodic controller from generating trajectories that optimize the ergodic metric over the task set asymptotically. Instead, we show that agents can be safe and ergodic with respect to an effective task set contained in the safe area.

Next, we define safe sets and task sets for both agent types. For delivery agents, the static set of safe areas in the environment is $\mathcal{S}_{E,D}$ (in Fig. 1(a), this would be the rocks and lakes that are permanently inaccessible). Areas unobserved by the scouting agents are unsafe by default. Let $\mathcal{B}_S(t) = \mathcal{B}(x_S(t), r_c)$ be the set of points around scouting agent trajectories $x_S(t)$ with radius r_c . Choose r_c such that it contains the states reachable by an agent with locally Lipschitz dynamics in some dt to ensure that the union of sets $\mathcal{B}_S(t)$ is connected for all t . Then, we define

$$\mathcal{C}_S(t) = \begin{cases} \mathcal{B}_S(t) & x_D(t) \in \mathcal{B}_S(t) \\ \emptyset & \text{else} \end{cases}$$

where \emptyset is the empty set and $x_D(t)$ are the delivery agents' trajectories. The delivery agents' safe set at time t will be

$$\mathcal{S}_D(t) = \left[\bigcup_{i=0}^t \mathcal{C}_S(i) \right] \cap \mathcal{S}_{E,D}. \quad (7)$$

Since $\mathcal{S}_D(t)$ evolves according to $x_S(t)$, scouting agents specify the safe area of the task set for delivery agents based on their observations and $\mathcal{S}_D(t)$ grows over time. We next define an effective task set for the delivery agents,

$$\mathcal{T}_{eff}^D(t) = \mathcal{T} \cap \mathcal{S}_D(t), \forall t. \quad (8)$$

The scouting agents' safe set is \mathcal{S}_S with effective task set

$$\mathcal{T}_{eff}^S = \mathcal{T} \cap \mathcal{S}_S. \quad (9)$$

We now establish that all delivery agents will converge to the same effective task set. To do this, scouting and delivery agents must be sufficiently close at the same time for the delivery agents' to update their safe set. We say agents will "cross paths" when they come within r_c of one another.

Lemma 1: Given that the trajectories $x_D(t)$ and $x_S(t)$ of the joint, decentralized system, are each independently ergodic with respect to $\mathcal{T}_{eff}^D(t)$ and \mathcal{T}_{eff}^S , the limit of the delivery agents' effective task set is the following:

$$\lim_{t \rightarrow \infty} \mathcal{T}_{eff}^D(t) = \mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D} \quad (10)$$

Proof: First, note that according to its definition, $\mathcal{T}_{eff}^D(t)$ only updates in time when $x_D(t)$ and $x_S(t)$ cross paths at some time t' , which occurs at discrete intervals in time. In between these intervals, $x_D(t)$ is ergodic with respect to $\mathcal{T}_{eff}^D(t')$, and $x_D(t)$ is guaranteed to visit every non-zero measure set of $\mathcal{T}_{eff}^D(t')$ as $t \rightarrow \infty$. Next, take $[x_S(t), x_D(t)]$ to be the states of the joint system. Given each system is independently ergodic, the joint system will be ergodic with respect to the product measure of $\mathcal{T}_{eff}^S \times \mathcal{T}_{eff}^D(t')$ [25], specifically implying that $x_S(t)$ and

$x_D(t)$ are guaranteed to be within r_c of each other as $t \rightarrow \infty$. Therefore, the set $\mathcal{T}_{eff}^D(t')$ will update an infinite number of times. Moreover, when $x_D(t)$ and $x_S(t)$ communicate at time t' , $\mathcal{T}_{eff}^D(t) \subseteq \mathcal{T}_{eff}^D(t')$, with equality once $x_S(t)$ have detected all safe states (guaranteed by the ergodicity of $x_S(t)$). ■

With these definitions and the safety properties discussed in Section III, Lemma 1 provides the only extra information needed to collect these formal properties into the following theorem. We know that \mathcal{A}_S and \mathcal{A}_D are constrained by CBFs which confine agents to safe sets \mathcal{S}_S and $\mathcal{S}_D(t)$ at every point in time. We assume that all agents in these sets are initialized in non-zero measure sets of the spatial domain. By definition, the exteriors of \mathcal{T}_{eff}^S and $\mathcal{T}_{eff}^D(t)$ are zero. Then, agents in \mathcal{A}_S are safe and forward invariant in \mathcal{T}_{eff}^S for all $x \in \text{Int}(\mathcal{S}_S)$, and agents in \mathcal{A}_D are guaranteed to be safe and forward invariant in $\mathcal{T}_{eff}^D(t)$ for all $x \in \text{Int}(\mathcal{S}_D(t)), \forall t$. From Lemma 1, we know that $\mathcal{T}_{eff}^D(t)$ converges to $\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}$ as $t \rightarrow \infty$. Then, delivery agents will be ergodic with respect to $\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}$, and scouting agents will be ergodic with respect to \mathcal{T}_{eff}^S , resulting in the theorem below.

Theorem 1: Let \mathcal{A}_S and \mathcal{A}_D be sets of scouting and delivery agents, each with ergodic optimal controllers as in (4). Additionally, let \mathcal{T}_{eff}^S and $\mathcal{T}_{eff}^D(t)$ be their respective task sets. Then, \mathcal{A}_S and \mathcal{A}_D will be safe with respect to the safe sets \mathcal{S}_S and $\mathcal{S}_D(t)$ for all t and ergodic with respect to the task sets \mathcal{T}_{eff}^S and $\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}$ as $t \rightarrow \infty$.

Note that the topology of the task set is key. If the task set is connected all pairs of points are reachable and the ergodic metric can be minimized. This assumption can be relaxed if all connected components of the task set contain at least one agent in their interior at initialization.

B. Communication

For multi-robot systems, communication between agents enables decentralization of objectives. In this case, communication is essential to the success of both groups of agents: if agents do not periodically communicate their trajectory histories within their group, the assumptions of decentralized ergodic control are violated and task success is jeopardized. We now prove that any two in-group agents using the decentralized ergodic control algorithm are guaranteed to communicate if their paths cross (or their communication radii, r_c , are greater than zero).

Again, the analysis of scouting agents is simpler because \mathcal{S}_S is not time-varying. Therefore we can rely on their ergodicity to asymptotically guarantee that their paths will cross. Heterogeneity presents a challenge for the analysis of delivery agents, whose effective task sets are time-varying. However, we know that delivery agents are asymptotically ergodic with respect to the same limiting effective task set, so we can guarantee that their paths will cross asymptotically, ensuring their communication.

Theorem 2: For any $r_c > 0$, let $\{a_1, a_2\} \subset \mathcal{A}_S$ and $\{a_3, a_4\} \subset \mathcal{A}_D$ be any pair of optimal ergodic agents (as in (4)) with respect to the measures of \mathcal{T}_{eff}^S and $\mathcal{T}_{eff}^D(t)$, respectively. Then, agent $\{a_1, a_2\}$ are asymptotically guaranteed to communicate as $t \rightarrow \infty$. Additionally, $\{a_3, a_4\}$ are asymptotically guaranteed to communicate as $t \rightarrow \infty$.

Proof: For any pair of agents to communicate, they must be simultaneously co-located (up to r_c) at some point in time. Let $\{a_1, a_2\} \subset \mathcal{A}_S$ with agent states $x_1(t), x_2(t) \in T \forall t$ be independently ergodic with respect to the same task set, \mathcal{T}_{eff}^S . Then, the joint system, with states $[x_1(t), x_2(t)]$, is ergodic with respect to the Cartesian product $\mathcal{T}_{eff}^S \times \mathcal{T}_{eff}^S$ and agents will cross paths. Next, agents $\{a_3, a_4\} \subset \mathcal{A}_D$ with states $x_3(t), x_4(t) \in T \forall t$ are independently ergodic with respect to $\mathcal{T}_{eff}^D(t)$, which varies in time. Per Lemma 1, we know that $\mathcal{T}_{eff}^D(t)$ converges to $\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}$ as $t \rightarrow \infty$. By construction, scouting agents are capable of efficient ergodic coverage and $|\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}|$ is finite. Thus we assume the limit of $\mathcal{T}_{eff}^D(t)$ converges in finite time. Then, agents a_3, a_4 with joint states $[x_3(t), x_4(t)]$ will be ergodic with respect to $\mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D} \times \mathcal{T}_{eff}^S \cap \mathcal{S}_{E,D}$. This guarantees their trajectories will cross paths and that they will communicate. ■

V. SAFE DECENTRALIZED ERGODIC COVERAGE

Coverage of unknown environments is common in disaster response and requires adaptable control strategies. Maps of an area before a disaster may be outdated while aftershocks or flooding may cause ongoing harm. This informs our use of adaptive safety constraints and coverage specifications.

A. Simulation Setup

The simulated system implements the setup from Section III-C, where delivery agents' safe areas depend on scouting agents' observations. Each agent type has a distinct task distribution: scouting agents cover the full domain, while delivery agents focus on areas near buildings seen in Fig. 1 to deliver supplies. This is represented as a 2-D spatial distribution over the coverage domain states, noted as $\phi(x)$ in Section III-A. Despite different safety constraints and objectives, both agent types generate trajectories using the controller described in Section III-A with an additional safety optimizer. Agents account for obstacles in a small radius from their position, meaning the safety optimization only affects the agents' trajectories when obstacles are detected. This optimization applies domain boundary conditions for all agents, but environmental debris restricts only delivery agents. Coordinates of unsafe areas are updated in every control loop, using OpenCV for image processing to generate CBFs. When modeling dynamic communication networks, we determine connectivity first with a distance only criterion and later using the modified two-ray path loss model [26], [27]. More detail can be found in Algorithm 1. The problem setup presented here also applies to results Section VI.

B. Scalability

To simulate 50–100 agents in parallel, we used an Intel Xeon(R) Platinum 8380 CPU at 2.30 GHz \times 160 server with 4TB RAM and 500 GB of disk memory running Ubuntu 18.04 and ROS Melodic. We also confirmed that our algorithm runs in real-time on a laptop as indicated by tests where a single agent runs in a 40 agent environment using 20% of the laptop's CPU

Algorithm 1: Safety-Constrained Ergodic Control.

- 1: **Initialize Agent** j : type label $l \in \{A_S, A_D\}$, initial state $x_j(0)$, dynamics \dot{x}_j , target distribution $\phi_{k,j}$, trajectory history $c_{k,j}$, safe set \mathcal{S}_j , barrier functions $h_j(x_j)$, time horizon τ , received message $M_r = (l_r, c_{k,r}, \mathcal{S}_r)$
 - 2: **while** task ongoing **do**
 - 3: Calculate $u_{erg,j} = \mathbf{ErgCtrl}(x_j(t), \dot{x}_j, \phi_{k,j}, c_{k,j}, \tau)$
 - 4: Update $h_j(x_j) = \mathbf{CtrlBarrFun}(\mathcal{S}_j, x_j(t)) \triangleright \mathbf{Safety}$
 - 5: Minimize $\|u_{erg,j} - u_{safe,j}\|$ subject to $h_j(x_j)$
 - 6: Apply control $u_{safe,j}$
 - 7: Broadcast $(l_j, c_{k,j}, \mathcal{S}_j) \triangleright \mathbf{Communication}$
 - 8: $M_r \leftarrow (l_r, c_{k,r}, \mathcal{S}_r) = \mathbf{Listen}()$
 - 9: **if** M_r received **then**
 - 10: **if** l_j is A_D **then** $\mathcal{S}_j \leftarrow \mathcal{S}_j \cup \mathcal{S}_r$ **end if**
 - 11: **if** l_j is same as l_r **then** store $c_{k,r}$ **end if**
 - 12: Average $c_{k,j}$ and stored $c_{k,r}$
 - 13: **end if**
 - 14: **end while**
-

and 1GB of RAM. The laptop's GPU was not utilized. The server simulated experimental conditions by creating independent instances for each agent, running their algorithms simultaneously, simulating the environment, and modeling the communication network for the agents all on one machine. Notably, this does not reflect the requirements of individual agents in a distributed hardware implementation. In this case, an agent would manage its own computations and communication locally, reducing the overall memory demands. The above test could be executed on the field test hardware from [12], [13].

C. Results: Coverage With 100 Agents

Fig. 1 shows results with 100 agents covering the cluttered environment described above, with 30 scouting agents and 70 delivery agents. Scouting agents freely cover the environment and broadcast safe areas to the delivery agents. At the start of the trial, each delivery agent is confined to a small area around its initial position. As scouting agents cover the environment, the area that they have observed (in this case, an area within a specified radius around each trajectory point) is communicated as safe or unsafe to the delivery agents, expanding the safe areas of the map, shown in Fig. 1(a). The delivery group navigates the irregular and dynamically growing safe space while maintaining safety (Fig. 1(c)). Simultaneously, the scouting group provides persistent coverage of the entire environment. The coverage cost of the delivery group is expected to stay higher than the scouting group, Fig. 1(b), because the delivery group cannot reach the permanently unsafe areas of the map (e.g., inside obstacles).

D. Results: Static Network Topologies

For simplicity, we consider a 10 agent collective with static, connected communication networks. We investigate four standard communication topologies: a star, two distinct clusters, a chain, and a fully connected network. Each stays connected during the coverage task as in Fig. 2(a). The nodes of the

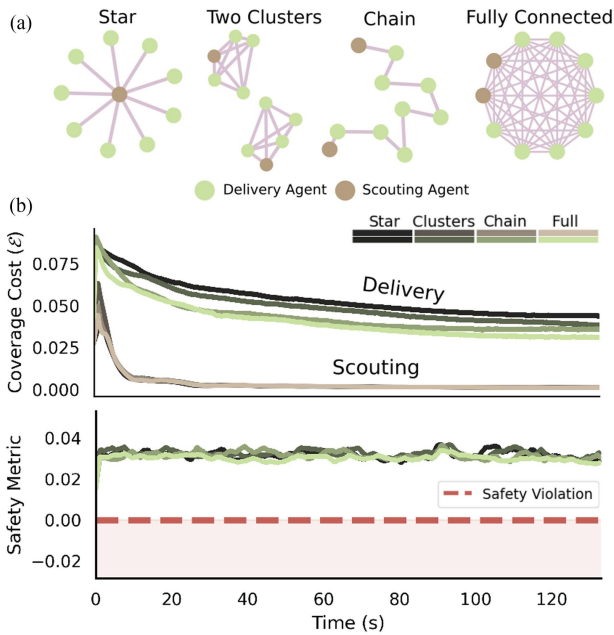


Fig. 2. Comparing static communication topologies. (a) We test four standard communication graphs. Edges indicate communication. Delivery agents obey safety constraints from scouting agents' observations. (b) (upper) Average coverage cost (ergodic metric) over 10 simulation runs, with randomized initial conditions for 10 agents. Performance depends on communication topology, with the fully connected network performing best. Delivery agents still minimize the ergodic metric. Scouting agents show nearly equivalent performance between cases. (lower) The system stays safe.

communication graphs (Fig. 2(a)) are color coded to reflect the heterogeneous makeup of the system: brown for scouting agents, green for delivery agents. Fig. 2(b) shows the ergodic metric over 10 runs of the simulation with randomized initial conditions for each agent. System performance varies with topology, with the fully connected network performing the best. However, all cases show successful minimization of the ergodic metric. Note that scouting agents show nearly equivalent coverage performance between cases so their coverage cost averages overlap. As expected, the system did not violate any barrier functions throughout the trials, confirming that safety is maintained with different network topologies while the ergodic metric is minimized.

VI. SAFE COVERAGE WITH LIMITED COMMUNICATION

Designing robotic systems to depend on local information is key for emergency deployment, as reliable communication networks are often unrealistic [28]. Theorem 2 proved that communication is an inevitable consequence of decentralized ergodic control. To study the effect of limited (but guaranteed) communication on system performance, we explore dynamic communication network topologies for a coverage task. The simulation setup from Section V-A applies here.

A. Results: Distance-Based Dynamic Networks

We consider a dynamic communication network in a 50 agent system where agents share their trajectory history when they are within a specified distance from each other. We first explore distance-based communication, chosen to approximate

real-world environments in which signals decay as transmitters and receivers move further apart. Here agents can broadcast and receive information in a limited area around their positions. Scouting and delivery agents are both subject to these communication constraints with the size of the communication radius held constant between air-air, air-ground, and ground-ground communication. For this set of results, the altitude of the scouting agents was not considered when calculating the communication distance.

Given this communication model, we performed an empirical study of performance under varying communication radii. These radii were defined as a percentage of the domain width (here a finite, square domain from zero to one). A communication radius of 5% corresponds to a radius of 0.05 if the domain side length is one. Therefore the communication area for each agent is $\frac{0.05^2 \pi}{1^2} = 0.00785$. If two agents are placed uniformly randomly in the domain the probability that they would communicate is 0.0157. The results in Fig. 3(b)–(d) are the average of ten trials for each communication radius, with agents having uniformly random initial positions over the safe area in each trial. The communication graphs for the 5% and 100% cases are shown in Fig. 3(a). Our empirical study indicates that the 5% range case, though sparsely connected (Fig. 3(a), left), leads to similar coverage performance to the fully connected 100% range case (Fig. 3(a), right). This can be seen in the agents' average coverage cost (Fig. 3(b)), which for communication radii as low as 5% of the domain size converged to the performance of the fully connected network.

We used two measures of graph connectivity to assess the effect of different communication radii: graph density and algebraic connectivity. Many algorithms for multi-robot systems use spectral properties and density based metrics to assess connectedness of communication graphs [29]. Graph density is the ratio of current edges to the total possible number of edges in a graph, an average of which is shown in Fig. 3(b). Despite having close to 1% of the possible edges in the graph, the 5% communication radius performed as well as the 100% case, with the maximum amount of edges. The algebraic connectivity of the collective indicates that the 5% radius resulted in communication topologies that were disconnected throughout the entirety of the trial. We show the change in average coverage performance and algebraic connectivity with eight different communication radii in Fig. 3(d). As performance steeply increases, the algebraic connectivity stays at zero (i.e., the graph is disconnected for these trials). As performance levels off, the algebraic connectivity is positive and increasing, thus the communication graph is increasingly connected as the communication radius grows.

B. Results: Modified Two-Ray Path Loss Model

Next, we use the same 50 agent system whose communication network is now governed by the modified two-ray path loss model, an altitude and distance inclusive model suitable for communication between air and ground agents [26]. To simulate radio communication with our system, we used the Python package UAVradio [27], a comprehensive tool for modeling air-ground and air-air communications which contains several well

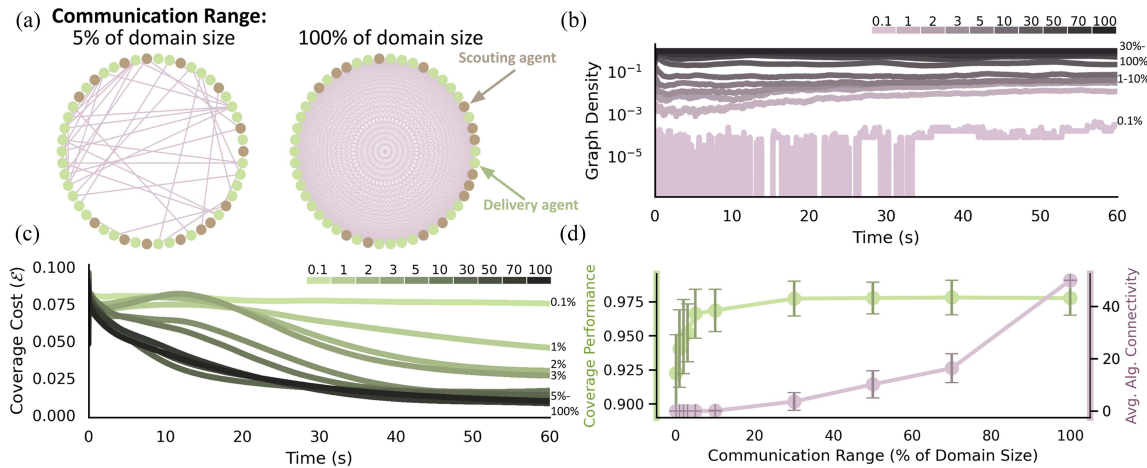


Fig. 3. Coverage performance with a dynamic communication network. (a) Communication networks with 50 agents at the end of two trials. Edges denote communication. A 5% range means an agent’s communication radius is 5% of the side length of the square coverage domain. (b) Average graph density over 10 trials for ten radii. Graph density measures sparsity of the network. The largest range enables all agents to communicate at all times. (c) Average coverage cost (ergodic metric) of 10 trials of delivery agents. All curves, except the 0.1% case, show the ergodic metric can be minimized. Curves from the 5% to 100% (fully connected) show approximately equivalent coverage performance. Notably, the 5% case has an order of magnitude less connections in its communication network when compared to the fully connected case. (d) As coverage performance changes rapidly, algebraic connectivity stays at zero (disconnected graph). As coverage performance levels off, algebraic connectivity shows the graph connectivity increases as the communication range grows. Error bars show standard deviation. Here, coverage is independent of connectivity for a wide range of communication radii.

established models. We account for transmission frequency, path loss, and both air and ground agent communication altitude with ground reflection. The scouting agents are at an altitude of 60 m, taken from the air vehicles used after Hurricane Harvey [30], while the delivery agents have an antenna height of 1m. For radio specifications we reference the Digi X-Bee S2C data sheet [31]. Our radio transmits at 1 dBm and 2.4 GHz with a receiver sensitivity of -100 dBm over a coverage area of 6.25 km². To determine if agents can communicate, we take the difference of the modified two-ray path loss and transmission power to find the received power. If the received power is greater than the receiver sensitivity, the two agents will communicate, subject to noise. We test several path loss exponents, factors that reflect path loss in different environments. We use path loss exponents that were determined experimentally [27] with additional extreme values to test our algorithm with higher loss than commonly documented. The model adds normally distributed zero mean noise with unit variance.

Results using the two-ray model are shown in Fig. 4 for 13 path loss exponents. The labeled path loss exponents in Fig. 4(a) were experimentally determined for various environments [32], [33], [34], [35], [36]. We see a similar pattern to Fig. 3: the system performs comparably in high and low loss environments. In Fig. 4(a), the coverage performance of the system increases while the algebraic connectivity shows that the network is disconnected on average. Once the algebraic connectivity begins to increase, coverage performance stays constant. The findings here and in Section VI-A suggest that maintaining a connected network may not be the most effective method to increase coverage performance for dynamic communication networks. Future algorithms for heterogeneous robots may consider reducing the importance of maintaining connected networks, instead focusing on an agent’s ability to act autonomously with local infrequent communication.

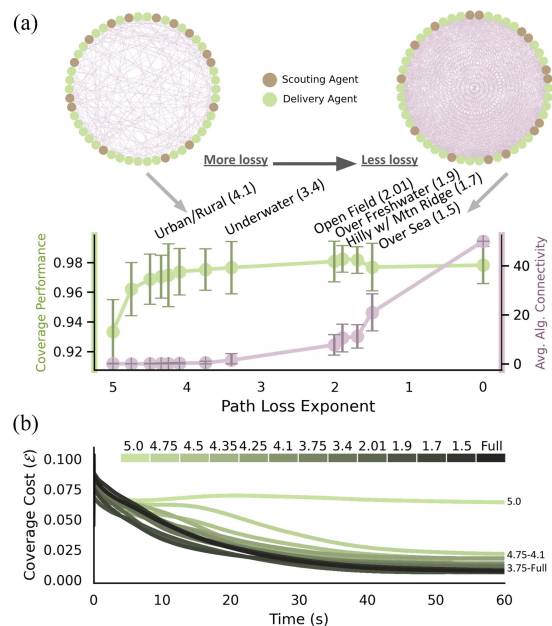


Fig. 4. Modified two-ray path loss model. (a) We show algebraic connectivity and coverage performance with varying path loss exponent. Error bars show standard deviation. We show communication networks for a 50 agent system for two path loss exponents at the end of a trial. The labeled exponents are determined experimentally. Algebraic connectivity increases only after performance levels out. Here, we find coverage performance and connectivity are unrelated in a critical regime. (b) Average coverage cost (ergodic metric) over 10 trials of delivery agents. All cases result in minimization of the ergodic metric except the highest loss case.

Work in [37] provides a possible explanation for the similar coverage results between low and high communication capability. That work used a voter model for a swarm consensus problem. The authors found fewer connections between agents improved the ability of a swarm to respond to new information

because the collective acted with an informed minority after communication. This effect was discussed for a more general set of complex systems such as vehicle traffic, social dynamics, and evolutionary adaptation in [38], referred to as the “slower is faster” effect. Here, the high number of interactions with a larger communication radius may lead to a choke point at which the agents run out of bandwidth to both receive and act on new information. This type of scalability challenge may appear in future hardware experiments.

VII. CONCLUSION

We presented an algorithm enabling heterogeneous multi-agent systems to manage safety constraints with limited communication. We updated safety constraints with evolving CBFs and provided guarantees on communication and safe coverage. We provided simulations for multi-agent teams with different tasks and capabilities, showing that a system with a limited communication range can perform as effectively as a fully connected communication network. This work is a step towards adaptable control architectures for heterogeneous robots to work autonomously in extreme conditions. Future work will extend theoretical results to find minimal communication requirements for successful coverage in resource-limited environments.

REFERENCES

- [1] Y. Rizk, M. Awad, and E. W. Tunstel, “Cooperative heterogeneous multi-robot systems: A survey,” *CSUR*, vol. 52, no. 2, pp. 1–31, 2019.
- [2] R. R. Murphy, S. Tadokoro, and A. Kleiner, “Disaster robotics,” in *Springer Handbook of Robotics*. Berlin, Heidelberg, Germany: Springer-Verlag, 2016, pp. 1577–1604.
- [3] D. S. Drew, “Multi-agent systems for search and rescue applications,” *Curr. Robot. Rep.*, vol. 2, pp. 189–200, 2021.
- [4] R. Almadhoun, T. Taha, L. Seneviratne, and Y. Zweiri, “A survey on multi-robot coverage path planning for model reconstruction and mapping,” *SN Appl. Sci.*, vol. 1, pp. 1–24, 2019.
- [5] L. Lindemann and D. V. Dimarogonas, “Control barrier functions for multi-agent systems under conflicting local signal temporal logic tasks,” *IEEE Control Syst. Lett.*, vol. 3, no. 3, pp. 757–762, Jul. 2019.
- [6] G. Bevacqua, J. Cacace, A. Finzi, and V. Lippiello, “Mixed-initiative planning and execution for multiple drones in search and rescue missions,” in *Proc. Int. Conf. Automated Plann. Scheduling*, 2015, vol. 25, no. 1, pp. 315–323.
- [7] C. Pinciroli and G. Beltrame, “Buzz: An extensible programming language for heterogeneous swarm robotics,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 3794–3800.
- [8] A. Pierson and M. Schwager, “Adaptive inter-robot trust for robust multi-robot sensor coverage,” in *Proc. 16th Int. Symp. Robot. Res.*, 2016, pp. 167–183.
- [9] Y. Diaz-Mercado, S. G. Lee, and M. Egerstedt, “Distributed dynamic density coverage for human-swarm interactions,” in *Proc. Amer. Control Conf.*, 2015, pp. 353–358.
- [10] K. Lee and K. Lee, “Adaptive centroidal voronoi tessellation with agent dropout and reinsertion for multi-agent non-convex area coverage,” *IEEE Access*, vol. 12, pp. 5503–5516, 2024.
- [11] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, “Ergodic exploration of distributed information,” *IEEE Trans. Robot.*, vol. 32, no. 1, pp. 36–52, Feb. 2016.
- [12] A. Prabhakar et al., “Ergodic specifications for flexible swarm control: From user commands to persistent adaptation,” in *Proc. Robot.: Sci. Syst.*, 2020, Paper 67.
- [13] J. Meyer et al., “Scale-invariant specifications for human-swarm systems,” *Field Robot.*, vol. 3, pp. 368–391, 2023.
- [14] I. Abraham and T. D. Murphey, “Decentralized ergodic control: Distribution-driven sensing and exploration for multiagent systems,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 2987–2994, Oct. 2018.
- [15] H. Pham and Q.-C. Pham, “A new approach to time-optimal path parameterization based on reachability analysis,” *IEEE Trans. Robot.*, vol. 34, no. 3, pp. 645–659, Jun. 2018.
- [16] I. Saha, R. Ramaithitima, V. Kumar, G. J. Pappas, and S. A. Seshia, “Automated composition of motion primitives for multi-robot systems from safe LTL specifications,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 1525–1532.
- [17] H. Salman, E. Ayvali, and H. Choset, “Multi-agent ergodic coverage with obstacle avoidance,” in *Proc. Int. Conf. Automated Plann. scheduling*, 2017, vol. 27, pp. 242–249.
- [18] S. Patel et al., “Multi-agent ergodic coverage in urban environments,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 8764–8771.
- [19] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *Proc. IEEE 18th Eur. Control Conf.*, 2019, pp. 3420–3431.
- [20] M. Rauscher, M. Kimmel, and S. Hirche, “Constrained robot control using control barrier functions,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 279–285.
- [21] P. Glotfelter, J. Cortés, and M. Egerstedt, “Nonsmooth barrier functions with applications to multi-robot systems,” *IEEE Control Syst. Lett.*, vol. 1, no. 2, pp. 310–315, Oct. 2017.
- [22] C. Lerch, D. Dong, and I. Abraham, “Safety-critical ergodic exploration in cluttered environments via control barrier functions,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 10205–10211.
- [23] D. Dong, H. Berger, and I. Abraham, “Time optimal ergodic search,” in *Proc. Robot.: Sci. Syst.*, 2023, Paper 82.
- [24] G. Mathew and I. Mezić, “Metrics for ergodicity and design of ergodic dynamics for multi-agent systems,” *Physica D: Nonlinear Phenomena*, vol. 240, no. 4/5, pp. 432–442, 2011.
- [25] E. G. Flytzanis, “Ergodicity of the cartesian product,” *Trans. Amer. Math. Soc.*, vol. 186, pp. 171–176, 1973.
- [26] N. Goddemeier, K. Daniel, and C. Wietfeld, “Role-based connectivity management with realistic air-to-ground channels for cooperative UAVs,” *IEEE J. Sel. Areas Commun.*, vol. 30, no. 5, pp. 951–963, Jun. 2012.
- [27] D. Aláez, M. Celaya-Echarri, L. Azpilicueta, and J. Villadangos, “UAV-radio: Radio link path loss estimation for UAVs,” *SoftwareX*, vol. 25, 2024, Art. no. 101628.
- [28] G.-Z. Yang et al., “The grand challenges of science robotics,” *Sci. Robot.*, vol. 3, no. 14, 2018, Art. no. eaar7650.
- [29] M. M. Alam, M. Y. Arafat, S. Moh, and J. Shen, “Topology control algorithms in multi-unmanned aerial vehicle networks: An extensive survey,” *J. Netw. Comput. Appl.*, vol. 207, 2022, Art. no. 103495.
- [30] O. Fernandes, R. Murphy, J. Adams, and D. Merrick, “Quantitative data analysis: CRASAR small unmanned aerial systems at hurricane harvey,” in *Proc. IEEE Int. Symp. Saf., Secur. Rescue Robot.*, 2018, pp. 1–6.
- [31] *XBee*, Digi, 3 2022, rev. P.
- [32] W. Newhall et al., “Wideband air-to-ground radio channel measurements using an antenna array at 2 GHz for low-altitude operations,” in *Proc. IEEE Mil. Commun. Conf.*, 2003, vol. 2, pp. 1422–1427.
- [33] U. M. Qureshi et al., “RF path and absorption loss estimation for underwater wireless sensor networks in different water environments,” *Sensors*, vol. 16, no. 6, 2016, Art. no. 890.
- [34] E. Yanmaz, R. Kuschnig, and C. Bettstetter, “Achieving air-ground communications in 802.11 networks with three-dimensional aerial mobility,” in *Proc. IEEE Int. Conf. Comput. Commun.*, 2013, pp. 120–124.
- [35] X. Cai et al., “Low altitude UAV propagation channel modelling,” in *Proc. IEEE 11th Eur. Conf. Antennas Propag.*, 2017, pp. 1443–1447.
- [36] R. Sun and D. W. Matolak, “Air-ground channel characterization for unmanned aircraft systems part II: Hilly and mountainous settings,” *IEEE Trans. Veh. Technol.*, vol. 66, no. 3, pp. 1913–1925, Mar. 2017.
- [37] M. S. Talamali, A. Saha, J. A. Marshall, and A. Reina, “When less is more: Robot swarms adapt better to changes with constrained communication,” *Sci. Robot.*, vol. 6, no. 56, 2021, Art. no. eabf1416.
- [38] C. Gershenson and D. Helbing, “When slower is faster,” *Complexity*, vol. 21, no. 2, pp. 9–15, 2015.