

CoCoPlan: Adaptive Coordination and Communication for Multi-robot Systems in Dynamic and Unknown Environments

Xintong Zhang^{2,†}, Junfeng Chen^{1,†}, Yuxiao Zhu², Bing Luo², and Meng Guo¹

Abstract—Multi-robot systems can greatly enhance efficiency through coordination and collaboration, yet in practice, full-time communication is rarely available and interactions are constrained to close-range exchanges. Existing methods either maintain all-time connectivity, rely on fixed schedules, or adopt pairwise protocols, but none adapt effectively to dynamic spatio-temporal task distributions under limited communication, resulting in suboptimal coordination. To address this gap, we propose CoCoPlan, a unified framework that co-optimizes collaborative task planning and team-wise intermittent communication. Our approach integrates a branch-and-bound architecture that jointly encodes task assignments and communication events, an adaptive objective function that balances task efficiency against communication latency, and a communication event optimization module that strategically determines when, where and how the global connectivity should be re-established. Extensive experiments demonstrate that it outperforms state-of-the-art methods by achieving a 22.4% higher task completion rate, reducing communication overhead by 58.6%, and improving the scalability by supporting up to 100 robots in dynamic environments. Hardware experiments include the complex 2D office environment and large-scale 3D disaster-response scenario.

Index Terms—Multi-robot System, Task Planning, Adaptive Communication.

I. INTRODUCTION

Multi-agent systems improve efficiency through concurrent execution and collaboration, particularly in complex scenarios. Traditional approaches focus either on spatial task planning [1], or on collective control such as formation maintenance [2], while assuming unrestricted communication for instantaneous information sharing. This assumption breaks down in real deployments, where communication is limited by obstacles [3], or range constraints [4]. Such limitations necessitate coordination of communication events. The challenge intensifies in dynamic unknown environments with unpredictable tasks and workspace models. Existing methods remain limited: static schedules [1] fail under unexpected tasks; dynamic approaches [5] require prior knowledge of distributions; and pairwise protocols [6] lack team-wide coordination, leading to degraded information flow. This creates a critical gap in jointly handling unknown spatiotemporal task distributions and team-wide intermittent communication under constraints. Three

Manuscript received 29 August 2025; Revised 29 November 2025; Accepted 29 December 2025.

This article was recommended for publication by Editor M. Ani Hsieh upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the National Natural Science Foundation of China (NSFC) under grants U2241214 and T2121002.

[†]Equal contribution. The authors are with ¹the School of Advanced Manufacturing and Robotics, Peking University, Beijing 100871, China; and ²the Division of Natural and Applied Sciences, Duke Kunshan University, Suzhou 215316, China.

Corresponding author: Meng Guo, meng.guo@pku.edu.cn.

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

©2026 IEEE

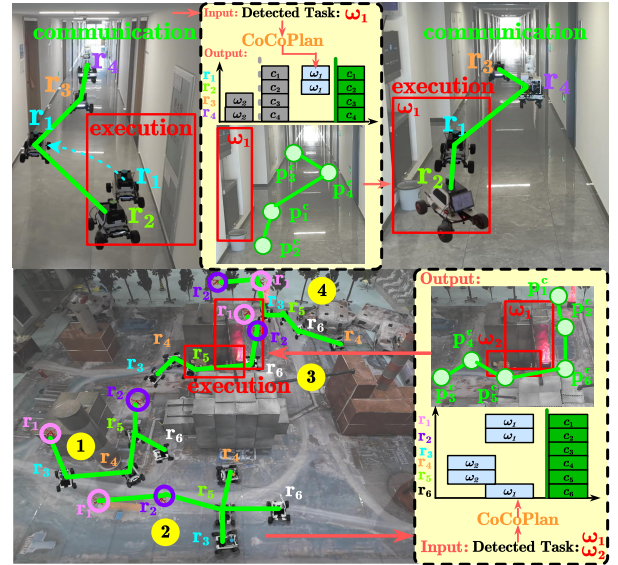


Fig. 1. Hardware demonstration of the proposed adaptive coordination and communication scheme (r_i denotes an individual hardware agent i , $\forall i \in \mathcal{N}$): 4 UGVs to fulfil office errands such as delivery and cleaning as requested online (top); 2 UAVs and 4 UGVs to detect fire hazards, search and rescue victims during a disaster-response mission (bottom).

fundamental challenges arise: the online discovery demands rapid replanning of assignments and communication when tasks emerge [7]; the sparse connectivity management requires strategic use of limited communication windows to preserve coordination [8]; the temporal constraints among tasks must be maintained between task deadlines and intermittent communication opportunities [9]. These challenges amplify with large fleets, where co-optimizing tasks and communication becomes combinatorially complex.

A. Related work

1) *Coordination under Limited Communication*: Communication is critical for multi-agent systems, especially in scenarios such as exploration and inspection where agents must exchange information to maintain alignment and improve outcomes. Existing approaches either employ an end-to-end connectivity model with constant communication, or an intermittent model where agents communicate only when necessary. Both involve trade-offs in efficiency, flexibility, and adaptability. All-time communication usually focuses on connectivity control, often using graph theory [10], [11], either by maintaining all initial links [10] or dynamically adding and removing them while preserving connectivity [12]. While this approach guarantees reliable data transfer, it often overlooks the communication needs of high-level, collaborative tasks. Moreover, it can be impractical in the presence of wireless channel uncertainties. Some studies address the

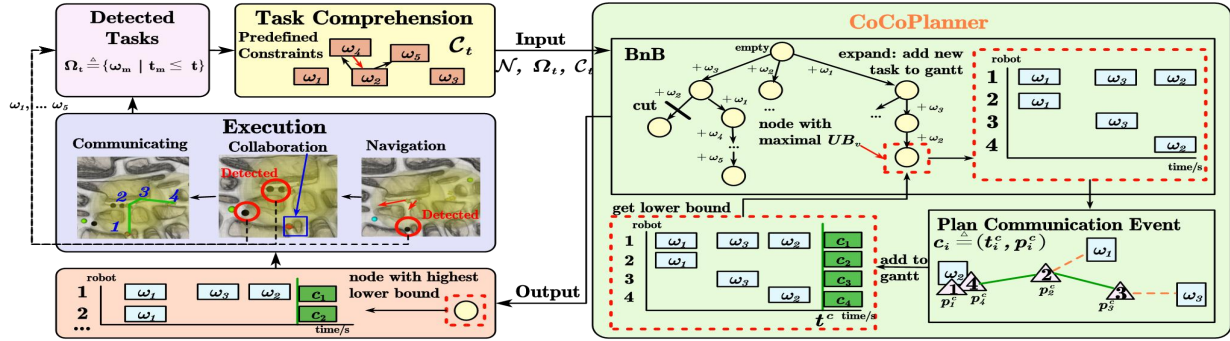


Fig. 2. Overview of the **CoCoPlan** framework, which integrates real-time task detection, a branch-and-bound planner for joint task and communication scheduling, and an adaptation scheme to unknown spatio-temporal task distributions.

joint constraints of connectivity and temporal tasks [13], yet they typically impose strict proximity requirements, reducing overall efficiency in large-scale deployments.

2) *Intermittent Communication under Complex Temporal Tasks*: To overcome these limitations, intermittent communication has been proposed as an alternative solution [6], [9], [14]. In this model, the communication network can be temporarily interrupted, and agents exchange information only at certain times. To this end, several related works adopt intermittent communication by predefined fixed intervals and optimizing meeting locations accordingly, including centralized [15], distributed [16] and predetermined meeting-point approaches [4]. However, these works typically fix communication times while optimizing locations, lacking simultaneous optimization of both aspects. Recent work mainly studies pairwise intermittent communication [6], [9], [14], where communication occurs only when necessary. While this improves efficiency, it enforces rigid patterns such as fixed schedules [6], [14] or predetermined meeting points [9], and introduces delays due to asynchronous execution, which is problematic for tasks that need timely coordination. Thus, all-time communication sacrifices task efficiency, while pairwise protocols slow information propagation. To mitigate these drawbacks, team-wise intermittent communication was proposed, in which all agents periodically re-establish global connectivity [8]. This differs from sub-team or clustered communication [17], [18], where “team-wise” denotes coordination within sub-teams or dynamically clustered groups. Nevertheless, existing team-wise approaches still assume known environments and task distributions, limiting applicability in uncertain scenarios.

B. Our Method

In this work, a multi-agent system handles complex temporally constrained tasks continuously released online as shown in Fig. 1. We formulate an online co-optimization problem to synthesize the collaborative task execution and communication strategies simultaneously. Unlike periodic team-wise or intermittent pair-wise communication protocols, our approach dynamically optimizes *when*, *where* and *how* agents communicate or collaborate based on real-time task specifications gathered online. To achieve this, we develop a branch-and-bound (BnB) framework with joint encoding of task planning and communication events within unified search nodes. A novel objective function balances the execution efficiency

against communication latency for robustness across spatio-temporal task distributions. Moreover, communication events are optimized via an iterative algorithm that strategically determines the communication locations and timings. As shown in Fig. 1, the method is validated through large-scale simulations and hardware experiments against several strong baselines.

The main contributions are three-fold: (I) A general online co-design of task planning and team-wise communication strategies, ensuring the timely information propagation and efficient task fulfillment; (II) An adaptive coordination mechanism via a novel objective function robust to unknown task distributions and dynamic environments; (III) Extensive validation in simulations and hardware experiments demonstrating the effectiveness in complex real-world scenarios.

II. PROBLEM DESCRIPTION

A. Workspace and Robot Model

Consider a team of N agents that share the common workspace $\mathcal{W} \subset \mathbb{R}^2$. Each agent $i \in \mathcal{N} \triangleq \{1, \dots, N\}$ is described by its position $p_i \in \mathcal{W}_i$, sensor range $d_i > 0$, and the action $a_i \in \mathcal{A}_i$, where $\mathcal{W}_i \subseteq \mathcal{W}$ is the allowed workspace; and \mathcal{A}_i is the set of primitive actions. More specifically, the actions \mathcal{A}_i require different capabilities from agent i . For instance, repair and rescue actions necessitate fine-grained manipulation capabilities, while delivery requires transportation capacity. The motion of each agent is characterized by its velocity $v_i \in \mathbb{R}^2$ with an upper limit $\|v_i\| \leq v_i^{\max}$, note that $\|v_i\|$ can be adjusted during navigation. Thus, the local plan of an agent is given by a sequence of timed goal positions and performed actions, i.e., $\tau_i \triangleq (t_i^1, g_i^1, a_i^1)(t_i^2, g_i^2, a_i^2) \dots$ with $t_i^\ell \geq 0$, $g_i^\ell \in \mathcal{W}_i$, $a_i^\ell \in \mathcal{A}_i$ being the time instant, goal position and action, $\forall \ell \geq 1$. In other words, under this local plan, agent $i \in \mathcal{N}$ should navigate to g_i^ℓ with velocity v_i and start performing a_i^ℓ from time t_i^ℓ , for all $\ell \geq 1$. The collective plan for all agents is denoted by $J \triangleq \{\tau_i\}, \forall i \in \mathcal{N}$.

B. Task Specification with Temporal Constraints

Collaborative tasks are dynamically released with unknown and possibly time-varying spatio-temporal distributions, each defined as $\omega_m \triangleq (S_m, \eta_m, \{(n_j, a_j)\}_{j=1}^{J_m})$ where and each pair $\{(n_j, a_j)\}$ indicates that n_j agents are required to perform action a_j within S_m . A task is detected only when $\min_{i \in \mathcal{N}} \|p_i(t) - S_m\| \leq d_i$ and no obstacle between $p_i(t)$ and S_m at t_m , where $\|\cdot\|$ is the Euclidean distance. The




Fig. 3. (Left): Ad-hoc local communication devices used in the hardware experiments; (Right): Map of the predicted communication quality used in the coordination of communication events.

detected tasks will later become known to all robots during the team-wise communication process. The set of detected tasks by time t is $\Omega_t \triangleq \{\omega_m \mid t_m \leq t\}$. For each task ω_r , its execution begins at $t_r^s > 0$ and finishes at $t_r^f > 0$. Three types of temporal constraints are predefined between tasks ω_p and ω_q : (1) precedence ($\omega_p \leq \omega_q$) holds if $t_p^f \leq t_q^s$; (2) mutual exclusion ($\omega_p \parallel \omega_q$) holds if their execution intervals are disjoint but no specific ordering is imposed, i.e., $[t_p^s, t_p^f] \cap [t_q^s, t_q^f] = \emptyset$; (3) concurrency ($\omega_p \sim \omega_q$) holds if parallel execution with overlapping time intervals is permitted, i.e., $[t_p^s, t_p^f] \cap [t_q^s, t_q^f] \neq \emptyset$. The set \mathcal{P}_t contains all predefined temporal relations between pairs of tasks in Ω_t , i.e.,

$$\mathcal{P}_t \subseteq \{(\omega_i, \omega_j, \kappa) \mid \omega_i, \omega_j \in \Omega_t, \kappa \in \{\leq, \parallel, \sim\}\}, \quad (1)$$

which can only be known at time $t \geq 0$. A task ω_m completes at finish time t_m^f if two conditions hold: (I) at time t_m^s , agents arrive at S_m and for each required action (n_j, a_j) , a team of n_j agents continuously performs a_j within S_m for duration η_m ; and (II) the execution interval $[t_m^s, t_m^f]$ satisfies all relevant constraints in \mathcal{P}_t involving ω_m .

C. Communication Constraints

Agents cannot maintain the continuous real-time communication due to physical limitations. As illustrated in Fig. 3, successful communication between agents $i, j \in \mathcal{N}$ requires sufficient communication quality $Q_{ij}(t) > \delta$, where $\delta > 0$ is the minimum quality threshold. The communication quality $Q_{ij}(t)$ is defined as: $Q_{ij}(t) \triangleq P_t - PL(d_{ij}(t)) - \alpha d_{ij}^{\text{obs}}(t)$, for locations $p_i, p_j \in \mathcal{W}_i, \mathcal{W}_j$, where P_t denotes the transmission power in dB, and $PL(d_{ij}(t)) \triangleq PL(d_0) + 10n \log_{10} \left(\frac{d_{ij}(t)}{d_0} \right)$ represents the free-space path loss, with d_0 as the reference distance, n as the path loss exponent, and $d_{ij}(t) \triangleq \|p_i(t) - p_j(t)\|$ as the distance between agents i and j at time t . Note that $d_{ij}^{\text{obs}}(t)$ denotes the total length of obstacles along the straight line connecting $p_i(t)$ and $p_j(t)$, with $\alpha > 0$ as the attenuation coefficient per meter of obstacle. This constrained connectivity is modeled as a time-varying graph $G(t) \triangleq (\mathcal{N}, \mathcal{E}(t))$ with the edge set $\mathcal{E}(t) \triangleq \{e_{ij} \mid Q_{ij}(t) > \delta\}$, where e_{ij} is the communication link between agents i and j . The agent team \mathcal{N} achieves the global communication at the time $t > 0$ if $G(t)$ is connected, namely: $G(t) \in \mathcal{G}_c$, where \mathcal{G}_c is the set of connected graphs. We refer to this as team-wise intermittent communication, requiring global connectivity across all agents at each event to enable centralized-style replanning. To establish the global connectivity, agents synchronize at the designated *communication events* defined as $\mathbf{C} \triangleq \{\mathbf{c}_i \mid \forall i \in \mathcal{N}\}$, where each agent-specific event $\mathbf{c}_i \triangleq$

(t_i^c, p_i^c) specifies the communication time t_i^c and position p_i^c . The global event \mathbf{C} occurs when all agents simultaneously reach their respective communication positions at their common times t^c , i.e., $t^c = t_i^c, \forall i \in \mathcal{N}$. A valid communication event \mathbf{C} must ensure that:

$$G(t^c) \in \mathcal{G}_c, \quad (2)$$

where t^c is the communication time. Once $G(t^c)$ is connected, agents exchange critical information including: their individually detected tasks Ω_t^i , locally generated plans τ_i , and future communication events \mathbf{c}_i . Subsequently, the local plan for agent i extends to $\tau_i \triangleq (t_i^1, g_i^1, a_i^1)(t_i^2, g_i^2, a_i^2) \cdots \mathbf{c}_i$.

D. Problem Statement

The overall objective is to synthesize the collective plans $J \triangleq \{\tau_i\}_{i \in \mathcal{N}}$ so as to maximize the long-run average rate of task completion over an infinite horizon. For a given collective plan J and time $T > 0$, define $N_T(J) \triangleq |\{\omega_m \in \Omega_T \mid t_m^f(J) \leq T\}|$, which denotes the number of tasks completed by time T under plan J . The long-run average task completion rate of J is then $\eta(J) \triangleq \lim_{T \rightarrow \infty} \frac{N_T(J)}{T}$. Thus, the overall problem is as follows:

$$\begin{aligned} & \max_J \eta(J) \\ & \text{s.t.} \quad (1) - (2), \forall i \in \mathcal{N}, \end{aligned} \quad (3)$$

where the collective plan J that maximizes the average rate $\eta(J)$; (1) ensures that all temporal relations between tasks are satisfied; and (2) guarantees that the team achieves intermittent global connectivity at communication events.

III. PROPOSED SOLUTION

A unified framework co-optimizes collaborative task planning and communication in dynamic, unknown environments. It combines a Branch-and-Bound method for temporally constrained planning and communication scheduling in Sec. III-A, an iterative approach for optimizing communication events in Sec. III-B, and an online execution scheme with complexity analysis and theoretical guarantees in Sec. III-C.

A. Adaptive Coordination Framework

1) *Design of Adaptive Optimization Objective*: To handle unknown spatiotemporal task distributions while still achieving the global objective (3), an adaptive optimization objective is formulated to maximize task completion efficiency within each execution horizon. The coordination framework is triggered at every communication event occurring at time t . Based on the currently detected task set Ω_t , the framework jointly optimizes the future local plans J together with the scheduling of the next communication event \mathbf{c} . The team-wise intermittent communication protocol guarantees that all tasks assigned in $\{\tau_i\}$ during the interval $[t, t^c]$ are completed before the next re-planning phase at t^c . The adaptive objective is expressed as the maximization of the task completion rate per unit time in the interval $[t, t^c]$, as an *local approximation* of $\eta(J)$ in (3):

$$\begin{aligned} & \max_{\{\tau_i\}, t^c} \frac{|\{\omega_m \in \Omega_t \mid t_m^f \leq t^c\}|}{t^c - t} \\ & \text{s.t.} \quad G(t^c) \in \mathcal{G}_c, \mathcal{C}_t \text{ holds}; t^c \geq t_{i,\ell}^f > t, \\ & \quad t_{i,\ell+1}^s \geq t_{i,\ell}^f + T_{\text{travel}}(\omega_i^\ell, \omega_i^{\ell+1}), \omega_i^\ell \in \Omega_t \end{aligned} \quad (4)$$

Algorithm 1: CoCoPlan(\cdot)

Input: Robot team \mathcal{N} , Detected tasks Ω_t , Time budget T'

Output: Best collective plan J^*

- 1 Initialize $\mathcal{T} \leftarrow \emptyset$, \mathcal{Q} (max-heap by UB)
- 2 Create ν_0 : $J_{\nu_0} \leftarrow \emptyset$
- 3 $(LB_{\nu_0}, J_{\nu_0}) \leftarrow \text{LowBound}(\nu_0, \Omega_t, \mathcal{N})$
- 4 $UB_{\nu_0} \leftarrow \text{UpBound}(\nu_0, \Omega_t, \mathcal{N})$
- 5 $\mathcal{T} \leftarrow \mathcal{T} \cup \{\nu_0\}$, $\mathcal{Q}.\text{insert}(\nu_0, UB_{\nu_0})$
- 6 $LB^* \leftarrow LB_{\nu_0}$, $J^* \leftarrow J_{\nu_0}$, $t_{\text{start}} \leftarrow$ current time
- 7 **while** $\mathcal{Q} \neq \emptyset$ **and** $(t - t_{\text{start}} < T')$ **do**
- 8 $\nu \leftarrow \mathcal{Q}.\text{extractMax}()$
- 9 **if** $UB_{\nu} > LB^*$ **then**
- 10 **if** $LB_{\nu} > LB^*$ **then**
- 11 $LB^* \leftarrow LB_{\nu}$, $J^* \leftarrow J_{\nu}$
- 12 $\mathcal{F} \leftarrow \text{GetTasks}(\nu, \Omega_t)$
- 13 **foreach** $\omega \in \mathcal{F}$ **do**
- 14 $\nu_+ \leftarrow \text{ExpNode}(\nu, \omega)$
- 15 $(LB_{\nu_+}, J_{\nu_+}) \leftarrow \text{LowBound}(\nu_+, \Omega_t, \mathcal{N})$
- 16 $UB_{\nu_+} \leftarrow \text{UpBound}(\nu_+, \Omega_t, \mathcal{N})$
- 17 $\mathcal{T} \leftarrow \mathcal{T} \cup \{\nu_+\}$
- 18 **if** $LB_{\nu_+} > LB^*$ **then**
- 19 $LB^* \leftarrow LB_{\nu_+}$, $J^* \leftarrow J_{\nu_+}$
- 20 **if** $UB_{\nu_+} > LB^*$ **then**
- 21 $\mathcal{Q}.\text{insert}(\nu_+, UB_{\nu_+})$
- 22 **return** J^*

where t_m^{c} is the completion time of task ω_m , t^{c} denotes the scheduled time of the next communication event, and $G(t^{\text{c}}) \in \mathcal{G}_{\text{c}}$ ensures that global communication occurs at t^{c} .

2) *Algorithm Description:* The proposed CoCoPlan in Alg. 1 uses a branch-and-bound (BnB) search to jointly optimize task assignments and communication events for the entire team. Each node ν in the BnB tree corresponds to a collective plan $J_{\nu} \triangleq [(\omega_{i_1}^1, \omega_{i_1}^2, \dots, \mathbf{c}_{i_1}), (\omega_{i_2}^1, \dots, \mathbf{c}_{i_2}), \dots]$, i.e., agent-specific task sequences interleaved with communication events \mathbf{c} . For brevity, the local plan elements $(t_i^{\text{c}}, g_i^{\text{c}}, a_i^{\text{c}})$ are associated with the task ω_i^{c} in J_{ν} . For each node ν , we maintain a lower bound LB_{ν} and an upper bound UB_{ν} on the objective value; nodes are pruned if $UB_{\nu} < LB^*$, where LB^* is the current best lower bound. The objective of each node is given by (4), representing the task completion rate up to the next communication event. The search continues until all nodes are expanded or pruned, or a time limit is reached. It then returns the best collective plan J^* found within the computation budget. The algorithm follows five key stages:

(I) Initialization (Lines 1–6): The root node ν_0 is created with an empty plan, and its bounds LB_{ν_0} and UB_{ν_0} are computed by `LowBound` and `UpBound`, providing pessimistic and optimistic estimates of achievable performance. These values initialize the max-heap \mathcal{Q} , which stores nodes ordered by UB_{ν} , and the incumbent solution (LB^*, J^*) . A timer t_{start} is started to enforce the computation budget T' .

(II) Selection (Lines 7–9): At each iteration, the node ν with the largest UB_{ν} is extracted from \mathcal{Q} (Line 8), following a best-first strategy that prioritizes nodes with the highest potential




Fig. 4. Illustration of the iterative process for computing lower bounds, after inserting new tasks in the collective plan.

to improve the incumbent. If $UB_{\nu} \leq LB^*$, then ν and its subtree are pruned, since they cannot yield a better solution. This pruning significantly reduces the explored search space.

(III) Expansion (Lines 12–21): For each surviving node, feasible tasks \mathcal{F} that respect the temporal constraints \mathcal{C}_t are identified (Line 12). The `ExpNode` function (Line 14) generates child nodes ν_+ by assigning each task $\omega \in \mathcal{F}$ to eligible agents and inserting it in their plans before the next communication event \mathbf{c}_{i_k} . This couples task allocation with communication scheduling. Each child is then evaluated to obtain new bounds LB_{ν_+} and UB_{ν_+} .

(IV) Bounding (Lines 3–16): Each child ν_+ undergoes iterative bound refinement starting from its plan J_{ν_+} . As illustrated in Fig. 4, at iteration k a feasible task $\omega^{(k)}$ is chosen from $\Omega_t \setminus \Omega_{\nu_+}$ and assigned to an eligible agent group $\mathcal{A} \in \mathcal{A}_k$ that minimizes the maximum travel time to the task location:

$$\min_{\mathcal{A} \in \mathcal{A}_k} \left\{ \max_{i \in \mathcal{A}} T_{\text{travel}}(p_i^{\text{c}}, \omega^{(k)}) \right\}, \quad (5)$$

yielding an updated plan $J^{(k)}$ for node $\nu^{(k)}$. Its completion time is computed in two phases: (1) *Task phase*: non-communication tasks are scheduled by solving: $\min \max_{\omega_m \in \Omega_{\nu^{(k)}}} t_m^{\text{c}}$, using a standard LP solver such as OR-Tools [19], temporarily ignoring communication; and (2) *communication phase*: given the resulting task schedule, `ComOpt` optimizes communication events and returns $t_{(k)}^{\text{c}}$. The corresponding candidate rate is $b^{(k)} = |\omega_m \in \Omega_{\nu^{(k)}}| / t_{(k)}^{\text{c}}$. Tasks are added while $b^{(k)}$ improves, and the lower bound LB_{ν_+} is set to $\max_j b^{(j)}$. The upper bound UB_{ν_+} is obtained by applying the same procedure but assuming instantaneous agent movement, giving an optimistic estimate for pruning.

(V) Termination (Line 22): The algorithm stops when the heap \mathcal{Q} is empty or the time budget T' is reached, and returns the incumbent plan J^* as the best solution found. Since bounds and solutions are refined throughout the search, the method is anytime: even early termination yields a feasible plan.

B. Optimization of Communication Events

1) *Problem of Communication Events Optimization:* Given the non-communication solution J_{ν_+} obtained in the previous step, let t_i^{c} denote the completion time of the last non-communication task in the plan of agent i as τ_i , and let $p_i^{\text{c}} \in \mathcal{W}$ be its position at time t_i^{c} . The subsequent communication event must occur at a time $t^{\text{c}} > \max_{i \in \mathcal{N}} t_i^{\text{c}}$, with agents located at positions $\{p_i^{\text{c}}\}$, where $p_i^{\text{c}} \in \mathcal{W}_i$, and the induced communication graph $G(t^{\text{c}})$ satisfies the connectivity requirement $G(t^{\text{c}}) \in \mathcal{G}_{\text{c}}$. The optimization problem aims to minimize the maximum transition time required for agents to travel from

Algorithm 2: ComOpt(\cdot)

Input: Robot team \mathcal{N} , set of last tasks $\{t_i^f\}$, time budget T' , threshold δ

Output: Communication event \mathbf{C}_t

```

1  $i_\ell \leftarrow \operatorname{argmax}_i t_i^f$ 
2  $p_0 \leftarrow p_{i_\ell}^f, t^c \leftarrow t_{i_\ell}^f, t^b \leftarrow t^c$ 
3  $\mathbf{C}_t \leftarrow \{(t^b, p_0) \mid \forall i \in \mathcal{N}, \forall j > i\}$ 
4  $t_{\text{start}} \leftarrow \text{current time}$ 
5 while  $(t - t_{\text{start}} < T')$  and  $(|t^c - t^b| \geq \Delta)$  do
6    $i_e \leftarrow \operatorname{argmin}_i (t_i^f + A^*(p_i^f, p_0))$ 
7    $j_n \leftarrow \operatorname{argmin}_{j \neq i_e} A^*(p_{i_e}^f, p_j^c)$ 
8    $p^c \leftarrow \operatorname{SelCom}(p_{i_e}^f, p_{j_n}^c)$ 
9    $t_a \leftarrow t_{i_e}^f + A^*(p_{i_e}^f, p^c)$ 
10   $t_+^c \leftarrow \max(t_a, \max_{k \neq i_e} (t_k^f + A^*(p_k^f, p_k^c)))$ 
11  if  $(Q_{c, p_{j_n}^c} > \delta)$  and  $(t_+^c < t^b)$  then
12     $p_{i_e}^c \leftarrow p^c$ 
13    if  $|t^b - t_+^c| < \delta$  then
14       $t^c \leftarrow t_+^c, t^b \leftarrow t_+^c$ ; break
15    else
16       $t^c \leftarrow t_+^c, t^b \leftarrow t_+^c$ 
17 return  $\mathbf{C}_t = \{(t^b, p_i^c) \mid \forall i\}$ 

```

their last task locations to the communication point, ensuring both feasibility and efficiency:

$$\begin{aligned}
 & \min_{t^c, \{p_i^c\}} \left\{ \max_{i \in \mathcal{N}} (t^c - t_i^f) \right\} \\
 \text{s.t. } & t^c \geq t_i^f + T_{\text{travel}}(p_i^f, p_i^c), \quad \forall i \in \mathcal{N}, \\
 & G(t^c) \in \mathcal{G}_c,
 \end{aligned} \quad (6)$$

where the first constraint guarantees sufficient time for each agent to navigate from its last task location p_i^f to its communication position p_i^c ; and the second enforces that the team remains connected at the communication time t^c . This formulation explicitly balances task feasibility with connectivity, ensuring that communication is synchronized across the team while minimizing idle time.

2) *Algorithm Description:* The ComOpt(\cdot) algorithm in Alg. 2 addresses the optimization problem in (6). The procedure begins by identifying the agent i_ℓ that finishes its last task the latest in Line 1, since all agents must eventually wait for this agent before communication. The candidate communication location p_0 is set to $p_{i_\ell}^f$, and the initial communication time is initialized to $t^b = t^c = t_{i_\ell}^f$ in Line 2. Accordingly, the event \mathbf{C}_t is initialized that all agents converge to p_0 at t^b .

Within the computational budget T' , the algorithm iteratively refines the initial guess until the gap $|t^c - t^b|$ is below the convergence threshold Δ as stated in Line 5. As illustrated in Fig. 5, each iteration proceeds as follows. First, the earliest-arriving agent i_e is selected in Line 6 by minimizing $t_i^f + A^*(p_i^f, p_0)$. Among the agents that finish last and the agents whose communication has already optimized, we identify the nearest neighbor j_n in Line 7 using the shortest A^* path from $p_{i_e}^f$ to $p_{j_n}^c$. A candidate position p^c is sampled between $p_{i_e}^f$ and $p_{j_n}^c$ in Line 8 using the SelCom function, and the corresponding arrival time t_a is computed in Line 9. The potential communication time t_+^c is then defined in Line 10 as the maximum of t_a and the arrival times of all other

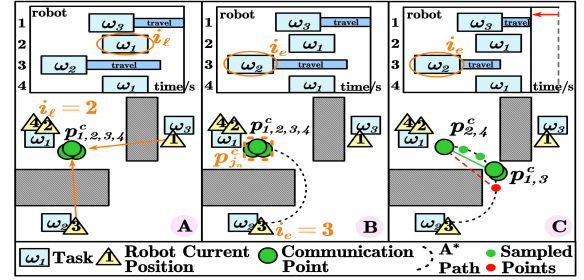


Fig. 5. Optimization of communication events, showing the selection of the earliest-arriving agent and the sampling of candidate communication points to achieve synchronized connectivity.

agents at their current positions. If the the communication quality between candidate p^c and $p_{j_n}^c$ exceeds the threshold δ , and the updated communication time t_+^c is earlier than the current best t^b , then the location p^c is accepted as the new position of agent i_e in Line 12. The communication time is updated accordingly in Line 14 and Line 16, and convergence is reached once the difference $|t^b - t_+^c|$ falls below δ , which terminates the loop early. Lastly, the algorithm outputs the optimized communication event \mathbf{C}_t , including the agent positions $\{p_i^c\}$ and the synchronized time t^b . This event satisfies $G(t^b) \in \mathcal{G}_c$ and minimizes the maximum delay $\max_i (t^b - t_i^f)$, as required by (6).

C. Overall Analysis

1) *Online Execution:* To handle dynamically emerging tasks, we adopt an online execution scheme. As illustrated in Fig. 2, agents discover new tasks while moving, incrementally updating the task set Ω_t . Each agent executes its current local plan τ_i and progresses toward its next communication location.

To handle uncertainties such as varying execution times and navigation delays, we adopt simple adaptation mechanisms: for each assigned task, the corresponding agents wait at the task region until all members have arrived before execution; at each communication event, all agents wait until every agent has reached its designated communication location, ensuring $G(t^c) \in \mathcal{G}_c$. This maintains both temporal and communication constraints, even when timing deviations occur. At each communication event, agents exchange information about newly discovered tasks. A designated agent aggregates these into Ω_t and invokes COCOPLAN to compute updated collective plans and the next communication event \mathbf{C} . The team then follows the new plan, and this cycle of execution, communication, and replanning repeats until the mission ends.

2) *Complexity Analysis:* The complexity of Alg. 1 is dominated by its branch-and-bound (BnB) search. With m tasks and N agents, a single expansion step can generate up to $\mathcal{O}(m \cdot N!)$ child nodes, leading to a worst-case total of $\mathcal{O}((m \cdot N!)^m)$ nodes. Per node, task expansion and bound evaluation are dominated by $\mathcal{O}(m \cdot N!)$, so the overall worst-case complexity remains $\mathcal{O}((m \cdot N!)^m)$. In practice, pruning substantially reduces the number of explored nodes.

Lemma 1. *Alg. 2 returns communication locations such that the communication graph $G(t^c)$ is connected.*

Proof. (Sketch) All agents are initialized at the same communication point p_0 , so $G(t^c)$ is connected. At each iteration, the

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

algorithm updates a single agent i_e to a new point p_{i_e} that preserves a communication link to already-fixed point p_{j_n} , thus adding a node adjacent to the existing connected set. Repeating this for all agents preserves connectivity by induction. ■

Theorem 2. Assume that communication events recur indefinitely and that at every communication time Alg. 1 is solved to optimality for the adaptive objective (4). Then on each communication cycle $[t_k^c, t_{k+1}^c)$ the returned plan is feasible w.r.t. (1) and (2), and locally maximizes the task completion rate over that cycle. Moreover, (4) provides a local approximation of the global infinite-horizon objective (3); Under stationary task arrival and ergodicity assumptions, the two optima coincide.

Proof. (Sketch) Feasibility on each cycle follows because Alg. 1 enforces the temporal constraints (1) and calls Alg. 2, which guarantees connectivity via Lemma 1. Solving the branch-and-bound search to optimality ensures that the returned plan maximizes (4) over that cycle, giving a locally optimal approximation of (3). Under stationary, constant-rate task discovery and ergodic cycles, the renewal-reward theorem implies that the long-run average completion rate equals the time average of the per-cycle rates, so local and global optimality are equivalent. ■

IV. NUMERICAL EXPERIMENTS

To validate the proposed method, extensive numerical experiments are conducted across multiple scenarios. The method is implemented in Python3 within the ROS framework, and evaluated on a workstation equipped with an Intel(R) i7 24-core CPU and an RTX-4070Ti GPU. Simulation and hardware videos are provided in the supplementary material.

A. Workspace and Task Description

Experiments employ $N = 10$ heterogeneous robots in two challenging environments. (I) The DARPA SubT challenge environment, measuring $40 \text{ m} \times 60 \text{ m}$, contains narrow passages that test maneuverability under constrained conditions. (II) The subterranean caves, with a size of $40 \text{ m} \times 70 \text{ m}$, include irregular obstacles that serve as a complex setting for validation. All robots operate with identical parameters: velocity range with maximum $v_i^{\max} = 2 \text{ m/s}$, sensing range $d_i = 8 \text{ m}$, and communication quality determined by the model in Section II-C. The method activates during communication events. Task sequences τ_i are executed using the native SLAM-based navigation package.

B. Simulation Results

The proposed method is evaluated in both environments, as shown in Fig. 6. Each algorithm invocation takes 15s of computation with 100 nodes. In the SubT scenario, four task types are considered: Type 1 fault detection precedes Type 2 cleaning, and Type 3 indoor delivery precedes Type 4 equipment maintenance. In Fig. 6(a), at 232s, the robots establish communication and synchronize their plans. By 258s, execution begins, with 2 robots performing fault detection. At 319s, the communication topology is reorganized for the next planning cycle. In the subterranean caves, the task set

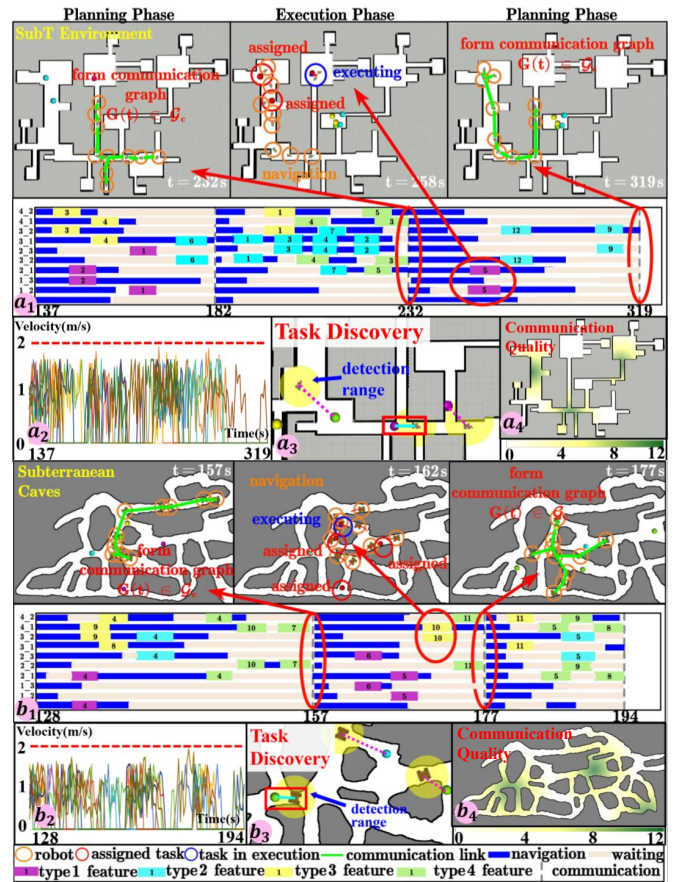


Fig. 6. Results of the proposed method in simulation, including representative keyframes and task planning outcomes in the DARPA SubT challenge environment (top) and the subterranean caves (bottom)

includes: Type 1 rescue precedes Type 2 cargo transportation, and Type 3 exploration precedes Type 4 mineral sampling. In Fig. 6(b), a communication graph forms at 157s to initiate task allocation. By 162s, robots navigate to their goals, with exploration tasks completed before mineral sampling due to temporal constraints. At 177s, the communication topology is reorganized for continued coordination. These results demonstrate real-time task allocation and communication topology adaptation. The system achieves 4 communication events and 17 tasks in SubT, and 4 communication events with 16 tasks in subterranean caves.

C. Comparison with Baselines

To evaluate the proposed method, it is compared against six baselines. (I) FIX triggers planning when accumulated tasks reach a threshold N , and is adapted to enforce temporal constraints. (II) FPMR [20] uses a fixed communication point; for fairness, it is applied only to communication optimization while retaining the proposed task allocation strategy. (III) FRDT [14] gathers robots around a fixed leader to communicate, and is likewise used only for communication optimization. (IV) FIMR [6] plans at regular intervals T_c , used only for communication optimization to preserve adaptivity. (V) RING [9] pairs robots with fixed partners and propagates information along the chain; it is modified to handle temporal and capability constraints. (VI) Greedy communicates and executes any feasible tasks whenever robots meet, and

TABLE I
 COMPARISON OF BASELINES ACROSS TWO SCENARIOS.

Env.	Strategy	Finished Tasks (#)	Comm. Num. (#)	Comm. Int. (s)	$t_c - t_f$ (s)
SubT	FIX (N=3)	74.7±3.1	34.0±0.0	22.2±10.4	2.6±4.8
	FIX (N=10)	95.3±6.6	12.0±0.0	51.8±18.2	3.0±4.6
	FPMR	89.7±1.0	9.3±1.3	78.7±60.3	19.3±8.3
	FRDT	94.3±3.1	10.3±1.9	69.8±54.8	17.2±12.9
	FIMR ($T_c = 35$)	96.0±5.0	17.7±1.3	35.0±0.0	6.2±6.1
	FIMR ($T_c = 80$)	87.7±4.0	9.0±0.0	80.0±0.0	16.4±18.9
	RING	52.3±6.8	7.0±0.8	156.1±75.1	85.0±42.7
	Greedy	30.0±10.6	166.3±10.8	—	—
	Ours	99.0±0.8	13.0±0.8	46.7±44.0	2.8±2.5
Caves	FIX (N=3)	81.7±5.9	34.7±0.4	21.6±11.6	2.6±5.2
	FIX (N=10)	95.7±4.8	12.0±0.0	52.8±17.8	4.2±6.1
	FPMR	96.7±2.5	10.3±1.7	68.3±55.6	19.6±9.9
	FRDT	84.3±1.0	10.0±0.8	82.0±76.6	18.2±13.1
	FIMR ($T_c = 35$)	83.7±4.5	20.7±0.4	35.0±0.0	5.9±5.9
	FIMR ($T_c = 80$)	90.7±6.8	8.7±0.4	80.0±0.0	19.8±22.6
	RING	49.0±8.3	7.0±0.8	183.0±84.8	106.2±53.6
	Greedy	57.7±3.3	166.7±17.5	—	—
	Ours	98.7±0.4	13.7±0.4	45.1±43.4	2.2±4.6

is adapted to account for capability constraints during task allocation. Four metrics are reported: finished tasks, communication count, communication interval, and $t_c - t_f$, the gap between the last task completion and the next communication. All methods are tested in both scenarios over three trials, each combining spatial patterns (clustered, uniform, sparse) and temporal patterns (spiky, uniform, low-frequency). Since the selected baselines were developed under slightly different (often simpler) settings, the comparison is intended to highlight performance in complex unknown scenarios rather than claim narrow superiority.

1) *Performance in Nominal Scenarios*: As shown in Table I, the proposed method achieves nearly double the task completion of RING, reaching 99.0 vs. 52.3 in SubT and 98.7 vs. 49.0 in the caves. It also completes nearly three times and twice as many tasks as Greedy, while requiring 92.2% and 91.8% fewer communication events in the two scenarios. Compared to FIX with $N = 3$, substantially more tasks are completed (99.0 vs. 74.7 and 98.7 vs. 81.7). Against FIX with $N = 10$, the adaptive objective markedly reduces variance (0.8 vs. 6.6 and 0.4 vs. 4.8). FIMR with $T_c = 35$ s exhibits instability across environments, with a 14% drop in completion, whereas the proposed method remains consistently high. FIMR with $T_c = 80$ s incurs long idle times and large variance. FPMR and FRDT are constrained by fixed communication, causing significant delays and exposing the drawback of rigid synchronization relative to adaptive coordination. Overall, the proposed method consistently delivers high task throughput while limiting unnecessary coordination overhead across diverse task distributions.

2) *Robustness under Varying Task Distributions*: To assess robustness under unknown spatiotemporal task distributions, a 600s experiment is conducted with three phases. During the first 200s, tasks follow a spatially sparse and temporally low-frequency pattern. Over [200s, 400s], they shift to a spatially clustered and temporally bursty pattern. From 400s to 600s,




Fig. 7. Robustness analysis: instantaneous variance of completed tasks (left) and task-completion efficiency (right).




Fig. 8. Snapshots of extensive experiments showing operation in a 3D urban environment (left) and a large-scale deployment with 100 robots (right).

the distribution becomes spatially and temporally uniform. Fig. 7(a) reports the instantaneous variance of completed tasks. The proposed method maintains consistently low variance under changing conditions, whereas baselines exhibit sharp fluctuations. Fig. 7(b) compares task-completion efficiency, defined as the slope of the task-completion curve. The proposed method sustains high mean efficiency with low variance in both environments (0.173 ± 0.005 in SubT and 0.176 ± 0.006 in caves), while other baselines show substantial volatility. Greedy yields the lowest mean efficiency, and FIX with $N = 10$ exhibits high variance, underscoring the weakness of fixed thresholds.

D. Generalizations

1) *3D Environment*: To evaluate performance in complex three-dimensional spaces, an urban environment of $60 \times 40 \times 20 \text{ m}^3$ is designed, deploying a heterogeneous team of five UAVs and five UGVs. Task types are designed to leverage aerial capabilities, such as high-altitude damage inspection and building fire-rescue missions. Figure 8 (left) illustrates the resulting communication topologies that emerge during task execution. The system successfully completes 45 tasks within 415s while triggering only 11 communication events, highlighting scalability of our method in 3D settings.

2) *Large-scale Fleets*: Scalability is further assessed in a large-scale deployment of $N = 100$ robots performing tasks analogous to those in the SubT environment. As shown in Fig. 8, the system completes 160 tasks within 198s, requiring only six communication events to maintain coordination. All tasks satisfy their temporal constraints, and the computation times is 5-15s. These results demonstrate the capability to operate for large-scale robotic fleets.

E. Hardware Experiments

To assess practicality, hardware experiments are conducted in office service and 3D disaster-response environments, demonstrating feasibility in real-world settings.




Fig. 9. Snapshots of hardware experiments: (a) Experimental setup for service robots in office; (b) Key snapshots of task execution; (c) Hardware deployment in 3D disaster response; (d) Key snapshots of collaborative task execution.

1) *Service Robots in Office*: The first experiment uses a $138 \times 114 \times 3, \text{m}^3$ office environment (Fig. 9) with four UGVs: one delivery, one maintenance, and two cleaning robots. Each robot is equipped with ad-hoc networking units (Fig. 3) and navigation modules. Five task types are considered with temporal priorities: fault detection (highest), indoor delivery and garbage collection (medium), and stain removal and equipment maintenance (lowest). At 80s, a floor stain is detected along with a coffee delivery request. By 91s, a communication graph is formed and both tasks are assigned. The delivery robot finishes at 129s (purple block). A cleaning robot then handles a stain detected at 137s, starting removal at 150s (blue block). At 176s, a new communication topology is established for subsequent planning. Overall, 11 tasks are completed within 537s with only 5 communication events, indicating responsive replanning under dynamic and human-triggered task arrivals.

2) *3D Disaster Response*: The second experiment takes place in a $16 \times 14 \times 26, \text{m}^3$ disaster environment (Fig. 9). Six task types are considered, including firefighting before rescue and chimney servicing before building repair, along with resource collection and water protection. The team includes two UAVs for monitoring/maintenance and four UGVs: two for firefighting/repair and two for rescue. The first planning round completes at 23s, after which UAVs execute chimney repair, completed at 51s (pink block). Maintenance requires UAV-UGV collaboration at 61s and again at 90s, with communication re-established after each phase to synchronize actions. In total, 12 tasks are completed in 248s, demonstrating timely responsiveness and effective ground-air coordination in emergency conditions.

V. CONCLUSION & FUTURE WORK

This work presents a branch-and-bound framework that co-optimizes task collaboration and team-wise communication online, for multi-agent systems operating in dynamic and unknown environments. The proposed approach achieves efficient information exchange while meeting temporal task requirements, as validated through extensive simulations and hardware experiments. Future research will explore the hierarchical combination of global and local coordination protocols.

REFERENCES

- [1] J. Huang, X. Li, and L. Gao, "A novel ga-cp method for fixed-type multi-robot collaborative scheduling in flexible job shop," *IEEE Transactions on Automation Science and Engineering*, vol. 22, no. 3, pp. 13531–13543, 2025.
- [2] W. Ren and Y. Cao, *Distributed coordination of multi-agent networks: emergent problems, models, and issues*. Springer, 2011, vol. 1.
- [3] X. Yu, D. Saldaña, D. Shishika, and M. A. Hsieh, "Resilient consensus in robot swarms with periodic motion and intermittent communication," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 110–125, 2021.
- [4] A. R. Da Silva, L. Chaimowicz, T. C. Silva, and M. A. Hsieh, "Communication-constrained multi-robot exploration with intermittent rendezvous," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 3490–3497.
- [5] A. T. Buyukkocak, D. Aksaray, and Y. Yazıcıoğlu, "Planning of heterogeneous multi-agent systems under signal temporal logic specifications with integral predicates," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1375–1382, 2021.
- [6] 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.
- [7] Z. Chen and Z. Kan, "Real-time reactive task allocation and planning of large heterogeneous multi-robot systems with temporal logic specifications," *The International Journal of Robotics Research*, vol. 44, no. 4, pp. 640–664, 2025.
- [8] M. Saboia, L. Clark, V. Thangavelu, J. A. Edlund, K. Otsu, G. J. Correa, V. S. Varadharajan, A. Santamaria-Navarro, T. Touma, A. Bouman *et al.*, "Achorde: Communication-aware multi-robot coordination with intermittent connectivity," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10184–10191, 2022.
- [9] 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.
- [10] K. Griparic, M. Polic, M. Krizmancic, and S. Bogdan, "Consensus-based distributed connectivity control in multi-agent systems," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 3, pp. 1264–1281, 2022.
- [11] B. Nguyen, T. X. Nghiem, L. Nguyen, H. M. La, and T. Nguyen, "Connectivity-preserving distributed informative path planning for mobile robot networks," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2949–2956, 2024.
- [12] A. Derbakova, N. Correll, and D. Rus, "Decentralized self-repair to maintain connectivity and coverage in networked multi-robot systems," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 3863–3868.
- [13] W. Luo and K. Sycara, "Minimum k-connectivity maintenance for robust multi-robot systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 7370–7377.
- [14] 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.
- [15] J. Wang, M. Guo, and Z. Li, "Multi-agent coordination under temporal logic tasks and team-wise intermittent communication," in *IEEE Conference on Decision and Control (CDC)*, 2023, pp. 3174–3179.
- [16] Y. Kantaros and M. M. Zavlanos, "Distributed intermittent connectivity control of mobile robot networks," *IEEE Transactions on Automatic Control (TAC)*, vol. 62, no. 7, pp. 3109–3121, 2016.
- [17] Z. Ren, C. Zhang, S. Rathinam, and H. Choset, "Search algorithms for multi-agent teamwise cooperative path finding," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [18] M. Yang, K. Zhao, Y. Wang, R. Dong, Y. Du, F. Liu, M. Zhou, and L. H. U, "Team-wise effective communication in multi-agent reinforcement learning," *Autonomous Agents and Multi-Agent Systems*, 2024.
- [19] G. L. O. Solver, <https://developers.google.com/optimization/lp>.
- [20] G. Zijlstra, K. L. Aplin, and E. R. Hunt, "Multi-robot strategies for communication-constrained exploration and electrostatic anomaly characterization," in *International Conference on Space Robotics (iSpara)*. IEEE, 2024, pp. 76–83.