Decentralized Multiple Mobile Depots Route Planning for Replenishing Persistent Surveillance Robots

Yifan Ding

CMU-RI-TR-19-17

May 2019

Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Katia Sycara, Chair
Maxim Likhachev
Wenhao Luo

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

For my parents and friends

Abstract

Persistent surveillance of a target space using multiple robots has numerous applications. The continuous operation in these applications is challenged by limited onboard battery capacity of the persistent robots. We consider the problem for replenishing persistent robots using mobile depots, where mobile depots collectively compute a set of tours to drop off batteries for replenishing all persistent robots with the minimum total cost. Compared to other works, persistent robots are not required to detour for recharging or battery swapping. We formulate this problem as generalized multiple depots traveling salesman problem (GMDTSP) on a complete graph. An efficient centralized heuristic-based algorithm Multiple Depots Random Select (MDRS) is proposed, which has been proved to have an analytical constant upper bound in the worst case scenario. To provide scalability and robustness, a fully decentralized asynchronous MDRS (dec-MDRS) is proposed, with the analysis of its correctness and efficiency. We also propose a post-processing heuristic (MDRS-IM) to improve the solution quality. We demonstrate the efficiency and effectiveness of our algorithm via benchmark on multiple instances from TSPLIB and GTSPLIB. The simulation results show that the complexity of dec-MDRS grows linearly as the number of robots increases. The simulation also shows that MDRS and MDRS-IM perform significantly faster than the state of the art heuristic solver LKH with only a loss of about 10% of solution quality.

Acknowledgments

Firstly, I would like to express my sincere gratitude to my advisor Prof. Katia Sycara for the continuous support of my master study and related research. Her patience, enthusiasm and immense knowledge are invaluable to this thesis. Without her guidance and persistent help, this thesis would not have been possible.

I would also like to thank the rest of my thesis committee, Prof. Maxim Likhachev, and Wenhao Luo for their insightful comments and encouragement, but also for the hard questions which incented me to widen my research from various perspectives.

A very special gratitude goes out to all my lab mates in the Advanced Agent - Robotics Technology Lab. It was great sharing laboratory with all of you during the last two year. Working with them has been an illuminating experience for me; this accomplishment would not have been possible without them. In particular, I am grateful to Dr. Changjoo Nam for enlightening me the first glance of research.

Finally, I would like to thank my family and friends for supporting me and encouragement throughout my years of study, throughout researching and writing this thesis, as well as my life in general.

Contents

x

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Persistent surveillance of a target space using multiple robots has numerous applications such as geographical surveys, air quality monitoring, and security monitoring [14, 16, 18]. Limited onboard battery capacity is one of the challenges for persistent robots to execute persistent tasks continuously. To address the long term operation for persistent surveillance tasks, route planning for periodic recharging or battery swapping becomes one of the popular research questions.

There is an extensive literature on replenishing persistent robots by placing static depots, where the persistent robots need to detour for replenishment [2, 10, 15]. These literature discuss the optimal quantity and placement for the static depots which lead to a minimum cost for the persistent robots to recharge. However, persistent robots need to be removed from the tasks for replenishment, which becomes a challenge for safety-critical missions. Moreover, due to the dynamics for the environment, the optimal placement always changes to time and is relevant to the persistent missions.

In this paper, multiple mobile depots are deployed to replenish persistent robots so that they are not required to detour for replenishment. Fig. 1.1 shows an example for our formulation, where there are two types of robots, persistent robots and mobile depots. Persistent robots, with limited onboard battery capacity, move on a pre-planned fixed route to monitor their prescribed areas. Before persistent robots run out of battery, they determine a set of potential preferred battery swapping locations. Mobile depots, which assumed to carry unlimited batteries, move among persistent robots to drop off batteries at some battery swapping locations. To minimize detour for persistent robots, we assume all potential preferred battery swapping locations are on its pre-planned route. Once persistent robots travel to the places where the battery is dropped off, they swap the battery to replenish itself. The persistent robots only accept to swap battery at one of their preferred locations. The problem we need to solve is to find tours for the mobile depots to replenish all persistent robots within their preferred battery swapping locations with the minimum total cost. We formulate this problem as Generalized Multiple Depots Traveling Salesman Problem (GMDTSP), which turns out to be an NP-hard problem. The challenge for this problem is 1) how to assign persistent robots to different mobile depots; 2) what is the optimal order to visit the persistent robots; 3) what is the optimal location to drop off batteries given each

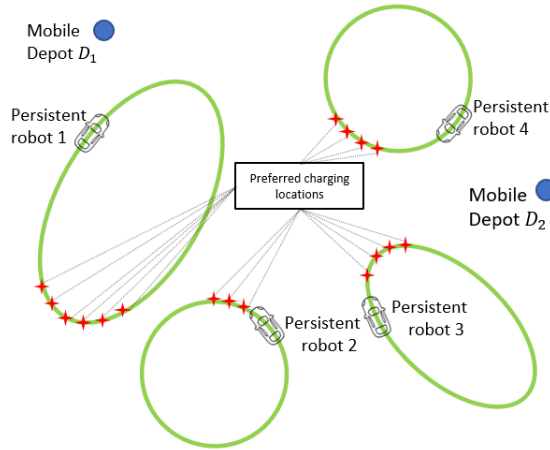persistent robot has multiple preferred locations.

Figure 1.1: Example for the multiple mobile depots route planning problem, with 2 mobile depots and 4 persistent robots. The red asterisks are the preferred battery swapping locations. The green paths are the pre-planned routes for persistent robots.

This paper presents three main contributions. First, a heuristic-based algorithm Multiple Depots Random Select (MDRS) is proposed to efficiently solve the route planning problem for multiple mobile depots for replenishing persistent surveillance robots. The problem is modeled as a GMDTSP on a complete graph. The MDRS has proved to have an analytical constant upper bound in the worst case scenario. Second, a fully decentralized asynchronous MDRS (dec-MDRS) is proposed to provide scalability and robustness. Simulations show that the computation time and the number of messages of dec-MDRS grows linearly as the number of robots increases. Third, a post-processing heuristic (MDRS-IM) is proposed to improve the solution quality further. The simulation results show that MDRS and MDRS-IM perform significantly faster than the state of the art heuristic solver LKH [6] with only a loss of about 10% of solution quality.

1.2   Organization

As is introduced in Section 1.1, this thesis focuses on introducing a decentralized framework for the multiple mobile depots route planning in the application of replenishing persistent surveillance robots. The rest of this thesis is organized as follows.

Chapter 2 discusses several existing approaches to multiple depots route planning problem. Three categories of existing methods are mainly discussed. Optimization-based solutions utilize the mixed integer linear programming, and use varies techniques such as branch and cut to obtain the optimal tours. The disadvantage of optimization-based solutions is that it cannot scale well when the robot number increases given the combinatorial problem. The second category is the heuristic-based algorithm. Although these algorithms obtain the results efficiently, most of them do not have suboptimality guarantees. There exist some worst cases which lead to bad performance. The third category is the transformation-based algorithms. These solutions transfer the

multiple depots problem into an asymmetrical traveling salesman problem, where many existing solutions could be used to solve the transformed problem. However, dummy vertices and edges need to be added to the transformed graph, which is computationally expensive.

Chapter 4 presents the skeleton of our centralized strategy to plan the mobile depots. To make it easier to understand, we first present the single mobile depot scenario in Section 4.1, followed by the generalization to the multiple depots scenario in Section 4.2. For these two scenarios, We prove the upper bound for the worst-case scenario. In Section 4.3, we propose a post-processing heuristic which improve the quality of the tours.

In Chapter 5, we present a decentralized framework as a generalization from Chapter 4. We divide our strategy into several sub-problems discussed in Chapter 4, and develop an asynchronous decentralized algorithm for two sub-components in the framework. We show the correctness of the decentralized algorithms.

Chapter 6 shows the simulation results of the proposed heuristics. Section 6.2 shows an illustrative example, followed by quantitative analysis. We also benchmark our proposed heuristics with the state of the art heuristic LKH. Section 6.3 discusses computational time and number of messages as the robot number increases.

In Chapter 7, we present conclusions and future work.

Chapter 2

Related Work

Recharging or swapping batteries for persistent robots using mobile depots has different formulations in the literature. Different than route planning for multiple mobile depots, [1, 8] formulate the problem as a single mobile depot planning a route for recharge the UAVs. The single mobile depot problem could be solved by the transformation method, mixed integer programming (MIP) or heuristic-based algorithm. In [11, 20, 21], a battery swapping system has been modeled, with the assumption that the swap could happen instantly without the charging duration. However, the battery swapping needs to have fully charged battery in stock which may significantly increase the operation cost.

A recent work [9] has a similar formulation, where a heuristic-based solution performs efficiently, but the effectiveness is only shown based on empirical simulation results instead of an analytical analysis. Also, the algorithm did not admit the periodic structure. Similar to this work, [13] discretize the state space for periodic recharges and cast this problem as a generalized traveling salesman problem (GTSP) on a partitioned directed acyclic graph. However, in our work, the problem is formulated as a GMDTSP on a complete graph. Moreover, the work [13] use the modified noon-bean transformation [17] to transform the problem to a traditional TSP problem, which increases the size of the vertex set.

We formulate the problem as a GMDTSP, where an exact MIP formulation exists [19]. However, the computation time is not acceptable for the real-time application. An efficient and effective heuristic-based algorithm is needed for solving this problem.

Chapter 3

Problem Formulation

Consider n persistent ground robots moving in $\mathbb{R}^2$, following their pre-planned routes to continuously monitor their prescribed areas. To address the long term operation for the persistent surveillance tasks, the persistent robots need to swap battery at some locations before running out of battery due to the limited onboard battery capacity. Assume that each persistent robots $i \in \{1, 2, \cdots, n\}$ can figure out a set of potential preferred battery swapping locations C_i based on their knowledge, which consists of a set of discretized vertices on its pre-planned route. The persistent robot i only accept to swap battery at one of its own preferred locations in C_i since no detour is allowed. Consider k holonomic mobile depots $D = \cup D_i$, $\forall i \in \{1, 2, \cdots, k\}$ carrying unlimited batteries. The mobile depots collectively plan a set of tours to replenish all persistent robots with the minimum total cost. The mobile depots drop off a battery at each stop of the tour so that the persistent robot could swap battery once it travels to this preferred location.

Fig. 1.1 shows an example for the multiple mobile depots route planning problem, with two mobile depots ($k = 2$) and four persistent robots ($n = 4$). The red asterisks are the preferred battery swapping locations, i.e., for a valid battery swap, the mobile depots should only drop off batteries at these red asterisks.

The multiple mobile depots route planning problem could be formulated on a complete undirected graph $G = (V, E, c)$ with vertex set $V = C \cup D$. The cost of an edge $e(p, q) \in E$ is assigned to be the Euclidean distance between vertices $p, q \in V$. The mobile depots collectively compute a set of at most k tours $TOUR_i, i \in \{1, 2, \cdots, k\}$. The tour is defined as a simple circle on G where each mobile depot ends at its starting location, i.e., $TOUR_i^{\text{end}} = TOUR_i^{\text{start}}$. At least one vertex from each vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ is visited by at least one mobile depot. The formal formulation for this problem then becomes,

$$\text{minimize} \quad \sum_{i=1}^{k} c(TOUR_i),$$

$$\text{subject to} \quad TOUR_i^{\text{end}} = TOUR_i^{\text{start}}, \quad i = 1, \cdots, k,$$

$$| \cup_{i=1}^{k} TOUR_i \cap C_j| \geq 1, \quad j = 1, \cdots, n.$$

Note that this problem formulation falls into the category of GMDTSP. When the vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ only contains a single vertex, then this problem degenerate to a MDTSP. When $k = 1$, then this problem degenerate to a GTSP.

Remark 1. *The graph G is a symmetric graph, where $c(p,q) = c(q,p)$ for all vertices $p, q \in V$. For non-holonomic mobile depot, the edge cost may not be symmetric due to the dynamic constraints. In our problem formulation, all mobile depots are holonomic.*

Remark 2. *The complete graph G satisfies the triangle inequality since the edge cost is defined by Euclidean distance between vertices. Based on the problem formulation, it is valid that multiple mobile depots drop off batteries to the same persistent robot. However, in fact, to minimized to total cost for all k tours, each persistent robot will only be served by one and only one mobile depot due to the triangular inequality [3].*

Remark 3. *The current cost only reflect the Euclidean distance between vertices, some other factors such as travel time and battery swapping time could also be associated with the cost as long as the new cost satisfy the triangle inequality.*

Chapter 4

Methodology

4.1   Single Mobile Depot Scenario

As discussed in Chapter 3, a single mobile depot ($k = 1$) degenerates the problem from a GMDTSP into a GTSP. The simplified version will shed light on a more complicated multiple depots scenario. We will first discuss the intuition of the proposed Random Select (RS) heuristics, and then we will describe the steps of RS, and finally a proof of bound will be provided for this heuristics.

The underlying nature for RS is that every vertex in the vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ has correlation with others, given every vertices are the potential charging location for the same persistent robot. This implies if we randomly choose one vertex from the vertex set, the selected vertex could represent the group of vertices. With one vertex in each of the vertex set, the GTSP problem degenerate to a classic symmetric TSP, where many polynomial time approximate solutions exists.

Based on this assumption, Algorithm 1 describe the procedures to executes RS.

Theorem 1. *For any existed TSP heuristic HEUR bounded by constant B, RS(G) is the cost of the route constructed by RS using HEUR on G, and OPT(G) be the cost of the optimal route. Then,*

$$\frac{RS(G)}{OPT(G)} \leq \mathcal{B}(1 + 2d/\rho)$$

where $d := \max_{i=1}^{n}\{\max_{p,q \in C_i} dist(p, q)\}$ is the maximum diameter inter-vertex set, and $\rho := \min dist(p, q)$, where $p \in C_i, q \in C_j, i \neq j$ is the minimum distance intra-vertex sets.

Proof. Given the problem formulation, $G = (V, E, c)$ is a complete undirected graph. Define another complete undirected graph $G' = (P, E', c')$, which is a subgraph of G. Note that the GTSP problem on G is degenerate to a TSP problem on G'.

Since G' is one of the selected subgraph for G, which means using the same TSP heuristic HEUR, $RS(G) \leq HEUR(G')$. Assume OPT(G') is the cost of the optimal route on G'. For any existed TSP heuristic HEUR bounded by $\mathcal{B}$, the cost for HEUR(G') is no greater than $\mathcal{B}\cdot$OPT(G'). Therefore, combine these two inequalities, we have

$$RS(G) \leq \mathcal{B} \cdot OPT(G') \tag{4.1}$$

Algorithm 1: Random Select (RS) heuristics
---
1 Randomly select a single vertex from each of the disjoint vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ to form a subset $P \subseteq V$, and $|P| = k$.
2 Use any existed polynomial time TSP heuristic HEUR to solve the tour T on the selected subset P. Return T as an approximation for the problem.

The obtained optimal TSP sequence on G' is some permutation on vertex set P, and define this optimal sequence as S. Define S' as another sequence on P, which uses the vertex set visitation sequence on OPT(G). We could imply that $OPT(G') = cost(S) \leq cost(S')$. Since S' is the optimal group visitation sequence on set P, using the triangular inequality, we could have $cost(S') \leq OPT(G) + 2n \cdot d$. Based on the definition on ρ, we have $OPT(G) \geq n \cdot \rho$. Combine these inequalities, we have

$$OPT(G') \leq (1 + 2d/\rho)OPT(G) \tag{4.2}$$

The theorem is proved by combining the inequality (4.1) and (4.2). $\quad\square \quad\quad\quad \square$

Theorem 2. *If the TSP heuristic obtained a worst case route with its upper bound $\mathcal{B}$, then there exist a worst case scenario that*

$$\frac{RS(G)}{OPT(G)} = \mathcal{B}(1 + 2d/\rho)$$

Proof. This could be shown by constructing an example. Assume four vertices are co-aligned on a line with equal distance l separate apart. Separate the four vertices in the middle to two vertex sets, with two vertices each. In this setting, $d = \rho = 2l$. We could verify the optimal cost $OPT(G) = l$. Choose the points at the ends for both sets. Given the chosen TSP heuristic obtained a worst case route, $RS(G) = \mathcal{B} \cdot 6l$. Thus, $RS(G)/OPT(G)$ has the upper bound equal to $\mathcal{B}(1 + 2d/\rho)$. $\quad\square \quad\quad\quad \square$

Remark 4. *$d = 0$ is a special case, where the vertices in each vertex set degenerate to a single point. This implies the GTSP degenerates to a TSP. The derived bound could be justified since the bound becomes $\mathcal{B}$ for the chosen TSP heuristic HEUR.*

Remark 5. *The proposed algorithm could be run in $O(n^2)$ time, where n is the number of persistent robots. Note that this is independent to $|C|$, the number of potential charging locations.*

Remark 6. *The upper bound for RHS does not depend on a specific TSP polynomial algorithm. Some popular heuristics includes Christofides algorithm with $\mathcal{B} = 1.5$, nearest neighbor heuristics with $\mathcal{B} = 2$ etc.*

4.2 Multiple Depot Random Select and Analysis

We present a heuristic-based Multiple Depot Random Select (MDRS) in Algorithm 2. Before dive deep into the technical details, the design philosophies are two-fold. First, the multiple mobile depots could reduce to a single depot scenario by assigning battery drop off tasks to each

mobile depot so that the total route cost is minimized. Second, a single vertex could be randomly selected to represent this group of vertices to reduce complexity. The underlying nature is that every vertex in the vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ correlates with others, given every vertex are the preferred battery swapping locations for the same persistent robot.

With the two design philosophies in mind, algorithm 2 first randomly select a single vertex from each disjoint vertex set $C_i, \forall i \in \{1, 2, \cdots, n\}$ to form a subset $P \subseteq C$ on line 1 to 2, and $|P| = n$. Inspired by the well-known minimum spanning tree (MST) heuristic in TSP and multiple TSP (mTSP), line 3 build a weighted undirected graph on vertex set P, D and a dummy node v_{dummy}, where line 10 to 14 describes how we define the edges with cost. Prim's algorithm is used to find the minimum spanning tree (MST) on the newly built graph. The MST has a total of $|P| + |D| + 1$ nodes if we includes the dummy node v_{dummy}. For all k depot nodes $d_i \in D$, line 6 break the MST into k subtrees T_i rooted at d_i. A depth first search is executed on each subtree T_i, starting from the corresponding root d_i. The visitation sequence S is recorded. Line 8 bypass the duplicated vertices in sequence S, and $TOUR_i$ is returned.

Algorithm 2: Multiple depot random select (MDRS) heuristics

    Inputs : $C := \cup C_i, \forall i \in \{1, 2, \cdots, n\}, D$
    Outputs: k tours $TOUR_i$ for $i = 1, \cdots, k$
1 Randomly select a vertex $v_i \in C_i, \forall i \in \{1, 2, \cdots, n\}$
2 $P := \cup v_i, \forall i \in \{1, 2, \cdots, n\}$
3 G' = `buildGraph`(P, D)
4 M = `mst`(G')
5 **for** $i = 1 : k$ **do**
6   Break M into subtree T_i rooted at $d_i \in D$
7 S = `traverse`(T_i) using depth first search
8 $TOUR_i$ = `bypass`(S) by eliminating previously occurred vertices
9 **end**
    Function `buildGraph`(P, D)
10   Set $V = P \cup D \cup v_{dummy}$
11   Add $e(v_{dummy}, i)$ to E with $c = 0, \forall i \in D$
12   Add $e(i, j)$ to E with $c = dist(i, j), \forall i, j \in P$
13   Add $e(i, j)$ to E with $c = dist(i, j), \forall i \in D, j \in P$
14   **return** undirected graph $G = (V, E, c)$

Theorem 3. *MDRS(G) is the cost of the route constructed by MDRS on G, and OPT(G) is the cost of the optimal route. Then,*

$$\frac{MDRS(G)}{OPT(G)} \leq 2(1 + 2d/\rho)$$

where $d := \max_{i=1}^{n}\{\max_{p,q \in C_i} dist(p, q)\}$ is the maximum inter-vertex set distance, and $\rho := \min dist(p, q)$, where $p \in C_i, q \in C_j, i \neq j$ is the minimum intra-vertex distance.

Proof. In the multiple depots scenario, the minimum spanning tree heuristic is used, where $\mathcal{B} = 2$. The rest of the proof is similar to Theorem 1. ☐

11

Theorem 4. *MDRS returns a feasible set of tours satisfying the problem formulation. In other words, edges $e(v_{dummy}, i)$ for all $i \in D$ are edges for the MST on G' rooted at the dummy node v_{dummy}. Mobile depot i has a sub-tree $T_i, i \in \{1, 2, \cdots, k\}$, where $|T_i \cap T_j| = \varnothing$, and $|T_i \cup T_j| = C$, where $i, j \in \{1, 2, \cdots, k\}, i \neq j$.*

Proof. If we apply the Prim's greedy algorithm, then all mobile depots will be connected to the dummy node v_{dummy} since the edge cost between them is the smallest, zero. With the connection between the dummy node and all mobile depots, any vertex in C will not connect to dummy node since the edge cost between them is infinite. This implies $|T_i \cup T_j| = C$. Also, the sub-tree rooted at the mobile depots will not intersect with each other. Otherwise, a circle will be formed which contradict with the concept of tree. This implies $|T_i \cap T_j| = \varnothing$. □

4.3   Improvement Heuristics on Tours

In this section, a centralized improvement heuristics (MDRS-IM) is proposed to increase solution quality. MDRS randomly select a vertex from each of the vertex set to represent the whole set, the constructed tours could be used as an near-optimal vertex set visitation sequence to further optimize and improve the tour.

MDRS generates a set of tours for each mobile depots, where each mobile depot knows a set of persistent robots need to be served with a visitation sequence. Define a map $\mathcal{M} := \mathbb{R}^3 \to \mathbb{R}$, which maps the randomly selected vertex $v_i \in C_i$ to the corresponding vertex set index $i \in \{1, 2, \cdots, n\}$. Given a tour T, define $T(j)$ as the j^{th} vertex in T. Define two tour T_1 and T_2 are equal if and only if 1) $|T_1| = |T_2|$, and 2) $T_1(j) = T_2(j)$ for all $j = \{1, 2, \cdots, |T_1|\}$. Otherwise, $T_1 \neq T_2$. Define $TSP(T)$ as a tour obtained on tour T using TSP algorithm.

Algorithm 3 describes the MDRS-IM procedure to improve the solution quality. The input T_{d_i} is the tour generated by MDRS for mobile depot d_i. The output is the improved tour T_{IM} with a better quality for this mobile depot. The algorithm build a graph $G_{IM} = (V_{IM}, E_{IM}, c_{IM})$ from Line 5 to Line 20 based on the input group visitation sequence. Vertices in neighbor groups are fully connected, and all vertices in the first and last group are connected to the dummy vertex s and t. The shortest path search such as A* return the optimal tour given the current group visit sequence. Line 4 is for continuous improvement based on the previous post-processing result until the improved heuristics get the local minimum.

12

Algorithm 3: Improvement Heuristics on MDRS (MDRS-IM) for mobile depot d_i

---

Inputs : The original tour T_{d_i} from MDRS
Outputs: The improved tour T_{IM}

1 Initialize $G_{IM} = (V_{IM}, E_{IM}, c_{IM})$ with $V_{IM} = V$,
2 $T_{IM} = \varnothing$
3 $L := |T_{d_i}|$
4 **while** $T_{IM} == \varnothing \ OR \ TSP(T_{IM}) \neq T_{IM}$ **do**
5   **for** $j = 0, 1, \cdots, L - 1$ **do**
6   **if** $j = 0$ **then**
7   Insert dummy node s, $V_{IM} = V_{IM} \cup s$
8   Connect $E_{IM}(s, v), \forall v \in C_{\mathcal{M}(T(1))}$
9   Set $c_{IM}(s, v) = 0, \forall v \in C_{\mathcal{M}(T(1))}$
10   **end**
11   **else if** $j = L - 1$ **then**
12   Insert dummy node t, $V_{IM} = V_{IM} \cup t$
13   Connect $E_{IM}(v, t), \forall v \in C_{\mathcal{M}(T(L))}$
14   Set $c_{IM}(v, t) = 0, \forall v \in C_{\mathcal{M}(T(L))}$
15   **end**
16   **else**
17   Connect $E_{IM}(u, v)$ for all $u \in C_{\mathcal{M}(T(j))}$, $v \in C_{\mathcal{M}(T(j+1))}$
18   Set $c_{IM}(u, v) = dist(E_{IM}(u, v))$
19   **end**
20   **end**
21      $T_{IM} \leftarrow$ shortest path from dummy node s to t
22 **end**

Chapter 5

A Decentralized Framework

In this section, we will introduce a decentralized framework to implement our approach which provides scalability and robustness. Our proposed MDRS consists of several sub-components, including construction of the minimum spanning tree, traversal of the tree using depth first search and by-pass the repeated vertices to form a tour. The goal for the fully decentralized MDRS (dec-MDRS) heuristic is to obtain the solution by asynchronously passing messages among all persistent robots and mobile depots.

The dec-MDRS works as follows. 1) Each persistent robot will randomly select a battery swapping location from its preferred locations set. 2) GHS algorithm [4], a well known asynchronous distributed algorithm, will be used to construct minimum spanning tree. The graph G' is the same as in the centralized version. Each robot knows the adjacency list with weights for itself. Assume the graph G' has n nodes and m edges, GHS algorithm runs in $O(n \log_2 n)$ time and using at most $O(m + n \log_2 n)$ messages. The output for this step is a minimum spanning tree, which is represented as the adjacency list of vertices adj_{MST} stored in each robot. 3) With the MST built, the last step is to traverse the tree and bypass the repeated vertices to form a tour. We present the last step in Algorithm 4.

Algorithm 4 describes the traversal and bypassing for robot r_i, where the robot could be both mobile depots and persistent robots. Before this algorithm starts, r_i knows its adjacency list of vertices adj_{MST}, which is a list of its neighbor vertices on the built MST. Note that adj_{MST} could be unsorted. Also, if the robot is a mobile depot, then it knows the Boolean $first$ is true, meaning mobile depots are the robots that kick off the algorithm. The output is the next vertex $next$ on the constructed $TOUR$. Each robot r_i has three internal variables to keep track its own status, $count$, $started$ and $visited$. $count$ is an integer variable which points the position of the adj_{MST}, which is initialized to be 1, meaning the first element in adj_{MST} has been pointed. Boolean $visited$ means whether it has been visited before, which is initialized to be $false$. Another Boolean $started$ means whether this robot has been woken up, which is initialized to be $false$. The messages has the same protocol, where $msg = \langle type, src, [optional]other \rangle$. Robots asynchronously send and receive three message types to each other with the known protocol to construct the tour. When a robot is been visited for the first time, it marks itself as $visited$. Line 13 to 17 find the its parent in adj_{MST} and move it to the last position in the list. Once the counter hit the end of adj_{MST}, the robot traces back to its parent as in the depth first search. The robot then increases the counter and do the same traversal to the next vertex in adj_{MST}. If the

robot is visited before, this means the robot needs to be bypassed. This is achieved by sending messages to the two related bypass neighbors. Assume that (r_j, r_i) and (r_i, r_k) is edges on the built MST, and robot r_i is bypassed. r_i will send message to r_j telling r_k is next vertex $next$ on the constructed tour, and send another similar message to r_k with the vertex r_j. The algorithm is terminated once the counter $count$ points out of adj_{MST}.

Algorithm 4: Distributed Traverse and Bypass Algorithm, for robot r_i

    Inputs : Adjacency vertex list adj_{MST}, Boolean $first$
    Outputs: Next vertex $next$ on $TOUR_i$
1   Initialize $count = 1, visited = false, started = false$
    When receiving no message:
2 | **if** $started = false$ **then**
3 | | $started = true$
4 | | **if** $first = true$ **then** `proceed(-1)`
5 | **end**
    When receiving message $\langle A, src\rangle$**:**
6 | **if** $count > |adj_{MST}|$ **then** Terminate
7 | **if** $visited = false$ **then** `proceed(src)`
8 | **else** `bypass(src)`
    When receiving message $\langle B, src, vertex\rangle$**:**
9 | $next = vertex$
    When receiving message $\langle C, src, vertex\rangle$**:**
10 | **if** $count > |adj_{MST}|$ **then** Terminate
11 | **if** $visited = false$ **then** `proceed(src)`
12 | **else** `bypass(vertex)`
    Function `proceed(src)`
13 | $visited = true$
14 | **if** $src \neq -1$ **then**
15 | | Delete $adj_{MST}(adj_{MST} == src)$
16 | | Insert src to the end of adj_{MST}
17 | **end**
18 | $count = count + 1$
19 | $next = adj_{MST}(count)$
20 | `sendmsg(`$\langle A, r_i\rangle$`, `$adj_{MST}(count)$`)`
    Function `bypass(`r_j`)`
21 | $r_k = adj_{MST}(count)$
22 | `sendmsg(`$\langle B, r_i, r_k\rangle$`, `$r_j$`)`
23 | `sendmsg(`$\langle C, r_i, r_j\rangle$`, `$r_k$`)`
24 | $count = count + 1$
    Function `sendmsg(`msg, dst`)`
25 | Send message msg to destination dst

Theorem 5. *Algorithm 4 eventually terminates and construct the correct set of tours.*

16

Proof. Algorithm 4 traverses the undirected edge at most twice in adj_{MST}, which means finite number of edges will eventually leads to termination. In other words, since the counter $count$ always increases while the length of adj_{MST} is unchanged, the counter will eventually meet the termination condition.

In terms of correctness, Algorithm 4 starts on the set of mobile depots. Once the robot is been visited for the first time, the robot marks itself as visited and remains the status. The robot will also mark its parent during the first visit. This ensures the robot can trace back once all other neighbors has been visited. All later visitation on those robots who has already been visited will be bypassed. □

Chapter 6

Simulation Results

6.1   Simulation Setup

The proposed algorithms are implemented in C++ with an open source template graph library LEMON[1]. The library provides efficient implementations of common data structures and graph algorithms. Simulation results were computed on a laptop running a 64 bit Ubuntu 16.04 operating system with an Intel® Core™i7-8550U CPU @1.80GHz x 8 and 16GB of RAM.

6.2   Centralized Heuristics

Fig. 6.1 illustrates an example of the multiple mobile depots route planning problem. Here we have $k = 2$ mobile depots shown in black diamonds. $n = 9$ persistent robots with their pre-planned paths are shown in green dotted lines, with $|C| = 70$ total potential charging locations shown in red asterisks on the green path. MDRS-IM is used to construct the set of tours, with the solution shown in the blue solid line, where mobile depots end the tour at their initial positions. It could be verified that the obtained solution is global optimal in terms of minimizing the total costs. The problem is solved in less than 0.1s by using MDRS-IM.

Fig. 6.2 investigates the relation between the solution quality ratio versus d/ρ. The solution quality ratio is defined as the cost obtained by MDRS-IM versus the known optimal cost. The optimal cost is calculated using the MIP optimization formulation proposed by [19]. The MIP is solved by IBM CPLEX®. We run 200 instances in total to compile Fig. 6.2, which shows the mean, standard deviation and the extremes for various d/ρ. The instances are created as follows. First, select 20 instances in TSPLIB[2]. Then, for each point on the TSP instance, we use 2D Gaussian distribution to generate 20 points to form a corresponding vertex set. Each vertex set is treated as the set of preferred charging locations for a persistent robot. We randomly distribute $k = 2$ depots on the constructed graph. We could obtain various d/ρ by controlling the μ and covariance matrix for the Gaussian distribution. Fig. 6.2 shows that the mean of cost ratio will

---

[1]Open source library LEMON, `https://lemon.cs.elte.hu/trac/lemon`.
[2]TSPLIB instances collected by Gerhard Reinelt, `https://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/tsp/`.

Figure 6.1: An illustrative example showing the constructed tours (blue solid line) using MDRS-IM. Black diamonds are mobile depots, green dotted lines are the pre-planned routes for persistent robots. Red asterisks are the preferred battery swapping locations. It could be verified that this path is optimal in terms of minimizing the total costs. It could also be verified that dec-MDRS also produce the same tour set.

gradually increase with the increment of d/ρ. In most cases, the algorithm obtains a cost ratio within 1.2.


Figure 6.2: Relation between the solution quality ratio versus d/ρ.

Fig. 6.3 show how the increment of persistent robot number will affect the simulation time for centralized algorithms. The simulation benchmarks with 1) optimization method [19], which solve the exact solution based on branch and cut, 2) transformation methods [12, 13] which

transform the problem to a classic TSP problem, and then use the state of the art LKH heuristic to solve the approximate solution. In this simulation, each persistent robot has 10 preferred battery swapping locations, which means each vertex set has size of 10. We randomly put $k = 5$ mobile depots in this problem. The simulation results show the optimization method grow exponentially with the increment of robot number. It also illustrates that MDRS-IM algorithm reduces the simulation time significantly, which is about 10 times less than the LKH method, and 150 times less than the optimization method.
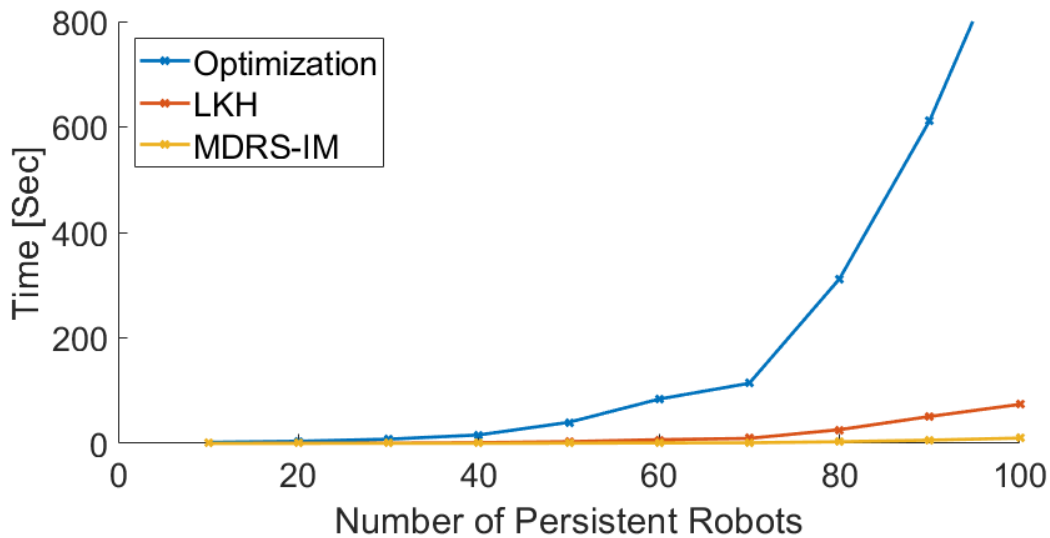

Figure 6.3: Comparison of computation time between three methods: optimization, LKH and MDRS-IM.

Table 6.1 shows the time and solution quality comparison between LKH with MDRS and MDRS-IM on GTSPLIB[3][5] instances. The Value and Time for LKH method are directed benchmarked using GLKH version 1.0 from [7]. For simplicity, all instances are route planning for single mobile depot. Both MDRS and MDRS-IM computes significantly faster than the LKH solver for these very large instances with little solution quality loss.

21

Table 6.1: Time and solution quality comparison between LKH with MDRS and MDRS-IM on GTSPLIB.

Name	Opt	LKH			MDRS (Our)			MDRS-IM (Our)		
		Value	Time [s]	Ratio	Value	Time [s]	Ratio	Value	Time [s]	Ratio
99d493	20023	20023.4	170.5	1.00	26168	0.0	1.31	22504	0.2	1.12
132d657	22498	22498	490.8	1.00	28136	0.0	1.25	25333	0.2	1.13
134gr666	163028	163028	162.3	1.00	204543	0.0	1.25	184782	0.2	1.13
145u724	17272	17272	145.4	1.00	21640	0.0	1.25	19285	0.2	1.12
157rat783	3262	3262.9	764.4	1.00	4172	0.0	1.28	3646	0.2	1.12
200dsj1000	9187884	9187884	794.4	1.00	11954184	0.1	1.30	10310418	0.2	1.12
201pr1002	114311	114311	164.8	1.00	150757	0.1	1.32	129730	0.3	1.13
207si1032	22306	22306	1202.5	1.00	28315	0.1	1.27	25135	0.2	1.13
212u1060	106007	106029.5	2054.9	1.00	136547	0.1	1.29	116941	0.3	1.10
217vm1084	130704	130704	209.0	1.00	169315	0.2	1.30	149974	0.3	1.15

6.3 Decentralized Heuristics

Fig. 6.4 shows our proposed dec-MDRS approach has good scalability. As the number of robots increases, the messages exchanged and the compulation time grows linearly. The correctness of dec-MDRS could also be verified in the example in Fig. 6.1, where the decentralized algorithm return the same set of tours as the centralized algorithm.
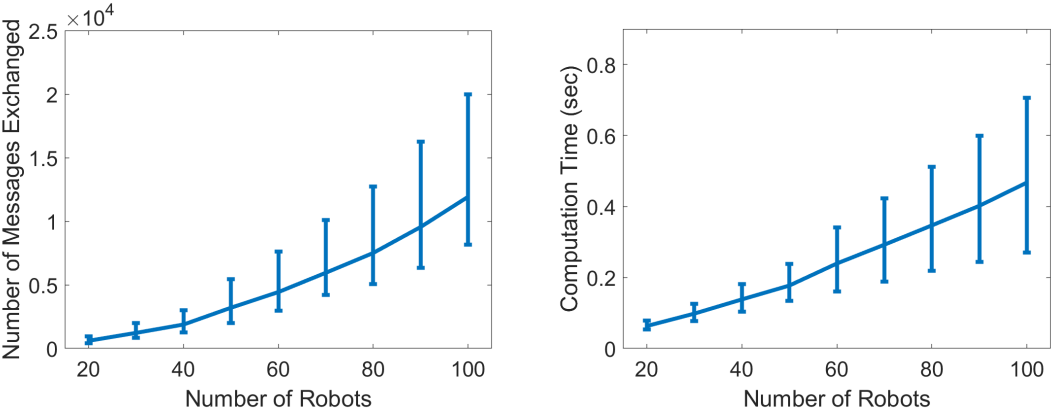

Figure 6.4: Results from our proposed dec-MDRS approach. Left figure shows number of messages exchanged to construct tours using dec-MDRS. The error bar shows the maximum and minimum number of messages exchanged. Right figure shows computation time of constructing dec-MDRS.

Chapter 7

Conclusion and Future Work

In this paper, we consider the problem for replenishing persistent robots using mobile depots. We formulate this problem as a GMDTSP on a complete graph. An efficient centralized heuristic-based algorithm MDRS is proposed. MDRS captures the nature for GMDTSP, where the vertices in the same vertex set have a close correlation, and thus a single vertex has a good representation for the whole vertex set. This nature dramatically reduces the problem complexity and computation time. dec-MDRS is then proposed to provide scalability and robustness for the centralized version of the algorithm. We also propose a post-processing heuristic (MDRS-IM) to improve the solution quality further.

In the future, we may incorporate temporal planning in the multiple mobile depots route planning problem formulation. In the current formulation, no temporal information is considered, and the persistent robots might need to wait on its preferred battery swapping location until the mobile depot deliver the battery. If the temporal information is considered, we could minimize the waiting time for persistent robots to swap the battery. We could also assign different replenishment priority for for different persistent robots. The other future work is to explore the optimal number and placement for mobile depots when initialization.

Bibliography

[1] J Wesley Barnes, Victor D Wiley, James T Moore, and David M Ryer. Solving the aerial fleet refueling problem using group theoretic tabu search. *Mathematical and computer modelling*, 39(6-8):617–640, 2004. 2

[2] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004. 1.1

[3] Matteo Fischetti, Juan José Salazar González, and Paolo Toth. A branch-and-cut algorithm for the symmetric generalized traveling salesman problem. *Operations Research*, 45(3): 378–394, 1997. 2

[4] Robert G Gallager, Pierre A Humblet, and Philip M Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Transactions on Programming Languages and systems (TOPLAS)*, 5(1):66–77, 1983. 5

[5] Gregory Gutin and Daniel Karapetyan. A memetic algorithm for the generalized traveling salesman problem. *Natural Computing*, 9(1):47–60, 2010. 6.2

[6] Keld Helsgaun. An effective implementation of the lin–kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106–130, 2000. 1.1

[7] Keld Helsgaun. Solving the equality generalized traveling salesman problem using the lin–kernighan–helsgaun algorithm. *Mathematical Programming Computation*, 7(3):269–287, 2015. 6.2

[8] Zhipu Jin, Tal Shima, and Corey J Schumacher. Optimal scheduling for refueling multiple autonomous aerial vehicles. *IEEE Transactions on Robotics*, 22(4):682–693, 2006. 2

[9] Nitin Kamra, TK Satish Kumar, and Nora Ayanian. Combinatorial problems in multirobot battery exchange systems. *IEEE Transactions on Automation Science and Engineering*, 15 (2):852–862, 2018. 2

[10] Jonghoe Kim and James R Morrison. On the concerted design and scheduling of multiple resources for persistent uav operations. *Journal of Intelligent & Robotic Systems*, 74(1-2): 479–498, 2014. 1.1

[11] Danny Lee, Joe Zhou, and Wong Tze Lin. Autonomous battery swapping system for quadcopter. In *2015 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 118–124. IEEE, 2015. 2

[12] Neil Mathew, Stephen L Smith, and Steven L Waslander. A graph-based approach to multi-

robot rendezvous for recharging in persistent tasks. In *2013 IEEE International Conference on Robotics and Automation*, pages 3497–3502. IEEE, 2013. 6.2

[13] Neil Mathew, Stephen L Smith, and Steven L Waslander. Multirobot rendezvous planning for recharging in persistent tasks. *IEEE Transactions on Robotics*, 31(1):128–142, 2015. 2, 6.2

[14] AJS McGonigle, A Aiuppa, G Giudice, G Tamburello, AJ Hodson, and S Gurrieri. Unmanned aerial vehicle measurements of volcanic carbon dioxide fluxes. *Geophysical research letters*, 35(6), 2008. 1.1

[15] Derek Mitchell, Micah Corah, Nilanjan Chakraborty, Katia Sycara, and Nathan Michael. Multi-robot long-term persistent coverage with fuel constrained robots. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1093–1099. IEEE, 2015. 1.1

[16] Nikhil Nigam and Ilan Kroo. Persistent surveillance using multiple unmanned air vehicles. In *2008 IEEE Aerospace Conference*, pages 1–14. IEEE, 2008. 1.1

[17] Charles E Noon and James C Bean. An efficient transformation of the generalized traveling salesman problem. *INFOR: Information Systems and Operational Research*, 31(1):39–44, 1993. 2

[18] Ethan Stump and Nathan Michael. Multi-robot persistent surveillance planning as a vehicle routing problem. In *2011 IEEE International Conference on Automation Science and Engineering*, pages 569–575. IEEE, 2011. 1.1

[19] Kaarthik Sundar and Sivakumar Rathinam. Generalized multiple depot traveling salesmen problempolyhedral study and exact algorithm. *Computers & Operations Research*, 70:39–55, 2016. 2, 6.2, 6.2

[20] Koji AO Suzuki, Paulo Kemper Filho, and James R Morrison. Automatic battery replacement system for uavs: Analysis and design. *Journal of Intelligent & Robotic Systems*, 65 (1-4):563–586, 2012. 2

[21] Kurt A Swieringa, Clarence B Hanson, Johnhenri R Richardson, Jonathan D White, Zahid Hasan, Elizabeth Qian, and Anouck Girard. Autonomous battery swapping system for small-scale helicopters. In *2010 IEEE International Conference on Robotics and Automation*, pages 3335–3340. IEEE, 2010. 2