

The Traveling Firefighter Problem

Majid Farhadi*
Georgia Tech
farhadi@gatech.edu

Alejandro Toriello
Georgia Tech
atorielle@isye.gatech.edu

Prasad Tetali
Georgia Tech
tetali@math.gatech.edu

Abstract

We introduce the L_p Traveling Salesman Problem (L_p -TSP), given by an origin, a set of destinations, and underlying distances. The objective is to schedule a destination visit sequence for a traveler of unit speed to minimize the Minkowski p -norm of the resulting vector of visit/service times. For $p = \infty$ the problem becomes a path variant of the TSP, and for $p = 1$ it defines the Traveling Repairman Problem (TRP), both at the center of classical combinatorial optimization.

L_p -TSP can be formulated as a convex mixed-integer program and enables a smooth interpolation between path-TSP and TRP, corresponding to optimal routes from the perspective of a server versus the customers, respectively. The parameter p can affect fairness or efficiency of the solution: The case $p = 2$, which we term the Traveling Firefighter Problem (TFP), models the scenario when the cost/damage due to a delay in service is quadratic in time.

We provide a polynomial-time reduction of L_p -TSP (losing a factor of $1 + \varepsilon$ in performance) to the segmented-TSP, a routing problem that defines a constant $O(1 + \varepsilon^{-2})$ number of deadlines by which given numbers of vertices should be visited. Subsequently we derive polynomial-time approximation schemes for L_p -TSP in the Euclidean metric and the tree metric (for which the problem is strongly NP-hard).

We also study the all-norm-TSP, in which the objective is to find a route that is (approximately) optimal with respect to the minimization of any norm of the visit times. We improve the approximation bound for this problem to 8, down from 16, and further prove an impossibility for an approximation factor better than 1.78, even in line metrics. Finally, we show the performance of our algorithm can be optimized for a specific norm, particularly yielding a 5.65-approximation for the TFP on general metrics.

We leave open several interesting directions to further develop this line of research.

1 Introduction

The L_p -TSP is a routing problem seeking to minimize the L_p norm of the vector of visit/service times to a set of customer locations. It generalizes and interpolates between two well-studied problems, the path variant of the TSP and the TRP, also known as the Minimum Latency Problem, which are the two extreme cases in which one minimizes either the largest or the sum of service times. By assuming the server's speed is constant, we use time and distance interchangeably in the remainder of the paper.

As one motivating example, the L_p -TSP for $p = 2$, which we call the Traveling Firefighter Problem (TFP), abstracts a macro-scale optimal strategy for dispatching a firefighter to minimize the total damage due to fires at various locations. Ride-sharing is another use-case of our problem. For example, devising the return route of a school bus, should we optimize the fuel consumption (i.e. the driver's time en route) or the average student waiting time? Is it fairer to further penalize larger waiting times, e.g. minimizing the sum of squares/cubes of the waiting times? Before formulating the problems under study we introduce some notation.

Notation. $[z]$ denotes the set $\{1, \dots, z\}$ for any positive integer z . $\tilde{O}(\cdot)$ is equivalent to $O(\cdot)$, treating $\varepsilon > 0$ as a constant. A metric over a set of nodes V is a distance function $d : V \times V \rightarrow \mathbb{R}_{\geq 0}$ that satisfies symmetry, $d(x, y) = d(y, x) \forall x, y \in V$, identity, $d(x, x) = 0$, and the triangle inequality, $d(x, y) \leq d(x, z) + d(z, y) \forall x, y, z \in V$.

The inputs to the problem are the set of vertices V , including both the destinations and the server's starting location, $s \in V$, and the underlying metric $d(\cdot, \cdot)$ over V , corresponding to distances (or times) between vertex pairs.

A feasible solution or route is a permutation σ over V that starts at the origin; σ_i denotes the i^{th} vertex to be visited, so we always have $\sigma_1 = s$. \mathcal{F} denotes the set of all feasible solutions. The i^{th} smallest visit time,

*Supported by aco.gatech.edu, triad.gatech.edu. Part of this work was conducted when authors visited Simons Institute for the Theory of Computing.

due to a solution $\sigma \in \mathcal{F}$, is denoted T_i^σ , i.e.

$$T_i^\sigma = \begin{cases} 0 & i = 1 \\ \sum_{j=2}^i d(\sigma_{j-1}, \sigma_j) & i \in \{2, \dots, n\} \end{cases}$$

The visit time for vertex v is denoted ℓ_v^σ . Similarly, $\ell_s^\sigma = 0$ and $\ell_v^\sigma = \sum_{i=2}^{\sigma_i=v} d(\sigma_{i-1}, \sigma_i) \quad \forall v \neq s$.

DEFINITION 1. (L_p -TSP) *The input of the optimization problem L_p -TSP is a set of destinations V , a starting vertex $s \in V$, and a metric $d : V \times V \rightarrow \mathbb{R}_{\geq 0}$. The objective is to find a feasible route, $\sigma \in \mathcal{F}$, starting at s and visiting all $v \in V$, that minimizes the Minkowski p -norm of the visit times, i.e. $\min_{\sigma \in \mathcal{F}} \|\ell^\sigma\|_p$, where*

$$\|\ell^\sigma\|_p := \left(\sum_{v \in V} |\ell_v^\sigma|^p \right)^{\frac{1}{p}}.$$

When the problem/objective is clear from context, we denote an optimal route as OPT and the answer of our algorithm as ALG.

A Better Objective. One can verify that the objective (norm) affects various aspects of the routing problem, such as efficiency. L_p -TSP enables a smooth transition between two extreme objectives. For larger p 's the objective is strongly affected by dominating (larger) entries of the delay vector and minimizing $\|\ell\|_p$ is more to the benefit of the server. In contrast, smaller p 's provide a further aggregated measure of the amount of time that the customers have waited. This trade-off can also be interpreted from a fairness perspective, as increasing p discourages the longest waiting time from becoming too large.

Firefighter Example. We further elaborate on why $p \notin \{1, \infty\}$ can be a useful objective by considering the routing of a firefighter.¹ Consider a set of wildfires in dispersed locations, and suppose a skilled firefighter extinguishes any fire the second they arrive at a location; the firefighter must choose the order in which to visit and extinguish the fires. One possible strategy is to choose a sequence in order to finish extinguishing all fires as soon as possible, which corresponds to solving the L_∞ -TSP (the path-TSP). However, this may not be the best solution for the firefighter: L_∞ -TSP minimizes the latest visit time for all fires, while the cost could be affected by all visit times; thus an aggregated measure could be a better objective.

¹Over the past decade, and for the first time on record, the annual number of acres burned in the United States exceeded 10 million; this occurred twice [Hoo18]. In July 2019, a record 2.4 million acres of the Amazon rainforest were torched [Bor19]. The optimal allocation, scheduling and routing of firefighting resources may help address this global challenge.




Figure 1: Different norms cause different routes.

For example, consider the following simple dynamics for the spread of wildfires over uniform territory, ignoring possible differences such as vegetation, weather and wind. After every second, a flammable point of territory is ignited if it is within a unit distance from the burning flames. Under these dynamics, the area of land scorched by the fire is a quadratic function of the elapsed time. Therefore, the damage due to the delay on the i^{th} visit can be better represented by $\ell_{\sigma_i}^2$ and minimizing $\sum_v \ell_v^2 = \|\ell\|_2^2$, or equivalently $\|\ell\|_2$, is a better objective for this scenario compared to other norms, particularly $\|\ell\|_1$ or $\|\ell\|_\infty$. Motivated by this example, we term L_2 -TSP as the Traveling Firefighter Problem (TFP).

In an applied setting, this approach may require some refinement but the basic idea still applies. Land and weather asymmetries can be modeled by a multiplicity of vertices: If a fire spreads twice as fast in area, we can represent it by two vertices overlapping in the metric. One could also generalize the objective to a weighted sum of the squared delays and/or discretize large fires into smaller ones. Moreover, the time required to extinguish a fire can be accounted for by adding a new edge, hanging from the original destination at a distance proportional to the time required to contain that fire, and moving the destination to the other endpoint of the new edge.

The Argument versus The Objective. The set of feasible routes for all L_p -TSP problems is the same, while the objectives are different. For any $p \neq q$, there exist instances where the two optimal routes are different. As a simple example, consider four vertices **S**, **A**, **B**, and **C** over a line metric at locations $0, -1-\varepsilon, +1$, and $+2$ respectively. Starting at **S**, the route **SABC** is optimal for L_2 -TSP with corresponding objective of $\|(0, 1+\varepsilon, 3+2\varepsilon, 4+2\varepsilon)\|_2 \simeq \sqrt{26}$, in contrast to the optimal route for L_1 -TSP that is **SBCA** with the objective $\|(0, 1, 2, 5+\varepsilon)\|_1 = 8+\varepsilon$. Figure 1 depicts this example.




Figure 2: Optimal path-TSP versus TFP routes for 10 U.S. locations; the areas of the discs at visit locations are proportional to the expected (quadratic) damage at each location, due to the delay along the route. Here, taking the TFP route (instead of the TSP) reduces total damage by 5%.

Consider the effect of a “wrong” norm on the optimal route. Figure 2 depicts an example considering 10 locations with simultaneous wildfires in the United States.² Dispatching a firefighter from Atlanta, and traveling 100 times faster than the spread of fire, TFP would look quite different than L_∞ -TSP. In particular, the TFP reduces the total damage by 5% compared to the path-TSP solution.

As another example, consider destinations on the x axis, with many fires located at $+1$ and a single fire at $-1 + \varepsilon$. Starting at $x = 0$, the optimal L_∞ -TSP route moves left first and then right, resulting in a solution roughly three times as expensive as the optimal TFP route, in terms of the L_2 objective. In fact, the relative performance of an L_∞ -TSP route for the L_2 objective can be *unbounded*. For instance, consider the Euclidean metric over the plane with $n - 2$ fires located at complex coded locations

$$\{e^{2\pi \frac{k}{n} i} : k \in \{1, \dots, n - 2\}\},$$

and m distinct fires located at $e^{2\pi \frac{n-1-\varepsilon}{n} i}$. Starting at $(1, 0) = e^{0i}$ and moving at unit speed, for $n \rightarrow \infty$ and $m/n \rightarrow \infty$ the damage (squared delay) due to L_∞ -TSP route (by walking in the wrong direction around the circle) converges to $4\pi^2$, while for the optimal TFP solution it goes to zero.

All-norm-TSP. We observed that TSP may not be a good solution for other L_p objectives. One may consider whether a different L_p -TSP is hopefully good for all norms. In this line, a natural question is whether there exists a single route that is approximately optimal concerning the minimization of any norm of the visit

times, and whether we can efficiently find one. This can be viewed also as an online problem where the adversary chooses the norm, e.g. L_p , and the objective is to provide a competitive solution with respect to the optimal route.

DEFINITION 2. (ALL-NORM-TSP) Given $s, V \ni s, d : V \times V \rightarrow \mathbb{R}_{\geq 0}$ as before, the objective is to choose a route that minimizes the maximum possible ratio between a symmetric norm of the visit time vector of the output route σ and the optimal route for that norm,

$$\min_{\sigma \in \mathcal{F}} \sup_{\|\cdot\|} \frac{\|\ell_\sigma\|}{\min_{\sigma' \in \mathcal{F}} \|\ell_{\sigma'}\|}.$$

Golovin et al. [GGKT08] introduced this problem as the *all-norm-TSP* and gave an algorithm that outputs a route that is a 16-approximation with respect to any norm of the visit time vector. The concept of all-norm minimization has been of interest in many applications, e.g., for routing, load balancing [KRT99], and machine scheduling [AERW04, BP03].

1.1 Main Results & Proof Ideas. Optimizing the non-linear objective of L_p -TSP can be computationally challenging. The problem is already NP-hard even in the linear case, $p = 1$, on a tree metric [Sit02], i.e. when the metric is pairwise distances over a graph on V that forms a tree. In contrast, TSP on trees is solvable, in linear time.

Archer and Williamson [AW03] showed that a $(1 + \varepsilon)$ -approximate solution for TRP exists that is a concatenation of $O(\log n \cdot \varepsilon^{-1})$ TSP paths. This can be generalized as the following Lemma, which enables a quasipolynomial time approximation scheme for weighted trees.

²Wildfires observed during the first week of Dec 2019, according to satellite data [NAS].

LEMMA 1.1. ([AW03]) L_p -TSP can be $(1 + \varepsilon)$ -approximated by a concatenation of $O(\log n \cdot \varepsilon^{-1})$ TSP paths.

Proof. Without loss of generality we can assume integral and polynomially bounded input distances, $d(i, j) \in \{0\} \cup [O(n^2/\varepsilon)] = \{0\} \cup [\tilde{O}(n^2)]$, as an appropriate quantization by rounding only adds a multiplicative error of $1 + O(\varepsilon/n^2)$ to any edge and keeps any norm of a valid tour, within a factor $(1 \pm \varepsilon)$.

Now break the optimal route according to time spots $1, (1 + \varepsilon), \dots, (1 + \varepsilon)^\gamma$ where $\gamma = O(\log n \cdot \varepsilon^{-1})$ because $O(n^3 \cdot \varepsilon^{-1})$ is a bound on the length of the optimal route. Replacing each sub-route between two consecutive time spots, with the shortest path TSP over the same points does not increase any visit time beyond a factor $1 + \varepsilon$, and hence preserves any norm of the visit times by a factor of $1 + \varepsilon$. \square

The above approach can only lead to a pseudo-polynomial time approximation algorithm, even for trees. Reducing the problem to many shortest paths cannot lead to an efficient polynomial-time reduction, because a concatenation of $o(\log n)$ path-TSP routes cannot approximate L_1 -TSP within a constant factor [Sit14]. To further reduce the number of paths, we use the notion of segmented-TSP, introduced by Sitters [Sit14], to enable some dependence between consecutive deadlines. This problem requires a (sequence of monotonically non-decreasing) number of destinations to be visited by a number of deadlines, formulated as follows.

DEFINITION 3. (SEGMENTED-TSP) Given $V \ni s, d : V \times V \rightarrow \mathbb{R}_{\geq 0}$ as before, in addition to integer numbers $n_1 \leq n_2 \leq \dots \leq n_k \leq |V| = n$ and fractional numbers $t_1 \leq t_2 \leq \dots \leq t_k$ as inputs, the segmented-TSP problem is a decision problem to verify whether a route exists that visits at least n_i distinct vertices by time t_i for all $i \in [k]$, starting at s .

Approximation of segmented-TSP, i.e. a decision problem, can be defined as follows.

DEFINITION 4. An α -approximate solution to a segmented-TSP instance, must visit the first n_i vertices by the modified deadline $\alpha \cdot t_i, \forall i \in [k]$, if an answer to the original segmented-TSP exists.

We generalize a main result of Sitters [Sit14] that showed TRP can be reduced to (a polynomially many number of approximate) segmented-TSP problems with a constant number of deadlines.

THEOREM 1. [L_p -TSP \Rightarrow segmented-TSP] Let $\varepsilon > 0$ be a constant and \mathcal{A} be an α -approximation algorithm for

segmented-TSP for some $k = O(1 + \varepsilon^{-2})$ of our choice. There is a $(1 + \varepsilon) \cdot \alpha$ -approximation algorithm for L_p -TSP that calls \mathcal{A} (on the same network (V, s, d)) for a strongly polynomial number of times.

Theorem 1, proved in Section 2, along with a PTAS for segmented-TSP [Sit14] on tree metrics and Euclidean metrics imply the following results.

COROLLARY 1.1. There exist polynomial-time approximation schemes (PTAS) for L_p -TSP on weighted trees and Euclidean plane metrics.

Note that due to our reduction of L_p -TSP to segmented-TSP, any approximation results for segmented-TSP, on specific metrics, would follow for L_p -TSP, at the cost of an additional factor of $(1 + \varepsilon)$ to the approximation bound. Our results can be generalized to the case of multiple travelers, starting from arbitrary locations, as discussed in Section 5.

For general metrics, a PTAS is unlikely, as the problem becomes Max-SNP-hard. On the other hand, the constant factor approximability of L_p -TSP for general metrics is immediate due to the 16-approximation of the all-norm-TSP by Golovin et. al. [GGKT08]. In this line, we improve the approximation bound for all-norm-TSP by a factor of 2.

THEOREM 2. [all-norm-TSP] There is a polynomial-time algorithm to find a route that is 8-approximate with respect to the minimization of any symmetric norm of the visit times (including L_p -TSP, TFP, and TRP).

Theorem 2 is proved in Section 3. Our algorithm builds on the partial covering idea, presented as Algorithm 1, that was pioneered by Blum et al. [BCC⁺94] for TRP, and was developed through subsequent studies [GK98, CGRT03, GGKT08, BKL21].

Algorithm 1 Routing via Partial Covering

```

1: procedure GEOMETRIC-COVERING( $V, s, d$ )
2:   Algorithm Parameters:  $b \in (0, \infty), c \in (1, \infty)$ 
3:    $i \leftarrow 0$ 
4:   while there remains destinations to visit do
5:      $\triangleright$  Conducting sub-tours
6:      $C_i \leftarrow$  a maximal route of length  $\leq b \cdot c^i$ .
7:     Travel through  $C_i$  (and return to the origin)
8:      $i \leftarrow i + 1$ 
9:   return an ordering  $\sigma$  of  $V$  according to their
   (first) visit time through the above loop.

```

The parameters b and c strongly affect the performance of the algorithm. For all-norm-TSP, we choose

$b = \min_{i,j \in V} d(i,j)$ and $c = 2$, as in [GGKT08]. The key difference of our algorithm is in line 6. Instead of an approximation algorithm for k -TSP, i.e. a route of minimum length that visits k -vertices, we utilize a milder relaxation, which is a tree (instead of a route/stroll) rooted at s , including k vertices, and of total length not larger than an optimal k -TSP. Such a *good k -tree* can be found in polynomial time using the primal-dual method that solves a Lagrangian relaxation of k -TSP [CGRT03, ALW08, PS14].

On the other hand, we provide a first impossibility result for all-norm-TSP, notably beyond known in-approximability bounds for specific L_p -TSP problems. The following is proved in Section 3.2.

THEOREM 3. *There is no approximation algorithm for all-norm-TSP with multiplicative factor better than 1.78, independent of $P = NP$ or other complexity hypotheses.*

The above result reaffirms the need for approximation algorithms specifically designed for each norm. Along this line, we present a randomized 5.65 approximation algorithm for the Traveling Firefighter Problem.

THEOREM 4. *There is a randomized, polynomial-time 5.65-approximation algorithm for TFP on general metrics.*

Theorem 4 is proved in Section 4, for which we adapt the ideas by Chaudhuri et. al. [CGRT03] and optimize the parameter c in line 2 of Algorithm 1. Choosing $b \in [1, c]$ at random, with a distribution of uniform density for $\log(b)$, simplifies the analysis, while one can efficiently de-randomize the algorithm by quantization of b .

1.2 Literature Review Traveling Salesman Problem is a principal problem in computer science, combinatorial optimization, and operations research, and its first formulations date back as early as 19th century (c.f. [ABCC06]). Since the celebrated 3/2-approximation algorithm of Christofides-Serdyukov [Chr76, Ser78], for its tour-variant on general metrics, TSP has been extensively studied for half a century [Wol80, SW90, BP90, Goe95, CV00, GLS05, BC11, SWVZ12, HNR19]. Very recently, Karlin, Klein, and Oveis Gharan [KKG20] showed TSP can be approximated strictly better than 3/2, while the problem remains NP-hard to approximate within a factor of 123/122 [KLS15].

The common ground in numerous variants of TSP is that a set of vertices are to be visited in the *fastest possible* way, i.e., optimizing the time spent by the server/traveler. In contrast, the Traveling Repair-

man Problem (TRP), a.k.a. Minimum Latency Problem, the school bus driver problem [Wil94], hidden treasure [BCC⁺94, KPY96, ALMS00], and the deliveryman problem [Min89, FLM93, MDZL08], optimizes the route purely from the perspective of clients, i.e., the total waiting time to be visited, and is another extensively studied combinatorial optimization problem [ACP⁺86, PY93, BCC⁺94, GK98, CGRT03, AW03] with a state-of-the-art approximation factor of $\simeq 3.59$ for general metrics [CGRT03, PS14].

Containment of fires can be abstracted from various perspectives. Hartnell [Har95] modeled a constant speed spread of fires through edges of a graph, along which the firefighters also displace. Many objectives, such as minimization of the number of burned vertices, or the required time to contain the fire(s) are studied in this model and the problem is an active area of research. See [FM09, KLL14, ABZ18, ABK20, DFH21] and references therein.

Generalizing the objective to Minkowski norm of the solution has allowed interpolating other classical problems, e.g., L_p set cover problem [GGKT08, BBFT20] that further united the greedy algorithms for set-cover and the minimum-sum-set-cover [FLT04] problems. Set cover is better approximable for $p = 1$ than $p = \infty$, while TSP ($p = \infty$) is currently better approximated than TRP ($p = 1$). Nevertheless, this order is not expected to be reversed as TRP is intrinsically a harder problem. Moreover, in contrast to the concordance among L_p set cover problems, for which the same greedy algorithm gives best (possible) bounds for any $p \in [1, \infty)$, state-of-the-art algorithms for TSP and TRP are significantly different, and potentially far from the best possible. This is yet another motivation to study L_p -TSP, ultimately towards unified best algorithms for all underlying problems.

1.3 Paper Organization. In Section 2 we provide a reduction from L_p -TSP to segmented-TSP and subsequent approximation schemes for Euclidean and weighted-tree metrics. In Section 3, we present the 8-approximation for all-norm-TSP in general metrics, along with a first inapproximability bound. This will also be a preliminary for optimized algorithm for L_2 -TSP in Section 4, where we present a 5.65-approximation for the Traveling Firefighter Problem on general metrics. We discuss generalizations to multiple vehicle scenarios in Section 5 and conclude the paper with a set of interesting open problems to continue this line of research.

2 Reducing L_p -TSP to Segmented TSP

In this section, we provide our reduction from L_p -TSP to polynomially many instances of segmented-TSP. Let us restate Theorem 1.

THEOREM 1. [L_p -TSP \Rightarrow segmented-TSP] *Let $\varepsilon > 0$ be a constant and \mathcal{A} be an α -approximation algorithm for segmented-TSP for some $k = O(1 + \varepsilon^{-2})$ of our choice. There is a $(1 + \varepsilon) \cdot \alpha$ -approximation algorithm for L_p -TSP that calls \mathcal{A} (on the same network (V, s, d)) for a strongly polynomial number of times.*

In particular, a corollary of the above Theorem (and the following Lemma) is a $(1 + \varepsilon)$ approximation algorithm for any L_p -TSP on weighted-tree metrics - where the problem becomes strongly NP-hard even for $p = 1$ - as well as the Euclidean plane.

LEMMA 2.1. ([SIT14]) *Segmented TSP, for any constant number of segments M , can be solved in polynomial time for weighted trees, and $1 + \varepsilon$ approximated for unweighted Euclidean metric.*

COROLLARY 1.1. *There exist polynomial-time approximation schemes (PTAS) for L_p -TSP on weighted trees and Euclidean plane metrics.*

In the rest of this section, we prove Theorem 1, by providing a dynamic programming algorithm that approximates L_p -TSP using polynomially many calls to (approximate) segmented-TSP. More precisely, we show if there is an α -approximate solution for segmented-TSP, the dynamic program guarantees an approximation factor of at most $\alpha \cdot (1 + \varepsilon)$ for arbitrary constant $\varepsilon > 0$.

The algorithm is presented in a few steps, each imposing no more than $1 + O(\varepsilon)$ multiplicative error. To achieve exact $1 + \varepsilon$ precision, one may run the algorithm for a constant fraction of the target ε .

For some k as large as $O(1 + \varepsilon^{-2})$ we can ensure $c \stackrel{\text{def}}{=} (1 + \varepsilon)^k \geq 3$. Let OPT^{λ_i} denote the maximal prefix of OPT route for L_p -TSP of length at most

$$\lambda_i \stackrel{\text{def}}{=} (1 + \varepsilon)^{-j} \cdot c^i, \quad \forall i \geq 0,$$

where j is a fixed random number, uniformly distributed over $\{0, \dots, k - 1\}$.

Let OPT' be a tour made of sub-tours consisting of traversing OPT^{λ_i} and returning to the origin and waiting until time $3\lambda_i$ before starting the next sub-tour. To confirm the above is feasible, we need to show sub-tour $i + 1$, being allowed to begin at $3\lambda_i$, does not leave before the return of previous sub-tour, i.e.,

$$3\lambda_{i-1} + 2\|T^{\text{OPT}^{\lambda_i}}\|_\infty \leq \lambda_i + 2\lambda_i = 3\lambda_i$$

which is immediate having $\lambda_i = c\lambda_{i-1} \geq 3\lambda_{i-1}$.

LEMMA 2.2. *The modified tour OPT' is approximately optimal in expectation, i.e., for some $k \in O(1 + \varepsilon^{-2})$*

$$\mathbb{E}_j \left[\|T^{\text{OPT}'}\|_p^p \right] \leq (1 + \varepsilon) \|T^{\text{OPT}}\|_p^p.$$

Proof. All vertices are visited (for the first time) in the same order, by OPT and OPT' . Let the d^{th} service time by the optimal solution be

$$T_d^{\text{OPT}} \in ((1 + \varepsilon)^\delta, (1 + \varepsilon)^{\delta+1}]$$

for some integer $\delta \geq 0$.

If this vertex is visited in the i^{th} sub-tour of OPT' , we can write

$$T_d^{\text{OPT}'} = T_d^{\text{OPT}} + 3\lambda_{i-1}.$$

We can bound this additional delay by

$$T_d^{\text{OPT}'} - T_d^{\text{OPT}} \leq 3(1 + \varepsilon)^{\delta-j'}$$

where j' has the same distribution as j .

We can prove the desired by bounding the per-vertex ratio by

$$\frac{\mathbb{E}_{j'} \left[(T_d^{\text{OPT}'})^p \right]}{(T_d^{\text{OPT}})^p} \leq (1 + \varepsilon)$$

because $\frac{\sum_i a_i}{\sum_i b_i} \leq \max_i \frac{a_i}{b_i}$ where a_1, \dots and b_1, \dots are positive real numbers.

Considering $p \geq 1$ and $T_d^{\text{OPT}'} = T_d^{\text{OPT}} + 3\lambda_{i-1}$, the left hand side of the target ratio is maximized for $T_d^{\text{OPT}} = (1 + \varepsilon)^\delta$, so it suffices to prove

$$\frac{\mathbb{E}_{j'} \left[((1 + \varepsilon)^\delta + 3\lambda_{i-1})^p \right]}{((1 + \varepsilon)^\delta)^p} \leq (1 + \varepsilon).$$

We will have

$$\begin{aligned} & \frac{\mathbb{E}_{j'} \left[((1 + \varepsilon)^\delta + 3\lambda_{i-1})^p \right]}{((1 + \varepsilon)^\delta)^p} \\ & \leq \frac{\mathbb{E}_{j'} \left[((1 + \varepsilon)^\delta + 3(1 + \varepsilon)^{\delta-j'})^p \right]}{((1 + \varepsilon)^\delta)^p} \\ & = \mathbb{E}_{j'} \left[(1 + 3(1 + \varepsilon)^{-j'})^p \right] \\ & = 1 + \mathbb{E}_{j'} \left[(1 + 3(1 + \varepsilon)^{-j'})^p - 1^p \right] \\ & \leq 1 + \frac{1}{k} \sum_{j'=0}^{k-1} (3p)^p (1 + \varepsilon)^{-j'} \\ & \leq 1 + \frac{(3p)^p}{k} \cdot \frac{1}{1 - (1 + \varepsilon)^{-1}} \\ & = 1 + \frac{(3p)^p}{k} \cdot \frac{1 + \varepsilon}{\varepsilon}. \end{aligned}$$

To get the desired (from the last inequality) it suffices to assume

$$k \geq \frac{(3p)^p(1+\varepsilon)}{\varepsilon^2}.$$

□

We can now complete the proof of Theorem 1, similar to the main idea of Sitters [Sit14], presented as follows.

Lemma 2.2 implies for some $j \in [k]$, where $k = O(1 + \varepsilon^{-2})$, there exists a near optimal routing, OPT' , that for each $i \in [\tilde{O}(n^2)]$, visits *new vertices* only during $[3\lambda_{i-1}, \lambda_i]$, and returns to the origin and remains there until

$$3\lambda_i = \frac{3}{(1+\varepsilon)^j} \cdot (1+\varepsilon)^{ki}.$$

We can search for such a path by reconstructing OPT^{λ_i} for all i and upper bounding the consequent $\|T^{\text{OPT}'}\|_p^p$ using dynamic programming.

Define $D[i][d]$ as (an upper bound on) the contribution of visit times of vertices that are visited by OPT^{λ_i} , to $\|T^{\text{OPT}'}\|_p^p$, further assuming the number of these vertices is d .

We can compute $D[i][d]$ considering $O(n^k)$ cases of (m_1, m_2, \dots, m_k) where m_r denotes the number of vertices that are visited by OPT^{λ_i} during

$$(3\lambda_{i-1} + \lambda_i \cdot (1+\varepsilon)^{r-k-1}, 3\lambda_{i-1} + \lambda_i \cdot (1+\varepsilon)^{r-k}].$$

Note that it is necessary to have $\sum_r m_r \leq d$ and let $d' = d - \sum_r m_r$ be the number of vertices visited by $\text{OPT}^{\lambda_{i-1}}$. We can write

$$D[i][d] = \min_{m_1, \dots, m_k} D[i-1][d'] +$$

$$\text{Seg-TSP}_i(d', m_{[r]}) \cdot \sum_r m_r \cdot (3\lambda_{i-1} + \lambda_i \cdot (1+\varepsilon)^{r-k})^p,$$

where $\text{Seg-TSP}_i(d', m_{[r]})$ has value 1 if segmented-TSP is feasible for visiting at least

$$d', d' + m_1, \dots, d' + m_1 + \dots + m_r$$

vertices by deadlines

$$\lambda_{i-1}, 3\lambda_{i-1} + \lambda_i \cdot (1+\varepsilon)^{-k}, \dots, 3\lambda_{i-1} + \lambda_i$$

is feasible, and otherwise has value ∞ . Note that we have an α approximate solver for Segmented-TSP, though for convenience we can alternatively assume the traveller goes at the speed of α instead of 1, to get a $1+\varepsilon$ approximate solution to $\|T^{\text{OPT}}\|_p$ by $(D[\tilde{O}(n^2)][n])^{1/p}$. In the end, moving at unit speed (instead of α) at every

stage can increase (any) norm of the delay vector $\|T\|$ by a factor α so we have an $\alpha \cdot (1+\varepsilon)$ approximation, that was promised by Theorem 1.

Finally it is worth to mention that the route (instead of the value) can be reconstructed using update (parent) information of $D[\cdot][\cdot]$ and a constructive approximate solver for Seg-TSP $[\cdot]$ and we can short-cut potential re-visits of vertices to have a valid Hamiltonian route.

3 All-norm TSP

Since the introduction of a first constant approximation for TRP by Blum et al. [BCC+94], partial covering through applying a geometric series of limits on the length of the sub-tours has been a core in the design of routing algorithms. In this section, we improve the 16-approximate/competitive solution for all-norm TSP by a factor of 2. We further provide a first lower bound for this problem.

3.1 Approximation for General Metrics. The idea is to iteratively cover more and more vertices by sub-tours of exponentially (geometrically) increasing length while trying to maximize the total number of vertices that are visited (not necessarily for the first time) in each iteration. Our algorithm uses the following milder relaxation of k -TSP, called a *good k -tree*, in place of line 6 in Algorithm 1.

DEFINITION 5. *A good k -tree is a tree of size k , including s , and with a total edge-weight of no more than that of the optimal k -TSP (starting from s).*

LEMMA 3.1. ([CGRT03]) *A good k -tree can be found in polynomial time.*

Chaudhuri et. al. [CGRT03] proved the above using a primal-dual approach [Gar96, AK00] that allows finding a feasible solution to the primal (integer) linear program of the k -tree problem paired with a feasible dual solution to k -TSP, that by weak duality has no less of a cost.

We are now ready to prove Theorem 2.

THEOREM 2. *[all-norm-TSP] There is a polynomial-time algorithm to find a route that is 8-approximate with respect to the minimization of any symmetric norm of the visit times (including L_p -TSP, TFP, and TRP).*

Proof. WLOG assume the nearest neighbor to s is at distance 1. For $k = 1, 2, \dots, n$ find a good- k -tree. Among these, name the largest tree (with respect to number of vertices) of total length at most 2^i as G_i for $i = 0, 1, 2, \dots$

Let C_i be a (randomized) depth-first traversal of G_i and let C be the concatenation of C_i 's for $i = 0, \dots$. The final tour ALG will visit vertices in the order that they appear in C , which does not increase the first-visit time for any vertex, due to triangle inequality of the metric, while short-cutting vertices that are being re-visited.

Let

$$T_k^{\text{OPT}} \in [2^i, 2^{i+1}).$$

This shows the shortest (length) k -path in G is no longer than 2^{i+1} . So our good- k -tree is no longer than 2^{i+1} , hence C_{i+1} has at least k distinct vertices, allowing us to upper bound our k^{th} visit time by

$$T_k^{\text{ALG}} \leq \sum_{j=0}^{i+1} |C_j| \leq \sum_{j=0}^{i+1} 2 \times 2^j < 2^{i+3}.$$

Together with $T_k^{\text{OPT}} \geq 2^i$ and the above inequality we have

$$T_k^{\text{ALG}} \leq 8 \times T_k^{\text{OPT}}.$$

We showed $T_k^{\text{ALG}} \leq 8 \cdot T_k^{\text{OPT}} \quad \forall i \in [n]$, i.e., T^{OPT} is 8-submajorized by T^{ALG} in terminology of [GGKT08, HLP88]. The result is that

$$\|T^{\text{ALG}}\| \leq 8 \cdot \|T^{\text{OPT}}\|$$

w.r.t. any norm $\|\cdot\|$. \square

One can verify the above algorithm performs asymptotically 3 times worse than the optimal TRP for the example with service points at $\{2^i : i \in \mathbb{N}\}$ and starting at $x = 0$.

3.2 Inapproximability. We conclude this section by providing a lower bound for all-norm TSP. We show even for line metrics, an α -approximate all-norm TSP cannot be guaranteed in general, for $\alpha < 1.78$.

THEOREM 3. *There is no approximation algorithm for all-norm-TSP with multiplicative factor better than 1.78, independent of $P = NP$ or other complexity hypotheses.*

Proof. Let the metric correspond to absolute differences over the real line, and start the walk from the origin $x = 0$. Our example has a similar structure as follows. There is a single destination at $x = -1$, in addition to n destinations at $x \in \{b^i - 1 : i \in [n]\}$. Choosing $b = 1 + \varepsilon$ such that $b^n \gg 1$, the optimal all-norm TSP visits the longer branch first, and is optimal with respect to TRP, while for TSP (which takes the shorter branch first) it achieves an approximation ratio of

$$\frac{2b^n + 1}{b^n + 2}.$$

For $b = 1.001$ and $n = 2256$ the above setup provides a bound of 1.71.

We further developed a similar, yet numerically verified example depicted as Figure 3, for which no 1.78-approximate all-norm-TSP route exists. For reproducibility we include this instance in the Appendix. \square

The above example, is yet another motivation to study and optimize routing algorithms specific to the appropriate objective/norm, as one solution *cannot* be good for all.

4 Traveling firefighter in general metrics

In this section we build upon our geometric partial-covering algorithm with good k -trees to improve approximation bound for a specific norm, i.e., the Traveling Firefighter Problem.

We achieve the improved approximation bound by randomization (of parameter b) and optimization of our analysis w.r.t. approximation guarantee for a specific norm. Our approach can provide approximation guarantees (better than 8) for other L_p -TSP problems. We present main ideas in the rest of the section by proving the following result.

THEOREM 4.1. *Traveling Firefighter Problem for general metrics can be 5.641-approximated in polynomial time.*

Proof. We present a randomized approximation algorithm for general metrics, that can be efficiently derandomized.

Among the set of good k -trees, pre-computed for all k , let G_i be the largest one of total length at most $b \cdot c^i$, and let C_i be a depth first traversal of that. Parameter $c > 1$ is a constant, to be optimized for performance guarantees, and

$$[1, c] \ni b = c^U$$

where U is a random variable distributed uniformly over the interval $[0, 1]$. Finally, we reverse each C_i with probability half, and concatenate C_i 's (and shortcut repeated visits) to achieve the output ordering ALG.

Let the latency of the k^{th} vertex visited by the optimal route be

$$T_k^{\text{OPT}} = ac^i,$$

for some $a \in [1, c]$ and integer $i \geq 0$.

It is easy to see that our $(i + \mathbb{1}[a \geq b])^{\text{th}}$ sub-tour contains at least k vertices, hence, we can bound

$$T_k^{\text{ALG}} \leq X_k + \sum_{j=0}^{i-1[\frac{a}{b}]} 2bc^j,$$



Figure 3: An example with no better than 1.78 all-norm TSP

Where X_k is zero if our tour visits its k^{th} vertex before the $(i-1[a < b])^{\text{th}}$ sub-tour, and

$$X_k \in [0, 2bc^{i+1[a \geq b]}],$$

depending on its location in the sub-tour.

Due to convexity of the quadratic function we can bound the expected damage, with

$$\begin{aligned} & \mathbb{E} [(T_k^{\text{ALG}})^2] \\ & \leq \frac{1}{2} \left(\sum_{j=0}^{i-1[a < b]} 2bc^j \right)^2 \\ & \quad + \frac{1}{2} \left(2bc^{i+1-1[a < b]} + \sum_{j=0}^{i-1[a < b]} 2bc^j \right)^2 \\ & = \frac{1}{2} \left(2b \frac{c^{i+1[a \geq b]} - 1}{c-1} \right)^2 \\ & \quad + \frac{1}{2} \left(2b \frac{c^{i+1+1[a \geq b]} - 1}{c-1} \right)^2 \\ & \leq \frac{1}{2} \left(\frac{2bc}{c-1} \right)^2 c^{2(i+1[a \geq b])} \\ & = c^{2i} \cdot \frac{2c^2}{(c-1)^2} \left(b^2 c^{2 \cdot 1[a \geq b]} \right). \end{aligned}$$

Bounding the expected damage at the i^{th} service, considering random variable b we have

$$\begin{aligned} & \mathbb{E} [(T_k^{\text{ALG}})^2] \\ & \leq \mathbb{E} \left[c^{2i} \cdot \frac{2c^2}{(c-1)^2} \left(b^2 c^{2[a \geq b]} \right) \right] \\ & = c^{2i} \cdot \frac{2c^2}{(c-1)^2} \left(c^2 \int_0^{\log_c d} c^{2U} dU + \int_{\log_c d}^1 c^{2U} dU \right) \\ & = c^{2i} \frac{2c^2}{(c-1)^2} \left(\frac{c^2(a^2-1)}{\ln c} + \frac{c^2-a^2}{\ln c} \right) \\ & = (ac^i)^2 \left(\frac{2c^2 \cdot (c^2-1)}{(c-1)^2 \ln c} \right) \\ & = (T_k^O)^2 \left(\frac{2c^2(c+1)}{(c-1) \ln c} \right). \end{aligned}$$

We can now choose c in order to minimize the multiplicative bound

$$\frac{c+1}{c-1} \cdot 2c^2 / \ln c \leq 31.82$$

that can be achieved for $c \simeq 2.54$. With this we have

$$\mathbb{E} [\|T^{\text{ALG}}\|_2^2] \leq 31.82 \|T^{\text{OPT}}\|_2^2.$$

We provided a randomized algorithm, that along law of large numbers, can be applied in practice, by repeating the algorithm and reporting the best performing route. On the other hand it can be simply de-randomized by exploring all values for a dense enough quantization of b . In the first case we will have a $\sqrt{31.82} \simeq 5.641$ approximate L_2 -TSP with high probability, and in latter scenario we will have deterministic result with negligible additional approximation error. \square

5 Generalizations to Multiple Vehicles

In many applications we have multiple vehicles, potentially dispatching from different hubs [KP20].

Our reduction of L_p -TSP to Segmented-TSP can be generalized to multiple vehicles, with arbitrary start locations. Similar to the approach in Section 2 we can convert any optimal multi-vehicle solution to repetition of $O(\varepsilon^{-2})$ of prefix routes for each vehicle, all synchronized with a single $j \in \{0, \dots, k-1\}$, picked uniformly at random. Lemma 2.2 can be subsequently adapted to allow limiting the search space to solutions that have all vehicles at starting locations simultaneously at all $3\lambda_i$, with negligible degrade of the optima.

Finally, adapting the dynamic programming, we can guarantee a multiplicative $O(\varepsilon)$ loss given an (approximate) solver for multi-vehicle segmented TSP with constant $O(\varepsilon^{-2})$ deadlines, and the corresponding total number of destinations to be visited (by at least one vehicle) until up to each. This is indeed fruitful as results on segmented-TSP also generalize to multi-vehicle variant, e.g., in tree-distance or Euclidean metric.

Our algorithms can be similarly adapted in the case where release dates are added for the destinations.

Generalizing results in Sections 3-4 to multi-vehicle variants of the problems is also possible. For this purpose, the k -stroll subproblem, for which we used the mild relaxation of good k -tree, can be generalized to bottleneck-stroll of Post and Swamy [PS14]. This will be at the cost of further degrades to the approximation constants and proposes interesting open problems for study. For general metrics, current best approximation

guarantees for multi vehicle {single, multi} depot L_1 -TSP is {7.183, 8.497} [PS14].

Discussion and Open Problems

We studied combinatorial optimization problems whose objectives can be more appropriate, efficient, fair, and adjustable, depending on enormous applications of optimal routing/scheduling.

For TSP and TRP, the analyses of approximation algorithms as well as complexity results heavily rely on the linearity of the objective function. Hence, TFP and more generally L_p -TSP pose further challenging problems and require new techniques to be developed.

We developed two approaches, towards high precision and scalable approximation of the answer.

First, we provided a high precision polynomial time reduction of L_p -TSP to segmented-TSP with only a constant number of deadlines for visiting the required number of destinations. Our reduction enables approximation schemes for L_p -TSP on Euclidean as well as weighted tree metrics; this is yet another motivation to further study the segmented-TSP problem.

The other approach relied mainly on the fact that the objective is a norm of the delay vector. In this line, we developed an algorithm for all-norm-TSP, on general metrics, of approximation factor 8. We also provided a first inapproximability result for all-norm-TSP.

Last but not least, we showed how the performance of the latter algorithm can be optimized for a specific norm, particularly approximating Traveling Firefighter Problem within a factor 5.65.

In the end, in addition to further improving the approximation bounds for the problems under study, we mention but a few of many interesting open problems and potential research directions regarding L_p -TSP and all-norm-TSP.

- L_1 -TSP, i.e., TRP is harder than L_∞ -TSP, at least on trees. For what p is the L_p -TSP problem the hardest?
- While TRP is strongly NP-hard on weighted trees, its complexity is unknown for caterpillars [Sit02]. Similarly TFP and L_p -TSP seem challenging even on such fundamental examples, that is yet to be resolved.
- While we theoretically claimed $p = 2$ to be ideal for the Traveling Firefighter Problem, this assumption should be given further justification / investigation in practice.

- Further applications of L_p -TSP can be inspected, e.g. in optimal containment of spread of pandemics [Har04, TCM20].
- We observe $c = 2$ to be optimal for the analysis of all-norm-TSP algorithm. For TRP, the current best result is due to a base $c \approx 3.59$ for the geometric series, while $c \approx 2.54$ is better for TFP, as we discussed. In this vein, one can inspect the best base for the geometric series used by the partial covering algorithm, depending on the objective, which naturally seems to be non-increasing on p .
- Stronger impossibility results for all-norm-TSP and even larger hardness of approximation bounds for this problem seem plausible.

Acknowledgements

We thank anonymous reviewers for valuable comments.

References

- [ABCC06] David L Applegate, Robert E Bixby, Vasek Chvatal, and William J Cook, *The traveling salesman problem: a computational study*, Princeton university press, 2006.
- [ABK20] Gideon Amir, Rangel Baldasso, and Gady Kozma, *The firefighter problem on polynomial and intermediate growth groups*, Discrete Mathematics **343** (2020), no. 11, 112077.
- [ABZ18] David Adjiashvili, Andrea Baggio, and Rico Zenklusen, *Firefighting on trees beyond integrality gaps*, ACM Transactions on Algorithms (TALG) **15** (2018), no. 2, 1–33.
- [ACP⁺86] Foto Afrati, Stavros Cosmadakis, Christos H Papadimitriou, George Papageorgiou, and Nadia Papakostantinou, *The complexity of the travelling repairman problem*, RAIRO-Theoretical Informatics and Applications **20** (1986), no. 1, 79–87.
- [AERW04] Yossi Azar, Leah Epstein, Yossi Richter, and Gerhard J Woeginger, *All-norm approximation algorithms*, Journal of Algorithms **52** (2004), no. 2, 120–133.
- [AK00] Sanjeev Arora and George Karakostas, *A 2+ approximation algorithm for the k-MST problem*, Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms, 2000, pp. 754–759.
- [ALMS00] Giorgio Ausiello, Stefano Leonardi, and Alberto Marchetti-Spaccamela, *On salesmen, repairmen, spiders, and other traveling agents*, Italian Conference on Algorithms and Complexity, Springer, 2000, pp. 1–16.
- [ALW08] Aaron Archer, Asaf Levin, and David P Williamson, *A faster, better approximation algorithm for the minimum latency problem*, SIAM Journal on Computing **37** (2008), no. 5, 1472–1498.

- [AW03] Aaron Archer and David P Williamson, *Faster approximation algorithms for the minimum latency problem*, Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, 2003, pp. 88–96.
- [BBFT20] Nikhil Bansal, Jatin Batra, Majid Farhadi, and Prasad Tetali, *Improved approximations for min sum vertex cover and generalized min sum set cover*, arXiv preprint arXiv:2007.09172 (2020).
- [BC11] Sylvia Boyd and Robert Carr, *Finding low cost TSP and 2-matching solutions using certain half-integer subtour vertices*, Discrete Optimization **8** (2011), no. 4, 525–539.
- [BCC⁺94] Avrim Blum, Prasad Chalasani, Don Coppersmith, Bill Pulleyblank, Prabhakar Raghavan, and Madhu Sudan, *On the minimum latency problem*, Proceedings of the twenty-sixth annual ACM symposium on Theory of computing, pages 163–171. ACM (1994).
- [BKL21] Marcin Bienkowski, Artur Kraska, and Hsiang-Hsuan Liu, *Traveling repairperson, unrelated machines, and other stories about average completion times*, arXiv preprint arXiv:2102.06904 (2021).
- [Bor19] Alejandra Borunda, *See how much of the Amazon forest is burning, how it compares to other years*, National Geographic (2019).
- [BP90] Sylvia C Boyd and William R Pulleyblank, *Optimizing over the subtour polytope of the travelling salesman problem*, Mathematical programming **49** (1990), no. 1-3, 163–187.
- [BP03] Nikhil Bansal and Kirk Pruhs, *Server scheduling in the lp norm: a rising tide lifts all boat*, Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, 2003, pp. 242–250.
- [CGRT03] Kamalika Chaudhuri, Brighten Godfrey, Satish Rao, and Kunal Talwar, *Paths, trees, and minimum latency tours*, 44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings., IEEE, 2003, pp. 36–45.
- [Chr76] Nicos Christofides, *Worst-case analysis of a new heuristic for the Travelling Salesman Problem*, Tech. Report RR-388, February 1976.
- [CV00] Robert Carr and Santosh Vempala, *Towards a 4/3 approximation for the asymmetric traveling salesman problem*, 116–125.
- [DFH21] Arye Deutch, Ohad Noy Feldheim, and Rani Hod, *Multi-layered planar firefighting*, arXiv preprint arXiv:2105.03759 (2021).
- [FLM93] Matteo Fischetti, Gilbert Laporte, and Silvano Martello, *The delivery man problem and cumulative matroids*, Operations Research **41** (1993), no. 6, 1055–1064.
- [FLT04] Uriel Feige, László Lovász, and Prasad Tetali, *Approximating min sum set cover*, Algorithmica **40** (2004), no. 4, 219–234.
- [FM09] Stephen Finbow and Gary MacGillivray, *The firefighter problem: a survey of results, directions and questions.*, Australas. J Comb. **43** (2009), 57–78.
- [Gar96] Naveen Garg, *A 3-approximation for the minimum tree spanning k vertices*, Proceedings of 37th Conference on Foundations of Computer Science, IEEE, 1996, pp. 302–309.
- [GKKT08] Daniel Golovin, Anupam Gupta, Amit Kumar, and Kanat Tangwongsan, *All-norms and all-Lp-norms approximation algorithms*, IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2008.
- [GK98] Michel Goemans and Jon Kleinberg, *An improved approximation ratio for the minimum latency problem*, Mathematical Programming **82** (1998), no. 1-2, 111–124.
- [GLS05] David Gamarnik, Moshe Lewenstein, and Maxim Sviridenko, *An improved upper bound for the TSP in cubic 3-edge-connected graphs*, Operations Research Letters **33** (2005), no. 5, 467–474.
- [Goe95] Michel X Goemans, *Worst-case comparison of valid inequalities for the TSP*, Mathematical Programming **69** (1995), no. 1-3, 335–349.
- [Har95] Bert Hartnell, *Firefighter! an application of domination*, the 24th Manitoba Conference on Combinatorial Mathematics and Computing, University of Manitoba, Winnipeg, Canada, 1995, 1995.
- [Har04] Stephen G Hartke, *Attempting to narrow the integrality gap for the firefighter problem on trees.*, Discrete Methods in Epidemiology, 2004, pp. 225–231.
- [HLP88] GH Hardy, JE Littlewood, and G Polya, *Inequalities cambridge univ*, Press, Cambridge (1988).
- [HNR19] Arash Haddadan, Alantha Newman, and R Ravi, *Shorter tours and longer detours: uniform covers and a bit beyond*, Mathematical Programming (2019), 1–29.
- [Hoo18] K Hoover, *Wildfire statistics congressional research service*, 2018.
- [KKG20] Anna R Karlin, Nathan Klein, and Shayan Oveis Gharan, *A (slightly) improved approximation algorithm for metric TSP*, arXiv preprint arXiv:2007.01409 (2020).
- [KLL14] Rolf Klein, Christos Levcopoulos, and Andrzej Lingas, *Approximation algorithms for the geometric firefighter and budget fence problems*, Latin American Symposium on Theoretical Informatics, Springer, 2014, pp. 261–272.
- [KLS15] Marek Karpinski, Michael Lampis, and Richard Schmieid, *New inapproximability bounds for tsp*, Journal of Computer and System Sciences **81** (2015), no. 8, 1665–1677.
- [KP20] Miroslav Kulich and Libor Preucil, *Multi-robot search for a stationary object placed in a known environment*.
- [KPY96] Elias Koutsoupias, Christos Papadimitriou, and Mihalis Yannakakis, *Searching a fixed graph*, International Colloquium on Automata, Languages, and Programming, Springer, 1996, pp. 280–289.
- [KRT99] Jon Kleinberg, Yuval Rabani, and Éva Tardos, *Fairness in routing and load balancing*, 40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039), IEEE, 1999, pp. 568–578.

- [MDZL08] Isabel Méndez-Díaz, Paula Zabala, and Abilio Lucena, *A new formulation for the traveling deliveryman problem*, *Discrete applied mathematics* **156** (2008), no. 17, 3223–3237.
- [Min89] Edward Minieka, *The delivery man problem on a tree network*, *Annals of Operations Research* **18** (1989), no. 1, 261–266.
- [NAS] *Fire information for resource management system*, NASA, <https://firms.modaps.eosdis.nasa.gov>.
- [PS14] Ian Post and Chaitanya Swamy, *Linear programming-based approximation algorithms for multi-vehicle minimum latency problems*, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2014, pp. 512–531.
- [PY93] Christos H Papadimitriou and Mihalís Yannakakis, *The traveling salesman problem with distances one and two*, *Mathematics of Operations Research* **18** (1993), no. 1, 1–11.
- [Ser78] AI Serdyukov, *O nekotorykh ekstremal’nykh obkhodakh v grafakh*, *Upravlyayemye sistemy* **17** (1978), 76–79.
- [Sit02] René Sitters, *The minimum latency problem is NP-hard for weighted trees*, *International conference on integer programming and combinatorial optimization*, Springer, 2002, pp. 230–239.
- [Sit14] ———, *Polynomial time approximation schemes for the traveling repairman and other minimum latency problems.*, *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, SIAM, 2014, pp. 604–616.
- [SW90] David B Shmoys and David P Williamson, *Analyzing the held-karp TSP bound: A monotonicity property with application*, *Information Processing Letters* **35** (1990), no. 6, 281–285.
- [SWVZ12] Frans Schalekamp, David P Williamson, and Anke Van Zuylen, *A proof of the Boyd-Carr conjecture*, *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, SIAM, 2012, pp. 1477–1486.
- [TCM20] Guy Tennenholtz, Constantine Caramanis, and Shie Mannor, *Sequential vaccination for containing epidemics*, medRxiv (2020).
- [Wil94] Todd Gerald Will, *Extremal results and algorithms for degree sequences of graphs*.
- [Wol80] Laurence A Wolsey, *Heuristic analysis, linear programming and branch and bound*, *Combinatorial Optimization II*, Springer, 1980, pp. 121–134.

A From Section 3

Our example for 1.78 inapproximability of all-norm-TSP has similar structure as the exponential sequence presented in Section 3 achieving the lower bound of 1.71, yet is computationally tuned and verified. This suggests the best (worst) example can have a different structure. Follows the locations of the destinations over the line to be visited by a traveler, starting at $x = 200$. All points are over the x -axis.

$$V = \{0, 200, 202, 204, 206, 208, 210, 212, 214, 216, 217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 228, 230, 232, 234, 236, 238, 240, 242, 244, 246, 250, 254, 258, 262, 266, 270, 274, 278, 282, 286, 289, 292, 295, 298, 301, 304, 307, 310, 313, 316, 316, 316, 316, 316, 316, 316, 316, 316, 316, 316, 316, 316, 322, 328, 334, 340, 346, 352, 358, 364, 370, 376, 382, 388, 394, 400, 406, 412, 418, 424, 430, 436, 446, 456, 466, 476, 486, 496, 506, 516, 526, 536, 540, 544, 548, 552, 556, 560, 564, 568, 572, 576, 595, 614, 633, 652, 671, 690, 709, 728, 747, 766, 775, 784, 793, 802, 811, 820, 829, 838, 847, 856, 888, 920, 952, 984, 1016, 1048, 1080, 1112, 1144, 1176, 1199, 1222, 1245, 1268, 1291, 1314, 1337, 1360, 1383, 1406, 1519, 1632, 1745, 1858, 1971, 2084, 2197, 2310, 2423, 2536\}$$