# Miniature Robot Path Planning for Bridge Inspection: Min-Max Cycle Cover-Based Approach 

Michael Lin and Richard J. La<br>Department of Electrical and Computer Engineering<br>and the Institute for Systems Research<br>University of Maryland, College Park<br>$\{m \operatorname{lin} 1025$, hyongla $\}$ @umd.edu


#### Abstract

We study the problem of planning the deployments of a group of mobile robots. While the problem and formulation can be used for many different problems, here we use a bridge inspection as the motivating application for the purpose of exposition. The robots are initially stationed at a set of depots placed throughout the bridge. Each robot is then assigned a set of sites on the bridge to inspect and, upon completion, must return to the same depot where it is stored.

The problem of robot planning is formulated as a rooted minmax cycle cover problem, in which the vertex set consists of the sites to be inspected and robot depots, and the weight of an edge captures either (i) the amount of time needed to travel from one end vertex to the other vertex or (ii) the necessary energy expenditure for the travel. In the first case, the objective function is the total inspection time, whereas in the latter case, it is the maximum energy expenditure among all deployed robots. We propose a novel algorithm with approximation ratio of $5+\epsilon$, where $0<\epsilon<1$. In addition, the computational complexity of the proposed algorithm is shown to be $O\left(n^{2}+2^{m-1} n \log (n+k)\right)$, where $n$ is the number of vertices, and $m$ is the number of depots.


## I. Introduction

With aging infrastructure, ensuring the safety of existing civil structures, such as bridges, roads and tunnels, is becoming an important societal challenge. Inadequate monitoring of infrastructure can result in major incidents, such as the collapse of bridges, e.g., the failure of Ponte Morandi bridge in Italy in August 2018, which killed 43 people. According to a 2018 U.S. Federal Highway Administration (FHWA) report, more than 47,000 bridges are deemed to be in "poor" condition out of approximately 616,000 bridges, and nearly a half of all bridges are found to be in "fair" condition [1].
Unfortunately, many segments of a bridge are not easily accessible, making it difficult for human inspectors to perform frequent inspections. As a result, many bridges are not inspected frequently enough to maintain their structural health and safety, which is reflected in the U.S. FHWA report, thereby raising the possibility of suffering another major bridge collapse in the future.

Rapid advances in robotics technologies make it possible to employ small mobile robots to help with the inspection of different types of structures, including bridges. These robots will likely be battery powered to improve their mobility, thereby limiting their ranges and tasks that they can perform before their battery needs to be recharged. For this reason, in
order to complete a bridge inspection as quickly as possible, it is important to take into account their energy constraints when employing the robots for inspection.


Fig. 1: Bridge inspection by mobile robots.
We study a robot planning problem in which a group of battery-powered mobile robots are stored and recharged at a set of depots, and are utilized for inspecting a set of sites. These sites could, for instance, represent various points on a bridge that need to be inspected (e.g., joints) by a robot. This is illustrated in Fig. 1. In the figure, vertices represent the set of sites on a bridge to be inspected by the robots, and edges show paths robots can take to move between various points on the bridge.

In our problem, we are interested in determining, for each robot, (i) a depot where it is to be stored (and recharges its battery) and (ii) a set of sites for the robot to inspect. We require that, upon completing the inspection of all assigned sites, the robot must return to the same depot where it is stored.

The problem is formulated as a rooted min-max cycle cover problem. Each cycle in the cycle cover, which is rooted at a depot, is assigned to a robot and determines a subset of sites that the robot must inspect as well as the depot at which the robot is to be stationed. We propose a new algorithm for the rooted min-max cycle cover problem with approximation ratio of $5+\epsilon(0<\epsilon<1)$.

The rest of the paper is organized as follows. Section $\Pi$ summarizes some of the most closely related studies in the literature. Section III presents the model and formulation of our problem. Our proposed algorithm is described in Sections IV and V We discuss the complexity of the proposed algorithm in Section VI

## II. Related Work

The well-known traveling salesman problem (TSP) is NPhard. Since the TSP is a special case with one depot and one agent, the general rooted min-max cycle cover problem is also NP-hard. For this reason, researchers proposed approximation algorithms for related problems over the years.

Even et al. [2] studied the min-max cycle cover problem in the context of nurse station location problem. They proposed algorithms for both rooted and unrooted (or rootless) min-max tree cover problems with approximation ratio of $4+\epsilon(\epsilon>0)$. This provides an $8+\epsilon$ approximation algorithm for the rooted min-max cycle cover problem.

In a closely related study, Arkin et al. [3] also provided a $4+\epsilon$ approximation algorithm for the unrooted min-max tree cover problem. Khani and Salavatipour [4] improved the approximation ratio to $3+\epsilon$ for unrooted min-max tree cover problem, which in turn yields an approximation ratio of $6+\epsilon$ for the unrooted min-max cycle cover problem.

Rather than starting with a tree cover problem, Jorati [5] directly studied both rooted and unrooted min-max cycle cover problems and proposed algorithms with approximation ratio of $5 \frac{1}{3}+\epsilon$ and $7+\epsilon$ for unrooted and rooted problem, respectively. Similarly, in an independent study, Xu et al. [6] investigated the same cycle cover problems and proposed algorithms with approximation ratio of $5 \frac{1}{3}+\epsilon$ and $6 \frac{1}{3}+\epsilon$ for unrooted and (uncapacitated) rooted min-max cycle cover problems, respectively ${ }^{1}$ In addition, Yu and Liu [7] proposed algorithms with improved approximation ratio of $5+\epsilon$ and $6+\epsilon$ for unrooted and rooted min-max cycle cover problems, respectively, by utilizing the well-known Christofides algorithm [8] for the TSP problem.

We point out that there are other studies on special cases of the cycle cover problems with better approximation ratios, e.g., [9], [10], [11], [12], [13]. For example, Frederickson et al. [9] proposed a $2 \frac{1}{2}+\epsilon$ approximation algorithm for single depot case. Finally, Xu and Wen [14] proved that, unless $P=$ $N P$, there exists no polynomial time $\left(1 \frac{3}{17}-\epsilon\right)$-approximation algorithm for the min-max cycle cover problem with a single root. This result is generalized by Xu et al. [15] who showed that there does not exist a polynomial time algorithm for the unrooted and rooted min-max cycle cover problems with an approximation ratio less than $1 \frac{1}{3}$ unless $P=N P$.

It is worth mentioning that, although our rooted min-max cycle cover formulation is inspired by [6] and [7], our model is slightly different: they only required the union of cycles to cover all vertices, except for depots. In other words, the depots need not be covered by the cycle cover. In our problem, however, we require that all vertices, including depots, be covered by the cycle cover.

Instead of investigating the unrooted min-max cycle cover problem and applying the algorithm on the rooted version as in many, if not most, of related studies, we directly tackle the rooted min-max cycle cover problem. We propose an

[^0]approximation algorithm that runs in polynomial time with a fixed number of depots. The approximation ratio of $5+\epsilon$ is comparable to the state-of-the-art algorithm for unrooted problem [7]. Moreover, even though our formulation is somewhat different, our approximation ratio is better than the previous best algorithm for rooted cycle cover problem [7].

## III. Model and Formulation

We formulate the problem of robot planning as a rooted min-max cycle cover problem on a complete undirected graph $G=(V, E)$ : the vertex set $V$ consists of both (a) the sites to be inspected by the robots and (b) the depots where the robots are stored, and each undirected edge $e$ in $E$ has a weight associated with it. The goal of the problem is to find a set of cycles subject to following two constraints: (i) each cycle is rooted at a depot in that it starts and ends at the same depot, and (ii) the union of all cycles covers all vertices in $V$. The interpretation is that each cycle found in the problem is assigned to a unique robot and determines the set of sites to be visited by the robot as well as the depot at which the robot will be stored.
The weights assigned to the edges depend on our objective. We consider two different choices of weights. Obviously, it is possible to take a weighted sum of the two choices.

- Case I: The weight of an edge models the minimum amount of time needed to travel from one end vertex to the other vertex. In this case, the overall cost for a robot is equal to the total travel time for the robot to visit all assigned sites in the cycle and come back to the depot. The objective of our problem is to minimize the total amount of time needed for inspection, which is equivalent to minimizing the maximum cost among all robots ${ }^{2}$
- Case II: In the second case, the edge weight indicates the necessary energy expenditure for the travel from one end vertex to the other. The goal is then to minimize the maximum energy expenditure among all robots for the given assignments.
The second formulation allows us to determine whether or not the battery-powered robots can perform the inspection without having to recharge; if the optimal value of the optimization problem is larger than the amount of battery energy available to the robots, it suggests that recharging is necessary for some robots.

These edge weights can be obtained from the weights we can estimate from the physical structure. For example, Fig. 1 tells us the available paths between various points on the bridge. Once we estimate the weights of these available paths on the bridge, we can use a well-known algorithm, such as Dijkstra's algorithm [16], to compute the weights of the shortest paths between any pair of vertices in $V$.

[^1]
## A. Rooted min-max cycle cover problem

Suppose (i) $G=(V, E)$ is a complete undirected graph described earlier, (ii) $w: E \rightarrow \mathbb{R}_{+}:=[0, \infty)$ is an edge weight function, (iii) $D \subseteq V$ is a set of depots (where the robots are stored), and (iv) $k$ is a positive integer (which is equal to the number of robots in our problem). Given a subgraph $G^{\prime}$ of $G$, e.g., a cycle or a tree, let $\mathcal{E}\left(G^{\prime}\right)$ and $\mathcal{V}\left(G^{\prime}\right)$ be the set of edges and the set of vertices, respectively, in $G^{\prime}$.

Denote by $\mathcal{C}^{k}$ the set of all possible edge-disjoint, rooted cycle covers of $V$ with at most $k$ cycles. In other words, an element $\mathscr{C}=\left\{C_{1}, \ldots, C_{q}\right\}$ of $\mathcal{C}^{k}$ consists of $q$ cycles in $G$ satisfying the following:

$$
\begin{array}{ll}
\text { c1. } & q \leq k \\
\text { c2. } & \mathcal{E}\left(C_{i}\right) \cap \mathcal{E}\left(C_{j}\right)=\emptyset \text { for all } i \neq j \\
\text { c3. for every cycle } C_{j} \in \mathscr{C},\left|\mathcal{V}\left(C_{j}\right) \cap D\right|=1 \text {; and } \\
\text { c4. } \cup_{j=1}^{q} \mathcal{V}\left(C_{j}\right)=V
\end{array}
$$

Given a cycle cover $\mathscr{C}$ in $\mathcal{C}^{k}$, we denote the number of cycles in $\mathscr{C}$ by $\xi(\mathscr{C})$.

The rooted min-max cycle cover problem we adopt is the following optimization problem $\sqrt[3]{3}$

## Rooted Min-Max Cycle Cover Problem

$$
\begin{equation*}
\operatorname{minimize}_{\mathscr{C} \in \mathcal{C}^{k}} \max _{j=1, \ldots, \xi(\mathscr{C})}\left(\sum_{e \in \mathcal{E}\left(C_{j}\right)} w(e)\right) \tag{1}
\end{equation*}
$$

where, with a little abuse of notation, $C_{j}$ denotes the $j$-th cycle in $\mathscr{C}$. Therefore, the goal of the rooted min-max cycle cover problem is to find $q$ edge-disjoint cycles $\mathscr{C}=\left\{C_{1}, \ldots, C_{q}\right\}$ such that (a) $q \leq k$, (b) each cycle contains exactly one depot in $D$, (c) the union of these cycles includes all vertices in $V$, and (d) the maximum weight of cycles is minimized. As mentioned earlier, a feasible solution we obtain from the optimization problem in (1) also determines the (minimum) number of robots that will be stored at each depot.

Throughout the paper, we assume that the edge weight function $w$ is a metric and, hence, satisfies symmetry and triangle inequality: for all distinct vertices $v_{1}, v_{2}, v_{3} \in V$, we have

$$
w\left(v_{1}, v_{2}\right) \leq w\left(v_{1}, v_{3}\right)+w\left(v_{3}, v_{2}\right)
$$

This is a natural assumption for both choices of edge weights discussed earlier.

Remark 1: Recall that the rooted min-max cycle cover problem in (1) permits only edge-disjoint cycles. Our problem of robot planning does not explicitly require that the cycles be edge-disjoint. However, using the assumption that the edge weight function $w$ is a metric, one can easily show that the following holds: given a feasible non-edge disjoint cycle cover $\mathscr{C}$, we can construct an edge-disjoint cycle cover $\mathscr{C}^{\text {ed }}$ such that the maximum cycle weight of $\mathscr{C}^{\text {ed }}$ is less than or equal to that of $\mathscr{C}$. For this reason, without loss of generality, we

[^2]can focus on edge-disjoint cycle covers in the rooted min-max cycle cover problem in (1).

Remark 2: It is clear from the constraints that, if $k<$ $|D|=: m$, then there is no feasible solution; since each cycle can include only one depot (constraint c3), at most $k$ depots can be contained in the union $\cup_{j=1}^{k} C_{j}$ and, as a result, constraint c4 cannot be satisfied. Although this issue can be dealt with by changing the formulation slightly, we assume that there are more robots than the depots, i.e., $k \geq m$. This is a reasonable assumption for our problem as the number of depots is expected to be small with each depot housing many robots.

Remark 3: From the problem formulation in 11, the cycle cover we are looking for must cover not only all the sites of interest, but also all depots in $D$. In other words, we require that at least one robot be stationed at each depot. Although we make this assumption explicit, this is likely to be satisfied in practice even without making it explicit; depots should be spread out across the bridge for storage and recharging, and sites close to each depot should be assigned to robot(s) stationed at the depot. Otherwise, the depot should be removed.

## IV. Preliminaries: Key Steps of the Proposed Algorithm

The proposed algorithm (Algorithm 3 in Section V) consists of several steps that we discuss in detail in this section. The input to the algorithm comprises the information for the rooted min-max cycle cover problem: (a) a complete undirected graph $G=(V, E)$, (b) a metric edge weight function $w$, (c) the number of cycles $k$, and (d) the depot set $D$.

## A. Step 1: Construction of a rooted spanning forest $F^{*}$


(a) Depots (red) and other vertices
in $V^{-}$(black).

(b) Merging of all depots into a single vertex $\hat{d}$ (red circle).

(c) Construction of a minimum spanning tree $T^{\prime}$.
(d) Construction of a rooted span-

Fig. 2: Illustration of Step 1 for constructing a rooted spanning forest.

The first step generates a rooted forest with $m$ trees and consists of the following two steps $\stackrel{4}{4}^{4}$

[^3]Step 1-i: First, we collapse all depots in $D$ into a single node, which we denote by $\hat{d}$. This is shown in Fig. 2(b). Second, we create an edge from $\hat{d}$ to every vertex $v$ in $V \backslash D=: V^{-}$with an edge weight

$$
w(\hat{d}, v)=\min _{d \in D} w(d, v)
$$

We denote the complete undirected graph with vertex set $V^{\prime}:=V^{-} \cup\{\hat{d}\}$ by $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$. Third, we compute a minimum spanning tree $T^{\prime}$ of $G^{\prime}$, using Prim's algorithm [16], as shown in Fig. 2(c).

Step 1-ii: We uncouple the node $\hat{d}$ back into the original depots to create a forest $F^{*}$ as follows: first, let $F^{*}=T^{\prime}$. Second, add the depots in $D$ and replace each edge $(v, \hat{d}), v \in$ $V^{-}$, with an edge $\left(d^{\star}, v\right)$, where $d^{\star} \in \arg \min _{d \in D} w(d, v)$. When $\left|\arg \min _{d \in D} w(d, v)\right|>1$, we arbitrarily choose a depot in the set. Finally, we remove $\hat{d}$. The resulting rooted forest $F^{*}$ has exactly $m$ trees in it, which are denoted by $T_{1}^{*}, \ldots, T_{m}^{*}$. This is illustrated in Fig. 2 d ).

## B. Step 2: Virtual Forest $(F)$ and an Exhaustive Search

The second step takes $F^{*}$ generated in Step 1 as the input and produces a set of forests with trees that are created by merging the trees in $F^{*}$ in different ways. In particular, it considers $2^{m-1}$ different possible ways in which the $m$ trees in $F^{*}$ can be connected with each other. The reason for this is to ensure that we consider at least one forest that will lead to a rooted cycle cover with a provable upper bound on the maximum cycle weight (Steps 3 and 4). Specifically, we want to make sure that the forest described below is considered in Steps 3 and 4 through an exhaustive search.

Step 2-i: Consider the construction of a spanning tree using Algorithm 1 below, starting with the forest $F^{*}$.

```
Algorithm 1 Construction of Edge Set \(E^{\dagger}\)
    Let \(F_{\mathrm{tmp}}=F^{*}\) and \(E^{\dagger}=\emptyset\)
    while there is more than one tree in \(F_{\mathrm{tmp}}\) do
        Find an edge \(e^{c} \in E\) that joins two distinct trees \(T_{1}\)
    and \(T_{2}\) in \(F_{\mathrm{tmp}}\) with the smallest weight
        Add \(e^{c}\) to \(E^{\dagger}\)
        Remove \(T_{1}\) and \(T_{2}\) from \(F_{\mathrm{tmp}}\), and add the new tree
    formed after connecting \(T_{1}\) and \(T_{2}\) with \(e^{c}\) to \(F_{\text {tmp }}\)
    end while
```

Note that the edge set $E^{\dagger}$ produced by Algorithm 1 contains $m-1$ edges that connect the $m$ trees in $F^{*}$ into a single spanning tree.
a) Forest $F_{\text {opt }}$ : Suppose that $\mathscr{C}^{*}=\left\{C_{1}^{*}, \ldots, C_{q^{*}}^{*}\right\}$, where $q^{*} \leq k$, is a solution to (1), i.e., an optimal cycle cover, and the maximum weight of cycles in $\mathscr{C}^{*}$ is $\lambda^{*}$. Note that $\lambda^{*}$ is the optimal value of (1). Throughout the remainder of the paper, with a little abuse of notation, we denote the total weight of edges in a subgraph $\tilde{G}$, such as trees and cycles, by $w(\tilde{G})$.

Assume for now that the optimal value $\lambda^{*}$ and the optimal cycle cover $\mathscr{C}^{*}$ are known. We classify a tree $T$ into three categories on the basis of the optimal value $\lambda^{*}$ :

T1. Heavy tree : $w(T) \geq \lambda^{*}$.
T2. Light tree : $w(T)<\lambda^{*}$ and there exists a cycle $C \in \mathscr{C}^{*}$ such that $\mathcal{V}(C) \subseteq \mathcal{V}(T)$, i.e., $T$ contains all the vertices in at least one cycle in the optimal cycle cover $\mathscr{C}^{*}$.
T3. Bad tree : $w(T)<\lambda^{*}$ and there is no cycle $C \in \mathscr{C}^{*}$ such that $\mathcal{V}(C) \subseteq \mathcal{V}(T)$.
Based on this classification of trees, starting with the forest $F^{*}$, we perform a procedure in Algorithm 2 to eliminate bad trees and construct a new forest free of bad trees.

```
Algorithm 2 Elimination of Bad Trees
    Let \(F_{\mathrm{tmp}}=F^{*}\)
    while there is at least one bad tree in \(F_{\mathrm{tmp}}\) do
        Choose a bad tree \(T_{b}\) in \(F_{\mathrm{tmp}}\)
        Pick an edge \(e_{\text {new }}=\left(v_{\text {new }}, v_{\text {new }}^{\prime}\right)\) from \(E^{\dagger}\left(T_{b}\right)\) with
    the smallest weight, where \(E^{\dagger}\left(T_{b}\right):=\left\{\left(v, v^{\prime}\right) \in E^{\dagger} \mid v \in\right.\)
    \(\left.\mathcal{V}\left(T_{b}\right), v^{\prime} \notin \mathcal{V}\left(T_{b}\right)\right\}\)
        Connect \(T_{b}\) to \(T_{c}\) using edge \(e_{\text {new }}\) to create a new tree
    \(T_{\text {new }}\), where \(T_{c}\) is the tree in \(F_{\mathrm{tmp}}\) which includes vertex
    \(v_{\text {new }}^{\prime}\)
        Remove \(T_{b}\) and \(T_{c}\) from \(F_{\mathrm{tmp}}\), and add the new tree
    \(T_{\text {new }}\) to \(F_{\text {tmp }}\)
    end while
```

When we connect a bad tree $T_{b}$ to $T_{c}$ using edge $e_{\text {new }}$ in Algorithm 2 (line 5), the possible types of the new tree $T_{\text {new }}$ depend on the type of tree $T_{c}$.

P1 $T_{c}$ is a heavy tree - When we connect $T_{b}$ to a heavy tree, the total weight of the new tree obviously exceeds $\lambda^{*}$, and the new tree $T_{\text {new }}$ is a heavy tree.
P2 $T_{c}$ is a light tree - After connecting the two trees, $T_{\text {new }}$ can be either light or heavy, depending on its total weight.
P3 $T_{c}$ is a bad tree - When two bad trees are connected, the resulting new tree $T_{\text {new }}$ could be any of the three types (bad, light or heavy).
Note that a new tree $T_{\text {new }}$ can be a bad tree only if $T_{c}$ is also a bad tree (case P3).

The order in which we choose bad trees in Algorithm 2 is not important. In addition, Algorithm 2 terminates after at most $m-1$ rounds; when all $m$ trees $T_{1}, \ldots, T_{m}$ are combined into a single spanning tree after at most $m-1$ rounds, the resulting tree contains all vertices in $V$ and, hence, cannot be a bad tree, thereby terminating the algorithm.

Denote by $F_{\text {opt }}$ the final forest $F_{\text {tmp }}$ produced by Algorithm 2 Let $n_{L T} \in\{0,1, \ldots, m\}$ be the number of light trees in $F_{\mathrm{opt}}$.

Lemma 1: Suppose that $F_{h}$ is the collection of heavy trees in $F_{\text {opt }}$. Then,

$$
w\left(F_{h}\right) \leq\left(k-n_{L T}\right) \lambda^{*}
$$

Proof A proof of the lemma can be found in Appendix A

Step 2-ii: Unfortunately, in practice, we do not have access to $\lambda^{*}$ or the optimal cycle cover $\mathscr{C}^{*}$. Hence, we cannot
determine which trees in the forest $F^{*}$ are bad trees and execute Algorithm 2.

For this reason, we consider all possible ways in which the $m$ trees in $F^{*}$ can be connected to form a new forest, using the edges in $E^{\dagger}$. Since $\left|E^{\dagger}\right|=m-1$, the number of possible forests we need to consider, including the case with a single spanning tree, is equal to $2^{m-1}$, and one of these possible forests coincides with $F_{\text {opt }}$. We denote by $\mathscr{F}$ the set of $2^{m-1}$ forests we consider, and use $F_{\text {cand }}$ to refer to a forest in $\mathscr{F}$.

We provide these $2^{m-1}$ forests in $\mathscr{F}$ as an input to Steps 3 and 4, one forest at a time. However, we are primarily interested in the forest $F_{\text {opt }}$ for finding an approximation ratio for the proposed algorithm.

## C. Step 3: Decomposition of Heavy Trees

In addition to the input to the proposed algorithm, namely the information regarding the rooted min-max cycle cover problem, Steps 3 and 4 described in this and following subsections require another parameter $\lambda$. The output of these two steps depends on whether or not the parameter $\lambda$ is greater than or equal to $\lambda^{*}$ : if $\lambda \geq \lambda^{*}$, they return a cycle cover with at most $k$ cycles and the maximum cycle weight less than or equal to $5 \lambda$.

Suppose that the forest under consideration (out of $2^{m-1}$ possible forests in $\mathscr{F})$ is $F_{\text {cand }}$, and $\lambda$ is a constant satisfying $\max _{e \in \mathcal{E}\left(F_{\text {cand }}\right)} w(e) \leq \lambda$. We decompose the trees in $F_{\text {cand }}$ with the help of the following lemma.

Lemma 2 (Lemma 2 of [4]): Fix a tree $T$ and $\lambda \geq$ $\max _{e \in \mathcal{E}(T)} w(e)$. Then, the tree $T$ can be decomposed into subtrees $T_{1}^{s}, \ldots, T_{m_{s}}^{s}$, such that (i) $w\left(T_{i}^{s}\right)<2 \lambda$ for all $i=1, \ldots, m_{s}$, (ii) $\mathcal{V}\left(\cup_{i=1}^{m_{s}} T_{i}^{s}\right)=\mathcal{V}(T)$, and (iii) the number of subtrees, $m_{s}$, satisfies

$$
m_{s} \leq \max \left(\left\lfloor\frac{w(T)}{\lambda}\right\rfloor, 1\right) .
$$

The goal of Step 3 is to construct a new forest $\hat{F}_{\text {cand }}$ that will be used in Step 4 to find a rooted cycle cover. In particular, when $F_{\text {cand }}=F_{\text {opt }}$, the cycle cover found in Step 4 will possess a provable upper bound on the maximum cycle weight.

To this end, we first put all trees in $F_{\text {cand }}$ whose weight is less than $2 \lambda$ in $\hat{F}_{\text {cand }}$. Second, for trees whose weight is greater than or equal to $2 \lambda$, we first decompose them into subtrees as described in Lemma 2, using the algorithm proposed in [2], and then put the subtrees in $\hat{F}_{\text {cand }}$.

The resulting forest $\hat{F}_{\text {cand }}$ when $F_{\text {cand }}=F_{\text {opt }}$, which we denote simply by $\hat{F}_{\text {opt }}$, has the following two properties. Given a subgraph $G^{\prime}$, we define $w_{\max }\left(G^{\prime}\right):=\max _{e \in \mathcal{E}\left(G^{\prime}\right)} w(e)$ to be the maximum edge weight in the subgraph.

Lemma 3: Suppose $\lambda \geq w_{\max }\left(F_{\mathrm{opt}}\right)$. Then, the maximum weight of trees in $\hat{F}_{\text {opt }}$ is at most $2 \lambda$.
Proof The lemma follows directly from Lemma 2 and the construction of $\hat{F}_{\text {opt }}$.

Lemma 4: If $\lambda \geq \lambda^{*}$, the number of trees in $\hat{F}_{\text {opt }}$ is at most k.

Proof Please see Appendix B for a proof of the lemma.

## D. Step 4: Generation of a Rooted Cycle Cover

Recall that, even though every tree in forest $F_{\text {cand }}$ contains at least one depot, due to the decomposition of trees $T$ with $w(T) \geq 2 \lambda$ in Step 3, some of the trees in $\hat{F}_{\text {cand }}$ may not include any depot. We partition the trees in the forest $\hat{F}_{\text {cand }}$ into two forests, $\hat{F}_{r}$ and $\hat{F}_{n r}$. Forest $\hat{F}_{r}$ consists of trees that contain a depot, and $\hat{F}_{n r}$ comprise trees that do not cover any depot.

We connect each tree $T$ in $\hat{F}_{n r}$ to a nearest depot $\tilde{d}(T)$ in $\arg \min _{d \in D}\left(\min _{v \in \mathcal{V}(T)} w(d, v)\right)$. Note that $\min _{d \in D}\left(\min _{v \in \mathcal{V}(T)} w(d, v)\right) \leq \lambda^{*} / 2$; for every vertex in $V^{-}$, we have $2 \min _{d \in D} w(v, d) \leq \lambda^{*}$ from the assumption that the edge weight function $w$ is a metric and the optimal cycle cover must cover all vertices in $V^{-}$. For each tree $T$ in $\hat{F}_{n r}$, we denote the resulting tree we obtain after connecting it to a nearest depot by $T^{+}$, and let $\hat{F}_{n r}^{+}:=\left\{T^{+} \mid T \in \hat{F}_{n r}\right\}$. Note that some trees in $\hat{F}_{n r}^{+}$may share a depot with other trees in $\hat{F}_{r}$ or $\hat{F}_{n r}^{+}$.


Fig. 3: An example of generating a cycle for a tree
Let us consider the case when $\hat{F}_{\text {cand }}=\hat{F}_{\text {opt }}$. Recall from Lemma 4 that the number of trees in $\hat{F}^{+}:=\hat{F}_{r} \cup \hat{F}_{n r}^{+}$does not exceed $k$ if $\lambda \geq \lambda^{*}$. Each tree in $\hat{F}^{+}$contains one depot, and its weight is at most $2 \lambda+\lambda^{*} / 2$ from its construction. By finding the Eulerian cycle and performing short cutting (as shown in Fig. 3) for each tree in $\hat{F}^{+}$, we can find a cycle cover $\mathscr{C}_{\text {cand }}$. When $\hat{F}_{\text {cand }}=\hat{F}_{\text {opt }}$, the cycle cover $\mathscr{C}_{\text {cand }}$ has the maximum cycle weight of at most $4 \lambda+\lambda^{*}$.

Lemma 5: When $F_{\text {cand }}=F_{\text {opt }}$, for $\lambda \geq w_{\max }\left(F_{\text {opt }}\right)$, Steps 3 and 4 together generate a cycle cover with the maximum cycle weight of at most $4 \lambda+\lambda^{*}$. In addition, if $\lambda \geq \lambda^{*}$, there are at most $k$ cycles in the cycle cover.

## V. The Proposed Algorithm

We are now ready to present our proposed algorithm for (11). The pseudocode of the proposed algorithm is shown in Algorithm 3 We denote the maximum cycle weight of a cycle cover produced in Step 4 (line 12), namely $\mathscr{C}_{\text {cand }}$, by $O_{\text {cand. }}$. It is clear from the pseudocode, if the number of trees in $\hat{F}_{\text {cand }}$ is larger than $k$, we do not need to execute Step 4 as it will not yield a feasible solution.

Given a forest $F_{\text {cand }}$ in $\mathscr{F}$, a binary search for a suitable value of $\lambda$ (lines $8-26$ ) is performed over the interval $\left[w_{\max }\left(F_{\text {cand }}\right),(n+k) w_{\max }(G)\right]$. The goal of the binary search

```
Algorithm 3 Rooted Min-Max Cycle Cover Algorithm
Input: (i) a complete undirected graph \(G=(V, E)\); (ii) a
    weight function \(w: E \rightarrow \mathbb{R}_{+}\); (iii) depot set \(D \subseteq V\); (iv)
    maximum number of cycles \(k\); (v) constant \(\epsilon \in(0,1)\)
Output: A cycle cover with the maximum weight of cycles
    less than \((5+\epsilon) \lambda^{*}\)
    Construct a rooted spanning forest \(F^{*} \quad \triangleright\) (Step 1)
    Find the edge set \(E^{\dagger}\)
                                \(\triangleright\) (Step 2-i)
    Generate \(\mathscr{F}\) with \(2^{m-1}\) possible forests
    Set \(\mathscr{C}_{\text {tmp }}=\emptyset\) and \(O_{\text {tmp }}=\infty\)
    for each forest \(F_{\text {cand }}\) in \(\mathscr{F}\) do \(\quad \triangleright\) (Step 2-ii)
        \(\ell \leftarrow 1\)
        \(a_{\ell} \leftarrow w_{\max }\left(F_{\text {cand }}\right)\) and \(b_{\ell} \leftarrow(n+k) w_{\max }(G)\)
        while true do
            \(\lambda_{\ell} \leftarrow 0.5\left(a_{\ell}+b_{\ell}\right)\)
            Construct forest \(\hat{F}_{\text {cand }}\) with \(\lambda=\lambda_{\ell} \quad \triangleright\) (Step 3)
            if number of trees in \(\hat{F}_{\text {cand }} \leq k\) then
                Find a cycle cover \(\mathscr{C}_{\text {cand }} \quad \triangleright(\) Step 4\()\)
                if \(O_{\text {cand }}<O_{\text {tmp }}\) then
                        \(\mathscr{C}_{\text {tmp }} \leftarrow \mathscr{C}_{\text {cand }}\)
                        \(O_{\mathrm{tmp}} \leftarrow O_{\text {cand }}\)
                    end if
                    \(a_{\ell+1} \leftarrow a_{\ell}\) and \(b_{\ell+1} \leftarrow \lambda_{\ell}\)
                else
                    \(a_{\ell+1} \leftarrow \lambda_{\ell}\) and \(b_{\ell+1} \leftarrow b_{\ell}\)
            end if
            if \(b_{\ell}-a_{\ell}<0.5 \epsilon a_{\ell}\) then
                Break \(\quad \triangleright\) break out of while loop
            else
                \(\ell \leftarrow \ell+1\)
            end if
        end while
    end for
```

is not to find the smallest value of $\lambda$ greater than or equal to $\lambda^{*}$. Instead, it is to find a value of $\lambda$ for which Step 3 produces a feasible forest $\hat{F}_{\text {cand }}$ with at most $k$ trees (line 11) and the termination condition in line 21 is satisfied.

Lemma 4 guarantees that our algorithm will produce at least one feasible solution with at most $k$ cycles as long as $\lambda^{*}$ lies in $\left[w_{\max }\left(F_{\text {opt }}\right),(n+k) w_{\max }(G)\right)$. Moreover, Lemma 5 ensures that the maximum cycle weight of the cycle cover found by the algorithm is at most $4 \lambda+\lambda^{*}$. Therefore, by selecting a suitable value of $\lambda$, we can prove the approximation ratio of $5+\epsilon$ for the algorithm, where $\epsilon$ is a constant in the interval $(0,1)$, which is an input to the algorithm we select.

Theorem 1: For fixed $\epsilon \in(0,1)$, Algorithm 3 returns a rooted cycle cover with at most $k$ cycles and the maximum cycle weight less than or equal to $(5+\epsilon) \lambda^{*}$.
Proof A proof of the theorem can be found in Appendix C.

## VI. Complexity of the Proposed Algorithm

In this section, we study the complexity of the proposed algorithm (Algorithm 3). For a fixed number of depots, $m$, we will show that the overall complexity is $O\left(n^{2}\right)$, where $n$ is the number of vertices in $V$, by examining the computational requirements of each step.

- Step 1: The Prim's algorithm for finding a minimum spanning tree $T^{\prime}$ (Step 1-i) has complexity of $O\left(n^{2}\right)$. Step 1-ii simply requires finding the closest depot for each vertex in $V^{-}$and has complexity of $O(n)$.
- Step 2: A naive way to find the edge set $E^{\dagger}$ is to construct a minimum spanning tree for $G$ (using, for instance, Prim's algorithm [16]) and remove the edges in $F^{*}$. This has complexity $O\left(n^{2}\right)$.
- Step 3: For every forest $F_{\text {cand }}$ in $\mathscr{F}, w_{\max }\left(F_{\text {cand }}\right) \geq$ $w_{\max }\left(F^{*}\right)$. Hence, the while loop in lines 8 through 26 terminates after at most

$$
\begin{aligned}
& \log _{2}\left(\frac{(n+k) w_{\max }(G)}{0.5 \epsilon w_{\max }\left(F^{*}\right)}\right) \\
& =O\left(\log _{2}(n+k)\right)+O\left(\log _{2}(1 / \epsilon)\right) \\
& \quad+O\left(\log _{2}\left(w_{\max }(G) / w_{\max }\left(F^{*}\right)\right)\right)
\end{aligned}
$$

iterations. In addition, decomposing trees whose weight is greater than or equal to $2 \lambda$ using 'splitting' described in [2], [4] has complexity of $O(n)$.

- Step 4: Both connecting trees in $\hat{F}_{n r}$ to the closest depots and finding the Eulerian cycles and associated rooted cycle cover have complexity of $O(n)$.
From the above discussion, the overall complexity of the proposed algorithm is $O\left(n^{2}\right)+O\left(2^{m-1} n \log (n+\right.$ $k))+O(n \log (1 / \epsilon))$, provided that $w_{\max }(G) / w_{\max }\left(F^{*}\right)=$ $O(\exp (n))$. In practice, we have $k=O(n)$ and, as a result, the complexity is in fact $O\left(n^{2}\right)+O\left(2^{m-1} n \log (n)\right)+$ $O(n \log (1 / \epsilon))$. Therefore, when both the number of depots $(m)$ and $\epsilon$ are fixed, the complexity is $O\left(n^{2}\right)$. Furthermore, since the $2^{m-1}$ possible forests in $\mathscr{F}$ can be considered independently, the execution of the for loop in the algorithm (lines 5-27) can be parallelized.

Finally, recall that $F^{*}$ is a spanning forest consisting of $m$ trees, which are constructed from the minimum spanning tree $T^{\prime}$ of $G^{\prime}$ in Step 1-i and are rooted at the $m$ depots. Hence, unless all remaining $n-m$ vertices get closer to the $m$ depots with increasing $n$, the assumption $w_{\max }(G) / w_{\max }\left(F^{*}\right)=$ $O(\exp (n))$ will hold. For instance, it is well known that the longest edge of the minimum spanning tree covering $n$ independent and identically distributed points in a unit ball in $\mathbb{R}^{d}, d \geq 2$, converges almost surely to $c(\log (n) / n)^{1 / d}$ as $n$ goes to $\infty$, where $c$ is some constant [17]. Hence, for a random geometric graph $G$ with $d=3$, we have $w_{\max }(G) / w_{\max }\left(F^{*}\right)=O\left((n / \log (n))^{1 / 3}\right)$.

## VII. Conclusion

We studied the problem of robot deployment planning for inspection of civil structures, with an emphasis on bridges.

Each mobile robot is stationed at a depot where it recharges the battery and is tasked with the inspection of a subset of points or segments of the structure. The problem is formulated as a (variant of the) uncapacitated rooted min-max cycle cover problem. We proposed a new algorithm with approximation ratio of $5+\epsilon$.

Our formulation as an uncapacitated rooted min-max cycle cover problem assumes that depots are large enough to house as many robots as needed. We suspect that any good solution will distribute the robots evenly across a bridge to minimize the max cycle weight. In some cases, however, the depots may have limited slots for the robots for housing and recharging. This may require reformulating the problem as a capacitated rooted min-max cycle cover, similar to that studied by the authors of [6]. We are currently working on this problem.

## Acknowledgment

This work was supported in part by an NSF grant ECCS 1446785 and a NIST grant 70NANB16H024.

## Appendix A

## Proof of Lemma 1

Let $V_{h}:=\mathcal{V}\left(F_{h}\right)$ be the set of vertices in $F_{h}$ and $F_{h}^{*}$ the smallest subset of trees in $F^{*}$, which covers $V_{h}$, i.e., $V_{h}=$ $\mathcal{V}\left(F_{h}^{*}\right)$. Similarly, we define $\mathscr{C}_{h}^{*}$ to be the smallest subset of cycles in $\mathscr{C}^{*}$ such that the union of vertices in the cycles contain $V_{h}$, i.e., $V_{h} \subseteq \cup_{C \in \mathscr{C}_{h}^{*}} \mathcal{V}(C)=: V_{\mathscr{C}_{h}^{*}}$.

First, note that, by deleting an edge with the largest weight in each cycle $C \in \mathscr{C}_{h}^{*}$, we can obtain a rooted forest $F_{\mathscr{C}_{h}^{*}}$ that covers $V_{\mathscr{C}_{h}^{*}}$ and each tree in the forest contains a depot in $D$ because each cycle $C \in \mathscr{C}_{h}^{*}$ must be rooted at a depot. Moreover, from the construction of $F^{*}$ in Step 1 (subsection IV-A), $F_{h}^{*}$ is a minimum weight rooted spanning forest of $V_{h}$. Therefore, because $V_{h} \subseteq V_{\mathscr{C}_{h}^{*}}$, we have

$$
\begin{aligned}
w\left(F_{h}^{*}\right) & \leq w\left(F_{\mathscr{C}_{h}^{*}}^{*}\right) \\
& =w\left(\mathscr{C}_{h}^{*}\right)-\sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C) \\
& \leq\left(k-n_{L T}\right) \lambda^{*}-\sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C)
\end{aligned}
$$

where $w_{\max }\left(G^{\prime}\right):=\max _{e \in \mathcal{E}\left(G^{\prime}\right)} w(e)$ for any subgraph $G^{\prime}$, and the second inequality follows from the fact that (i) there are at most $k-n_{L T}$ cycles in $\mathscr{C}_{h}^{*}$ because at least $n_{L T}$ cycles are covered by $n_{L T}$ light trees and (ii) for every cycle $C$ in $\mathscr{C}_{h}^{*}, w(C) \leq \lambda^{*}$.

In order to bound the total weight of the trees in $F_{h}$, in addition to the bound for $w\left(F_{h}^{*}\right)$, we need to account for the weights of the edges that were introduced in the the trees of $F_{h}$ through the merging process of bad trees in $F_{h}^{*}$. For this, we consider two cases based on how the heavy tree $T_{h}$ in $F_{h}$, which does not belong to $F_{h}^{*}$, was created. Let $T_{b}$ be the bad tree that was connected to another tree to form $T_{h}$.
C 1 . The bad tree $T_{b}$ is a bad tree in $F_{h}^{*}$.
C 2 . The bad tree $T_{b}$ is not a bad tree in $F_{h}^{*}$. In this case, $T_{b}$ must have been created as a result of merging at least
two bad trees in previous round(s) of connecting bad trees.
Case C1: Let $d$ be the depot covered by $T_{b}$, and choose a cycle $C_{d}^{*}$ in $\mathscr{C}_{h}^{*}$, which includes $d$. The existence of such a cycle is guaranteed because the cycles in $\mathscr{C}_{h}^{*}$ must cover all vertices in $V_{h}$, including depots in $V_{h}$.

Since $T_{b}$ is a bad tree, there exists at least one vertex $v^{*}$ in $C_{d}^{*}$ which is not covered by $T_{b}$ and instead belongs to a different tree; otherwise, $T_{b}$ would be a light tree. As a result, there is at least one edge $e^{\prime}$ that connects a vertex in $T_{b}$ to another tree and $e^{\prime}$ belongs to $C_{d}^{*}$. Since Algorithm 1 uses an edge with the smallest weight to connect $T_{b}$ to the rest of spanning tree in the process of finding the edge set $E^{\dagger}$, the edge $e^{+} \in E^{\dagger}$ used to connect $T_{b}$ to another tree (hence, is a new edge introduced in $T_{h}$ ) satisfies

$$
\begin{equation*}
w\left(e^{+}\right) \leq w\left(e^{\prime}\right) \leq w_{\max }\left(C_{d}^{*}\right) \tag{2}
\end{equation*}
$$

Case C2: We first consider the construction of the bad tree $T_{b}$. Since $T_{b}$ is a bad tree, it must have been produced as a result of connecting two or more bad trees in $F_{h}^{*}$. Assume that $T_{b}$ is constructed as a result of $L$ rounds of merging bad trees, and for each $\ell=1, \ldots, L$, let $T_{b_{\ell}}$ be the tree chosen to be connected to another tree $T_{b_{\ell}}^{\prime}$ for the $\ell$-th round of merging. Furthermore, for each $\ell=1, \ldots, L$, denote the edge chosen to connect the two bad trees $T_{b_{\ell}}$ and $T_{b_{\ell}}^{\prime}$ by $e_{\ell}^{+} \in E^{\dagger}$.

For each $\ell=1, \ldots, L$, suppose that $d_{\ell}$ is a depot in $T_{b_{\ell}}$ which has not been considered in the previous rounds, i.e., $d_{\ell} \notin\left\{d_{1}, \ldots, d_{\ell-1}\right\}$, and a cycle $C_{\ell}^{*}$ in $\mathscr{C}_{h}^{*}$ covers $d_{\ell}$. We can always find a depot $d_{\ell}$ that meets the above condition because $T_{b_{\ell}}$ contains at least $r+1$ depots, where $r$ is the number of rounds of merging process needed to form $T_{b_{\ell}}$ before (hence, $T_{b_{\ell}}$ includes at least $r+1$ bad trees in $F_{h}^{*}$ ) and exactly $r$ of these depots were considered in the $r$ rounds. This also implies that $C_{\ell}^{*} \neq C_{l}^{*}$ for all $l=1, \ldots, \ell-1$, because each cycle in $\mathscr{C}_{h}^{*}$ contains exactly one depot.

Following the same argument used in the previous case (C1), for each $\ell=1, \ldots, L$, we can find an edge $e_{\ell}^{\prime}$ in the cycle $C_{\ell}^{*}$, which connects a vertex in $T_{b_{\ell}}$ to another tree. Consequently,

$$
\begin{equation*}
w\left(e_{\ell}^{+}\right) \leq w\left(e_{\ell}^{\prime}\right) \leq w_{\max }\left(C_{\ell}^{*}\right) \tag{3}
\end{equation*}
$$

Finally, the bad tree $T_{b}$ is connected to another tree to form the heavy tree $T_{h}$. By the same argument used in case C 1 , we can find a depot $d$ in $T_{b}$ such that $d \notin\left\{d_{1}, \ldots, d_{L}\right\}$ and a cycle $C_{d}^{*}$ containing $d$ so that the weight of the edge $e^{+}$that connects $T_{b}$ to another tree to form $T_{h}$ satisfies

$$
\begin{equation*}
w\left(e^{+}\right) \leq w\left(e^{\prime}\right) \leq w_{\max }\left(C_{d}^{*}\right) \tag{2}
\end{equation*}
$$

where $e^{\prime}$ is an edge in $C_{d}^{*}$ which connects $T_{b}$ to another tree as described in case C 1 .

Note from the above discussion that, for each added edge $e^{\star}$ during the construction of a heavy tree $T_{h}$ that is not in $F_{h}^{*}$, we can find a distinct cycle in $\mathscr{C}_{h}^{*}$ such that the weight of the new added edge is upper bounded by the maximum edge weight of the cycle. Let $E^{+}\left(\subseteq E^{\dagger}\right)$ denote the set of edges
that were added to connect trees in $F_{h}^{*}$ to create $F_{h}$. Then, we have $\sum_{e \in E^{+}} w(e) \leq \sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C)$ because no cycle $C \in \mathscr{C}_{h}^{*}$ is considered more than once during the process. Therefore, we have

$$
\begin{aligned}
& w\left(F_{h}\right)=w\left(F_{h}^{*}\right)+\sum_{e \in E^{+}} w(e) \\
& \leq\left(k-n_{L T}\right) \lambda^{*}-\sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C)+\sum_{e \in E^{+}} w(e) \\
& \leq\left(k-n_{L T}\right) \lambda^{*}-\sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C)+\sum_{C \in \mathscr{C}_{h}^{*}} w_{\max }(C) \\
& \leq\left(k-n_{L T}\right) \lambda^{*} .
\end{aligned}
$$

## Appendix B

## Proof of Lemma 4

First, note that all $n_{L T}$ light trees in $F_{\text {opt }}$ belong to $\hat{F}_{\text {opt }}$ since the weight of light trees is less than $\lambda^{*}$, which is less than or equal to $\lambda$ by assumption.

Order the heavy trees in $F_{h}$ by decreasing weight: $T_{1}, \ldots, T_{n_{1}}, T_{n_{1}+1}, \ldots, T_{n_{1}+n_{2}}$, where $n_{1}$ is the number of heavy trees whose weight is greater than or equal to $2 \lambda$, i.e., $w\left(T_{i}\right) \geq 2 \lambda$ for $i=1, \ldots, n_{1}$, and $\lambda^{*} \leq w\left(T_{i}\right)<2 \lambda$ for $i=n_{1}+1, \ldots, n_{1}+n_{2}$. Let $\hat{F}_{h}$ be the set of trees in $\hat{F}_{\mathrm{opt}}$ which come from the heavy trees in $F_{h}$.

Using Lemma 2 we can upper bound $\left|\hat{F}_{h}\right|$ as follows.

$$
\begin{aligned}
\left|\hat{F}_{h}\right| & \leq \sum_{i=1}^{n_{1}} \max \left(\left\lfloor\frac{w\left(T_{i}\right)}{\lambda}\right\rfloor, 1\right)+n_{2} \\
& \leq \sum_{i=1}^{n_{1}} \frac{w\left(T_{i}\right)}{\lambda}+n_{2} \\
& =\frac{w\left(F_{h}\right)-\sum_{i=n_{1}+1}^{n_{1}+n_{2}} w\left(T_{i}\right)}{\lambda}+n_{2} \\
& \leq \frac{w\left(F_{h}\right)-n_{2} \lambda^{*}}{\lambda}+n_{2}
\end{aligned}
$$

where the last inequality follows from $\lambda^{*} \leq w\left(T_{i}\right)$ for all $i=n_{1}+1, \ldots, n_{1}+n_{2}$. Using the bound for $w\left(F_{h}\right)$ in Lemma 1. we obtain

$$
\begin{aligned}
\left|\hat{F}_{h}\right| & \leq \frac{\left(k-n_{L T}\right) \lambda^{*}-n_{2} \lambda^{*}}{\lambda}+n_{2} \\
& \leq\left(k-n_{L T}\right)-n_{2}+n_{2} \\
& =k-n_{L T}
\end{aligned}
$$

where $n_{L T}$ is the number of light trees in $F_{\mathrm{opt}}$, and the second inequality is a consequence of the assumption $\lambda^{*} \leq \lambda$. Thus, the number of trees in $\hat{F}_{\text {opt }}$ can be upper bounded by $k$ because $\left|\hat{F}_{\text {opt }}\right|=\left|\hat{F}_{h}\right|+n_{L T} \leq k$.

## Appendix C <br> Proof of Theorem 1

In order to prove the theorem, we only need to show that at least one of the $2^{m-1}$ forests in $\mathscr{F}$ leads to a solution that satisfies the approximation ratio in the theorem. To this end, we consider the forest $F_{\text {opt }}$ generated from the bad tree
elimination process (Algorithm 2) with an optimal cycle cover $\mathscr{C}^{*}$. Recall that $F_{\text {opt }}$ always belongs to $\mathscr{F}$.

For forest $F_{\text {opt }}$, the binary search in Algorithm 3 is performed over the interval $\left[w_{\max }\left(F_{\mathrm{opt}}\right),(n+k) w_{\max }(G)\right]$. It is clear that $(n+k) w_{\max }(G)>\lambda^{*}$ since the optimal cycle cover cannot have more than $(n+k-m)$ edges. The following lemma demonstrates that $w_{\max }\left(F_{\text {opt }}\right)$ is a lower bound on the optimal value $\lambda^{*}$.

Lemma 6: The optimal value $\lambda^{*}$ is greater than or equal to $w_{\text {max }}\left(F_{\text {opt }}\right)$.
Proof: Let $e^{*}:=(u, v) \in \arg \max _{e \in \mathcal{E}\left(F_{\mathrm{opt}}\right)} w(e)$.
We shall consider following two cases.
C-i $e^{*} \in \mathcal{E}\left(F^{*}\right)$ - Suppose that $e^{*}$ belongs to a tree $T^{*}$ in $F^{*}$. Without loss of generality, we assume that, when we remove $e^{*}$ from $T^{*}$ and divide it into two subtrees, vertex $v$ belongs to the subtree with the depot in $T^{*}$.
Suppose that the claim is false and $w\left(e^{*}\right)>\lambda^{*}$. Let $C_{u}$ be a cycle in the optimal cycle cover $\mathscr{C}^{*}$ which covers vertex $u$ and $d_{u}$ be the depot in the cycle $C_{u}$. Denote by $e_{u}$ the edge $\left(u, d_{u}\right)$. Note that $w\left(e_{u}\right) \leq \lambda^{*} / 2$ because $2 w\left(e_{u}\right) \leq w\left(C_{u}\right) \leq \lambda^{*}$ because $w$ satisfies the triangle inequality.
Consider a new spanning forest $F^{\prime}$ we can generate by replacing the edge $e^{*}$ in $F^{*}$ with the edge $e_{u}$. It is clear that $F^{\prime}$ is also a rooted spanning forest of $G$ since every tree in $F^{\prime}$ includes exactly one depot and $F^{\prime}$ covers all vertices in $V$. However,

$$
w\left(F^{\prime}\right)=w\left(F^{*}\right)-w\left(e^{*}\right)+w\left(e_{u}\right)<w\left(F^{*}\right)
$$

because $w\left(e_{u}\right)<\lambda^{*}<w\left(e^{*}\right)$. This contradicts the earlier assumption that $F^{*}$ is a minimum rooted spanning forest.
C-ii $e^{*} \notin \mathcal{E}\left(F^{*}\right)$ - This means that $e^{*}$ is an edge added by Algorithm 2 during the bad tree elimination process. From (2) and (3),

$$
w\left(e^{*}\right) \leq \max _{\ell=1, \ldots, q^{*}} w_{\max }\left(C_{\ell}^{*}\right) \leq \max _{\ell=1, \ldots, q^{*}} w\left(C_{\ell}^{*}\right) \leq \lambda^{*}
$$

Since the claim of the lemma holds in both cases, this completes the proof of the lemma.

An important observation is that, during the binary search in Algorithm 3, the lower bound of the interval, $a_{\ell}$, increases for the next round only if the returned forest $F_{\text {cand }}$ in the $\ell$-th round has more than $k$ trees (line 19). However, since Steps 3 and 4 guarantee a rooted cycle cover with at most $k$ cycles when $\lambda \geq \lambda^{*}$ (Lemma 55) and $a_{1} \leq \lambda^{*}$ by Lemma 6, we have $a_{\ell} \leq \lambda^{*}$ for all $\ell=1,2, \ldots$

Suppose that the process terminates after $N$ rounds, i.e., $b_{N}-a_{N}<\epsilon a_{N} / 2$. We denote the returned rooted cycle cover after the $N$-th round by $\mathscr{C}^{\star}$. Using Lemma 5, we can show that the weight of the cycle cover $\mathscr{C}^{\star}$ is upper bounded by $(5+\epsilon) \lambda^{*}$ as follows.

$$
\begin{aligned}
w\left(\mathscr{C}^{\star}\right) & \leq 4 \lambda_{N}+\lambda^{*}=4\left(a_{N}+b_{N}\right) / 2+\lambda^{*} \\
& <2\left(a_{N}+a_{N}+\epsilon a_{N} / 2\right)+\lambda^{*} \\
& \leq(5+\epsilon) \lambda^{*}
\end{aligned}
$$

The last inequality is a consequence of the earlier observation that $a_{\ell} \leq \lambda^{*}$ for all $\ell=1,2, \ldots, N$. Obviously, the maximum cycle weight of $\mathscr{C}^{\star}$ is upper bounded by $w\left(\mathscr{C}^{\star}\right)$. Therefore, this proves the theorem.

## REFERENCES

[1] https://www.fhwa.dot.gov/bridge/nbi/no10/condition18.cfm
[2] G. Even, N. Garg, J. Könemann, R. Ravi, and A. Sinha, "Min-max tree covers of graphs," Operations Research Letters, vol. 32, no. 4, pp. 309315, 2004.
[3] E. M. Arkin, R. Hassin, and A. Levin, "Approximations for minimum and min-max vehicle routing problems," Journal of Algorithms, vol. 59, no. 1, pp. 1-18, 2006.
[4] M. R. Khani and M. R. Salavatipour, "Improved approximation algorithms for the min-max tree cover and bounded tree cover problems," Algorithmica, vol. 69, no. 2, pp. 443-460, 2014.
[5] A. Jorati, Approximation algorithms for some min-max vehicle routing problems. University of Alberta (Canada), 2013.
[6] W. Xu, W. Liang, and X. Lin, "Approximation algorithms for min-max cycle cover problems," IEEE Transactions on Computers, vol. 64, no. 3, pp. 600-613, 2015.
[7] W. Yu and Z. Liu, "Improved approximation algorithms for some min-max and minimum cycle cover problems," Theoretical Computer Science, vol. 654, pp. 45-58, 2016.
[8] N. Christofides, "Worst-case analysis of a new heuristic for the travelling salesman problem," tech. rep., Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.
[9] G. N. Frederickson, M. S. Hecht, and C. E. Kim, "Approximation algorithms for some routing problems," in Foundations of Computer Science, 1976., 17th Annual Symposium on, pp. 216-227, IEEE, 1976.
[10] B. Farbstein and A. Levin, "Min-max cover of a graph with a small number of parts," Discrete Optimization, vol. 16, pp. 51-61, 2015.
[11] H. Nagamochi, "Approximating the minmax rooted-subtree cover problem," IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, vol. 88, no. 5, pp. 1335-1338, 2005.
[12] H. Nagamochi and K. Okada, "Approximating the minmax rooted-tree cover in a tree," Information Processing Letters, vol. 104, no. 5, pp. 173178, 2007.
[13] S. Karakawa, E. Morsy, and H. Nagamochi, "Minmax tree cover in the euclidean space," in International Workshop on Algorithms and Computation, pp. 202-213, Springer, 2009.
[14] Z. Xu and Q. Wen, "Approximation hardness of min-max tree covers," Operations Research Letters, vol. 38, no. 3, pp. 169-173, 2010.
[15] Z. Xu, D. Xu, and W. Zhu, "Approximation results for a min-max location-routing problem," Discrete Applied Mathematics, vol. 160, no. 3, pp. 306-320, 2012.
[16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, Introduction to Algorithms. McGraw-Hill, second ed., 2001.
[17] M. D. Penrose, "A strong law for the longest edge of the minimal spanning tree," The Annals of Probability, vol. 27, no. 1, pp. 246-260, 1999.


[^0]:    ${ }^{1}$ The authors of [6] also studied the capacitated rooted min-max cycle cover problem and proposed an algorithm with approximation ratio of $7+\epsilon$.

[^1]:    ${ }^{2}$ Here, we implicitly assume that the amount of time it takes to inspect the sites is negligible compared to the travel time. However, the case with non-zero weights associated with vertices can be handled by constructing a new graph with only edge weights as described in [15].

[^2]:    ${ }^{3}$ In some cases, it may make sense to limit the number of robots that can be stored at each robot. This would give rise to a capacitated rooted min-max cover problem. Here, we assume no such constraints on depot sizes.

[^3]:    ${ }^{4}$ The forests we construct are in fact tree covers for $G$. For consistency, we shall refer to them as forests in the remainder of the paper.

