

Approximation Algorithms for Movement Repairmen

MohammadTaghi Hajiaghayi^{1*}, Rohit Khandekar², M. Reza Khani^{3*}, and Guy Kortsarz^{4**}

¹ University of Maryland, College Park, MD; and AT&T Labs, Florham Park, NJ. hajiagha@cs.umd.edu.

² Knight Capital Group, Jersey city, NJ. rkhandekar@gmail.com.

³ University of Maryland, College Park, MD. khani@cs.umd.edu.

⁴ Rutgers University, Camden, NJ. guyk@camden.rutgers.edu.

Abstract. In the *Movement Repairmen (MR)* problem we are given a metric space (V, d) along with a set R of k repairmen r_1, r_2, \dots, r_k with their start depots $s_1, s_2, \dots, s_k \in V$ and speeds $v_1, v_2, \dots, v_k \geq 0$ respectively and a set C of m clients c_1, c_2, \dots, c_m having start locations $s'_1, s'_2, \dots, s'_m \in V$ and speeds $v'_1, v'_2, \dots, v'_m \geq 0$ respectively. If t is the earliest time a client c_j is collocated with any repairman (say, r_i) at a node u , we say that the client is served by r_i at u and that its latency is t . The objective in the (SUM-MR) problem is to plan the movements for all repairmen and clients to minimize the sum (average) of the clients latencies. The motivation for this problem comes, for example, from Amazon Locker Delivery [Ama10] and USPS gopost [Ser10]. We give the first $O(\log n)$ -approximation algorithm for the SUM-MR problem. In order to solve SUM-MR we formulate an LP for the problem and bound its integrality gap. Our LP has exponentially many variables, therefore we need a separation oracle for the dual LP. This separation oracle is an instance of Neighborhood Prize Collecting Steiner Tree (NPCST) problem in which we want to find a tree with weight at most L collecting the maximum profit from the clients by visiting at least one node from their neighborhoods. The NPCST problem, even with the possibility to violate both the tree weight and neighborhood radii, is still very hard to approximate. We deal with this difficulty by using LP with geometrically increasing segments of the time line, and by giving a *tricriteria* approximation for the problem. The rounding needs a relatively involved analysis. We give a constant approximation algorithm for SUM-MR in Euclidean Space where the speed of the clients differ by a constant factor. We also give a constant approximation for the makespan variant.

1 Introduction

In the well-known *Traveling Repairman (TR)* problem, the goal is to find a tour to cover a set of clients such that the sum of latencies seen by the clients is minimized. The problem is also known as the *minimum latency problem*, see [GK98], the *School-bus driver problem*, see [WAW93] etc. This problem is well studied in the operations research literature and has lots of applications in the real world, see for example [BYCR93]. The problem is NP-Hard even in tree metrics [Sit06]. Blum *et al.* [BCC⁺94] give the first constant-factor approximation algorithm for the TR problem. They also observe that there is no PTAS ($(1 + \epsilon)$ -approximation algorithm for an arbitrary constant $\epsilon > 0$) for the problem unless $P = NP$. After a sequence of improvements, Chaudhuri et al. [CGRT03]

* Supported in part by NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, DARPA/AFOSR grant FA9550-12-1-0423, and a University of Maryland Research and Scholarship Award (RASA).

** Supported in part by NSF award number 434923

give a 3.59-approximation algorithm for TR which is the current best approximation factor for this problem.

Fakcharoenphol *et al.* [FHR03] generalize the TR problem to the k -Traveling Repairman (k -TR) problem in which instead of one repairman, we can use k repairmen to service the clients where all the repairmen start from the same depot. They give a 16.994-approximation algorithm for k -TR. Chekuri and Kumar [CK04] give a 24-approximation algorithm for the “multiple-depot” version of the k -TR problem where the repairmen can start from different depots.

Friggstad *et al.* [FSS13] are first to solve the TR problem by formulating an LP. They give an $O(\log n)$ -approximation for this problem. Their work is significant as it is the first LP approach to approximate the problem. Chakrabarty and Swamy [CS11] give a constant-factor approximation algorithm for the TR problem by introducing two new LPs.

We generalize [CS11] for the k -TR problem by allowing the repairmen to start from different starting depots and to have different speeds. More importantly, we give the clients ability to move with different speeds, which makes the problem significantly harder. We formally define SUM-MR as follows.

Definition 1. *In the SUM-MR problem the inputs are given as follow.*

- A metric space $\mathcal{M} = (V, d)$ where V is the set of nodes and $d : (V \times V) \rightarrow \mathbb{Q}^+$ is the distance function.
- A set R of k repairmen r_1, r_2, \dots, r_k . Each repairman r_i has a start depot $s_i \in V$ and speed $v_i \in \mathbb{Q}^+$.
- A set C of m clients c_1, c_2, \dots, c_m . Each client c_j has a start location $s'_j \in V$ and speed $v'_j \in \mathbb{Q}^+$.

A solution to the problem consists of the following.

- A pair (u_j, t_j) for each client c_j such that c_j can reach node u_j by time t_j considering its speed v'_j (i.e., $d(s'_j, u_j) \leq v'_j \cdot t_j$). Here t_j is the latency of u_j .
- A path p_i for each repairman r_i . In general p_i may not be a simple path and can contain a node or an edge multiple times. Repairman r_i can travel along p_i with any speed starting from zero up to its maximum speed v_i .
- For each pair (u_j, t_j) assigned to client c_j there has to be at least one repairman (r_i) such that r_i visits u_j at time t_j when it travels path p_i .

The objective for SUM-MR is to minimize $\sum_{j=1}^m t_j$.

The problem is very natural and is also motivated by the following real-world scenario. Amazon Locker Delivery is an optional shipping method in Amazon. In this method clients have an option to select a certain locker location to pick up their purchased items. Afterwards, Amazon puts the items into a locker in the specified location and sends the locker number and its key combination, to the customer. The package can be picked up by the customer who can go to the locker location by her own means. A very similar delivery option is also offered by the United States Postal Service which is known as gopost [Ser10].

Our algorithms can be used directly in order to plan the movements to minimize the average latency or the maximum latency. Here the locations of the Amazon stores, clients’ homes, and locker locations can be thought as the nodes of the metric space in SUM-MR and the repairmen are the shipping vehicles starting from the Amazon stores with different speeds. Moreover, we can take as the input how customers are going to pick up their packages (*e.g.* by a car, public transportation, bike,

etc.) which realizes the different speeds for the clients. Note that unlike SUM-MR in this scenario it is not necessary for both a repairman and a client to meet at the same node and the same time in order to serve; but if a repairman visits a node at time t a client can visit the node at any time after t and still gets served. We formalize these methods of serving and show that the difference in the objective of SUM-MR for the two methods is at most 3 in Section 4.1.

In Section 2 we describe connection of our problem to the movement framework, neighborhood TSP problems, and orienteering problems respectively. In Section 3 we give the outlines of our techniques and summarize all our results. Section 4 contains the detailed explanation of our method to solve SUM-MR in three subsections. Subsection 4.1 contains the necessary preliminaries, in Subsection 4.2 we give our LP formulation, in Subsection 4.3 we solve the LP, and finally in Subsection 4.4 we show how to round a fractional solution to the LP to get an integral solution to SUM-MR. It turns out that the separation oracle for our LP is a generalization of the Neighborhood TSP problem (to be defined in Subsection 2.2), we give the results related to the separation oracle problem in Section 5. In Section 6 we give our result for the Euclidean space. Section 7 contains all the materials related to the Max-MR problem in which instead of minimizing the sum of the latencies seen by the clients we want to minimize the latency of the last client that is served.

2 Connection of MR to the Other Class of Problems

Chakrabarty and Swamy [CS11] define a general problem called Minimum Latency Uncapacitated Facility Location (MLUFL). In MLUFL we are given a set F of n facilities with opening costs $\{f_i\}$, a set D of m clients, a root node r , and connection costs $\{c_{ij}\}$ for connecting client j to facility i . The objective is to select a subset F' of facilities to open, find a path p starting from r to visit all the facilities in F' , and assign each client j to a facility $\Phi(j)$, to minimize $\sum_{i \in F'} f_i + \sum_{j \in D} (c_{\Phi(j)j} + t_j)$ where t_j is the distance of $\Phi(j)$ from r in the path p . In the *related* MLUFL problem the connection cost $c_{\Phi(j)j}$ is the distance of client j to facility $\Phi(j)$ in the metric space. Chakrabarty and Swamy [CS11] provide a constant factor approximation algorithm for the related MLUFL problem. They also generalize it to the case when instead of one activating path p , we can have k activating paths. This result gives a constant factor approximation algorithm for the special case of SUM-MR when all the repairmen have the same speed and start from the same depot. In SUM-MR we want to minimize $\sum_{j \in D} \max(c_{\Phi(j)j}, t_j)$ as opposed to $\sum_{j \in D} (c_{\Phi(j)j} + t_j)$ in related MLUFL where the facility opening costs are zero; but note that they have a multiplicative difference of at most 2.

In the rest of this section we show how the MR problem is related to the other well-studied problems in computer science. In particular we show its connections to the movement framework, neighborhood TSP problems, and Orienteering problems.

2.1 Connection to the movement framework

SUM-MR can be defined under the *movement* framework first introduced by Demaine et al. [DHM⁺09b]. In the movement framework, we are given a general weighted graph and pebbles with different colors are placed on the nodes of the graph forming a starting configuration. The goal is to move the pebbles such that the final configuration of the pebbles meets a given set of properties. We can think of the repairmen as blue pebbles and the clients as red pebbles on the metric completion of the graph. Each red or blue pebble can move with different speeds. The latency of a

red pebble is the earliest time it is collocated with a blue pebble. The objective in SUM-MR is to minimize the total latency.

The paper by Demaine et al. [DHM⁺09b] inspired several other papers on the movement problems and there are a handful of recent approximation algorithms for them as well. Of primary relevance to this paper, is the work of Friggstad and Salavatipour [FS08] who consider minimizing movement in the facility location setting in which both facilities and clients are mobile and can move. The quality of a solution can be measured either by the total distance (average distance) clients and facilities travel or by the maximum distance traveled by any client or facility. They obtain constant-factor approximation algorithms for these problems and recently, Friggstad *et al.* [AFS13] improve the constant factor to $(3 + \epsilon)$ for a slightly more general problem. Very recently, Berman, Demaine, and Zadimoghaddam [BDZ11] obtain a constant-factor approximation algorithm for minimizing maximum movement to reach a configuration in which the pebbles form a connected subgraph. This result is very interesting since Demaine et al. [DHM⁺09b] show with the total sum movement object function, the connectivity movement problem is $\Omega(n^{1-\epsilon})$ inapproximable, and for which there is only an $\tilde{O}(n)$ -approximation algorithm. Demaine, Hajiaghayi, and Marx [DHM09a] consider the problem when we have a relatively small number of mobile agents (e.g., a team of autonomous robots, people, or vehicles) moving cooperatively in a vast terrain or complex building to achieve some task. They find optimal solutions for several movement problems in polynomial time where the number of pebbles is constant. Finally there are various specific problems considered less formally in practical scenarios [CHP⁺04a, CHP⁺04b, HAB⁺03, LaV06, RW95, SPS03, DMM⁺97, JBQZ04, SW01, JQQ⁺03].

2.2 Connection to the Neighborhood TSP problem

Suppose we guess in advance that a client c is going to be served by time t_c ⁵. Client c can go to any node which is reachable from its starting location by time t_c in order to be served. The set of such reachable nodes defines a neighborhood for each client. On the other hand, repairmen’s task is to visit a node in each client’s neighborhood in order to serve him. In fact, giving the clients ability to move in SUM-MR can be thought as assigning a neighborhood to each client.

Neighborhood problems are well-studied in the theoretical computer science community. A well-studied problem in this category is the Neighborhood TSP (NTSP) problem where we are given a set of clients each with a neighborhood and the goal is to find a minimum tour to visit at least one node from the neighborhood of each client. Visiting a neighborhood instead of a node makes the problem considerably harder; we will show an $\Omega(\log^{2-\epsilon} n)$ -hardness for the NTSP. This hardness is indeed the source of subtlety in SUM-MR.

NTSP is especially studied for the Euclidean Space in the CS community, as it has routing-related and VLSI design applications [Mit00, Chapter 15] [RW90]. Euclidean Neighborhood TSP (ENTSP) is known to be APX-hard [DBGK⁺05, SS06, EFS06]. A neighborhood is called fat if we can fit a disc inside the neighborhood such that the radius of the disc is at least a constant fraction of the radius of the neighborhood, where the radius of a neighborhood is half of the distance between its two farthest points. A PTAS for ENTSP is known, when the neighborhoods are fat, roughly the same size (their sizes differ in at most a constant factor), and overlap with each other in at most a constant number of times [DM01, FG04, Mit07]. When neighborhoods are connected and disjoint Mitchell [Mit10] gives a constant-factor approximation algorithm for the problem. Elbassioni *et al.* [EFS06] give a constant-

⁵ We are able to estimate time t_c for each client with the help of our LP for SUM-MR.

factor approximation for ENTSP when the neighborhoods are roughly the same size, convex, fat, and can intersect.

In this paper we generalize ENTSP to the prize-collecting version. In the prize collecting ENTSP we are given a budget B and each neighborhood is assigned a profit, the objective is to find a tour of maximum length B to maximize the sum of profits of the neighborhoods the tour intersects. We give a constant-factor approximation algorithm for the prize-collecting ENTSP when the neighborhoods are fat, roughly the same size, and can intersect. As a result of this algorithm, we obtain a constant factor approximation algorithm to SUM-MR in Euclidean Space.

2.3 Connection to the Orienteering Problem

It turns out that the separation oracle for the dual of our LP is closely related to the *orienteering problem*. In the orienteering problem we are given a graph $G = (V, E)$, two nodes $s, t \in V$, and a length bound B and the goal is to find a tour starting from s and ending in t with length at most B that visits maximum number of nodes. Orienteering is shown to be NP-Hard via an easy reduction from the TSP problem and it is also APX-hard [BCK⁺07]. Blum *et al.* [BCK⁺07] give the first constant-factor approximation algorithm with ratio 4 for orienteering⁶ which is improved to 3 by Bansal *et al.* [BBCM04]. Chekuri *et al.* [CKP08] give a $(2 + \epsilon)$ -approximation algorithm for the problem in the undirected graphs which is the current best approximation factor for orienteering and an $O(\log^2 \text{OPT})$ -approximation algorithm for the directed case. Nagarajan and Ravi [NR1] give a $O(\log^2 n / \log \log n)$ -approximation algorithm for the directed orienteering problem. Therefore, depending on OPT the best approximation factor for the directed orienteering problem is $O(\log^2 \text{OPT})$ or $O(\log^2 n / \log \log n)$.

Another closely related work is the submodular orienteering problem introduced by Pal and Chekuri [CP05] in undirected graphs. In this problem instead of maximizing the number of visited nodes, we want to maximize a submodular function of the set of visited nodes. Later in this paper (Section 5.4) we show that using this result we can obtain an interesting approximation algorithm for the separation oracle of our dual LP and consequently a quasi polynomial time approximation algorithm for SUM-MR with the same $O(\log n)$ factor.

3 Results and techniques

In this section we summarize all our results along with the overview of their proofs.

Our main result is an $O(\log n)$ -approximation algorithm for SUM-MR. More precisely we prove the following theorem.

Theorem 1. *There is an $O(\log n)$ -approximation algorithm for the SUM-MR problem which also upper bounds the integrality gap of its LP formulation.*

We present the novel properties (in this context) of our techniques. First we relax conditions on serving the clients. If a client collocates with a repairman at a certain node of the metric space during the movements we say it gets served *perfectly*. On the other hand if a client visits a node through which a repairman has passed no later than the arrival of the client, we say it gets served *indirectly*. We design a procedure that transforms any solution to the SUM-MR problem where the

⁶ Their algorithm is not point-to-point, *i.e.*, their algorithm cannot control what would be the ending vertex.

clients are served indirectly to a solution where all the clients are served perfectly by increasing the total latency with a multiplicative factor at most 3. We solve SUM-MR for the case when we serve the clients indirectly and use this procedure to serve the clients perfectly. We give an LP formulation for SUM-MR (to serve the clients indirectly) and bound its integrality gap. However there are two major challenges we must overcome in order to do so. First, solving the LP which has exponentially many variables and second, rounding a solution to the LP efficiently to an integral solution.

In order to solve the LP we need a separation oracle for its dual which turns out to be the following problem.

Definition 2. *Neighborhood Prize Collecting Steiner Tree (NPCST): An instance of the NPCST problem consists of an ordered tuple (V, d, r, C, L) where V is the set of nodes, d is a metric distance function on set V , $r \in V$ is the root node, C is the set of clients, and L is the cost budget. Each client $c \in C$ is associated with a profit θ_c and a neighborhood ball $\mathcal{B}(c, t_c)$ which contains all the nodes u with $d(u, c) \leq t_c$. The goal is to find a Steiner tree T_{OPT} rooted at r such that $\text{cost}(T_{\text{OPT}}) \leq L$ and the sum of the profits of the clients whose \mathcal{B} -balls hit T_{OPT} is maximized.*

The vehicle routing problems become significantly harder when instead of visiting a node, it is sufficient to visit a neighborhood around it. For example in the Neighborhood Steiner Tree (NST) problem, we are given a graph G with a set of clients C where each client c is associated with a neighborhood ball. The objective for NST is to find a tree T with minimum weight that serves at least one node from each client's neighborhood ball. We will prove the following hardness result about the NST problem which shows the source of difficulty in our problem.

Theorem 2. *There is no $O(\log^{2-\epsilon} n)$ -approximation algorithm for the NST problem unless NP has quasi-polynomial Las-Vegas algorithms.*

To avoid this hardness we allow relaxing the NPCST constraints. More formally we accept a tri-criteria approximation algorithm for NPCST as our separation oracle defined formally below.

Definition 3. *A (σ, ϕ, ω) -approximation algorithm for the instance (V, d, r, C, L) of the NPCST problem finds a Steiner tree T with the following properties. T is said to hit a client c with \mathcal{B} -ball $\mathcal{B}(c, t_c)$ if T has at least one node in $\mathcal{B}(c, t_c \cdot \sigma)$, the weight of T is at most $\phi \cdot L$, and sum of the profits of the clients got hit by T is at least $\frac{1}{\omega} \text{OPT}$ where OPT is the amount of profit an optimum tree collects with no violation in any bound.*

Accepting a tri-criteria approximation algorithm for NPCST has two benefits. Firstly, it reduces the difficulty of solving the NPCST problem to avoid the hardness results similar to Theorem 2. Secondly, later when we transform the solution of the algorithm to a solution of SUM-MR, it allows the approximation factor on the traveling time for a client and a repairman to reach to a certain node (latency) to get split between both the client (violating its neighborhood) and the repairman (violating the weight of the tree). However a solution to NPCST resulting from a tri-criteria approximation algorithm is harder to transform to a solution of SUM-MR. We prove the following general theorem to transform any tri-criteria approximation algorithm to the NPCST problem to an efficient approximation algorithm for SUM-MR.

Theorem 3. *Given a (σ, ϕ, ω) -approximation algorithm for NPCST, we can find an $O(\max(\sigma, \phi) \cdot \omega)$ -approximation algorithm for SUM-MR.*

Proving Theorem 3 has two parts. The first part is to use the tri-criteria approximation algorithm to find a feasible solution for our LP and the second part is to round the feasible solution. For the

first part we introduce a new relaxed LP for SUM-MR to absorb the violations of the tri-criteria approximation algorithm while keeping the optimal value of the relaxed LP to be at most the optimal value of the original LP. Then we show that using a (σ, ϕ, ω) -approximation algorithm for NPCST, we can find a feasible solution to the relaxed LP with the objective value at most the optimum value of the original LP.

For the second part of the proof, we round the feasible solution found in the previous part to an integral solution for SUM-MR with the total latency at most $O(\max(\sigma, 2\phi) \cdot \omega)$ times the optimal value of our LP. Our algorithm (later given in Figure 1) for the rounding part is easy to state and implement but needs a relatively complicated analysis. The algorithm runs in several steps where each step represents a time-stamp. The time-stamps increase geometrically, *i.e.*, the time-stamp of a step is twice as the time-stamp of the previous step. At each step we randomly select a tour for each repairman from the set of all tours with the length at most the time-stamp times the repairman's speed. The random selection is done using the LP values. The output of our algorithm is the concatenation of all the tours selected at each step. The idea for the analysis of our algorithm is as follows. Let the time-stamp for a certain step be 2^a and F be the (fractional) number of clients that are served by time 2^a according to the LP values. We show the expected number of clients that our algorithm serves in the step is at least $\frac{3F}{4}$. We show that this condition is enough to bound the total latency of the clients. Finally, we show that our algorithm can be derandomized. The derandomization is done by a recursive algorithm which takes an arbitrary subset (R') of R and selects a path for a repairman r in R' and calls itself with parameter $R' \setminus \{r\}$. It selects a path for r which covers the maximum number of clients from the set of clients that are served fractionally by repairmen of R' in the LP solution but not served by the paths we have selected till now. We prove that the greedy algorithm serves at least $\lceil \frac{3F}{4} \rceil$ clients by induction on the size of set R' .

We prove the following theorem about the NPCST problem for the general metrics which is of independent interest and non-trivial. In order to prove Theorem 1, we plug this result about the NPCST problem in Theorem 3.

Theorem 4. *There is an $(O(\log n), O(\log n), 2)$ -approximation algorithm for the NPCST problem in general metrics.*

Remember that in the NPCST problem we have to find a tree T to maximize the number of clients whose \mathcal{B} -balls contain a node of T . To prove the above theorem we embed the graph into a distribution of tree metrics [FRT04]. Note that the \mathcal{B} -balls for the clients are not preserved in the tree metrics. We define a new problem on the tree metrics as follows. Given a budget L' we want to find a tree T with weight at most L' to maximize the size of the set of served clients C' where T serves set C' if the sum of distances of the clients in C' to T is at most L' . We solve the new problem efficiently in the tree metrics with a dynamic programming. Finally we show by violating the radii of the \mathcal{B} -balls of the clients by a constant factor, T actually serves a good fraction of the clients in C' (by Markov's Inequality) through hitting their \mathcal{B} -balls.

An anonymous referee pointed out that we can obtain an $(O(1), O(1), O(\log n))$ -approximation algorithm for the NPCST problem in general metrics using the ideas in [GKK⁺01] and [SK04]. This algorithm is interesting since as opposed to the algorithm of Theorem 4 there are constant violations on the cost of the tree and the neighborhoods' radii but it collects an $O(\log n)$ fraction of the optimal profit. Note that by plugging this algorithm into Theorem 3 we get the same result for the NPCST problem as in Theorem 1. The description of the algorithm and an outline of the proof given by the referee, is brought in Section 5.3. Another anonymous referee pointed out that we can obtain a quasi polynomial time $(1, 1, O(\log n))$ -approximation algorithm for the NPCST problem using the

algorithm of [CP05]. Although the running time is quasi polynomial but, interestingly, there is no violation on the budget and the neighborhood radii. The outline of this is brought in Section 5.4.

Motivated from the application of SUM-MR in Amazon Locker Delivery [Ama10] and USPS gopost [Ser10] which occurs in the Euclidean space, we prove the following theorem to get a constant-factor approximation algorithm for SUM-MR in the Euclidean space. Here, we assume that one can freely move in the Euclidean plane. In fact the neighborhood problems are especially studied in the geometric settings. The usual assumption in the neighborhood TSP problems for getting a constant factor approximation (see Section 2.2) is to assume the radius of the biggest neighborhood is at most a constant factor larger than the smallest one. We plug the following theorem in Theorem 3 to get a constant factor approximation algorithm for SUM-MR. Here the radius constraints means the maximum speed of the clients is at most a constant factor larger than the minimum speed which is an acceptable constraint for the package delivery application motivating SUM-MR.

Theorem 5. *There is an $(O(P), O(1), O(1))$ -approximation algorithm for the NPCST problem in the Euclidean space where the radius of the greatest neighborhood is at most P times larger than the radius of the smallest neighborhood.*

Our last result is a constant-factor approximation algorithm for Max-MR.

Theorem 6. *There is a constant-factor approximation algorithm for the Max-MR problem in general metrics when the repairmen have the same speed.*

4 The Sum Movement Repairmen Problem

4.1 Preliminaries

First note that without loss of generality we assume that the minimum latency of a client is at least 1. Here, if the starting depot of a client collocates with a repairman's starting depot, we ignore this client as its latency is zero. Let v^* be the maximum speed of all the clients and repairmen and f^- be the smallest distance between two nodes in metric d . We multiply all the edges of the graph by v^*/f^- and throughout the paper assume that the minimum service time a client can see is at least 1.

Now we relax the conditions that have to be met in order to serve the clients. In our standard definition of SUM-MR a client gets served with latency t , if it collocates with a repairman in a node at time t . Here we call that the client is served *perfectly* with latency t . If a repairman visits a node u at time t' and a client visits u at time $t \geq t'$, we say the client is served *indirectly* with latency t . Now we define an *indirect solution* (sol) to be similar to Definition 1 except that the last condition is changed to the following. For each pair (u_j, t_j) assigned to client c_j there has to be at least one repairman (r_i) such that r_i visits u_j at time t'_j when $t'_j \leq t_j$.

The following lemma shows that an indirect solution can be changed to a standard solution such that all the clients get served perfectly while the total latency does not change by more than a constant factor.

Lemma 1. *Suppose an indirect solution (sol) to SUM-MR has sum of latencies l where all the clients are served indirectly, then sol can be transformed to a standard solution (sol') in which all the clients are served perfectly with sum of latencies at most $3 \cdot l$.*

Proof. Remember in Definition 1 that indirect solution sol assigns a path p_i to each repairman r_i and a node u_j to each client c_j such that c_j can go to u_j by time t_j while a repairman has visited u_j before or at time t_j . Suppose repairman r_i can travel p_i by time t_i considering its speed. In other words, the length of p_i is at most $v_i \cdot t_i$ where v_i is the speed of r_i . When r_i serves the clients indirectly it is better for him to travel p_i as fast as possible and does not wait for the clients to arrive since the clients can arrive in the nodes of p_i later and are still served indirectly.

We plan the movements in sol' as follows. The movement for the clients are the same as sol , each client c_j is assigned to same node u_j as in sol and go to the assigned node by time t_j . Each repairman r_i starts from depot s_i (its starting node) and goes one unit of time along path p_i with its maximum speed v_i and comes back to s_i , we refer to this as round 1, then it goes 2 units of time and comes back to s_i (round 2). In general at each round ℓ it travels $2^{\ell-1}$ units of time along p_i and comes back. If at round ℓ the given time $2^{\ell-1}$ is enough to travel p_i completely, repairman r_i travels p_i completely, stays there until $2^{\ell-1}$, and then comes back to s_i ⁷.

Now we prove that if a client is served indirectly with latency q in sol , it will be served perfectly with latency $3q$ in sol' . Suppose an arbitrary client c_j is served indirectly with r_i at time q in a node u_j of path p_i . Note that q either represents the time when both r_i and c_j arrive at u_j or the time when c_j arrives at u_j but r_i is already passed u_j . In sol' , when c_j reaches u_j it waits for repairman r_i to visit u_j after or at time q during its back and forth travels.

Note that each round ℓ takes 2^ℓ units of time, for going forward and back, and round ℓ finishes after time $2^{\ell+1} - 2$, the geometric sum over the previous rounds. Now consider the finishing time ($t_{\lfloor \log q \rfloor}$) of round $\lfloor \log q \rfloor$ which is $2^{\lfloor \log q \rfloor + 1} - 2$. This value is at least $q - 2$ and at most $2q$. If $t_{\lfloor \log q \rfloor} \geq q$ then in the next round r_i by traveling at most q units of time visits u_j because it was served perfectly by time q . The latency for this case is $3q$. If $q - 2 \leq t_{\lfloor \log q \rfloor} < q$ then the length of the next round is at least q and at most $q + 2$. Therefore, r_i visits u_j in the next run, possibly on the way back, and the latency in this case is at most $2q + 2$. Remember that the earliest visiting time cannot be less than 1. Hence the latency for this case is also upper bounded by $3q$. \square

We focus on finding an indirect solution to the SUM-MR problem where the clients are served indirectly and then transform it to a standard solution which serves the clients perfectly using Lemma 1. Therefore, from now on whenever we use serving or solution we mean an indirect solution or serving indirectly.

We start with some important definitions.

Definition 4. Let neighborhood $\mathcal{B}(c, t_c)$ denotes the set of all nodes whose distances from s'_c , the starting node of client c , are at most t_c .

Definition 5. Let $\mathcal{P}(r, t_r)$ denotes the set of all non-simple paths (i.e., they can visit nodes or edges multiple times) with length at most t_r starting from s_r , the starting depot of r .

Using the above two definitions we can formalize the notion of serving as follows.

Definition 6. We say a repairman r serves client c or client c gets served by r at time t if the path selected for r hits neighborhood $\mathcal{B}(c, v'_c \cdot t)$ by time t where v'_c is the speed of client c .

⁷ In fact if r_i returns as soon as it finishes traveling p_i , the latencies might decrease. But we avoid this for the sake of simplicity and the fact that in worst case it does not matter.

Remember that the minimum latency is at least 1. Let T be the largest service time a client can see, here we upper bound T to be $\frac{2 \cdot \text{MST}(G)}{\min_i v_i}$ which is the units of time required to travel all the edges by the slowest repairman and hence serving all the clients. We use set $Q = \{1, 2, \dots, 2^i, \dots, 2^{\lceil \log T \rceil}\}$ to index geometrically increasing time-stamps. Note that we have $\lceil \log T \rceil$ elements in Q and hence its size is polynomially bounded by the size of input.

4.2 LP formulation for Sum-MR

In this section we introduce an LP formulation for the SUM-MR problem and show how to solve this LP approximately. We use the following LP for SUM-MR inspired by the ideas from LPs introduced by Chakrabarty and Swamy [CS11].

$$\begin{aligned}
\min \quad & \sum_{c \in C} \sum_{t \in Q} t \cdot y_{c,t} && \text{(PLP)} \\
\text{s.t.} \quad & \sum_{p \in \mathcal{P}(r, v_r \cdot t)} x_{r,p,t} \leq 1 && \forall t \in Q, \forall r \in R \quad (1) \\
& \sum_{r \in R} \sum_{p \in \mathcal{P}(r, v_r \cdot t): p \cap \mathcal{B}(c, v'_c \cdot t) \neq \emptyset} x_{r,p,t} \geq \sum_{t' \leq t} y_{c,t'} && \forall c \in C, \forall t \in Q \quad (2) \\
& \sum_{t \in Q} y_{c,t} \geq 1 && \forall c \in C \quad (3) \\
& x, y \geq 0 && (4)
\end{aligned}$$

The variable $x_{r,p,t}$ is the indicator variable showing whether repairman r travels path $p \in \mathcal{P}(r, v_r \cdot t)$ completely by time t . Note that if p is in set $\mathcal{P}(r, v_r \cdot t)$, from the definition of $\mathcal{P}(r, v_r \cdot t)$, r can complete traveling p within time t . Variable $y_{c,t}$ is the indicator variable showing if client c is served at time t .

Constraints (3) guarantee that every client gets served. Constraints (1) require each repairman r to travel at most one path by time t . The amount $\sum_{t' \leq t} y_{c,t'}$ shows the fraction of service, client c demands until time t and the amount $\sum_{r \in R} \sum_{p \in \mathcal{P}(r, v_r \cdot t): p \cap \mathcal{B}(c, v'_c \cdot t) \neq \emptyset} x_{r,p,t}$ shows the fraction of service c gets from the repairmen until time t . Constraints (2) guarantee that the total service from all the repairmen is at least as large as the demand from client c . Note that Constraints (2) just require the total demand by client c should be served by the repairmen at anytime before time t which is the case for serving indirectly.

Note that we only consider times that are in the set Q but in a solution of SUM-MR the clients may be served at any time.

Lemma 2. *The optimal value of PLP is at most twice the optimal solution of SUM-MR.*

Proof. Before proving the lemma, note that by an appropriate scaling we assumed that the smallest element of Q (the smallest latency seen by the clients) is one. Note that this scaling does not make the algorithm pseudo polynomial since the time-stamps in Q grow exponentially.

Intuitively the factor two comes from the fact that we only consider the powers of 2 as the time-stamps in set Q . We show that an optimal solution (\hat{s}) to the SUM-MR problem with total latency ℓ^* can be transformed to a feasible solution (\hat{x}, \hat{y}) to PLP with the objective value at most $2\ell^*$. For each repairman r_i and time $t_i \in Q$ let p_i be the path travelled by r_i by time t_i in \hat{s} , we

set $\hat{x}_{r_i, p_i, t_i} = 1$. If client c is served at time t in \hat{sol} , we set $y_{c, t'} = 1$ where $t' = 2^{\lceil \log t \rceil}$. We set all the other non-set entries of (\hat{x}, \hat{y}) to zero. As \hat{sol} is a valid solution to SUM-MR one can check that (\hat{x}, \hat{y}) satisfies all the constraints of **PLP**.

If a client (c) is served at time t in \hat{sol} , c contributes at most t' to the objective value of (\hat{x}, \hat{y}) in **PLP**, where $t' < 2t$ since $2^{\lceil \log t \rceil} < 2^{\log t + 1}$. Thus the optimal value of **PLP** is at most twice the total latency of an optimal solution for SUM-MR. \square

In the next subsection we show that we are able to find a solution to the SUM-MR problem which has total latency at most $O(\log n)$ times the optimum value of **PLP** which along with Lemma 2 upper bounds the integrality gap of the LP, where $n = |V|$ is the number of nodes in the metric space.

4.3 Solving PLP in polynomial time

The first difficulty to solve **PLP** is that it has exponentially many variables. In order to solve the LP we formulate its dual. The dual LP has exponentially many constraints but polynomially many variables therefore we need a separation oracle for the constraints in order to solve the dual LP in polynomial time. The dual LP for **PLP** is as follows.

$$\max \quad \sum_{c \in C} \lambda_c - \sum_{r \in R, t \in Q} \beta_{r, t} \quad (\mathbf{DLP})$$

$$s.t. \quad \sum_{c: p \cap \mathcal{B}(c, v'_c \cdot t) \neq \emptyset} \theta_{c, t} \leq \beta_{r, t} \quad \forall r \in R, \forall t \in Q, \forall p \in \mathcal{P}(r, v_r \cdot t) \quad (5)$$

$$\lambda_c \leq t + \sum_{t \leq t'} \theta_{c, t'} \quad \forall c \in C, \forall t \in Q \quad (6)$$

$$\lambda, \beta, \theta \geq 0 \quad (7)$$

We have exponentially many Constraints (5), therefore we need a separation oracle for them in order to use Ellipsoid algorithm to solve **DLP**. Given a candidate solution (λ, β, θ) for any repairman $r_i \in R$ and time-stamp $t \in Q$ we define Separation Oracle Problem $SOP(r_i, t)$ as follows. Assume that each client c has profit $\theta_{c, t}$ and \mathcal{B} -ball $\mathcal{B}(c, v'_c \cdot t)$. The objective is to find a path in $\mathcal{P}(r_i, v_i \cdot t)$ (has maximum length $t \cdot v_i$) which collects the maximum profit. Here a path collects the profit of any client whose \mathcal{B} -ball is hit by it. If for all $r_i \in R$ and $t \in Q$ the optimal path collects at most $\beta_{r_i, t}$ profits, there is no violating constraint and (λ, β, θ) is a feasible solution; otherwise there exists a separating hyperplane.

The separation oracle explained above is NP-Hard since it contains the orienteering problem as a special case where the radius of all the \mathcal{B} -balls are zero. Therefore, we can only hope for an approximate solution for the separation oracle unless $P = NP$.

Note that $SOP(r_i, t)$ is the same as instance $(V, d, r_i, C, t \cdot v_i)$ of the NPCST problem (Definition 2) except instead of finding an optimum tree we have to find an optimum path. Because paths are the special cases of the trees, the optimum value for the NPCST instance is at least the optimum value of $SOP(r_i, t)$. Therefore, if we solve the NPCST instance we collect at least the same amount of profit. We will use the $(O(\log n), O(\log n), 2)$ -approximation algorithm in Theorem 4 to solve the NPCST instance and obtain a resulting tree. Then we transform it to a path by doubling the edges and taking an Eulerian tour which makes the length of the path to be at most 2 times larger than

the tree. In fact, we approximately solve $SOP(r_i, t)$ by violating the budget on the path, the radius of clients' \mathcal{B} -balls, and not collecting the maximum profit.

Due to all the violations explained above on the constraints of $SOP(r_i, t)$ we cannot bound the objective value of the feasible solution resulting from the $(O(\log n), O(\log n), 2)$ -approximation algorithm. To this end, we introduce a relaxation of **PLP** ($PLP^{(\mu, \omega)}$) in the following, when μ, ω are constant integers greater than or equal to 1.

$$\min \quad \sum_{c \in C} \sum_{t \in Q} t \cdot y_{c,t} \quad (PLP^{(\mu, \omega)})$$

$$s.t. \quad \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} x_{r,p,t} \leq \omega \quad \forall r \in R, \forall t \in Q \quad (8)$$

$$\sum_{r \in R} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} \geq \sum_{t' \leq t} y_{c,t'} \quad \forall c \in C, \forall t \in Q \quad (9)$$

$$\sum_{t \in Q} y_{c,t} \geq 1 \quad \forall c \in C \quad (10)$$

$$x, y \geq 0 \quad (11)$$

Constraint (8) is the same as Constraint (1) except instead of $\mathcal{P}(r, v_r \cdot t)$ we have $\mathcal{P}(r, \mu \cdot v_r \cdot t)$ which allows repairman r to take a path which is μ times longer than a path in $\mathcal{P}(r, v_r \cdot t)$. Moreover by putting ω instead of 1 we allow each repairman to take ω routes instead of one. Constraint (9) is the same as Constraint (2) except instead of $\mathcal{B}(c, v'_c \cdot t)$ we have $\mathcal{B}(c, \mu \cdot v'_c \cdot t)$ and instead of $\mathcal{P}(r, v_r \cdot t)$ we have $\mathcal{P}(r, \mu \cdot v_r \cdot t)$ which allow both the repairmen and clients to take paths that are μ times longer.

In the following lemma we show a (σ, ϕ, ω) -approximation algorithm for NPCST can be used to find a feasible solution to $PLP^{(\mu, \omega)}$ whose cost is at most the optimum solution of **PLP**. The proof of this lemma is relatively involved and is more general than a lemma used in [CS11].

Lemma 3. *Given a (σ, ϕ, ω) -approximation algorithm for NPCST, one can find a feasible solution to $PLP^{(\mu, \omega)}$ in polynomial time, where $\mu = \max(\sigma, 2 \cdot \phi)$, with objective value at most $\text{OPT}(1 + \epsilon)$ for any $\epsilon > 0$ where OPT is the optimal value of **PLP**.*

Proof. Consider the following polytope ($\text{DLP}(\xi; \mu, \omega)$).

$$\sum_{c \in C} \lambda_c - \sum_{r \in R, t \in Q} \beta_{r,t} \geq \xi \quad (12)$$

$$\sum_{c: p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} \theta_{c,t} \leq \frac{1}{\omega} \beta_{r,t} \quad \forall r \in R, \forall t \in Q, \forall p \in \mathcal{P}(r, \mu \cdot v_r \cdot t) \quad (13)$$

$$\lambda_c \leq t + \sum_{t \leq t'} \theta_{c,t'} \quad \forall c \in C, \forall t \in Q \quad (14)$$

$$\lambda, \beta, \theta \geq 0 \quad (15)$$

$\text{DLP}(\xi; \mu, \omega)$ is feasible if the optimal value of the dual LP of $PLP^{(\mu, \omega)}$ is greater than or equal to ξ since Constraints (13), (14), and (15) are the constraints of the dual of $PLP^{(\mu, \omega)}$ and Constraint

(12) lower bounds the objective value of it. Therefore $DLP(\xi; 1, 1)$ is feasible if the optimum value of **DLP** is greater than or equal ξ .

First we prove the following claim.

Claim 1. *Given a real value ξ and triple (β, λ, θ) as the candidate solution to $DLP(\xi; 1, 1)$, there is a polynomial time separation oracle that either: (1) shows $(\beta, \lambda, \theta) \in DLP(\xi; 1, 1)$, or (2) finds a hyperplane separating (β, λ, θ) and $DLP(\xi; \mu, \omega)$.*

Proof. First we check if triple (β, λ, θ) satisfies all Constraints (12), (14) and (15). The checks can be done in polynomial time as there are polynomially many Constraints (12), (14) and (15). If a constraint does not satisfy, then we find a hyperplane separating (β, λ, θ) and $DLP(\xi; \mu, \omega)$ and the claim follows.

We might have exponentially many Constraints (13). In order to check if all of Constraints (13) are satisfied, for every value $t \in Q$ and each repairman $r \in R$ we define instance $\mathcal{I}_{t,r} = (V, d, s_r, C, v_r \cdot t)$ of NPCST (see Definition 2) as follows. Node set V and metric d in $\mathcal{I}_{t,r}$ is the same as graph G in the input of SUM-MR, the root node s_r is the starting depot of r , the cost budget for the tree is $v_r \cdot t$ (v_r is the speed of r), and each client c has profit $\theta_{c,t}$ and neighborhood $\mathcal{B}(c, v'_c \cdot t)$. The separation oracle is the following. We run the given (σ, ϕ, ω) -approximation algorithm on $\mathcal{I}_{t,r}$ and find a tree whose cost is at most $\phi \cdot v_r \cdot t$ and collects at least $\frac{1}{\omega}$ fraction of the optimum profit while violating the radius of the clients' \mathcal{B} -ball by σ factor. Here by violating a \mathcal{B} -ball by factor σ we mean that instead of visiting a node within the \mathcal{B} -ball's radius, we visit a node that can be σ times farther. We transform the resulting tree to a path by doubling the edges and taking an Eulerian tour which makes the length of the tour at most $2 \cdot \phi \cdot v_r \cdot t$. Note that because $\mu = \max(\sigma, 2 \cdot \phi)$ the resulting tour is in $\mathcal{P}(r, \mu \cdot v_r \cdot t)$ and it collects the profit of client c by visiting a node in its $\mathcal{B}(c, \mu \cdot v'_c \cdot t)$ neighborhood ball.

If there exists repairman $r \in R$ and time $t \in Q$ such that the path resulting from the separation oracle collects profits greater than $\frac{1}{\omega} \beta_{r,t}$, the corresponding Constraint (13) gives a separating hyperplane and the claim follows. If not, we prove by contradiction that $(\beta, \lambda, \theta) \in DLP(\xi; 1, 1)$. If $(\beta, \lambda, \theta) \notin DLP(\xi; 1, 1)$ then at least one of the constraints of $DLP(\xi; 1, 1)$ has to not hold for (β, λ, θ) . As we check all Constraints (12), (14) and (15) for $DLP(\xi; \mu, \omega)$ and they are the same in $DLP(\xi; 1, 1)$ the violating constraint (\mathcal{V}) is one of the Constraint (13). Let the violating constraint \mathcal{V} happens for repairman r and time t and $p_{r,t}^*$ be the path returned by our separation oracle algorithm on $\mathcal{I}_{t,r}$. Constraint \mathcal{V} being a violating constraint means that there exists a path with length at most $v_r \cdot t$ which collects profits greater than $\beta_{r,t}$ without any violation in the clients' \mathcal{B} -ball. Therefore, as $p_{r,t}^*$ is the path resulting from the (σ, ϕ, ω) -approximation algorithm, $p_{r,t}^*$ is in $\mathcal{P}(r, \mu \cdot v_r \cdot t)$ (remember μ is at least $2 \cdot \phi$) and collects total profits at least $\frac{1}{\omega} \cdot \beta_{r,t}$ while violating the clients' \mathcal{B} -ball by a factor of at most μ (remember μ is at least σ). This means r, t , and $p_{r,t}^*$ is a hyperplane separating (β, λ, θ) and $DLP(\xi; \mu, \omega)$ which cannot happen. Therefore there is no such violating constraint \mathcal{V} and hence (β, λ, θ) is in polytope $DLP(\xi; 1, 1)$. \square

The following proposition is obtained by running Ellipsoid algorithm [GLS93] using the separation oracle of Claim 1.

Proposition 1. *For any value ξ , the Ellipsoid algorithm in polynomial time either shows $DLP(\xi; \mu, \omega)$ is empty or finds a feasible solution (β, λ, θ) in the polytope $DLP(\xi; 1, 1)$.*

We find the largest value (ξ^*) , within an error factor of at most $(1 + \epsilon)$, for which there exists a feasible solution $(\beta^*, \lambda^*, \theta^*) \in DLP(\xi^*; 1, 1)$ by a binary search over ξ on Proposition 1. Because

$\text{DLP}(\xi^*; 1, 1)$ is feasible, it means that **DLP** has the optimum value at least ξ^* (since $\text{DLP}(\xi^*; 1, 1)$ has the same constraints as (**DLP**) with an extra constraint to lower bound the objective value). Thus, from the LP duality theorem we conclude $\xi^* \leq \text{OPT}$ (Fact 1).

Because ξ^* is the largest value within a factor of $(1 + \epsilon)$ for which $\text{DLP}(\xi; 1, 1)$ is non-empty, Ellipsoid algorithm in Proposition 1 with $\xi = \xi^*(1 + \epsilon)$ terminates in polynomial time certifying infeasibility of $\text{DLP}(\xi^*(1 + \epsilon); \mu, \omega)$. Thus it generates a collection of constraints of type (13), (14), (15), and (12) which all together constitute an infeasible system of constraints. Lets denote this infeasible system of constraints by τ . Note that τ consists of polynomially many constraints as Ellipsoid algorithm of Proposition 1 runs in polynomial time and at each step it finds one violating constraint. We apply the Farkas's lemma to the constraints in τ .

Lemma 4 (Farkas). *For a matrix $A \in \mathbb{R}^{n \times m}$ and a vector $b \in \mathbb{R}^m$, exactly one of the following holds.*

1. $\exists w \in \mathbb{R}^m \geq 0$ such that $Aw \geq b$.
2. $\exists z \in \mathbb{R}^n \geq 0$ such that $A^T z \leq 0$ and $b^T z > 0$.

Here we embed the constraints of type (13), (14), and (15) of τ by $A^T z \leq 0$. More precisely z is concatenation of all the variables of λ, β, θ and A^T is selected such that $A^T z \leq 0$ is equivalent to Constraints (13), (14), and (15). Similarly Constraint (12) of τ is embedded in $b^T z > 0$. Therefore as the constraints of τ are infeasible, Case 2 of Farkas' lemma is false which implies that Case 1 is true. Because $Aw \geq b$ is the dual constraints of $A^T z \leq 0$ which is τ , $Aw \geq b$ is just constraints of $PLP(\mu, \omega)$. Thus, Case 1 of Farkas' lemma implies existences of a solution (x, y) that is feasible for $PLP(\mu, \omega)$ with objective value at most $\xi^*(1 + \epsilon)$. Because the number of constraints in τ (size of A in Case 1 of Farkas Lemma) is polynomially bounded, we can actually find a feasible solution (x, y) to $PLP(\mu, \omega)$ with objective value at most $\xi^*(1 + \epsilon)$ by the Ellipsoid algorithm. From Fact 1 we know that $\xi^* \leq \text{OPT}$, therefore the objective value of solution (x, y) to $PLP(\mu, \omega)$ is at most $\text{OPT}(1 + \epsilon)$. \square

4.4 Rounding the LP

We show how to use feasible solution (x, y) taken from Lemma 3 to obtain an integral solution to SUM-MR with the total latency at most $O(\max(\sigma, 2\phi) \cdot \omega \cdot \text{OPT})^8$.

Sum-Movement Repairmen Algorithm (SUM-MRA) shown in Figure 1 is the randomized algorithm for doing so. Later in this subsection we derandomize it in order to prove the total latency bound and hence finish the proof of Theorem 3.

As explained earlier Q in SUM-MRA is the set $\{1, 2, \dots, 2^i, \dots, 2^{\lceil \log T \rceil}\}$ where T is the latest service time a client can see which was upper bounded by $\frac{2 \cdot \text{MST}(G)}{\min_i v_i}$. Moreover, μ and ω are the constants in $PLP(\mu, \omega)$.

SUM-MRA serves clients in multiple steps. It starts serving clients with paths that have the maximum latency $1 \cdot \mu$ then it concatenate paths of maximum latency $2 \cdot \mu$, then $4 \cdot \mu$ and so on. These paths come from the set $\mathcal{P}(r, \mu \cdot v_r \cdot q)$ for $q \in Q$ and the selection is done using $PLP(\mu, \omega)$ variables $x_{r,p,q}$. In fact, for each $q \in Q$ and $f \in [4 \cdot \omega]$ we execute $\text{Append-path}(q, f)$ in order to find $4 \cdot \omega$ paths for r . This is because we want to have independence between the paths selected at each execution of $\text{Append-path}(q, f)$ which helps us to better analyze the number of clients get served in the execution.

We use the following definitions to refer to the clients served by SUM-MRA.

⁸ This suggests that the integrality gap of the LP is constant but the difficulty is in solving it.

<ol style="list-style-type: none"> 1. For each $q \in Q$ in an ascending order do: <ol style="list-style-type: none"> (a) For each $f \in [4 \cdot \omega]$ do: <ol style="list-style-type: none"> i. Perform Append-path(q, f). 2. Serve all the remaining clients by a fastest repairman using a path obtained by doubling the edges of a MST.
<p>Append-path(q, f):</p> <ol style="list-style-type: none"> 1. For each repairman r, select a path ($p_r^f \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$) independently at random where the probability of path p gets selected is equal to value $\frac{1}{\omega} x_{r,p,q}$ computed from $PLP^{(\mu, \omega)}$. Each repairman r travels its selected path and comes back to its starting depot. 2. Serve all the clients, that are yet unserved but can reach a node in a selected path by time $q \cdot \mu$.

Figure 1: Movement Repairmen Algorithm (SUM-MRA)

Definition 7. Let $A^{q,f}$ denotes the set of non-served clients getting served at Instruction 2 of Append-path(q, f). i.e., $A^{q,f}$ is the set of clients c such that c is not served before the execution of Append-path(q, f) but gets served by repairman r in a node $v \in p_r^f$ (remember p_r^f is the path selected for r in Append-path(q, f)).

Definition 8. Set $\mathcal{A}^{q,f} = \bigcup_{(q',f') \leq (q,f)} A^{q',f'}$ is the set of all clients served by SUM-MRA up to and including the execution of Append-path(q, f). Here the operator \leq is the lexicographic ordering for the ordered pairs where the first entry has more priority than the second one.

We define function $prev(q, f)$ as follows.

$$prev(q, f) = \begin{cases} (q, f - 1) & f \neq 1 \\ (\frac{q}{2}, 4 \cdot \omega) & f = 1 \end{cases}$$

Definition 9. Let $(q', f') = prev(q, f)$ and Append-path(q', f') be the predecessor of Append-path(q, f). Let $F^{q,f}$ denotes the value of $\sum_{c \in C \setminus \mathcal{A}^{q',f'}} \sum_{t \leq q} y_{c,t}$. Intuitively, $F^{q,f}$ can be taught as the fractional number of clients that are (fractionally) served in feasible solution (x, y) by the time q , but not served by SUM-MRA before the execution of Append-path(q, f).

We would like $A^{q,f}$ be a large fraction of $F^{q,f}$. First we prove the following lemma to lower bound the probability of a client getting served in the execution of Append-path(q, f).

Lemma 5. Let q be any element of Q and c be any client in C . If we randomly select a path for each repairman $r \in R$ such that the probability of selecting $p \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$ is $\frac{1}{\omega} \cdot x_{r,p,q}$, then the probability of c getting served (a selected path visits a node from $\mathcal{B}(c, \mu q)$) is at least $\frac{1}{2\omega} \cdot \sum_{q' \leq q} y_{c,q'}$.

Proof. The probability of a client c getting served by an arbitrary repairman $r \in R$ is $D_r = \frac{1}{\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v_r \cdot t) \neq \emptyset} x_{r,p,t}$ because of the probability distribution used in the rounding. To simplify the notations, let $B = \sum_{r \in R} D_r$ and $Y_c = \frac{1}{\omega} \sum_{q' \leq q} y_{c,q'}$.

The probability that a client c is not served by any repairman in R is $\prod_{r \in R} (1 - D_r)$.

$$\begin{aligned}
\prod_{r \in R} (1 - D_r) &\leq \left(\frac{|R| - \sum_{r \in R} D_r}{|R|} \right)^{|R|} && \text{Arithmetic and Geometric} \\
& && \text{Means Inequality }^9 \\
&= \left(1 - \frac{B}{|R|} \right)^{|R|} && \text{(replacing by } B) \\
&= \left(1 - \frac{1}{\frac{|R|}{B}} \right)^{\frac{|R|}{B} B} \\
&\leq e^{-B} \\
&\leq e^{-Y_c} && \text{Constraint (9)}
\end{aligned}$$

From the above inequality we conclude that client c gets served with probability at least $1 - e^{-Y_c}$. The following inequalities finish the proof of the lemma.

$$\begin{aligned}
1 - e^{-Y_c} &= 1 - \left(\sum_{i=0}^{\infty} \frac{(-Y_c)^i}{i!} \right) && \text{by Taylor Expansion} \\
&\geq Y_c - \frac{Y_c^2}{2} && \text{as } 0 \leq Y_c \leq 1 \\
&\geq \frac{1}{2} Y_c && \text{as } 0 \leq Y_c \leq 1 \\
&\geq \frac{1}{2\omega} \cdot \sum_{q' \leq q} y_{c,q'} && \text{definition of } Y_c
\end{aligned}$$

□

We use the following lemma to derandomize selections of the paths in $\text{Append-path}(q, f)$ and to show that $A^{q,f}$ is at least $\left\lceil \frac{F^{q,f}}{2 \cdot \omega} \right\rceil$.

Lemma 6. *We can derandomize $\text{Append-path}(q, f)$ to deterministically select a path $p_r^f \in \mathcal{P}(r, q \cdot v_r \cdot \mu)$ for each repairman r , such that the set of newly served clients ($A^{q,f}$ as defined in Definition 7) to be at least $\left\lceil \frac{F^{q,f}}{2 \cdot \omega} \right\rceil$.*

Proof. If we select a path $p_r^f \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$ with probability $x_{r,p_r^f,q}/\omega$ for each repairman r , the probability of an arbitrary client c getting served is at least $\frac{1}{2\omega} \cdot \sum_{t \leq q} y_{c,t}$ (by Lemma 5). By linearity of expectations, we conclude that the expected number of clients served with these paths is at least $\frac{1}{2\omega} \cdot \sum_{c \in C} \sum_{t \leq q} y_{c,t}$. Therefore the expected number of new clients that are served with these paths is at least $\frac{F^{q,f}}{2\omega}$ by definition of $F^{q,f}$ (see Definition 9).

⁹ For any set of n non-negative numbers x_1, \dots, x_n we have $\frac{x_1 + \dots + x_n}{n} \geq \sqrt[n]{x_1 \cdot x_2 \cdot \dots \cdot x_n}$

Now, we derandomize the random selection of the paths in $\text{Append-path}(q, f)$ so that we guarantee serving at least $\left\lceil \frac{F^{q,f}}{2\omega} \right\rceil$ number of new clients. This guarantee is important since we run this procedure in multiple rounds (variable f denotes the round number) and in each round we need to serve at least this fraction of clients. Let $R' \subseteq R$ be an arbitrary subset of the set of the repairmen. For each client $c \in C$ and time $q \in Q$ we define variable $Y_{c,q}^{R'}$ as follows.

$$Y_{c,q}^{R'} = \frac{1}{\omega} \sum_{r \in R'} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t) : p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} \quad (16)$$

Intuitively, $Y_{c,q}^{R'}$ represents the fractional service client c receives from the Repairmen in R' in solution (x, y) , if each repairman could take at most one path. Here, we use $\frac{1}{\omega}$ because in (x, y) each repairman can take up to ω paths (see Constraint (1) in **PLP**). Consequently, for an arbitrary subset $C' \subseteq C$, $R' \subseteq R$ and, time $q \in Q$, we define the variable $\mathcal{Y}_q^{C', R'}$ as follows.

$$\mathcal{Y}_q^{C', R'} = \sum_{c \in C'} Y_{c,q}^{R'} \quad (17)$$

Intuitively, we can think of $\mathcal{Y}_q^{C', R'}$ to be the (fractional) amount of service that clients in C' receive from the repairmen in R' by time q in feasible solution (x, y) . We prove the following claim.

Claim 2. *Let C' be an arbitrary subset of C and R' be an arbitrary subset of R . We can deterministically select one path for each repairman in R' such that they serve at least $\left\lceil \frac{1}{2} \mathcal{Y}_q^{C', R'} \right\rceil$ clients from C' .*

Proof. We prove this claim by induction on $|R'|$. For the base case when $|R'| = 1$ assume that repairman r is the only member of R' . For each path $p \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$, let C'_p be the set of clients c in C' whose neighborhood ball $(\mathcal{B}(c, v'_c \cdot \mu \cdot t))$ gets hit by p . Let path $p_r^* \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$ be the path that intersects with the maximum number of neighborhood balls of the clients in C' (i.e., $|C'_{p_r^*}|$ is the maximum). We select path p^* for r . Consider the following inequalities.

$$\begin{aligned} \mathcal{Y}_q^{C', \{r\}} &= \sum_{c \in C'} Y_{c,q}^{\{r\}} \\ &= \sum_{c \in C'} \frac{1}{\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t) : p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} && \text{definition of } Y_{c,q}^{C', \{r\}} \\ &\leq \frac{1}{\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} \sum_{c \in C' : p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} \\ &\leq \frac{1}{\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} |C'_p| x_{r,p,t} && \text{definition of } C'_p \\ &\leq \frac{1}{\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} |C'_{p^*}| x_{r,p,t} \\ &\leq \frac{1}{\omega} \omega \cdot C'_{p^*} && \text{Constraint (8)} \\ &\leq C'_{p^*} \end{aligned}$$

The above inequality shows $\mathcal{Y}_q^{C',\{r\}} \leq C'_{p^*}$ and hence $\left\lceil \frac{1}{2}\mathcal{Y}_q^{C',\{r\}} \right\rceil \leq C'_{p^*}$ as C'_{p^*} is an integer value which completes the proof for the base case.

Assume that the claim holds for any subset of repairmen with size k' as the induction hypothesis. We prove that the claim holds for an arbitrary subset $R' \subseteq R$ of size $k'+1$. Let $r \in R'$ be a repairman in R' . Similar to the base case let path $p^* \in \mathcal{P}(r, \mu \cdot v_r \cdot q)$ be the path that intersects with the maximum number of \mathcal{B} -balls of the clients in C' (i.e., $|C'_{p^*}|$ is the maximum). We select path p^* for r which serves $|C'_{p^*}|$ new clients from C' . By the induction hypothesis we can select one path for each of the remaining repairmen $R' \setminus \{r\}$ such that they serve at least $\left\lceil \frac{1}{2}\mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right\rceil$ many clients. In the following we prove that $\left\lceil \frac{1}{2}\mathcal{Y}_q^{C', R'} \right\rceil \leq C'_{p^*} + \left\lceil \frac{1}{2}\mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right\rceil$ which completes the proof of the induction for R' and hence the claim.

$$\begin{aligned}
& \left\lceil \frac{1}{2}\mathcal{Y}_q^{C', R'} \right\rceil \\
&= \left\lceil \frac{1}{2} \sum_{c \in C'} Y_{c,q}^{R'} \right\rceil \\
&= \left\lceil \frac{1}{2} \left(\sum_{c \in C' \setminus C'_{p^*}} Y_{c,q}^{R'} + \sum_{c \in C'_{p^*}} Y_{c,q}^{R'} \right) \right\rceil \\
&= \left\lceil \frac{1}{2} \sum_{c \in C' \setminus C'_{p^*}} \frac{1}{\omega} \sum_{r' \in R'} \sum_{p \in \mathcal{P}(r', \mu \cdot v_{r'} \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r',p,t} \right. \\
&\quad \left. + \frac{1}{2} \sum_{c \in C'_{p^*}} Y_{c,q}^{R'} \right\rceil \tag{def. of } Y_{c,q}^{R'} \\
&= \left\lceil \frac{1}{2} \sum_{c \in C' \setminus C'_{p^*}} \frac{1}{\omega} \sum_{r' \in R' \setminus \{r\}} \sum_{p \in \mathcal{P}(r', \mu \cdot v_{r'} \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r',p,t} \right. \\
&\quad \left. + \frac{1}{2\omega} \sum_{c \in C' \setminus C'_{p^*}} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} + \frac{1}{2} \sum_{c \in C'_{p^*}} Y_{c,q}^{R'} \right\rceil \tag{excluding } r \text{ from } R' \\
&= \left\lceil \frac{1}{2}\mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right. \\
&\quad \left. + \frac{1}{2\omega} \sum_{c \in C' \setminus C'_{p^*}} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} + \frac{1}{2} \sum_{c \in C'_{p^*}} Y_{c,q}^{R'} \right\rceil \tag{def. of } \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}}
\end{aligned}$$

$$\begin{aligned}
&\leq \left[\frac{1}{2} \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right. \\
&+ \left. \frac{1}{2\omega} \sum_{c \in C' \setminus C'_{p^*}} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} + \frac{1}{2} |C'_{p^*}| \right] \quad \text{Noting that } Y_{c,q}^{R'} \leq 1 \\
&= \left[\frac{1}{2} \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right. \\
&+ \left. \frac{1}{2\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} \sum_{c \in C' \setminus C'_{p^*}: p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} + \frac{1}{2} |C'_{p^*}| \right] \quad \text{reordering the} \\
&\hspace{15em} \text{sum functions} \\
&\leq \left[\frac{1}{2} \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} + \frac{1}{2\omega} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t)} |C'_{p^*}| x_{r,p,t} + \frac{1}{2} |C'_{p^*}| \right] \quad p^* \text{ serves the maximum number of} \\
&\hspace{15em} \text{clients from } C' \text{ hence } |C'_{p^*}| \text{ is larger} \\
&\hspace{15em} \text{than what } p^* \text{ covers in } C' \setminus C'_{p^*} \\
&\leq \left[\frac{1}{2} \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} + \frac{1}{2\omega} \omega |C'_{p^*}| + \frac{1}{2} |C'_{p^*}| \right] \quad \text{Constraint (8)} \\
&\leq \left[\frac{1}{2} \mathcal{Y}_q^{C' \setminus C'_{p^*}, R' - \{r\}} \right] + |C'_{p^*}|
\end{aligned}$$

□

Let $\mathcal{A}^{q', f'}$ be the set of served clients before execution of $\text{Append-path}(q, f)$ (remember Definition 8). We prove the following inequality about $\mathcal{Y}_q^{C', R'}$ where $R' = R$ and $C' = C \setminus \mathcal{A}^{q', f'}$.

$$\begin{aligned}
&\mathcal{Y}_q^{C \setminus \mathcal{A}^{q', f'}, R} \\
&= \sum_{c \in C \setminus \mathcal{A}^{q', f'}} \frac{1}{\omega} \sum_{r \in R} \sum_{p \in \mathcal{P}(r, \mu \cdot v_r \cdot t): p \cap \mathcal{B}(c, \mu \cdot v'_c \cdot t) \neq \emptyset} x_{r,p,t} \quad \text{Inequality (17)} \\
&\geq \sum_{c \in C \setminus \mathcal{A}^{q', f'}} \frac{1}{\omega} \sum_{q' \leq q} y_{c,q'} \quad \text{Constraint (9) of } PLP^{(\mu, \omega)} \\
&= \frac{F^{q, f}}{\omega} \quad \text{Definition 9}
\end{aligned}$$

The above inequality proves that $\left[\frac{1}{2} \mathcal{Y}_q^{C \setminus \mathcal{A}^{q', f'}, R} \right] \geq \left[\frac{F^{q, f}}{2\omega} \right]$. Therefore by using Claim 2 where $R' = R$ and $C' = C \setminus \mathcal{A}^{q', f'}$, we serve at least $\frac{F^{q, f}}{2\omega}$ clients which finishes the proof of the lemma. □

We prove the following lemma which combined with Lemma 3 finishes the proof of Theorem 3.

Lemma 7. A feasible solution (x, y) to $PLP^{(\mu, \omega)}$ with objective value OPT (the optimum value for PLP) can be rounded to an integral solution to SUM-MR with total latency $O(\mu \cdot \omega) \cdot \text{OPT}$.

Proof. For the proof, we show SUM-MRA serves all the clients with total latency $O(\mu \cdot \omega) \cdot \text{OPT}$. The following definition is our last definition in this section.

Definition 10. For any $q \in Q$, let h_q be $\sum_{c \in C} y_{c,q}$. As $y_{c,q}$ denotes how much client c is served in time q , h_q can be thought as the total amount of fractional service clients receive at time q in feasible solution (x, y) .

Note that after executing $\text{Append-path}(q, 4 \cdot \omega)$ we jump to the next element in Q , i.e., we execute $\text{Append-path}(2 \cdot q, 1)$ and append paths of length $2q \cdot \mu$. Although our algorithm does not execute $\text{Append-path}(q, 4 \cdot \omega + 1)$, we use value $F^{q, 4 \cdot \omega + 1}$ (see Definition 9) to denote the fractional number of clients that are served in feasible solution (x, y) by time q but are not served by SUM-MRA with paths of length at most $q \cdot \mu$ (executions of $\text{Append-path}(q, f)$ for all f in $[4 \cdot \omega]$). We use the following claim to upper bound $F^{q, 4 \cdot \omega + 1}$.

Claim 3. For any $q = 2^a \in Q$ we have $F^{q, 4 \cdot \omega + 1} \leq \sum_{s=0}^a \left(\frac{1}{4^{a-s+1}} \right) h_{(2^s)}$.

Proof. First we prove that for any $q \in Q$ we have $F^{q, 4 \cdot \omega + 1} \leq \frac{F^{q,1}}{4}$. By Lemma 6 we conclude that performing $\text{Append-path}(q, f)$ serves at least $\left\lceil \frac{F^{q,f}}{2 \cdot \omega} \right\rceil$ new clients for any $f \in [4 \cdot \omega]$. In other words, $F^{q, f+1}$ is at most $F^{q,f} \left(1 - \frac{1}{2 \cdot \omega}\right)$ which implies that after the execution of $\text{Append-path}(q, f)$ we drop the number of non-served clients by at least factor of $\left(1 - \frac{1}{2 \cdot \omega}\right)$. By iterating f over set $[4 \cdot \omega]$ we conclude that after executing $\text{Append-path}(q, 4 \cdot \omega)$ the total number of non-served clients is at most $F^{q,1} \cdot \left(1 - \frac{1}{2 \cdot \omega}\right)^{4 \cdot \omega}$ which shows that $F^{q, 4 \cdot \omega + 1} \leq \left(F^{q,1} \cdot \left(1 - \frac{1}{2 \cdot \omega}\right)^{4 \cdot \omega}\right)$.

$$\begin{aligned}
F^{q, 4 \cdot \omega + 1} &\leq F^{q,1} \cdot \left(1 - \frac{1}{2 \cdot \omega}\right)^{4 \cdot \omega} && \text{since } \left(1 - \frac{1}{x}\right)^x \leq \frac{1}{e} \\
&\leq F^{q,1} \cdot \left(\frac{1}{e}\right)^2 \\
&\leq \frac{1}{4} \cdot F^{q,1}
\end{aligned} \tag{18}$$

We prove the claim by induction on a . The base case when $q' = 2^0$ is the direct result of Inequality (18) by noting that $F^{1,1} = h_1$ (from the definition of $F^{1,1}$). Assume that for $q' = 2^{a'}$ the claim holds as the induction hypothesis. The following inequalities prove the claim's statement for $2q' = 2^{a'+1}$ which finish the proof of the induction and hence the claim.

$$\begin{aligned}
F^{2q', 4 \cdot \omega + 1} &\leq \frac{1}{4} F^{2q', 1} && \text{Due to Inequality (18) with } q = 2q' \\
&\leq \frac{1}{4} \left(F^{q', 4 \cdot \omega + 1} + h_{2q'} \right) && \text{Definition 9 and Definition 10} \\
&\leq \sum_{0 \leq s \leq a'+1} \left(\frac{1}{4^{(a'+1)-s+1}} \right) h_{(2^s)} && \text{Induction hypothesis}
\end{aligned}$$

□

Now we bound the total latency of SUM-MRA. We start with calculating the latency of the clients that are served at Step 2 of SUM-MRA. These are the remaining clients after executing all the Append-path procedures. Remember that $T = \frac{2 \cdot MST(G)}{\min_i v_i}$ is the time required to visit all the nodes of G by the slowest repairman, hence, is the latest time a client can be served. This means that h_t is equal to zero for all $t > T$. The greatest element of Q is $2^{\lceil \log T \rceil}$. After executing $\text{Append-path}(2^{\lceil \log T \rceil}, 4 \cdot \omega)$ we have at most the following number of clients to serve by Claim 3.

$$\sum_{0 \leq s \leq \lceil \log T \rceil} \left(\frac{1}{4^{\lceil \log T \rceil - s + 1}} \right) h_{(2^s)}$$

We serve each of these remaining clients by latency at most T at Step 2 of SUM-MRA. Also note from definition of h_q (Definition 8), we can rewrite the objective value of (x, y) in $PLP^{\mu, \omega}$ as $\sum_{0 \leq s \leq \lceil \log T \rceil} 2^s h_{(2^s)}$. We show that the total latency of Step 2 of SUM-MRA is at most OPT .

$$\begin{aligned} \sum_{0 \leq s \leq \lceil \log T \rceil} \left(\frac{T}{4^{\lceil \log T \rceil - s + 1}} \right) h_{(2^s)} &= \sum_{0 \leq s \leq \lceil \log T \rceil} \left(\frac{T}{4^{\lceil \log T \rceil + 1}} \right) 2^s h_{(2^s)} \\ &\leq \sum_{0 \leq s \leq \lceil \log T \rceil} 2^s h_{(2^s)} \\ &< \text{OPT} \end{aligned}$$

Now we bound the sum of latencies produced by SUM-MRA on the clients at Step 1a of SUM-MRA. At each $\text{Append-path}(q, f)$, we add a path for each repairman r_i who can travel its path and come back by time $2 \cdot q \cdot \mu$. Therefore repairman r_i travels at most $8 \cdot \omega \cdot q \cdot \mu$ units of time for the paths added at $\text{Append-path}(q, f)$ for $1 \leq f \leq 4 \cdot \omega$. Repairman r_i before starting to travel the path added at $\text{Append-path}(q, 1)$ has to travel all the paths added before. Traveling previous paths takes the total of $\sum_{i=0}^{\log q - 1} 8 \cdot \omega \cdot 2^i \cdot \mu \leq 8 \cdot \omega \cdot \mu \cdot q$. Therefore we can assume that all the clients that are served at $\text{Append-path}(q, f)$ for $1 \leq f \leq 4 \cdot \omega$ have latency at most $16 \cdot \mu \cdot \omega \cdot q$ (Fact 1). This is because $8 \cdot \omega \cdot q \cdot \mu$ for the paths added at $\text{Append-path}(q, \cdot)$ and $8 \cdot \omega \cdot q \cdot \mu$ for the paths added before $\text{Append-path}(q, 1)$.

From Claim 3 we know for $q = 2^a$ we have $F^{q, 4 \cdot \omega + 1} \leq \sum_{0 \leq s \leq a} \left(\frac{1}{4^{a-s+1}} \right) h_{(2^s)}$. Remember that set $\mathcal{A}^{q, 4 \cdot \omega}$ is the set of all the clients that SUM-MRA has served in or before executing $\text{Append-path}(q, 4 \cdot \omega)$. From the definition of $\mathcal{A}^{q, 4 \cdot \omega}$ we conclude the following inequality.

$$\begin{aligned} |\mathcal{A}^{q, 4 \cdot \omega}| &\geq \sum_{0 \leq s \leq a} h_{(2^s)} - \sum_{0 \leq s \leq a} \left(\frac{1}{4^{a-s+1}} \right) h_{(2^s)} && \text{Definition 8} \\ &= \sum_{0 \leq s \leq a} \left(\frac{4^{a-s+1} - 1}{4^{a-s+1}} \right) h_{(2^s)} && (19) \end{aligned}$$

Note that the size of $\mathcal{A}^{q, 4 \cdot \omega}$ is integral. In fact, Inequality (19) shows SUM-MRA serves at least $\left\lceil \sum_{0 \leq s \leq a} \left(\frac{4^{a-s+1} - 1}{4^{a-s+1}} \right) h_{(2^s)} \right\rceil$ clients after it executes $\text{Append-path}(q, 4 \cdot \omega)$. Note that from Fact 1 each client that is served in or before execution of $\text{Append-path}(q, 4 \cdot \omega)$ sees a latency at most $16 \cdot \mu \cdot \omega \cdot q$ and each client that is served in the succeeding executions of Append-path sees a higher latency. For the sake of explanation and to avoid dealing with the ceiling function,

we slightly abuse the notation $|\mathcal{A}^{q,4\cdot\omega}|$ to denote possibly a fractional value. More formally, we assume that SUM-MRA serves exactly $\sum_{0 \leq s \leq a} \left(\frac{4^{a-s+1}-1}{4^{a-s+1}}\right) h_{(2^s)}$ (possibly fractional) clients after it executes $\text{Append-path}(q, 4 \cdot \omega)$ and assume that it serves the rest (if there are more) in the succeeding executions of Append-path . Note that this way we do not decrease the total latency. We use the following claim to upper bound the total latency of SUM-MRA.

Claim 4. *After SUM-MRA executes $\text{Append-path}(q, 4 \cdot \omega)$, the total latency of $\sum_{0 \leq s \leq a} \left(\frac{4^{a-s+1}-1}{4^{a-s+1}}\right) h_{(2^s)}$ clients that are served is at most $\sum_{0 \leq s \leq a} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{a-s+1}-1}{4^{a-s+1}}\right) h_{(2^s)}$.*

Proof. Let $q = 2^a$, we prove the claim by induction on a . For the base case when $a = 0$, Inequality (19) implies that SUM-MRA serves at least $\frac{3}{4}h_1$ clients after it executes $\text{Append-path}(1, 4 \cdot \omega)$. From Fact 1 the total latency for these clients is $16 \cdot 2^0 \cdot \omega \cdot \mu \cdot \frac{3}{4}h_0$.

For $a = k'$ when SUM-MRA executes $\text{Append-path}(2^{k'}, 4 \cdot \omega)$, we assume that it serves $\sum_{0 \leq s \leq k'} \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)}$ clients with total latency $\sum_{0 \leq s \leq k'} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)}$ as the induction hypothesis.

For $a = k' + 1$, Inequality (19) shows SUM-MRA serves at least $\sum_{0 \leq s \leq k'+1} \left(\frac{4^{k'+1-s+1}-1}{4^{k'+1-s+1}}\right) h_{(2^s)}$ clients after it executes $\text{Append-path}(2^{k'+1}, 4 \cdot \omega)$. From the induction hypothesis we serve $\sum_{0 \leq s \leq k'} \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)}$ clients with latency $\sum_{0 \leq s \leq k'} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)}$. Each of the remaining $\sum_{0 \leq s \leq k'+1} \left(\frac{4^{k'+1-s+1}-1}{4^{k'+1-s+1}}\right) h_{(2^s)} - \sum_{0 \leq s \leq k'} \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)}$ clients is served with latency $16 \cdot \mu \cdot \omega \cdot 2^{k'+1}$ (from Fact 1). We bound the total latency in the following inequalities which finishes the proof of the induction and hence the claim.

$$\begin{aligned}
& \sum_{0 \leq s \leq k'} 32 \cdot \mu \cdot \omega \cdot 2^s \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)} + \\
& \left(\sum_{0 \leq s \leq k'+1} \left(\frac{4^{k'+1-s+1}-1}{4^{k'+1-s+1}}\right) h_{(2^s)} - \sum_{0 \leq s \leq k'} \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)} \right) \cdot 16 \cdot \mu \cdot \omega \cdot 2^{k'+1} \\
& = \sum_{0 \leq s \leq k'} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{k'-s+1}-1}{4^{k'-s+1}}\right) h_{(2^s)} + \\
& \left(\sum_{0 \leq s \leq k'} \left(\frac{3}{4^{k'+1-s+1}}\right) h_{(2^s)} \right) \cdot 16 \cdot \mu \cdot \omega \cdot 2^{k'+1} + \left(\frac{3}{4}\right) h_{(2^{k'+1})} \cdot 16 \cdot \mu \cdot \omega \cdot 2^{k'+1} \\
& = \sum_{0 \leq s \leq k'} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{k'+1-s+1}-1}{4^{k'+1-s+1}}\right) h_{(2^s)} + \left(\frac{3}{4}\right) h_{(2^{k'+1})} \cdot 16 \cdot \mu \cdot \omega \cdot 2^{k'+1} \\
& \leq \sum_{0 \leq s \leq k'+1} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot \left(\frac{4^{k'+1-s+1}-1}{4^{k'+1-s+1}}\right) h_{(2^s)}
\end{aligned}$$

□

Note that from Definition 10 for h_q , we can rewrite the objective value of (x, y) in $PLP^{(\mu, \omega)}$ as $\sum_{0 \leq s \leq \lceil \log T \rceil} 2^s h_{(2^s)}$. The total latency of SUM-MRA is at most $\sum_{0 \leq s \leq \lceil \log T \rceil} 32 \cdot \mu \cdot \omega \cdot 2^s \cdot h_{(2^s)}$ from

Claim 4. Therefore SUM-MRA has the total latency at most $32 \cdot \mu \cdot \omega \cdot \text{OPT} + \text{OPT}$ where the last OPT is from Step 2. \square

5 Neighborhood Prize Collecting Steiner Tree

In this section we present the results for the NPCST problem.

5.1 Proof of Theorem 2

Proof. As described in the introduction section, serving neighborhoods instead of a single node in the covering problems like neighborhood TSP makes the problem significantly harder. We show an $\Omega(\log^{2-\epsilon} k)$ hardness for Neighborhood Steiner Tree (NST) where k is the number of clients.

Definition 11. An instance $\mathcal{I} = (\mathcal{M}, C)$ of NST consists of a metric space $\mathcal{M} = (V, d)$ where V is the node set and $d : (V \times V) \rightarrow \mathbb{Q}^+$ is the distance function, and a set of clients C where each client c has a neighborhood ball $\mathcal{B}(c, t_c)$ which is the set of all nodes u such that $d(u, c) \leq t_c$. The objective for NST is to find a tree T with minimum cost such that for each client $c \in C$, T serves at least one node of $\mathcal{B}(c, t_c)$.

Note that as opposed to NPCST, in NST we do not have bound on the resulting tree and profits for the clients but we have to serve all the clients.

We reduce from the Group Steiner Tree (GST) problem. An instance $\mathcal{I}' = (\mathcal{M}', \mathcal{G}')$ of GST consists of a metric space $\mathcal{M}' = (V', d')$ where V' is the node set and $d' : (V' \times V') \rightarrow \mathbb{Q}^+$ is the distance function and a collection of groups $\mathcal{G}' = \{G'_1, G'_2, \dots, G'_k\}$ where each group G'_i is a subset of nodes in V' . The objective for GST is to find a tree T' with minimum cost such that T' contains at least one node from each group. We use the following result of Halperin and Krauthgamer [HK03] about the group Steiner tree problem.

Theorem 7. [HK03]: For every fixed $\epsilon > 0$, group Steiner tree cannot be approximated within ratio $\log^{2-\epsilon} k$, unless $NP \subseteq ZTIME(n^{\text{poly} \log(n)})$; this holds even for trees.

We show a polynomial-time reduction from any instance \mathcal{I}' of GST to an instance \mathcal{I} of NST where any α -approximation for \mathcal{I} results in an α -approximation algorithm for \mathcal{I}' . We build \mathcal{I} from \mathcal{I}' as follows. Let us assume the distance between the farthest pair of points in V' is M' . The node set in \mathcal{I} is V' plus k dummy nodes y_1, y_2, \dots, y_k where each y_i corresponds to G'_i in \mathcal{G}' , i.e., $V = V' \cup \{y_1, y_2, \dots, y_k\}$. The distances between nodes of V that correspond to the nodes in V' is the same as d' . Each dummy node y_i is connected to all the nodes u whose corresponding node in V' is in G'_i with an edge of length M' . The set of clients C in \mathcal{I} is $\{y_1, \dots, y_k\}$ where neighborhood ball of each y_i has radius M' ($\mathcal{B}(y_i, M')$).

Note that this reduction can be done in polynomial time. Any solution T' for GST in \mathcal{I}' can be transformed to a solution T for NST in \mathcal{I} with the same cost. Here T contains all corresponding nodes and edges of T' in \mathcal{I} . As T' serves at least one client from each group in \mathcal{I}' , T serves one node from each neighborhood in \mathcal{I} . Now we prove that any solution T for NST in \mathcal{I} can be transformed to a solution T' for GST in \mathcal{I}' with at most the same cost. T' contains each node v'_i whose corresponding node v_i is in T and is not a dummy node. Firstly, if a dummy node y_i is a leaf in T we simply ignore y_i in T' without affecting servicing the groups since y_i is not a node of any group in \mathcal{G}' . Secondly, if

T uses a dummy node to go from node v_i to a node v_j which costs $2M$, we instead use the shortest path between v'_i and v'_j (the corresponding nodes to v_i, v_j) in V' which costs at most M . Thus, T can be transformed to a tree T' in \mathcal{I}' using only nodes in V' that serves all the groups, with at most the same cost as T .

Because each solution for \mathcal{I} can be transformed to a solution to \mathcal{I}' with at most the same cost and vice versa, we conclude that the optimum solution for both \mathcal{I}' and \mathcal{I} have the same cost. Moreover, we can get the output of an α -approximation algorithm for \mathcal{I} and transform it to a solution to the GST with at most the same cost which is an α -approximation algorithm for the GST problem. Thus, from Theorem 7 we conclude that there is no $O(\log^{2-\epsilon} k)$ -approximation algorithm for NST, unless $NP \subseteq ZTIME(n^{\text{poly} \log(n)})$ for every fixed $\epsilon > 0$. \square

5.2 Proof of Theorem 4

We use the result of Fakcharoenphol *et al.* [FRT04] and Charikar *et al.* [CCG⁺98] to embed metric space $\mathcal{M} = (V, d)$ into a distribution of dominating trees with a distortion at most $O(\log n)$ to be explained in the following¹⁰. Here, a tree T is dominating metric d if for any two nodes $u, v \in V$ we have $d_T(u, v) \geq d(u, v)$ where $d_T(u, v)$ is the length of the unique path between u and v in T . Fakcharoenphol *et al.* show that their embedding can be done into at most $O(n \log n)$ dominating trees where n is the number of the nodes using the technique first introduced by Charikar *et al.* [CCG⁺98]. More formally, they proved that the metric d can be embedded into a distribution π of p trees T_1, T_2, \dots, T_p where $p \in O(n \log n)$ such that the following holds.

$$\forall u, v \in V, \quad \sum_{i=1}^p \pi(T_i) d_{T_i}(u, v) < O(\log n) \cdot d(u, v)$$

In order to be more specific we assume that $O(\log n)$ in the above inequality is $A \log n$ for an appropriate constant value A , more formally we use the following inequality.

$$\forall u, v \in V, \quad \sum_{i=1}^p \pi(T_i) d_{T_i}(u, v) < A \cdot \log n \cdot d(u, v) \quad (20)$$

Note that the average distortion in the distances of all the pairs are bounded, however we cannot ensure that there is a single tree in which all the distances are distorted with at most a factor of $A \cdot \log n$. Moreover, the neighborhood balls of the clients are also not preserved in the trees. Therefore, instead of hitting the neighborhoods of the clients we try to minimize the total service cost of all the clients where the service cost of a client (c) in a tree (H) is its profit θ_c times its distance to H divided by t_c (the radius of its neighborhood ball). More formally, we define the following alternative problem to solve for each tree T_i which helps us solve the NPCST problem in the original metric d .

Definition 12. Instance $\mathcal{I} = (T = (V, E), r, C, B, B')$ of the Total Service Cost Steiner Tree (TSCST) problem consist of a tree T on the node set V and edges E rooted at r where each client

¹⁰ The Metric embedding problems are studied well in the theory community. Bartal [Bar96] first defined probabilistic embeddings and gives the distortion ratio $O(\log^2 n)$. Bartal [Bar98] improved this ratio to $O(\log n \log \log n)$ by using the ideas inspired from Seymours work on feedback arc set [Sey95]. The $O(\log n)$ distortion ratio of Fakcharoenphol *et al.* [FRT04] is the best possible distortion ratio one can hope for.

$c \in C$ has profit θ_c and is assigned to a neighborhood ball with radius t_c , length bound $B \in \mathbb{Z}^+$, and service cost $B' \in \mathbb{Q}^+$. The objective is to find a tree H (subtree of T) to serve a subset $C' \subseteq C$ such that $\sum_{c \in C'} \theta_c$ is maximized. The constraints on H are as follow: (1) the total length of H has to be at most B and (2) the service cost of C' which is defined as $\sum_{c \in C'} \theta_c \cdot \frac{d_T(c, H)}{t_c}$ has to be at most B' where $d_T(c, H)$ is the distance of the location of c to its nearest node in H in metric d_T .

Intuitively, in TSCST as opposed to NPCST, we try to minimize the overall violations on the radius of the neighborhoods instead of having hard capacity on the neighborhoods' radii.

Lemma 8. *For an arbitrary small positive value $\epsilon > 0$, we can design an algorithm whose running time is polynomial in n and $\frac{1}{\epsilon}$ which solves the TSCST problem efficiently while violates service cost bound B' by a factor at most $(1 + \epsilon)$.*

Note that the profits in the NPCST problem (θ) come from the dual variables of our LP. We are interested in the service costs that are integers and polynomially bounded by the size of input in order to design an efficient algorithm to solve TSCST. Therefore, we define value X to be $\frac{B' \cdot \epsilon}{|C|}$ and scale and round the service cost for each client c to value $\lfloor \frac{\theta_c \cdot d_T(c, H)}{t_c \cdot X} \rfloor$. More formally, we define the following scaled version of the TSCST problem in which all the service costs are positive integers as opposed to rational numbers.

Definition 13. *Instance $\hat{\mathcal{I}} = (T = (V, E), r, C, B, \hat{B}')$ of the Scaled Total Service Cost Steiner Tree (STSCST) problem consists of tree T rooted at r on the node set V and edges E , set of clients C where each client $c \in C$ has profit $\theta_c \in \mathbb{Q}^+$, length bound $B \in \mathbb{Z}^+$, and service cost bound $\hat{B}' \in \mathbb{Z}^+$. The objective is to find a tree (subtree of T) (H) rooted at r to serve a subset $C' \subseteq C$ such that $\sum_{c \in C'} \theta_c$ is maximized. The constraints on H are as follow: (1) the total length of H has to be at most B and (2) the service cost of C' which is defined as $\sum_{c \in C'} \lfloor \frac{\theta_c \cdot d_T(c, H)}{t_c \cdot X} \rfloor$ has to be at most \hat{B}' where $d_T(c, H)$ is the distance of location of c to its nearest node in H in metric d_T and $X = \frac{B' \cdot \epsilon}{|C|}$ for a fixed $\epsilon > 0$.*

In order to prove Lemma 8, for a given instance $\mathcal{I} = (T = (V, E), r, C, B, B')$ of TSCST we create a corresponding instance ($\hat{\mathcal{I}} = (T = (V, E), r, C, B, \hat{B}')$) of STSCST where $\hat{B}' = \lfloor \frac{B'}{X} \rfloor$. Clearly every solution to \mathcal{I} is a solution to $\hat{\mathcal{I}}$ as we round down all the service costs. Therefore, the optimum value of $\hat{\mathcal{I}}$ is at least the optimum value \mathcal{I} . In the following lemma we show that STSCST can be solved efficiently. We find an optimal solution to $\hat{\mathcal{I}}$ and announce it as a solution to TSCST, however the solution might violate the service cost bound B' . The violation is at most $X \cdot |C|$ since we have at most $|C|$ clients and for each client the scaled value of service cost can be X . As $X \cdot |C|$ is equal to $B' \cdot \epsilon$, the solution we find in $\hat{\mathcal{I}}$ collects the maximum profit and violates the service cost by at most a factor of ϵ in \mathcal{I} and hence the proof of Lemma 8 follows.

Lemma 9. *There is an algorithm with running time $O(n \cdot B^2 \cdot \hat{B}'^2 + n \cdot |C| \cdot B \cdot \hat{B}'^3)$ that solves the STSCST problem efficiently.*

Proof. Note that since we assume in the input of NPCST all the distances are polynomially bounded, value B is also polynomially bounded. Value \hat{B}' is at most $|C| \frac{1}{\epsilon}$. Therefore, the running time of our algorithm is polynomial in terms of the inputs.

We replace every node with at least 3 children by a complete binary tree of edges of cost zero (except for the last edge connecting a leaf that has the weight) thus imposing at most 2 children for

any node. Therefore, we assume that every node in T has at most two children. The *level number* of a node v in T is the number of edges in the unique path from v to r . We denote by $sub(v)$ the subtree containing node v as its root and all of its children.

We design a Dynamic Programming (DP) to solve the problem. The subproblems to our DP algorithm are defined for the subtrees rooted at each node and all possible combination of length bounds and service costs.

Definition 14. For each node v , value B_v ($0 \leq B_v \leq B$), and value \hat{B}'_v ($0 \leq \hat{B}'_v \leq \hat{B}'$), sub-problem $(DP(v, B_v, \hat{B}'_v))$ to our dynamic programming is instance $(sub(v), r, C_v, B_v, \hat{B}'_v)$ of STSCST where C_v is the set of all clients whose starting locations are in $sub(v)$.

We call that sub-problem $DP(u, B_u, \hat{B}'_u)$ is smaller than $DP(v, B_v, \hat{B}'_v)$ if u is a children of v . The following claim shows that each sub-problem $DP(v, B_v, \hat{B}'_v)$ can be solved efficiently using only smaller sub-problems.

Claim 5. Each subproblem $DP(v, B_v, \hat{B}'_v)$ can be solved in polynomial time given we can find an optimum solution for all of its sub-problems.

Proof. In order to solve $DP(v, B_v, \hat{B}'_v)$ we need to use an algorithm which solves the Knapsack problem. In the Knapsack problem we are given a set of items i_1, i_2, \dots, i_q where each item i_j has weight $w(i_j) \in \mathbb{Z}^+$ and value $t(i_j) \in \mathbb{Q}^+$ and a knapsack which can hold items with total weight at most $W \in \mathbb{Z}^+$. The objective is to pack a subset of the items into the knapsack so that their total value is maximized and their total weight is at most W . The following is a well-known result about the Knapsack problem.

Theorem 8. ([IK75]) There is an algorithm for the Knapsack problem which finds the packing of items with maximum total value such that their total weight is at most W . The algorithm running time is $O(W \cdot q)$.

Note that the objective for sub-problem $DP(v, B_v, \hat{B}'_v)$ is to find a tree (H_v) , a subtree of $sub(v)$, and a subset of clients (C'_v) such that $\sum_{c \in C'_v} \theta_c$ is maximized with constraints that the weight of H_v is at most B_v and $\sum_{c \in C'_v} \theta_c \cdot \frac{d_T(c, H_v)}{t_c} \leq \hat{B}'_v$. Consider an optimum tree H_v^* , there are only four different cases regarding if H_v^* contains children of v or not. More formally, H_v^* lies in one of the following cases: (Case 1) if H_v^* contains none of the children of v , (Case 2) if v has only one children and H_v^* contains it, (Case 3) if v has two children and H_v^* contains only one of them, and finally (Case 4) if v has two children and H_v^* contains both of them. In the following we show how to find an optimum tree for each case. We select the one which collects the most profit as the solution of our algorithm. Because H_v^* lies in one of the cases our solution is an optimum tree and the claim follows.

Case 1: If H_v^* does not contain any children of v then resulting tree H_v also contain none of the children of v and hence H_v contains only v . Therefore for any client $c \in sub(v)$ we have $d_T(c, H_v) = d_T(c, v)$ and hence the service cost of c is $\theta_c \cdot \frac{d_T(c, v)}{t_c}$. Thus, the problem is only finding a subset of clients with maximum profit such that their total service cost is at most \hat{B}'_v . In order to solve this problem we define the corresponding instance of the knapsack problem as follows. For each client $c \in C_v$ we create an item c with value $t(c) = \theta_c$ and weight $w(c) = \theta_c \cdot d_T(c, v)$ in the corresponding knapsack instance. The weight bound W is equal to \hat{B}'_v . We use Theorem 8 to find a

subset of items (clients) with maximum total profit where their total weight (service cost) is at most \hat{B}'_v . The running time of our algorithm is $O(\hat{B}'_v \cdot |C|)$ as there are at most $|C|$ nodes in $sub(v)$.

Case 2: Let v_1 be the only children of v . Because H_v^* contains v_1 , the resulting tree H_v must contain v_1 and the length bound B_v has to be at least $d_T(v, v_1)$, otherwise in this case the resulting tree is an empty tree with zero profit. Therefore an optimum solution for this case is the tree resulting from adding v and its connecting edge $e(v, v_1)$ to an optimum solution of $DP(sub(v_1), B_v - d_T(v, v_1), \hat{B}'_v)$. The set of served clients are the ones covered by $DP(sub(v_1), B_v - d_T(v, v_1), \hat{B}'_v)$ union all the clients reside in v . Note that as $DP(sub(v_1), B_v - d_T(v, v_1), \hat{B}'_v)$ is a smaller sub-problem we can assume that we have an optimum solution for this sub-problem. The running time for this case is $O(1)$.

Case 3: Let v_1, v_2 be the two children of v . We explain for the case when H_v^* contains v_1 but not v_2 , the alternative is symmetric. We build H_v as follows. Similar to Case 2, B_v has to be at least $d_T(v, v_1)$ in order to compensate for connecting v_1 to v . In this case we have the total budget of \hat{B}'_v for the service costs of both the clients in $sub(v_1)$ and clients in $sub(v_2)$. If we spend \hat{B}'_{v_1} for the service costs of the clients in $sub(v_1)$ then the budget for the service costs of clients in $sub(v_2)$ is $\hat{B}'_{v_2} = \hat{B}'_v - \hat{B}'_{v_1}$. Note that the service cost for the clients reside in v is zero therefore we serve them for free. Given \hat{B}'_{v_1} we can find an optimum tree and hence the optimum set of served clients in $sub(v_1)$ by taking the solution of the smaller sub-problem $DP(sub(v_1), B_v - d_T(v, v_1), \hat{B}'_{v_1})$. Note that the nearest node in H_v to the clients in $sub(v_2)$ is actually v as we know H_v does not contain v_2 . Therefore given \hat{B}'_{v_2} , we solve the corresponding instance of the Knapsack problem for the nodes in $sub(v_2)$ with weight bound \hat{B}'_{v_2} similar to Case 1. We iterate over all possible values of \hat{B}'_{v_1} from the set $\{0, 1, \dots, \hat{B}'_v\}$ and take the value for which we can collect the most profit from the clients in $sub(v_1)$ and clients in $sub(v_2)$. The running time for this case is $O(\hat{B}'_v^2 \cdot |C|)$ because we have a loop of $O(\hat{B}'_v)$ to iterate over the values of \hat{B}'_{v_1} where for each iteration we have to solve an instance of the Knapsack problem with $O(\hat{B}'_v \cdot |C|)$.

Case 4: Let v_1, v_2 be the two children of v . Similar to Case 2, B_v has to be at least $d_T(v, v_1) + d_T(v, v_2)$ in order to compensate for connecting both v_1 and v_2 to v in the resulting tree H_v . Here, we have total service cost bound \hat{B}'_v and length bound $B_v - d_T(v, v_1) + d_T(v, v_2)$ to spend in the nodes in $sub(v_1)$ and $sub(v_2)$. If we spend \hat{B}'_{v_1} for the service costs for the clients in $sub(v_1)$ and B_{v_1} for the length bound in $sub(v_1)$ we can spend at most $\hat{B}'_v - \hat{B}'_{v_1}$ for the service costs and $B_v - \hat{B}'_{v_1} - d_T(v, v_1) - d_T(v, v_2)$ for the length in $sub(v_2)$. Given \hat{B}'_{v_1} and B_{v_1} we can find an optimum tree and hence the optimum set of served clients in $sub(v_1)$ by taking the solution of smaller sub-problem $DP(sub(v_1), B_{v_1}, \hat{B}'_{v_1})$. Similarly we find an optimum tree and hence the optimum set of served clients in $sub(v_2)$ by taking the solution of smaller sub-problem $DP(sub(v_2), B_v - B_{v_1} - d_T(v, v_1) - d_T(v, v_2), \hat{B}'_v - \hat{B}'_{v_1})$. We iterate over all possible values of B_{v_1} from the set $\{0, \dots, B_v - d_T(v, v_1) - d_T(v, v_2)\}$ and all possible values of \hat{B}'_{v_1} from the set $\{0, \dots, \hat{B}'_v\}$ and take the combination which collects the maximum profit from the clients in $sub(v_1)$ and $sub(v_2)$. The running time for this case is $O(B_v \cdot \hat{B}'_v)$.

Note that the running time of our algorithm is the sum of running time of all the cases and hence $O(B_v \cdot \hat{B}'_v + \hat{B}'_v^2 \cdot |C|)$. \square

Note that if node u is a leaf then sub-problem $DP(u, B_u, \hat{B}'_u)$ does not depend on any smaller sub-problem and can be solved efficiently by Claim 5. In order to avoid the dependency on the smaller sub-problems in Claim 5 we solve the smaller sub-problems first. More formally, we run the following algorithm.

1. For each node $u \in V$ in non-increasing order of the level number do:
 - (a) For each value B_u in $\{0, \dots, B\}$ do:
 - i. For each value \hat{B}'_u in $\{0, \dots, \hat{B}'\}$ do:
 - A. Solve sub-problem $DP(sub(u), B_u, \hat{B}'_u)$ using Claim 5 and store the solution in the memory.

Note that in the above algorithm when we use Claim 5 to solve sub-problem $DP(sub(u), B_u, \hat{B}'_u)$ all of its sub-problems are solved beforehand. The solution to the STSCST problem is the solution for the sub-problem $DP(r, B, \hat{B}')$ by its definition (see Definition 14). As we have the total of $O(n \cdot B \cdot \hat{B}')$ sub-problem and each of them take $O(B_v \cdot \hat{B}'_v + \hat{B}'_v^2 \cdot |C|)$ time to solve by Claim 5, the total running time of our algorithm is $O(n \cdot B^2 \cdot \hat{B}'^2 + n \cdot |C| \cdot B \cdot \hat{B}'^3)$. \square

In the following we show how Lemma 8 helps us solve the NPCST problem and finishes the prove of Theorem 4. Note that the maximum possible value of B' is $\sum_{c \in C} \frac{M}{\min_c t_c} \cdot \theta_v$ where M is the distance between the farthest pair of nodes over all the trees in the set $\{T_1, T_2, \dots, T_p\}$. In order to solve instance $\mathcal{I} = (V, d, r, C, L)$ of the NPCST problem, for each tree T_i in distribution π and each value 2^j where $0 \leq j \leq \left\lceil \log \left(\sum_{c \in C} \frac{M}{\min_c t_c} \cdot \theta_v \right) \right\rceil$, we define a corresponding instance $\mathcal{I}_{i,j} = (T_i, r, C, 4 \cdot A \cdot \log n \cdot L, 2^j)$ of TSCST as follows. The nodes set, root, and clients in $\mathcal{I}_{i,j}$ are the same as \mathcal{I} but instead of distance function d we have the distance function d_{T_i} . Let $H_{i,j}$ be the solution for instance $\mathcal{I}_{i,j}$ using Lemma 8. We transform tree $H_{i,j}$ to the original metric $(\mathcal{M}(V, d))$ to obtain tree $\hat{H}_{i,j}$. Let $C_{i,j}$ be the set of all served clients by $\hat{H}_{i,j}$ in d with violation $16 \cdot A \cdot \log n$ in the neighborhoods, *i.e.*, $C_{i,j}$ contains each client $c \in C$ such that the neighborhood ball $\mathcal{B}(c, 16 \cdot A \cdot \log n \cdot t_c)$ is hit by $\hat{H}_{i,j}$. Our solution to the NPCST problem is tree $\hat{H}_{i,j}$ for $i \in [p]$ and $j \in [A \log n]$ whose $\sum_{c \in C_{i,j}} \theta_c$ is maximized.

In the following we show that $\hat{H}_{i,j}$ indeed serves $\frac{1}{2} \text{OPT}$. Let tree T^* be an optimum solution to the NPCST problem in the original metric d which has maximum length L and hits \mathcal{B} -ball of a subset $(C^* \subseteq C)$ of the clients such that their total profit is optimum ($\text{OPT} = \sum_{v \in C^*} \theta_v$). Similar to Definition 12, we define the service cost of T^* for the clients in C^* to be $\sum_{c \in C^*} \theta_v \cdot \frac{d(v, T^*)}{t_c}$. Note that as T^* touches the neighborhood ball $(\mathcal{B}(c, t_c))$ of each client c that it serves, $\sum_{c \in C^*} \theta_v \cdot \frac{d(v, T^*)}{t_c}$ can be at most OPT (Fact 1).

The expected service cost of T^* in the distribution π of the trees is at most $A \cdot \log n \cdot \text{OPT}$ by the following equations.

$$\begin{aligned}
\sum_{i=1}^p \pi(T_i) \sum_{c \in C^*} \frac{d_{T_i}(T^*, c)}{t_c} \cdot \theta_c &= \sum_{c \in C^*} \sum_{i=1}^p \pi(T_i) \frac{d_{T_i}(T^*, c)}{t_c} \cdot \theta_c && \text{Linearity of the expectations} \\
&< A \cdot \log n \cdot \sum_{c \in C^*} \frac{d(T^*, c)}{t_c} \theta_c && \text{By Inequality (20)} \\
&= A \cdot \log n \cdot \text{OPT}
\end{aligned}$$

Similar to the above we can conclude that the expected length of T^* in distribution π of the trees is $A \cdot \log n \cdot L$. By Markov's inequality, we can conclude that there are at least $\frac{3}{4} \cdot p$ trees in distribution π in which length of T^* is at most $4 \cdot A \cdot \log n \cdot L$. Similarly, there are at least $\frac{3}{4} \cdot p$ (possibly different) trees in distribution π in which the service cost of T^* is at most $4 \cdot A \cdot \log n \cdot \text{OPT}$. The intersection

of these two sets of trees contains at least $\frac{1}{2}p$ trees. Therefore, among T_1, T_2, \dots, T_p there exists at least one tree T_k whose corresponding instance $(T_k = (V, E), r, C, 4 \cdot A \cdot \log n \cdot L, 4 \cdot A \cdot \log n \cdot \text{OPT})$ of TSCST (see Definition 12) has a solution which collects at least OPT profits. Hence, the optimal value of instance $\mathcal{I}_{k,j} = \lceil \log(4 \cdot A \cdot \log n \cdot \text{OPT}) \rceil$ of TSCST that we run in our algorithm is at least OPT (Fact 2).

First note that the total number of instances $\mathcal{I}_{i,j}$ s, we solve for TSCST is at most polynomial in terms of input as $i \in O(n \log n)$ and $0 \leq j \leq \left\lceil \log \left(\sum_{c \in C} \frac{M}{\min_c t_c} \cdot \theta_v \right) \right\rceil$. Consider our algorithm when it solves instance $\mathcal{I}_{i,j}$ where $i = k$ and $j = \lceil \log(4 \cdot A \cdot \log n \cdot \text{OPT}) \rceil$ which are taken from Fact 2. We prove that the total profit of the clients in $C_{i,j}$ served by tree $\hat{H}_{i,j}$ is at least $\frac{1}{2}\text{OPT}$ which completes the proof of Theorem 4. From Fact 2 we know that $\hat{H}_{i,j}$ collects OPT profits with total service cost at most $8 \cdot A \cdot \log n \cdot \text{OPT}$. Now consider set $C'_{i,j}$ of all the clients c whose $\frac{d(c, \hat{H}_{i,j})}{t_c}$ is at most $16 \cdot A \log n$, from Markov's inequality we can conclude $\sum_{c \in C'_{i,j}} \theta_c$ is at least $\frac{1}{2} \sum_{c \in C_{i,j}} \theta_c = \frac{1}{2}\text{OPT}$. Therefore, $\hat{H}_{i,j}$ whose length is at most $4 \cdot A \cdot \log n \cdot L$ collects at least $\frac{\text{OPT}}{2}$ profits while violates the neighborhoods by factor $16 \cdot A \log n$.

5.3 Description of an $(O(1), O(1), O(\log n))$ -Approximation Algorithm for the NPCST Problem

Let B_c denotes $\mathcal{B}(c, t_c)$. We ignore clients c for which $d(r, B_c) > L$ (these nodes are not covered in any optimal solution). Let θ_{max} be the largest profit of a client among the remaining clients. Note that $\theta_{max} \leq \text{OPT}$ since an optimal tree can just visit the client with profit θ_{max} . Consider the following natural LP-relaxation. We have edge variables $\{x_e\}$ indicating if e is in the resulting tree or not, and a variable z_c , $0 \leq z_c \leq 1$, for every client c indicating if B_c is touched by the tree. The objective function of the LP is to maximize $\sum_{c \in C} \theta_c \cdot z_c$, and the constraints are as follow. The first set of constraints are $\sum_{e \in \delta(S)} x_e \geq z_c$, for every node c and set $S \subseteq V$ containing B_c and not containing r where $\delta(S)$ is the set of all the edges that have exactly one endpoint in S . The second set of constraints are $\sum_e d_e \cdot x_e \leq L$ ensuring the cost budget of the tree.

This LP can be solved efficiently by noting that there is a separation oracle for the first set of constraints in which we contract each set B_c and use min-cut max-flow algorithm between the contracted super node and root r . Let (x^*, z^*) be an optimal solution and $K = \sum_c z_c^*$. We can ignore each node c in which $z_c^* \leq \frac{1}{2 \cdot n}$; this decreases the objective value by at most $\frac{\theta_{max}}{2}$. Therefore the remaining clients have total objective value of $O(\text{OPT})$ since $\text{OPT} > \theta_{max}$. The remaining clients with value z_c^* can be bucketed into $O(\log n)$ groups as $z_c^* \leq 1$, where the z_c^* values in each group are within a factor of 2 of each other. The contribution from one of these buckets is at least $\text{OPT}/\log n$, and we focus on such a bucket $C' \subseteq C$.

We perform a facility-location style clustering of the B_c balls for the clients in C' (We use similar ideas for Max-MR in Section 7). Repeatedly pick the client with smallest neighborhood radius t_c , and remove all terminals w such that $d(c, w) \leq 10 \max(t_c, t_w) = 10t_w$ (the value 10 is an arbitrary large enough constant), and set the representative of w ($rep(w)$) to be c . Now, we can pretend that w is co-located with c , and that $z_w^* = z_c^*$ and hence the first set of constraints still holds; this loses only a constant factor in the profit (since z_c^* and z_w^* are within a factor of 2 of each other) and gives a constant-factor violation in the radius. We show the subset of C' that we pick as the cluster centers by C'' . Now, we contract the ball B_c for the cluster centers, and set the profit of each contracted node c to be $\sum_{w \in C': rep(w)=c \text{ or } w=c} \theta_w$.

Note that $(x^*, 1 - z^*)$ induces a solution to the Prize Collecting Steiner Tree (PCST) problem instance defined by the contracted nodes. We can build a tree T of cost $O(L)$ such that the expected profit of the contracted nodes connected by T is in fact at least $\sum_{c \in C''} \theta_c \cdot z_c^*$ which is in $\Omega(\text{OPT}/\log n)$. This can be done by using results of Bang-Jensen et al. [BJFJ95], which gives another way of obtaining a 2-approximation algorithm for Prize Collecting Steiner Tree (PCST). The result of Bang-Jensen et al. [BJFJ95] implies that from x^* we can build a convex combination of trees such that each client c appears in z_c^* fractional number of the trees. Note that $\sum_{c \in C''} \theta_c \cdot z_c^* = \Omega(\text{OPT}/\log n)$ and the average cost of the trees is $\sum_e x_e^* \leq L$. By Markov inequalities we can conclude that at least $\frac{3}{4}$ of the trees have cost at most $4 \cdot L$. Similarly, by Markov Inequality at least $\frac{3}{4}$ of the trees collect at least $\frac{1}{4} \Omega(\text{OPT}/\log n)$ amount of profit. Therefore we conclude that at least half of the trees have length at most $4 \cdot L$ and collect at least $\frac{1}{4} \Omega(\text{OPT}/\log n)$ amount of profit. Let T be one of them.

Now we take T and uncontract the super nodes. We need to make T connected in the uncontracted graph, which can be done by connecting all the nodes of T incident to B_c to c , where c is a cluster center. Since any two cluster centers are far apart, the extra cost incurred in this step can be charged to the edges of T incident to the contracted node corresponding to B_c (The argument here is the same as the argument for Max-MR in Section 7).

5.4 Description of a Quasi Polynomial Time $(1, 1, O(\log n))$ -Approximation Algorithm for the Separation Oracle of DLP

Chekuri and Pal [CP05] introduced the submodular orienteering problem. The inputs to the submodular orienteering problem are directed graph G , source node s , destination node t , budget B , and a submodular function f over the subsets of nodes. The objective is to find a walk starting from s and ending at t with length at most B such that the value of f over the visited nodes is maximized. Chekuri and Pal [CP05] give an $O(\log \text{OPT})$ -approximation algorithm to this problem where OPT is the optimum value of function f .

Recall that the separation oracle for our dual LP (DLP) is the following. We want to find a walk starting at a root node r with length at most B in order to hit clients' balls such that the sum of profits of clients that are hit is maximized. First note that the sum of profits of clients that are hit by a path is a submodular function of the nodes in the path. Second, we scale down all the profits such that the sum of profits is polynomially bounded in the number of nodes (n), hence the approximation factor of Chekuri and Pal [CP05]'s algorithm will turn into $O(\log n)$.

Now we can solve the separation oracle using Chekuri and Pal's algorithm by looping over all the nodes as a possible destination node and take a walk which collects the highest sum of profits.

6 NPCST in Euclidean Space (Proof of Theorem 5)

We prove formally Theorem 5 in the 2-dimensional Euclidean space. We believe that our method can be generalized to the D -dimensional space; the changes needed to do so are brought at the end of this section. We ignore optimizing the constants to keep the algorithm and its analysis easier to explain.

Let us assume that the radius of the biggest neighborhood is $\frac{M}{2}$. We tile the plane with regular hexagons of side length M so that every point in the plane is covered with exactly one tile and root node r is at the center of a hexagon (as shown in Figure 2).

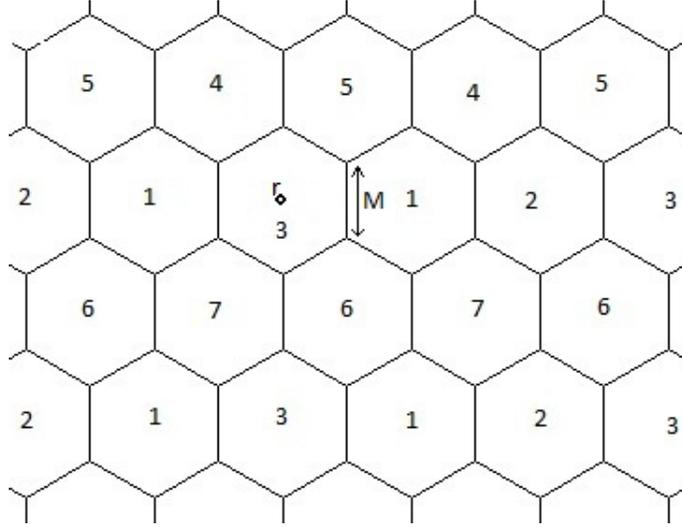


Figure 2: Tiling the plane with regular hexagons. The digits in each hexagon denotes its color. Root node r is centered at one of the tiles and the length of the edges of the hexagons is M .

We color the hexagons with 7 colors such that no two neighboring hexagons get the same color. This coloring can be done by coloring each non-colored hexagon with one of the 7 colors that is not used in a neighboring colored hexagon. The coloring is possible for all the non-colored hexagon since it has 6 neighboring hexagons and there are 7 colors (see Figure 2).

We construct an auxiliary graph from G in order to assign the profit of each neighborhoods to a single node which relaxes the problem to the Budgeted Prize Collecting Steiner Tree (BPCST) [JMP00]. In the BPCST problem we are given a graph where all the nodes have a profit, a root node r , and a budget L ; the objective is to find a tree with length at most L which the sum of profits of the nodes it contains is maximum. There is a $(4 + \epsilon)$ -approximation algorithm for BPCST due to Chekuri *et al.* [CKP08].

We call a hexagonal tile an *occupied-tile* if it contains at least one node of G . We correspond each occupied-tile to a node in V . We call the nodes that are corresponded to the occupied tiles *center* nodes. If the occupied-tile contains root r then its center node is r otherwise the center node is an arbitrary node which is inside the tile. We construct the auxiliary graph $(\hat{G} = (\hat{V}, \hat{E}))$ as follow. $\hat{V} \subseteq V$ is the set of all the center nodes that correspond to occupied-tiles so that there is a bijection between nodes in \hat{V} and the occupied-tiles. Note that root r is a center node and is in \hat{V} since r is assigned to an occupied-tile. $\hat{E} \subseteq E$ is the subset of the edges in E that have both endpoints in \hat{V} . We assign each client to all the occupied tiles its \mathcal{B} -ball intersects since each client resides in a node of G at least one occupied-tile exists for each client. The profit of each center node in \hat{V} is the sum of profits of the clients assigned to its corresponding occupied-tile.

Neighborhood Prize Collecting Steiner Tree Algorithm (NPCSTA) showed in Figure 3 is our $(4P + 1, 35, 12 + \epsilon)$ -approximation algorithm for Euclidean NPCST where P is the ratio of the largest neighborhood radius over the shortest one. Remember NPCSTA being $(4P + 1, 35, 12 + \epsilon)$ -approximation algorithm for the instance (V, d, r, C, L) of the ENPCST problem, means it finds a

Steiner tree T with the following properties. T is said to hit a client c with neighborhood ball $\mathcal{B}(c, t_c)$ if T has at least one node in $\mathcal{B}(c, t_c \cdot (4P+1))$, the weight of T is at most $35 \cdot L$, and sum of the profits of the clients got hit by T is at least $\frac{1}{12+\epsilon} \text{OPT}$ where OPT is the amount of profit an optimum tree collects with no violation in any bound. We use the $(4 + \epsilon')$ -approximation algorithm of [CKP08] for the BPCST problem as a black box in our algorithm (here we set $\epsilon' = \frac{\epsilon}{3}$).

1. For each color $i = 1, 2, \dots, 7$:
 - (a) Find a tree $\hat{\mathbf{T}}_i$ rooted at r in $\hat{\mathbf{G}}$ with weight at most $5 \cdot L$ which collects the maximum profit only from the nodes that correspond to the occupied-tiles of color i using the $(4 + \epsilon')$ -approximation algorithm of [CKP08] for the BPCST problem.
2. Return $T = \hat{\mathbf{T}}_1 \cup \dots \cup \hat{\mathbf{T}}_7$ as the result. The set of covered clients are all the clients that are assigned to the occupied-tiles whose center node is in T .

Figure 3: Neighborhood Prize Collecting Steiner Tree Algorithm

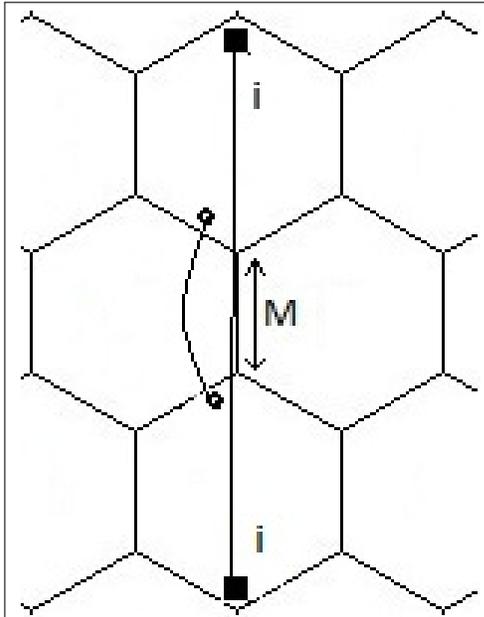
We prove the following lemma which implies Theorem 5.

Lemma 10. *NPCSTA (showed in Figure 3) is a $(4P + 1, 35, 12 + \epsilon)$ -approximation algorithm for the ENPCST problem where P is the ratio of the largest neighborhood radius over the shortest one.*

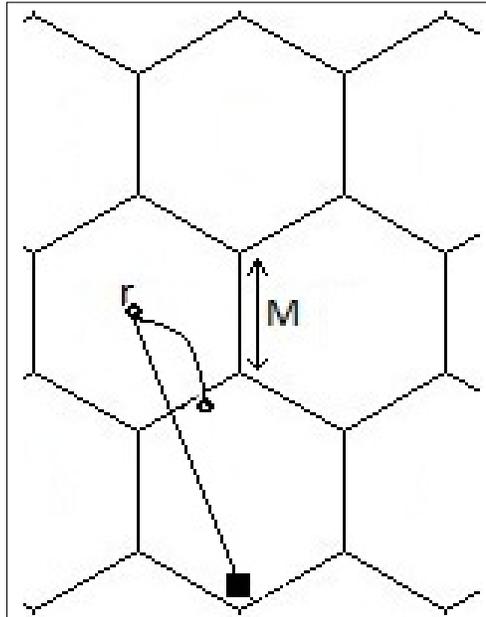
Proof. We use the structure of an optimum solution in order to show that there is also a tree in graph $\hat{\mathbf{G}}$ which also collects the optimal profit with small violations in the neighborhood balls and the weight bound. Let T^* be an optimal tree collecting OPT profit. Let $\hat{\mathbf{T}}_i^*$ be the projection of T^* on the set of nodes in $\hat{\mathbf{G}}$ that correspond to the tiles of color i plus root node r . More specifically, we build $\hat{\mathbf{T}}_i^*$ as follows. If T^* contains a node in a tile of color i then $\hat{\mathbf{T}}_i^*$ contains the node in $\hat{\mathbf{G}}$ that corresponds to that tile. $\hat{\mathbf{T}}_i^*$ also contains root node r . There is an edge ($\hat{e} = \hat{v}_1 \hat{v}_2$) between two nodes (\hat{v}_1, \hat{v}_2) of $\hat{\mathbf{T}}_i^*$ if and only if there is a node v_1 in T^* that resides in the tile corresponding to \hat{v}_1 and a node v_2 in T^* that resides in the tile corresponding to \hat{v}_2 such that the unique path between v_1 and v_2 in T^* does not pass through any other node that resides in a tile that corresponds to a node of $\hat{\mathbf{T}}_i^*$. We call the unique path which was the reason for adding an edge between \hat{v}_1 and \hat{v}_2 a *creator* path. In the following claim we prove that the weight of $\hat{\mathbf{T}}_i^*$ is at most 5 times larger than T_i^* .

Claim 6. *The weight of $\hat{\mathbf{T}}_i^*$ is at most 5 times larger than the weight of T_i^* .*

Proof. The weight of $\hat{\mathbf{T}}_i^*$ can be larger than T_i^* as we select the center nodes of the tiles instead of the original nodes in T_i^* . By the construction of $\hat{\mathbf{T}}_i^*$, each of its edge is added to $\hat{\mathbf{T}}_i^*$ because of a path in T_i^* . Moreover each path of T_i^* contribute in adding at most one edge in $\hat{\mathbf{T}}_i^*$ since only the endpoints of the path are in the occupied-tiles whose center is in $\hat{\mathbf{T}}_i^*$ and all the other nodes are inside the other occupied-tiles (the ones that neither colored i nor contain r). Therefore, we can charge each edge of $\hat{\mathbf{T}}_i^*$ to their corresponding path. There are two types of edges in $\hat{\mathbf{T}}_i^*$. The first type (Type One) contains the edges between two center nodes of tiles with color i (see Figure 4a) and the second type (Type Two) contains the edges between r and a center node of a tile with color i (see Figure 4b).



(a) An edge between two center nodes in the tiles of color i compared to its corresponding path in T^* (creator path) shown as the curved line.



(b) An edge between the root node (r) and a center node in \hat{T}_i^* compared to its corresponding path in T^* (creator path) shown as the curved line.

Figure 4: Edges in \hat{T}_i^* compared to paths in T^* . The squares represents the center nodes in \hat{V} and circles represent the non-center nodes in V . The straight lines are the edges of \hat{T}_i^* and the curvy lines are the paths of T^* (the creator paths for the edges of \hat{T}_i^*).

Note that we require no two tiles of color i be neighbor of each other (This paragraph indeed shows the crucial point of coloring). Therefore, each creator path in T_i^* has length at least M if its endpoints are the closest points in the two tiles (see Figure Figure 4a). The corresponding edge in \hat{T}_i^* can have length $5M$ since its two center nodes can be the farthest points of the two tiles as opposed to the creator path (see Figure 4a). Therefore in the worst case the length of the Type One edges in \hat{T}_i^* are 5 times larger than its creator path.

The end points of creator paths for the Type Two edges can be in the neighboring tiles. The length of the creator paths for the Type Two edges is at least $\frac{\sqrt{3}}{2}M$ which is the distance of the center of the tile (node r) to the middle of a side (see Figure 4b). The length of the corresponding edge in \hat{T}_i^* can be at most $2M$ units of length larger than the creator path since the center node of the tile with color i can be at the farthest point of the tile (note that the distance between two farthest point in a tile is $2M$). Therefore in the worst case the maximum multiplicative inflation in the Type Two edges of \hat{T}_i^* is $\frac{(2+\sqrt{3}/2)M}{\sqrt{3}/2 \cdot M} \approx 3.31$ (see Figure 4b, in fact the multiplicative factor is smaller but we upper bound it by 3.31).

Because the multiplicative increase compared to the creator paths for the Type One edges is at most 5 and for the Type Two edges is at most 3.31, the weight of \hat{T}_i^* is at most 5 times larger than the weight of T_i^* .

In order to prove lemma we show that the resulting tree of NPCSTA (T) satisfies all the three criteria in the following three claims. We start with the first criteria to show that the multiplicative increase in the neighborhoods' radius is $4P + 1$.

Claim 7. *For each client c that is in the set of served clients by NPCSTA, T visits a node whose distance from c is $4P + 1$ times the radius of the neighborhood of c (t_c).*

Proof. Let us assume that c is in the set of served clients because its neighborhood intersects with a tile of color i whose center is in T_i . Similar to proof of Claim 6 the center node of the tile can be far from the point where neighborhood of c intersects. Note that this distance is at most $2M$. Therefore T_i visits a node which is at most $2M + t_c$ away from c where t_c is the radius of c 's neighborhood. Because the ratio of the largest neighborhood to the smallest one is P , T_i and hence T contain a node which has distance $(4P + 1) \cdot t_c$ from c . \square

In the following we prove that the weight of the output of NPCSTA (T) is at most $35L$ (the second criterion of the tri-criteria approximation algorithm).

Claim 8. *The weight of T is at most $35L$.*

Proof. The weight of each tree T_i where $i \in \{1, \dots, 7\}$ is $5L$ and all of them contain root r and hence connected to each other. Weight of T is at most $35L$ because it is the union of all the T_i s. \square

The following claim proves the bound of the third criterion of NPCSTA and finishes the proof of the lemma.

Claim 9. *T collects at least $\frac{1}{12+\epsilon}$ fraction of the optimal profit.*

Proof. Let OPT_i be the amount of profit T^* collects from the clients (C_i) that are assigned to a tile of color i . Because the side length of each hexagon is M and no two tiles with color i intersects; the distance between two tiles of color i is at least M . As the radius of the largest neighborhood is

$\frac{M}{2}$ the \mathcal{B} -ball of each client in C_i is intersecting to exactly one tile of color i . Therefore, there is a unique way of assigning clients in C_i to the tiles of color i . From Claim 6 we can conclude that there is a tree \hat{T}_i^* which visits all the center nodes of the tiles with color i that contain at least one node of T^* . Therefore \hat{T}^* collects at least OPT_i profit from the i colored tiles and has weight $5L$. Thus, T_i which is the result of $(4 + \epsilon')$ -approximation algorithm of [CKP08] collects at least $\frac{1}{4+\epsilon'} \text{OPT}_i$.

Note that the radius of the maximum neighborhood is $\frac{M}{2}$ and the side length of the hexagons is M . Each neighborhood can intersect with at most 3 hexagons since if it intersects with 4 tiles two of them cannot be neighbor of each other and have distance M . Therefore the profit of each client contributes in at most 3 colors. From each color i , tree T collects at least $\frac{1}{4+\epsilon'} \text{OPT}_i$ and each client's profit can appear in at most 3 colors; the proof of the Claim follows by summing over all the 7 colors and setting $\epsilon' = \frac{\epsilon}{3}$. \square

The approach for generalizing the above algorithm to the D -dimensional space is the same as NPC-STA (Figure 3). We decompose the space into tiles using D -dimensional cubes and proceed as before. Each cube touches at most $O(2^D)$ other cubes so we have at most $O(2^D)$ colors. The farthest points between two points in a D -dimensional cube is $O(2^D)$. Therefore, we believe that our algorithm can be generalized to the D -dimensional Euclidean space and get $(O(2^D), O(2^D), O(2^D))$ -approximation algorithm.

7 Max-MR problem

The other objective which is usually considered in the movement framework is to minimize the maximum latency. This objective is taken from the applications when there is a deadline by which all the clients have to be served. In our application the deadline specifies the latest time we can serve the last client which can be referred to as minimizing the maximum latency. We refer to this objective in our setting as the Max-MR problem. We give a constant-factor approximation algorithm for the case when all repairmen have the same speed and prove Theorem 6. Client serving problems with max objective are studied thoroughly for lots of different scenarios [FHK76, LSLD92, EGK⁺04, AHL06, CVH08, XX09, XW10, KS11, NR12].

Proof. We start by guessing the optimum maximum latency (OPT) by which all the clients will be served. Our algorithm for a given guessed value (T) for the maximum latency, either finds paths for repairmen that serve all the clients with latency at most $10T$ or announce that $T \leq \text{OPT}$. Therefore, we can find an appropriate value T such that $\text{OPT} \leq T \leq \text{OPT}(1 + \epsilon)$ for a small positive constant ϵ by binary search over the range $[1, \frac{2 \cdot \text{MST}(G)}{\min_i v_i}]$. The upper value in the range is indeed an upper bound for the maximum latency which is the time required for the slowest repairman to visit all the nodes in the graph with a path obtained by doubling the edges of a MST.

Our algorithm proceed as follows. For a given T we assign to each client c a neighborhood $\mathcal{B}(c, t_c) = \{v | d(v, c) \leq t_c\}$ of radius $t_c = v'_c \cdot T$ where v'_j is the speed of c_j . Now, for each client c at least one repairman should visit at least one node in $\mathcal{B}(c_j, t_c)$. We run a clustering algorithm to cluster these neighborhoods as follows. We start with a client (c) whose neighborhood has the smallest radius and tag it as a *leader* client. We assign all the clients c' for which $\mathcal{B}(c, t_c) \cap \mathcal{B}(c', 9t_{c'}) \neq \emptyset$ to the client c and tag them *slave* clients, *i.e.*, a client is slave if its stretched neighborhood intersects with the neighborhood of a leader client. After that, we discard all the tagged clients and make a non-tagged client with the smallest neighborhood radius a leader and proceed as before. We continue tagging until all the clients get tagged.

Claim 10. *The distance between the neighborhoods of any two leader clients (c and c') is at least 8 times the radius of the larger neighborhood. Here the distance between neighborhood $\mathcal{B}(c, t_c)$ and neighborhood $\mathcal{B}(c', t_{c'})$ is the distance between the closest pair of nodes one in $\mathcal{B}(c, t_c)$ and the other in $\mathcal{B}(c', t_{c'})$.*

Proof. Without loss of generality assume $t_c \geq t_{c'}$. We prove the claim by contradiction. Assume the distance is less than $8t_c$, then $\mathcal{B}(c, 9t_c)$ intersects with the neighborhood of c' . This makes c tagged as a slave client of c' in the process of tagging when we tagged c' as a leader client. \square

We contract each leader ball to a super-node and take the metric completion of it in order to obtain graph $\hat{\mathbf{G}}$. We define an instance of the rooted version of the min-max k -tree cover [EGK⁺04, AHL06] over $\hat{\mathbf{G}}$. In the min-max k -tree cover we are given a set of terminals in a metric space along with k root nodes. The objective is to find k trees rooted in the root nodes to cover all the terminals in the graph such that the maximum length of the trees is minimized. In $\hat{\mathbf{G}}$ the terminals are the super-nodes and the root nodes are the starting locations of the repairmen (which may now be contracted into a super-node).

We use the 4-approximation algorithm of Even et al. [EGK⁺04] for the min-max k -tree cover to find k -trees in $\hat{\mathbf{G}}$ rooted at the root nodes and covering all the super-nodes. The set of paths for the repairmen are constructed as follows. For each repairman r we take the tree ($\hat{\mathbf{T}}_r$) rooted at the starting location of r found by the algorithm of Even et al. [EGK⁺04]. We double the edges of $\hat{\mathbf{T}}_r$, build an Eulerian walk, and obtain path $\hat{\mathbf{O}}_r$. Now, consider the edges of $\hat{\mathbf{O}}_r$ in the graph G when we uncontract each super-node to the neighborhood ball of its corresponding leader client. We reconnect $\hat{\mathbf{O}}_r$ in G by adding an extra edge for each pair of separated nodes that were a single super-node in $\hat{\mathbf{O}}_r$. Moreover, if the root node is in a super-node and is disconnected from the rest of the path, we reconnect it by adding an extra edge to the node that was in the same super node as the root in $\hat{\mathbf{O}}_r$. Let the connected path in G obtained from $\hat{\mathbf{O}}_r$ be P_r . If length of P_r is less than $10 \cdot v \cdot T$ for each repairman r where v is the common speed of all the repairmen, then set $\{P_r\}_{r \in R}$ is the result of our algorithm for the given T . Moreover, each client gets service from the closest node in one of the paths in $\{P_r\}_{r \in R}$. Otherwise if length of P_r is greater than $10 \cdot v \cdot T$ our algorithm announces that T is less than OPT (the optimum solution of Max-MR defined over G). We prove the following lemma which proves the correctness of our algorithm and finishes the proof of Theorem 6.

Lemma 11. *If the guessed value T is not smaller than OPT then the length of $P_r \leq 10 \cdot v \cdot T$ for each repairman r and the maximum latency of the clients is $10T$.*

Proof. Let us assume $\text{OPT} \leq T$. In the following claim we prove that the length of the paths of the repairmen is at most $10 \cdot v \cdot T$.

Claim 11. *Let repairman r be the repairman whose P_r has the maximum length. The length of P_r is at most $10 \cdot v \cdot T$.*

Proof. Note that every path assigned to the repairmen in an optimal solution to the Max-MR has length at most $v \cdot \text{OPT}$. The optimal solution of the min-max k -tree cover problem defined over $\hat{\mathbf{G}}$ is at most $v \cdot \text{OPT}$ because the paths in an optimal solution of Max-MR after contracted into the super nodes, are a candidate solution to the min-max k -tree cover in $\hat{\mathbf{G}}$. Therefore the length of the maximum tree in the solution we obtain from the 4-approximation algorithm is at most $4 \cdot v \cdot \text{OPT}$ and hence it is at most $4 \cdot v \cdot T$. Therefore the length of $\hat{\mathbf{O}}_r$ is at most $8 \cdot v \cdot T$ as it is obtained by doubling the edges of the tree assigned to r .

Path P_r obtained from $\hat{\mathbf{O}}_r$ by adding extra edges inside the neighborhood of the leader clients. We start from the starting location of r and go along path P_r . Whenever we encounter an edge inside a leader client's neighborhood which is not in $\hat{\mathbf{O}}_r$ we charge its length to the path in $\hat{\mathbf{O}}_r$ that goes from the neighborhood of the current leader client to the neighborhood of the next leader client. From Claim 10 we know that the length of the path to which we charged the length of the extra edge is at least 8 times the radius of the neighborhood. Because the length of the extra edge is at most twice the radius of the neighborhood, the length of the extra edge is at most $\frac{1}{4}$ of the path it is charged to. Therefore the total length of P_r is at most $10 \cdot v \cdot T$. \square

In the following claim we prove that for each client c there exist a node with distance at most $10t_c$ from c which resides in the path of a repairman.

Claim. For each client c there exist at least one repairman r such that there is at least one node which is in $\mathcal{B}(c, 10t_c)$ and in P_r .

Proof. If c is a leader client then the claim follows directly as the paths of the repairmen visit at least one node from the neighborhood ($\mathcal{B}(c, t_c)$) of each leader client. If c is a slave client $\mathcal{B}(c, 9t_c)$ intersects with the neighborhood of a leader client c' . Note that $t_c \geq t_{c'}$ since we tag the clients as leaders in the increasing order of their neighborhood radius. Therefore the distance of c from any node in $\mathcal{B}(c', t_{c'})$ is at most $10t_c$. As c' is a leader client the paths of repairmen visit at least one node from $\mathcal{B}(c', t_{c'})$. Therefore, there exists a node which is in $\mathcal{B}(c, 10t_c)$ and resides in the path of a repairman.

From Claim 11 we conclude that all the repairman can travel their assigned paths by time $10T$. From Claim 7 we conclude that each client by time $10T$ can go to a node which is in the path of a repairman. Therefore, the maximum latency of the clients is $10T$. \square

8 Acknowledgment

The authors would like to thank anonymous reviewers of Approx+Random 2013 conference and Transaction on Algorithms journal who provided us with helpful comments and more importantly alternative proofs which give an $(O(1), O(1), O(\log n))$ -approximation algorithm (Section 5.3) and a quasi polynomial time $(1, 1, O(\log n))$ -approximation algorithm (Section 5.4) for the NPCST problem.

References

- [AFS13] S. Ahmadian, Z. Friggstad, and C. Swamy. Local-search based approximation algorithms for mobile facility location problems. *To appear in ACM-SIAM Symposium on Discrete algorithms*, 2013.
- [AHL06] E. M. Arkin, R. Hassin, and A. Levin. Approximations for minimum and min-max vehicle routing problems. *Journal of Algorithms*, 59:1–18, 2006.
- [Ama10] Amazon. Amazon locker delivery. <http://www.amazon.com/gp/help/customer/display.html?nodeId=200689010>, 2010.
- [Bar96] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *IEEE Symposium on Foundations of Computer Science*, pages 184–193, 1996.

- [Bar98] Y. Bartal. On approximating arbitrary metrics by tree metrics. In *ACM symposium on Theory of computing*, pages 161–168, 1998.
- [BBCM04] N. Bansal, A. Blum, S. Chawla, and A. Meyerson. Approximation algorithms for deadline-TSP and vehicle routing with time-windows. In *ACM symposium on Theory of computing*, pages 166–174, 2004.
- [BCC⁺94] A. Blum, P. Chalasani, D. Coppersmith, B. Pulleyblank, P. Raghavan, and M. Sudan. The minimum latency problem. In *ACM Symposium on Theory of Computing*, pages 163–171, 1994.
- [BCK⁺07] A. Blum, S. Chawla, D.R. Karger, T. Lane, A. Meyerson, and M. Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM Journal on Computing*, 37(2):653–670, 2007.
- [BDZ11] P. Berman, E. D. Demaine, and M. Zadimoghaddam. $O(1)$ -approximations for maximum movement problems. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 62–74, 2011.
- [BJFJ95] Jorgen Bang-Jensen, Andras Frank, and Bill Jackson. Preserving and increasing local edge-connectivity in mixed graphs. *SIAM J. Discret. Math.*, 8(2):155–178, 1995.
- [BYCR93] R.A. Baeza-Yates, J.C. Culberson, and G.J.E. Rawlins. Searching in the plane. *Information and Computation*, 106:234–234, 1993.
- [CCG⁺98] M. Charikar, C. Chekuri, A. Goel, S. Guha, and S. Plotkin. Approximating a finite metric by a small number of tree metrics. In *IEEE Symposium on Foundations of Computer Science*, pages 379–388, 1998.
- [CGRT03] K. Chaudhuri, B. Godfrey, S. Rao, and K. Talwar. Paths, trees, and minimum latency tours. In *IEEE Symposium on Foundations of Computer Science*, pages 36–45, 2003.
- [CHP⁺04a] P. Corke, S. Hrabar, R. Peterson, D. Rus, S. Saripalli, and G. Sukhatme. Autonomous deployment and repair of a sensor network using an unmanned aerial vehicle. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 3602–3608, 2004.
- [CHP⁺04b] P. Corke, S. Hrabar, R. Peterson, D. Rus, S. Saripalli, and G. Sukhatme. Deployment and connectivity repair of a sensor net with a flying robot. In *International Symposium on Experimental Robotics*, 2004.
- [CK04] C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. *Approximation, Randomization, and Combinatorial Optimization*, pages 72–83, 2004.
- [CKP08] C. Chekuri, N. Korula, and M. Pal. Improved algorithms for orienteering and related problems. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 661–670, 2008.
- [CP05] C. Chekuri and M. Pal. A recursive greedy algorithm for walks in directed graphs. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 245–253, Oct 2005.
- [CS11] D. Chakrabarty and C. Swamy. Facility location with client latencies: linear programming based techniques for minimum latency problems. *Integer Programming and Combinatorial Optimization*, pages 92–103, 2011.
- [CVH08] A. M. Campbell, D. Vandenbussche, and W. Hermann. Routing for relief efforts. *Transportation Science*, 42:127–145, 2008.
- [DBGK⁺05] M. De Berg, J. Gudmundsson, M.J. Katz, C. Levkopoulos, M.H. Overmars, and A.F. van der Stappen. TSP with neighborhoods of varying size. *Journal of Algorithms*, 57(1):22–36, 2005.
- [DHM09a] E. D. Demaine, M. Hajiaghayi, and D. Marx. Minimizing movement: Fixed-parameter tractability. In *European Symposium on Algorithms*, volume 5757, pages 718–729, 2009.
- [DHM⁺09b] E.D. Demaine, M.T. Hajiaghayi, H. Mahini, A.S. Sayedi-Roshkhar, S. Oveisgharan, and M. Zadimoghaddam. Minimizing movement. *ACM Transactions on Algorithms*, 5(3):30, 2009.
- [DM01] A. Dumitrescu and J.S.B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. In *ACM-SIAM Symposium on Discrete algorithms*, pages 38–46, 2001.
- [DMM⁺97] S. Doddi, M. V. Marathe, A. Mirzaian, B. M. E. Moret, and B. Zhu. Map labeling and its generalizations. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 148–157, 1997.

- [EFS06] K. Elbassioni, A. Fishkin, and R. Sitters. On approximating the TSP with intersecting neighborhoods. *International Symposium on Algorithms and Computation*, pages 213–222, 2006.
- [EGK⁺04] G. Even, N. Garg, J. Konemann, R. Ravi, and A. Sinha. Min-max tree covers of graphs. *Operations Research Letters*, 32(4):309–315, 2004.
- [FG04] C. Feremans and A. Grigoriev. *Approximation schemes for the generalized geometric problems with geographic clustering*. Maastricht research school of Economics of TEchnology and ORganizations, 2004.
- [FHK76] G.N. Frederickson, M.S. Hecht, and C.E. Kim. Approximation algorithms for some routing problems. In *IEEE Symposium on Foundations of Computer Science*, pages 216–227, 1976.
- [FHR03] J. Fakcharoenphol, C. Harrelson, and S. Rao. The k -traveling repairman problem. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 655–664, 2003.
- [FRT04] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004.
- [FS08] Z. Friggstad and M. R. Salavatipour. Minimizing movement in mobile facility location problems. In *IEEE Symposium on Foundations of Computer Science*, pages 357–366, 2008.
- [FSS13] Zachary Friggstad, Mohammad R. Salavatipour, and Zoya Svitkina. Asymmetric traveling salesman path and directed latency problems. *SIAM Journal on Computing*, 42(4):1596–1619, 2013.
- [GK98] M. Goemans and J. Kleinberg. An improved approximation ratio for the minimum latency problem. *Mathematical Programming*, 82(1):111–124, 1998.
- [GKK⁺01] Anupam Gupta, Jon Kleinberg, Amit Kumar, Rajeev Rastogi, and Bulent Yener. Provisioning a virtual private network: a network design problem for multicommodity flow. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 389–398, 2001.
- [GLS93] M. Grotschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimizations*. Springer-Verlag, 1993.
- [HAB⁺03] T.-R. Hsiang, E. M. Arkin, M. A. Bender, S. P. Fekete, and Joseph S. B. Mitchell. Algorithms for rapidly dispersing robot swarms in unknown environments. In *Algorithmic Foundations of Robotics V*, pages 77–94. 2003.
- [HK03] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *ACM symposium on Theory of computing*, pages 585–594, 2003.
- [IK75] O. H. Ibarra and C. E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of ACM*, 22(4):463–468, 1975.
- [JBQZ04] M. Jiang, S. Bereg, Z. Qin, and B. Zhu. New bounds on map labeling with circular labels. In *International Symposium on Algorithms and Computation*, volume 3341, pages 606–617, 2004.
- [JMP00] D. S. Johnson, M. Minkoff, and S. Phillips. The prize collecting steiner tree problem: theory and practice. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 760–769, 2000.
- [JQQ⁺03] M. Jiang, J. Qian, Z. Qin, B. Zhu, and R. Cimikowski. A simple factor-3 approximation for labeling points with circles. *Information Processing Letters*, 87(2):101–105, 2003.
- [KS11] M.R. Khani and M.R. Salavatipour. Approximation algorithms for min-max tree cover and bounded tree cover problems. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 302–314, 2011.
- [LaV06] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [LSLD92] C. Li, D. Simchi-Levi, and M. Desrochers. On the distance constrained vehicle routing problem. *Operations Research*, 40:790–799, 1992.
- [Mit00] J.S.B. Mitchel. *Handbook of computational geometry*. Elsevier, 200.
- [Mit07] J.S.B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *ACM-SIAM symposium on Discrete algorithms*, pages 11–18, 2007.
- [Mit10] J. S.B. Mitchell. A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane. In *Symposium on Computational Geometry*, pages 183–191, 2010.

- [NR1]
- [NR12] V. Nagarajan and R. Ravi. Approximation algorithms for distance constrained vehicle routing problems. *Networks*, 59(2):209–214, 2012.
- [RW90] G. Reich and P. Widmayer. Beyond steiner’s problem: a vlsi oriented generalization. In *International Workshop on Graph-theoretic Concepts in Computer Science*, pages 196–210, 1990.
- [RW95] J. H. Reif and H. Wang. Social potential fields: a distributed behavioral control for autonomous robots. In *Workshop on Algorithmic Foundations of Robotics*, pages 331–345, 1995.
- [Ser10] The United States Postal Service. gopost. <https://tools.usps.com/go/EPLAction!input>, 2010.
- [Sey95] P.D. Seymour. Packing directed circuits fractionally. *Combinatorica*, 15:281–288, 1995.
- [Sit06] R. Sitters. The minimum latency problem is NP-hard for weighted trees. In *Integer Programming and Combinatorial Optimization*, pages 230–239, 2006.
- [SK04] Chaitanya Swamy and Amit Kumar. Primal–dual algorithms for connected facility location problems. *Algorithmica*, 40(4):245–269, 2004.
- [SPS03] A. C. Schultz, L. E. Parker, and F. E. Schneider. *Multi-Robot Systems: From Swarms to Intelligent Automata*. Springer, 2003.
- [SS06] S. Safra and O. Schwartz. On the complexity of approximating TSP with neighborhoods and related problems. *Computational Complexity*, 14(4):281–307, 2006.
- [SW01] T. Strijk and A. Wolff. Labeling points with circles. *International Journal of Computational Geometry and Applications*, 11(2):181–195, 2001.
- [WAW93] T.G. Will and D.B. Adviser-West. Extremal results and algorithms for degree sequences of graphs. 1993.
- [XW10] Z. Xu and Q. Wen. Approximation hardness of min-max tree covers. *Operations Research Letters*, 38:169–173, 2010.
- [XX09] Z. Xu and L. Xu. Approximation algorithms for min-max path cover problems with service handling time. *International Symposium on Algorithms and Computation*, pages 383–392, 2009.