

Large-Scale Multi-Robot Coverage Path Planning via Local Search*

Jingtao Tang, Hang Ma

Simon Fraser University
{jingtao.tang, hangma}@sfu.ca

Abstract

We study graph-based Multi-Robot Coverage Path Planning (MCPP) that aims to compute coverage paths for multiple robots to cover all vertices of a given 2D grid terrain graph G . Existing graph-based MCPP algorithms first compute a tree cover on G —a forest of multiple trees that cover all vertices—and then employ the Spanning Tree Coverage (STC) paradigm to generate coverage paths on the decomposed graph D of the terrain graph G by circumnavigating the edges of the computed trees, aiming to optimize the makespan (i.e., the maximum coverage path cost among all robots). In this paper, we take a different approach by exploring how to systematically search for good coverage paths directly on D . We introduce a new algorithmic framework, called LS-MCPP, which leverages a local search to operate directly on D . We propose a novel standalone paradigm, Extended-STC (ESTC), that extends STC to achieve complete coverage for MCPP on any decomposed graphs, even those resulting from incomplete terrain graphs. Furthermore, we demonstrate how to integrate ESTC with three novel types of neighborhood operators into our framework to effectively guide its search process. Our extensive experiments demonstrate the effectiveness of LS-MCPP, consistently improving the initial solution returned by two state-of-the-art baseline algorithms that compute suboptimal tree covers on G , with a notable reduction in makespan by up to 35.7% and 30.3%, respectively. Moreover, LS-MCPP consistently matches or surpasses the results of optimal tree cover computation, achieving these outcomes with orders of magnitude faster runtime, thereby showcasing its significant benefits for large-scale real-world coverage tasks.

Introduction

Coverage path planning (CPP) is a fundamental problem (Galceran and Carreras 2013) in robotics, which aims to find a path for a robot to completely cover a terrain of interest, such as indoor floors (Bormann et al. 2018) and outdoor fields (Torres et al. 2016). 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 completely cover the given terrain. With improved task efficiency and system robustness, MCPP has facilitated diverse



Figure 1: Graph-based CPP and MCPP: Gray squares, black circles, and black stars represent terrain graph vertices, decomposed graph vertices, and initial vertices of robots, respectively; Solid lines and dashed lines represent coverage paths and spanning edges, respectively. (a) Terrain graph with uniform edge weights. (b) The single-robot coverage path generated by STC. (c)(d) Suboptimal and optimal 2-robot coverage paths with makespans 2 and 1.5, respectively.

real-world applications, including environmental monitoring (Collins et al. 2021) and search and rescue (Song et al. 2022). A fundamental challenge of MCPP lies in generating cost-balancing coverage paths to optimize task efficiency, commonly quantified by the *makespan*, which is the maximum path cost of all robots. This challenge is further compounded when dealing with large-scale applications where the number of robots and the size of the terrain increase.

In this paper, we follow existing graph-based MCPP algorithms (Zheng et al. 2010; Li et al. 2023) that represent the terrain to be covered as a 4-connected 2D grid graph G , where each edge connects horizontally or vertically adjacent vertices. The robots are required to start at and return to their respective initial vertices, as in the *cover and return* setting (Zheng and Koenig 2007). The foundation of these graph-based MCPP algorithms lies in the Spanning Tree Coverage (STC) paradigm (Gabriely and Rimon 2001, 2002), initially developed for (single-robot) CPP. STC operates on the terrain graph G but finds a coverage path with minimal makespan on the decomposed graph D derived from G . The decomposed graph D is also a 4-connected 2D grid graph, resulting from decomposing each vertex of G into four decomposed vertices. Fig. 1 shows the terrain graph G and its corresponding decomposed graph D of an example terrain to be covered, where STC generates a single-robot coverage path on D by circumnavigating (i.e., always moving along the right side of the spanning edges) the minimum spanning tree of G .

Like STC, existing graph-based MCPP algorithms operate on the given terrain graph G exclusively to build a tree cover—a forest of multiple trees, each rooted at the initial

*Code: <https://github.com/reso1/LS-MCPP>

vertex of a robot, that jointly cover all vertices of G . The coverage path for each robot is then obtained by circumnavigating its corresponding tree. In essence, these algorithms reduce MCPP to the NP-hard Min-Max Rooted Tree Cover problem (Even et al. 2004; Nagamochi and Okada 2007) on G that aims to optimize the weight of the largest-weighted tree in the tree cover since it determines the makespan of the resulting coverage paths on D . However, operating on the terrain graph G exclusively has two disadvantages. Firstly, it does not work for an incomplete terrain graph G where some of the four decomposed vertices of a vertex might be blocked and thus absent in the decomposed graph D . Secondly, an optimal tree cover on G does not necessarily result in an optimal MCPP solution (as illustrated in Fig. 1-(c) and (d)), which yields an asymptotic suboptimality ratio of four for makespan in the worst case (Zheng et al. 2010), since circumnavigating the trees in a tree cover explores only a portion of the solution space that encompasses all possible sets of coverage paths on the decomposed graph D .

Therefore, we take a different route to explore how to systematically search for good coverage paths directly on the decomposed graph. Our algorithmic contribution is through the following key aspects: (1) We propose a novel standalone algorithmic paradigm called Extended-STC (ESTC), an extension of STC, to address coverage planning problems on both complete and incomplete terrain graphs by directly operating on decomposed graphs. Importantly, we demonstrate that ESTC guarantees complete coverage for both single- and multi-robot settings, rendering it an efficient and versatile solution to coverage path planning. (2) We propose three types of specialized neighborhood operators to facilitate an effective local search process by identifying cost-efficient subgraphs of the decomposed graph that are used to generate coverage paths for the robots. The strategic integration of these operators enhances the efficiency of the local search in exploring the solution space. (3) We demonstrate how to combine these neighborhood operators with iterative calls to the ESTC paradigm to establish our proposed LS-MCPP framework for solving MCPP. To validate the effectiveness of LS-MCPP, we conduct extensive experiments, comparing it with three state-of-the-art baseline graph-based MCPP algorithms that operate on complete terrain graphs only. The results show that LS-MCPP achieves makespans that are smaller by up to 35.7% and 30.3% than two of the baseline algorithms, respectively, which compute suboptimal tree covers on the terrain graph. Additionally, LS-MCPP consistently achieves makespans comparable to or better than those achieved by the remaining baseline algorithm, which employs mixed integer programming (MIP) to compute optimal tree covers on the terrain graph. While the baseline algorithm takes more than tens of hours to complete, LS-MCPP accomplishes the same task in just a matter of minutes, showcasing its efficiency and practicality for large-scale real-world coverage problems.

Related Work on MCPP

Existing MCPP algorithms first partition the given terrain into multiple regions and then use a single-robot CPP algorithm to generate a coverage path on each region for

a robot. Without loss of generality, we categorize existing MCPP algorithms into decomposition-based and graph-based. Decomposition-based algorithms first partition the terrain using geometric critical points found on the polylines of the terrain to be covered (Karapetyan et al. 2017; Vandermeulen, Groß, and Kolling 2019) and then generate zigzag paths covering each partitioned region. Despite their simplicity in relatively open environments, they are unsuitable for obstacle-rich terrains (e.g., mazes) due to their reliance on geometric partitioning and cannot consider non-uniform traversal costs for weighted terrains.

Our paper focuses on graph-based MCPP algorithms (Zheng et al. 2010; Li et al. 2023) that are more versatile as they model the terrain to be covered as a graph and account for non-uniform traversal costs by assigning weights to the edges or vertices. Solving MCPP on a graph for makespan minimization has been proved to be NP-hard (Zheng et al. 2010). Existing approaches include polynomial-time approximation algorithms (Hazzon and Kaminka 2005; Tang, Sun, and Zhang 2021), numerical optimization (Kapoutsis, Chatzichristofis, and Kosmatopoulos 2017), and an MIP formulation (Tang and Ma 2023). A closely related problem is the multi-traveling salesman problem (mTSP) with a minmax objective (França et al. 1995), which aims to find optimal routes for multiple salesmen to visit all given cities. However, existing scalable mTSP algorithms are designed for Euclidean spaces (He and Hao 2023). We are unaware of any mTSP algorithms directly applicable to the graph-based MCPP problem.

Problem Formulation

In an MCPP instance, we are given a connected 2D grid terrain graph $G = (V_g, E_g)$ and its corresponding decomposed graph $D = (V_d, E_d)$, where each terrain vertex in V_g is decomposed into four small adjacent vertices in V_d as shown in Fig. 1-(a). For clarity, we let $\delta_v \in V_g$ denote the corresponding terrain vertex of a decomposed $v \in V_d$ and $\varepsilon = (\delta_u, \delta_v) \in E_g$ denote the edge connecting δ_u and δ_v . Each $\varepsilon \in E_g$ has a weight of w_ε . Note that the definition of weighted edges is different from the literature that considers weighted vertices (Zheng et al. 2010), yet it is rather convenient to convert weighted edges to weighted vertices by evenly distributing each edge weight to its two endpoint vertices. In that case, we have $w_{\delta_v} = \sum_{\varepsilon \sim \delta_v} w_\varepsilon$ for each $\delta_v \in V_g$, where $\varepsilon \sim \delta_v$ represents that δ_v is an endpoint of ε . Each edge $e = (u, v) \in E_d$ of D is thereby defined to have a weight $w_e = \frac{1}{4} \cdot \frac{1}{2} \cdot [\sum_{\varepsilon \sim \delta_u} w_\varepsilon + \sum_{\varepsilon \sim \delta_v} w_\varepsilon]$ by evenly splitting the weight w_δ of each $\delta \in V_g$ into $w_\delta/4$ for each of its four decomposed vertices (Tang and Ma 2023).

For any path $\pi = (e_1, e_2, \dots, e_{|\pi|})$ on D where each $e_i \in E_d$, we define its cost as $c(\pi) = \sum_{i=1}^{|\pi|} w_{e_i}$. Given a set $I = \{1, 2, \dots, k\}$ robots with a set $R = \{r_i\}_{i \in I} \subseteq V_d$ of initial root vertices, the graph-based MCPP problem is to find a set $\Pi = \{\pi_i\}_{i \in I}$ of k paths such that each $v \in V_d$ is visited by at least one coverage path for complete coverage and each π_i starts and ends at r_i . The solution quality is measured by the makespan $\tau = \max\{c(\pi_1), c(\pi_2), \dots, c(\pi_k)\}$.

Notation: We use the tuple (G, D, R) to denote an MCPP




Figure 2: A CPP instance with a root vertex (black star) and incomplete terrain vertices (yellow cells), where STC has no complete coverage. Full-STC is sub-optimal with undesirable routes (red circled area), and ESTC is optimal based on the augmented terrain graph G' .

instance. For each robot $i \in I$, we differentiate its corresponding subgraph $D_i = (V_{d,i}, E_{d,i})$ of D and its corresponding subgraph $G_i = (V_{g,i}, E_{g,i})$ of G .

Extended-STC

In this section, we present the new Extended-STC (ESTC) paradigm to address coverage problems on incomplete terrain graphs. We define that $\delta_v \in V_g$ is *complete* if all its four decomposed vertices are present in V_d and *incomplete* otherwise. A terrain graph G is *complete* if all its vertices are complete and *incomplete* otherwise. STC works only for complete terrain graphs due to its reliance on circumnavigating spanning edges of G . As in Fig. 1-(c)(d), this restriction of STC also prevents STC-based MCPP algorithms (Tang and Ma 2023) from finding good solutions even on complete terrain graphs. In this paper, we focus on covering a connected but possibly incomplete decomposed graph, as covering an unconnected decomposed graph only additionally requires a robot in each of its connected components.

In order to generate coverage paths through incomplete terrain vertices, our proposed ESTC paradigm cleverly integrates a path-deformation procedure of Full-STC (Gabieli and Rimon 2002) into the offline computation of STC. Full-STC is specifically designed for online (single-robot) CPP on an unknown terrain graph. It can be applied to incomplete terrain graphs but is suboptimal for CPP instances on known terrain graphs since it does not consider different path costs resulting from different spanning trees. Fig. 2 shows an example of the suboptimality of Full-STC.

Like Full-STC, ESTC operates on the augmented terrain graph G' where some edges in G are removed in G' to reflect the connectivity between its terrain vertices. Specifically, edge $\varepsilon = (\delta_u, \delta_v) \in E_g$ is removed in G' if each decomposed vertex u of terrain vertex δ_u is nonadjacent to each decomposed vertex v of terrain vertex δ_v . A special case is that each incomplete terrain vertex $\delta \in V_g$ with only two diagonally opposite decomposed vertices present is replaced with two nonadjacent terrain vertices in G' since its two decomposed vertices are nonadjacent. Unlike Full-STC, ESTC considers non-uniform edge weights. An edge $\varepsilon = (\delta_u, \delta_v)$ has the same weight as in G if both δ_u and δ_v are complete. Otherwise, ε has a manipulated edge weight of

$$w_{\max} \cdot \frac{1}{2} \cdot (\sum_{\varepsilon \sim \delta_u} w_\varepsilon + \sum_{\varepsilon \sim \delta_v} w_\varepsilon) \quad (1)$$

where w_{\max} is the maximal edge weights of E_g to prioritize using edges connecting complete terrain vertices. Fig. 2

shows an example of the augmented terrain graph G' .

By construction, G' is connected for a connected D . ESTC then builds a minimum spanning tree (MST) on G' instead of G and uses the same path-deformation rule defined as in Full-STC (Gabieli and Rimon 2002) to route through incomplete vertices. In essence, this rule reroutes the coverage path on the left side of a spanning edge whenever moving along the right side is impossible due to non-adjacency between decomposed vertices. The manipulated edge weight in Eqn. (1) prioritizes making each incomplete vertex a leaf of the MST, which prevents some decomposed vertices from repeatedly covering due to rerouting in path deformation and can thus improve the CPP solution quality (see Fig. 2). ESTC is guaranteed to generate a path that covers all vertices of a connected decomposed graph D based on the argument used in the proof of Lemma 3.1 in (Gabieli and Rimon 2002) since the argument is not affected by the addition of edge weights and prioritization of edges.

Replacing STC with ESTC for MCPP: Recall that for an MCPP instance (G, D, R) , STC-based MCPP algorithms first build k trees of G and then apply STC to generate the coverage paths. Therefore, they are restricted to generating each coverage path only on terrain graphs with complete vertices, which is suboptimal for most cases. We propose to extend ESTC to MCPP as follows. First, we generate a set $\{D_i = (V_{d,i}, E_{d,i})\}_{i \in I}$ of k connected subgraphs of D . Each D_i induces a subgraph G_i of G that is possibly incomplete. Then, we solve each sub-CPP instance (G_i, D_i, r_i) as described above by constructing an MST on the augmented G'_i to generate the coverage path for each robot (see Fig. 1-(d) for an example). The following theorem shows a sufficient condition for complete coverage.

Theorem 1. *If D_i is connected and $\bigcup_{i \in I} V_{d,i} = V_d$ for all i , ESTC is guaranteed to achieve complete coverage.*

Proof. Given that D_i is connected, ESTC is guaranteed to generate a path that covers all its vertices, analogous to the argument used in the proof of Lemma 3.1 in (Gabieli and Rimon 2002). Since $\bigcup_{i \in I} V_{d,i} = V_d$, the coverage paths generated for all D_i jointly cover all vertices of D . \square

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 D_i for the robots, each being connected, with the property that the union of all their vertices is the set of vertices D to ensure complete coverage.

Boundary Editing Operators

In this section, we introduce three types of boundary editing operators: grow operators, deduplicate operators, and exchange operators. These operators are designed to modify the boundaries of each subgraph D_i , involving the addition or removal of vertices. We define the *duplication* set $V^+ = \{v \in V_d \mid n_v > 1\}$ as the set of decomposed vertices covered by more than one robot, where $n_v = \sum_{i \in I} |\{x \in V_{d,i} \mid x = v\}|$ counts the occurrences of vertex $v \in V_d$ across all subgraphs. A boundary editing operator alters the set $\{D_i\}_{i \in I}$ of subgraphs but ensures that its property of complete coverage (i.e., $\bigcup_{i \in I} V_{d,i} = V_d$) and subgraph connectivity remains invariant. The operator aims to




Figure 3: An example of grow operator $o_g(i, e)$ and the coverage paths after it is executed. Blue circles and black circles represent vertices of the boundary vertex set B_i and vertices of $V_{d,i}$, respectively. (a) An invalid grow operator without any parallel edge in D_i . (b)(c) Two valid grow operators.

introduce only necessary duplication, which holds the potential to improve the resulting MCPP solution on the updated set of subgraphs with ESTC.

To allow for an efficient transition between different sets of subgraphs, we restrict our operator design to be edge-wise on any subgraph; that is, each operator involves the addition or removal of only one edge for each subgraph. This design is based on the observation that manipulating individual vertices can sometimes introduce unnecessary duplication to the coverage paths when ESTC is applied to the resulting subgraphs, as is exemplified by undesirable routes from a vertex $u \in V_{d,i}$ to its neighbor $v \in V_{d,i}$ and then back to u . In contrast, edge-wise operators modify two vertices simultaneously during each operation, eliminating such concerns and ensuring a more effective adjustment. Remarkably, edge-wise operators can achieve the same outcome as operators involving multiple edges. In addition, our operators add or remove only edges whose two endpoints are decomposed from the same terrain vertex since manipulating other edges would introduce unnecessary duplication. Formally, we denote the set of such edges of any D_i as

$$F_i = \{(u, v) \in E_{d,i} \mid \delta_u = \delta_v\}. \quad (2)$$

Grow Operators: A grow operator serves to expand a subgraph to cover an additional vertex that is already covered by other subgraphs. Let B_i denote the set of vertices that are not part of D_i but adjacent to a vertex of D_i , i.e., $B_i = \{v \in V_{d,i} \mid \exists (u, v) \in E_{d,i}, u \in V_{d,i}\}$. Formally, a *valid grow operator* $o_g(i, e)$ for D_i adds edge $e = (u, v) \in F_i$ with $u, v \in B_i$ and all relevant connecting edges into D_i , where $\exists (p, q) \in E_{d,i}$ such that $(u, p), (v, q) \in E_{d,i}$. In essence, a valid grow operator can only add an edge $e = (u, v)$ iff there exists a parallel edge (p, q) in D_i (see Fig. 3-(b)(c) for examples). This design choice avoids introducing undesirable routing to the coverage path, as one case exemplified in Fig. 3-(a). After a grow operator $o_g(i, e)$ for $e = (u, v)$ is executed, vertices u, v are added to V^+ since they are also covered by other subgraphs.

Deduplicate operators: A deduplicate operator serves to remove unnecessary duplication from a subgraph. For an edge $e = (u, v) \in F_i$, 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. 4-(a) demonstrates one example that can be rotated to various symmetric cases (with only δ_e^t potentially having its decomposed vertices adjacent to both u and v). Formally, a *valid deduplicate operator* $o_d(i, e)$ for D_i removes edge $e = (u, v) \in F_i$ with $u, v \in V^+ \cap V_{d,i}$ and all relevant connecting edges from D_i such that D_i remains




Figure 4: An example of deduplicate operator $o_d(i, e)$. Red circles and black circles represent vertices in V^+ and not in V^+ , respectively. (a) Four neighbors of δ_u . (b) An invalid operator violating condition (3). (c)(d) Two valid operators.

connected. This removal is constrained by three conditions if δ_u is complete: (1) δ_e^t is not in $V_{g,i}$; (2) all decomposed vertices of δ_e^b are in $V_{d,i}$; and (3) if $\delta = \delta_e^l, \delta_e^r$ is in $V_{g,i}$, then all decomposed vertices of both δ and its (only) common neighboring vertex with δ_e^b are in $V_{d,i}$. This design choice avoids introducing new duplication (such as the one shown in Fig. 4-(b)) to the coverage path when ESTC is applied to D_i . Otherwise, when δ_u is incomplete, the removal is not constrained by specific conditions. This design choice is based on the observation that the edge connecting an incomplete δ_u has low priority in the MST construction of ESTC. Consequently, δ_u is either positioned on or proximate to the periphery of the MST, and removing $e = (u, v)$ often does not introduce unnecessary or any duplication. We illustrate two valid deduplicate operators for an incomplete and complete δ_u in Fig. 4-(c) and Fig. 4-(d), respectively. After a deduplicate operator $o_d(i, e)$ for $e = (u, v)$ is executed, vertices u, v are removed from V^+ if they are not covered more than once by other subgraphs.

Exchange Operators: 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 the local search and speeding up the convergence of the local search. Formally, a *valid exchange operator* $o_e(i, j, e)$ adds $e = (u, v) \in F_i$ and all relevant connecting edges into D_i and removes them from D_j , where $o_g(i, e)$ and $o_d(j, e)$ are both valid, except that the condition $u, v \in V^+ \cap V_{d,j}$ for $o_d(j, e)$ is now replaced by $u, v \in V_{d,j}$ since u, v are not necessarily in V^+ . After an exchange operator $o_e(i, j, e)$ is executed, V^+ remains unchanged.

Properties: After any valid operator is executed, (1) D_i remains connected since only deduplicate operators can remove edges but a valid deduplicate guarantees that D_i remains connected, and (2) the property of the union of the vertices of all subgraphs being the vertex set D is not affected since no vertex coverage is lost.

LS-MCPP

In this section, we present LS-MCPP, a novel local search framework for effective MCPP solutions. LS-MCPP integrates the boundary editing operators, directly operating on the decomposed graph D to explore good coverage paths by adjusting subgraphs D_i . We categorize D_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} of all subgraphs; otherwise, it is a *heavy subgraph*. LS-MCPP em-

ploy grow operators on light subgraphs and deduplicate operators on heavy subgraphs to eliminate redundancy and exchange operators to balance path costs between subgraphs. The goal is to achieve a cost-equilibrium solution with minimal duplication. At a higher level, LS-MCPP employs a hierarchical sampling approach for efficient exploration of the constructed neighborhood in each iteration of its local search. It selects an operator pool using *roulette wheel* selection (Goldberg 1989) from three pools, each containing operators of the same type. Then, new heuristics are introduced to draw operators from the selected pool. LS-MCPP also calls a deduplication function periodically to exploit the current neighborhood and achieve a low-makespan solution.

Pseudocode (Alg. 1): LS-MCPP takes an initial MCPP solution as input. It starts with initializing three pools O_g, O_d, O_e of grow operators, deduplicate operators, and exchange operators, respectively (lines 1-3), and initializing the temperature scalar t and the pool weight vector \mathbf{p} (line 4). The scalar t originated from *simulated annealing* (Van Laarhoven et al. 1987) is widely used to adaptively accept non-improving operators in search-based algorithms. The vector \mathbf{p} represents the pool weight for each of the three aforementioned pools and is thereby used for roulette wheel selection to select a pool O from \mathcal{O} in each iteration (line 6), 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 pool weight $\mathbf{p}[O]$ is updated with a weight decay factor $\gamma \in [0, 1]$ in line 7. 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 9), where \mathbf{h} is the vector of heuristic function h evaluated on all operators in O (line 8), which we will describe in the next paragraph. Then, LS-MCPP evaluates operator o on its relating subgraph(s) and calculates the makespan increment $\Delta\tau$ after applying ESTC on the relating subgraph(s) (line 10) to determine whether to accept the new solution (i.e., applying o and updating Π). LS-MCPP accepts the new solution if it has a smaller makespan, and, otherwise, with a probability $\exp(-\Delta\tau/t) \in [0, 1]$ (lines 11-14). Such an adaptive acceptance criterion is dynamically adjusted by $\Delta\tau$ divided by temperature t that is scaled down by the decay factor $\alpha \in [0, 1]$ for every iteration (line 18) so that it gets harder to accept a non-improving operator over iterations, which is the simulated annealing strategy to skip the local minimum. LS-MCPP calls a function FORCEDDEDUPLICATION() (to be described later) every S iteration or when the current solution has the smallest makespan (line 16) and records the makespan-minimal solution in Π^* (line 17). If operator o is applied thereby updating its relating subgraph(s), then the pools of \mathcal{O} are updated with all the operator(s) related to the modified subgraph(s) (line 18). LS-MCPP terminates and returns the improved solution Π^* when it reaches the maximum number of iterations M (line 19).

Operator Sampling with Heuristics: LS-MCPP should carefully determine the probability of selecting each operator to better explore the neighborhood when sampling an operator from the selected pool (line 9 of Alg. 1). We propose three heuristic functions tailored to the three types of operators to evaluate their potential in guiding the neighbor-

Algorithm 1: LS-MCPP

Input: MCPP instance (G, D, R) , initial solution Π
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 $O_g \leftarrow \bigcup_{i \in I} \{o_g(i, e) \mid c(\pi_i) \leq \bar{c}\}$
- 2 $O_d \leftarrow \bigcup_{i \in I} \{o_d(i, e) \mid c(\pi_i) > \bar{c}\}$
- 3 $O_e \leftarrow \bigcup_{i, j \in I} \{o_e(i, j, e) \mid c(\pi_i) \leq \bar{c}\}$
- 4 $\Pi^* \leftarrow \Pi, t \leftarrow 1, \mathbf{p} \leftarrow [1, 1, 1]^T, \mathcal{O} = \{O_g, O_d, O_e\}$
- 5 **for** $i \leftarrow 1, 2, \dots, M$ **do**
- 6 $O \sim \mathcal{O}$ w/ probability function $P(O) = \sigma(\mathbf{p})$
- 7 $\mathbf{p}[O] \leftarrow (1 - \gamma) \cdot \mathbf{p}[O] + \gamma \cdot \max(-\Delta\tau, 0)$
- 8 $\mathbf{h} \leftarrow$ the heuristic value vector $[h(o)]_{o \in O}^T$
- 9 $o \sim O$ w/ probability function $P(O) = \sigma(\mathbf{h})$
- 10 $\Delta\tau \leftarrow$ makespan increment after using ESTC on updated subgraph(s) with o
- 11 **if** $\Delta\tau < 0$ **then**
- 12 | Apply operator o and update Π
- 13 **else**
- 14 | Apply o and update Π w/ probability of $\exp(-\frac{\Delta\tau}{t})$
- 15 **if** $i \% S = 0 \vee \Delta\tau < 0$ **then**
- 16 | FORCEDDEDUPLICATION(Π, O_d)
- 17 Assign Π to Π^* if Π is better than Π^*
- 18 Update each pool of \mathcal{O} and assign t to $\alpha \cdot t$
- 19 **return** Improved MCPP solution Π^*
- 20 **Function** FORCEDDEDUPLICATION(Π, O_d):
- 21 **for** $i \in \arg \text{sort}(\{c(\pi_i)\}_{i \in I})$ **do**
- 22 | **while** $(u, v) \leftarrow$ any U-turn in π_i **do**
- 23 | Remove u, v from D_i and update π_i
- 24 **for** $i \in \arg \text{sort}(\{c(\pi_i)\}_{i \in I})$ **do**
- 25 | **while** $\{o_d(j, e) \in O_d \mid i = j\} \neq \emptyset$ **do**
- 26 | $o \leftarrow \arg \min_{o \in \{o_d(j, e) \in O_d \mid i = j\}} ([h(o)]_{o \in \{o_d(j, e) \in O_d \mid i = j\}}^T)$
- 27 | Apply operator o , update Π and O_d

hood 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 small coverage path costs and (2) prioritizing covering vertices with less duplication. Formally, the heuristic value for a grow operator $o_g(i, e)$ with edge $e = (u, v)$ is defined as $h(o_g) = -k \cdot c(\pi_i) - (n_u + n_v)/2$. Note that $k \geq n_v$ holds for any $v \in V_d$, thereby prioritizing consideration (1) (i.e., $c(\pi_i)$) over consideration (2) (i.e., $(n_u + n_v)/2$). The heuristic function for deduplicate operators $o_d(i, e)$ is defined as the exact opposite of the one for grow operators: $h(o_d) = k \cdot c(\pi_i) + (n_u + n_v)/2$. The heuristic function for exchange operators $o_e(i, j, e)$ is defined as $h(o_e) = c(\pi_j) - c(\pi_i)$ to prioritize two subgraphs with a big difference in their coverage path costs.

Forced Deduplication: A key consideration towards the success of LS-MCPP is to restrict duplication to only those necessarily needed over the neighborhood construction; for instance, a narrow path connecting two separate regions of the terrain in the *flr-s* instance (Tang and Ma 2023). For




Figure 5: FORCEDDEDUPLICATION() executed on a coverage path π_i . Each red circle represents a duplication of $V_{d,i}$. In each frame, a U-turn $(u, v) \in \pi_i$ is found and u, v are removed from π_i until no U-turn exists.

Instance	Grid spec.	V_g	E_g	k	weighted
<i>AR0701SR</i>	107×117	4,860	8,581	20	✓
<i>Shanghai2</i>	128×128	11,793	22,311	25	×
<i>NewYork1</i>	128×128	11,892	21,954	32	✓

Table 1: MCPP instance specification.

this purpose, FORCEDDEDUPLICATION() (line 20 of Alg. 1) aims to deduplicate all possible duplication in two folds. The first part (lines 21-23) iterates through each coverage path $\pi_i \in \Pi$ in a cost-decreasing order to remove any U-turn in it (lines 22-23). A U-turn is an edge $(u, v) \in \pi_i$ with $u, v \in V^+ \cap V_{d,i}$, satisfying $\exists (p, q) \in \pi_i$ such that $(u, p), (v, q) \in \pi_i$. (so that $p \rightarrow u \rightarrow v \rightarrow q$ forms a ‘‘U-turn’’.) The second part (lines 24-27) tries to recursively apply all valid deduplicate operators in a descending order of the coverage path costs (line 24) and an ascending order of the heuristic values (line 26). Fig. 5 demonstrates FORCEDDEDUPLICATION() on a coverage path. It is worth noting that the first part can remove edges that do not constitute valid deduplicate operators, and some valid deduplicate operators can remove edges that are not U-turns (depending on the spanning tree generated by ESTC). The complementary relation between the two parts yields a better performance of the proposed LS-MCPP.

Empirical Evaluation

This section describes our experimental results on an *Intel® Xeon® Gold 5218 2.30 GHz* Linux server with 300 GB memory. All LS-MCPP experiments were conducted across 12 repetitions with seeds spanning from 0 to 11. **Instances:** We use MCPP instances from (Tang and Ma 2023), where their numbers of graph vertices, graph edges, and robots range from 46 to 824, 60 to 1495, and 4 to 12, respectively. As shown in Fig. 6 and Tab. 1, we design three additional larger instances whose terrain graphs are adopted from popular 2D pathfinding benchmarks (Sturtevant 2012). All the above instances have complete terrain graphs. **Baselines:** We use MFC (Zheng et al. 2010), MSTC* (Tang, Sun, and Zhang 2021), and MIP/MIP(H) (Tang and Ma 2023) as baseline algorithms that all use STC to transform a tree cover on G to an MCPP solution. Both MFC and MSTC* compute a suboptimal tree cover. MIP computes an optimal tree cover with a MIP model. MIP(H) computes a suboptimal tree cover with a reduced-size MIP model but is more efficient than MIP. We also design a new baseline VOR that first computes the Voronoi diagram (i.e., k subgraphs) of G based on the initial vertices of robots and then uses ESTC to generate coverage paths on the subgraphs. **Parameters:** For all instances solved with LS-MCPP, the pool weight decay




Figure 6: Three very large-scale MCPP instances adopted from 2D pathfinding benchmarks (Sturtevant 2012). Red circles represent the initial vertices of robots.




Figure 7: Single-robot coverage path costs (y-axis) with the percentages of incomplete terrain vertices (x-axis). Left and right sub-figures correspond to *flr-l* and *trn1-l*, respectively.

factor γ is set to 0.01, and the temperature decay factor α is set to $\exp \frac{\log 0.2}{M}$ to decrease the temperature t from 1 to 0.2. For instances from (Tang and Ma 2023), we set the max iteration M to $3e3$ and the forced deduplication step S to $1e2$. For the three additional larger instances in Tab. 1, we set M to $1.5e4$ and S to $5e2$.

Ablation Study

We use two instances *flr-l* and *trn1-l* to conduct an ablation study on different components of LS-MCPP.

ESTC vs Full-STC: For each instance, we make its terrain graph incomplete by randomly sampling a set $V'_g \subseteq V_g$ and removing the respective 1, 2, or 3 decomposed vertices from each $\delta \in V'_g$ randomly, without making the decomposed graph unconnected. Fig. 7 shows that ESTC consistently outperforms Full-STC on incomplete terrain graphs.

Initial Solution: We compare the performance of LS-MCPP with different initial solutions in Fig. 8. LS-MCPP shows convergence for all cases, except for *trn1-l* with VOR as a result of its low-quality solution due to its dependence on the initial vertices of robots. MSTC* tends to require more search iterations than MFC and MIP before reaching a good solution since it partitions a single MST to obtain a tree cover that often results in a different solution structure from the direct tree cover computation employed by MFC and MIP. As MIP might yield near-optimal solutions (e.g., *trn1-l*) at the cost of longer runtime, the potential for improvement using LS-MCPP is limited. Therefore, we select MFC to produce initial solutions for LS-MCPP, striking a balance between efficiency and solution quality.

Boundary Editing Operators: We assess the impact of different types of operators on LS-MCPP in Fig. 9. The first column removes grow operators, causing an early stop of LS-MCPP before the $3e3$ iteration limit. This can limit exploration of the solution space, especially when exchange operators struggle to grow light subgraphs effectively due to the graph structure (e.g., *flr-l* instance). The second column removes deduplicate operators, leading to instability and worse performance. These operators are crucial for redupli-




Figure 8: Path costs (y-axis) of 4 robots for LS-MCPP over iterations (x-axis) with different initial solutions. The upper and lower rows correspond to *flr-l* and *trn-l*, respectively.




Figure 9: Path costs (y-axis) of 4 robots for LS-MCPP over iterations (x-axis) with different operators removed. The upper and lower row correspond to *flr-l* and *trn-l*, respectively.

cation between two FORCEDDEDUPLICATION() calls. The third column removes exchange operators, hindering LS-MCPP convergence and solution quality improvement.

Operator Sampling: We use a uniform random operator sampling method (Rand) to compare with the proposed heuristic-based operator sampling method (Heur). In Fig. 10, we compare the first and third columns, and the second and fourth columns. We can see that Heur is much better than Rand in *trn-l*, possibly because *trn-l* is a relatively simpler instance for LS-MCPP thus the difference between the two sampling methods is trivial.

Forced Deduplication: In Fig. 10, we validate the function FORCEDDEDUPLICATION (FD) of LS-MCPP by comparing the first and second columns, and the third and fourth columns. The periodic calling of FD makes the subgraphs return to states with only necessary duplications remaining, guiding LS-MCPP to better solutions.

Performance Comparison

We compare LS-MCPP with the baselines in Tab. 2 with a runtime limit of 24 hours. In summary, LS-MCPP outperforms VOR, MFC, MSTC* for all instances within 20 minutes, demonstrating a makespan reduction of up to 67.0%, 35.7%, and 30.3%, and on average, 50.4%, 26.7%, and 13.4%, respectively. For the first six smaller instances where MIP(H) and MIP can almost compute the optimal tree cover, LS-MCPP achieves an average makepan reduction of -1.02% and 1.16% with orders of magnitude faster runtime. Furthermore, we use the same method as in the ESTC ablation study to make 20% of terrain vertices incomplete for each instance. The last column of Tab. 2 reports results for such instances where we generate covering subtrees of the




Figure 10: Path costs (y-axis) of 4 robots for LS-MCPP over iterations (x-axis) with different operator sampling and whether FORCEDDEDUPLICATION (FD) is used. The upper and lower row correspond to *flr-l* and *trn-l*, respectively.

	VOR	MFC	MSTC*	MIP(H)	MIP	LS-MCPP (Ours)			
<i>flr-s</i>	42.75 0.01s	23 0.03s	21 0.02s	16 13s	0.0%	16 20s	16.75±0.3 5.3±0.13s	21.5±2.1 3.6±2.9s	
<i>trn-m</i>	420.2 0.04s	368.2 0.58s	269.5 0.32s	246.7 24h	0.6%	246.7 24h	1.6%	17.8±0.1s 19.6±1.3s	
<i>mze-m</i>	134.8	67	69	52	0.0%	51 24h	20%	54.3±2.9 10.7±0.1s	
<i>flr-l</i>	303.8 0.07s	294 2.37s	212.5 0.09s	207 24h	8.3%	254 24h	25%	192.1±0.4 30.5±1.3s	207±20.2 35.2±1.6s
<i>mze-l</i>	193.8 0.04s	105 0.10s	139.5 0.58s	91.5 54s	0.9%	93 24h	2.9%	97.3±0.5 15.4±0.2s	104.8±7.6 12.9±5.8s
<i>trn-l-l</i>	1,303 0.07s	597.4 1.67s	468.6 1.93s	435.1 24h	9.7%	419.1 24h	9.7%	429.6±1.4 34.3±0.4s	444±11.2 37.3±1.6s
AR07	1,301	1,254	878.7	1,351	42%	1,333	5%	805.9±7.9	1,024±74
OISR	0.47s	36.9s	9.91s	24h		24h		6.3±0.1m	6.2±0.3m
Shan	1,451	754	576.5	1,509	59%	/	/	570.9±2.1	719±73.0
ghai2	1.21s	7.2m	2.21s	24h		/	/	17.6±2.1m	12.5±0.8m
New	1,735	1,530	1,208	4,736	80%	/	/	1,062±14	1,563±110
York1	1.13s	2.4m	20.2s	24h		/	/	15.0±4.1m	12.6±0.5m

Table 2: Performance comparison. Each instance-method cell reports the makespan at the top and the runtime at the bottom. The percentages for MIP(H)/MIP are the gap between the solution and the lower bound returned by the MIP solver. The last column is for incomplete terrain graphs.

terrain graph using RTC (Even et al. 2004) (used internally by MFC) followed by ESTC on the inducing subgraphs of the RTC subtrees to obtain the initial solution for LS-MCPP. This is necessary since no baseline can handle incomplete terrain graphs. We observe that removing decomposed vertices can alter the connectivity of its neighboring vertices, thus affecting the overall structure of both the decomposed graph and the augmented terrain graph, resulting in a worse LS-MCPP solution than that of the original instance.

Conclusions and Future Work

We introduced LS-MCPP, a new local search framework that integrates the novel standalone ESTC paradigm and boundary editing operators to systematically explore MCPP solutions directly on decomposed graphs for the first time. Future work includes enhancing operator selection with machine learning techniques, speeding up LS-MCPP for larger-scale instances with parallelization techniques, and addressing potential inter-robot collisions in occluded environments with multi-agent pathfinding (Stern et al. 2019) techniques.

Acknowledgements

This work was supported by the NSERC under grant number RGPIN2020-06540 and a CFI JELF award.

References

- Bormann, R.; Jordan, F.; Hampp, J.; and Hägele, M. 2018. Indoor coverage path planning: Survey, implementation, analysis. In *IEEE International Conference on Robotics and Automation (ICRA)*, 1718–1725.
- Collins, L.; Ghassemi, P.; Esfahani, E. T.; Doermann, D.; Dantu, K.; and Chowdhury, S. 2021. Scalable coverage path planning of multi-robot teams for monitoring non-convex areas. In *IEEE International Conference on Robotics and Automation (ICRA)*, 7393–7399.
- Even, G.; Garg, N.; Könemann, J.; Ravi, R.; and Sinha, A. 2004. Min–max tree covers of graphs. *Operations Research Letters*, 32(4): 309–315.
- França, P. M.; Gendreau, M.; Laporte, G.; and Müller, F. M. 1995. The m-traveling salesman problem with minmax objective. *Transportation Science*, 29(3): 267–275.
- Gabriely, Y.; and Rimon, E. 2001. Spanning-tree based coverage of continuous areas by a mobile robot. *Annals of mathematics and artificial intelligence*, 31: 77–98.
- Gabriely, Y.; and Rimon, E. 2002. Spiral-STC: An on-line coverage algorithm of grid environments by a mobile robot. In *IEEE International Conference on Robotics and Automation (ICRA)*, 954–960.
- Galceran, E.; and Carreras, M. 2013. A survey on coverage path planning for robotics. *Robotics and Autonomous systems*, 61(12): 1258–1276.
- Goldberg, D. E. 1989. Genetic algorithms in search. *Optimization, Machine Learning*.
- Hazon, N.; and Kaminka, G. A. 2005. Redundancy, efficiency and robustness in multi-robot coverage. In *IEEE International Conference on Robotics and Automation (ICRA)*, 735–741.
- He, P.; and Hao, J.-K. 2023. Memetic search for the minmax multiple traveling salesman problem with single and multiple depots. *European Journal of Operational Research*, 307(3): 1055–1070.
- Kapoutsis, A. C.; Chatzichristofis, S. A.; and Kosmatopoulos, E. B. 2017. DARP: Divide areas algorithm for optimal multi-robot coverage path planning. *Journal of Intelligent & Robotic Systems*, 86: 663–680.
- Karapetyan, N.; Benson, K.; McKinney, C.; Taslakian, P.; and Rekleitis, I. 2017. Efficient multi-robot coverage of a known environment. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1846–1852.
- Li, L.; Shi, D.; Jin, S.; Yang, S.; Lian, Y.; and Liu, H. 2023. SP2E: Online spiral coverage with proactive prevention extremum for unknown environments. *Journal of Intelligent & Robotic Systems*, 108(2): 30.
- Nagamochi, H.; and Okada, K. 2007. Approximating the minmax rooted-tree cover in a tree. *Information Processing Letters*, 104(5): 173–178.
- Song, H.; Yu, J.; Qiu, J.; Sun, Z.; Lang, K.; Luo, Q.; Shen, Y.; and Wang, Y. 2022. Multi-UAV disaster environment coverage planning with limited-endurance. In *International Conference on Robotics and Automation (ICRA)*, 10760–10766.
- Stern, R.; Sturtevant, N.; Felner, A.; Koenig, S.; Ma, H.; Walker, T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T.; et al. 2019. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *International Symposium on Combinatorial Search (SoCS)*, 151–158.
- Sturtevant, N. 2012. Benchmarks for grid-based pathfinding. *Transactions on Computational Intelligence and AI in Games*, 4(2): 144 – 148.
- Tang, J.; and Ma, H. 2023. Mixed integer programming for time-optimal multi-robot coverage path planning with efficient heuristics. *IEEE Robotics and Automation Letters*, 8(10): 6491–6498.
- Tang, J.; Sun, C.; and Zhang, X. 2021. MSTC*: Multi-robot coverage path planning under physical constrain. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2518–2524.
- Torres, M.; Pelta, D. A.; Verdegay, J. L.; and Torres, J. C. 2016. Coverage path planning with unmanned aerial vehicles for 3D terrain reconstruction. *Expert Systems with Applications*, 55: 441–451.
- Van Laarhoven, P. J.; Aarts, E. H.; van Laarhoven, P. J.; and Aarts, E. H. 1987. *Simulated annealing*. Springer.
- Vandermeulen, I.; Groß, R.; and Kolling, A. 2019. Turn-minimizing multirobot coverage. In *International Conference on Robotics and Automation (ICRA)*, 1014–1020.
- Zheng, X.; and Koenig, S. 2007. Robot coverage of terrain with non-uniform traversability. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3757–3764.
- Zheng, X.; Koenig, S.; Kempe, D.; and Jain, S. 2010. Multirobot forest coverage for weighted and unweighted terrain. *IEEE Transactions on Robotics*, 26(6): 1018–1031.