

# Multi-Agent Pathfinding Under Team-Connected Communication Constraint via Adaptive Path Expansion and Dynamic Leading

HOANG-DUNG BUI\*, George Mason University, USA

ERION PLAKU, U.S. National Science Foundation, USA

GREGORY J. STEIN, George Mason University, USA

This paper proposes a novel planning framework to handle a multi-agent pathfinding problem under a team-connected communication constraint, where all agents must have a connected communication channel to the rest of the team during their entire movements. Standard multi-agent pathfinding approaches (e.g., priority-based search) have potential in this domain but routinely fail when neighboring configurations at start and goal differ. Their single-expansion approach—computing each agent’s path from the start to the goal in just a single expansion—cannot reliably handle planning under communication constraints for agents as their neighbors change during navigating. Similarly, leader-follower approaches (e.g., platooning) are effective at maintaining team communication, but fixing the leader at the outset of planning can cause planning to become stuck in dense-clutter environments, limiting their practical utility. To overcome this limitation, we propose a novel two-level multi-agent pathfinding framework that integrates two techniques: *adaptive path expansion* to expand agent paths to their goals in multiple stages; and *dynamic leading* technique that enables the reselection of the leading agent during each agent path expansion whenever progress cannot be made. Simulation experiments show the efficiency of our planning approach, which can handle up to 25 agents across five environment types under a limited communication range constraint and up to 11–12 agents on three environments types under line-of-sight communication constraint, exceeding 90 % success-rate where baselines routinely fail.

**JAIR Track:** Multi-agent Path Finding

**JAIR Associate Editor:** Sven Koenig

## JAIR Reference Format:

Hoang-Dung Bui, Erion Plaku, and Gregory J. Stein. 2026. Multi-Agent Pathfinding Under Team-Connected Communication Constraint via Adaptive Path Expansion and Dynamic Leading. *Journal of Artificial Intelligence Research* 85, Article 46 (April 2026), 21 pages. DOI: [10.1613/jair.1.19915](https://doi.org/10.1613/jair.1.19915)

## 1 Introduction

We want a team of agents to navigate through an obstacle-rich environment to reach goals while maintaining constant team communication: a spanning tree created from range-limited or line-of-sight communication between pairs of agents. This problem is relevant to scenarios such as supply delivery during disasters or monitoring hostile environments, where maintaining constant communication can be a means of ensuring resiliency, allowing the team to immediately plan anew if the mission or environment were to change. In another scenario, such as wildfire monitoring, autonomous vehicles may form a dynamic perimeter around a spreading front within pre-mapped terrain, ensuring the network remains connected despite range-limited radio

\*Corresponding Author.

Authors’ Contact Information: Hoang-Dung Bui, ORCID: [0000-0003-1901-7120](https://orcid.org/0000-0003-1901-7120), [hbui20@gmu.edu](mailto:hbui20@gmu.edu), George Mason University, Fairfax, Virginia, USA; Erion Plaku, ORCID: [0000-0002-6622-386X](https://orcid.org/0000-0002-6622-386X), [eplaku@nsf.gov](mailto:eplaku@nsf.gov), U.S. National Science Foundation, Alexandria, Virginia, USA; Gregory J. Stein, ORCID: [0000-0003-1981-4154](https://orcid.org/0000-0003-1981-4154), [gjstein@gmu.edu](mailto:gjstein@gmu.edu), George Mason University, Fairfax, Virginia, USA.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.19915](https://doi.org/10.1613/jair.1.19915)




Fig. 1. **Fixed leader and single-expansion of agent paths fail in scenarios a) & b). Our proposed techniques: adaptive path expansion and dynamic leading handle both (c).** There are 3 agents (starting at  $s_1$ ,  $s_2$ , and  $s_3$ ) finding paths (solid lines) to their goals  $g_1$ ,  $g_2$ , and  $g_3$ , respectively while maintaining team communication connectivity (purple dotted lines). **Planning order** for (b) & (c): Agent 3  $\rightarrow$  Agent 1  $\rightarrow$  Agent 2.

constraints. In these high-stakes missions, a single link failure can partition the team, leading to a critical loss of coordinated decision-making. There is thus a need for algorithmic solutions to this problem, which can quickly and effectively plan trajectories for an entire team to reach its target destinations while constrained to maintain team communication.

Maintaining the spanning tree while the agents head in different directions with varied lengths of actions makes pathfinding challenging in continuous time and space, even for holonomic agents. The challenge is compounded by agents having varying neighboring configurations between start and goal positions which limits the efficacy of selecting a single *leader agent* that other agents can follow. Furthermore, agents must navigate through narrow passages and non-convex spaces without collisions to reach the goal configurations. Collectively, we refer this problem domain as Multi-agent Pathfinding under Team-Connected Communication Constraint (MAT3C).

The problem can theoretically be solved using composite state approaches (Standley 2010; Wagner and Choset 2015), in which planning selects between team actions, each simultaneously specifying an action for all agents at once. However, this approach quickly suffers from the curse of dimensionality, as the difficulty of planning increases exponentially with the number of agents, making planning intractable for even relatively small problems.

The decoupled approach (Andreychuk, Yakovlev, Boyarski, et al. 2021; Andreychuk, Yakovlev, Surynek, et al. 2022; Ma et al. 2019; Okumura, Bonnet, et al. 2023; Okumura, Machida, et al. 2022; Sigurdson et al. 2018; Solovey et al. 2016) mitigates this challenge by planning each agent in sequence and resolving conflicts among paths by enforcing constraints and replanning. However, incorporating the team communication constraint would require considerable algorithmic modifications to these planners, especially when an agent's local neighbors must change between the start and goal configurations, as in Figure 1b.

Another approach is *platooning leader-follower* (Gao et al. 2019; Huang et al. 2019; Qian et al. 2016; Shojaei and Yousefi 2019) that handles the communication constraint by selecting a leader that has no communication constraint and others act as followers, who plan so as to maintain communication with the agent that planned before them. However, establishing a fixed planning order for the leader agent can result in issues during planning. These issues include the follower agents finding no action to follow the leader within the communication range or becoming stuck once the leader reaches its goal (Figure 1a), or the team needing to spread out for the followers to reach their goals while maintaining the communication constraint (Figure 1b).

To address these challenges, this work proposes a novel planning algorithm for solving MAT3C, ensuring a connected communication channel in the team at any time from start to goal positions. Our planning algorithm integrates two techniques: *adaptive path expansion*, which extends agents' paths across multiple expansions as the agents' progress stalls or they encounter conflicts; and *dynamic leading*, which allows a new agent to assume the leader role if expanding via the current leader cannot make progress, enabling further expansion of each single-agent trajectory.

From these techniques, we have developed a planner, called Multi-Agent Pathfinding with Adaptive Path Expansion and Dynamic Leading (APEDL), for planning in multi-agent systems under team connected-communication constraint. We show our planner is effective under two types of communication constraint: *limited communication range* (LCR) and *line-of-sight communication* (LOS). Simulated experiments show that our planner results in fast and effective planning, handling up to 25 agents under LCR constraint across five environment types, and up to 12 agents with Random Forest and Rings environments, and up to 3–10 agents for Office, Waves, and Maze environments under LOS constraint with more than 90% success-rate where baselines routinely fail.

The paper is structured as follows: Section 1 introduces the MAT3C problem and its challenges. Section 2 discusses existing approaches, their potential to address the problem, and the obstacles they face. Section 3 presents the problem formulation. Section 4 provides an overview of our contributions, while Section 5 details the planning algorithm. Section 6 discusses the experiments and results, Section 7 analyzes the planner's completeness, and Section 8 evaluates our contribution and outlines future work.

## 2 Related Work

There exist multiple approaches that seek to solve multi-agent pathfinding problems, in the absence of a communication constraint, that could in theory be adapted to solve the MAT3C problem. To effectively scale to more agents, approaches in this space are typically designed around reducing the exponential growth of the planning tree with planning time horizon. One class of planners limits state space expansion by introducing composite states only if necessary. Standley (2010) proposed two techniques: *operator decomposition* (OD) reduces the branching factor by allowing only one agent to select an action at each timestep; and *independent detection* (ID), a collision-checking and replanning mechanism that determines whether composite states are required for planning. Wagner and Choset (2015) introduced the  $M^*$  algorithm, which used *individual policies* to guide agents for next actions if no collision occurs, yet builds composite states for agents by *collision set* when their actions are in collision. Though  $M^*$  and OD-ID algorithm work well if the agents are far apart, when agents are close to each other as in the MAT3C problem, collisions occur frequently, causing the planners eventually to use joint-state and struggle to handle the exponential growth of joint-state space; thus limiting the efficacy of these approaches for solving our limited-communication problem, which requires the agents to be close to one another.

Another class of planners uses the decoupled approach that plans each agent individually, deploying additional techniques to resolve conflicts among agents and replanning as needed. Andreychuck et al. (2021; 2022) developed *continuous conflict-based search* (CCBS)—a variant of *conflict-based search* (CBS) in continuous time—that defines constraints with time period for conflicting agents. Ma et al. (2019) proposed *priority-based search* (PBS) that searches for a priority of planning order among agents, which results in conflict-free among the single agent paths. Okumura, Bonnet, et al. (2023) proposed the *offline time-independent multi-agent path planning* (OTIMAPP) planner, which employs *deadlock-based search* (DBS) to resolve conflict among paths determined through prioritized planning. Although CCBS, PBS and OTIMAPP are robust to compute paths for standard *multi-agent pathfinding* (MAPF) problem, their approaches of single-expansion for agent paths—computing each agent's path from the start to the goal in just one expansion—fail to solve MAT3C problem in Figure 1b scenario.

*Priority inheritance with backtracking* (PIBT) (Okumura, Machida, et al. 2022) and *Bounded multi-agent A\** (BMAA\*) (Sigurdson et al. 2018) expand a single action for each agent in each planning iteration. These planners

can handle well deadlocks among the agents. However, there are two issues for these planners to solve the MAT3C problem. The first one is their myopic action selection strategy, which can result in situations where followers are unable to find actions that keep them within the leader's communication range. The second one is their fixed planning order causing followers to fail to plan when their goal positions are out of communication with the leader's goal position (Figure 1b). Wang and Botea (2011) introduced the MAPP planner, which requires that sufficient space exists for agents to trade places with one another to guarantee completeness, and also does not consider communication restrictions. In addition to the limitation of a fixed planning order (as mentioned above), MAPP would require additional algorithmic changes to modify the computed paths while adhering to deadlock and communication constraints.

Choudhury et al. (2022) proposed an online FV-MCTS-MP planner, which utilized coordinate graphs and the max-plus algorithm to address the challenge of action space growth as the number of agents increases. However, in the MAT3C problem, neighbors are dynamic, and the agent team only needs to maintain a spanning tree rather than a fully connected graph. These features complicate the setup of local payoff components, limiting the planner's efficiency in solving the MAT3C problem and needlessly restrict the space of possible solutions. Solovey, Salzman, and Halperin (2016) introduced the DRRT planner, which uses an implicit composite roadmap—a product of single-agent roadmaps. Their algorithm grows a multi-agent tree in this composite roadmap using RRT. However, using such a roadmap requires abstracting low-level details of each robot's motion over time, making it difficult to apply their approach for planning with communication constraints. In general, it is very challenging and not straightforward to adapt these approaches to solve the MAT3C problem in which each member of the team must be in constant communication with one another in continuous time.

Several algorithms (Bhattacharya et al. 2010; Díaz-Báñez et al. 2017; Pal et al. 2012) were developed for limited versions of the team communication constraint. Bhattacharya et al. (2010) employed a soft constraint with incremental penalties to enforce communication constraint on pairs of agents at specific predefined points en route to goals. However, the MAT3C problem requires the agents to maintain continuous team communication. It is also challenging to apply soft constraints on the entire agent team because the constraints of later agent pairs can break the communication constraints of earlier pairs. Pat et al. (2012) introduced a planning framework for exploration, but their team communication constraint applies only at agents' frontiers and goals, not during movement. To the best of our knowledge, no MAPF literature addresses the MAT3C problem.

Formation control (Aljassani et al. 2023; Garrido et al. 2013; Kowdiki et al. 2019; Qian et al. 2016; Sehn and Collier 2024) is a straightforward method to address the MAT3C problem if relaxing the formation requirement. In this approach, leader agents are planned first, followed by the paths for follower agents, which maintain predefined formations while following the leaders that ensure constant communication among the agents. However, this approach encounters challenges when the relative positions of the agents at start and goals are randomly arranged.

Specific to the MAT3C problem, *platooning* has emerged as a state-of-the-art solution (Shojaei and Yousefi 2019) in which the motion of a leader-agent is planned first and follower-agents planned one at a time in sequence to follow the agent that preceded them, and so is designed to maintain communication constraints between pairs of agents. Most existing work in platooning performs full motion planning only for the leader and uses a low-level controller to regulate followers to maintain a distance from the preceding agents (Agachi et al. 2024; Huang et al. 2019; Shojaei and Yousefi 2019). Zhao et al. (2017), and Gao et al. (2019) introduced a model predictive controller that generates trajectories for a virtual center, which followers need to track. However, recent work in the space of CBS (Ma et al. 2019) has proven that planning with a fixed vehicle order in general, including platooning, is incomplete and frequently becomes stuck when the leader moves in a direction different to the follower's goals, as we show in Figure 1a.

To overcome the limitations of using a fixed leader in a planning iteration and handling neighbor changing as navigating from start and goal, a planner should integrate two key capabilities: *dynamic leading*, in which the leader can be changed in a single-agent path expanding when the team fails to make progress towards the goals;

and *adaptive path expansion*, that allows an agent to resume its planning in next path expansions if expansion stalls.

### 3 Problem Definition

We formally define the MAT3C problem as follows. There are  $n$  agents  $\{a_1, \dots, a_n\}$  and a known map of the world  $\mathcal{W}$  with obstacles  $\mathcal{O} = \{O_1, \dots, O_o\}$ . The agents start at the initial positions  $s^{\text{init}} = \{s_1^{\text{init}}, \dots, s_n^{\text{init}}\}$  and head to the goals  $g = \{g_1, \dots, g_n\}$ , where  $s_i, g_i \in \mathcal{W}$  and  $s_i, g_i \notin \mathcal{O}$ . The configurations of initial positions and goals are chosen randomly but such that they still satisfy a team communication constraint, defined below.

The world  $\mathcal{W}$  is divided into sub-divisions  $\Delta$ , which are obstacle-free regions. A graph  $\mathcal{G}$  is created where the vertices are the centroids of  $\Delta$  and the edges represent connections of neighboring sub-divisions. An agent moves to a neighboring sub-division with constant velocity  $v_c$  by taking the action:  $Move(v, u)$ , where  $v, u$  are neighbor vertices in  $\mathcal{G}$ . As reaching its goal, the agent takes the action of  $Stop = Move(v, v)$  and comes to rest at its goal location. Due to the communication constraint, subsequent agents still need to check collision with agents at their goal locations. Agents moving over  $\mathcal{G}$  are guaranteed to be *collision-free* with respect to the static obstacles.

We define two levels of communication constraint for the agents. The first level is between two agents, referred to as *agents communication constraint* (ACOMM). The existence of ACOMM is defined by:

$$\text{ACOMM} = \begin{cases} 1 & \text{if IsConnected}(p_i(t), p_j(t)) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Where  $p_i(t)$  and  $p_j(t)$  are the positions of agents  $i$  and  $j$  at time  $t$ .

The function  $\text{IsConnected}(p_i(t), p_j(t))$  depends on communication types. In this work, we consider two different means of establishing communication between two agents: The first one is the *limited communication range* (LCR) constraint: if the Euclidean distance between the two agents' positions is less than or equal to the communication range  $r_c$ .

$$\text{IsConnected}(p_i(t), p_j(t)) = |p_i(t) - p_j(t)| \leq r_c \quad (2)$$

The second type of constraint is the *line-of-sight communication* (LOS) constraint: the two agents have an unobstructed line of sight to each other.

$$\text{IsConnected}(p_i(t), p_j(t)) = \text{LoS}(p_i(t), p_j(t), \mathcal{O}) \quad (3)$$

where  $\text{LoS}(\dots, \mathcal{O})$  is a boolean *Line-of-Sight* function that returns true if the straight line between  $p_i$  and  $p_j$  does not intersect any obstacle in the set  $\mathcal{O}$ .

The second level of communication constraint applies to the entire team of agents and is referred to as *team communication constraint* (TCOMM). This constraint requires the agents to form a connected communication graph in which the nodes represent the agents and the edges correspond to pairwise connections that satisfy ACOMM. In our work, we verify the existence of TCOMM by computing a spanning tree over the agent communication graph.

An action  $Move(v, u)$  satisfies the ACOMM constraint if during the movement  $Move(v, u)$  the agent has at least one neighboring agent within a distance of  $r_c$ . Due to varied lengths of actions and TCOMM requirement, time and agents' positions are continuous in MAT3C, even though the actions are discrete.

A **collision** occurs between two agents,  $a_i$  and  $a_j$ , when their distance is less than a threshold  $d_c$  at timestep  $t$ . A **path** is a sequence of waypoints to transition an agent from a position  $p_s$  to position  $p_g$  with constant velocity  $v_c$ . We say a path is **reach-goal** if it can lead the agent to the goal, is collision-free, and the movement between two sequential waypoints satisfies ACOMM constraint; When an agent stops at its goal, we still consider collision

and ACOMM constraints. A path is **valid** if it is *reach-goal* and after reaching the goal at time  $t_g$  forever stays at the goal, satisfies ACOMM, and is collision-free.

The objective is to compute *valid* paths  $\{\zeta_1, \dots, \zeta_n\}$ , one for each agent, so that  $\zeta_i$  starts at  $s_i^{\text{init}}$ , reaches  $g_i$ , and the agents satisfy TCOMM constraint from the start time to the time the last agent reaches its goal. We develop a MAPF planner that seeks to reduce the overall planning time and the sum of travel distances.

## 4 Overview of Contributions

To address the MAT3C problem, we need to solve several challenges, including maintaining a team communication channel at any time, changing of neighbors as navigating, and planning getting stuck due to the leader agent stalling. We propose two methods, *adaptive path expansion* (APE) and *dynamic leading* (DL), to resolve these challenges. This section is organized as follows: Section 4.1 introduces APE—that allows multistage growing of agent paths; Section 4.2 presents DL, which enables seamless leader transitions within each single agent path expansion, ensuring planning progress with minimal reliance on planning-order and leader selection. Section 4.3 discusses a *team communication tree* (TCT) that manages planning progress and ensures the leader agent maintains team communication in case some of its neighbor agents have already reached their goals.

### 4.1 Adaptive Path Expansion

In our MAT3C experiments, the start and goal positions are randomly chosen while adhering to the TCOMM constraint. Agents have different neighbors at the start and the goal, forcing them to rearrange and change their nearest neighbors as they navigate from the start to the goal positions. When an agent follows the leader—who by definition is not affected by the ACOMM constraint—to a branching point where both head to different goals, the follower agent will stall because no action satisfies the ACOMM constraint (Figure 1b). If deploying single-expansion of agent paths (computing each agent’s path from the start to the goal in a single expansion) as in PBS, CCBS, OD-ID, OTIMAPP, or *platooning* approaches, the planning fails because no feasible path can be found during the current expansion (Figure 1b). If each agent selects a single action based on its own interests in each iteration, as in PIBT and BMAA\*, it would lead to situations where a follower cannot find a feasible action to maintain communication constraint with its leader (e.g. due to obstacles), resulting in failure of planning.

To address this issue, we introduce *adaptive path expansion* (APE): a method that enables agents to iteratively refine and expand their paths multiple times if they fail to reach their goals in prior expanding attempts (Fig. 2.1–5). The term *adaptive* refers to two features: (1) that the number of path expansions can vary for each agent, with additional attempts to expand the plan potentially resolving conflicts borne of communication constraint violations, and (2) that the planned paths can be modified if *collision-at-goal* situations occur in which an agent coming to rest at its goal location impedes progress of other agents along their own planned paths (Bui et al. 2024). When expanding, different from PBS and OTIMAPP, the planned paths in our planning framework must satisfy the ACOMM constraint with prior planned paths. As *collision-at-goal* situations happen, rather than replanning the entire paths, we refine the prior planned paths by trimming them to the point of collision and then re-expanding from that point. This approach avoids breaking the communication constraints of follower agents, which would otherwise necessitate replanning for all subsequent agents. The APE alone does not fully address the MAT3C problem if the leader stalls or achieves goals before the follower agents. To resolve this, we propose the second technique *dynamic leading*.

### 4.2 Dynamic Leading

State-of-the-art MAPF planners, such as PBS, OTIMAPP, and PIBT, allow the planning order to be modified before each planning iteration but keep it fixed during the iteration. This fixed planning order with the first agent as the

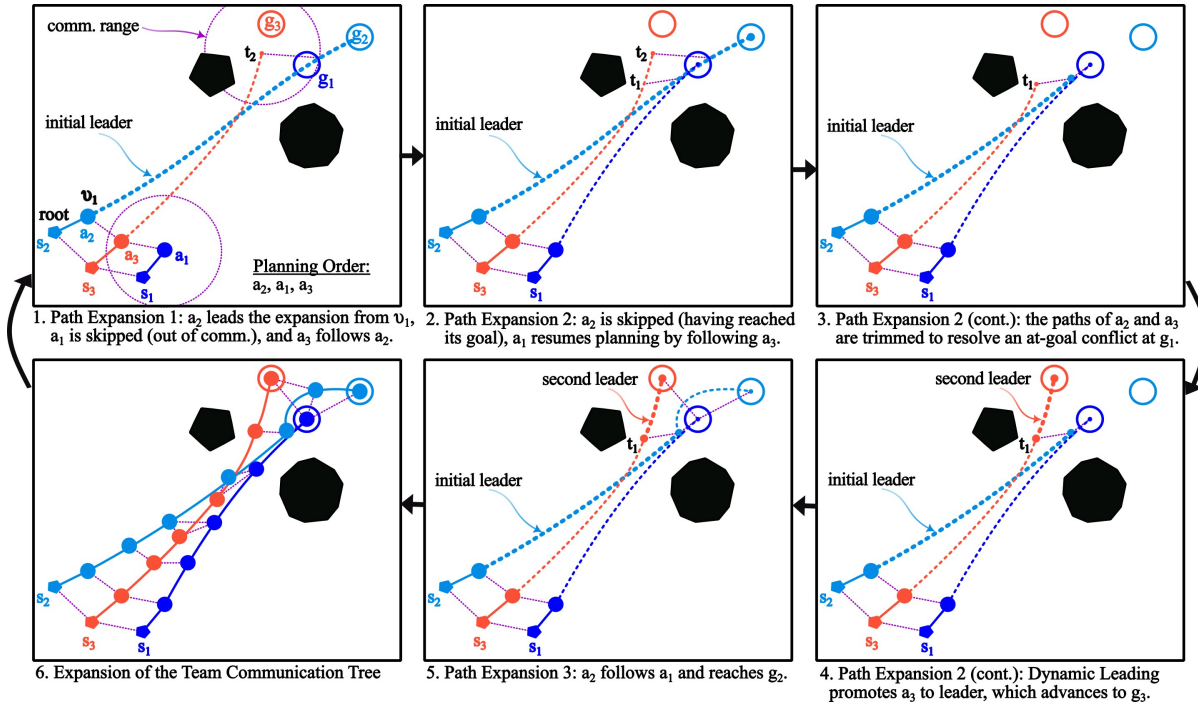


Fig. 2. **An illustration of one-step TCT's expansion with three agents.** Agents  $a_1, a_2,$  and  $a_3$  start at  $s_1, s_2,$  and  $s_3$ , and aim to reach their respective goals  $g_1, g_2,$  and  $g_3$  while avoiding obstacles (black) and each other. In **step 1** (Path Expansion 1),  $a_2$  leads and moves toward  $g_2$  without communication constraints, passing through  $g_1$ . Agent  $a_1$  is skipped as it cannot establish a communication link with  $a_2$ . Agent  $a_3$  follows  $a_2$  but stops at  $t_3$  after losing communication. In **step 2**,  $a_2$  is skipped since it has reached its goal. Agent  $a_1$  resumes planning by following  $a_3$  and reaches  $g_1$  at  $t_1$ . However, stopping at  $g_1$  creates a conflict with  $a_2$ , whose path passes through this location at a later time. In **step 3**, this conflict is resolved by trimming all paths beyond  $t_1$ . In **step 4**, *Dynamic Leading* promotes  $a_3$  to leader, allowing it to advance to  $g_3$ . In **step 5**,  $a_2$  is replanned to reach  $g_2$  while avoiding  $g_1$ . **Steps 1–5** constitute our second contribution: *adaptive path expansion*. From the resulting trajectories, **step 6** expands the team communication tree  $\mathcal{T}_{tc}$ .

leader can lead to failures in solving the MAT3C problem as followers must continuously communicate with the leader, causing them to stall if the leader halts (Figure 1a).

To address the issue, we propose *dynamic leading* (DL), a technique that dynamically reassigns the leader role to the agent that has progressed the farthest for guiding the team within the *same planning order* during single-agent path expansion. Choosing a new leader based solely on position or proximity to goals can cause the team to become stuck, especially when the current leader reaches its target. The DL method allows an agent to take the leader role if it has made the most temporal progress in planning. Leader switching occurs when the planning agent takes an action at time  $t$  while other agents have not yet reached  $t$  (Fig. 2.4–5), allowing the planning agent to assume the leader role and planning the action without adhering to the ACOMM constraint. Experiments demonstrate that the DL technique performs efficiently, even with randomized selection of planning orders and leaders.

### 4.3 Team Communication Tree

*Dynamic leading* allows to switch the leader that can move without ACOMM constraint. However, the new leader may expand to a new direction that disrupts communication with agents already at their goals. To address this challenge, we introduce a *team communication tree* (TCT) (Fig. 2) to manage planning progress and avoid recomputing the entire agent paths from the starts. If the leader's actions break the team's communication channel, the paths under its leadership are excluded from the tree. Instead of computing the paths from the starts, the TCT tree allows the planner to recompute from the closest node on the TCT for replanning. The TCT comprises agent paths, with nodes containing the agent locations, a timestep, an initial  $f$ -value assigned as follows:

$$f = \alpha g + (1 - \alpha)h, \quad (4)$$

where  $g$  and  $h$  are the sum of  $g_i$  and  $h_i$  values from all agents.  $\alpha$  is a weight factor that balances the trade-off between prioritizing shorter paths and faster goal convergence ( $\alpha = 0.1$  in our experiments).

A node is added into the tree, if the agent's positions at the node satisfy the TCOMM constraint. An agent's position at a given time is interpolated from its path based on time. Due to continuous time under the TCOMM constraint and varied length of actions, the interpolated agent's position can fall anywhere along its paths, and thus anywhere within a sub-division. Consequently, expanded paths can result in agent positions being in continuous-space locations throughout the workspace. If a node is selected to expand (Alg. 1.4), its  $f$ -value is penalized to encourage the planner to choose other nodes in the next iteration.

From the proposed techniques, we develop the planner APEDL which combines all three of these techniques to address the MAT3C problem.

## 5 Planning Framework

Our planning framework, which we call *Multi-agent Pathfinding with Adaptive Path Expansion and Dynamic Leading* (APEDL), consists of two levels: High Level, which manages multi-stage growing of agents paths, and Low Level: Single-Agent Pathfinding with Dynamic Leading (SAPF-DL). The major components of the framework are described as follows. Section 5.1 introduces the High Level (Alg. 1, Figure 2) that consists of selecting a node from the tree (TCT), expanding agent's paths and modifying them if necessary via APE, then growing the tree until reaching the goals or running out of runtime. Section 5.2 presents the Low Level: a single-agent pathfinding with dynamic leading (Alg. 2) that continuously expands the shortest collision-free and ACOMM-satisfied paths to goals for the agents one by one. Section 5.3 describes how the tree TCT is expanded, and Section 5.4 analyzes the time complexity of the planner.

### 5.1 High Level: Growing the Team Communication Tree with Adaptive Path Expansion

The high level grows a TCT with nodes composed of agent paths. The nodes are collision-free and each have a configuration of agent positions that includes a spanning tree of communication among the agents. A node consists of a timestep, positions of agents,  $f$ -value (defined by Eq. 4). The  $f$ -value is increased by a penalty factor when a node is reselected (in `SELECTNODE()` function at line Alg. 1.4) to encourage the planner to explore other nodes. The high-level module develops the tree TCT  $\mathcal{T}_{tc}$  by the following steps: (1) selecting a node to expand, (2) randomizing the order in which to plan trajectories for the agents, (3) running the APE loop, then (4) growing the tree from the returned paths.

The module gets the number of agents  $n$ , the goals  $g$ , initial start  $s^{\text{init}}$ , and the world  $\mathcal{W}$  as inputs and returns the *valid* paths or, if valid paths cannot be found, those closest to the goals. The algorithm (Alg. 1) starts by initializing a TCT  $\mathcal{T}_{tc}$  with the root consisting of all initial positions  $s^{\text{init}}$  (Alg. 1:1). The world  $\mathcal{W}$  is divided into sub-divisions  $\Delta$  of a predefined area, with obstacle-overlapping sub-divisions further subdivided along their largest dimension until they are clear or reach a size threshold. A graph  $\mathcal{G}$  is created from the centroids of

**Algorithm 1** High Level: Growing a Team Communication Tree

---

INPUT:  $n$ : number of agents,  $g = \{g_1, \dots, g_n\}$ : goals,  $\mathcal{W}$ : world,  $s^{\text{init}} = \{s_1^{\text{init}}, \dots, s_n^{\text{init}}\}$ : starts,  $t_{\text{ds}}$ : max runtime  
OUTPUT: paths  $\zeta = \{\zeta_1, \dots, \zeta_n\}$

```

1:  $\mathcal{T}_{\text{tc}} \leftarrow \text{INITTCTREE}(s^{\text{init}}, t = 0)$ ;  $\Delta \leftarrow \text{SUBDIVISION}(\mathcal{W})$ ;  $\mathcal{G} \leftarrow \text{CREATEGRAPH}(\Delta, g)$ 
2:  $\mathcal{H} \leftarrow \{\text{CALHEURISTIC}(\mathcal{G}, g_1), \dots, \text{CALHEURISTIC}(\mathcal{G}, g_n)\}$ 
3: while  $\text{TIME}() < t_{\text{ds}}$ 
4:    $v \leftarrow \text{SELECTNODE}(\mathcal{T}_{\text{tc}})$ ;  $p \leftarrow \text{RANDOMORDER}()$ ;  $\zeta^s \leftarrow \text{INITPATHS}(v)$  //planning iter.(line:4-14)
5:   for  $m$  iterations //adaptive path expansion (line:5-12)
6:     allRG  $\leftarrow$  true
7:     for  $1 \leq i \leq n$  //single-expansions for agent paths (line:7-11)
8:       if  $\text{FOUNDGOAL}(\zeta_{p_i}^s)$  then continue
9:          $\zeta^t \leftarrow \text{SAPF-DL}(p_i, g_{p_i}, \zeta^s, \mathcal{G}, \mathcal{H}_{p_i})$ ;  $\zeta_{p_i}^s \leftarrow \text{INSERT}(\zeta^t)$  //a single-path expansion
10:        if  $\text{FOUNDGOAL}(\zeta_{p_i}^s)$  then  $\text{MODIFYIFOVERLAP}(p_i, \zeta^s, \text{allRG})$ 
11:        else  $\text{allRG} = \text{false}$ 
12:        if  $\text{allRG}$  then break
13:         $\text{EXPANDTCTREE}(\mathcal{T}_{\text{tc}}, \zeta^s, v)$ 
14:        if  $\text{allRG}$  then break
15:         $\text{vid} = \text{CLOSESTNODETOGOAL}(\mathcal{T}_{\text{tc}})$ 
16: return  $\text{GETPATHS}(\mathcal{T}_{\text{tc}}(\text{vid}))$ 

```

---

the obstacle-free sub-divisions, while the edges represent the connections between neighboring sub-divisions (sharing boundaries or corners). We compute the shortest paths from all sub-divisions to each agent's goal  $g_i$  used as a planning heuristic by the function  $\text{CALHEURISTIC}()$  and store them in the array  $\mathcal{H}$  (Alg. 1:2), a one-time cost that greatly accelerates plan expansion. The main loop (Alg. 1:3–14) runs multiple planning iterations starting by selecting a node  $v$  with smallest  $f$ -value (defined by Eq. 4) from  $\mathcal{T}_{\text{tc}}$  (Fig. 2.1). This  $f$ -value is then increased by factor  $c$  ( $c = 0.05$  in our experiments) to encourage exploring other nodes. A planning order  $p$  is created by a random function  $\text{RANDOMORDER}()$ . A list of  $n$  paths  $\zeta^s$  is initialized from  $v$  (Alg. 1:4).

The *adaptive path expansion* (APE) technique—the *outer loop* (Alg. 1:5–12/ Fig. 2.1–5) with  $m$  expanding attempts to find valid paths for the agents from node  $v$  (Fig. 2.1). Making repeated planning attempts before backtracking enable paths to resume expansion from stalled positions and avoid local conflicts; our experiments (Figure 5) show that even just a few replanning attempts greatly helps to improve success rates with negligible additional computation. This outer loop begins by setting the variable  $\text{allRG}$  (which stands for *all reach goal*) to true (Alg. 1:6), then it goes into the *inner loop* (Alg. 1:7–11).

The inner *for* loop plans for each agent by calling the module  $\text{SAPF-DL}$  then inserting each returned path into the trajectory array  $\zeta_{p_i}^s$  (Alg. 1:9). (As demonstrating in Fig. 2, there are three inner loops consisting of step 1a, step 2-3-4, and step 5). If the returned path  $\zeta_{p_i}^s$  is *reach-goal*, its validity is checked by the function  $\text{MODIFYIFOVERLAP}()$ . The path is invalid if the agent stops at its goal at time  $t$  and blocks the future movement (after  $t$ ) of an already-planned agent (according to planning order  $p$ ), a situation called *collision-at-goal* shown in Figure 2.2–3. If this situation occurs, the function determines the collision time  $t_c$ , trims the already-planned paths (those which extend beyond  $t_c$ ) to  $t_c$ ; then sets  $\text{allRG}$  to false (Alg. 1:10). If the path is not *reach-goal*,  $\text{allRG}$  is set to false (Alg. 1:11).

After the planning for all agents has completed, the tree  $\mathcal{T}_{\text{tc}}$  is expanded from  $v$  by the function  $\text{EXPANDTCTREE}()$  using the agent's paths  $\zeta^s$  (Alg. 1:13). If all agents reach the goals ( $\text{allRG}$  is still true), the *while* loop is broken (Alg. 1:14), then the planner returns paths for all agents from  $\mathcal{T}_{\text{tc}}$  (Alg. 1:15–16).

**Algorithm 2** Low Level: Single-Agent Pathfinding with Dynamic Leading (SAPF-DL)

INPUT:  $\rho$ : agent's ID,  $g_\rho$ : goal,  $\mathcal{G}$ : init. graph,  $\zeta^s$ : previous planned paths,  $t_{sa}$ : max runtime,  $\mathcal{H}_\rho$ : shortest path heuristic

OUTPUT: a path  $\zeta_\rho$

```

1:  $\mathcal{G}^s = \text{MODIFYGRAPH}(\mathcal{G}, g_\rho, \zeta_\rho^s, \text{end})$ ;  $v_g \leftarrow \text{GETNODE}(g_\rho, \mathcal{G}^s)$ ;  $v_s \leftarrow \text{GETNODE}(\zeta_\rho^s, \text{end}, \mathcal{G}^s)$ 
2:  $\text{openList.ADD}(v_s)$ ;  $\text{closedSet} \leftarrow \{\}$ ;  $v_{\text{best}} \leftarrow v_s$ 
3: while  $\text{openList}$  not empty and  $\text{TIME}() < t_{sa}$ 
4:    $v \leftarrow \text{openList.POP}()$ ;  $\text{closedSet.ADD}(v)$ 
5:   if  $v = v_g$  then return  $\text{GETPATH}(v_s, v, \mathcal{G}^s)$ 
6:   for  $u$  in  $\text{GETNEIGHBORS}(v)$ 
7:     if  $u$  is not in  $\text{closedSet}$  then
8:        $u.\text{parent} \leftarrow v$ ;  $d_{uv} \leftarrow \text{DISTANCE}(v, u)$ ;  $u.g \leftarrow v.g + d_{uv}$ 
9:        $u.t \leftarrow v.t + d_{uv}/v_c$ ;  $u.h \leftarrow \text{GETHEURISTIC}(u, \mathcal{H}_\rho)$ ;  $u.f = u.g + u.h$ 
10:      if  $\text{ISACTIONVALID}(\rho, \zeta^s, u, v, r_c)$  then
11:         $\text{openList.ADD}(u)$ 
12:        if  $u = v_g$  then return  $\text{GETPATH}(v_s, u, \mathcal{G}^s)$ 
13:        if  $v_{\text{best}}.f > u.f$  then  $v_{\text{best}} \leftarrow u$ 
14:      else if  $v.g + d_{uv} < u.g$  then
15:         $u_1 \leftarrow \text{COPY}(u)$ ;  $u_1.t \leftarrow v.t + d_{uv}/v_c$ 
16:        if  $\text{ISACTIONVALID}(\rho, \zeta^s, u_1, v, r_c)$  then
17:           $u.\text{parent} \leftarrow v$ ;  $u.t \leftarrow v.t + d_{uv}/v_c$ ;  $u.g \leftarrow v.g + d_{uv}$ ;  $u.f = u.g + u.h$ 
18:          if  $v_{\text{best}}.f > u.f$  then  $v_{\text{best}} \leftarrow u$ 
19: return  $\text{GETPATH}(v_s, v_{\text{best}}, \mathcal{G}^s)$ 

```

**5.2 Low Level: Single-Agent Pathfinding with Dynamic Leading**

Under the continuous-time ACOMM constraint, Single-Agent Pathfinding with Dynamic Leading (SAPF-DL) planner—an A\*-based search algorithm—needs several key features. First, a timestep must be computed for each node added into the openList. The input also include the start and goal positions for each agent. Second, the ISACTIONVALID() function (Alg. 3) verifies whether the current action is in collision-free and satisfies ACOMM constraint in continuous time, i.e., along the entirety of each trajectory segment. Our proposed *dynamic leading* technique is integrated within this function to dynamically reassign the leader role to the planning agent if it is the most progressive one. Third, to improve the search efficiency in complex obstacle structures and cluttered spaces, we employ shortest paths to the goals from all sub-divisions as a heuristic, obtained via the GETHEURISTIC() function.

The inputs to this low-level planner include the agent's ID  $\rho$ , goal position  $g_\rho$ , a graph  $\mathcal{G}$  representing the environment, the planned paths of other agents  $\zeta^s$ , and the shortest path heuristic  $\mathcal{H}_\rho$ . The input graph  $\mathcal{G}$  must also include the agent's current position  $\zeta_\rho^s.\text{end}$  and its goal  $g_\rho$ . The modification process involves: (1) identifying the two subdivisions that contain the current position and goal, (2) locating the corresponding vertices in graph  $\mathcal{G}$ , and (3) updating the vertices' positions—originally the subdivision centroids—to the  $\zeta_\rho^s.\text{end}$  and the goal  $g_\rho$ , respectively. All these steps are implemented by the MODIFYGRAPH() function (Alg. 2.1). The variables  $v_s$  and  $v_g$  represent the current and goal nodes, respectively (Alg. 2:1). The node  $v_s$  is then added into openList, a priority queue that sorts elements based on their cost-to-goal. The closedSet and the variable representing the closest-to-goal node,  $v_{\text{best}}$ , are also initialized (Alg. 2:2).

The main search loop (Alg. 2:3–18) operates as follows. The node with lowest cost is popped out from the openList and inserted into the closedSet (Alg. 2:4). If this node is the goal, search terminates successfully (Alg. 2:5).

**Algorithm 3**  $\text{IsACTIONVALID}(\rho, \zeta^s, u, v, r_c)$  (Dynamic leading employed within this function)

---

```

1: lead  $\leftarrow$  true; comm  $\leftarrow$  false
2: for  $1 \leq i \leq |\zeta^s|$  //dynamic leading (line:2-3)
3:   if  $u.t \leq \zeta_i^s.\text{maxtime}$  then lead  $\leftarrow$  false; //is this agent most progressive?
4: if lead then lead  $\leftarrow$   $\text{IsCOMMATGOAL}(u, \zeta^s, \rho)$ 
5: for  $1 \leq i \leq |\zeta^s|$ 
6:   if  $\text{IsCOLLISION}(u, v, \zeta_i^s, \rho)$  then return false
7:   if  $\text{IsCOMMS}(u, v, \zeta_i^s, \rho, r_c)$  then comm  $\leftarrow$  true
8: return (lead || comm) //leader moves without acomm constraint

```

---

Otherwise, all of its neighbors are iterated through (Alg. 2:6–18) to identify valid nodes, compute or update their costs, parent nodes, and timestep, then insert them into openList. For each node  $u$  not in the closedSet, we check its validity by determining if the *move* action from parent  $v$  satisfies both the collision-free and the ACOMM constraints, as evaluated by the  $\text{IsACTIONVALID}()$  function (Alg. 2:10). If  $u$  is already in the closedSet, its estimated cost is updated only if the new cost is smaller than the existing one (Alg. 2:14-17). The  $\text{GETNEIGHBORS}()$  returns both physical neighbors and, if a *wait* action is included in the set of available actions, the current node. When the *wait* action is selected, the waiting time is one second, with a corresponding cost  $v_c$  (Alg. 2:6,8).

**IsACTIONVALID() Function.** This function (Alg. 3) checks the validity of action  $\text{Move}(v, u)$  for an agent  $\rho$  starting at time  $v.t$ . To be valid, the action must maintain ACOMM constraint and collision free as the agent implements the action. The ACOMM constraint is satisfied by one of two conditions: having a communication link with at least one neighbor or being a leader.

The *dynamic leading* technique allows the agent  $\rho$  to take the leader role if its action goes furthest in time (Alg. 3:2–3). The leader is exempt from the ACOMM constraint. The leader agent is determined by comparing the timestep of current action to the max-time of all the paths (Alg. 3:2–3). For agent-collision detection,  $\text{IsCOLLISION}()$  (Alg. 3:6) checks if any segmented point along the action  $\text{Move}(v, u)$  collides with the planned paths in  $\zeta^s$  (Alg. 3:5) at the corresponding timesteps. The action is also checked by function  $\text{IsCOMMS}()$  (Alg. 3:7) for the ACOMM constraint, that determines whether agent  $\rho$  can communicate with a neighbor during the move.

As deploying *dynamic leading*, a situation called *out-of-communication-at-goal* can occur, which is shown in Figure 3a. At  $t_i$  agent  $a_3$  is the leader, and reaches its goal in two steps ( $t_{i+1}$  and  $t_{i+2}$ ).  $a_2$  follows  $a_3$  and also stops at its goal by  $t_{i+2}$ . Agent  $a_1$  moves along  $a_3$  until  $t_{i+2}$ , then becomes the leader, planning without ACOMM constraints. Agents  $a_2$  and  $a_3$  reaches the goals and stop there, thus the TCOMM constraint breaks at  $t_{i+6}$ . To resolve this issue, we propose the function  $\text{IsCOMMATGOAL}()$ .

**IsCOMMATGOAL() Function.** The function  $\text{IsCOMMATGOAL}()$  guides  $a_1$  to remain in communication as shown in Figure 3b. It is triggered when the leader is changed (Alg. 3:4) to check if any of the new leader's neighbors have reached their destination. If so, the leader's action  $\text{Move}(v, u)$  must meet the ACOMM constraint with an at-goal agent. If it cannot, its leader status is revoked.

### 5.3 Expand Team Communication Tree

Function  $\text{EXPANDTCTREE}()$  (Alg. 4) expand the TCT  $\mathcal{T}_{tc}$  by integrating the agents' paths  $\zeta^s$ . Each node on  $\mathcal{T}_{tc}$  has an associated timestep,  $f$ -value, and positions are interpolated to make them match for agent paths. A new node must satisfy the TCOMM constraint (function  $\text{IsTCOMM}()$ ) to be added to the tree. The *for* loop (Alg. 4:2–3) collects all timesteps of the waypoints on the agent paths and inserts into a priority queue timeList.

The second *for* loop (Alg. 4:4–14) goes through all elements in timeList to create new nodes and adds valid nodes into the tree  $\mathcal{T}_{tc}$ . Each new node  $v_n$  inherits the travel cost from its parent with timestep  $t$  (Alg. 4:5). We

**Algorithm 4** EXPANDTCTREE( $\mathcal{T}_{tc}, \zeta^s, v$ )

---

```

1:  $v_p \leftarrow v$ ; timeList  $\leftarrow \{\}$  //priority queue
2: for  $0 \leq i < |\zeta^s|$ 
3:   for  $0 \leq j < |\zeta_i^s \cdot \text{times}|$  do timeList.INSERT( $\zeta_i^s \cdot \text{times}[j]$ )
4: for  $t$  in timeList
5:    $v_n \leftarrow \text{NEWNODE}()$ ;  $v_n(\text{parent}, g, t) \leftarrow (v_p, v_p.g, t)$ 
6:   for  $0 \leq j < |\zeta^s|$ 
7:     if  $t > \zeta_j^s \cdot \text{end.t}$  then
8:       if FOUNDGOAL( $\zeta_j^s$ ) then pos  $\leftarrow \zeta_j \cdot \text{end.pos}$ 
9:       else return
10:    else pos  $\leftarrow \text{GETPOSATTIME}(\zeta_j^s, t)$ 
11:     $v_n \cdot \text{state} \cdot \text{INSERT}(\text{pos})$ ;  $d = \text{DIST}(v_p \cdot \text{state}_j, v_n \cdot \text{state}_j)$ ;
12:     $v_n \cdot g += d$ ;  $\zeta_j^s \cdot \text{costgoal} -= d$ ;  $v_n \cdot h += \zeta_j^s \cdot \text{costgoal}$ 
13:    if ISTCOMM( $v_n$ ) then  $v_n \cdot f \leftarrow \alpha v_n \cdot g + (1 - \alpha) v_n \cdot h$ ;  $\mathcal{T}_{tc} \cdot \text{APPEND}(v_n)$ ;  $v_p \leftarrow v_n$ 
14:    else return

```


---

then go to each agent's path, interpolate to recover its position at time  $t$  (Alg. 4:7–12), and insert it into the state of node  $v_n$ . If  $t$  is larger than the max-timestep of the path  $\zeta_j^s$  (Alg. 4:7) and the path reaches the goal, the final position of  $\zeta_j^s$  is returned (Alg. 4:8). If the position is not the goal, the tree's expansion stops (Alg. 4:9).


If  $t$  is smaller or equal to the last timestep  $\zeta_j^s$ , the agent's position at  $t$  is interpolated by function GETPOSATTIME() (Alg. 4:10). The lines in Alg. 4:11–12 add the agent's positions, cost, and heuristic of each agent into the node state. The new node is then verified the TCOMM constraint (preventing the scenario in Figure 3a) before adding to the tree  $\mathcal{T}_{tc}$  (Alg. 4:13). If the TCOMM constraint is not satisfied, the tree expansion is terminated (Alg. 4:14).

#### 5.4 Algorithm Complexity Analysis

The runtime of the High-Level module (Alg. 1) is  $O(k)$  where  $k$  is the number of TCT expansions. For each expansion, the low level SAPF-DL (Alg. 2), based on A\* search, is called  $m \cdot n$  times, where  $m$  is the max number of single-agent path expanding attempts per agent (Alg. 1:5) and  $n$  is the number of agents. The A\* search with the shortest path heuristic has a complexity of  $O(|E| \cdot \log(|E|))$ , where  $|E|$  is the number of edges, modified by a valid action check whose runtime, dependent on collision checking and the communication model with worst case is  $O(n^2)$  for all pairwise constraint checks. Thus, the overall complexity is approximately  $O(k \cdot m \cdot n \cdot |E| \log |E| \cdot n^2) =$



(a) Out of communication at-goal situation



(b) IsCOMMATGOAL() function

#### Diagram Key

**Planning order:** Agent 3  $\rightarrow$  Agent 2  $\rightarrow$  Agent 1. The current leader—Agent 3—reaches its goal at  $t_{i+2}$ .

**(a) Out of communication at-goal:** Agent 1 resumes the leading role at  $t_{i+2}$  and proceeds without communication constraints, however, team connectivity breaks at  $t_{i+6}$ .

**(b) With IsCOMMATGOAL():** Agent 1's leading role is revoked since its neighbors have already reached their goals; it must therefore satisfy the communication constraint.

Fig. 3. Out-of-communication-at-goal situation (a) and how IsCOMMATGOAL() works (b).

$O(k \cdot m \cdot n^3 \cdot |E| \log |E|)$ . The performance of APEDL planner depends on the selection of TCT's nodes and whether the *out-of-communication-at-goal* situation occurs. The current heuristic—sum of agents' shortest paths—is efficient to identify nodes with the shortest travel distance to goals but fails to consider the communication constraints, making it ineffective at selecting nodes to prevent the *out-of-communication-at-goal* issue. In future research, it may be possible to develop a more informative, if more computationally demanding, heuristic to select nodes from TCT that incorporates communication costs and avoids the *out-of-communication-at-goal* situation.

## 6 Experiments and Results

Experiments are conducted on five obstacle-rich environments (Figure 4). The planning performance is measured by three metrics: (1) success-rate, (2) runtime, and (3) per-agent travel distance. Our planner, APEDL, is evaluated with five baselines: a centralized approach with composite states, a platooning leader-follower approach, modified versions for the MAT3C problem of three planners: PIBT, OD-ID, and PBS. All agents have 8 actions: to move in the four cardinal directions and in their diagonal directions. Although the actions are discrete, the agent's positions are continuous, with random starts, goals, and the interpolated positions from agent paths. We also compare the performance of our APEDL when a *wait* action is included (referred to as APEDL-wait) and without it (APEDL). We evaluate the planner's performance on two types of communication constraints: *limited communication range* (LCR) and *line-of-sight communication* (LOS); and the results are described in Section 6.2 and Section 6.3, respectively. We also test the planners versus runtime, goal configurations, and difficulty of environments.

### 6.1 Experimental Setup

**6.1.1 Baselines.** We have implemented five baselines: *centralized planning with composite state* (COMP), *platooning leader-follower* (PLF), *operator decomposition-independent detection* under TCOMM (ODID-COMM), PIBT under TCOMM (PIBT-COMM), and PBS under TCOMM (PBS-COMM). COMP is selected because it theoretically can solve the MAT3C problem. PLF is chosen because it is the state-of-the-art algorithm to handle MAT3C problem in the leader-follower approach. Meanwhile, OD-ID, PBS, and PIBT are state-of-the-art algorithms in standard MAPF instances, which were shown efficient to plan for a large number of agents. Both PBS and PIBT belong to decoupled approach with different single-agent path expansion strategies: single-expansion and single-action in each planning iteration. We select OD-ID to represent for the composite state approach over  $M^*$  because agents in  $M^*$  select a single action per planning iteration based on their own interests (similar to PIBT), which can cause follower agents to fail in maintaining communication constraints with the leader agent. Difference to general MAPF instances, in MAT3C problem, actions are only valid if they satisfy the ACOMM. Consequently, planners must establish planning orders, ensuring that later-planned agents maintain the ACOMM constraint with previously planned agents, and if one agent is replanned, all agents subsequent to it in the planning order must also be replanned. Thus, these planners are algorithmically modified to fit TCOMM requirement, and referred as ODID-COMM, PBS-COMM, PIBT-COMM, respectively.

- *Centralized planning with composite state* (COMP) grows a multi-agent tree with composite states of all agent's positions via  $A^*$  with a heuristic that sums shortest-to-goal paths from each agent. A node is added into the tree if it satisfies both TCOMM and collision-free among agents. This planner is the optimal planner but inefficient due to the exponential expansion of state-space with the number of agents.
- *Platooning leader-follower* (PLF) also builds a multi-agent tree and uses priority planning to grow the tree. At the root of tree expansion, the planning order is shuffled and so a random leader is chosen; planning order is then fixed for downstream nodes. The followers plan to follow another agent if its preceding agent reached its goal. If the current expansion fails after some iterations, the planning starts again at the root with a new random leader and a corresponding planning order.

- *Operator decomposition-independent detection* under TCOMM (ODID-COMM) initially computes individual paths for each agent, then uses the ID technique to identify conflicts and replan one agent's path to avoid collisions while keeping other paths fixed. If no valid path is found, the planner merges the conflicting agents and plan on their composite states with the OD technique. Different to OD-ID, ODID-COMM needs to select the first agent randomly to plan with ACOMM-free, requiring subsequent follower agent's actions to satisfy ACOMM constraint relative to the already planned paths. As planning on composite states, composite actions must satisfy the TCOMM constraint and collision free.
- *Priority inheritance with backtracking* (PIBT) is a state-of-the-art MAPF algorithm that uses the *priority inheritance* scheme to move agents by single actions respect to their own interest in each iteration. To work on the MAT3C problem, we modified PIBT by adding ACOMM into the action selection of the agents except the first planning agent. We refer to this planner as PIBT-COMM.
- *Priority-based search* (PBS) uses a priority tree to resolve conflicts among agents. To handle the communication constraint in the MAT3C problem, we develop a modified version of PBS, named PBS-COMM. In MAT3C, neighbors at starting positions are possible not neighbors at goals, thus PBS-COMM needs multiple iterations to find the feasible planning order for the entire agent team with ACOMM at the root node. In many cases, the planner fails because no feasible planning order exists or planning runs out of time with large number of agents. The second change is to add a planning order due to the communication constraint. That leads to a modification of replanning implementation: if an agent with priority  $p_i$  is replanned, all agents with lower priority than the agent are also recomputed due to the communication dependence. These two changes, necessary to apply PBS to the MAT3C problem, significantly increases runtime.

**6.1.2 Environments and Instances.** We evaluate APEDL in five obstacle-rich environment types: Random Forest, Office, Waves, Rings, and Maze (Figure 4). All environments are square shapes of size  $114\text{m} \times 114\text{m}$ , which are divided into sub-divisions with resolution of  $1\text{m} \times 1\text{m}$ . Without considering the obstacles, the sub-divisions create a graph with approximate 12700 vertices and 51000 edges. On each environment type, there are one hundred maps generated with random locations of obstacles.

For each number of agents, one instance is generated on each environment map. So, there are total of 12000 instances for the experiments. The start and goal configurations are chosen randomly, subject to the TCOMM constraint (satisfying both LCR and LOS constraints), on alternate sides of the maps (except the ring environment), distributed with a rectangular area. Communication between two agents depends on the communication model. For LCR, agents can communicate if the distance between them is within 15 m. For LOS, communication is established if the agents have an unobstructed view of each other.

*Env. Type 1: Random Forest.* Obstacles with random shapes and sizes are distributed randomly occupying 10% of the environment's area.

*Env. Type 2: Office.* The environments consist multiple rooms and hallways. Each room has a fixed width and varying length from 9–13 m. There are three long hallways along the building, 2–3 short hallways connecting them. Hallway width varies from 7–9 m.

*Env. Type 3: Waves.* The environments have wave-like obstacles which are separated at regular distances. Gaps are placed in each wave with random widths. The number of waves is set to 10.

*Env. Type 4: Rings.* The environments are featured by concentric rings with six breaks of random widths within 6–8 m. The separation between the rings is set to 8 m. For each instance, the starts are randomly placed at the center, and the goals are on one of four corners of the maps.

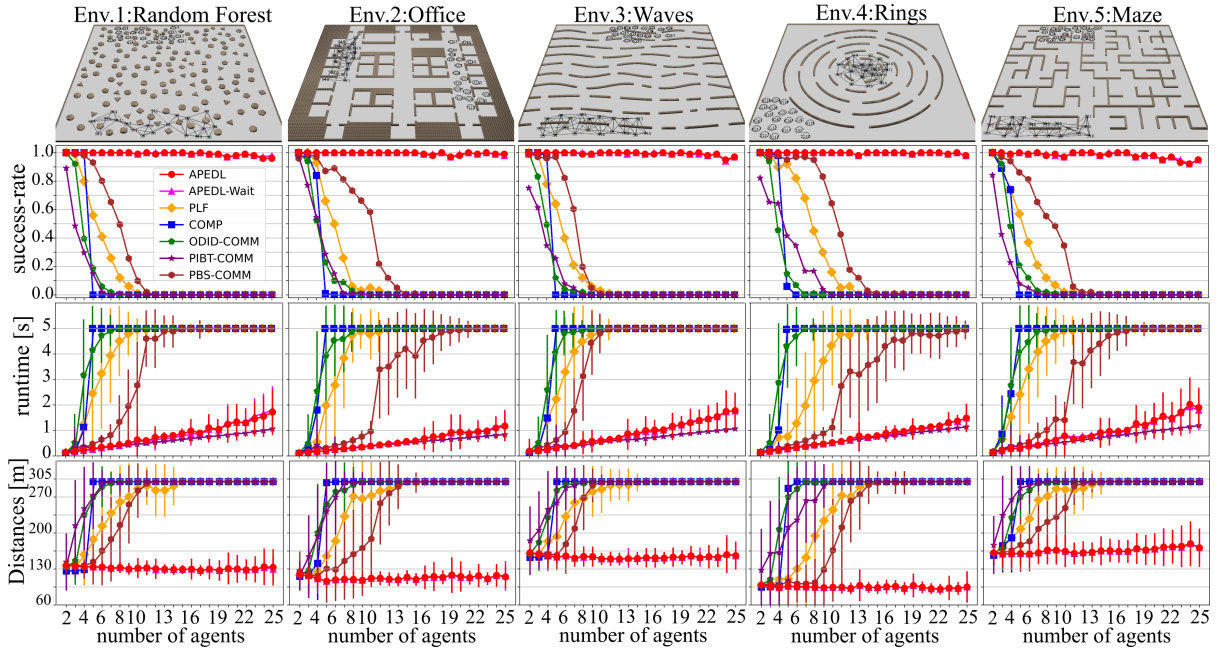


Fig. 4. **Success-rate, Runtime, and Travel Distances of our planners: APEDL and APEDL-E with PLF, COMP, ODID-COMM, PIBT-COMM, and PBS-COMM on 5 environment types under LCR constraint.** The comm. range is 15 m and max-runtime is 5 s. The runtime and travel distances are shown by means and standard deviations.

*Env. Type 5: Maze.* The environments consist of mazes generated by Kruskal’s algorithm with size of  $14 \times 14$ . To ease in generation of the starting and goal configurations, boundary walls connecting to the top-most and bottom-most rows are removed.

**6.1.3 Measuring Performance.** We evaluate APEDL and APEDL-E against COMP, PLF, ODID-COMM, PIBT-COMM, PBS-COMM baselines on 100 instances for each environment type and number of agents, measuring success-rate, average runtime, and average distance traveled per agent. Planning is considered successful if all agents reach their goals within 5 s of runtime. Runs that exceed that planning ceiling are assigned a travel distance of 300 m. The framework is evaluated with variations over number of agents, goal configurations, runtime, and environment difficulty levels.

**6.1.4 Computing Resources.** The experiments ran on HOPPER, a computing cluster provided by GMU’s Office of Research Computing. Each planning instance was run single threaded, yet experiments were run in parallel across 48 cores on a 2.40GHz processor. Our code was developed in C++ and compiled with g++-9.3.0.

## 6.2 Results Under Limited Communication Range Constraint

In these experiments, we deployed LCR as the communication constraint on the agent team. Communication between two agents is established if their distance is smaller than a threshold  $d_c$  ( $d_c=15$  m in the experiments). We evaluate the planner’s robustness versus the number of agents (Section 6.2.1), runtime (Section 6.2.4), environment difficulty (Section 6.2.3), and goal configurations (Section 6.2.5).

*6.2.1 Results when varying the agent number.* We evaluated the planners with 2–25 agents, allocating 5 seconds of runtime, and the results are shown in Figure 4. The APEDL planner performs well up to 25 agents exceeding 90% success-rate in all environments. The APE technique iteratively grows and refines the agent paths, enabling the APEDL planner to effectively handle arbitrary start and goal configurations. The DL technique allows leader to be reassigned in the same planning order during each planning iteration, making followers resume the leader roles and progress to their goals as the current leader stalls. We additionally show that the performance of APEDL with the *wait* action (APEDL-wait) is similar, suggesting that there is little benefit in our approach to having later agents pause during their motion.

All baseline planners struggle as the number of agents increases. PLF manages up to 5 agents in Rings Env. and only 3–4 agents in other environments with 90% success-rate. The heuristic of shortest-to-goal path sum guides COMP well to handle up to 4 agents achieving 90% success-rate in Random Forest, Rings, and Waves Env. With a small number of agents, COMP could compute optimal paths with shortest travel distances comparing to all other planners. However, with greater than 4 agents, the high dimensionality of the search space means that planning cannot reach the goal in the allotted time and success-rate quickly declines.

The PBS-COMM baseline outperforms the other baselines in the MAT3C problem. Its priority tree effectively determines feasible planning orders for conflicting agents and efficiently identifies the best order. However, relying solely on planning order is insufficient to fully address the MAT3C problem, as neighbors of agents change as navigating from start and goal configurations, often rendering feasible planning orders unattainable. Consequently, the performance of PBS-COMM decreases from 5–7 agents (with 90 % success-rate) in Waves, Random Forest, and Rings environments and with as few as 3–4 agents in the more constrained Office and Maze environments.

PIBT-COMM exhibits the poorest performance due to its myopic action selection, which prioritizes actions based on a shortest path heuristic in each planning iteration. This approach frequently leads to situations where follower agents cannot select actions to maintain the communication channel with the leader in subsequent iterations. As a result, PIBT-COMM failed even with as few as three agents across all tested environments.

The ODID-COMM algorithm first plans paths for individual agents, then addresses collisions by prioritized planning, and if necessary, employs the OD technique for composite states. In the MAT3C problem, where agents are closely positioned and collision are frequent, ODID-COMM allocates part of its runtime to single-agent planning and the remainder to composite-state planning. In contrast, COMP planner focuses entirely on composite states, leading to better performance in high-collision scenarios. Consequently, ODID-COMM generally underperforms compared to COMP planner in the MAT3C problem due to this runtime distribution. However, in simpler cases with up to five agents, ODID-COMM successfully completes planning, while the COMP planner completely fails due to the exponential growth of its composite state space.

In addition to average success rate, we also report the mean and standard deviation for runtime and travel distance in Figure 4. If a planner fails, the runtime and traveled distance penalties are applied to the results. The runtime and distance results aligned with the success-rate results across all planners except PIBT-COMM planner. The PIBT-COMM planner exhibits low runtime due to early termination when follower agents fail to maintain communication with the leader, limiting its effectiveness in the MAT3C problem. In contrast, as the number of agents increases, the PBS-COMM planner allocates significant runtime to computing a feasible ACOMM planning order from all possible agent sequence combinations, hindering its efficiency in resolving the MAT3C problem.

*6.2.2 Results when varying number of path expansions.* This experiment evaluated the impact of number of path expansions  $m$  to the performance of APEDL on two goal configurations within the Maze environment: rectangle shape and long-thin shape. The results include the success-rate and runtime and are illustrated in Figure 5. When the number of path expansion attempts is 1 ( $m = 1$ ), the APEDL only computes paths from each node of  $\mathcal{T}_{ic}$  only once (similar to PBS, OTIMAPP). With more than 13 agents or in challenging goal configurations, making only a

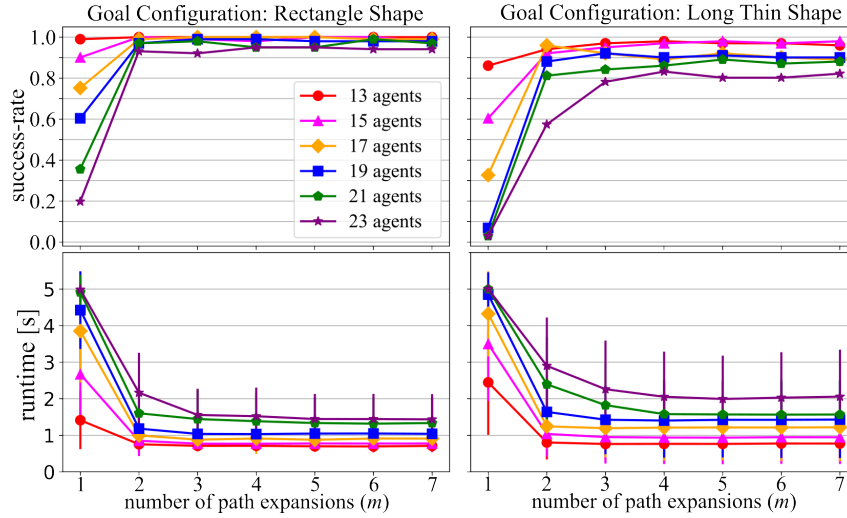


Fig. 5. Performance of APEDL as varying the number of path expansions in Maze environments (Env.5). One-time path expansion ( $m = 1$ ) fails as the number of agents increase. Adaptive path Expansion (APE) ( $m > 1$ ) can compute paths efficiently even with Goal Configuration in Long-Thin Shape.

single attempt considerably limits the success-rate. By contrast, increasing  $m$  improves performance without increasing planning cost and in fact decreasing plan cost since suitable paths are found more quickly. In all other experiments, we use a value of  $m = 5$ .

6.2.3 Results when varying environment difficulty. To demonstrate the robustness of our framework, we run APEDL on variations of the Rings environment type with three difficulty levels: *easy*, *medium*, and *hard*. The difficulty is controlled by the number of concentric rings, the disorderance between consecutive rings, and the number of breaks/gates on the rings. The features controlling difficulty are shown in Table 1. An easy and a hard environment are illustrated in Figure 6.b. The performances on three environment types with 23 agents and 5 s runtime are shown in Figure 6.b with success-rate decreasing from easy to hard level. With medium or hard levels, APEDL planner needs more than 5 s runtime to find valid paths for the agents.




Fig. 6. Performance of APEDL as varying the runtime (a) and environment difficulty (b) with 23 agents.

Table 1. Environment Difficulty Features On Rings Environments

Difficulty Level	Ring Count	Ring Distance	Break Count
Easy	4–5	8.0 m	6–7
Medium	5	7.0 m	5–6
Hard	6	5.5 m	4–5

**6.2.4 Results when varying runtime.** In this experiment, we ran the planners with various runtime from 1 s to 512 s. We selected the Maze Env., the most challenging Env. type for the experiments. The results are illustrated in Figure 6.a, showing that success-rate reaches 100% as the runtime grows.

**6.2.5 Results with Different Goal Configurations.** The goal configuration distribution impacts the planner’s performance, especially when managing diverse agent movement. A thin, long goal distribution increases runtime as the planner must search longer for a suitable planning order. We run APEDL with an alternate goal configuration distributions: long and thin versus rectangular (Figure 7) to test its robustness. The DL technique makes the APEDL robust in handling diverse movement directions across most environments (similar to Random Forest Env.-Figure 7.left). However, in Maze environments with narrow passages and only one way to go, frequent collisions and path modifications significantly increase difficulty of the planning problem, reducing planning performance somewhat.

### 6.3 Results under Line-of-Sight Communication Constraint

We test the performance of our APEDL planner under LOS constraint, which poses significant challenges due to the reason: maintaining the LOS constraint is difficult in obstacles-rich environments. In experiments using the same environments and instances as those with LCR constraint, we replace LCR with LOS constraint, with results presented in Figure 8.

Under the LOS constraint, the APEDL planner still outperform baseline methods, demonstrating the robustness of our proposed techniques: *adaptive path expansion* and *dynamic leading*. The APEDL planner achieves superior performance, successfully computing paths for 11–12 agents in Office and Rings environments and 4–6 agents in Random Forest, Waves, and Maze environments.

Maintaining the LOS communication in long, narrow passages or obstacle-rich environments is challenging, as a leader pursuing its goal (e.g., entering a room or navigating around obstacles) may inadvertently break communication with followers.




Fig. 7. Long-Thin Shape (Top) vs Rectangle Shape (Bottom) Goal Configuration.

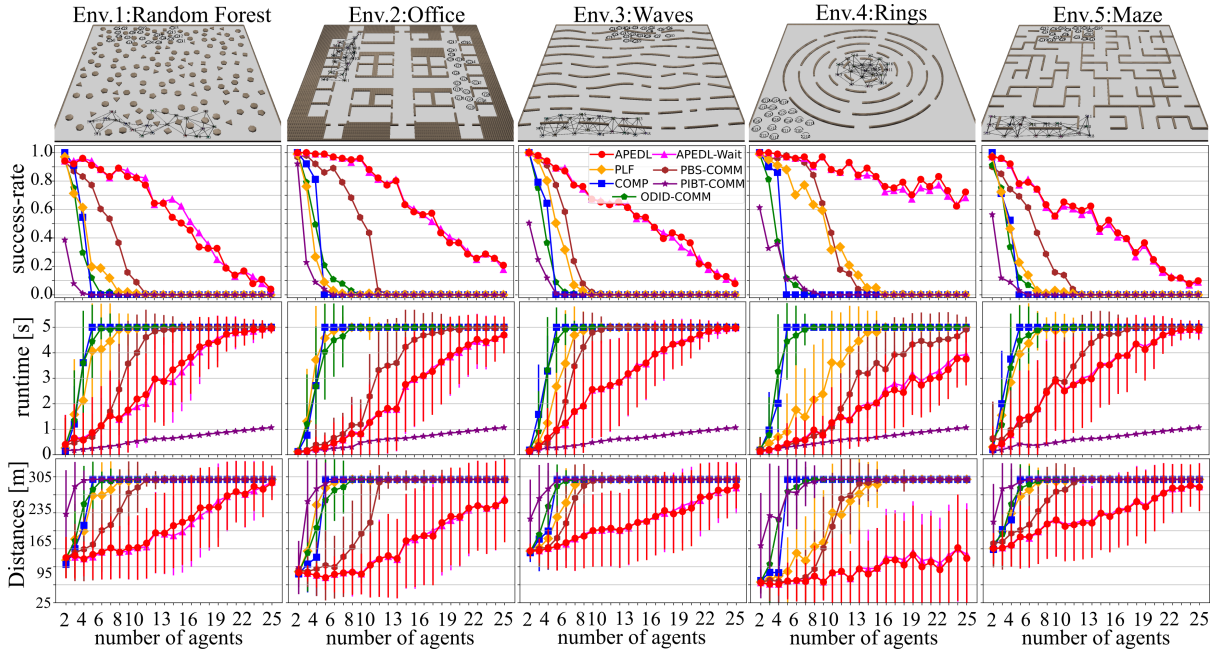


Fig. 8. **Success-rate, Runtime, and Travel Distances of our proposed planner: APEDL with PLF, COMP, ODID-COMM, PIBT-COMM, and PBS-COMM under LOS constraint.** The LOS constraint poses significant challenges for the planners to compute paths for agents. The results highlight the relative effectiveness of our proposed techniques: *adaptive path expansion* and *dynamic leading*.

## 7 Completeness Analysis

In our planning framework, the low-level single-agent planning (SAPF-DL) operates similarly to the  $A^*$  search algorithm, maintaining an *openList* and a *closedSet* to manage newly generated nodes and previously visited ones. The search systematically explores all reachable nodes until either the goal is reached or *openList* is empty (then returning the closest feasible path toward the goal). By design, the low-level single-agent planner is complete.

Though our high-level planner, APEDL, represents an advancement over leader-follower and PBS planners by eliminating their reliance on planning order that results in their incompleteness in solving the MAT3C problem, the APEDL planners is also *incomplete* in this domain. This is a consequence of the greedy nature of SAPF-DL, which always immediately returns upon finding a valid single agent path. We illustrate a scenario in Figure 9 that APEDL planner fails to solve: though the solution requires that agent  $a_3$  selects a longer path to allow other agents to pass by, the SAPF-DL planner returns only the shortest path and so the team becomes stuck.

Future work could explore extending our planning framework to achieve completeness. For this to be possible, the SAPF-DL single-agent planner would need to be extended to eventually generate all possible paths when expanding each agent, a modification that would require careful consideration and detailed study so that it would not dramatically slow planning performance. In this research, our focus is on developing a fast and practical planner that can efficiently handle the MAT3C problem, outperforming existing approaches. We plan to develop a complete version of the APEDL planner in future research.




Fig. 9. **A scenario demonstrating the incompleteness of planning framework.** Due to the greedy nature of SAPF-DL (find the shortest path), all agents attempt to go through  $G_1$  without trying  $G_2$  and cannot all reach their respective goals.

## 8 Conclusion

This paper presents a novel two-level planning framework, APEDL, to address the MAT3C problem. The core advancements of our approach are twofold: (1) at the high level, we deploy an *adaptive path expansion* technique to expand and refine agent paths multiple times to enable the planner to compute paths with arbitrary start and goal configurations, where state-of-the-art MAPF planners such as PBS, PIBT, OD-ID fail; and (2) at the low level, single-agent pathfinding with *dynamic leading* enables the dynamic reselection of the leading agent during single-agent path expansions whenever progress stalls, helping our planner to outperform state-of-the-art priority-based search approaches. We have tested our planner in multiple environments with features such as obstacles-richness, narrow passages, non-convex spaces, and long hallways. The results demonstrate the robustness of our planning framework APEDL capable of planning up to 25 agents under the LCR constraint, and up to 11–12 agents in Rings and Office environments and up to 3–10 agents in other environments under the LOS constraint within 5 seconds of runtime. Using the APEDL as a step-stone, in future work, we hope to extend these contributions to enable a version of the planner that has the completeness property. There are also potential performance improvements to be gained through the design of a new heuristic that improves the expansion of the team communication tree. Furthermore, we plan to extend the framework to a continuous action space to better support robot teams with rich kinodynamic constraints, our work a step towards communication-constrained planning more broadly. In this setting, SIPP is a promising choice for the low-level planner because it maintains a smaller state space by reasoning over time interval rather than time step, as in  $A^*$ . Incorporating continuous actions and kinodynamic constraints would bring APEDL a step closer to real-world applications.

## Acknowledgments

The work by E. Plaku is supported by (while serving at) the U.S. National Science Foundation. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors, and do not necessarily reflect the views of the National Science Foundation.

This material is based upon work supported by the National Science Foundation under Grant No. 2232733. G. J. Stein acknowledges support from this grant. G. J. Stein is also supported by Army Research Lab under award W911NF2520011.

## References

- C. Agachi, F. Arvin, and J. Hu. 2024. “RRT\*-Based Leader-Follower Trajectory Planning and Tracking in Multi-Agent Systems.” In: *2024 IEEE 12th International Conference on Intelligent Systems (IS)*, 1–6. doi:10.1109/IS61756.2024.10705259.
- A. M. Aljassani, S. N. Ghani, and A. M. Al-Hajjar. 2023. “Enhanced multi-agent systems formation and obstacle avoidance (EMAFOA) control algorithm.” *Results in Engineering*, 18, 101151.
- A. Andreychuk, K. Yakovlev, E. Boyarski, and R. Stern. 2021. “Improving continuous-time conflict based search.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 13. Vol. 35, 11220–11227.
- A. Andreychuk, K. Yakovlev, P. Surynek, D. Atzmon, and R. Stern. 2022. “Multi-agent pathfinding with continuous time.” *Artificial Intelligence*, 305, 103662.

- S. Bhattacharya, M. Likhachev, and V. Kumar. 2010. "Multi-agent path planning with multiple tasks and distance constraints." In: *2010 IEEE International Conference on Robotics and Automation*. IEEE, 953–959.
- H.-D. Bui, E. Plaku, and G. J. Stein. 2024. "Multi-Robot Guided Sampling-Based Motion Planning with Dynamics in Partially Mapped Environments." *IEEE Access*, 12.
- S. Choudhury, J. K. Gupta, P. Morales, and M. J. Kochenderfer. 2022. "Scalable Online planning for multi-agent MDPs." *Journal of Artificial Intelligence Research*, 73, 821–846.
- J.-M. Díaz-Báñez, L. E. Caraballo, M. A. Lopez, S. Bereg, I. Maza, and A. Ollero. 2017. "A general framework for synchronizing a team of robots under communication constraints." *IEEE Transactions on Robotics*, 33, 3, 748–755.
- L. Gao, D. Chu, Y. Cao, L. Lu, and C. Wu. 2019. "Multi-lane Convoy Control for Autonomous Vehicles based on Distributed Graph and Potential Field." In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2463–2469. doi:10.1109/ITSC.2019.8917409.
- S. Garrido, L. Moreno, J. V. Gómez, and P. U. Lima. 2013. "General path planning methodology for leader-follower robot formations." *International Journal of Advanced Robotic Systems*, 10, 64.
- Z. Huang, D. Chu, C. Wu, and Y. He. 2019. "Path Planning and Cooperative Control for Automated Vehicle Platoon Using Hybrid Automata." *IEEE Transactions on Intelligent Transportation Systems*, 20, 959–974. doi:10.1109/TITS.2018.2841967.
- K. H. Kowdiki, R. K. Barai, and S. Bhattacharya. 2019. "Autonomous leader-follower formation control of non-holonomic wheeled mobile robots by incremental path planning and sliding mode augmented tracking control." *International Journal of Systems, Control and Communications*, 10, 191–217.
- H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig. 2019. "Searching with consistent prioritization for multi-agent path finding." In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 33, 7643–7650.
- K. Okumura, F. Bonnet, Y. Tamura, and X. Défago. 2023. "Offline time-independent multiagent path planning." *IEEE Transactions on Robotics*, 39, 4, 2720–2737.
- K. Okumura, M. Machida, X. Défago, and Y. Tamura. 2022. "Priority inheritance with backtracking for iterative multi-agent path finding." *Artificial Intelligence*, 310, 103752.
- A. Pal, R. Tiwari, and A. Shukla. 2012. "Multi-robot exploration in wireless environments." *Cognitive Computation*, 4, 4, 526–542.
- X. Qian, A. de La Fortelle, and F. Moutarde. 2016. "A hierarchical Model Predictive Control framework for on-road formation control of autonomous vehicles." In: *2016 IEEE Intelligent Vehicles Symposium (IV)*, 376–381. doi:10.1109/IVS.2016.7535413.
- J. Sehn and J. Collier. 2024. "Hard Act to Follow: A Practical Leader-Follower System using Curvilinear Formation Motion Plans." In: *2024 IEEE 4th International Conference on Human-Machine Systems (ICHMS)*. IEEE, 1–6.
- K. Shojaei and M. R. Yousefi. 2019. "Tracking control of a convoy of autonomous robotic cars with a prescribed performance." *Transactions of the Institute of Measurement and Control*, 41, 3725–3741.
- D. Sigurdson, V. Bulitko, W. Yeoh, C. Hernández, and S. Koenig. 2018. "Multi-agent pathfinding with real-time heuristic search." In: *2018 IEEE conference on computational intelligence and games (CIG)*. IEEE, 1–8.
- K. Solovey, O. Salzman, and D. Halperin. 2016. "Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning." *The International Journal of Robotics Research*, 35, 501–513.
- T. Standley. 2010. "Finding optimal solutions to cooperative pathfinding problems." In: *Proceedings of the AAAI Conference on Artificial Intelligence 1*. Vol. 24, 173–178.
- G. Wagner and H. Choset. 2015. "Subdimensional expansion for multirobot path planning." *Artificial intelligence*, 219, 1–24.
- K. Wang and A. Botea. 2011. "MAPP: a scalable multi-agent path planning algorithm with tractability and completeness guarantees." *Journal of Artificial Intelligence Research*, 42, 55–90.
- X. Zhao, W. Yao, N. Li, and Y. Wang. 2017. "Design of leader's path following system for multi-vehicle autonomous convoy." In: *2017 IEEE international conference on unmanned systems (icus)*. IEEE, 132–138.

Received 07 July 2025; accepted 17 March 2026