Lines: 494

A PTAS for TSP with Neighbourhoods Over Parallel Line Segments

Benyamin Ghaseminia
Department of Computing Science, University of Alberta, Edmonton, Canada

Mohammad R. Salavatipour
Department of Computing Science, University of Alberta, Edmonton, Canada

Abstract

We consider the Travelling Salesman Problem with Neighbourhoods (TSPN) on the Euclidean plane ($\mathbb{R}^2$) and present a Polynomial-Time Approximation Scheme (PTAS) when the neighbourhoods are parallel line segments with lengths between $[1, \lambda]$ for any constant value $\lambda \geq 1$. In TSPN (which generalizes classic TSP), each client represents a set (or neighbourhood) of points in a metric and the goal is to find a minimum cost TSP tour that visits at least one point from each client set. In the Euclidean setting, each neighbourhood is a region on the plane. TSPN is significantly more difficult than classic TSP even in the Euclidean setting, as it captures group TSP. A notable case of TSPN is when each neighbourhood is a line segment. Although there are PTASs for when neighbourhoods are fat objects (with limited overlap), TSPN over line segments is **APX**-hard even if all the line segments have unit length. For parallel (unit) line segments, the best approximation factor is $3\sqrt{2}$ from more than two decades ago. The PTAS we present in this paper settles the approximability of this case of the problem. Our algorithm finds a $(1+\varepsilon)$-factor approximation for an instance of the problem for n segments with lengths in $[1, \lambda]$ in time $n^{O(\lambda/\varepsilon^3)}$.

2012 ACM Subject Classification Theory of computation → Routing and network design problems; Theory of computation → Approximation algorithms analysis

Keywords and phrases Approximation Scheme, TSP Neighbourhood, Parallel line segments

Related Version *Full Version*: https://arxiv.org/abs/2504.02190

Funding *Mohammad R. Salavatipour*: NSERC

1 Introduction

The Travelling Salesman Problem (TSP) is one of the most fundamental and well-studied problems in combinatorial optimization due to its wide range of applications. In TSP, one is given a set of points in a metric space and the goal is to find a (closed) tour (or walk) of minimum length visiting all the points. For several decades, the classic algorithm by Christofides [10] and independently by Serdyukov [32] which implies a $\frac{3}{2}$-approximation was the best-known approximation for TSP until a recent result in [24] which shows a slight improvement. Several generalizations (or special cases) of TSP have been studied as well, the most notable is when the points are given in fixed dimension Euclidean space. Arora and Mitchell [4, 30] presented different PTASs for (fixed dimension) Euclidean TSP. There have been many papers that have extended these results. Arkin and Hassin [3] introduced the notion of TSP with neighbourhoods (TSPN).

An instance of TSPN is a set of neighbourhoods (or regions) given in a metric space, and the goal is to find a minimum length (or cost) tour that visits all these regions. Each region can be a single point or could be defined by a subset of points. They gave several $O(1)$-approximations for the geometric settings where each region is some well-defined shape on the plane, e.g. disks, and parallel unit length segments. Several papers have studied TSPN for various classes of neighbourhoods and under different metrics.

TSPN is much more difficult than TSP in general and in special cases, just as group Steiner tree is much more difficult than Steiner tree (one can consider each neighbourhood as a group/set from which at least one point needs to be visited). In group Steiner tree or group TSP, one is given a metric along with groups of terminals (each group is a finite set). The goal is to find a minimum cost Steiner tree (or a tour) that contains (or visits) at least one terminal from each group. TSPN generalizes group TSP by allowing infinite size groups. Using the hardness result for group Steiner tree [22], it follows that general TSPN is hard to approximate within a factor better than $\Omega(\log^{2-\varepsilon} n)$ for any $\varepsilon > 0$ even on tree metrics. The algorithms for group Steiner tree on trees in [20], and embedding of metrics onto tree metrics in [18], imply an $O(\log^3 n)$-approximation for TSPN in general metrics. Unlike Euclidean TSP (which has a PTAS), TSPN is **APX**-hard on the Euclidean plane (i.e. $\mathbb{R}^2$) [11]. The special case when each region is an arbitrary finite set of points in the Euclidean plane (group TSP) has no constant approximation [31] and the problem remains **APX**-hard even when each region consists of exactly two points [13].

Focusing on Euclidean metrics, most of the earlier work have studied the cases where the regions (or objects) are *fat*. Roughly speaking, it usually means the ratio of the smallest enclosing circle to the largest circle fitting inside the object is bounded. There are some work on when regions are *not* fat, most notably when the regions are (infinite) lines or line segments or in higher dimensions when they are hyperplanes. For the case of infinite line segments on $\mathbb{R}^2$, the problem for n lines can be solved exactly in $O(n^4 \log n)$ time by a reduction to the Shortest Watchman Route Problem (see [12, 23]). For the same setting, Dumitrescu and Mitchell [14] presented a linear time $\frac{\pi}{2}$-approximation, which was improved to $\sqrt{2}$ by Jonsson [23] (again in linear time). For infinite lines in higher dimensions (i.e. ≥ 3), the problem is proved to be **APX**-hard (see [2] and references there). For neighbourhoods being hyperplanes and dimension being $d \geq 3$, Dumitrescu and Tóth [15] present a constant factor approximation (which grows exponentially with d). For arbitrary d, they present an $O(\log^3 n)$-approximation. For any fixed $d \geq 3$, authors of [1] present a PTAS.

For parallel (unit) line segments on the plane ($\mathbb{R}^2$) Arkin and Hassin [3] presented a $(3\sqrt{2}+1)$-approximation which was improved to $3\sqrt{2}$ by [14], which remains the best-known approximation for this case as far as we know for over two decades. Authors of [17] proved that TSPN for unit line segments (in arbitrary orientation) is **APX**-hard. In this paper, we settle the approximability of TSPN when regions are parallel line segments of similar length (unit length is a special case) and present a PTAS for it. As mentioned above, the best-known approximation for unit length parallel segments has ratio $3\sqrt{2}$ [14]. We first focus on the case of unit line segments and show how our result extends to when line segments have bounded length ratio. This is in contrast with the **APX**-hardness of [17] for unit line segments with arbitrary orientation. Our result also implies a $(2+\varepsilon)$-approximation for the case where the segments can be both vertical and horizontal.

1.1 Related Work

The work on TSPN is extensive, we list a subset of the most notable and relevant work here and refer to the references of them for earlier works. All works listed are for $\mathbb{R}^2$ metric. Arkin and Hassin [3] presented constant factor approximations for several TSPN cases including when the neighbourhoods are parallel unit-length line segments (with ratio $3\sqrt{2}+1$). Very recently, PTASs were proposed for the case of unit disks and unit squares in [5]. Mata and Mitchell [25] presented $O(\log n)$-approximation for general connected polygonal neighbourhoods. Mitchell [28] gave a PTAS the case that regions are *disjoint* fat objects. This built upon his earlier work on PTAS for Euclidean TSP [30]. However, it is still open whether the problem is **APX**-hard for disjoint (general) shapes. Dumitrescu and Mitchell [14] presented several results, including $O(1)$-approximation for TSPN where objects are connected regions of the same or similar diameter. This implies $O(1)$-approximation for line segments of the same length in arbitrary orientation. They also improved the bounds for cases of parallel segments of equal length, translates of a convex region, and translates of a connected region. For parallel (unit) line segments they obtain a $3\sqrt{2}$-approximation which remained the best ratio for this class until now. Feremans and Grigoriev [19], gave a PTAS for the case that regions are disjoint fat polygons of similar size. A similar result was obtained by [7] which gave a PTAS for TSPN for disjoint fat regions of similar sizes. Mitchell [29] presented an $O(1)$-approximation for planar TSPN with pairwise-disjoint connected neighbourhoods of any size or shape (see also [11]). Subsequently, [16] gave $O(1)$-approximation for discrete setting where regions are overlapping, convex, and fat with similar size.

These results were further generalized by Chan and Elbassioni [8] who presented a QPTAS for fat weakly disjoint objects in doubling dimensions (this allows limited amount of overlapping of the objects). This was further improved to a PTAS for the same setting [9]. The TSPN problem is even harder if the neighbourhoods are disconnected. See the surveys of Mitchell [26, 27]. If the metric is planar and each group (or neighbourhood) is the set of vertices around a face, [6] presents a PTAS for the group Steiner tree and a $(2+\varepsilon)$-approximation for TSPN.

1.2 Our Results and Technique

The main result of this paper is the following theorem.

▶ **Theorem 1.** *Given a set of n parallel (say vertical) line segments with lengths in $[1,\lambda]$ for a fixed λ as an instance of TSPN, there is an algorithm that finds a $(1+\varepsilon)$-approximation solution in time $n^{O(\lambda/\varepsilon^3)}$.*

The algorithm we present is randomized but can be easily derandomized (like the PTAS for classic TSP). To simplify the presentation, we give the proof for the case of unit line segments first.

This problem generalizes the classic (point) TSP (at a loss of $(1+\varepsilon)$ factor). Given a point TSP instance, scale the plane so that the minimum distance between the points is at least $1/\varepsilon$; call this instance $\mathcal{I}$. Obtain instance $\mathcal{I}'$ for line TSP by placing a vertical unit line segment over each point. It is easy to see that the optimum solutions of $\mathcal{I}$ and $\mathcal{I}'$ differ by at most an ϵ factor.

The difficult cases for line TSP are when the line segments are not too far apart (for e.g. they can be packed in a box of size $O(\sqrt{n})$ or smaller). There are two key ingredients to our proof that we explain here. One may try to adapt the hierarchical decomposition by Arora [4] to this setting. Following that hierarchical decomposition, the first issue is that some line segments might be crossing the horizontal dissecting lines, and so we don't have independent sub-instances, and it is not immediately clear in which subproblem these crossing segments must be covered. Note that the number of line segments crossing a dissecting line can be large. Our first insight is the following:

Insight 1: *At a loss of $(1+\varepsilon)$, we can drop the line segments crossing horizontal dissecting lines and instead requiring a subset of portals of each square (of the dissection) to be visited, provided we continue the quad-tree decomposition until each square has size $\Theta(1/\varepsilon)$.*

In other words, assuming all the squares in the decomposition have height at least $\Omega(1/\varepsilon)$, then at a small loss we can show that a solution for the modified instance where line segments on the boundary of the squares are dropped, can be extended to a solution for the original instance. So, proving this property allows us to work with the hierarchical (quad-tree) decomposition until squares of size $\Theta(1/\varepsilon)$. This can be proved by a proper packing argument. But then we need to be able to solve instances where the height of the sub-problem instance is bounded by $O(1/\varepsilon)$. Let's define the notion of shadow of a solution (or in general, shadow of a collection of paths on the plane) as the maximum number of times a vertical sweeping line Γ (that moves from left to right) intersects any of these paths. Our second insight is the following:

Insight 2: *If we consider a window that is a horizontal strip of height $O(h)$ and move this window vertically anywhere over an optimum solution, then the shadow of the parts of optimum visible in this strip is at most $O(h)$.*

In other words, one expects that in the base case of the decomposition (where squares have height $\Theta(1/\varepsilon)$), the shadow is bounded by $O(1/\varepsilon)$. Despite our efforts, proving this appears to be more difficult than thought and it seems there are examples where, even in the unit length segments, the shadow may be large (see Figure 1). We were not able to prove this nor come up with an explicit counterexample. However, we are able to prove the following slightly weaker version that still allows us to prove the final result:

(Revised) Insight 2: *There is a $(1+\varepsilon)$-approximate solution such that the shadow of any strip of height h over that solution is bounded by $O(f(\varepsilon)\cdot h)$. for some function $f(\cdot)$.*

The proof of this insight forms the bulk of our work. We characterize specific structures that would be responsible for having a large shadow in a solution and show how we can modify the solution to a $(1+\varepsilon)$-approximate one with shadow $O(1/\varepsilon)$. Consider the unit line segments case, and suppose opt is the cost of an optimum solution.

▶ **Theorem 2.** *Given any $\varepsilon > 0$, there is a solution $\mathcal{O}'$ of cost at most $(1+\varepsilon)\cdot opt$ such that in any strip of height 1, the shadow of $\mathcal{O}'$ is $O(1/\varepsilon)$.*

We will show that this near-optimum solution has in fact further structural properties that allow us to solve the bounded height cases at the base cases of the hierarchical decomposition using a Dynamic Program (DP) which, later on, is referred to as the inner DP. Proof of this theorem is fairly long and involves multiple steps that gradually proves structural properties for specific configurations.

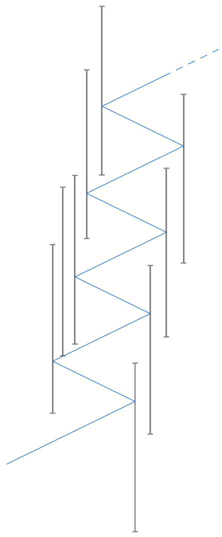

Figure 1 A potential arrangement of line segments where the solution has a large shadow

To get a very high level idea of the proof of this theorem, consider some fixed optimum solution OPT. We decompose the problem into horizontal *strips* by drawing horizontal lines, called *cover-lines* that are 1-unit apart. The region of the plane between two consecutive cover-lines is called a *strip*. Note that each line segment crosses one cover-line (or might touch exactly two consecutive cover-lines). Let's consider one strip, say S, and consider the intersection of OPT with this strip. This intersection looks like a collection of paths that enter/exit this strip. We define the shadow of this strip similarly: consider a hypothetical vertical sweep line that moves left to right along the x-axis, the maximum number of intersections of this sweep line with these pieces of OPT restricted to S is the shadow in S. We show that we can modify OPT to a near-optimum solution of cost at most $(1+\varepsilon)$ times that of OPT so that the shadow in each strip is at most $O(1/\varepsilon)$. To prove this, we show that there are certain potential structures that can cause OPT having a large shadow in a strip, one of which we call a zig-zag (see Figure 5). We show that we can modify OPT (at a small increase to the cost) so that the shadow becomes bounded along each zig-zag (or similar structures).

Organization of the paper: We start by some preliminaries in the next section. In Section 3, we present some structural properties of a near-optimum solution and prove Theorem 2. We describe the main algorithm in Section 4, which includes the outer DP and inner DP. All the missing proofs appear in the full version of the paper [21].

2 Preliminaries

Suppose we are given n vertical line segments $s_1, \ldots, s_n$ of lengths in the range $[1, \lambda]$ for some constant $\lambda \geq 1$, and the top and bottom points of each s_i are denoted by s_i^t and s_i^b, respectively. These end-points are also called *tips* of the segment. For any point p, let $x(p)$ and $y(p)$ denote the x and y-coordinates of p, respectively. Similarly, for any segment or vertical line s, let $x(s)$ denote its x-coordinate. For two points p, q, we use $||pq||$ to denote the Euclidean distance between them. A feasible (TSP) tour is specified by a sequence of points where each of these points is on one of the segments of the instance and each line segment has at least one such point, and the tour visits these points consecutively using straight lines. The line that connects two consecutive points in a tour is called a *leg* of the tour. In our problem, the goal is to find a TSP tour of minimum total length. As mentioned earlier, we focus on the case where all the line segments have length 1 and then show how the proof easily extends to the setting where they have lengths in $[1, \lambda]$. Fix an optimum solution, which we refer to by OPT and use opt to refer to its cost. Our goal is to show the existence of a near-optimum (i.e. $(1+\varepsilon)$-approximate) structured solution that allows us to find it using dynamic programming.

First, we show at a small loss we can assume all the line segments have different x-coordinates. We assume that the minimal bounding box of these line segments has length L and height H. For now, assume $H > 3$ (case of $H \leq 3$ is easier, see Theorem 4). Let $B = \max\{L, H-2\}$. So opt $\geq 2B$; we can also assume $B \leq \frac{n}{\varepsilon}$, because otherwise opt $\geq 2n/\varepsilon$ and if we consider an arbitrary point on each line segment (say the lower tip) and use a PTAS for the classic TSP for these points, then it will be a PTAS for our original instance as well (because we pay at most an extra $+2$ for each line for a total of $2n$ which is $O(\varepsilon \cdot \text{opt})$). For a given $\epsilon > 0$, consider a grid on the plane with side length $\frac{\epsilon B}{n^2}$. Now move each line segment (parallel to the y-axis) so that the lower tip of each s_i is moved to the nearest grid point where there is no other line segment s_i with that x-coordinate. By doing this, all the segments will have different x-coordinates and each segment would move at most $\frac{\sqrt{2}}{2} \cdot \frac{\varepsilon B}{n} < \frac{\epsilon B}{n}$, and in total, all segments would move at most a distance of ϵB. So the optimum value of the new instance has cost at most $(1+\varepsilon) \cdot \text{opt}$. For simplicity of notations, from now on we assume the original instance has this property and let OPT (and opt) refer to an optimum (and its value) of this modified instance.

3 Properties Of A Structured Near-Optimum Solution

One special instance of the problem is when there is a horizontal line that crosses all the input segments. This special case can be detected and solved easily. Otherwise, any optimum solution will visit at least 3 points that are not colinear. In such cases, like in the classic (point) TSP [4], we can assume the optimum does not cross itself, i.e. there are no two legs of the optimum ℓ (between points p, q) and ℓ' (between points p', q') that intersect, as otherwise removing these two and adding the pair of $pq', p'q$ or pp', qq' will be a feasible solution of smaller cost.

▶ **Definition 3.** *Given a collection $\mathcal{P}$ of paths on the plane and a vertical line at point $x_0 \in \mathbb{R}$, the* ***shadow*** *at x_0 is the number of legs of the paths in $\mathcal{P}$ that have an intersection with the vertical line at x_0. The shadow of a given range $[a, b]$ is defined to be the maximum shadow of values $x_0 \in [a, b]$.*

Note that if a solution is self-crossing, the operation of uncrossing (which reduces the cost) does not increase the shadow

Suppose the sequence of points of OPT is $p_1, p_2, \ldots, p_\sigma$ and the straight lines connecting these points (i.e. legs of OPT) are $\ell_1, \ell_2, \ldots, \ell_\sigma$ where ℓ_i connects two points p_i, p_{i+1} (with $p_{\sigma+1} = p_1$), and each s_i has at least one point p_j on it. We consider OPT oriented in this order, i.e. going from p_i to p_{i+1}. Since all segments have distinct x-coordinates, we can assume No two consecutive points p_i, p_{i+1} are on the same line segment by short-cutting (so no leg ℓ_i is vertical) and all points p_i on distinct line segments have distinct x-coordinates.

As mentioned before, let the length of the sides of the minimal bounding box of an instance of the problem be $L \times H$. By proving the following special case, we show we can instead focus on the cases that H is large.

▶ **Theorem 4.** *If $H \leq 3$, then the shadow of an optimum solution is at most 2.*

From now on we assume that $H > 3$. Our main goal is to prove Theorem 2.

▶ **Definition 5.** *Given a segment s and a leg ℓ incident to a point on s, we say ℓ is to the* ***left*** *of s if ℓ is entirely in the subplane $x \leq x(s)$; and ℓ is to the* ***right*** *of s if ℓ is entirely in the subplane $x \geq x(s)$.*

Since there are no vertical legs, there is no leg that is both to the left and to the right of a segment of the instance at the same time. Consider any segment s_i and suppose that ℓ_j, ℓ_{j+1} are the two legs of OPT with common end-point p_j that is on s_i. Let s_i^t and s_i^b denote the top and the bottom tips of s_i. We consider 3 possible cases for the location of p_j and the arrangement of ℓ_j, ℓ_{j+1}. Informally, one possibility is that the two legs ℓ_j, ℓ_{j+1} form a straight line that crosses s_i at p_j; one possibility is that the two legs are touching s_i at one of its tips (i.e. $p_j = s_i^t$ or $p_j = s_i^b$) such that one is to the left and one is to the right of s_i and they don't make a straight line, and the third possibility is that the two legs ℓ_j, ℓ_{j+1} are on the same side (both left or both right) of s_i.

▶ **Observation 6.** *Consider any segment s_i (with top/bottom points s_i^t,s_i^b) and suppose that ℓ_j, ℓ_{j+1} are the two legs of OPT with common end-point p_j that is on s_i. Then either:*

- *The subpath of OPT going through p_{j-1}, p_j, p_{j+1} forms a straight line (i.e. $\angle \ell_j p_j \ell_{j+1} = \pi$), and ℓ_j, ℓ_{j+1} are on two sides (left/right) of s_i; then we call p_j a* ***straight point****, or*
- *p_j is a tip of s_i (i.e. $p_j = s_i^t$ or $p_j = s_i^b$), $\angle \ell_j p_j \ell_{j+1} \neq \pi$ and ℓ_j and ℓ_{j+1} are on two sides of s_i (one left and one right); in this case p_j is called a* ***break point****, or*
- *both ℓ_j, ℓ_{j+1} are on the left or on the right of s_i; in this case p_j is called a* ***reflection point****.*

For the case of a reflection point p_j with two legs ℓ_j, ℓ_{j+1}, if both legs are to the left of the segment it is called a *left reflection* point and otherwise it is a *right reflection* point. Also note that if ℓ_j, ℓ_{j+1} are on the two sides of s_i and $\angle \ell_i p_j \ell_{i+1} \neq \pi$, then p_j must be a tip, or else we could move p_j slightly up or down and reduce the length of OPT (see Figure 2).

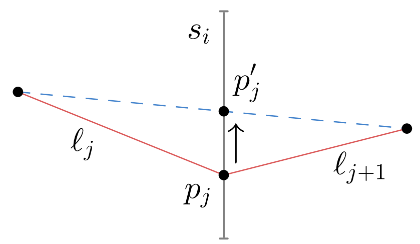

Figure 2 If p_j isn't a tip of s_i, then ℓ_j, ℓ_{j+1} must be collinear

▶ **Lemma 7.** *If P is a subpath of OPT with end-points p, q where both are to the right of a vertical line Γ, and if P crosses Γ, then the left-most point on P to the left of Γ is a right reflection point (symmetric statement holds for opposite directions).*

▶ **Definition 8.** *Consider an arbitrary reflection point r on a segment s. Let the two legs of OPT incident to r visited before and after r (on the orientation of OPT) be ℓ^- and ℓ^+, respectively. ℓ^- is said to be* ***on top of*** *ℓ^+ if all the points of ℓ^- have larger y-coordinate than all of the points of ℓ^+. In this case we also call ℓ^- the* upper leg *and ℓ^+ the* lower leg*. Also, in this case, r is called a* ***descending*** *reflection point. If ℓ^+ is on top of ℓ^-, then r is called an* ***ascending*** *reflection point.*

If ℓ_j, ℓ_{j+1} are two legs incident to a reflection point p on a segment s, if the angle between ℓ_j and s is the same as the angle between ℓ_{j+1} and s (i.e. ℓ_{j+1} is like the reflection of a ray ℓ_j on mirror s) then p is called a *pure reflection* point.

▶ **Lemma 9.** *Any reflection point that is not a tip of a segment is a pure reflection point.*

We say a subpath contains a reflection point p_j if p_j is not the start or end vertex of the subpath (i.e. both legs of incident to p_j belong to that subpath.)

▶ **Lemma 10.** *If a sweeping vertical line Γ moves left to right on the x-axis, the only values of x for which the shadow at Γ changes will be when Γ hits a reflection point on that x-coordinate. Specifically, this means that any subpath of OPT that doesn't contain a reflection point, must have a shadow of 1 throughout its length.*

▶ **Definition 11.** *Let P_1 and P_2 be any two subpaths of OPT. We say P_1 is **above** P_2 in range $I = [x_0, x_1]$ if for every vertical line Γ with $x(\Gamma) \in I$, the top-most intersection of Γ with these two paths is a point on P_1. We say P_2 is **below** P_1 if the bottom-most intersection of Γ with P_1, P_2 is a point on P_2. Similarly, we say L_1 is to the **left** of L_2 in range $I' = [y_0, y_1]$ if for every horizontal line Λ with $y(\Lambda) \in I'$, the left-most intersection point of Λ with L_1, L_2 (i.e. one with the least x value) always belongs to L_1. We say L_2 is to the **right** of L_1 if the right-most intersection of Λ is with L_2.*

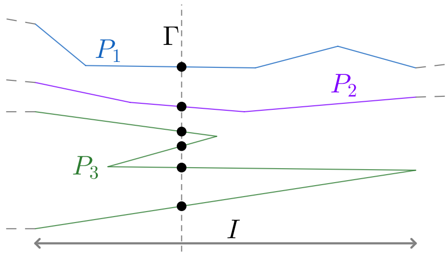

Figure 3 In range I, P_1 is above P_2, P_3, and P_2 is above P_3

▶ **Lemma 12.** *For any distinct points p_j and $p_{j'}$ on OPT, following OPT according to its orientation, either the path from p_j to $p_{j'}$ or the path from $p_{j'}$ to p_j must contain at least one reflection point.*

▶ **Lemma 13.** *Among the set of points visited by OPT following its orientation, suppose $p_j, p_{j'}$, $j < j'$ (on segments s_i, $s_{i'}$, respectively) are two consecutive reflection points (i.e. no other reflection point exists in between them). Then p_j and $p_{j'}$ cannot be both left or both right reflection points. Furthermore, if s_i is to the left of $s_{i'}$ then p_j is a right reflection and $p_{j'}$ is a left reflection (the opposite holds if $s_{i'}$ is to the left of s_i).*

▶ **Corollary 14.** *Consecutive reflection points in OPT alternate between left and right reflections.*

3.1 Strips, Zig-zag, Sink

We decompose the problem into horizontal *strips* by drawing some parallel horizontal lines, which we call *cover-lines*. Starting from the bottom tip of the top-most segment, draw parallel horizontal lines that are 1-unit apart, these are our *cover-lines*. Each input segment is considered "covered" by the top-most (i.e. the first) cover-line that intersects it. Let's label these cover-lines by $C_1, C_2, \ldots$.

▶ **Definition 15** (strip, top/bottom segments)**.** *The region of the plane between two consecutive cover-lines $C_\tau, C_{\tau+1}$ is called a **strip** and denoted by S_τ. We consider $C_\tau, C_{\tau+1}$ part of S_τ as well. The input line segments that are intersecting the top cover-line of S_τ (C_τ) are called **top segments,** and the segments covered by the bottom cover-line ($C_{\tau+1}$) are called **bottom segments** of the strip.*

We show the near-optimum solution guaranteed by Theorem 2 has more structural properties that will be defined later. Note that once we prove that theorem, it follows that if we restrict a solution to $h > 1$ many strips, then the shadow is bounded by $O(h/\varepsilon)$ as well.

For now, let us focus on an (arbitrary) strip S_τ and imagine we cut the plane along $C_\tau, C_{\tau+1}$ and look at the pieces of line segments of the instance left inside this strip, along with pieces of OPT inside S_τ. Each top segment is now a partial segment in S_τ that has one end on C_τ; and each bottom segment has one end on $C_{\tau+1}$. Let OPT_τ be the restriction of OPT to S_τ. For each leg of OPT that intersects C_τ or $C_{\tau+1}$, we add a dummy point at the intersection(s) of that leg with C_τ and $C_{\tau+1}$ (so that the components of OPT_τ become consistent with our definition of legs). So OPT_τ can be seen as a collection of subpaths within S_τ (possibly along C_τ or $C_{\tau+1}$). Following the orientation of OPT, each subpath of OPT_τ is formed by it intersecting with S_τ, traveling within S_τ (possibly along one of the cover-lines), and then exiting S_τ. Using the dummy points added, each path in OPT_τ is a subpath of OPT that is between two points on cover-lines (these are called the entry points of the path with the strip. A formal definition is provided below).

▶ **Definition 16** (entry points, loops, ladders)**.** *For each subpath P_j of OPT_τ, let e_j and o_j be the first and last intersections of P_j with the interior of S_τ. Points e_j and o_j are called the **entry points** of P_j.*
*If both e_j and o_j lie on the same cover-line (either C_τ or $C_{\tau+1}$), then P_j is called a **loop**, otherwise it's called a **ladder**. If a subpath of OPT_τ enters S_τ at e_j on a cover-line and follows on that cover-line to point o_j and exits the strip, it is a special case of loop that we refer to as a **cover-line loop**.*

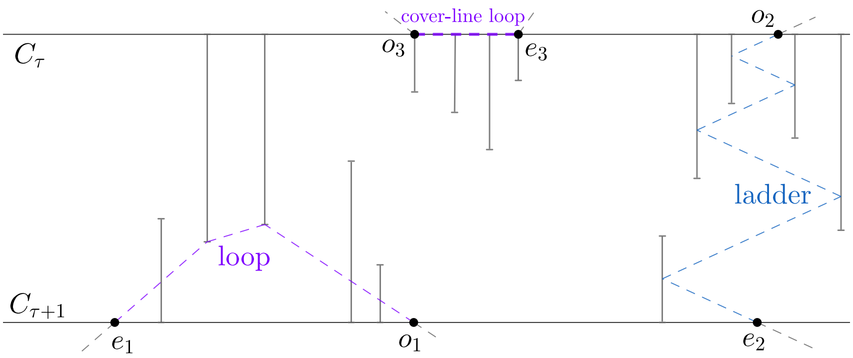

Figure 4 An example of loops and ladders in a strip S_τ (i.e. area between cover-lines $C_\tau, C_{\tau+1}$)

Since we're assuming $H > 3$ (see Theorem 4), we get that OPT is not limited to a single strip, and that it has to indeed enter and exit any given strip that it intersects with (i.e. there is no strip that OPT completely lies inside it). Note that if a path of OPT_τ is a cover-line loop, i.e. a section of the line C_τ or $C_{\tau+1}$, then the entry points of that path must be the two end-points of this section. In other words, if for a cover-line loop of OPT_τ the first point is e_j on (say) C_τ, and the last point is o_j on C_τ, then this subpath must be travelling straight from e_j to o_j without any change of direction. This is true because otherwise, that cover-line loop would have to go back and forth on some portion on a cover-line, which is only possible if it's self-intersecting; but this is against our assumption that OPT is not self-crossing. The two structures defined below (called a zig-zag and a sink) are the two configurations that can cause a large shadow.

▶ **Definition 17** (Zig-zag/Sink). *Consider any loop or ladder of OPT_τ, call it P. Let $\mathcal{R} = r_1, r_2, \ldots, r_t$ be the sequence of points of P that are reflection points (indexed by the order they're visited). Consider any maximal sub-sequence $r_j, r_{j+1}, \ldots, r_q$ of $\mathcal{R}$ (with $q \geq j+1$) such that all are ascending or all are descending reflections, and the segments containing them alternate between top and bottom segments; then the subpath P that starts at r_j and ends at r_q is called a **zig-zag**.*
*If $r_j, r_{j+1}, \ldots, r_q$ is a maximal sub-sequence of $\mathcal{R}$ that are all ascending or all descending, and all belong to top segments or all belong to bottom segments; then the subpath P that starts at r_j and ends at r_q is called a **sink** (see Figure 5).*

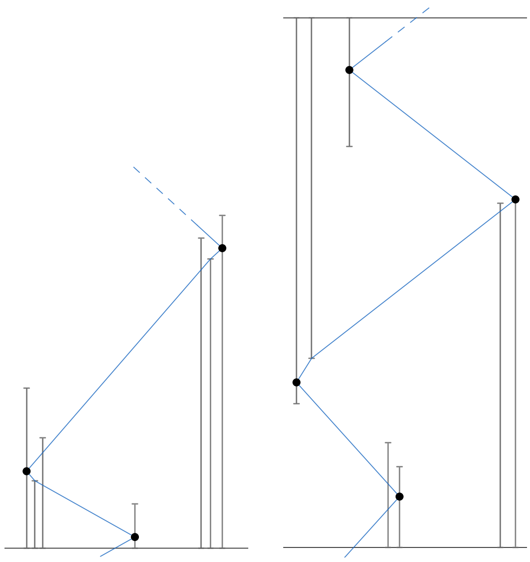

Figure 5 Examples of a sink (left) and a zig-zag (right). Bold dots indicate the reflection points.

Note that each zig-zag has at least two reflection points (or else it will be a sink). Also, using Corollary 14, the reflection points in a zig-zag or sink should alternate between left and right reflection. The next important lemma is used critically to show that very specific structures (namely zig-zags and sinks) are responsible for having large shadow along a ladder or loop in OPT_τ. Additionally, we can partition each ladder/loop into parts (subpaths), such that the shadow of the ladder/loop is equal to the maximum shadow among these parts, and each part is a path that consists of up to three sinks and/or zig-zags. So the shadow of a loop/ladder is within $O(1)$ of the maximum shadow of zig-zag/sinks along it.

▶ **Lemma 18.** *Consider any strip S_τ and any ladder or loop $P \in OPT_\tau$ within S_τ. Suppose the sequence of reflection points of P is $r_1, \ldots, r_q$. Then these reflection points can be partitioned into disjoint parts, say part i consists of reflection points $r_{a_i}, r_{a_i+1}, \ldots, r_{a_j}$, where the subpath of P from r_{a_i} to r_{a_j} is concatenation of up to three sections in the following order: **a)** A sink , followed by **b)** A zig-zag , followed by **c)** A sink, where any of these three sections can possibly be empty, and the last reflection of a section is common with the first* *reflection of the next section. Furthermore, for any vertical line Γ, there is at most one of these parts (of the partition) that intersects with it, i.e. the shadow of the ladder/loop is the maximum shadow among the parts plus 2.*

The proof of this lemma is rather involved (see [21] for details). To give an idea of the proof, we essentially show that for any loop/ladder in any strip, the vertical line at which the largest shadow for that loop/ladder happens, can intersect with at most two sinks and a zig-zag. So the shadow of a loop or ladder is within $O(1)$ of the maximum shadow of the zig-zags and sinks along itself. The following lemma is one of the main components of proof of Theorem 2.

▶ **Lemma 19.** *Consider OPT_τ for an arbitrary strip S_τ, and let opt_τ be the total cost of OPT_τ. Given any $\varepsilon > 0$, we can change OPT_τ to a solution of cost at most $(1 + O(\varepsilon)) \cdot opt_\tau$ where the shadow of each zig-zag and sink is at most $O(1/\varepsilon)$.*

The following corollary immediately follows from Lemmas 18 and 19:

▶ **Corollary 20.** *There is a $(1 + \varepsilon)$-approximate solution in which all loops/ladders have shadow $O(1/\varepsilon)$.*

▶ **Definition 21.** *Let $\mathcal{R} = p_i, p_{i+1}, \dots, p_q$ be any sequence of consecutive points in OPT such that p_i and p_q are reflection points. If none of p_j's ($i < j < q$) in $\mathcal{R}$ is a tip of a segment, then $\mathcal{R}$ is called a* ***pure reflection sequence****.*

So each point in $\mathcal{R}$ is either a straight point or a pure reflection according to Lemma 9.

▶ **Lemma 22.** *Consider OPT_τ for an arbitrary strip S_τ and suppose the total length of legs of OPT_τ is opt_τ. Given $\varepsilon > 0$, we can change OPT_τ to a solution of cost at most $(1+\varepsilon) \cdot opt_\tau$ in which the size of any pure reflection sequence is bounded by $O(\frac{1}{\epsilon})$.*

Our next goal is to show that for any vertical line, it can intersect at most $O(1)$ many loops or ladders of OPT_τ in a strip S_τ. This together with the above corollary implies there is a $(1 + \varepsilon)$-approximate solution where the shadow in each strip S_τ is $O(1/\varepsilon)$.

▶ **Definition 23.** *A collection of loops and or ladders are said to be* overlapping *with each other if there is a vertical line that intersects all of them.*

▶ **Lemma 24.** *Consider OPT_τ, the restriction of OPT to any strip S_τ. We can modify the solution (without increasing the shadow or the cost) such that there are at most $O(1)$ loops or ladders in OPT_τ that all are overlapping with each other.*

Assuming the correctness of main lemmas defined above (i.e. Lemmas 18, 19, 22, and 24), we can now prove Theorem 2.

Proof of Theorem 2. If the height of the bounding box is at most 3, simply use Theorem 4. Otherwise, consider any strip S_τ (to be more precise, S_τ can be any arbitrary strip of height 1 in the plane). Using Lemma 19, there is a solution $\mathcal{O}''$ of cost at most $(1 + \varepsilon) \cdot \mathrm{opt}$ where the shadow of each sink and zig-zag is bounded by $O(1/\varepsilon)$. By Lemma 18, each loop or ladder in S_τ has a shadow that is at most 3 times the maximum shadow of a sink or zig-zag in it, plus two. So each loop or ladder has shadow $O(1/\varepsilon)$. Finally, Lemma 24 shows that there can be at most $O(1)$ overlapping loops or ladders in a strip. Thus, the overall shadow of $\mathcal{O}''$ in S_τ is bounded by $O(1/\varepsilon)$. Furthermore, we apply Lemma 22 on O'' to get a solution $\mathcal{O}'$. This new solution has the property that with an additional cost of factor $(1 + O(\varepsilon))$ compared to $\mathcal{O}''$, the size of any pure reflection sequence is bounded by $O(1/\varepsilon)$. The total cost of $\mathcal{O}''$ is at most $(1 + O(\varepsilon)) \cdot \mathrm{opt}$ and the shadow is bounded by $O(1/\varepsilon)$. ◀

Note: Although having a bound on the length of pure reflection sequences is not in the statement of the theorem, we use this extra property crucially in designing our DP to find a near-optimum solution.

4 Main Algorithm and Reduction to Structured Bounded Height Instances

As mentioned in the introduction, we follow the paradigm of Arora [4] for designing a PTAS for classic Euclidean TSP with some modifications. We focus more on defining the modifications that we need to make to that algorithm. First, we describe the main algorithm and how it reduces the problem into a collection of instances with a constant-height bounding box. We show how those instances can be solved using another DP (referred to as the *inner DP*), and how we can combine the solutions for them using another DP (referred to as the *outer DP*) to find a near-optimum solution of the original instance. Recall that in Section 2, we assumed the minimal bounding box of the instance has length L and height H and we defined $B = \max\{L, H-2\}$, and also we can assume that $B \leq \frac{n}{\varepsilon}$. We moved each line segment to be aligned with a grid point with side length $\frac{\epsilon B}{n^2}$ (at a loss of $(1+\varepsilon)$ at approximation). Now, we scale the grid (as well as the line segments of the instance) by a factor of $\rho = \frac{4n^2}{\epsilon B}$ so that each grid cell has size 4. We obtain an instance where each line segment has length ρ, all have even integer coordinates, any two segments are at least 4 units apart, and the bounding box has size $N = O(n^2/\varepsilon)$. Let this new instance be $\mathcal{I}$. Note that if we define cover-lines as before but with a spacing of ρ, all the arguments for the existence of a near-optimum solution with a bounded shadow in any strip (the area between two consecutive cover-lines) still hold. We will present a PTAS for this instance. It can be seen that this implies a PTAS for the original instance of the problem. From now on, we use OPT to refer to an optimum solution of instance $\mathcal{I}$, and opt to refer to its cost. Note that since the bounding box has side length N, then opt $\geq 2N$. Similar to Arora's approach, we do the hierarchical dissectioning of the instance into nested squares using random axis-parallel dissectioning lines, and put portals at these dissecting lines. We continue this dissectioning process until the distances between horizontal (and so vertical) dissecting lines is $h \cdot \rho$ for $h = \lceil 1/\varepsilon \rceil$. So at the leaf nodes of our recursive decomposition quad-tree, each square is $(h \cdot \rho) \times (h \cdot \rho)$, and the height of the decomposition is $\log(N/\rho h) = O(\log n)$ since $B \leq \frac{n}{\varepsilon}$. We choose vertical dissecting lines only at odd x-coordinates so no line segment of the instance will be on a vertical dissecting line.

We define our cover-lines C_τ based on these horizontal dissecting lines carefully. Consider the first (horizontal) dissecting line we choose, this will be a cover-line, and then moving in both up and down directions from this line, we draw horizontal lines that are ρ apart. These will be all the cover-lines. Label the cover-lines from the top to bottom by $C_1, C_2, \ldots, C_\sigma$ in that order. As before, the smallest index τ such that C_τ crosses a line segment is the cover-line that "covers" that line segment. We partition the cover-lines into h groups based on their indices: Group G_j contains all those cover-lines with index τ where $j = \tau \pmod h$. Let G_{j*} be the group of cover-lines that includes the first horizontal dissecting line, and hence all the other horizontal dissecting lines as well. The arguments for the case of unit-length line segments to show there is a near-optimum solution in which the shadow in each strip of height 1 is $O(1/\varepsilon)$ (Theorem 2), also imply the same for the scaled instance $\mathcal{I}$. Furthermore, if we consider h consecutive strips, i.e. the area between two consecutive cover-lines in the same group G_j, then there is a near-optimum solution that has shadow $O(h/\varepsilon) = O(1/\varepsilon^2)$. Our goal is to show that, at a $(1+\varepsilon)$-factor loss, we can simply drop the line segments that are intersecting the horizontal dissecting lines (i.e. all those intersecting cover-lines in G_{j^*}) with appropriate consideration of portals (to be described). Removing the line segments that cross the dissecting lines allows us to decompose the instance into "independent" sub-instances that interact only via portals. The details of how to remove these segments which leads to proof of Lemma 25 below are deferred to the full version [21].

Similar to Arora's scheme for TSP, for $m = O(\frac{1}{\varepsilon}\log(N/\rho h))$, we place portals at all 4 corners of a square in the decomposition, plus an additional $m-1$ equally distanced portals along each side (so a total of $4m$ portals on the perimeter of a square of the dissection). For simplicity, we assume m is a power of 2 and at least $\frac{4}{\varepsilon}\log(N/\rho h)$. We say a tour is *portal respecting* if it crosses between two squares in our decomposition only via portals of the squares. A tour is *r-light* if it crosses the portals on each side of a square of the dissection at most r times. For classic (point) TSP, it can be shown that there is a near-optimum solution that is portal respecting and r-light for $r = O(1/\varepsilon)$. Our goal is to show a similar statement, except that we want the restriction of the tour to each "base" square of side length $O(h \cdot \rho)$ to have bounded (by $O(h/\varepsilon) = O(1/\varepsilon^2)$) shadow as well. We then show that we can find an optimum solution with a bounded shadow for the base cases using a DP. This will be our *inner DP*. We then show how the solutions of for the 4 sub-squares of a square in our decomposition can be combined into a solution for the bigger subproblem using the outer DP. We show that at a small loss in approximation (i.e. $O(\varepsilon \cdot \text{opt})$), we can drop all the line segments of input that are intersecting the horizontal dissecting lines (i.e. covered by a cover-line in group G_{j^*}), solve appropriate subproblems, and then extend the solutions to cover those dropped segments. This modification requires certain portals of each square in the decomposition to be visited in the solution for that square. We show there is a feasible solution that visits all the remaining segments as well as the "required" portals, of total cost at most $(1+\varepsilon) \cdot \text{opt}$, and that such a solution can be extended to a feasible solution visiting all the segments of the original instance (i.e. including the ones that we dropped) at an extra cost of $O(\varepsilon \cdot \text{opt})$.

▶ **Lemma 25.** *Given instance $\mathcal{I}$, there is another instance $\mathcal{I}'$ that is obtained by removing all the segments that are crossing cover-lines in G_{j^*} (i.e. intersecting horizontal dissecting lines), and instead some of the portals around (more precisely, the top and bottom sides of) each square of quad-tree dissection are required to be covered (visited); such that there is a solution for $\mathcal{I}'$ of cost at most $(1+O(\varepsilon)) \cdot opt$, and such a solution can be extended to a feasible solution of $\mathcal{I}$ of cost at most $(1+O(\varepsilon)) \cdot opt$. Furthermore, the shadow of the solution for $\mathcal{I}'$ between any two consecutive cover-lines in G_{j^*} is at most 4 more than the shadow of OPT between those two lines.*

The outer DP based on the quad-tree dissection is similar to the classic PTAS for Euclidean TSP. One can show that for $r = O(1/\varepsilon)$, there is a r-light portal respecting tour for $\mathcal{I}'$ with cost at most $(1+\varepsilon) \cdot \text{opt}'$ where opt' is the cost of an optimum solution fo $\mathcal{I}'$. The DP will also "guess" in the recursion for each square, which portals are the "required" portals around it. The base case of this DP will be instances with bounding box of size $\rho \cdot h$. For such instances, we solve the problem using an inner DP. Informally, the inner DP is a nontrivial generalization of the DP for the classic (and textbook example) bitonic TSP in which the shadow is 2. In our case, the shadow is $O(1/\varepsilon^2)$. We defer the details of both the inner and outer DP to the full version [21].

References

1. Antonios Antoniadis, Krzysztof Fleszar, Ruben Hoeksma, and Kevin Schewior. A ptas for euclidean tsp with hyperplane neighborhoods. *ACM Trans. Algorithms*, 16(3), jun 2020. doi:10.1145/3383466.
2. Antonios Antoniadis, Sándor Kisfaludi-Bak, Bundit Laekhanukit, and Daniel Vaz. On the approximability of the traveling salesman problem with line neighborhoods. In Artur Czumaj and Qin Xin, editors, *18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, June 27-29, 2022*, volume 227 of *LIPIcs*, pages 10:1–10:21, Tórshavn, Faroe Islands, 2022. Schloss Dagstuhl - Leibniz-Zentrum für Informatik. doi:10.4230/LIPICS.SWAT.2022.10.
3. Esther M. Arkin and Refael Hassin. Approximation algorithms for the geometric covering salesman problem. *Discret. Appl. Math.*, 55(3):197–218, 1994. doi:10.1016/0166-218X(94)90008-6.
4. Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998. doi:10.1145/290179.290180.
5. Sayan Bandyapadhyay, Katie Clinch, William Lochet, Daniel Lokshtanov, Saket Saurabh, and Jie Xue. *PTASes for Euclidean TSP with Unit Disk and Unit Square Neighborhoods*, pages 2326–2356. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611978322.78, arXiv:https://epubs.siam.org/doi/pdf/10.1137/1.9781611978322.78, doi:10.1137/1.9781611978322.78.
6. MohammadHossein Bateni, Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dániel Marx. A ptas for planar group steiner tree via spanner bootstrapping and prize collecting. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 570–583, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2897518.2897549.
7. Hans L. Bodlaender, Corinne Feremans, Alexander Grigoriev, Eelko Penninkx, René Sitters, and Thomas Wolle. On the minimum corridor connection problem and other generalized geometric problems. *Comput. Geom.*, 42(9):939–951, 2009. doi:10.1016/J.COMGEO.2009.05.001.
8. T.-H. Hubert Chan and Khaled M. Elbassioni. A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics. *Discret. Comput. Geom.*, 46(4):704–723, 2011. doi:10.1007/S00454-011-9337-9.
9. T.-H. Hubert Chan, Shuguang Hu, and Shaofeng H.-C. Jiang. A PTAS for the steiner forest problem in doubling metrics. In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016*, pages 810–819, Hyatt Regency, New Brunswick, New Jersey, USA, 2016. IEEE Computer Society. doi:10.1109/FOCS.2016.91.
10. Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.
11. Mark de Berg, Joachim Gudmundsson, Matthew J. Katz, Christos Levcopoulos, Mark H. Overmars, and A. Frank van der Stappen. TSP with neighborhoods of varying size. *J. Algorithms*, 57(1):22–36, 2005. doi:10.1016/J.JALGOR.2005.01.010.
12. Moshe Dror, Alon Efrat, Anna Lubiw, and Joseph S. B. Mitchell. Touring a sequence of polygons. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003*, pages 473–482, San Diego, CA, USA, 2003. ACM. doi:10.1145/780542.780612.
13. Moshe Dror and James B. Orlin. Combinatorial optimization with explicit delineation of the ground set by a collection of subsets. *SIAM J. Discret. Math.*, 21(4):1019–1034, 2008. doi:10.1137/050636589.

14 Adrian Dumitrescu and Joseph S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *J. Algorithms*, 48(1):135–159, 2003. doi:10.1016/S0196-6774(03)00047-6.

15 Adrian Dumitrescu and Csaba D. Tóth. The traveling salesman problem for lines, balls, and planes. *ACM Trans. Algorithms*, 12(3), apr 2016. doi:10.1145/2850418.

16 Khaled M. Elbassioni, Aleksei V. Fishkin, Nabil H. Mustafa, and René Sitters. Approximation algorithms for euclidean group TSP. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 1115–1126, Lisbon, Portugal, 2005. Springer. doi:10.1007/11523468_90.

17 Khaled M. Elbassioni, Aleksei V. Fishkin, and René Sitters. Approximation algorithms for the euclidean traveling salesman problem with discrete and continuous neighborhoods. *Int. J. Comput. Geom. Appl.*, 19(2):173–193, 2009. doi:10.1142/S0218195909002897.

18 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004. doi:10.1016/J.JCSS.2004.04.011.

19 Corinne Feremans and Alexander Grigoriev. Approximation schemes for the generalized geometric problems with geographic clustering. In *(Informal) Proceedings of the 21st European Workshop on Computational Geometry, March 9-11, 2005*, pages 101–102, Eindhoven, The Netherlands, 2005. Technische Universiteit Eindhoven. URL: http://www.win.tue.nl/EWCG2005/Proceedings/26.pdf.

20 Naveen Garg, Goran Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group steiner tree problem. In Howard J. Karloff, editor, *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, 25-27 January 1998*, pages 253–259, San Francisco, California, USA, 1998. ACM/SIAM. URL: http://dl.acm.org/citation.cfm?id=314613.314712.

21 Benyamin Ghaseminia and Mohammad R. Salavatipour. A ptas for travelling salesman problem with neighbourhoods over parallel line segments of similar length, 2025. URL: https://arxiv.org/abs/2504.02190, arXiv:2504.02190.

22 Eran Halperin and Robert Krauthgamer. Polylogarithmic inapproximability. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003*, pages 585–594, San Diego, CA, USA, 2003. ACM. doi:10.1145/780542.780628.

23 Håkan Jonsson. The traveling salesman problem for lines in the plane. *Inf. Process. Lett.*, 82(3):137–142, 2002. doi:10.1016/S0020-0190(01)00259-9.

24 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved bound on the integrality gap of the subtour LP for TSP. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, October 31 - November 3, 2022*, pages 832–843, Denver, CO, USA, 2022. IEEE. doi:10.1109/FOCS54457.2022.00084.

25 Cristian S. Mata and Joseph S. B. Mitchell. Approximation algorithms for geometric tour and network design problems (extended abstract). In Jack Snoeyink, editor, *Proceedings of the Eleventh Annual Symposium on Computational Geometry, June 5-12, 1995*, pages 360–369, Vancouver, B.C., Canada, 1995. ACM. doi:10.1145/220279.220318.

26 Joseph S. B. Mitchell. Geometric shortest paths and network optimization. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 633–701. North Holland / Elsevier, Amsterdam, NL, 2000. doi:10.1016/B978-044482537-7/50016-4.

27 Joseph S. B. Mitchell. Shortest paths and networks. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry, Second Edition*, pages 607–641. Chapman and Hall/CRC, London, UK, 2004. doi:10.1201/9781420035315.CH27.

28 Joseph S. B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In Nikhil Bansal, Kirk Pruhs, and Clifford Stein, editors, *Proceedings of the Eighteenth Annual*

ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, January 7-9, 2007, pages 11–18, New Orleans, Louisiana, USA, 2007. SIAM. URL: http://dl.acm.org/citation.cfm?id=1283383.1283385.

29 Joseph S. B. Mitchell. A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane. In David G. Kirkpatrick and Joseph S. B. Mitchell, editors, *Proceedings of the 26th ACM Symposium on Computational Geometry, June 13-16, 2010*, pages 183–191, Snowbird, Utah, USA, 2010. ACM. doi:10.1145/1810959.1810992.

30 Joseph SB Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric tsp, k-mst, and related problems. *SIAM Journal on computing*, 28(4):1298–1309, 1999.

31 Shmuel Safra and Oded Schwartz. On the complexity of approximating TSP with neighborhoods and related problems. In Giuseppe Di Battista and Uri Zwick, editors, *Algorithms - ESA 2003, 11th Annual European Symposium, September 16-19, 2003, Proceedings*, volume 2832 of *Lecture Notes in Computer Science*, pages 446–458, Budapest, Hungary, 2003. Springer. doi:10.1007/978-3-540-39658-1_41.

32 A. I. Serdyukov. O nekotorykh ekstremal'nykh obkhodakh v grafakh. *Upravlyaemye sistemy 17*, page pp. 76–79, 1978. URL: http://nas1.math.nsc.ru/aim/journals/us/us17/us17_007.pdf.