

Area Coverage With Multiple Capacity-Constrained Robots

Saurav Agarwal[✉] and Srinivas Akella[✉]

Abstract—The area coverage problem is the task of efficiently servicing a given two-dimensional surface using sensors mounted on robots such as unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs). We present a novel formulation for generating coverage routes for multiple capacity-constrained robots, where capacity can be specified in terms of battery life or flight time. Traversing the environment incurs demands on the robot resources, which have capacity limits. The central aspect of our approach is transforming the area coverage problem into a line coverage problem (i.e., coverage of linear features), and then generating routes that minimize the total cost of travel while respecting the capacity constraints. We define two modes of travel: (1) servicing and (2) deadheading, which correspond to whether a robot is performing task-specific actions or not. Our formulation allows separate and asymmetric travel costs and demands for the two modes. Furthermore, the cells computed from cell decomposition, aimed at minimizing the number of turns, are not required to be monotone polygons. We develop new procedures for cell decomposition and generation of service tracks that can handle non-monotone polygons with or without holes. We establish the efficacy of our algorithm on a ground robot dataset with 25 indoor environments and an aerial robot dataset with 300 outdoor environments. The algorithm generates solutions whose costs are 10% lower on average than state-of-the-art methods. We additionally demonstrate our algorithm in experiments with UAVs.

Index Terms—Path planning for multiple mobile robots or agents, aerial systems: applications, computational geometry, area coverage, line coverage.

I. INTRODUCTION

THIS paper addresses the *area coverage problem*—the task of efficiently servicing a given planar surface. There are several applications of the area coverage problem; these include mapping and inspection of large regions using a team of aerial robots (i.e., UAVs), and vacuuming, lawn mowing and harvesting with ground robots (i.e., UGVs). The area coverage problem also applies to CNC-based machining operations [1]. The problem is widely studied in the robotics literature (see recent surveys [2], [3]). However, relatively few approaches for the area coverage problem consider multiple robots. Furthermore,

Manuscript received September 9, 2021; accepted January 3, 2022. Date of publication January 31, 2022; date of current version February 15, 2022. This letter was recommended for publication by Associate Editor L. Pimenta and Editor M. Ani Hsieh upon evaluation of the reviewers' comments. This work was supported in part by NSF under Award IIP-1919233 and in part by a UNC IPG Award. (*Corresponding author:* Saurav Agarwal.)

The authors are with the Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC 28223 USA (e-mail: sagarw10@uncc.edu; sakella@uncc.edu).

Digital Object Identifier 10.1109/LRA.2022.3146952

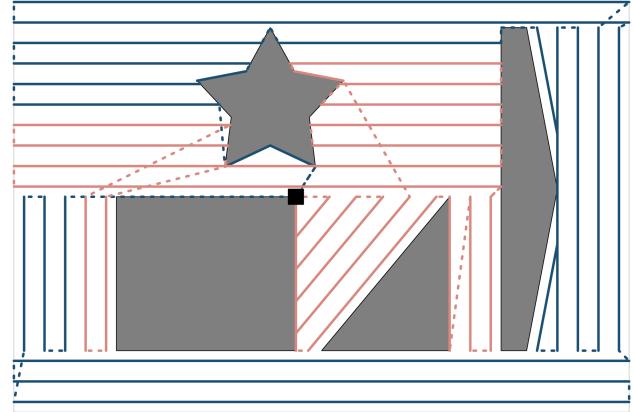


Fig. 1. Area coverage of an environment with a team of capacity-constrained robots. The grey regions represent obstacles in the environment. The black square (near the center) represents the depot location—the robots start and end their routes at the depot. The solution consists of two routes, shown in dark blue and light red. The solid lines represent the service tracks. The dashed lines represent deadheading travel—the robots can turn off the sensors and travel at faster speeds along these line segments.

practical constraints such as limited battery capacity and the effect of wind or uneven terrain are usually not considered. The paper presents a method for area coverage that addresses these challenges. Even when these constraints are not modeled, the algorithms in this paper, in comparison to recent work, generate higher quality solutions.

We consider two modes of travel for a robot. A robot is said to be *servicing* when it performs task-specific actions, such as taking images or vacuuming using its sensors or tools. A robot may travel from one location to another at faster speeds without performing task-specific actions—referred to as *deadheading*—to optimize the mission time, conserve energy, or reduce the amount of sensor data for analysis. The robots usually have a finite amount of resources, such as battery charge, which can be specified in terms of energy or a time limit, referred to hereafter as *capacity*. The robots must return to their home location before the resource consumed exceeds the capacity. The goal is to find efficient routes for a team of robots such that the entire environment is serviced while respecting the capacity constraints. Fig. 1 shows an example environment and routes generated for capacity-constrained robots.

Our formulation for solving the area coverage problem consists primarily of three components: 1) *Cell decomposition* of the environment, 2) *Service track generation* for individual cells,

and 3) *Routing* to traverse the service tracks. The central aspect is to transform the area coverage problem into a line coverage problem—the coverage of linear features in an environment [4]. The service tracks form the linear features that the robots must service, and an efficient algorithm for the line coverage problem is used to generate routes for the team of robots. This allows us to model the cost of travel (e.g., time), the demands on resources (e.g., battery), and asymmetric costs and demands for travel due to wind or uneven terrain. Our formulation facilitates a significant generalization of the cell decomposition component to reduce the number of turns that the robots must take. In particular, the cells are no longer required to be *monotone* polygons [5] with respect to the service direction. This generalization enables additional service directions for the cells to minimize the number of turns. Furthermore, allowing cells to be non-monotone polygons with holes enables the additional merging of adjacent cells with the same service directions. Merging adjacent cells reduces the number of service tracks by avoiding overlapping sensor coverage regions at the common boundary of a pair of adjacent cells. Additionally, we observe that a simple back-and-forth (i.e., boustrophedon) pattern does not always guarantee complete coverage. We mitigate this issue in the new service track generation algorithm.

The contributions of the paper are:

- 1) A cell decomposition algorithm that allows non-monotone polygons and optimizes the number of turns that the robots need to take.
- 2) A new service track generation algorithm capable of handling non-monotone polygons with obstacles. The algorithm improves coverage of the environment over the traditional boustrophedon pattern.
- 3) A new formulation to transform an instance of the area coverage problem into that of the line coverage problem.
- 4) We minimize the total cost of coverage routes for multiple robots while respecting their capacity constraints.
- 5) An open-source implementation¹ of our algorithms.

This is the first method for cell decomposition and service track generation that minimizes the number of turns while allowing non-monotone polygons with holes. Furthermore, this is the first approach for the area coverage problem that allows two modes of travel, capacity constraints, and asymmetric travel costs and demands.

II. RELATED WORK

Area coverage has a large body of work that has been covered extensively in recent survey papers [2], [3]. The area coverage problem is related to the lawn mowing problem, which was shown to be NP-hard [1]. Consequently, several approximation and heuristic algorithms have been proposed. The approaches for area coverage problems can be broadly classified into approximate and exact methods.

Grid-based approaches, which fall under approximate methods, were some of the earliest techniques for solving the area

coverage problem [6]. These methods typically discretize the environment into small cells based on a given resolution. Thus, the quality of the results depends on the resolution [7]. Moreover, the computational complexity increases rapidly with environment size. Vandermeulen *et al.* [8] address coverage with multiple robots using turn minimization as the objective for cell decomposition. The environment is contracted into a rectilinear polygon with integer side lengths. This new polygon is then decomposed into rectangles of unit width (called *ranks*) such that the sum of altitudes is minimized. An m -TSP algorithm is used to find paths for the robots. While the algorithm works well for rectilinear environments, it is not designed for complex non-rectilinear environments.

There has been recent interest in learning-based strategies. However, they are not yet generalizable to large complex environments. Usually, very small grid sizes are used to benchmark the results—a 16×16 grid was used in [9] and a 7×7 grid in [10]. Retraining of the neural network was required for each environment in [10]. Moreover, these do not consider multiple robots. In contrast to the above grid-based approaches, our formulation can handle environments with non-rectilinear boundaries and obstacles. We also allow capacity constraints, asymmetric costs and demands, and two different modes of travel.

In this paper, we focus our attention on exact methods. These methods typically use computational geometry and graph theory. A common approach is to decompose the environment into cells, known as cell decomposition. Choset [11] presented the widely used boustrophedon cell decomposition (BCD), an efficient way to decompose a given environment with obstacles. The key idea is to generate *monotone polygons* [5] with respect to a given service direction using a sweep-line based algorithm.

In most mobile robotics applications, it is desirable to have long paths with as few turns as possible. Turns can be very expensive both in terms of time and battery consumption, as the robot may need to slow down, take a turn, and then accelerate again. Huang [12] presented a minimum sum of altitudes (MSA) formulation. The MSA corresponds to the number of turns required for a robot to service the environment. For both convex and non-convex polygons, the service track orientation that minimizes the number of turns is parallel to one of the polygon edges. A dynamic programming algorithm with an exponential running time was presented to compute an optimal decomposition. In contrast, the BCD is computationally very efficient but does not consider the number of turns. Hence, several heuristic algorithms have been developed that trade off optimizing the number of turns and computational efficiency.

A trapezoidal decomposition was used in [13] to obtain an initial set of cells that are then merged to reduce the number of cells. Service directions are determined by using a bisection search. A sweep-line based algorithm, similar to BCD, was presented in [14] to obtain an initial decomposition of the environment. A service direction is determined independently for each cell. Adjacent cells that have the same service direction are merged if they remain monotone even after merging. In [15], an initial decomposition of the environment is obtained by extending interior edges of concave vertices. An integer programming formulation and a heuristic algorithm are proposed to obtain

¹Source code available at:
<https://github.com/UNCCharlotte-CS-Robotics/AreaCoverage-library>.

solutions efficiently. An approach based on vehicle routing problems is used to route multiple robots. All these techniques require a set of monotone polygons. In contrast, our approach removes this requirement, enabling us to improve the cell decomposition procedure.

In [16] and [17], a Reeb graph is generated from the BCD, where the cells are represented by edges and the connectivity of cells is represented using vertices. An algorithm for the Chinese postman problem (CPP) finds a tour on the Reeb graph, which provides a sequence of cells to be visited by the robot. A single service direction is assumed for the entire environment, determined by simple heuristics such as longest edge or wind direction. This work was extended in [18] to multiple robots using clustering and the k -CPP algorithm for routing. These methods do not consider the minimization of the number of turns. Furthermore, each cell is treated as a unit. Assuming a fixed path for individual cells or treating each cell as a unit can be very restrictive, especially for capacity-constrained robots. The robots might not be able to cover multiple cells or even a single large cell, resulting in a large number of inefficient routes.

Algorithms for the generalized traveling salesperson problem (GTSP) are often used for computing routes. In [19], the BCD is used to obtain a set of cells with the same service direction. A GTSP instance is generated with two vertices for each service track for the two travel directions, forming a cluster. A GTSP algorithm generates a tour such that a single vertex is traversed from each cluster. In [20], the cell decomposition starts with any convex decomposition of the polygon and is improved by adding cuts at reflex vertices. Finally, the GTSP is used to generate a tour on an auxiliary graph, similar to [19]. In a recent paper [21], the BCD is computed for each edge direction, and cells are assigned independent service directions. The BCD that has the least MSA is selected. For each cell and edge direction, four patterns for servicing are provided based on where the robot starts and ends. Each pattern forms a vertex in the GTSP instance graph, and vertices corresponding to the same cell are grouped in a cluster. A visibility graph is used to form edges between vertices. A GTSP tour then traverses a vertex, representing a pattern, from each cluster. In GTSP [19], [20] and m -TSP [8] based approaches, the cells are usually not treated as a unit (except for [21]) and thus, are more efficient. These procedures do not consider capacity constraints and are designed for a single robot. Algorithms for vehicle routing problems (VRP) [22] allow capacity constraints and multiple robots. Although the costs of the edges in both VRP and GTSP graphs can be asymmetric, the service tracks are represented as nodes, which do not have costs or demands. Thus the nodes cannot model asymmetric costs and demands of the service tracks. Our approach uses the line coverage problem to closely model the area coverage problem, with the edges in the graph representing the service tracks. This enables modeling of capacity constraints, asymmetric costs and demands, and two modes of travel.

III. THE AREA COVERAGE PROBLEM

Given a region $\mathcal{R} \subset \mathbb{R}^2$, the area coverage problem is to find a set of routes for a team of robots such that the total cost of the

routes is minimized, and all the points in the region are serviced by the robots. Limited battery life is one of the most critical restrictions on mobile robots, especially for aerial robots. Thus, in our formulation, we incorporate an additional constraint that the total *demand* on resources for each route should not exceed a given capacity for the robots. The capacity can be specified in terms of energy, time limit, or travel length. We have two modes of travel for the robots: (1) servicing and (2) deadheading. A robot is said to be servicing if it performs task-specific actions, such as taking images, as it traverses a path. A robot may travel from one location to another while not performing servicing tasks, such as returning to the home/depot location. Such travel is known as deadheading. Functions for service and deadhead costs and demands are given as input to the problem. Our formulation can handle separate and asymmetric costs and demands.

We model the environment with a set of polygons. The environment may have obstacles or sub-regions that are not required to be serviced. These sub-regions are referred to as *holes*. Depending on the application, the robots may be permitted to travel across the holes, e.g., an aerial robot flying at a high altitude may optimize its path by flying over a hole representing a building. We treat the robots as point robots, unless otherwise specified. For finite-sized robots, we compute the free workspace using techniques for computing configuration space, such as the Minkowski sum [5].

We now describe our approach to solve the area coverage problem with multiple capacity-constrained robots. We break the problem into three components: (1) Cell decomposition, (2) Service track generation, and (3) Routing.

A. Cell Decomposition

The primary motivation for the cell decomposition component is to minimize the number of turns that the robots need to take. This is done by decomposing the environment into smaller polygons, referred to as cells, and computing a *service direction* that minimizes the number of turns for each cell independently. Such an optimal direction is related to the minimum sum of altitudes (MSA) of a polygon and is parallel to one of the edges of the boundary or hole of the cell [12]. The service tracks are generated parallel to the corresponding service direction for a cell. Our cell decomposition method is a culmination of experimentation with various existing methods. However, we deviate significantly in one crucial aspect—we allow the cells to be non-monotone with respect to the direction perpendicular to the service direction, i.e., the intersection of a cell (its interior) and a line parallel to the service direction need not be a connected line segment. Allowing non-monotone cells increases the feasible solution space, enabling cell decompositions that can potentially reduce the number of turns.

Our cell decomposition method is composed of three steps (1) initial decomposition, (2) greedy improvement, and (2) cell merging.

Initial Decomposition: We use an approach similar to [21] for the initial decomposition.

- 1) Obtain a set of directions corresponding to the edges of the environment, i.e., edges of the outer boundary and the

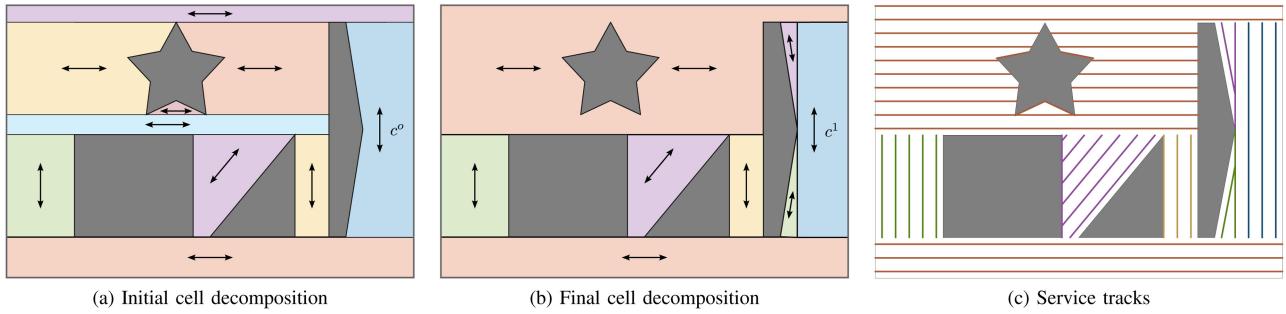


Fig. 2. An environment with four obstacles. The double head arrows indicate the service directions. (a) In the initial decomposition with ten cells, the cell c^o has an optimal service direction for which it is not monotone. (b) The final decomposition, with eight cells, is obtained after the greedy improvement and then merging adjacent cells with the same service direction. The cells in the final decomposition are also not necessarily monotone polygons. (c) The service tracks are generated for each cell independently, and overlapping segments are removed. Service tracks include those belonging to the two scenarios described in Fig. 3, e.g., the vertical purple and green tracks on the left edge of cell c^1 and some of the edges of the star-shaped obstacle.

holes. Parallel directions are ignored. Let \mathcal{V} denote the set of all such directions.

- 2) For each direction $v \in \mathcal{V}$:
 - a) Perform BCD with a line parallel to v and sweeping perpendicular to itself. Let \mathcal{C} denote the set of cells obtained from the BCD.
 - b) For each cell $c \in \mathcal{C}$: Compute the MSA a_c and the corresponding service direction u_c , even if the cell c is non-monotone with respect to the direction perpendicular to u_c .
 - c) Compute the total MSA for the direction v : $\alpha_v = \sum_{c \in \mathcal{C}} a_c$.
- 3) Select the decomposition that gives the minimum sum of altitudes, i.e., $\alpha^* = \min_{v \in \mathcal{V}} \alpha_v$.

Fig. 2(a) shows an initial decomposition consisting of ten cells for an environment with four obstacles. The double arrows indicate the optimal service direction. Note that the right cell marked c^o is non-monotone with respect to the direction perpendicular to the optimal service direction. Such directions would have been eliminated in [21].

The running time for this step is $\mathcal{O}(n^2 \log n)$, where n is the number of vertices in the environment, and it dictates the overall complexity for cell decomposition.

Greedy Improvement: Improvements to the decomposition have been shown by further decomposition of the cells by splitting them along edges corresponding to a non-convex vertex of a polygon [12], [15], [20]. Thus, we apply a greedy strategy to split the cells in the initial decomposition further. We identify the non-convex vertices for a cell and compute a set of splitting lines \mathcal{L}_s . There are two types of splitting lines: (1) The set of lines corresponding to an edge adjacent to a non-convex vertex, and (2) The set of lines parallel to an edge of the cell and passing through a non-convex vertex such that both its adjacent edges lie on the same side of the line. We ensure that the splitting lines are neither parallel nor anti-parallel to each other. For each line $l \in \mathcal{L}_s$, we split the cell polygon to obtain a new set of polygons \mathcal{C}_l . Now obtain the total MSA for \mathcal{C}_l . If this total MSA is less than the MSA for the original cell, then the line l is a valid candidate for splitting. Of all the splitting lines $l \in \mathcal{L}_s$, we greedily select the one that gives the least total MSA after

splitting. The new polygons are then recursively improved using the same procedure. The cell marked c^o in Fig. 2(a) has been split further by a vertical line of Type 2 to create three new cells, shown in Fig. 2(b). Note that the splitting line does not lie within the *cone of bisection* described in [20]. As the objective is to reduce the number of turns, the greedy improvement may split a non-monotone cell into monotone cells if doing so reduces the number of turns. We used a greedy approach instead of a dynamic programming approach to reduce computation costs.

Cell Merging: Adjacent cells that have the same (or similar) service directions can be merged to reduce the total number of cells [13], [14]. Merging adjacent cells reduces the overlapping regions of sensor coverage. This decreases the number of service tracks and the number of turns. We allow the merging of adjacent cells even when the resulting cell is non-monotone and contains holes; this significantly reduces the number of cells. Fig. 2(b) shows the final decomposition with eight cells. The total length of the service tracks for the initial decomposition is 2146 m. The greedy improvement reduces it to 2103 m, and cell merging reduces it further to 2003 m, an improvement of 6.7% over the initial decomposition. The most significant reduction comes from merging the cells around the star-shaped obstacle into a non-monotone cell with a hole.

B. Service Track Generation

The next step is to generate the service tracks for each cell. Since the cells are not required to be monotone polygons with respect to the service direction, we develop a new algorithm to generate the service tracks. Existing approaches usually generate a path, in the form of a boulevard or lawn mower pattern, within each cell that a robot must follow. We identify two scenarios wherein a boulevard pattern does not guarantee complete coverage, as shown in Fig. 3. Assuming a square sensor field-of-view, these can happen (1) when an edge is oriented at a very small angle with the service direction, and (2) when an edge is inclined at an angle smaller than $\pi/4$ with the service direction and intersects the service tracks. These scenarios can be extended to other types of sensor field-of-view as well. The second scenario was discussed for disc-shaped sensors in [8].

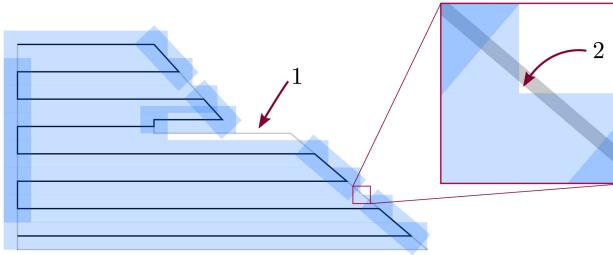


Fig. 3. Two scenarios, identified by the red arrows, for which a boustrophedon pattern does not always guarantee a complete coverage even when the polygon is monotone with respect to the service direction (horizontal here). The blue shaded region represents coverage of the environment with a square sensor field-of-view.

We now discuss a new algorithm, based on the sweep-line algorithm [5], for generating the service tracks for a cell obtained from the cell decomposition step. The algorithm can handle non-monotone cells with holes and resolves the issues shown in Fig. 3. Without loss of generality, we assume the service direction is parallel to the X -axis. The sweep line is parallel to the service direction and sweeps vertically from the lowest to the highest vertex while keeping track of the edges it encounters. The vertices of the edges correspond to the events in the sweep-line algorithm. We assume a square sensor field-of-view for clarity.

Initialization:

- Represent a cell vertex u by its coordinates (u_x, u_y) , and represent each cell edge by a pair of vertices (u, v) such that $u_y \leq v_y$, i.e., the lower vertex appears first.
- Let \mathcal{E} be a list of all the edges in the cell. Sort the edges in \mathcal{E} in ascending order by u_y , the y -coordinate of the lower vertex u , for faster computation in Step (2).
- Let \mathcal{E}_c represent the set of current edges used for generating service tracks parallel to the service direction.
- Let \mathcal{E}_s represent the set of special edges of the cell that correspond to scenario 1 shown in Fig. 3.
- Let \mathcal{S} represent the set of service tracks.
- Initially, the sets \mathcal{E}_c , \mathcal{E}_s , and \mathcal{S} are all empty.
- Let the sweep line L be the line parallel to the service direction (X -axis) and passing through the lowest vertex $u^l = (u_x^l, u_y^l)$ of the first edge in the sorted list \mathcal{E} .
- Set the offset $o = u_y^l + f/2$, where f is the lateral sensor field-of-view.

Iteration:

while the edge list \mathcal{E} is not empty

- 1) Offset the line L to be at a distance o from the X -axis.
- 2) For each edge $e \in \mathcal{E}$ with $u_y \leq o$, remove e from \mathcal{E} :
 - if $v_y \leq o$, add e to \mathcal{E}_s ,
 - else, add e to the set of current edges \mathcal{E}_c . If the edge subtends an angle smaller than $\pi/4$, add e to the list of service tracks \mathcal{S} as well (scenario 2).
- 3) Compute the intersections of the line L with the edges in \mathcal{E}_c . Sort these intersection points in ascending order of the x -coordinate. Generate service tracks using the intersection points such that the tracks do not overlap with the interior of the obstacles. Add the service tracks to the set \mathcal{S} .
- 4) For each edge $e \in \mathcal{E}_s$, check if e lies within the field-of-view of any of the service tracks generated in the current and the

previous iterations. If not, add e to the set of service tracks \mathcal{S} . Remove all edges from \mathcal{E}_s .

5) Offset the sweep line by f and let $o = o + f$.

The set \mathcal{S} gives the set of service tracks for a cell. We compute the service tracks for each cell in the decomposition. Finally, overlapping components of any pair of service tracks for the entire environment are iteratively removed. Fig. 2(c) shows the generated service tracks. The running time complexity is $\mathcal{O}(n_c \text{MSA}_c)$, where n_c and MSA_c are the number of vertices and the MSA for a cell c , respectively.

C. Routing

Once the service tracks are computed, the problem can be transformed into the line coverage problem—the coverage of linear features. Here the service tracks correspond to the linear features in the environment. We have recently addressed the line coverage problem and presented an efficient algorithm for multiple capacitated robots [4]. We use a graph as the underlying data structure for the transformation.

Vertices V : The set of vertices V consists of:

- the endpoints of the service tracks;
- the vertices of the environment polygons;
- the depot: A special vertex in the graph from where the routes start and end. For aerial robots, it corresponds to the home location for take off and landing.

Required Edges E_r : A required edge is an edge that needs to be serviced exactly once, and therefore the set of required edges E_r is precisely the set of service tracks. There is a cost and a demand associated with each of the required edges. Furthermore, the edges are considered to have asymmetric costs and demands. Techniques from [23]–[25] can be used to obtain demands on battery life. Demands can also be specified in terms of time. Our algorithm can handle arbitrary non-negative input values for costs and demands.

Non-required Edges E_n : We add a non-required edge between each pair of vertices $u, v \in V$ such that the line segment (u, v) does not pass through any of the obstacles and remains within the interior of the outer boundary. We compute a visibility graph to determine if a pair of vertices forms a valid edge for travel. In applications where the robots may travel across the holes, such as aerial robots flying at high altitudes, we do not need to check if the line segment crosses the holes. We may still need to compute the visibility graph as the outer boundary may be non-convex. The robots are not required to traverse the non-required edges. They may, however, use these edges to travel quickly from one vertex to another. A robot is said to be deadheading when traveling along a non-required edge. There is a cost and demand associated with deadheading also. As task-specific actions such as taking images need not be performed, the robots may travel faster than when servicing. Thus, the deadheading costs and demands can differ from those for servicing. Energy-efficient operating speeds for servicing and deadheading can be obtained through experiments [24]; these show that as the speed increases, the power consumption first decreases and then increases rapidly as the speed approaches the upper limit.

Capacity Q: A fixed capacity such as battery life or maximum flight time is specified for the robots. The total demand on the resources accumulated from traversing required and non-required edges along each route should not exceed the capacity Q . This constraint is critical for safe operations, particularly for UAVs, as the battery life can be very limited.

Let $G = (V, E, E_r)$ be the graph created from the area coverage problem, where $E = E_r \cup E_n$ is the set of all the edges. The line coverage problem is then to find a set of routes such that the total cost of travel is minimized and each of the required edges is serviced, while ensuring that the total demand for each route is less than the capacity of the robots [4]. We use the Merge-Embed-Merge (MEM) algorithm [4] to solve the line coverage problem as it is fast and efficient for robotics applications.

The MEM algorithm is composed of four elements: (1) *initialization* of routes, (2) computation of *savings*, (3) *merging* of two routes to form a new route, and (4) *embedding* the new merged route. A route is initialized for each of the required edges (i.e., service tracks) by deadheading from the depot to one of the vertices of the edge, servicing the required edge, and then deadheading back to the depot via the other vertex. Since the costs are asymmetric, the service direction that gives a lower route cost is selected. Two routes can be merged to form a new route with potentially lower cost than the sum of the costs of the two routes. The difference in the costs by performing such a merge is called *savings*. There are eight possible ways of merging two routes, and only the ones that satisfy the capacity constraint are considered. For each pair of initial routes, the optimal savings that respects the capacity constraint is inserted into a max-heap data structure. The routes with the maximum savings are extracted from the max-heap and are merged to form a new route. The new route is inserted into the set of routes, and the individual routes are set to invalid. This new route is then embedded into the max-heap by computing savings with the other valid routes. The merging and the embedding operations are performed iteratively until no further valid merge is possible. Since the capacity constraints are checked before creating a new route, the algorithm always maintains a set of feasible routes. The running time of the algorithm is $\mathcal{O}(m^2 \log m)$, where m is the number of required edges (i.e., service tracks).

Using the line coverage problem, instead of a node routing problem such as the GTSP or the vehicle routing problem [22], allows direct modeling of the service tracks as graph edges. More importantly, asymmetric costs and demands for the edges can be modeled in the line coverage problem, along with the capacity constraints. The MEM algorithm for the line coverage problem rapidly computes routes of high quality.

The service track generation and the routing components are independent modules. The modular nature of these components allows making independent improvements. Tracks can be generated by other decomposition methods, and a line coverage algorithm can be used for routing.

IV. SIMULATIONS AND EXPERIMENTS

The algorithms for the area coverage problem are implemented in C++. We use the computational geometry algorithms library (CGAL) [26] for precise numerical computations



Fig. 4. Coverage of a large indoor environment with 107 m^2 area and 151 vertices, given in [8]. The left figure shows our solution for two robots with a tool width of 0.1 m and a capacity of 700 m, with tour costs 616 m and 620 m. The right figure shows our solution for four robots with a capacity of 320 m, with tour costs 312 m, 313 m, 318 m, and 300 m.

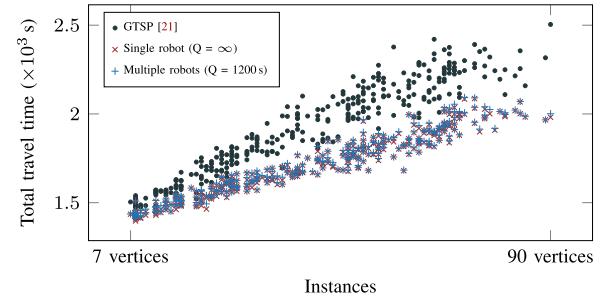


Fig. 5. Comparison of total travel time cost of routes computed by our algorithm and the GTSP based algorithm [21]. The instances are arranged in increasing order of the number of vertices in the environments. We compute the results for two cases: (1) Single robot with infinite capacity (red crosses), and (2) Multiple robots with a capacity of 1200 s (blue pluses). The sum of costs for a single robot over all 300 instances is 577,218 s for [21] and 512,619 s for our method. Our algorithm consistently performs better than [21], with an average cost reduction of 10%.

and geometry functionalities. The program is executed on a desktop with an Intel 7thGen Core i9-7980XE processor.

A. Dataset With 25 Indoor Environments

Simulation results on a dataset² with 25 large indoor environments for vacuuming robots were presented in [8]. The environments are primarily rectilinear in structure. The robots have a tool width of 0.1 m. We use the path length for both the cost and demand functions for direct comparison with [8]. We compute the free workspace by taking a Minkowski sum of the obstacle polygons with the square robot geometry [5]. The robots must graze the boundaries to vacuum thoroughly. Thus, the entire boundary of the free workspace is added to the set of service tracks. Thereafter, we run the three components of our algorithm to obtain the coverage routes for the robots.

The cumulative lengths and number of turns for the 25 environments are presented in Table I. Fig. 4 shows our solution with 2 and 4 robots for the largest environment.

²The datasets and our detailed results are available at: <https://github.com/UNCCharlotte-CS-Robotics/AreaCoverage-dataset>.

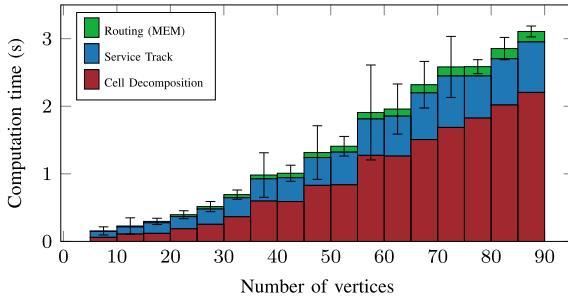


Fig. 6. Average computation times for the 300 outdoor environment dataset from [21]; the route costs are shown in Fig. 5. Each environment is placed in bins of size 5 according to the number of vertices in the environment. The time is averaged over 100 runs for all the environments in a bin. The bars indicate the standard deviation of the computation time. The computation time increases with the number of vertices in the environment. The cell decomposition is the most time-consuming step of the algorithm, while the MEM routing algorithm is very fast.

TABLE I
CUMULATIVE RESULTS FOR THE 25 INDOOR ENVIRONMENTS DATASET

r	Capacity	[8]		This Paper		Improvement (%)	
		l (m)	η	l (m)	η	l	η
1	∞	15,195	11,377	14,781	10,183	2.72	10.49
2	0.75	15,303	11,380	14,793	10,191	3.33	10.45
3	0.50	15,461	11,533	14,823	10,211	4.13	11.46
4	0.30	15,564	11,586	14,939	10,274	4.02	11.32
5	0.25	15,715	11,663	15,030	10,308	4.36	11.62

The first column r indicates the number of robots. The length and the number of turns are denoted by l and η , respectively. The capacity is set as a fraction of the route cost for a single robot. The average computation time is 0.42 s, over all 25 environments and over 100 runs. The computation time does not vary much with the number of robots.

B. Dataset With 300 Outdoor Environments

To benchmark our results for outdoor coverage with aerial robots, we use the dataset (provided with the source code) from [21] consisting of 300 unique environments with 1 to 15 holes derived from buildings. The outer boundary has an area of 10,000 m², and an aerial robot with a 3 m sided square sensor field-of-view is used. The trajectories of the robots are defined by a velocity ramp model with a maximum acceleration a_{\max} and velocity v_{\max} of 1 m/s² and 3 m/s, respectively. The travel time t is used as the cost and demand functions and is given in terms of segment length d as:

$$t = \begin{cases} \sqrt{\frac{4d}{a_{\max}}}, & \text{if } d < d_a \\ \frac{v_{\max}}{a_{\max}} + \frac{d}{v_{\max}}, & \text{if } d \geq d_a \end{cases}, \text{ where } d_a = \frac{v_{\max}^2}{a_{\max}}$$

We ran the simulations for two scenarios: (1) infinite capacity, representing coverage with a single robot, and (2) capacity of the robots set to 20 minutes (1200 s). The comparison of the total cost of the routes is shown in Fig. 5. Our algorithm generates lower cost solutions than that of [21] for both single and multiple robots and for all the instances, with an average improvement of 10% and a standard deviation of 4%. The total cost for the multiple robot solutions is the sum of the costs of the individual routes, and the solutions are better than [21], even though a limited battery capacity reduces the feasible space considerably.

The computation time is shown in Fig. 6, and is similar for both single and multiple robots; the only difference is in the running time of the MEM algorithm, which converges faster as the capacity decreases.

C. Outdoor Experiment With Aerial Robots

We selected a 19,000 m² area in the UNC Charlotte campus, shown in Fig. 7, for coverage with UAVs. An appropriate launch site was assigned as a depot. A subregion corresponding to the footprint of a building was selected as a hole, consisting of 45 vertices, in the environment. As aerial robots can fly at high altitudes, we allow non-required edges that cross the hole. The servicing and deadheading speeds were set to 3.33 m/s and 5 m/s, respectively. A wind of 1.39 m/s at an angle of 225 degrees (from NE), for the day of the experiment, was incorporated into the cost and demand functions, making the edges asymmetric in the two directions of travel. The costs and the demands are based on the edge travel times. A conservative capacity of 600 s was set, and two routes, computed in 2.4 s, were obtained. A DJI Phantom 4 drone was used to autonomously fly the two routes sequentially. The routes and the orthomosaic obtained from the collected images are shown in Fig. 7. The computed costs for the routes, shown in blue and red in Fig. 7, were 336 s and 394 s, and the flight times were 317 s and 369 s.

Discussion: In our simulations and experiments, we considered three types of scenarios: (1) The ground robots of finite size cannot intersect with the obstacles, (2) The aerial robots are not permitted to fly over obstacles, and (3) The aerial robots can fly over obstacles. Using a visibility graph, we can address any combination of the above scenarios by permitting non-required edges only over the obstacles that the robots can traverse. We can also compute the Minkowski sum for the obstacles that a finite-sized robot is not allowed to overlap. Furthermore, non-overlapping disconnected regions of environments can also be addressed by performing cell decomposition and service track generation for each such region individually and computing routes using the MEM algorithm for the service tracks in a unified manner.

V. CONCLUSION

We presented a novel approach for solving the area coverage problem with multiple capacity-constrained robots by transforming it into the line coverage problem. This allowed us to generate routes that minimize the total cost of travel while respecting the capacity constraints. The formulation enables two modes of travel—servicing and deadheading—with distinct and asymmetric costs and demands that can have arbitrary non-negative values. Travel time, travel length, or battery consumption can be used to model costs and demands. A depot, from where the robots start and end their routes, can be specified. These features were demonstrated in an outdoor experiment using a commercial UAV.

The cell decomposition permits non-monotone polygons, thus increasing the feasible solution space for the service directions to further minimize the number of turns. Allowing non-monotone polygons with holes enables further merging of adjacent cells. A new service track generation algorithm generates tracks for

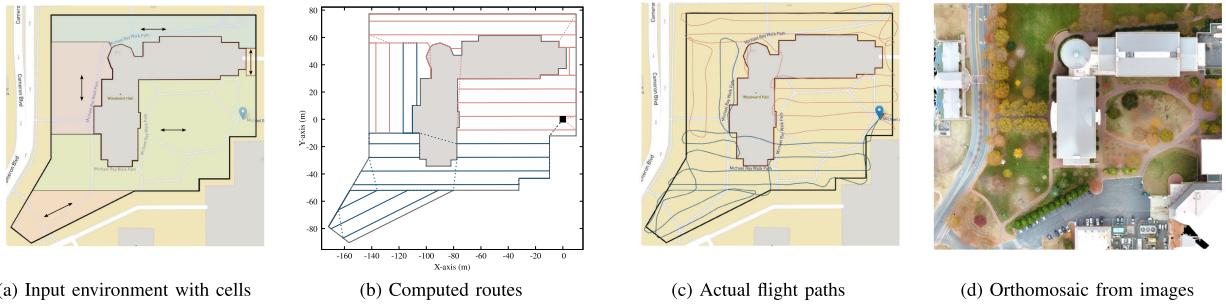


Fig. 7. Area coverage by autonomous aerial robots: (a) The region surrounding a building is to be covered. The blue marker indicates the depot location for the robots to take off and land. The cell decomposition is shown with double head arrows indicating service directions. Note that the cells are non-monotone. (b) The generated routes for the aerial robots, distinguished by color. The dashed lines correspond to deadheading travel. Here deadheading is permitted over the building for efficiency. (c) The actual paths taken by the aerial robots. (d) Orthomosaic generated from images taken during the flights.

non-monotone polygons with or without holes. We benchmarked the approach on a ground robot dataset with 25 indoor environments and an aerial robot dataset with 300 environments, with an average cost improvement of 10%.

Since we establish that the cells from cell decomposition are no longer required to be monotone, our work raises the following questions: Is there a better strategy for cell decomposition to minimize the number of turns? Is a polynomial-time optimal algorithm possible, or is this cell decomposition problem NP-hard?

Future work includes incorporating turning costs and non-holonomic constraints, and using a more realistic energy model for the demands and the capacity. Another interesting direction is maximizing the area covered by the robots when there are limits on the number of routes or robots.

ACKNOWLEDGMENTS

We thank Isaac Vandermeulen for sharing the ground robot dataset from [8], and Rik Bähnemann for publishing the source code and aerial robot dataset from [21]. Fig. 7 uses map data from Mapbox and OpenStreetMap.

REFERENCES

- [1] E. M. Arkin, S. P. Fekete, and J. S. B. Mitchell, “Approximation algorithms for lawn mowing and milling,” *Comput. Geometry*, vol. 17, no. 1, pp. 25–50, 2000.
- [2] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robot. Auton. Syst.*, vol. 61, no. 12, pp. 1258–1276, Dec. 2013.
- [3] T. M. Cabreira, L. B. Brisolara, and P. R. Ferreira Jr., “Survey on coverage path planning with unmanned aerial vehicles,” *Drones*, vol. 3, no. 1, 2019, Art. no. 4.
- [4] S. Agarwal and S. Akella, “Line coverage with multiple robots,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 3248–3254.
- [5] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed. Berlin, Germany: Springer-Verlag, 2008.
- [6] Y. Gabriely and E. Rimon, “Spanning-tree based coverage of continuous areas by a mobile robot,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2001, vol. 2, pp. 1927–1933.
- [7] M. Wei and V. Isler, “Coverage path planning under the energy constraint,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 368–373.
- [8] I. Vandermeulen, R. Groß, and A. Kolling, “Turn-minimizing multirobot coverage,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 1014–1020.
- [9] K. G. S. Apuroop, A. V. Le, M. R. Elara, and B. J. Sheu, “Reinforcement learning-based complete area coverage path planning for a modified hTrihex robot,” *Sensors*, vol. 21, no. 4, Feb. 2021, Art. no. 1067.
- [10] M. Theile, H. Bayerlein, R. Nai, D. Gesbert, and M. Caccamo, “UAV coverage path planning under varying power constraints using deep reinforcement learning,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 1444–1449.
- [11] H. Choset, “Coverage of known spaces: The boustrophedon cellular decomposition,” *Auton. Robots*, vol. 9, no. 3, pp. 247–253, 2000.
- [12] W. H. Huang, “Optimal line-sweep-based decompositions for coverage algorithms,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2001, vol. 1, pp. 27–32.
- [13] T. Oksanen and A. Visala, “Coverage path planning algorithms for agricultural field machines,” *J. Field Robot.*, vol. 26, no. 8, pp. 651–668, 2009.
- [14] X. Yu and J. Y. Hung, “Coverage path planning based on a multiple sweep line decomposition,” in *Proc. 41st Annu. Conf. IEEE Ind. Electron. Soc.*, 2015, pp. 4052–4058.
- [15] L. D. Nielsen, I. Sung, and P. Nielsen, “Convex decomposition for a coverage path planning for autonomous vehicles: Interior extension of edges,” *Sensors*, vol. 19, no. 19, 2019, Art. no. 4165.
- [16] R. Mannadiar and I. Rekleitis, “Optimal coverage of a known arbitrary environment,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2010, pp. 5525–5530.
- [17] A. Xu, C. Viriyasuthee, and I. Rekleitis, “Efficient complete coverage of a known arbitrary environment with applications to aerial operations,” *Auton. Robots*, vol. 36, no. 4, pp. 365–381, Apr. 2014.
- [18] N. Karapetyan, K. Benson, C. McKinney, P. Taslakian, and I. Rekleitis, “Efficient multi-robot coverage of a known environment,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 1846–1852.
- [19] J. S. Lewis, W. Edwards, K. Benson, I. Rekleitis, and J. M. O’Kane, “Semi-boustrophedon coverage with a dubins vehicle,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 5630–5637.
- [20] S. Bochkarev and S. L. Smith, “On minimizing turns in robot coverage path planning,” in *Proc. IEEE Int. Conf. Automat. Sci. Eng.*, 2016, pp. 1237–1242.
- [21] R. Bähnemann, N. Lawrence, J. J. Chung, M. Pantic, R. Siegwart, and J. Nieto, “Revisiting boustrophedon coverage path planning as a generalized traveling salesman problem,” in *Proc. Field Service Robot.*, 2021, pp. 277–290.
- [22] D. Vigo and P. Toth, Eds. *Vehicle Routing*, 2nd ed. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2014.
- [23] Y. Mei, Y.-H. Lu, Y. C. Hu, and C. S. G. Lee, “Deployment of mobile robots with energy and timing constraints,” *IEEE Trans. Robot.*, vol. 22, no. 3, pp. 507–522, Jun. 2006.
- [24] C. D. Franco and G. Buttazzo, “Energy-aware coverage path planning of UAVs,” in *Proc. IEEE Int. Conf. Auton. Robot Syst. Competitions*, 2015, pp. 1111–1117.
- [25] T. M. Cabreira, C. D. Franco, P. R. Ferreira, and G. C. Buttazzo, “Energy-aware spiral coverage path planning for UAV photogrammetric applications,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3662–3668, Oct. 2018.
- [26] CGAL The Project, CGAL user and reference manual, 5.2.1 ed. CGAL editorial board, 2021. [Online]. Available: <https://doc.cgal.org/5.2.1/Manual/packages.html>