

Large-Scale Multi-Robot Coverage Path Planning on Grids with Path Deconfliction

Jingtao Tang, Zining Mao, Hang Ma

Abstract—We study Multi-Robot Coverage Path Planning (MCPP) on a 4-neighbor 2D grid G , which aims to compute paths for multiple robots to cover all cells of G . Traditional approaches are limited as they first compute coverage trees on a quadrant coarsened grid \mathcal{H} and then employ the Spanning Tree Coverage (STC) paradigm to generate paths on G , making them inapplicable to grids with partially obstructed 2×2 blocks. To address this limitation, we reformulate the problem directly on G , revolutionizing grid-based MCPP solving and establishing new NP-hardness results. We introduce Extended-STC (ESTC), a novel paradigm that extends STC to ensure complete coverage with bounded suboptimality, even when \mathcal{H} includes partially obstructed blocks. Furthermore, we present LS-MCPP, a new algorithmic framework that integrates ESTC with three novel types of neighborhood operators within a local search strategy to optimize coverage paths directly on G . Unlike prior grid-based MCPP work, our approach also incorporates a versatile post-processing procedure that applies Multi-Agent Path Finding (MAPF) techniques to MCPP for the first time, enabling a fusion of these two important fields in multi-robot coordination. This procedure effectively resolves inter-robot conflicts and accommodates turning costs by solving a MAPF variant, making our MCPP solutions more practical for real-world applications. Extensive experiments demonstrate that our approach significantly improves solution quality and efficiency, managing up to 100 robots on grids as large as 256×256 within minutes of runtime. Validation with physical robots confirms the feasibility of our solutions under real-world conditions. A project page with code, demo videos, and additional resources is available at: <https://sites.google.com/view/lsmcpp>.

Index Terms—Multi-Robot System, Automated Planning, Heuristic Search.

I. INTRODUCTION

Coverage Path Planning (CPP) addresses the problem of determining a path that fully covers a designated workspace [1]. This problem is essential for a broad spectrum of robotic applications, from indoor tasks like vacuum cleaning [2] and inspection [3] to outdoor activities such as automated harvesting [4], planetary exploration [5], and environmental monitoring [6]. Multi-Robot Coverage Path Planning (MCPP) is an extension of CPP tailored for multi-robot systems, aiming to coordinate the paths of multiple robots to collectively cover the given workspace, thereby enhancing both task efficiency [7] and system robustness [8]. A fundamental challenge of MCPP is to balance the costs across multiple robots, often quantified by the *makespan*—the maximum path cost among all robots. This challenge intensifies in large-scale applications where the

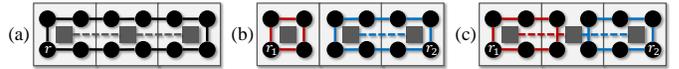


Figure 1. Grid-based CPP and MCPP: black circles, dark gray squares, solid lines, and dashed lines represent vertices of G , vertices of \mathcal{H} , coverage paths on G , and spanning edges of \mathcal{H} , respectively. Starting at r , the robot(s) need to collectively cover each vertex of G and return to r , with a cost of 1 to traverse between adjacent vertices of G . (a) Single-robot coverage path generated by STC. (b) Suboptimal 2-robot coverage paths with a makespan of 8, resulting from 2 coverage trees that minimize the maximum tree weights. (c) Optimal 2-robot coverage paths with a makespan of 6.

number of robots and the size of the workspace significantly increase.

This article tackles MCPP on a 4-neighbor 2D (edge-)weighted grid graph G . The problem is to compute paths for a given set of k robots, each starting from and returning to their given root vertices, to jointly cover all vertices of G , a problem analogous to the NP-hard Multiple Traveling Salesman Problem (m-TSP) [9] and Vehicle Routing Problem (VRP) [10]. Existing grid-based MCPP methods [11]–[15] operate on a coarsened grid \mathcal{H} , derived from G by grouping every 2×2 adjacent vertices, to compute coverage trees that jointly cover all vertices of \mathcal{H} . These methods then apply the Spanning Tree Coverage (STC) [16] paradigm to generate coverage paths on G by circumnavigating the computed coverage trees. However, operating exclusively on the coarsened grid \mathcal{H} has two significant disadvantages. First, it fails in cases where \mathcal{H} is *incomplete*—that is, when any 2×2 blocks contain obstructed grid cells absent from G .¹ Second, even optimal coverage trees on \mathcal{H} do not necessarily result in an optimal MCPP solution (as illustrated in Fig. (1b) and (1c)), as evidenced by an asymptotic suboptimality ratio of four for makespan minimization [13], since the paths derived from circumnavigating coverage trees of \mathcal{H} constitute only a subset of all possible sets of coverage paths on G . Furthermore, existing grid-based MCPP methods neglect to resolve inter-robot conflicts [17] in their computed solutions, leading to potential collisions during execution and

¹The grid graph G is technically a vertex-induced subgraph of an infinite grid. For clarity and consistency, we refer to the vertices in the infinite grid as “grid cells” throughout this paper. These grid cells can be obstructed, resulting in their absence as vertices in G . This distinction helps articulate the difference between the physical presence of cells in the infinite grid and their representation (or lack thereof) in the subgraph G . Such terminology is particularly useful in scenarios where contracting a 2×2 block of grid cells does not yield a complete set of four vertices, or a “hypervertex”, due to the obstruction or absence of one or more grid cells. This approach enables us to discuss the structure and topology of G more precisely, especially in the context of MCPP problems where the presence or absence of grid cells affects path planning and coverage strategies.

The authors are with the School of Computing Science, Simon Fraser University, Burnaby, BC V5A1S6, Canada. {jingtao_tang, zining_mao, hangma}@sfu.ca.

thus limiting their practical application in diverse real-world multi-robot scenarios.

A. Contributions

We revolutionize solving MCPP on grid graphs, overcoming the above limitations through a two-phase approach that first systematically searches for good coverage paths directly on G and subsequently resolves inter-robot conflicts in a post-processing procedure, which is illustrated in Fig. (2). Our algorithmic contribution are detailed as follows:

- 1) We propose a novel standalone paradigm called Extended-STC (ESTC), which extends STC to address coverage path-planning problems on any grid graph G , applicable both when the quadrant coarsening \mathcal{H} is complete and incomplete. Importantly, we demonstrate that ESTC guarantees complete coverage for both single- and multi-robot settings with bounded suboptimality, rendering it an efficient and versatile solution for coverage path planning.
- 2) We design three types of specialized neighborhood operators to facilitate an effective local search process by identifying cost-efficient subgraphs of G that are then used to generate coverage paths for the robots. The strategic integration of these operators significantly improves the efficiency of exploring the solution space. We then demonstrate how to combine these neighborhood operators with iterative calls to the ESTC paradigm to establish our LS-MCPP framework for solving MCPP.
- 3) We address inter-robot conflicts in the coverage paths computed by LS-MCPP or other MCPP methods by formulating the problem as a variant of Multi-Agent Path Finding [17]. This variant asks to compute conflict-free continuous-time trajectories for robots as each of them visits an ordered sequence of vertices specified by the MCPP solution. We develop new approaches for solving this variant based on existing MAPF techniques. These approaches integrate the novel Multi-Label Safe Interval Path Planning search that plans a continuous-time trajectory for an individual robot through its ordered sequence of vertices and, as a side benefit, allows our post-processing step to accommodate important kinematic constraints, such as turning costs—a critical factor in real-world coverage applications [18]–[20],

We showcase the benefits of our algorithmic pipeline through extensive quantitative experiments on large-scale instances involving up to 100 robots on a 256×256 grid. Key findings include:

- 1) Our ESTC paradigm results in coverage paths with significantly smaller makespans than existing STC-based paradigms and is effective on any grid graphs G , including those with incomplete quadrant coarsening \mathcal{H} , for both CPP and MCPP tasks.
- 2) Our LS-MCPP framework achieves significantly smaller makespans than state-of-the-art MCPP methods that rely on suboptimal tree coverage computations on \mathcal{H} and requires far less runtime to achieve makespans com-

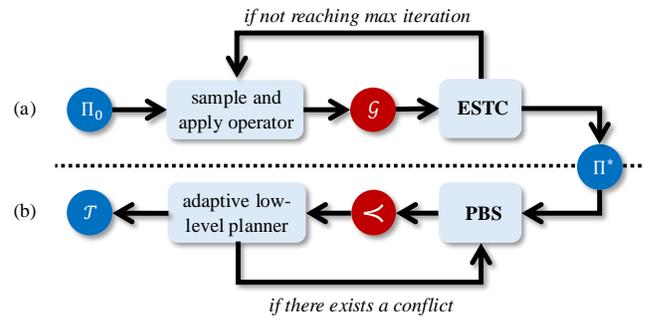


Figure 2. The proposed two-phase approach for MCPP in a grid graph with k robots. (a) The LS-MCPP phase: Starting with an initial solution Π_0 , the algorithm iteratively samples operators to update the subgraph set \mathcal{G} and generates k coverage paths using ESTC, ultimately producing a makespan-minimized solution Π^* . (b) The deconflicting phase: Given Π^* (or any other MCPP solution) as input, Priority-Based Search (PBS) explores a tree of priority ordering \prec while invoking an adaptive low-level planner to generate conflict-free paths, considering only high-priority robot trajectories for each robot. The process terminates with a set \mathcal{T} of conflict-free coverage paths.

parable to or better than those achieved by the optimal tree coverage computation.

- 3) Our post-processing procedure can be applied to any MCPP solutions, including not only those generated by LS-MCPP but also those generated by existing MCPP methods, to effectively resolve conflicts between robots and accounts for turning costs, further enhancing the practicability of the solutions.

Additionally, we validate the feasibility of our approach with physical robot deployments, confirming its practical applicability in real-world robotics applications.

This article significantly extends its preliminary conference version [21] by:

- 1) Providing a reformulation of MCPP directly on an (edge-)weighted grid graph G and a new NP-hardness result.
- 2) Improving the original ESTC paradigm by encoding cost differences between circumnavigating paths, with a new theoretical analysis of the properties of ESTC.
- 3) Incorporating new vertex-wise operators into the original LS-MCPP framework to facilitate a more effective exploration of the solution space.
- 4) Introducing the new post-processing procedure to resolve conflicts between robots.
- 5) Expanding the numerical results and including physical robot experiments to demonstrate the practical applicability of our approach.

B. Organization

We organize the remainder of this article as follows: We review related problems and methods in Sec. II. We present the problem formulations and complexity results in Sec. III. We introduce the ESTC paradigm in Sec. IV. We detail the specialized neighborhood operators and the LS-MCPP framework in Sec. V. We present our proposed post-processing procedure for resolving inter-robot conflicts in Sec. VI. We present our numerical results in Sec. VII and discuss our physical robot experiments in Sec. VII-D. We conclude our findings and outline future directions in Sec. VIII.

II. RELATED WORK

In this section, we survey related work on MCPP and MAPF.

A. Multi-Robot Coverage Path Planning (MCP)P

Following the existing taxonomy on CPP [22], we categorize MCP)P methods into grid-based, cellular decomposition, and global methods based on their approach to workspace decomposition.

Grid-Based Methods: Grid-based methods conceptualize the workspace as a grid graph G [23]–[25], allowing for the application of various graph algorithms. These methods traditionally operate exclusively on a quadrant coarsened \mathcal{H} derived from G and fail when \mathcal{H} is incomplete. They generalize the Spanning Tree Coverage (STC) paradigm for single-robot CPP in two ways. Single-tree methods [11], [14] compute a spanning tree on \mathcal{H} , generate a coverage path on G that circumnavigate this tree, and then segment this path to distribute among multiple robots. These methods do not offer guarantees on solution quality due to their simplistic approach to path distribution. Multi-tree methods [12], [13], [15] compute multiple coverage trees, each rooted at the root vertex of a specific robot, to jointly cover all vertices of \mathcal{H} , and then generate the coverage paths on G for the robots by circumnavigating their respective trees. In essence, multi-tree methods reduce MCP)P to the NP-hard Min-Max Rooted Tree Cover problem [26], [27], yielding an asymptotic suboptimality ratio of four in makespan minimization on grid graphs. It is important to note that, unlike our direct formulation of MCP)P on an edge-weighted G , existing grid-based methods typically formulate MCP)P on a vertex-weighted \mathcal{H} , making a restricted assumption that the weight of each vertex in \mathcal{H} is evenly distributed among its corresponding four vertices in G .

Cellular Decomposition Methods: Cellular decomposition methods [28]–[30] divide the workspace into sub-regions by detecting geometric critical points, using techniques such as trapezoidal decomposition [31], boustrophedon decomposition [32], and Morse decomposition [33]. After dividing the workspace, these methods generate a back-and-forth patterned path to cover each sub-region, often optimizing the path alignment orientation to enhance coverage efficiency [34]. While these methods are effective in structured environments, their performance diminishes in areas dense with obstacles or in non-rectilinear workspaces due to their reliance on geometric partitioning. Furthermore, they typically assume uniform traversal costs across the workspace, limiting their applicability in scenarios where robots have varying movement costs or task-specific coverage demands.

Global Methods: Global methods find significant application in additive manufacturing [35]–[37] (also known as layered fabrication), where a 3D object is printed layer by layer. Each layer poses a unique coverage problem, necessitating the planning of a space-filling curve to ensure the nozzle covers the entire 2D plane effectively. These methods focus on minimizing path curvature and discontinuities without explicitly decomposing the workspace. Recent advancements include the adaptation of the connected Fermat spiral [38] to reduce

energy consumption in CPP tasks over undulating terrains [39] and to minimize the makespan in general MCP)P scenarios [40]. The coverage quality of global methods largely depends on the characteristics of the space-filling curve generated, often necessitating additional trajectory optimization in a post-processing step.

B. Multi-Agent Path Finding (MAPF)

MAPF aims to compute conflict-free plans for multiple agents to move from their start vertices to their goal vertices. It is NP-hard to solve optimally for makespan minimization on grid graphs [41]. We highlight only relevant MAPF variants where each agent is given multiple goal vertices and direct readers interested in broader MAPF discussions to detailed surveys [17], [42]. OMG-MAPF [43] specifies a total order on the goal vertices given to each agent, while MAPF-PC [44] specifies precedence constraints—some goal vertex must be visited before another. MG-MAPF [45]–[47] does not specify a specific visiting order for the goal vertices given to each agent but instead aims to solve the goal sequencing problem jointly with MAPF. Multi-agent Pickup and Delivery (MAPD) [48]–[50] does not pre-assign goal vertices to agents and aims to solve the goal allocation and sequencing problem jointly with MAPF. All the above MAPF variants assume unit action costs and discrete time steps and do not directly apply to the problem of deconflicting MCP)P solutions on graphs with nonuniform edge weights.

III. MCP)P AND ITS NP-HARDNESS

We consider a set $I = \{1, 2, \dots, k\}$ of k robots operating on a four-neighbor connected undirected 2D grid graph $G = (V, E)$, where V is the set of vertices and E is the set of edges connecting each vertex to its top, bottom, left, and right neighbors if they exist. We assign each edge $e \in E$ with a non-negative weight w_e and simply set $w_e = 1$ for every $e \in E$ when considering an unweighted graph G . We specify a set $R = \{r_i\}_{i=1}^k \subseteq V$ of root vertices, where each robot i starts at vertex r_i . We use the tuple (G, I, R) to denote an MCP)P instance.

A path π_i for robot i is defined as a finite ordered sequence of vertices $(v_1, v_2, \dots, v_{|\pi_i|})$ such that $v_1 = v_{|\pi_i|} = r_i$ ² and $(v_{j-1}, v_j) \in E$ for each $j = 2, 3, \dots, |\pi_i|$. The cost of any path π is defined as $c(\pi) = \sum_{e \in \pi} w_e$. For clarity, we let $V(\pi)$ and $\mathcal{E}(\pi)$ denote the set of unique vertices in π and the multiset of edges in π that possibly contains repeated edges, respectively. The MCP)P problem is defined as follows:

Problem 1 (MCP)P. Find a set $\Pi = \{\pi_i\}_{i=1}^k$ of paths such that $\bigcup_{i=1}^k V(\pi_i) = V$.

The above problem formulation aligns with the *cover with return* coverage setting [13], where each robot must start from and return to its given root vertex. The quality of an MCP)P

²The requirement that each robot i must start and end at vertex r_i is also known as the *cover-and-return* setting later in [13], which naturally aligns with STC (and also our ESTC) algorithm in Sec. IV since its coverage path forms a “loop”. A post-processing procedure can be found in [12] to convert a STC-like path for the *cover-without-return* setting.

solution is evaluated using the *makespan* metric, denoted by $\theta(\Pi) = \max\{c(\pi_1), c(\pi_2), \dots, c(\pi_k)\}$.

Remark 1. *Following conventional formulations, our MCPP problem definition intentionally excludes inter-robot deconfliction (see Definition 1). However, practical deployment demands collision-free paths. Previous coupled approaches for multi-goal sequencing MAPF [46], [51], a variant of our problem, typically scale only to a few dozen goals (vertices to be covered) and often have low success rates. Our two-phase approach addresses these limitations by first computing an MCPP solution without enforcing collision-free paths (Sec.V) and then deconflicting the solution via MAPF techniques (Sec.VI). This decoupled approach not only scales efficiently but is also more generally applicable: Any MCPP method can be used in the first phase, and our post-processing procedure can subsequently handle inter-robot conflicts.*

MCPP can be viewed as an extension of min-max m-TSP on grid graphs where multiple depots exist and each city is allowed to be visited more than once. Given that the Hamiltonian cycle problem on grid graphs is NP-complete [52], both (single-robot) CPP and thus MCPP are NP-hard, which can be derived through a straightforward reduction from the Hamiltonian cycle problem.

Theorem 1. *It is NP-hard to find a coverage path with the smallest cost on an unweighted grid graph.*

Proof. The Hamiltonian cycle problem, known to be NP-complete on a grid graph $G = (V, E)$ [52], can be directly reduced to the decision version of CPP that asks to determine whether there is a coverage path on the same grid graph G with a cost not exceeding $|V|$. ■

We now examine MCPP for a more restricted class of grid graphs. Although the Hamiltonian cycle problem can be solved in polynomial time [53] on solid grid graphs, which are grid graphs without holes, the scenario changes when considering MCPP. Specifically, inspired by the NP-hardness proof for a specialized MCPP variant defined on a vertex-weighted grid graph [13], we extend this complexity result to our more broadly defined MCPP on solid grid graphs through the following theorem.

Theorem 2. *It is NP-hard to find an MCPP solution with the smallest makespan, even on a solid unweighted grid graph.*

Proof. We prove by showing the NP-completeness of a decision version of MCPP, that is, determining whether a given MCPP instance has a solution with a makespan not exceeding a given positive integer. One can evaluate whether the makespan of a given solution exceeds the given positive integer in polynomial time, implying that the MCPP problem is in NP. We now reduce the MCPP problem to the NP-complete 3-Partition problem [54] in polynomial time. Specifically, given a positive integer B and a set $S = \{a_i\}_{i=1}^{3k}$ of positive integers, where each $a_i \in (\frac{B}{4}, \frac{B}{2})$ and $\sum_{i=1}^{3k} a_i = kB$, the problem of 3-Partition is to determine if there exists an evenly partition S into k subsets S_1, S_2, \dots, S_k , each summing to B .

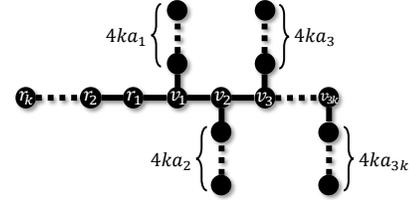


Figure 3. An MCPP instance reducible from the 3-Partition problem.

Given a 3-Partition instance, we construct an MCPP instance on an unweighted grid graph as shown in Fig. (3). This graph has a central horizontal “corridor” with $4k$ vertices $r_k, r_{k-1}, \dots, r_1, v_1, v_2, \dots, v_{3k}$ and $3k$ vertical “tunnel”, each connected to a vertex v_i . The i -th tunnel contains $4ka_i$ vertices and extends v_i upwards if i is odd and downwards if even. Robots $1, 2, \dots, k$ start at their root vertices r_1, r_2, \dots, r_k , respectively.

We now show that the makespan of the above MCPP instance is at most $8kB + 8k - 2$ if and only if S can be evenly partitioned into k subsets:

- “If” direction: Given an even partition of S into k subsets S_j where $\sum_{i \in S_j} a_i = B$, we assign each robot j to cover the i -th tunnel for each $i \in S_j$. Each robot needs at most $\sum_{i \in S_j} 4ka_i \times 2 = 8kB$ steps to cover its tunnels and at most $(4k - 1) \times 2$ steps to traverse the corridor (between r_k to v_{3k}). Therefore, the cost of any robot is at most $8kB + 8k - 2$, ensuring the makespan of the resulting MCPP solution does not exceed $8kB + 8k - 2$.
- “Only if” direction: Given an MCPP solution where the cost of any robot is at most $8kB + 8k - 2$, we define S_j as the set of tunnel indices where the innermost vertex of each tunnel is first covered by robot j , resulting in a partition of k subsets S_1, S_2, \dots, S_k . Since each $a_i \in (\frac{B}{4}, \frac{B}{2})$ and $|S| = 3k$, each robot must cover exactly 3 tunnels. Therefore, the cost for any robot is at least $\sum_{i \in S_j} (4ka_i \times 2) + 3 \times 2 = 8k \sum_{i \in S_j} a_i + 6$, calculated as the cost of robot 1 covering the first three tunnels. On the other hand, the cost of any robot is at most $8kB + 8k - 2$ in the given MCPP solution. Since a_i, B and k are all positive integers, chaining the above two inequalities yields $\sum_{i \in S_j} a_i \leq B + 1 - \frac{1}{k} \leq B$ for any j . Since $\sum_{j=1}^k \sum_{i \in S_j} a_i = \sum_{i=1}^{3k} a_i = kB$, $\sum_{i \in S_j} a_i = B$ must hold to satisfy $\sum_{i \in S_j} a_i \leq B$, for any j . Therefore, S_1, S_2, \dots, S_k is an even partition of S . ■

IV. EXTENDED SPANNING TREE COVERAGE (ESTC)

In this section, we present our new ESTC paradigm that advances standard STC-based paradigms by addressing coverage problems on any given connected grid graph, even when the quadrant coarsening \mathcal{H} is incomplete. We first describe how ESTC solves (single-robot) CPP, a special case of MCPP (Problem 1) when $k = |I| = 1$. Solving CPP efficiently is fundamental to our local search framework for MCPP detailed in Sec. V. We then analyze the properties and time complexity of ESTC and introduce two local optimizations to improve ESTC solutions. We finally discuss how ESTC can be

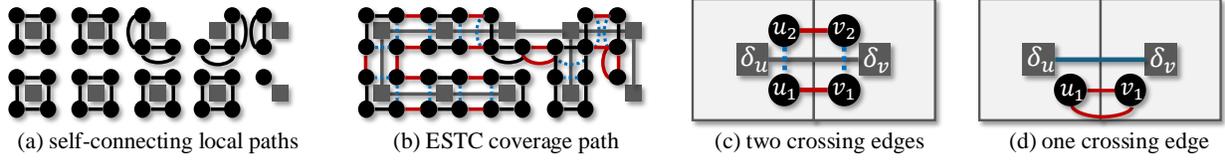


Figure 4. Illustration of the ESTC algorithm. (a) The initial local paths (black lines). (b) The resulting ESTC coverage path by applying the rerouting rule on each hyperedge of a hypergraph spanning tree. (c)(d) For each hyperedge (gray solid line) $\varepsilon = (\delta_u, \delta_v) \in E_h$, the rerouting rule adds E_ε^+ (red solid lines) to and removes E_ε^- (blue dashed lines) from the local paths of δ_u and δ_v , depending on the number of the crossing edges (u_1, v_1) and (u_2, v_2) of ε .

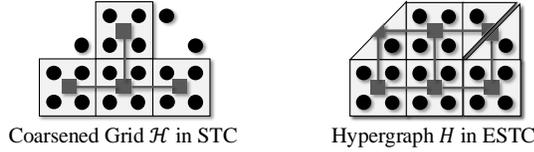


Figure 5. The coarsened grid \mathcal{H} in STC and the hypergraph H in ESTC (both in gray markers and lines) for an input grid graph (black circles). The former yields incomplete coverage while the latter ensures complete coverage.

integrated into existing MCPP methods to effectively handle incomplete \mathcal{H} .

ESTC extends both the STC paradigm for offline CPP and the Full-STC paradigm [55] for online CPP, which operate on a quadrant coarsened grid \mathcal{H} derived from G . These paradigms typically evaluate the traversal costs by decomposing the given hypervertex weights of \mathcal{H} into vertex weights for G to maintain solution optimality. However, this approach does not adequately capture the actual traversal costs incurred by a robot, which are better represented by the edges traversed, as defined in our formulation.

A. Algorithm Description

ESTC operates on a *hypergraph* graph $H = (V_h, E_h)$, derived from the grid graph $G = (V, E)$. As shown in Fig. (5), similar to standard STC-based paradigms, ESTC attempts to contract neighboring vertices of V corresponding to a 2×2 block of grid cells into a *hypervertex* of V_h . However, unlike the quadrant coarsened grid \mathcal{H} used by standard STC-based paradigms, the hypergraph H effectively addresses the disconnections caused by partially obstructed 2×2 blocks. A *hyperedge* connects two hypervertices only if there is at least one adjacent vertex pair from these hypervertices. A special case in hypervertex construction occurs when a 2×2 block contains only two diagonal vertices. In this case, each diagonal vertex forms its own hypervertex, resulting in two nonadjacent hypervertices (see Fig. (5)). A hypervertex δ is *incomplete* if it is formed by fewer than four vertices; otherwise, it is *complete*. The graph G is contracted left to right and bottom to top by default during the hypergraph construction. For clarity, we let $\delta_v \in V_h$ denote the hypervertex contracted from $v \in V$ and $\varepsilon = (\delta_u, \delta_v) \in E_h$ denote the hyperedge connecting δ_u, δ_v .

We describe the pseudocode for ESTC in Alg. 1. ESTC begins with the hypergraph construction [Line 1] and finds the Minimum Spanning Tree (MST) of H . Specifically, ESTC employs Kruskal's algorithm [56] to find the MST [Line 2] since the hyperedge weights (will be defined later) can be

Algorithm 1: Extended-STC (ESTC)

Input: grid graph $G = (V, E)$

- 1 construct hypergraph $H = (V_h, E_h)$ from G
- 2 $M \leftarrow$ find MST of H via Kruskal's algorithm
- 3 $\pi_\delta \leftarrow$ initial self-connecting local path for each $\delta \in V_h$
- 4 $\pi \leftarrow$ an empty path
- 5 **for** $\varepsilon = (\delta_u, \delta_v) \in$ the set of hyperedges in M **do**
- 6 \lfloor connect $\pi_{\delta_u}, \pi_{\delta_v}$ via rerouting rule and add it to π
- 7 **return** ESTC coverage path π

negative. It then initializes a self-connecting local path π_δ for each $\delta_v \in V_h$ [Line 3], as shown in Fig. (4a). Note that π_δ is the optimal coverage path for the vertices with respect to hypervertex δ . Then, for each hyperedge $\varepsilon = (\delta_u, \delta_v)$ in the MST M , a rerouting rule is applied to connect π_{δ_u} and π_{δ_v} to construct a coverage path for the vertices in δ_u and δ_v [Lines 5-6]. The resulting ESTC coverage path π is returned after all the local paths are connected one to another [Line 7], as illustrated in Fig. (4b).

Rerouting Rule: The rerouting rule aims to connect pairwise self-connecting local paths of a hyperedge. We categorize each hyperedge $\varepsilon = (\delta_u, \delta_v) \in E_h$ based on the number of *crossing edges* between the corresponding vertices within the blocks of δ_u, δ_v . A crossing edge is an edge $(u, v) \in E$ where u and v are contracted into adjacent hypervertices δ_u and δ_v , respectively. Denote x_ε as the number of crossing edges for a hyperedge ε , we define two multisets for any hyperedge ε , $E_\varepsilon^+ = \{(u_1, v_1), (u_2, v_2)\}$ and $E_\varepsilon^- = \{(u_1, u_2), (v_1, v_2)\}$ if $x_\varepsilon = 2$; otherwise when $x_\varepsilon = 1$, $E_\varepsilon^+ = \{(u_1, v_1), (v_1, u_1)\}$ and $E_\varepsilon^- = \emptyset$. Note that $x_\varepsilon = 0$ implies $(\delta_u, \delta_v) \notin E_h$. As illustrated in Fig. (4c) and (4d), E_ε^+ and E_ε^- are represented in red solid lines and blue dashed lines, respectively. The rerouting rule for every ε is then defined to add E_ε^+ to connect the two local paths while removing E_ε^- . It ensures that each vertex has an even degree with the least vertex duplication, making the connected path a valid path. The above rerouting rule can be interpreted as a combination of the standard STC circumnavigating pattern when $x_\varepsilon = 2$ and the Full-STC rerouting rule when $x_\varepsilon = 1$. It is always feasible as the construction of H guarantees no vertex within any incomplete hypervertex is isolated, thereby ensuring the path remains feasible and continuous across any grid graphs.

Hyperedge Weights: The ESTC algorithm has a good property: Given any two spanning trees T and T' with their resulting ESTC paths π and π' , respectively, their path difference (represented using the multisets of their edges) is a constant,

given as follows:

$$\begin{aligned} \mathcal{E}(\pi) \setminus \mathcal{E}(\pi') &= \bigcup_{\varepsilon \in \Delta E_h} E_\varepsilon^+ \cup \bigcup_{\varepsilon' \in \Delta E'_h} E_{\varepsilon'}^- \quad \text{and} \\ \mathcal{E}(\pi') \setminus \mathcal{E}(\pi) &= \bigcup_{\varepsilon' \in \Delta E'_h} E_{\varepsilon'}^+ \cup \bigcup_{\varepsilon \in \Delta E_h} E_\varepsilon^-, \end{aligned} \quad (1)$$

where $\Delta E_h = E_h(T) \setminus E_h(T')$ and $\Delta E'_h = E_h(T') \setminus E_h(T)$ represent the two hyperedge difference sets between T and T' . We then define the weight w_ε for hyperedge ε as below:

$$w_\varepsilon = \sum_{e \in E_\varepsilon^+} w_e - \sum_{e \in E_\varepsilon^-} w_e. \quad (2)$$

The definition of w_ε implicitly encodes the ESTC path cost difference of whether ε exists in the spanning tree. In Theorem 5, we validate that with the defined hyperedge weights in Eqn. (2), an MST of hypergraph H always results in an optimal ESTC circumnavigating coverage path on G .

B. Theoretical Analysis

ESTC is guaranteed to generate a path that covers all vertices of a connected grid graph G based on the argument used in the proof of Lemma 3.1 in [55] since the argument is not affected by the addition of edge weights.

Theorem 3. *ESTC achieves complete coverage on a connected grid graph G .*

We now analyze the time complexity of ESTC below.

Theorem 4. *The time complexity of ESTC (Alg. 1) on a grid graph $G = (V, E)$ is $O(|V| \log |V|)$.*

Proof. To construct the hypergraph $H = (V_h, E_h)$ from $G = (V, E)$ as described in Sec. IV-A, ESTC first iterates through V to build V_h and then through V_h to build E_h . Given that $|V_h|$ is proportional to $|V|$, the time complexity of this construction is $O(|V|)$. To find the MST of H , ESTC uses Kruskal's algorithm, which incurs a time complexity of $O(|E| \log |E|)$. Initializing the local paths and connecting them via the rerouting rule entail a time complexity of $O(|V| + |E|)$. As $O(|E|) = O(|V|)$ for grid graphs, the overall time complexity of ESTC combines to $O(|V| \log |V|)$. ■

The following theorem validates that our definition of hyperedge weights ensures that the ESTC circumnavigating path derived from an MST of H has the smallest cost among those derived from any spanning trees of H .

Theorem 5. *The cost of the circumnavigating path derived from a minimum spanning tree (MST) of the hypergraph H is no greater than that from any other spanning tree of H .*

Proof. For hypergraph $H = (V_h, E_h)$, consider its two spanning trees T and T' and their corresponding ESTC circumnav-

igating paths π and π' , respectively. With Eqn. (1), the path cost difference Δc between π and π' can be computed as:

$$\begin{aligned} \Delta c &= c(\pi) - c(\pi') = \sum_{e \in \mathcal{E}(\pi) \setminus \mathcal{E}(\pi')} w_e - \sum_{e \in \mathcal{E}(\pi') \setminus \mathcal{E}(\pi)} w_e \\ &= \sum_{\varepsilon \in \Delta E_h} \sum_{e \in E_\varepsilon^+} w_e + \sum_{\varepsilon' \in \Delta E'_h} \sum_{e \in E_{\varepsilon'}^-} w_e \\ &\quad - \sum_{\varepsilon' \in \Delta E'_h} \sum_{e \in E_{\varepsilon'}^+} w_e - \sum_{\varepsilon \in \Delta E_h} \sum_{e \in E_\varepsilon^-} w_e. \end{aligned}$$

With Eqn. (2) substituted into the above, we have:

$$\begin{aligned} \Delta c &= \sum_{\varepsilon \in \Delta E_h} \left[\sum_{e \in E_\varepsilon^+} w_e - \sum_{e \in E_\varepsilon^-} w_e \right] \\ &\quad - \sum_{\varepsilon' \in \Delta E'_h} \left[\sum_{e \in E_{\varepsilon'}^+} w_e - \sum_{e \in E_{\varepsilon'}^-} w_e \right] \quad (3) \\ &= \sum_{\varepsilon \in \Delta E_h} w_\varepsilon - \sum_{\varepsilon' \in \Delta E'_h} w_{\varepsilon'}. \end{aligned}$$

If T is an MST and T' is an arbitrary spanning tree, we know that $\forall \varepsilon \in \Delta E_h \subseteq E_h(T)$ and $\varepsilon' \in \Delta E'_h \subseteq E_h(T')$, we have $w_\varepsilon \leq w_{\varepsilon'}$, where $E_h(T)$ and $E_h(T')$ are the edge sets of T and T' , respectively. Substituting it into Eqn. (3), we have $\Delta c \leq 0$ for an MST T and an arbitrary spanning tree T' . ■

The following theorem shows that the optimal circumnavigating path on H yields an optimal CPP solution when all hypervertices are complete and a bounded suboptimal CPP solution in general.

Theorem 6. *ESTC returns a coverage path with a cost at most $2 \cdot \frac{w_{\max}}{w_{\min}} \cdot (1 + \frac{n_c - 1}{|V|})$ times the cost of the optimal single-robot coverage path, where n_c is the number of complete hypervertices and w_{\max} and w_{\min} are the maximum and minimum hyperedge weights, respectively.*

Proof. Consider a CPP instance (G, I, R) , where G is contracted into hypergraph $H = (V_h, E_h)$ with n_c complete hypervertices. For any optimal valid coverage path π^* , it follows that $V(\pi^*) = V$ and both the first and last vertices of π^* are the root vertex, implying $|\mathcal{E}(\pi^*)| \geq |V|$, with equality only if there are no incomplete hypervertices in H . The suboptimality ratio ρ for any ESTC path π resulted from MST T is given by:

$$\rho = \frac{c(\pi)}{c(\pi^*)} = \frac{\sum_{e \in \mathcal{E}(\pi)} w_e}{\sum_{e \in \mathcal{E}(\pi^*)} w_e} \leq \frac{|\mathcal{E}(\pi)| \cdot w_{\max}}{|\mathcal{E}(\pi^*)| \cdot w_{\min}} \leq \frac{|\mathcal{E}(\pi)| \cdot w_{\max}}{|V| \cdot w_{\min}}.$$

We now categorize V_h into four disjoint sets such that $V_h = \bigcup_{i=1}^4 V_h^{(i)}$ and each $V_h^{(i)}$ contains hypervertices with i vertices located in it. As illustrated in Fig. (4a), for any $\delta \in V_h^{(i)}$, the initial ESTC self-connecting local path $|\pi_\delta| = 4$ if $i = 4$ or 3, $|\pi_\delta| = 2$ if $i = 2$, and $|\pi_\delta| = 0$ if $i = 1$. We categorize the hyperedge set $E_h(T)$ of MST T into $E_h^{(2)}(T)$ and $E_h^{(1)}(T)$ according to the number of crossing edges as

described in Fig. (4c) and (4d). By the initialization of the self-connecting local paths and the rerouting rule, the cardinality of the multiset $\mathcal{E}(\pi)$ can be computed as:

$$|\mathcal{E}(\pi)| = 4|V_h^{(4)}| + 4|V_h^{(3)}| + 2|V_h^{(2)}| + 2|E_h^{(1)}(T)|. \quad (4)$$

As $|V| = \sum_{i=1}^4 i \cdot |V_h^{(i)}|$, the bound can be further refined as:

$$\rho \leq \frac{w_{\max}}{w_{\min}} \cdot \left[1 + \frac{|V_h^{(3)}| - |V_h^{(1)}| + 2|E_h^{(1)}(T)|}{|V|} \right]. \quad (5)$$

Substituting $|E_h^{(1)}(T)| \leq |E_h(T)| = |V_h| - 1 = \sum_{i=1}^4 |V_h^{(i)}| - 1$ and $n_c = |V_h^{(4)}|$ into Eqn. (5), we obtain a simpler yet looser bound:

$$\rho \leq 2 \cdot \frac{w_{\max}}{w_{\min}} \cdot \left(1 + \frac{n_c - 1}{|V|} \right). \quad (6)$$

Remark 2. With the tighter bound in Eqn. (5), we can see that ESTC is optimal ($\rho = 1$) when G is unweighted ($\frac{w_{\max}}{w_{\min}} = 1$) and all hypervertices are complete ($|V_h^{(3)}| = |V_h^{(1)}| = |E_h^{(1)}(T)| = 0$). The bound in Eqn. (6) may appear relatively loose. However, the last multiplier $(1 + \frac{n_c - 1}{|V|}) \in [1, 2)$ as $n_c \leq |V|$. In fact, it can be seen that as n_c approaches 0, the bound in Eqn. (6) gets tighter and closer to $2 \cdot \frac{w_{\max}}{w_{\min}}$. This is because when H has more incomplete hypervertices, the optimal solution π^* necessitates a rerouting rule, such as that in Full-STC, to revisit edges through incomplete hypervertices to achieve complete coverage.

C. Local Optimizations

We introduce two efficient local optimizations for ESTC, namely *Parallel Rewiring* and *Turn Reduction*. Parallel Rewiring operates as a post-processing procedure to refine the coverage path produced by ESTC, by strategically adjusting path segments locally to reduce the path cost. Turn Reduction serves as a secondary objective during the execution of ESTC (see Alg. 1) by reducing the number of turns in the path, which is particularly beneficial in environments where turning is costly or undesirable. While these local optimizations do not theoretically improve the upper bound on the suboptimality of ESTC in Theorem 6, empirical observations (see Sec. IV-A) often demonstrate a notable improvement in solution quality.

Parallel Rewiring: As the original STC only explores the solution space of circumnavigating coverage paths, we develop two parallel rewiring procedures to refine any ESTC solution π using non-circumnavigating local path segments. Type-A Parallel Rewiring searches for every subsequence in π matching the pattern $(\dots, v, u, \dots, v, t, s, \dots, t, \dots)$ and collapses it into $(\dots, v, u, \dots, s, \dots, t, \dots)$ if the resulting cost difference $w_{(u,s)} - [w_{(u,v)} + w_{(v,t)} + w_{(t,s)}] < 0$. As shown in Fig. (6a), each collapsing operation effectively removes the intermediate vertices $v, t \in V$ without losing any vertex coverage. Type-B parallel rewiring identifies every pair of two vertically (or horizontally) parallel edges $(u_1, u_2), (v_1, v_2) \in \mathcal{E}(\pi)$ where $(u_1, u_3), (u_2, u_4), (s_1, u_3), (s_2, u_4) \in \mathcal{E}(\pi)$ are aligned horizontally (or vertically). As shown in Fig. (6b),

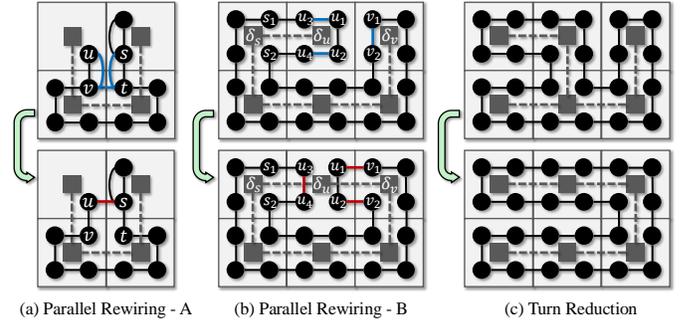


Figure 6. ESTC coverage path before (top row) and after local optimizations (bottom row). (a)(b) The two types of parallel rewiring. (c) Turn reduction with horizontally aligned hyperedges.

it then replaces the edges $(u_1, u_3), (u_2, u_4), (v_1, v_2)$ with $(u_3, u_4), (u_1, v_1), (u_2, v_2)$ if the resulting cost difference $[w_{(u_3, u_4)} + w_{(u_1, v_1)} + w_{(u_2, v_2)}] - [w_{(u_1, u_3)} + w_{(u_2, u_4)} + w_{(v_1, v_2)}] < 0$. Both types have a time complexity of $O(|V|^2)$ with a naive implementation involving two-layer nested iterations through V .

Turn Reduction: ESTC can be seamlessly integrated with a secondary objective of turn reduction in addition to its primary objective of path cost minimization. Specifically, when ESTC constructing the MST, it prioritizes the alignment of hyperedges. As the hypergraph $H = (V_h, E_h)$ is a four-neighbor 2d grid, its hyperedges can be categorized by their orientation as vertical or horizontal. To align with a given target orientation, the ESTC prioritizes the hyperedges traversal in Kruskal's algorithm (line 2 of Alg. 1) to favor a hyperedge where (1) its orientation matches the given target orientation and (2) sum of the degrees of its two endpoints is the smallest. The first prioritization ensures orientation consistency as much as possible, whereas the second prioritization aims to extend the MST into less connected regions of H for better alignment. Turn Reduction requires adding only two priority keys when constructing the MST and thus maintains the same time complexity for ESTC as demonstrated in Theorem 4. Fig. (6c) shows an example of applying Turn Reduction to an ESTC coverage path.

D. Integration into Existing MCPP Methods

Existing grid-based MCPP methods often focus on special cases of MCPP that assume complete hypervertices and vertex-weighted graphs. However, our more generally formulated MCPP, as defined in Problem 1, addresses a broader scope that remains relatively unexplored in the literature, despite its significant practical relevance. The proposed ESTC paradigm can be seamlessly integrated into these existing methods, allowing them to effectively solve MCPP even without the above constraints typically assumed. To demonstrate how existing MCPP methods can be adapted using the ESTC paradigm to solve an MCPP instance on any grid graph G , even when some hypervertices are incomplete, we use the following examples categorized into multi-tree and single-tree methods.

1) **Multi-Tree Methods:** For MCPP instance (G, I, R) , we first generate a set $\{G_i = (V_i, E_i)\}_{i=1}^k$ of k connected

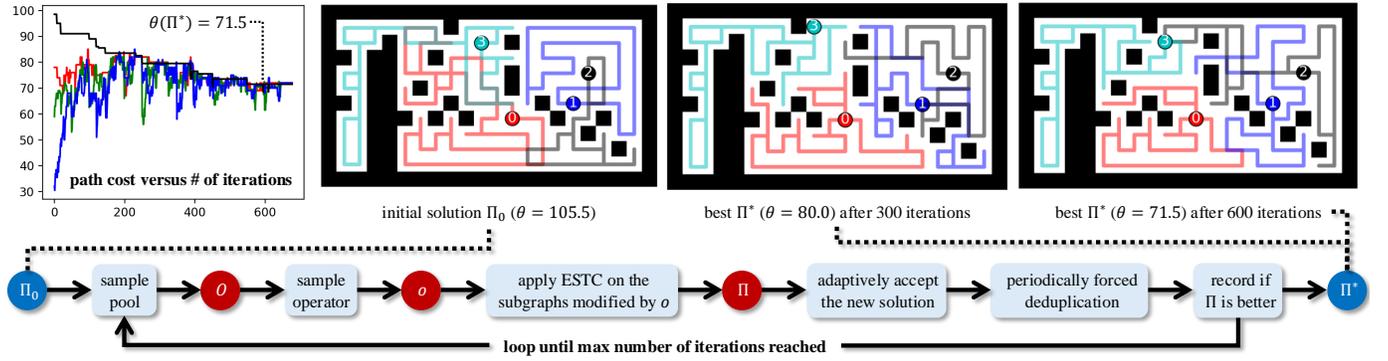


Figure 7. The flowchart of our LS-MCPP algorithmic framework for instance *floor-small*. Given an initial MCPP solution Π_0 , LS-MCPP iteratively samples a pool O from three pools of dedicated operators and then samples an operator o from O , to transform a set \mathcal{G} of subgraphs from one to another. It then applies ESTC on the subgraphs to evaluate the new MCPP solution Π , which is adaptively accepted according to its makespan θ . A periodic subroutine of forced deduplication is called at the end of each iteration, and the best solution Π^* is recorded if Π is better with a lower makespan.

subgraphs of G . Each G_i contains the root vertex r_i . We then solve each sub-CPP instance for robot i on G_i as described in Alg. 1 to generate its coverage path. The following theorem shows a sufficient condition for complete coverage.

Theorem 7. *If each G_i is connected and $\bigcup_{i=1}^k V_i = V$, then using ESTC is guaranteed to achieve complete coverage of G .*

Proof. Given that G_i is connected, ESTC is guaranteed to generate a path that covers all its vertices due to Theorem 3. If $\bigcup_{i=1}^k V_i = V$, then the coverage paths generated for all G_i jointly cover all vertices of G . ■

The design of our LS-MCPP framework is based on the principle expounded in the above theorem. It aims to search for a good set of subgraphs G_i for the robots, each being connected and containing the root vertex r_i , with the property that the union of all their vertices comprises the entire vertex set of G for complete coverage. It now suffices to outline the details below for adapting existing methods to compute the k subgraphs, with comprehensive numerical results in Sec. VII. **VOR:** VOR is a baseline method introduced in [21] that is inherently compatible with ESTC. It partitions G into k subgraphs $\{G_i\}_{i=1}^k$ as a Voronoi diagram [57], using the set R of root vertices as the pivot points and shortest path distances on G as the distance metric.

MFC: Multi-Robot Forest Coverage (MFC) [12], [13] employs the Rooted-Tree-Cover procedure [26] to heuristically compute a *rooted tree cover* on the quadrant coarsened grid \mathcal{H} —a forest of k trees $\{T_i\}_{i=1}^k$, each rooted at the hypervertex containing the root vertex r_i , which jointly cover all hypervertices of \mathcal{H} . MFC then generates a coverage path for each robot i by applying STC on its tree T_i . To integrate MFC with ESTC, we apply the Rooted-Tree-Cover procedure on the hypergraph H (constructed as per Sec. IV-A) instead of \mathcal{H} to generate a rooted tree cover on H that induces the set of k subgraphs $\{G_i\}_{i=1}^k$ of G .

MIP: MIP [15] solves a mixed-integer program to compute an optimal rooted tree cover on the quadrant coarsened grid \mathcal{H} that minimizes the maximum tree weight. To integrate MIP with ESTC, we adapt the mixed-integer program to compute an optimal rooted tree cover on the hypergraph H (constructed

as per Sec. IV-A) instead of \mathcal{H} , which induces the set of k subgraphs $\{G_i\}_{i=1}^k$ of G .

2) **Single-Tree Methods:** Single-tree methods, unlike multi-tree methods, apply the STC paradigm to only a single tree: STC first computes an MST on the quadrant coarsened grid \mathcal{H} and subsequently generates a coverage path that circumnavigates the MST. The path is then partitioned into k segments, one for each robot. MSTC [11] partitions the path such that each segment starts at a unique root vertex. Its improvement MSTC* [14] greedily identifies balanced partitioning points along the path to potentially reduce the makespan and then connects each root vertex to its respective segment via shortest paths. We can adapt MSTC* to solve MCPP for scenarios with incomplete hypervertices by simply replacing STC with ESTC.

V. THE LS-MCPP FRAMEWORK

In this section, we introduce LS-MCPP, our local search algorithmic framework designed to effectively address MCPP as formulated in Problem 1. This framework integrates three types of neighborhood operators and utilizes ESTC as a subroutine, enhancing its ability to solve MCPP even in complex environments with incomplete hypervertices. LS-MCPP aims to optimize coverage paths by iteratively refining subgraphs and their corresponding coverage paths, ensuring both efficiency and completeness in coverage solutions.

A. Overview

As shown in Fig. (7), LS-MCPP begins with an initial solution for the MCPP instance (G, I, R) with its corresponding set $\mathcal{G}_0 = \{G_i = (V_i, E_i)\}_{i=1}^k$ of k connected subgraphs induced by the solution paths. LS-MCPP operates as an iterative procedure that employs a two-layer hierarchical sampling scheme, progressively adjusting the subgraph set \mathcal{G} to efficiently explore the constructed neighborhood: The first layer selects an operator pool using *roulette wheel* selection [58] from three operator pools, each containing a different type of operators; The second layer heuristically selects an operator from the chosen pool, as detailed in Sec. V-D.

Algorithm 2: LS-MCPP

Input: MCPP instance (G, I, R) , initial solution Π_0
Param: max iteration $M \in \mathbb{Z}^+$, forced deduplication step $S \in \mathbb{Z}^+$, temperature decay factor $\alpha \in [0, 1]$, pool weight decay factor $\gamma \in [0, 1]$

- 1 $\mathcal{G} = \{G_i\}_{i=1}^k \leftarrow$ set of subgraphs extracted from Π_0
- 2 $\Pi \leftarrow \Pi_0, \Pi^* \leftarrow \Pi_0 \triangleright$ current and optimal solutions
- 3 $\mathcal{O} \leftarrow \{\emptyset, \emptyset, \emptyset\} \triangleright$ set of pools of operators
- 4 *Update-Pools*($\mathcal{G}, \mathcal{O}, \Pi$) \triangleright initialize the pools
- 5 $t \leftarrow 1 \triangleright$ temperature in *simulated annealing*
- 6 $\mathbf{p} \leftarrow [1, 1, 1] \triangleright$ pool weights for *roulette wheel* selection
- 7 **for** $i \leftarrow 1, 2, \dots, M$ **do**
- 8 $O \sim \mathcal{O} \triangleright$ sample a pool by the probability $\sigma(\mathbf{p})$
- 9 update the weight $\mathbf{p}[O]$ of the selected pool O to $(1 - \gamma) \cdot \mathbf{p}[O] + \gamma \cdot \max(-\Delta\theta, 0)$
- 10 $\mathbf{h} \leftarrow$ stacked heuristic values for all operators in O
- 11 $o \sim O \triangleright$ sample an operator by probability $\sigma(\mathbf{h})$
- 12 $\mathcal{G}' \leftarrow$ apply o to its relating subgraph(s) of \mathcal{G}
- 13 $\Pi' \leftarrow$ apply ESTC on updated subgraph(s) in \mathcal{G}'
- 14 **if** $\Delta\theta \leftarrow \theta(\Pi') - \theta(\Pi) < 0$ **then**
- 15 update \mathcal{G} to \mathcal{G}' and Π to Π'
- 16 **else**
- 17 update \mathcal{G} to \mathcal{G}' and Π to Π' with a probability of $\exp(-\frac{\Delta\theta}{t})$
- 18 call *Update-Pools*($\mathcal{G}, \mathcal{O}, \Pi, o$) if \mathcal{G} is updated
- 19 **if** $i \% S = 0$ or $\Delta\theta < 0$ **then**
- 20 call *Forced-Deduplication*($\mathcal{G}, \mathcal{O}, \Pi$)
- 21 update Π^* to Π if $\theta(\Pi) < \theta(\Pi^*)$
- 22 update the temperature t to $\alpha \cdot t$
- 23 **return** improved MCPP solution Π^*

We categorize G_i as a *light subgraph* if its coverage path cost $c(\pi_i)$ from ESTC is no larger than the average coverage path cost \bar{c} over all subgraphs; otherwise, it is a *heavy subgraph*. We define the *duplication set* $V^+ = \{v \in V \mid n_v > 1\}$ as the set of vertices included in more than one subgraph, where $n_v = \sum_{i=1}^k |\{x \in V_i \mid x = v\}|$ counts the occurrences of vertex $v \in V$ in all subgraphs.

LS-MCPP employs *grow operators* on light subgraphs to assign them new vertices to cover, *deduplicate operators* on heavy subgraphs to eliminate unnecessary duplication, and *exchange operators* to balance path costs between subgraphs with a large cost difference. The operators are detailed in Sec. V-C, designed to attain a cost-equilibrium MCPP solution with a low makespan. LS-MCPP then employs ESTC on each subgraph G_i to generate its coverage path and evaluate the makespan. LS-MCPP also periodically calls a deduplication function to exploit the current subgraphs to achieve a low-makespan solution.

B. Algorithm Description

Pseudocode (Alg. 2): LS-MCPP begins by computing the corresponding set \mathcal{G}_0 of subgraphs on the input initial MCPP solution Π_0 [Line 1]. It then initializes a set of three operator pools via a subroutine *Update-Pools* described in Sec. V-C, containing only grow operators, deduplicate operators, and

Algorithm 3: Forced-Deduplication

Input: set $\mathcal{G} = \{G_i\}_{i=1}^k$ of subgraphs,
 set $\mathcal{O} = \{O_g, O_d, O_e\}$ of pools,
 set $\Pi = \{\pi_i\}_{i=1}^k$ of coverage paths

- 1 **for** $i \in I$ sorted by descending $c(\pi_i)$ **do**
- 2 **while** $(u, v) \leftarrow$ any U-turn in π_i **do**
- 3 remove u, v from G_i and π_i
- 4 **for** $i \in I$ sorted by descending $c(\pi_i)$ **do**
- 5 **for** $o \in O_d$ sorted by descending $h(o)$ **do**
- 6 apply o on G_i and update π_i via ESTC
- 7 call *Update-Pools*($\mathcal{G}, \mathcal{O}, \Pi$)

exchange operators, respectively [Lines 3-4]. It also initializes the temperature scalar t and the pool weight vector \mathbf{p} [Lines 5-6]. The scalar t is scaled down by the decay factor $\alpha \in [0, 1]$ for every iteration [Line 22] to dynamically adjust the probability for LS-MCPP to become increasingly less likely to accept a non-improving solution over iterations, which is the *simulated annealing* strategy [59] to skip local minima. The vector \mathbf{p} represents the pool weight for each of the three aforementioned pools. During each iteration [Line 7], vector \mathbf{p} is used for roulette wheel selection [58] to select a pool O from \mathcal{O} [Line 8], where $\sigma(\mathbf{p})$ is the *softmax* function of \mathbf{p} determining the probability of selecting each pool. Once the pool O is selected, its corresponding weight $\mathbf{p}[O]$ is updated with a weight decay factor $\gamma \in [0, 1]$ [Line 9]. Similar to the pool selection, LS-MCPP samples an operator o from the selected pool O with $\sigma(\mathbf{h})$ as the probability function [Line 11]. Here, \mathbf{h} is the vector of heuristic function h evaluated on all operators in O [Line 10], detailed in Sec. V-D. Following this, LS-MCPP applies operator o to its relating subgraph(s) to obtain a new set \mathcal{G}' of subgraphs and a new solution Π' [Lines 12-13] and calculates the makespan increment $\Delta\theta$ to determine whether to accept the new solution Π' . LS-MCPP accepts the new solution Π' if it has a smaller makespan, and, otherwise, accepts it with a probability of $\exp(-\Delta\theta/t) \in [0, 1]$ [Lines 14-17]. Upon the acceptance of Π' , LS-MCPP updates the set \mathcal{O} of operator pools via *Update-Pools* (Alg. 4) [Line 18]. LS-MCPP also calls a subroutine *Forced-Deduplication* (Alg. 3) every S iteration or the makespan has been reduced in the current iteration [Line 20]. The iterative search procedure terminates when it reaches the maximum number of iterations M , and LS-MCPP returns the best solution Π^* found [Line 23].

Forced Deduplication: A critical aspect contributing to the success of LS-MCPP is restricting duplication to only those necessary over the neighborhood exploration, such as those in a narrow path connecting two separate regions of graph G . For this purpose, LS-MCPP uses the *Forced-Deduplication* subroutine as described in Alg. 3 to deduplicate all possible duplication in two folds. The first part iterates through each path $\pi_i \in \Pi$ in a cost-decreasing order to remove any U-turns [Lines 1-3]. A U-turn of π_i is defined as an edge traversal $(u \rightarrow v)$ in a ‘‘U-turn’’ pattern segment $p \rightarrow u \rightarrow v \rightarrow q$ in path π_i . If $u, v \in V_i \cap V^+$, then we can safely remove $(u \rightarrow v)$

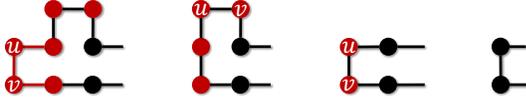


Figure 8. Subroutine *Forced-Deduplication* executed on a segment of coverage path π_i , where red and black circles represent vertices from $V_i \cap V^+$ and $V_i \setminus V^+$, respectively. Frames from left to right: a U-turn $(u, v) \in E(\pi_i)$ is found, and u, v are removed from π_i until no U-turn exists.

to reduce $c(\pi)$ without losing any vertex coverage or introducing any new duplication. Fig. (8) illustrates the procedure to remove U-turns for a path segment progressively. A special case of U-turn often occurs for incomplete hypervertices where the ‘‘U-turn’’ pattern becomes $p = v \rightarrow u \rightarrow v \rightarrow q$. The second part [Lines 4-6] iteratively process each graph G_i in descending order of path cost π_i [Line 4], applying each deduplicate operator o from O_d in descending order of its heuristic value $h(o)$ [Line 5]. After the above two parts, the function re-initializes all the pools in \mathcal{O} by calling *Update-Pools* (Alg. 4) [Line 7]. It is worth noting that the first part can remove edges that do not constitute deduplicate operators, and some deduplicate operators can remove edges that are not U-turns depending on the MST generated by ESTC. This complementary relation between the two parts contributes to the improved performance of LS-MCPP.

Initial Solution: For the choice of the initial solution Π_0 in LS-MCPP, we limit ourselves to existing multi-tree MCPP methods that can generate high-quality solutions efficiently. Specifically, we choose our initial solution from VOR or MFC, depending on which provides the better solution quality. Both methods are integrated with the proposed ESTC paradigm, as described in Sec. IV-D.

Remark 3. *Our LS-MCPP framework is inherently better suited to utilizing a multi-tree method for generating an initial MCPP solution, which aligns with its design to iteratively improve a set of k subgraphs using a local search strategy and apply ESTC to obtain coverage paths for evaluating the solution cost in each iteration. While a single-tree method can also provide initial subgraphs for LS-MCPP, induced by its solution coverage paths, our empirical finding [21] indicates that this approach typically results in significant overlaps among subgraphs and tends to reduce the effectiveness of the neighborhood operators within our LS-MCPP framework.*

C. Pools of Operators

We introduce three types of operators: grow operators, deduplicate operators, and exchange operators. These operators are designed to modify the boundaries of each connected subgraph $G_i = (V_i, E_i) \in \mathcal{G}$ of the input graph $G = (V, E)$. An operator alters the set \mathcal{G} while ensuring that its property of complete coverage (i.e., $\bigcup_{i=1}^k V_i = V$) and subgraph connectivity remains invariant. Each operator aims to introduce only essential duplication, with the potential to enhance solution quality.

Specifically, we consider both *edge-wise* and *vertex-wise* operations for these operators. An edge-wise operator alters a single edge at a time, whereas a vertex-wise operator alters a single vertex at a time. Altering a crossing edge $(u, v) \in E$

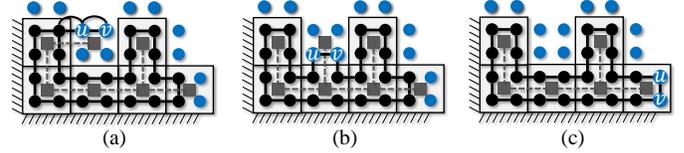


Figure 9. ESTC path after applying an edge-wise $o_g^{(e)}(i, \{u, v\})$ on subgraph $G_i = (V_i, E_i)$. Blue circles and black circles represent vertices of the boundary vertex set B_i and vertices of V_i , respectively. (a) An invalid operator without any parallel edge in G_i . (b)(c) Two valid edge-wise grow operators.

with $\delta_u \neq \delta_v$ often introduces unnecessary vertex duplications. Therefore, we restrict the edge-wise operators to only altering intra-hypervertex edges $E^{\text{intra}} = \{(u, v) \in E \mid \delta_u = \delta_v\}$. In practice, edge-wise operators generally perform well but sometimes cannot be constructed in configurations with many incomplete hypervertices, which may lead to early termination of the local search (Alg. 2). Vertex-wise operators are less restricted than edge-wise operators and tend to introduce unnecessary vertex duplications if applied to complete hypervertices. Therefore, vertex-wise operators serve only as complementary tools when no edge-wise operators are available, as reflected in the *Update-Pools* function (Alg. 4), which initializes or updates the operator pools.³

Grow Operator: A grow operator serves to expand a subgraph to cover additional vertices already covered by other subgraphs. Let B_i denote the set of ‘‘boundary’’ vertices that are not part of G_i but adjacent to a vertex of G_i , defined as $B_i = \{v \in V \setminus V_i \mid \exists (u, v) \in E, u \in V_i\}$. A vertex-wise grow operator $o_g^{(v)}(i, \{v\})$ adds vertex $v \in B_i$ and all its adjacent edges to subgraph G_i . An edge-wise grow operator $o_g^{(e)}(i, \{u, v\})$ adds edge $e = (u, v) \in E_i \cap E^{\text{intra}}$ with $u, v \in B_i$ and all its adjacent edges into G_i , where $\exists (p, q) \in E_i$ such that $(u, p), (v, q) \in E_i$. In essence, a valid $o_g(i, e)$ can only add an edge (u, v) if there exists a parallel edge (p, q) in G_i , as shown in the two examples in Fig. (9b) and (9c). This design choice avoids introducing undesirable routing to the coverage path, as shown in the example in Fig. (9a). After its execution, a grow operator $o_g^{(v)}(i, \{v\})$ (or $o_g^{(e)}(i, \{u, v\})$) effectively adds vertex v (or vertices u and v) to V^+ .

Deduplicate Operator: A deduplicate operator serves to remove unnecessary duplication from a subgraph. A vertex-wise deduplicate operator $o_d^{(v)}(i, \{v\})$ removes vertex $v \in V_i \cap V^+$ and all its adjacent edges from G_i such that G_i remains connected. For an edge $e = (u, v) \in E_i \cap E^{\text{intra}}$, let $\delta_e^t, \delta_e^b, \delta_e^l,$ and δ_e^r denote the four neighbors of δ_u (also equal to δ_v) based on the positioning of u and v . Fig. (10a) demonstrates one example that can be rotated to various symmetric cases (with only δ_e^t potentially containing both a vertex adjacent to u and one adjacent to v). An edge-wise deduplicate operator $o_d^{(e)}(i, \{u, v\})$ removes edge $(u, v) \in E_i \cap E^{\text{intra}}$ with $u, v \in V_i \cap V^+$ and all its adjacent edges from G_i such that G_i remains connected. If δ_u contains only u and v , then this removal is always valid and reduces the cost of the corresponding ESTC path of G_i , as shown in Fig. (10c).

³Our preliminary conference version [21] uses only edge-wise operators, which may lead to early termination issues. See Sec. VII-B for a detailed ablation study on the impact of vertex-wise operators.

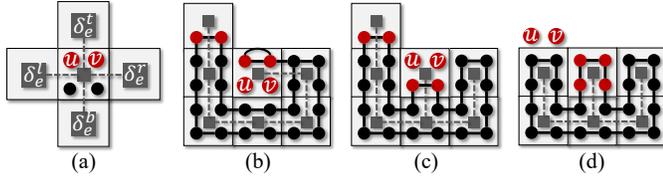


Figure 10. ESTC path after applying an edge-wise $o_d^{(e)}(i, \{u, v\})$ on subgraph $G_i = (V_i, E_i)$. Red circles and black circles represent vertices in $V_i \cap V^+$ and not in $V_i \setminus V^+$, respectively. (a) Four neighbors of the hypervertex $\delta_u (= \delta_v)$. (b) An invalid operator violating the third condition when δ_u is complete. (c)(d) Two valid edge-wise deduplicate operators.

Otherwise, this removal is constrained by three conditions: (1) any vertex contracted to δ_e^t is not in V_i ; (2) all the four vertices contracted to δ_e^b are in V_i ; and (3) if any $v \in V_i$ is contracted to δ_e^l (or δ_e^r), then both δ_e^l (or δ_e^r) and its (only) common neighboring hypervertex with δ_e^b must be complete, and all their vertices must also be in V_i . Such validity filtering on the edge-wise deduplicate operators is essential to avoid introducing unnecessary duplications to the corresponding ESTC path of G_i , as shown in Fig. (10). After its execution, a duplicate operator $o_d^{(v)}(i, \{v\})$ (or $o_d^{(e)}(i, \{u, v\})$) effectively removes vertex v (or vertices u and v) from V^+ if the vertex (or vertices) is not included in multiple subgraphs.

Exchange Operator: An exchange operator combines a grow operator and a deduplicate operator so that the deduplicate operator immediately removes the new duplication introduced by the grow operator, potentially improving the effectiveness of a single operator and speeding up the convergence of the local search. A vertex-wise exchange operator $o_e^{(v)}(i, j, \{v\})$ combines $o_g^{(v)}(i, \{v\})$ and $o_d^{(v)}(j, \{v\})$, where $v \in B_i \cap V_j$. An edge-wise exchange operator $o_e^{(e)}(i, j, \{u, v\})$ combines $o_g^{(e)}(i, \{u, v\})$ and $o_d^{(e)}(j, \{u, v\})$. Note that the operators o_g and o_d being combined must be both valid except that vertex v for a vertex-wise $o_d^{(v)}(j, \{v\})$ and vertices u and v for an edge-wise $o_d^{(e)}(j, \{u, v\})$ are not required to be in V^+ since they are not necessarily covered by subgraphs other than G_j before the grow operator is applied. After the execution of an exchange operator, V^+ remains unchanged.

Update Pools: The *Update-Pools* function in Alg. 4 initializes the set \mathcal{O} of pools of operators [Line 5] or updates the set \mathcal{O} whenever \mathcal{G} is changed by some operator o [Lines 2-3]. Given the set ΔV of dirty vertices (defined in line 3 or line 5), *Update-Pools* first removes all operators involving any $v \in \Delta V$ from each pool [Line 6]. For each $v \in \Delta V$, it then iterates each light subgraph to add a valid edge-wise grow operator to O_g if one exists; otherwise, it adds a vertex-wise grow operator [Lines 8-12]. Similarly, deduplicate operators are added to O_d by iterating each heavy subgraph [Lines 13-17]. Finally, the pool O_e of exchange operators is updated by combining the newly added operators in O_g and O_d [Line 18].

Properties: After any valid operator is executed, (1) every $G_i \in \mathcal{G}$ remains connected since only deduplicate operators can remove edges but a valid deduplicate operator guarantees that G_i remains connected, and (2) the property of the union of the vertices of all subgraphs being the vertex set V is not affected since no vertex coverage is lost. Therefore, an LS-MCPP solution ensures complete coverage due to Theorem 3.

Algorithm 4: Update-Pools

Input: set $\mathcal{G} = \{G_i\}_{i=1}^k$ of subgraphs,
 set $\mathcal{O} = \{O_g, O_d, O_e\}$ of pools,
 set $\Pi = \{\pi_i\}_{i=1}^k$ of coverage paths,
 [optional] most recent applied operator o

- 1 **if** o is provided as an input **then**
- 2 $V_o \leftarrow$ the set of involved vertices in o
- 3 $\Delta V \leftarrow V_o \cup \{u \in V \mid v \in V_o, (u, v) \in E\}$
- 4 **else**
- 5 $\Delta V \leftarrow V \triangleright$ update for all vertices in $G = (V, E)$
- 6 remove operators from each $O \in \mathcal{O}$ involving any $v \in \Delta V$
- 7 **for** $v \in \Delta V$ **do**
- 8 **for** $i \in \{j \in I \mid c(\pi_j) \leq \bar{c}\}$ **do**
- 9 **if** found a valid $o_g^{(e)} = (i, \{u, v\})$ **then**
- 10 $O_g \leftarrow O_g \cup \{o_g^{(e)}\}$
- 11 **else**
- 12 $O_g \leftarrow O_g \cup \{o_g^{(v)} = (i, \{v\})\}$
- 13 **for** $i \in \{j \in I \mid c(\pi_j) > \bar{c}\}$ **do**
- 14 **if** found a valid $o_d^{(e)} = (i, \{u, v\})$ **then**
- 15 $O_d \leftarrow O_d \cup \{o_d^{(e)}\}$
- 16 **else**
- 17 $O_d \leftarrow O_d \cup \{o_d^{(v)} = (i, \{v\})\}$
- 18 update O_e by iterating newly added operators in O_g, O_d

D. Heuristics for Operator Sampling

To better explore the solution neighborhood, LS-MCPP should carefully determine which operator to sample from the selected pool (line 11 of Alg. 2). We propose three heuristic functions tailored to the three types of operators to evaluate their potential in guiding the neighborhood search and improving the solution. An operator with a larger heuristic value results in a higher probability of being sampled. The heuristic function for grow operators has two considerations: (1) prioritizing growing light subgraphs with smaller path costs and (2) prioritizing covering vertices with less duplication. Formally, the heuristic value $h(o_g)$ for a grow operator $o_g(i, V_{o_g})$ is defined as $h(o_g) = -k \cdot c(\pi_i) - \frac{1}{|V_{o_g}|} \cdot \sum_{v \in V_{o_g}} n_v$. Note that $k \geq n_v$ holds for any $v \in V$, thereby prioritizing consideration (1) over consideration (2). The heuristic function for deduplicate operators $o_d(i, V_{o_d})$ is defined as the exact opposite of the one for grow operators: $h(o_d) = k \cdot c(\pi_i) + \frac{1}{|V_{o_d}|} \cdot \sum_{v \in V_{o_d}} n_v$. The heuristic function for exchange operators $o_e(i, j, V_{o_e})$ is defined as $h(o_e) = c(\pi_j) - c(\pi_i)$ to prioritize pairs of subgraphs with larger differences in their path costs.

VI. DECONFLICTING MCPP SOLUTIONS

In this section, we address inter-robot conflicts in coverage paths produced by LS-MCPP or other grid-based MCPP methods. Our MCPP formulation (Problem 1) does not inherently prevent collisions among robots. In an MCPP solution with a small makespan, robots typically cover different regions and are distributed across the grid, which lowers the likelihood of conflicts. However, certain regions remain prone to congestion

and conflicts, particularly around clustered root vertices or narrow corridors that only allow one robot to pass at a time.

To resolve these potential conflicts, we formulate the deconflicting task as a variant of MAPF, leveraging MAPF techniques to ensure each robot follows its designated coverage path while avoiding collisions with others. Specifically, we develop a two-level hierarchical planner, following standard MAPF practices, to effectively deconflict MCPP solutions: The high-level planner employs Priority-Based Search (PBS) [60] to assign priorities among pairs of conflicting robots; The low-level planner is a novel algorithm that plans a continuous-time path for each robot that avoids conflicting with any higher-priority robots. Our key extensions from standard MAPF lie in 1) our treatment of non-uniform edge weights, which translate to variable time costs for robot actions and 2) an adaptive approach of the low-level planner for time-efficient single-robot path planning with multiple goals.

A. Formulation of Deconflicted MCPP

Given an input solution $\Pi = \{\pi_i\}_{i=1}^k$ for the MCPP instance (G, I, R) , we define the *Deconflicted MCPP* problem, which aims to construct a set $\mathcal{T} = \{\tau_i\}_{i=1}^k$ of k conflict-free trajectories. A *trajectory* τ_i for robot i is an ordered sequence $(\mathbf{x}_1, \dots, \mathbf{x}_{|\tau_i|})$ of time-embedded states $\mathbf{x} = (v, t)$ indicating that the robot stays at vertex $v \in V$ at time $t \geq 0$. Given path $\pi_i = (v_1, \dots, v_{|\pi_i|})$ where $v_1 = v_{|\pi_i|} = r_i \in R$, trajectory τ_i must start from state $\mathbf{x}_1 = (v_1, 0)$, visit all vertices in the specified order of π_i , and end with state $\mathbf{x}_{|\tau_i|} = (v_{|\pi_i|}, \cdot)$. Formally, there exist indices $a_1 = 1 < a_2 < \dots < a_{|\pi_i|} = |\tau_i|$ such that $\mathbf{x}_{a_j} = (v_j, \cdot)$ for all $j = 1, \dots, |\pi_i|$. To transit from state (u, t) to the next state (v, t') in a trajectory, the robot performs a *wait-and-move* action, where it waits at v for $t' - w_e - t \geq 0$ and then moves from u to v along edge $e = (u, v) \in E$ with a time cost of $w_e > 0$. Let $\tau[j]$, $\mathbf{x}.v$ and $\mathbf{x}.t$ denote the j -th state in τ , the vertex and time in each state \mathbf{x} , respectively. The quality of a Deconflicted MCPP solution \mathcal{T} is evaluated using the *makespan* metric, denoted by $\theta(\mathcal{T}) = \max\{\tau_1[|\tau_1|].t, \dots, \tau_k[|\tau_k|].t\}$, which is consistent with that used for MCPP.

Definition 1 (Conflict). *A conflict between trajectory τ with states \mathbf{x}_j for $j = 1, \dots, |\tau|$ and another trajectory τ' with states $\mathbf{x}'_{j'}$ for $j' = 1, \dots, |\tau'|$ occurs if there exist states \mathbf{x}_j and $\mathbf{x}'_{j'}$ such that their vertices $\mathbf{x}_j.v = \mathbf{x}'_{j'}.v$ and the two time intervals $(\mathbf{x}_{j-1}.t, \mathbf{x}_{j+1}.t)$ and $(\mathbf{x}'_{j'-1}.t, \mathbf{x}'_{j'+1}.t)$ intersect. To ensure conflict checks are valid for all j , we assume two dummy bounds $\mathbf{x}_0.t = 0$ and $\mathbf{x}_{|\tau|+1}.t = +\infty$.*

In our definition of conflicts, a vertex is considered occupied by a robot not only when the robot stays at it (as in typical MAPF frameworks [61]) but also when a robot is moving into or out of it. This conservative definition subsumes all types of conflicts outlined in MAPF literature [17], including vertex conflicts, edge conflicts, following conflicts, and swapping conflicts. This generalized definition accommodates continuous time and non-uniform edge weights, unlike traditional MAPF definitions that assume discretized, uniform time steps.

Remark 4. *Instead of first solving MCPP and then deconflicting the resulting paths as a post-processing procedure, one could integrate the deconfliction at the end of each LS-MCPP iteration (Alg. 2), potentially improving performance beyond our proposed sequential approach. However, such an approach incurs significantly higher runtime due to the computational demands of repeated runs of the deconfliction (see Tab. II).*

B. The High-Level Planner

The high-level planner employs Priority-Based Search (PBS) [60] to resolve conflicts defined in Definition 1. PBS searches a priority tree where each node N stores a unique priority ordering given by a set \prec_N of ordered pairs of robots and a set $N.\mathcal{T}$ of k trajectories that respect the prioritized planning scheme specified by \prec_N . The root node contains an empty priority set and a set of potentially conflicting trajectories. When PBS expands a node N , it checks $N.\mathcal{T}$ for conflicts. If no conflicts are found, then N is a goal node, and $N.\mathcal{P}$ is the solution. If conflicts exist between the trajectories of robots i and j , PBS generates two child nodes, N_1 and N_2 , and adds the pair $i \prec j$ (robot i has a higher priority than robot j) to N_1 and $j \prec i$ to N_2 . For child node N_1 , PBS invokes a low-level planner to update the trajectories of robot j and other robots j' with $j \prec_{N_1} j'$, ensuring $N_1.\mathcal{T}$ respects \prec_{N_1} . PBS processes N_2 symmetrically. The low-level planner computes a trajectory for each robot that avoids conflicts with higher-priority trajectories. PBS performs Depth-First Search on the priority tree by selecting the child node whose trajectories have a lower makespan in the next iteration, thus exploring the space of all priority orderings.

PBS is incomplete in general; however, with a complete low-level planner, it is complete for well-formed instances of various MAPF problems [60]. Specifically, PBS with a complete low-level planner is complete for well-formed Deconflicting MCPP instances, defined as follows.

Definition 2 (Well-Formedness). *A Deconflicted MCPP instance with a given input solution Π of MCPP instance (G, I, R) is well-formed if, for all $i \in I$, path $\pi_i \in \Pi$ does not traverse any root vertex $r_j \in R$ of robot $j \neq i$.*

Since the given solution Π might not satisfy the well-formed condition, before deconflicting the given Π , we preprocess each $\pi_i \in \Pi$ to remove any $v \in R \setminus \{r_i\}$ from π_i , which preserves complete coverage while ensuring the completeness of PBS when, for all i , G remains connected with all vertices $v \in R \setminus \{r_i\}$ removed. A sufficient condition is that G remains connected with all vertices $v \in R$ removed, which is used to generate instances for our experiments in Sec. VII.

Remark 5. *Alternatively, Conflict-based Search (CBS) [61] can serve as the high-level planner. CBS is optimal for MAPF problems when an optimal low-level planner is used. However, CBS generally incurs significantly longer runtimes than PBS due to its Best-First Search on a tree whose size grows exponentially with the number of conflicts. This runtime increase is even more pronounced in our multi-goal MAPF setting, where an optimal low-level planner is notably more computationally intensive than in standard single-goal MAPF.*

According to [43] which applies CBS to OMG-MAPF, a special case of our Deconflicting MCPP problem with unweighted edges, CBS only scales to sequences of up to 5 goals on relatively small graphs.

C. The Low-level Planner

The lower-level computes a trajectory τ_i for robot i that sequentially visits vertices along the given π_i . It uses a *reservation table* that records the reserved time intervals for each vertex, ensuring that the resulting trajectory τ_i avoids conflicts with higher-priority trajectories.

Reservation Table: To accommodate planning in an edge-weighted graph, we adopt the idea from Safe Interval Path Planning (SIPP) [62], [63] to reserve continuous-time intervals, instead of discrete time steps in the classic vertex-time A* [64]. The reservation table is indexed by a vertex, with the corresponding entry recording the *reserved intervals* during which the vertex is occupied. Specifically, following the conflict definition (Definition 1), given state (v, t) , the reservation table records the reserved interval $[t_p, t_s)$ for vertex v , where (v_p, t_p) and (v_s, t_s) are the preceding and succeeding states of state (v, t) , respectively. Notably, the reservation also records $[0, t_s)$ for the first state $(r_j, 0)$ and $[t_p, +\infty)$ for the last state in each higher-priority trajectory τ_j . The *safe intervals* for each vertex v are obtained as the complements of its reserved intervals with respect to interval $[0, +\infty)$.

SIPP Node Expansion: The low-level planner uses A* search to generate a search tree, where each node n is uniquely identified by a pair $\langle v, [lb, ub) \rangle$, consisting of a vertex v and one of its safe intervals $[lb, ub)$. The g -value $g(n) \in [lb, ub)$ of a node n stores the earliest arrival time at which the robot can stay at vertex v within the safe interval $[lb, ub)$ along any path. When expanding node $n = \langle v, [lb, ub) \rangle$, the search generates a child node $n' = \langle u, [lb', ub') \rangle$ with $g(n')$ for each vertex u adjacent to v and each safe interval $[lb', ub')$ of u , corresponding to the action of waiting at v for the minimal duration $g(n') - w_{(v,u)} - g(n) \geq 0$, then departing from v at time $g(n') - w_{(v,u)}$ and arriving at u at the earliest available time $g(n')$. The search discards node n' if no valid arrival time $g(n')$ is available such that $g(n') \geq lb'$ with the corresponding departure time $g(n') - w_{(v,u)} < ub$. If the search generates a duplicate copy of an already generated node, it keeps only the dominating one with the smallest g -value.

Conflict-Free Trajectory Reconstruction: The trajectory τ_i for π_i is generated by recursively backtracking the search tree that starts from $\langle \pi_i[|\pi_i|], \cdot \rangle$ and ends at $\langle \pi_i[1], \cdot \rangle$. For each SIPP node $n = \langle v, \cdot \rangle$ along the search tree, we recover a state $x = (v, g(n))$ and append it to τ_i .

Given multiple ordered goals along a path π , we design three low-level planner approaches, *chaining*, *multi-label*, and *adaptive*.

Chaining Approach (Incomplete): This approach iteratively computes a trajectory by constructing (trajectory) segments between consecutive vertices in π . Starting from vertex $\pi[1]$ with time $t_1 = 0$, each j -th iteration (for $j = 1, \dots, |\pi| - 1$) computes a segment from vertex $\pi[j]$ with the given arrival time t_j to vertex $\pi[j + 1]$ with a computed arrival time t_{j+1} .

Each segment computation uses the aforementioned A* search, starting from the node with vertex $\pi[j]$ and a safe interval containing t_j . The computation succeeds if the search expands a goal node $n = \langle \pi[j + 1], \cdot \rangle$ with arrival time $t_{j+1} = g(n)$; otherwise, it fails. The search uses the precomputed shortest-path distance from any vertex v to $\pi[j + 1]$ as the heuristic value for all nodes associated with v . The final trajectory, if all segments are successfully computed, is constructed by concatenating the segments. However, this chaining approach is incomplete [43] because it does not allow backtracking. For example, a successful trajectory may require searching from alternate safe intervals of $\pi[j]$ rather than the one containing t_j , which this approach does not account for.

Multi-Label Approach (Complete and Optimal; Novel for Continuous Time): This approach extends Multi-Label A* [65]–[67], originally developed for unweighted graphs, to plan a time-minimal trajectory through ordered goals $\pi[j]$ ($j = 1, \dots, |\pi|$) in a single search. This approach adds a *label* dimension to the state space and thus searches with nodes uniquely identified by a tuple $\langle v, [lb, ub), l \rangle$, consisting of a vertex v , one of the safe intervals $[lb, ub)$ of v , and a label l indicating the current goal $\pi[l]$ —resulting in a novel path-planning algorithm named Multi-Label SIPP (ML-SIPP). The search of ML-SIPP follows the same node expansion rule as standard SIPP, letting a child node inherit its label l from its parent, except when its vertex is $\pi[l]$, in which case its label changes to $l + 1$. The search terminates when it expands a (final-)goal node $\langle \pi[|\pi|], \cdot, |\pi| + 1 \rangle$. ML-SIPP uses the precomputed concatenated shortest-path distance heuristic:

$$h(\langle v, \cdot, l \rangle) = \text{dist}(v, \pi[l]) + \sum_{j=l}^{|\pi|-1} \text{dist}(\pi[j], \pi[j + 1]).$$

Since this heuristic is admissible—never overestimating the cost of reaching the (final-)goal node—ML-SIPP is a complete and optimal low-level planner.

Adaptive Approach (Complete; Novel): The adaptive approach, shown in Alg. 5, seeks to enhance practicability by combining the efficiency of the chaining approach and the completeness of ML-SIPP. It sequentially applies ML-SIPP to subsequences of goals in path π , each time starting from a subsequence of a single goal [Line 9] and backtracking to include the previous goal whenever ML-SIPP fails to find a trajectory for the current subsequence [Lines 10-12]. The backtracking depth is limited by a small integer $b_{\max} \in \mathbb{Z}^+$ (we set it as 5 in Sec. VII). Upon successfully planning for a single goal, the approach resets current backtracking depth b [Lines 6-9]. Our adaptive approach is complete since it calls ML-SIPP for the entire sequence of goals in π (i.e., the multi-label approach) in the worst case when ML-SIPP fails with the window size b_{\max} [Lines 14-15]. In practice, we postpone the ML-SIPP call on the entire goal sequence and mark the PBS node N for the multi-label approach. If the high-level PBS tree reports failure once its search stack is empty, we pick one postponed node, unmark it, and call the multi-label approach for the node with children nodes branched. This strategy greatly improves the overall efficiency of the deconflicting procedure, as each call of the multi-label

Algorithm 5: Adaptive-Low-Level-Planner

Input: path $\pi = (v_1, \dots, v_{|\pi|})$ for robot i ,
reservation table rt w.r.t. \prec_N of PBS node N
Param: max backtracking size $b_{\max} \in \mathbb{Z}^+$

- 1 remove every $v \in R \setminus \{r_i\}$ from π
- 2 $\mathbf{x}_{\text{last}} \leftarrow (v_1, 0)$, $j = 2$, $b \leftarrow 1$
- 3 $\tau \leftarrow (\mathbf{x}_{\text{last}}, S_{\text{goal}} \leftarrow (v_j,)$
- 4 **while** $j \leq |\pi|$ **do**
- 5 find a trajectory from \mathbf{x}_{last} to S_{goal} via ML-SIPP
- 6 **if** a valid trajectory $\tilde{\tau}$ is found **then**
- 7 append $\tilde{\tau}$ to the end of τ
- 8 $j \leftarrow j + 1$, $b \leftarrow 1$
- 9 $S_{\text{goal}} \leftarrow (v_j,)$
- 10 **else if** $b \leq b_{\max}$ or τ is not empty **then**
- 11 pop the last state from τ until v_{j-b}
- 12 $S_{\text{goal}} \leftarrow (v_{j-b},) + S_{\text{goal}}$, $b \leftarrow b + 1$
- 13 **else**
- 14 mark N as postponed for multi-label approach
- 15 **return** \emptyset
- 16 $\mathbf{x}_{\text{last}} \leftarrow$ the last state in τ
- 17 **return** τ

approach is highly time-consuming, especially when dealing with large-scale instances.

D. Integrating Turning Costs

So far, the trajectory computation assumes the robots are holonomic, that is, they can move omnidirectional without the need to change their headings between consecutive move actions. We now show that one can remove this assumption for non-holonomic robots by adding an *orientation* $\psi \in \{E, N, W, S\}$ to the state space, restricted to East (E), North (N), West (W), and South (S) in the four-neighbor grid graph G . Correspondingly, we allow a robot to *turn* at vertex v from ψ to $\psi' \neq \psi$ at a time cost of $t' - t = C \cdot \frac{|\psi' - \psi|}{90^\circ} > 0$, where C is a constant cost for a 90° turns. Effectively, each robot performs a *turn-wait-move* action. With this modification, one can deconflict an MCPP solution for non-holonomic robots.

Additionally, the above modification requires the parallel rewiring optimization for ESTC to consider the turn costs when calculating the path cost difference (Sec. IV-C). Other components of LS-MCPP (Alg. 2), including ESTC (Alg. 1) and its turn reduction optimization, all remain unchanged.

VII. EXPERIMENTS AND RESULTS

This section describes our numerical results of simulations on an Ubuntu 22.04 PC with a 2.50GHz Intel[®] Core[®] i5-13490F CPU and 32GB RAM. Our proposed algorithms and the baseline methods are publicly available at Github⁴, with all the benchmark instances and reported numerical results provided as well. In Sec. VII-D, we deployed the proposed algorithms on two *Pepper*⁵ robots in an indoor environment to validate our planning pipeline for MCPP.

A. Experiment Setup

Instances: We generate a set of instances to benchmark different methods utilizing the four-way connected 2d undirected grid graphs from previous grid-based MCPP works [15], [21] as well as the 2D pathfinding benchmark [68]. As shown in Fig. 11, the graphs in the top row have unweighted edges with a uniform weight of 1, whereas the graphs in the bottom row have weighted edges whose weights are randomly generated from 1 to 3. Originally, each graph can be contracted to a hypergraph where all its hypervertices are complete. The original graphs are then used to generate 12 groups of 12 mutated MCPP instances (i.e., mutations) of k robots, resulting from a removal ratio $\rho = 0, 1, \dots, 11$ and a seed $s = 0, 1, \dots, 11$. Each mutation (G, I, R) with ρ and s is constructed by (1) randomly sampling (with seed s) a set of k vertices from G as R and (2) randomly removing (with seed s) vertices from G by $x = (\frac{100}{12} \cdot \rho)\%$, meanwhile ensuring the root vertex of any robot and all its neighbors are not removed. The root vertices in R are sampled to ensure they are not cut vertices for constructing well-formed instances [60], [69], that is, a robot can wait indefinitely at its root vertex without blocking any other robot. Note that the graph edge weights of each mutation are also randomly generated using its corresponding seed s . For all mutations, we set $C = 0.5$ as the 90° turning cost as described in Sec. VI-D.

Baselines: We describe all the methods being compared in our experiments. (1) *Single-Robot CPP*: To verify the effectiveness of our ESTC with the proposed hyperedge weight definition, we use a baseline method, namely ESTC-UW, that computes the circumnavigating coverage path on a spanning tree assuming the edges are unweighted. We also performed an ablation study on the two local optimizations for ESTC: parallel rewiring (PR) in Sec. IV-A and turn reduction (TR) in Sec. IV-C. (2) *MCPP*: As mentioned in Sec. IV-D, the proposed ESTC algorithm can be integrated into various grid-based MCPP methods, making them immediately applicable to MCPP. To reiterate, VOR, MFC, MIP, and MSTC* are the comparing baseline methods with our LS-MCPP. In addition, we also compare LS-MCPP with and without the vertex-wise operators that are new to our previous work [21]⁶, namely LS(+VO) and LS(-VO) respectively. Note that all the methods compared use ESTC with TR and PR applied by default. (3) *Deconflicted MCPP*: Given LS(+VO) solutions to produce the initial time-embedded trajectories, we compare the deconflicting performance using the chaining approach (CHA), the multi-label approach (MLA), and the adaptive approach (ADA) as the low-level planner for PBS.

Parameters: Given an MCPP instance (G, I, R) , we describe the settings of all the parameters used in the simulation experiments. For our LS-MCPP (Alg. 2), we set the maximum number of iterations $M = 1.0e^3 \times \sqrt{|V|/k}$ for a good tradeoff between planning efficiency and solution quality, the forced deduplication step size $S = \lfloor M/20 \rfloor$ to arrange at least 20 *Forced-Deduplication* calls, the temperature decay factor

⁴<https://github.com/reso1/LS-MCPP>

⁵<https://www.aldebaran.com/en/pepper>

⁶In the ablation study of our previous work [21], we have shown the effectiveness of grow, deduplicate, and exchange operators as well as the use of Forced Deduplication (Alg. 3) and heuristic sampling for LS-MCPP.

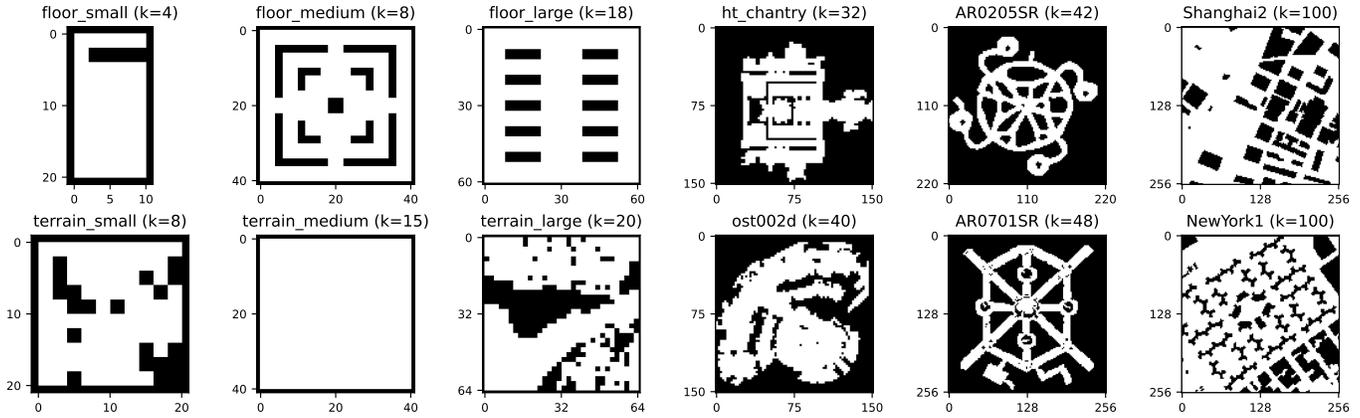


Figure 11. MCPP instances with k robots for the simulation experiments. Top row: unweighted grid graphs. Bottom row: weighted grid graphs.

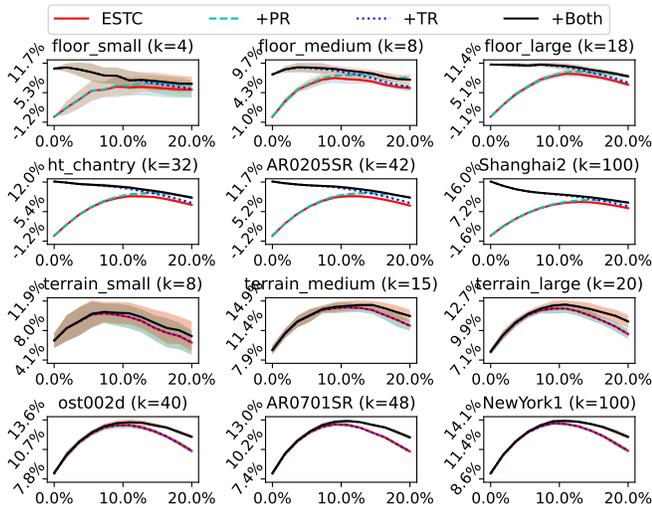


Figure 12. Ablation study for single-robot CPP. Each subplot displays the removal ratio ρ (x-axis) versus cost reductions compared to ESTC-UW (y-axis). The cost statistics (mean and variance) are computed across all 12 seeds.

$\alpha = \exp \frac{\log 0.2}{M}$ to decrease the temperature t from 1 to 0.2, and the pool weight decay factor $\gamma = 1.0e^{-2}$. For MIP, we set a runtime limit of one hour for the mixed-integer program solving. For the Deconflicted MCPP algorithms in Sec. VI-C, we set a total runtime limit of one hour for the overall procedure. Specifically for the adaptive approach (Alg. 5), we find that setting the max backtracking size $b_{\max} = 5$ is sufficient to resolve all local conflicts efficiently.

B. Ablation Study

Single-Robot CPP: In Fig. 12, we compare all four ESTC variants with the baseline ESTC-UW. We observe a more critical cost reduction on weighted graphs (upper 6 subplots) than on unweighted graphs (lower 6 subplots). For unweighted graphs, TR results in significant performance improvements since uniformly weighted edges can be aligned in one direction to minimize turning costs. For weighted graphs, TR is much less effective as there are few tie-breaking cases for reducing the turning costs, while PR achieves higher cost reductions

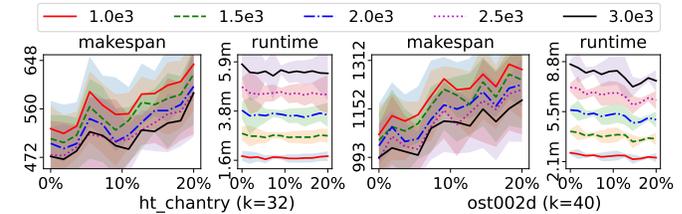


Figure 13. Performance of LS(+VO) solutions with varying iteration limits. The maximum number of iterations is scaled as $M = b \times \sqrt{|V|}/k$, where different colored lines represent different base scalars b .

as the vertex removal percentage increases. In general, ESTC shows its effectiveness for single-robot CPP, and both PR and TR contribute to substantial path cost reductions.

MCPP: We compare the solution quality of MCPP between different methods in Fig. 14. Generally speaking, the mutated instances become more complex as the vertex removal ratio increases. Our LS-MCPP consistently outperforms other methods across all instances, while the performance of the other methods varies, with each showing advantages in different instances. Note that due to the huge memory and runtime requirements of MIP, we only test it on relatively smaller instances (i.e., names with a suffix of *small* or *medium*), and its performance is highly restricted by the runtime limit of 1 hour. Specifically for our LS-MCPP, as the vertex removal ratio increases, LS(+VO) performs better with the additional vertex operators compared to LS(-VO). However, when the instances become more complex with a large map and more robots, the solution quality of LS(+VO) gets closer to LS(-VO) due to an inadequate number of local search iterations limited by M . To demonstrate that we can further improve the solution quality by saturating LS(+VO) with more iterations, we compare the LS(+VO) solutions with increasing maximum numbers of iterations, as shown in Fig. 13. The solution quality steadily improves as the maximum iteration limit M increases with a linear increase in runtime. In practice, M can be flexibly adjusted according to the available runtime budget, allowing a trade-off between better solution quality and faster planning.

Deconflicted MCPP: We explain the Deconflicted MCPP

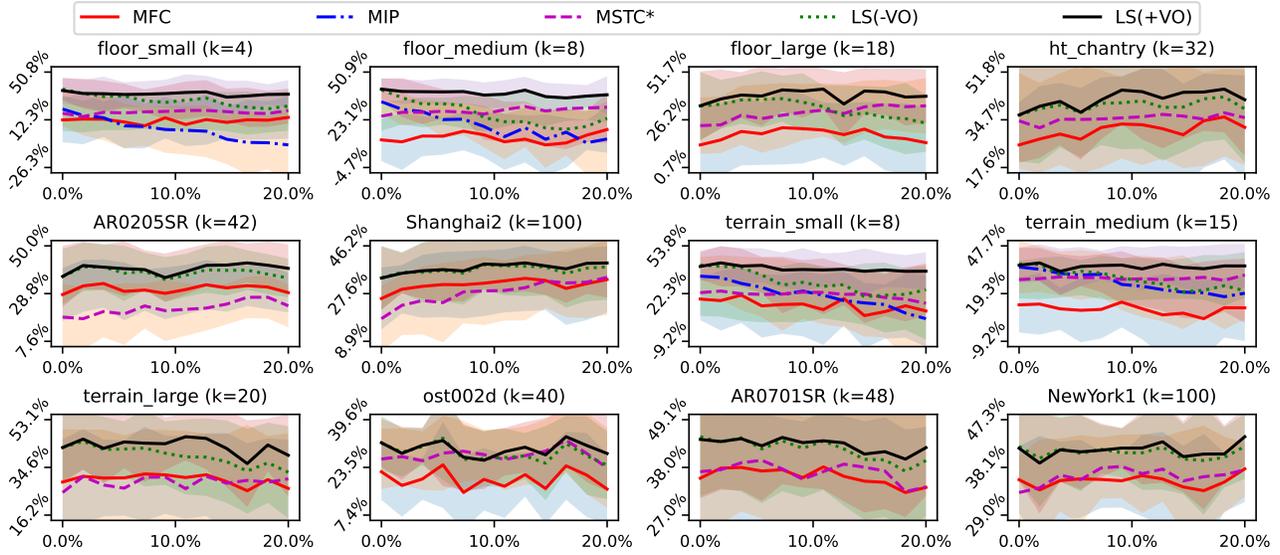


Figure 14. Ablation study for MCPP. Each subplot displays the removal ratio ρ (x-axis) versus makespan reductions compared to VOR (y-axis). The makespan statistics (mean and variance) are computed across all 12 seeds.

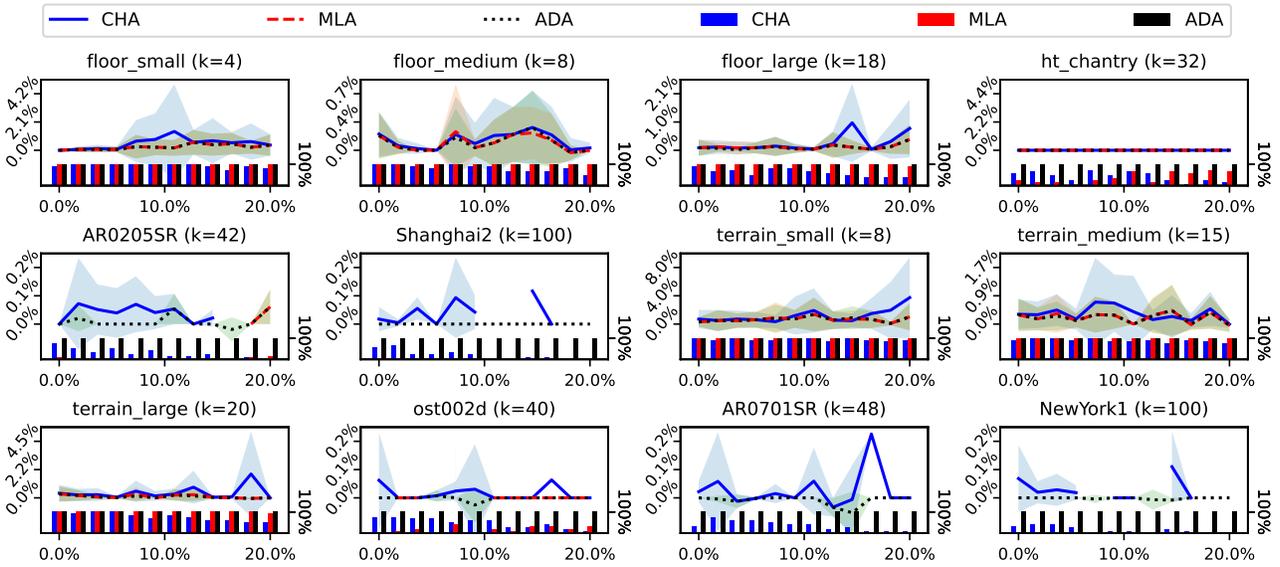


Figure 15. Ablation study for Deconflicted MCPP. Each subplot displays the removal ratio ρ (x-axis) versus makespan increments relative to the input LS(+VO) MCPP solution (left y-axis) and success rates (right y-axis). The makespan statistics (mean and variance) are computed across seeds where CHA, MLA, and ADA all succeed; when certain approach(es) fail in all seeds, statistics are computed across seeds where the remaining approach(es) succeed.

ablation study results in Fig. 15. CHA demonstrates the highest makespan increments and the lowest success rate due to its inability to backtrack to previous goal states, resulting in incompleteness. This limitation makes CHA susceptible to “dead-ends” when no valid succeeding states are available. Consequently, it reports a higher number of failure nodes during PBS search, which either prevents PBS from finding a valid priority ordering or significantly increases the number of PBS nodes explored (sometimes exceeding the runtime limit in large-scale instances). MLA produces high-quality trajectory sets, with no makespan increases for most of the “easier” instances it can handle within the runtime limit.

However, for more complex instances involving larger maps and more robots, the added label dimension across the entire coverage path expands the search space considerably, which significantly increases planning time for each PBS node and results in a lower success rate due to timeouts. Our ADA succeeds in all instances, achieving high solution quality with no makespan increases in most cases. This is due to its ability to restrict the label-augmented search space within a local window, rather than across the entire coverage path as MLA, and only when backtracking to previous states is necessary. In Fig. 16, we also compare the deconflicting results of different MCPP solutions for *floor_large* and *terrain_large* using ADA

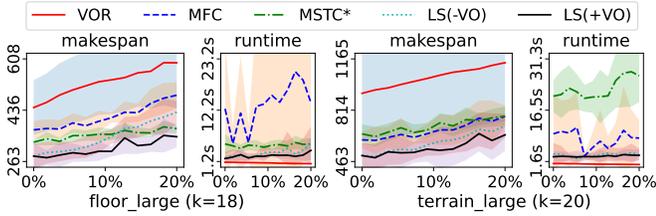


Figure 16. Deconflicting different MCPP solutions using ADA as the low-level planner for PBS. The left two subplots show the removal ratio ρ (x-axis) versus makespan (left y-axis) and success rates (right y-axis). The makespan statistics (mean and variance) are computed across seeds where ADA succeeds for all input MCPP solutions.

as the low-level planner for PBS. The results demonstrate the better solution quality of the proposed LS-MCPP when used as the input MCPP solution for deconfliction. In contrast, as MFC and MSTC* have larger makespans due to the higher overlapping covered regions across the map, the deconflicting procedures on them are more time-consuming. Notably, there is one failure case of deconflicting the MFC solution where the one-hour runtime limit is exceeded. For VOR, as there are no overlapping regions in its solution, PBS immediately returns the initial time-embedded trajectories without any conflict.

C. Planning Runtime Comparison

We report the planning runtime of all the MCPP methods and the Deconflicted MCPP methods across all the 12 groups of instances in Fig. 11. We show the planning runtime comparison of methods for MCPP in Tab. I, where our LS-MCPP requires more time to explore the solution neighborhood. In particular, adding vertex-wise operators only increases the runtime of LS-MCPP by a little if we compare LS(-VO) and LS(+VO). Nevertheless, it produces high-quality MCPP solutions within around 13 minutes even for large-scale instances such as *NewYork1* and *Shanghai2*, which is acceptable for offline planning in real-world robotics applications.

In Tab. II, we compare the planning runtime of the three low-level planners with PBS as the high-level planner, using the LS(+VO) solution as input. Each approach-instance data cell displays the runtime for instances where all three approaches succeed at the top, and the runtime for instances where each approach succeeds is shown at the bottom. Note that for the top item, when one or more approaches fail for all 12 instances, we also include the runtime for instances where the other approaches all succeed. We can see that for the relatively easier instances where all three approaches succeed at least one mutation, CHA is slightly faster than ADA, whereas MLA is the slowest. For more complex instances (e.g., *NewYork1*) where CHA only succeeds at a few mutations (see Fig. 15), ADA can still produce a set of valid trajectories for all of them within 10 minutes, and sometimes even faster than CHA who needs to explore more PBS nodes to find a successful set of trajectories.

Table I
RUNTIME COMPARISON OF MCPP METHODS.

	<i>floor small</i>	<i>floor medium</i>	<i>floor large</i>	<i>ht_ch antry</i>	<i>AR02 05SR</i>	<i>Shang hai2</i>
VOR	0.01±0.01s	0.03±0.01s	0.1±0.01s	0.2±0.03s	0.4±0.04s	1.5±0.1s
MFC	0.02±0.01s	0.21±0.02s	1.0±0.08s	6.8±0.11s	20±1.35s	4.7±0.3m
MIP	1h	1h	/	/	/	/
MSTC*	0.01±0.01s	0.12±0.05s	0.6±0.24s	3.5±1.63s	9.6±4.4s	1.7±0.7m
LS(-VO)	2.38±2.50s	0.57±0.2m	0.9±0.1m	1.8±0.2m	3.0±0.3m	9.9±1.5m
LS(+VO)	2.87±1.89s	0.69±0.1m	0.9±0.1m	2.1±0.2m	3.4±0.3m	13±0.8m
	<i>terrain small</i>	<i>terrain medium</i>	<i>terrain large</i>	<i>ost002d</i>	<i>AR07 01SR</i>	<i>New York1</i>
VOR	0.01±0.01s	0.04±0.01s	0.1±0.01s	0.4±0.04s	0.6±0.1s	1.6±0.1s
MFC	0.04±0.01s	0.40±0.04s	1.1±0.22s	15±1.46s	0.7±0.1m	5.4±0.2m
MIP	1h	1h	/	/	/	/
MSTC*	0.04±0.02s	0.50±0.21s	1.4±0.55s	15±6.71s	0.5±0.3m	2.6±0.8m
LS(-VO)	5.88±3.41s	0.46±0.1m	0.6±0.2m	2.6±0.3m	4.9±0.5m	9.1±0.5m
LS(+VO)	6.65±2.86s	0.51±0.1m	0.7±0.1m	3.0±0.3m	5.7±0.4m	13±0.5m

Table II
RUNTIME COMPARISON OF DECONFLICTED MCPP METHODS.

	<i>floor small</i>	<i>floor medium</i>	<i>floor large</i>	<i>ht_ch antry</i>	<i>AR02 05SR</i>	<i>Shang hai2</i>
CHA	0.10±0.1s	0.7±0.2s	6.9±2.7s	0.5±0.7m	1.6±1.5m	9.0±8.4m
	0.10±0.1s	0.7±0.2s	6.9±2.7s	0.9±2.7m	1.5±1.4m	9.0±8.4m
MLA	0.16±0.1s	4.8±3.9s	3.7±4.4m	25±23m	3.2±0.0m	/
	0.17±0.2s	5.2±4.5s	4.0±4.8m	34±21m	31±28m	/
ADA	0.10±0.1s	0.7±0.2s	7.3±3.2s	0.4±0.5m	1.4±1.4m	9.1±6.3m
	0.10±0.1s	0.8±0.2s	8.1±4.1s	0.6±0.6m	1.4±1.2m	7.0±5.2m
	<i>terrain small</i>	<i>terrain medium</i>	<i>terrain large</i>	<i>ost002d</i>	<i>AR07 01SR</i>	<i>New York1</i>
CHA	0.18±0.1s	1.3±0.6s	9.3±3.7s	1.0±0.6m	2.4±1.6m	5.8±7.0m
	0.18±0.1s	1.3±0.6s	9.4±3.8s	1.3±1.1m	2.4±1.6m	5.8±7.0m
MLA	0.27±0.2s	5.9±13s	5.4±12m	24±24m	/	/
	0.28±0.3s	5.8±12s	5.2±12m	38±21m	/	/
ADA	0.20±0.1s	1.3±0.5s	9.9±4.7s	0.9±0.5m	2.1±1.0m	4.4±3.0m
	0.20±0.2s	1.3±0.5s	10±4.6s	1.2±0.6m	2.1±0.8m	4.2±2.4m

D. Physical Robot Experiments

We constructed an MCPP instance of an indoor environment with a *Vicon*⁷ motion capture system for localization, as shown in Fig. 17. Our pipeline computes the conflict-free set of trajectories for the robots by first finding an MCPP solution via LS-MCPP and then deconflicting the MCPP solution via PBS and ADA. After offline planning, the trajectories are sent to the robots to execute the coverage task via a centralized scheduler implemented in *ROS*⁸. Fig. 17 demonstrates the execution procedure of the resulting trajectories using our planning pipeline. The detailed setup is described as follows. **Multi-Robot System:** We employ two Pepper robots, named *Cayenne* and *Cumin*, each with a square bounding box of 0.55 meters. The 2D grid graph represents the indoor environment as a 10×6 map with each grid unit sized at 0.65 meters to accommodate their bounding boxes. The initial root vertices of *Cayenne* and *Cumin* are the bottom-left and top-right cells, respectively. Both Peppers can have omnidirectional movement driven by velocity commands $u = (v_x, v_y, \omega)$. The linear velocities v_x and v_y are capped at 0.15 meters per second, while the angular velocity ω is restricted to 0.5 radians per second. To determine the non-uniform traversal

⁷<https://www.vicon.com>

⁸<https://www.ros.org>

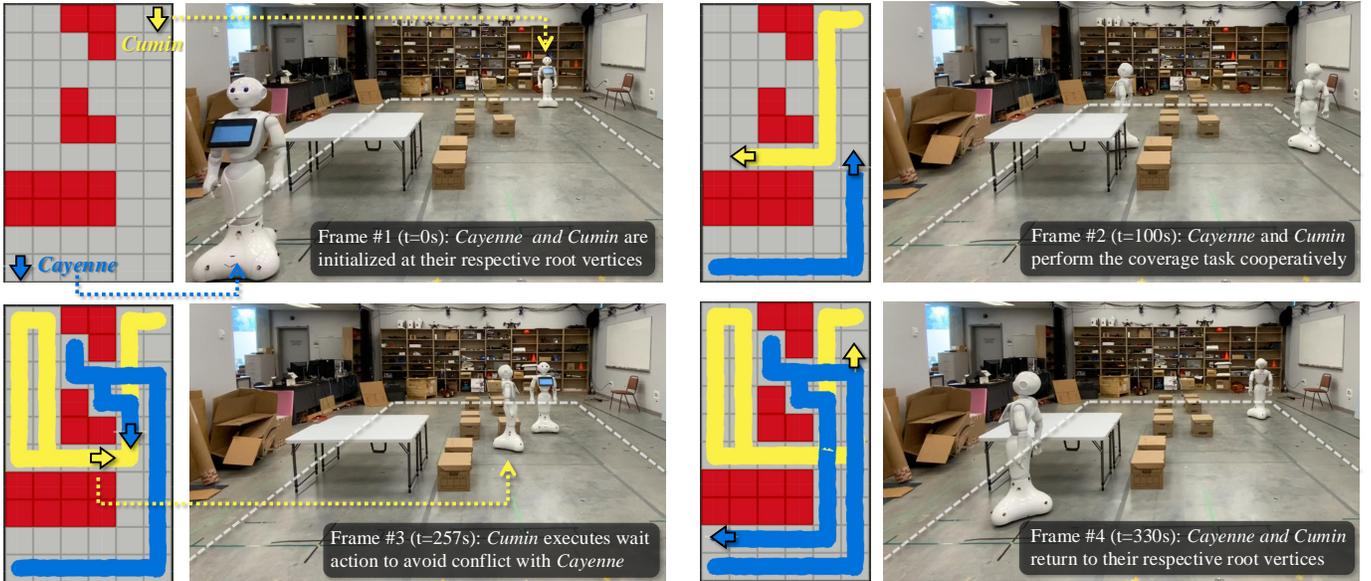


Figure 17. Demonstration of *Cayenne* and *Cumin* executing the coverage task cooperatively and without conflicts. Four frames from the execution are shown, with robot trajectories and the grid graph displayed on the left. Static obstacles and the trajectories of *Cayenne* and *Cumin* are highlighted in red, blue, and yellow, respectively. The entire planning pipeline was completed in 0.1 seconds, and the execution of the coverage task took approximately 6 minutes.

costs (i.e., the edge weights) for Peppers, we recorded the maximum and minimum time costs to traverse a single grid unit by performing random move actions across the graph. We then sampled weights randomly within this range and assigned them to the edges. Although Peppers can move omnidirectionally, we assume they always face the moving directions by integrating turn actions as described in Sec. VI-D. We set C as the maximum action time of 90° rotations for Peppers. In the final constructed graph, the time costs for move actions range from 7.2 to 10.8 seconds, while the time cost for a 90° turn is 6.0 seconds.

Centralized Scheduler: Given the set of conflict-free trajectories resulting from our planning pipeline, we generate an action sequence consisting of move, turn, and wait actions between consecutive states for each Pepper. A centralized scheduler is employed to sequentially send actions to the corresponding Pepper, maintaining a positional tolerance of 0.01 meters and a rotational tolerance of 0.02 radians. The velocity commands for each action are generated using a proportional controller operating at 20 Hz to align with the nominal action time cost. In practice, the move and turn action time costs are not precise for the Peppers, as we assume instantaneous velocity changes and do not account for their real-time dynamics or potential perturbations. This often leads to an overestimation of the actual action time costs, as the time for acceleration and deceleration is ignored, though underestimations can also occur occasionally. Therefore, we design the centralized scheduler to synchronize the discrepancies between the nominal time costs t and the actual time costs t' for the actions. Upon completing each action, each Pepper sends the time difference $\Delta t = t - t'$ back to the scheduler. If $\Delta t < 0$, the scheduler signals the Pepper to wait for a period of $-\Delta t$ to match the nominal time cost; otherwise, the scheduler instructs the other Pepper to add Δt to its waiting period after

completing its current action.

Remark 6. *Our centralized scheduler currently uses a basic approach to synchronize time differences between nominal and actual action time costs, which may increase the task completion time. More advanced techniques from MAPF execution research, such as constructing an action dependency graph [70], could address this issue. However, exploring these advanced methods is beyond the scope of this paper.*

VIII. CONCLUSIONS & FUTURE WORK

In this paper, we presented a comprehensive approach to MCP on 2D grid graphs, addressing key challenges including grids with incomplete quadrant coarsening, non-uniform traversal costs, inter-robot conflicts, and turning costs in large-scale, real-world scenarios. Our contributions began with the ESTC paradigm, designed to overcome limitations of standard STC-based paradigms by enabling path generation on any grid structure, even those with incomplete quadrant coarsening. We further enhanced ESTC with two local optimizations, improving both path quality and computational efficiency. Building on ESTC, we developed LS-MCPP, a local search framework that integrates ESTC to iteratively refine and optimize coverage paths, yielding high-quality solutions for MCP on arbitrary grid graphs. To address inter-robot conflicts, we introduced a new MAPF variant as a post-processing step for MCP. This approach marks the first integration of MAPF techniques into MCP, bridging two critical fields in multi-robot coordination. Our solver effectively resolves conflicts while preserving solution quality, demonstrating robust performance even in complex environments with up to 100 robots. Our experimental results confirm that the proposed pipeline is scalable, robust, and capable of delivering high-quality solutions, making it suitable for a wide range of real-world multi-robot applications across diverse environments.

This work lays a strong foundation for advancing multi-robot systems with implications for practical deployments in real-world applications. Future work will explore encoding action costs on grid graphs with more precise action dynamics, further accelerating the local search through parallelization and machine learning techniques, and integrating online replanning to handle real-time changes and execution deviations.

ACKNOWLEDGEMENT

This work was supported by the NSERC under grant number RGPIN2020-06540 and a CFI JELF award. We sincerely thank Chuxuan Zhang and Dr. Angelica Lim from SFU ROSIE lab for providing the two Pepper robots, Cayenne and Cumin, for our experiments.

REFERENCES

- [1] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [2] Z. Wang, X. Zhao, J. Zhang, N. Yang, P. Wang, J. Tang, J. Zhang, and L. Shi, "APF-CPP: An artificial potential field based multi-robot online coverage path planning approach," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 9199–9206, 2024.
- [3] R. Almadhoun, T. Taha, D. Gan, J. Dias, Y. Zweiri, and L. Seneviratne, "Coverage path planning with adaptive viewpoint sampling to construct 3d models of complex structures for the purpose of inspection," in *IROS*, 2018, pp. 7047–7054.
- [4] T. Chen, H. S. Ahn, W. Sun, J. Pan, Y. Liu, J. Cheng, and L. Xu, "Optimizing path planning for a single tracked combine harvester: A comprehensive approach to harvesting and unloading processes," *Computers and Electronics in Agriculture*, vol. 224, p. 109217, 2024.
- [5] S. Santra, K. Uno, G. Kudo, and K. Yoshida, "Risk-aware coverage path planning for lunar micro-rovers leveraging global and local environmental data," in *iSpaRo*. IEEE, 2024, pp. 42–47.
- [6] S. K. R. Sudha, D. Mishra, and I. A. Hameed, "A coverage path planning approach for environmental monitoring using an unmanned surface vehicle," *Ocean Engineering*, vol. 310, p. 118645, 2024.
- [7] J. Tang, Y. Gao, and T. L. Lam, "Learning to coordinate for a workstation multi-robot system in planar coverage tasks," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12315–12322, 2022.
- [8] C. Sun, J. Tang, and X. Zhang, "FT-MSTC: An efficient fault tolerance algorithm for multi-robot coverage path planning," in *RCAR*, 2021, pp. 107–112.
- [9] P. M. França, M. Gendreau, G. Laporte, and F. M. Müller, "The m-traveling salesman problem with minmax objective," *Transportation Science*, vol. 29, no. 3, pp. 267–275, 1995.
- [10] J. Carlsson, D. Ge, A. Subramaniam, A. Wu, and Y. Ye, "Solving min-max multi-depot vehicle routing problem," *Lectures on global optimization*, vol. 55, pp. 31–46, 2009.
- [11] N. Hazon and G. A. Kaminka, "Redundancy, efficiency and robustness in multi-robot coverage," in *ICRA*, 2005, pp. 735–741.
- [12] X. Zheng and S. Koenig, "Robot coverage of terrain with non-uniform traversability," in *IROS*, 2007, pp. 3757–3764.
- [13] X. Zheng, S. Koenig, D. Kempe, and S. Jain, "Multirobot forest coverage for weighted and unweighted terrain," *IEEE Transactions on Robotics*, vol. 26, no. 6, pp. 1018–1031, 2010.
- [14] J. Tang, C. Sun, and X. Zhang, "MSTC*: Multi-robot coverage path planning under physical constraint," in *ICRA*, 2021, pp. 2518–2524.
- [15] J. Tang and H. Ma, "Mixed integer programming for time-optimal multi-robot coverage path planning with efficient heuristics," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6491–6498, 2023.
- [16] Y. Gabriely and E. Rimon, "Spanning-tree based coverage of continuous areas by a mobile robot," *Annals of mathematics and artificial intelligence*, vol. 31, pp. 77–98, 2001.
- [17] R. Stern, N. Sturtevant, A. Felner, S. Koenig, H. Ma, T. Walker, J. Li, D. Atzmon, L. Cohen, T. Kumar *et al.*, "Multi-agent pathfinding: Definitions, variants, and benchmarks," in *SoCS*, 2019, pp. 151–158.
- [18] I. Vandermeulen, R. Groß, and A. Kolling, "Turn-minimizing multirobot coverage," in *ICRA*, 2019, pp. 1014–1020.
- [19] J. Lu, B. Zeng, J. Tang, T. L. Lam, and J. Wen, "TMSTC*: A path planning algorithm for minimizing turns in multi-robot coverage," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 5275–5282, 2023.
- [20] M. Ramesh, F. Imeson, B. Fidan, and S. L. Smith, "Optimal partitioning of non-convex environments for minimum turn coverage planning," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9731–9738, 2022.
- [21] J. Tang and H. Ma, "Large-scale multi-robot coverage path planning via local search," in *AAAI*, 2024, pp. 17567–17574.
- [22] C. K. Tomaszewski, "Constraint-based coverage path planning: A novel approach to achieving energy-efficient coverage," Ph.D. dissertation, Carnegie Mellon University, 2020.
- [23] A. C. Kapoutsis, S. A. Chatzichristofis, and E. B. Kosmatopoulos, "Darp: Divide areas algorithm for optimal multi-robot coverage path planning," *Journal of Intelligent & Robotic Systems*, vol. 86, pp. 663–680, 2017.
- [24] H. Azpúrua, G. M. Freitas, D. G. Macharet, and M. F. Campos, "Multi-robot coverage path planning using hexagonal segmentation for geophysical surveys," *Robotica*, vol. 36, no. 8, pp. 1144–1166, 2018.
- [25] L. Li, D. Shi, S. Jin, S. Yang, Y. Lian, and H. Liu, "Sp2e: Online spiral coverage with proactive prevention extremum for unknown environments," *Journal of Intelligent & Robotic Systems*, vol. 108, no. 2, p. 30, 2023.
- [26] G. Even, N. Garg, J. Könemann, R. Ravi, and A. Sinha, "Min-max tree covers of graphs," *Operations Research Letters*, vol. 32, no. 4, pp. 309–315, 2004.
- [27] H. Nagamochi and K. Okada, "Approximating the minmax rooted-tree cover in a tree," *Information Processing Letters*, vol. 104, no. 5, pp. 173–178, 2007.
- [28] I. Rekleitis, A. P. New, E. S. Rankin, and H. Choset, "Efficient boustrophedon multi-robot coverage: an algorithmic approach," *Annals of Mathematics and Artificial Intelligence*, vol. 52, pp. 109–142, 2008.
- [29] A. Xu, C. Viriyasuthee, and I. Rekleitis, "Efficient complete coverage of a known arbitrary environment with applications to aerial operations," *Autonomous Robots*, vol. 36, pp. 365–381, 2014.
- [30] N. Karapetyan, K. Benson, C. McKinney, P. Taslakian, and I. Rekleitis, "Efficient multi-robot coverage of a known environment," in *IROS*, 2017, pp. 1846–1852.
- [31] J.-C. Latombe and J.-C. Latombe, "Exact cell decomposition," *Robot Motion Planning*, pp. 200–247, 1991.
- [32] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, pp. 247–253, 2000.
- [33] E. U. Acar, H. Choset, A. A. Rizzi, P. N. Atkar, and D. Hull, "Morse decompositions for coverage tasks," *International Journal of Robotics Research*, vol. 21, no. 4, pp. 331–344, 2002.
- [34] S. Bochkarev and S. L. Smith, "On minimizing turns in robot coverage path planning," in *CASE*, 2016, pp. 1237–1242.
- [35] Y. Yang, H. T. Loh, J. Fuh, and Y. Wang, "Equidistant path generation for improving scanning efficiency in layered manufacturing," *Rapid Prototyping Journal*, vol. 8, no. 1, pp. 30–37, 2002.
- [36] F. Ren, Y. Sun, and D. Guo, "Combined reparameterization-based spiral toolpath generation for five-axis sculptured surface machining," *International Journal of Advanced Manufacturing Technology*, vol. 40, pp. 760–768, 2009.
- [37] I. Gibson, D. W. Rosen, B. Stucker, M. Khorasani, D. Rosen, B. Stucker, and M. Khorasani, *Additive manufacturing technologies*. Springer, 2021, vol. 17.
- [38] H. Zhao, F. Gu, Q.-X. Huang, J. Garcia, Y. Chen, C. Tu, B. Benes, H. Zhang, D. Cohen-Or, and B. Chen, "Connected fermat spirals for layered fabrication," *ACM Transactions on Graphics*, vol. 35, no. 4, pp. 1–10, 2016.
- [39] C. Wu, C. Dai, X. Gong, Y.-J. Liu, J. Wang, X. D. Gu, and C. C. Wang, "Energy-efficient coverage path planning for general terrain surfaces," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2584–2591, 2019.
- [40] J. Tang and H. Ma, "Multi-robot connected fermat spiral coverage," in *ICAPS*, 2024, pp. 579–587.
- [41] J. Banfi, N. Basilico, and F. Amigoni, "Intractability of time-optimal multirobot path planning on 2d grid graphs with holes," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 1941–1947, 2017.
- [42] H. Ma, "Graph-based multi-robot path finding and planning," *Current Robotics Reports*, vol. 3, no. 3, pp. 77–84, 2022.
- [43] G. Mouratidis, B. Nebel, and S. Koenig, "Fools rush in where angels fear to tread in multi-goal cbs," in *SoCS*, 2024, pp. 243–251.
- [44] H. Zhang, J. Chen, J. Li, B. Williams, and S. Koenig, "Multi-agent path finding for precedence-constrained goal sequences," in *AAMAS*, 2022.
- [45] P. Surynek, "Multi-goal multi-agent path finding via decoupled and integrated goal vertex ordering," in *AAAI*, 2021, pp. 12409–12417.

- [46] Z. Ren, S. Rathinam, and H. Choset, “Ms*: A new exact algorithm for multi-agent simultaneous multi-goal sequencing and path finding,” in *ICRA*, 2021, pp. 11 560–11 565.
- [47] M. Tang, Y. Li, H. Liu, Y. Chen, M. Liu, and L. Wang, “Mgcbs: An optimal and efficient algorithm for solving multi-goal multi-agent path finding problem,” in *IJCAI*, 2024, pp. 249–256.
- [48] H. Ma, J. Li, T. S. Kumar, and S. Koenig, “Lifelong multi-agent path finding for online pickup and delivery tasks,” in *AAMAS*, 2017, pp. 837–845.
- [49] M. Liu, H. Ma, J. Li, and S. Koenig, “Task and path planning for multi-agent pickup and delivery,” in *AAMAS*, 2019.
- [50] Q. Xu, J. Li, S. Koenig, and H. Ma, “Multi-goal multi-agent pickup and delivery,” in *IROS*, 2022, pp. 9964–9971.
- [51] Z. Ren, S. Rathinam, and H. Choset, “Cbss: A new approach for multiagent combinatorial path finding,” *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2669–2683, 2023.
- [52] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter, “Hamilton paths in grid graphs,” *SIAM Journal on Computing*, vol. 11, no. 4, pp. 676–686, 1982.
- [53] C. Umans and W. Lenhart, “Hamiltonian cycles in solid grid graphs,” in *FOCS*, 1997, pp. 496–505.
- [54] M. R. Garey and D. S. Johnson, “Computers and intractability: A guide to the theory of np-completeness,” 1979.
- [55] Y. Gabriely and E. Rimon, “Spiral-stc: An on-line coverage algorithm of grid environments by a mobile robot,” in *ICRA*, 2002, pp. 954–960.
- [56] J. B. Kruskal, “On the shortest spanning subtree of a graph and the traveling salesman problem,” *Proceedings of the American Mathematical Society*, vol. 7, no. 1, pp. 48–50, 1956.
- [57] F. Aurenhammer and R. Klein, “Voronoi diagrams,” *Handbook of Computational Geometry*, vol. 5, no. 10, pp. 201–290, 2000.
- [58] D. E. Goldberg, “Genetic algorithms in search,” *Optimization, Machine Learning*, 1989.
- [59] P. J. Van Laarhoven, E. H. Aarts, P. J. van Laarhoven, and E. H. Aarts, *Simulated annealing*. Springer, 1987.
- [60] H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, “Searching with consistent prioritization for multi-agent path finding,” in *AAAI*, 2019, pp. 7643–7650.
- [61] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, “Conflict-based search for optimal multi-agent pathfinding,” *Artificial intelligence*, vol. 219, pp. 40–66, 2015.
- [62] M. Phillips and M. Likhachev, “SIPP: Safe interval path planning for dynamic environments,” in *ICRA*, 2011, pp. 5628–5635.
- [63] H. Ma, W. Hönig, T. S. Kumar, N. Ayanian, and S. Koenig, “Lifelong path planning with kinematic constraints for multi-agent pickup and delivery,” in *AAAI*, 2019, pp. 7651–7658.
- [64] D. Silver, “Cooperative pathfinding,” in *AIIDE*, 2005, pp. 117–122.
- [65] F. Grenouilleau, W.-J. Van Hove, and J. N. Hooker, “A multi-label a* algorithm for multi-agent pathfinding,” in *ICAPS*, 2019, pp. 181–185.
- [66] J. Li, A. Tinka, S. Kiesel, J. W. Durham, T. S. Kumar, and S. Koenig, “Lifelong multi-agent path finding in large-scale warehouses,” in *AAAI*, 2021, pp. 11 272–11 281.
- [67] X. Zhong, J. Li, S. Koenig, and H. Ma, “Optimal and bounded-suboptimal multi-goal task assignment and path finding,” in *ICRA*, 2022, pp. 10 731–10 737.
- [68] N. Sturtevant, “Benchmarks for grid-based pathfinding,” *Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 2, pp. 144 – 148, 2012.
- [69] M. Čáp, J. Vokřínek, and A. Kleiner, “Complete decentralized method for on-line multi-robot trajectory planning in well-formed infrastructures,” in *ICAPS*, 2015, pp. 324–332.
- [70] W. Hönig, S. Kiesel, A. Tinka, J. W. Durham, and N. Ayanian, “Persistent and robust execution of MAPF schedules in warehouses,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1125–1131, 2019.



Jingtao Tang received the B.E. and M.E. degrees in software engineering from East China Normal University, Shanghai, China, in 2018 and 2021, respectively.

He is currently working toward the Ph.D. degree at the School of Computing Science, Simon Fraser University, Burnaby, Canada. His research interests include heuristic search, combinatorial optimization, and automated planning for multi-robot systems.



Zining Mao received B.E. in computer science from New York University Shanghai, Shanghai, China, in 2021.

He is currently working towards the MSc. Thesis degree at the School of Computing Science, Simon Fraser University, Burnaby, Canada. His research interests include heuristic search, reinforcement learning and automated planning for multi-robot systems.



Hang Ma is an Assistant Professor at the School of Computing Science, Simon Fraser University, Burnaby, Canada. He received the Ph.D. degree in computer science from the University of Southern California, Los Angeles, USA.

His research interests lie at the intersection of artificial intelligence, robotics, and machine learning. Specifically, he focuses on automated planning and multi-agent and multi-robot systems..