

CONSTANT-FACTOR APPROXIMATION FOR TSP WITH DISKS

Adrian Dumitrescu* Csaba D. Tóth†

Thursday 4th August, 2016

Abstract

We revisit the traveling salesman problem with neighborhoods (TSPN) and present the first constant-ratio approximation for disks in the plane: Given a set of n disks in the plane, a TSP tour whose length is at most $O(1)$ times the optimal can be computed in time that is polynomial in n . Our result is the first constant-ratio approximation for a class of planar convex bodies of arbitrary size and arbitrary intersections.

In order to achieve a $O(1)$ -approximation, we reduce the traveling salesman problem with disks, up to constant factors, to a minimum weight hitting set problem in a geometric hypergraph. The connection between TSPN and hitting sets in geometric hypergraphs, established here, is likely to have future applications.

Keywords: Traveling salesman problem, minimum weight hitting set, approximation algorithm.

1 Introduction

In the Euclidean Traveling Salesman Problem (ETSP), given a set of points in the Euclidean space \mathbb{R}^d , $d \geq 2$, one seeks a shortest closed curve (a.k.a. *tour*) that visits each point. In the TSP *with neighborhoods* (TSPN), each point is replaced by a point-set, called *region* or *neighborhood*, and the TSP tour must visit at least one point in each region, i.e., it must intersect each region. The oldest record that we could trace of this variant goes back to Arkin and Hassin [2]. Since the Euclidean TSP is known to be NP-hard in \mathbb{R}^d for every $d \geq 2$ [21, 22, 40], TSPN is also NP-hard for every $d \geq 2$. TSP is recognized as one of the corner-stone problems in combinatorial optimization. Other related problems in geometric network optimization can be found in the two (somewhat outdated) surveys by Mitchell [32, 33].

It is known that the Euclidean TSP admits a polynomial-time approximation scheme in \mathbb{R}^d , where $d = O(1)$, due to classic results of Arora [3] and Mitchell [31]. Subsequent running time improvements have been obtained by Rao and Smith [41]; specifically, the running time of their PTAS is $O(f(\varepsilon)n \log n)$, where $f(\varepsilon)$ grows exponentially in $1/\varepsilon$. In contrast, TSPN is generally harder to approximate. Typically, somewhat better approximations are available when the neighborhoods are *pairwise disjoint*, or *fat*, or have *comparable sizes*. We briefly review some of the previous work concerning approximation algorithms for TSPN.

*Department of Computer Science, University of Wisconsin–Milwaukee, WI, USA. Email: dumitres@uwm.edu.

†Department of Mathematics, California State University, Northridge, Los Angeles, CA; and Department of Computer Science, Tufts University, Medford, MA, USA. Email: cdtoth@acm.org. Research by this author was supported in part by the NSF award CCF-1423615.

Related work. Arkin and Hassin [2] gave constant-factor approximations for translates of a connected region, and more generally, for neighborhoods of pairwise *parallel* diameters, where the ratio between the longest and the shortest diameter is bounded by a constant. Dumitrescu and Mitchell [16] extended the above result to connected neighborhoods with comparable diameters. Bodlaender et al. [8] described a PTAS for TSPN with disjoint fat neighborhoods of about the same size in \mathbb{R}^d , where d is constant (this includes the case of disjoint unit disks in the plane). Earlier Dumitrescu and Mitchell [16] proposed a PTAS for TSPN with fat neighborhoods of about the same size and bounded depth in the plane, where Spirkl [44] recently reported and filled a gap; see also a follow-up note in [36].

Mata and Mitchell [29] gave a $O(\log n)$ -approximation for TSPN with n connected and arbitrarily intersecting neighborhoods in the plane; see also [7]. Elbassioni et al. [19] and Gudmundsson and Levcopoulos [24] improved the running time of the algorithm. The $O(\log n)$ -approximation relies on the following early result by Levcopoulos and Lingas [28]: Every (simple) rectilinear polygon P with n vertices, r of which are reflex, can be partitioned in $O(n \log n)$ time into rectangles whose total perimeter is $\log r$ times the perimeter of P .

Using an approximation algorithm due to Slavik [43] for Euclidean group TSP, de Berg et al. [6] obtained constant-factor approximations for disjoint fat convex regions (of arbitrary diameters) in the plane. Subsequently, Elbassioni et al. [19] gave constant-factor approximations for arbitrarily intersecting fat convex regions of comparable size. Preliminary work by Mitchell [34] gave a PTAS for planar regions of bounded depth and arbitrary size, in particular for disjoint fat regions. Chan and Jiang [12] gave a PTAS for fat, weakly disjoint regions in metric spaces of constant doubling dimension (combining an earlier QPTAS by Chan and Elbassioni [11] with a PTAS for TSP in doubling metrics by Bartal et al. [4]).

Disks and balls are undoubtedly among the simplest neighborhood types [2, 16, 25]. TSPN for disks is NP-hard, and it remains so for congruent disks, since when the disk centers are fixed and the radius tends to zero, the problem reduces to TSP for points. Regarding approximations, the case of congruent balls is relatively well understood: Given a set of n congruent (say, unit) balls in \mathbb{R}^d , a TSP tour whose length is at most $O(1)$ times the optimal can be computed in polynomial time, when d is constant [18]. However, for disks of arbitrary radii and intersections, no constant-ratio approximation was known. Some of the difficulties with disks of arbitrary radii in the plane were uncovered in [17].

Recent work of Dumitrescu and Tóth [18] focused on unbounded neighborhoods, such as lines or hyperplanes: They gave a constant-factor approximation for TSPN with n hyperplanes in \mathbb{R}^d in $O(n)$ time; and a $O(\log^3 n)$ -approximation for n lines in \mathbb{R}^d in time polynomial in n , where d is constant. In contrast, the current paper considers TSPN with arbitrary disks in \mathbb{R}^2 , which requires quite different approximation techniques and new ideas.

Degree of approximation. Regarding the degree of approximation achievable, TSPN with arbitrary neighborhoods is generally APX-hard [6, 15, 42], and it remains so even for segments of nearly the same length [19]. For disconnected neighborhoods, TSPN cannot be approximated within any constant ratio unless $P = NP$ [42]. Further, approximating TSPN for (arbitrary) connected neighborhoods in the plane within a factor smaller than 2 is NP-hard [42]. Delineating the class of neighborhoods for which constant-factor approximations are possible remains mysterious, at least at the moment. It is conjectured that approximating TSPN for disconnected regions in the plane within a $O(\log^{1/2} n)$ factor is intractable unless $P = NP$ [42]. Similarly, it is conjectured that approximating TSPN for connected regions in \mathbb{R}^3 within a $O(\log^{1/2} n)$ factor and for disconnected regions in \mathbb{R}^3 within a $O(\log^{2/3} n)$ factor [42] are probably intractable.

Our results. In this paper we present a polynomial-time (deterministic) algorithm that, given a set of n disks in \mathbb{R}^2 (with arbitrary radii and intersections), returns a TSP tour whose length is $O(1)$ times the optimal.

Theorem 1. *Given a set of n disks in the plane, a TSP tour whose length is at most $O(1)$ times the optimal can be computed in time polynomial in n .*

In their seminal paper on TSPN, Arkin and Hassin [2] suggested disks as the most natural type of neighborhood—in which the traveling salesman can meet each potential buyer at some point close to the respective buyer’s location. Here the radius of each disk indicates how much each potential buyer is willing to travel to meet the salesman. While constant-ratio approximations for disks of the same (or comparable) radius [2, 16] and for disjoint disks of arbitrary radii [6] have been obtained early on, the case of disks with arbitrary radii and arbitrary intersections has remained open until now.

2 Preliminaries

We achieve a $O(1)$ -approximation for TSP with disks by reducing the problem, up to constant factors, to a minimum weight hitting set problem in a geometric hypergraph, for which a constant-factor approximation algorithm was found only recently [13].

Hitting sets. Hitting sets are defined in general in terms of hypergraphs (i.e., set systems or range spaces). A hypergraph is a pair $\mathcal{G} = (V, E)$ where V is a finite vertex set and $E \subset 2^V$ is a finite collection of subsets of V (called edges). In a geometric (primal) hypergraph, the vertex set V is a finite set of n points in Euclidean space \mathbb{R}^d , and all sets in E are of the form $V \cap Q$ where Q is a certain geometric shape of bounded description complexity, e.g., halfspace, ball, triangle, axis-aligned rectangle, etc. Geometric hypergraphs often have nice properties such as bounded VC-dimension or bounded union complexity; see [20, 39].

A *hitting set* in a hypergraph $\mathcal{G} = (V, E)$ is a subset of vertices $H \subseteq V$ such that every hyperedge in E contains some point in H . The *minimum hitting set* (MHS) problem asks for a hitting set of minimum cardinality in a given hypergraph. The *minimum weight hitting set* (MWHS) problem asks for a hitting set of minimum weight in a given hypergraph with vertex weights $w : V \rightarrow \mathbb{R}^+$.

Brönnimann and Goodrich [9] gave a $O(\log \text{OPT})$ -approximation for MHS in geometric hypergraphs using LP-relaxations and the fact that geometric hypergraphs have bounded VC-dimension. Clarkson and Varadarajan [14] gave a $O(\log \log n)$ -approximation for *some* geometric hypergraphs, by observing a connection between hitting sets and the combinatorial complexity of the union of the corresponding geometric objects. Mustafa and Ray [38] gave a PTAS for MHS with disks and pseudo-disks in the plane using a local search paradigm; see also [1, 10]. However, this method does not seem to extend to the weighted version (MWHS).

Varadarajan [45] gave a $O(\log \log n)$ -approximation for MWHS, extending the results from [14]. His approach was further extended by Chan et al. [13] who obtained a randomized polynomial-time $O(1)$ -approximation algorithm for MWHS in geometric hypergraphs of linear union complexity, including geometric hypergraphs defined by disks in \mathbb{R}^2 [26]; their algorithm can be derandomized [13, Section 3]. Specifically, we use the following result due to Chan et al. [13].

Theorem 2. [13, Corollary 1.4 and Section 3] *There is a polynomial-time (deterministic) $O(1)$ -approximation algorithm for the minimum weight hitting set problem for disks in \mathbb{R}^2 .*

Definitions. Let \mathcal{R} be a set of regions (neighborhoods) in \mathbb{R}^2 . An optimal TSP tour for \mathcal{R} , denoted by $\text{OPT}(\mathcal{R})$, is a shortest closed curve in the plane that intersects every region in \mathcal{R} ; when \mathcal{R} is clear from the context, $\text{OPT}(\mathcal{R})$ and OPT are used interchangeably.

The Euclidean length of a curve γ is denoted by $\text{len}(\gamma)$. Similarly, the total (Euclidean) length of the edges of a geometric graph G is denoted by $\text{len}(G)$. The perimeter of a polygon P is denoted by $\text{per}(P)$; the boundary and the interior of a region R are denoted by ∂R and R° , respectively; the convex hull of a planar set S is denoted by $\text{conv}(S)$.

The distance between two planar point sets $S_1, S_2 \subset \mathbb{R}^2$, is $\text{dist}(S_1, S_2) = \inf\{\text{dist}(s_1, s_2) : s_1 \in S_1, s_2 \in S_2\}$. The distance between a point set S_1 and a geometric graph G is defined as $\text{dist}(S_1, G) := \text{dist}(S_1, S_G)$, where S_G is the set of all points at vertices and on the edges of G .

Algorithm Outline. Given a set S of n disks in Euclidean plane, we construct a connected geometric graph G that intersects every disk in S and such that $\text{len}(G) = O(\text{len}(\text{OPT}))$. An Eulerian tour of the multi-graph obtained by doubling each edge of G visits each disk and its length is $2\text{len}(G) = O(\text{len}(\text{OPT}))$, as desired.

The graph G is the union of three geometric graphs, G_1 , G_2 and G_3 . The graph G_1 is a $O(1)$ -approximation of an optimal tour for a maximal subset of pairwise disjoint disks in S ; this step is based on earlier results [6, 19] (Section 3). The graph G_2 connects G_1 to nearby disks that are guaranteed to be at distance at most $\text{len}(\text{OPT})/n$ from G_1 . The graph G_3 connects any remaining disks to G_1 ; this step is based on recent results on minimum weight hitting sets due to Chan et al. [13] (Section 5).

The interface between TSP and the hitting set problem is established by a quadtree subdivision [5, Ch. 14]. Previously, Arora [3] and Mitchell [35] used quadtrees for approximating Euclidean TSP and TSP with disjoint neighborhoods, respectively. The quadtree variety that we need, a so-called *stratified grid*, was introduced by Mitchell [35] for certain orthogonal polygons. Here we define stratified grids in a more general setting, for arbitrary geometric graphs (Section 4).

3 Preprocessing

Let S be a set of n disks in the plane. The algorithm first constructs the graphs G_1 and G_2 as follows (Fig. 1).

1. Select an *independent subset* I , $I \subseteq S$, of pairwise disjoint disks by the following greedy algorithm: Set $I := \emptyset$. Consider the disks in S in increasing order of radius (with ties broken arbitrarily), and successively place a disk $D \in S$ into I if it is disjoint from all previous disks in I .
2. Compute a constant-factor approximate TSP tour ξ_0 for I using the algorithm in [6] or [19]. (A PTAS for disjoint disks in the plane is available [16, 44] but not needed here.) It is clear that $\text{len}(\xi_0) = O(\text{len}(\text{OPT}))$.
3. Let R be a minimum axis-parallel square such that $\text{conv}(R)$ intersects every disk in S (i.e., every disk intersects ∂R or is contained in R°). The square R is determined by up to 3 disks in S , thus R can be trivially computed in $O(n^4)$ time: there are $O(n^3)$ squares that pairs and triples define, and each can be checked in $O(n)$ time as to whether it intersects all disks. Alternatively, finding R is an LP-type problem of combinatorial dimension 3 that can be solved in $O(n)$ time [30][Section 5]; see also [23] for a modern treatment of LP-type problems and violator spaces. Let r denote its side-length of R ; obviously, we have $r \leq \text{len}(\text{OPT})$.
4. Let G_1 be the union of ξ_0 , R , and a shortest line segment connecting ξ_0 and R (if disjoint).

The graph G_1 intersects all disks in I and possibly some disks in $S \setminus I$. Our primary interest is in the disks in S that are disjoint from G_1 .




Figure 1: Left: a set of disks in the plane; the independent set I selected by a greedy algorithm is highlighted; a TSP tour ξ_0 for I , and a minimum square R intersecting all disks are shown. Middle: G_1 is the union of ξ_0 , R , and a shortest line segment connecting them; only the disks in $S_2 \cup S_3$ that do not intersect G_1 are shown. Right: a stratified grid for R and G_1 .

Lemma 1. *For every disk $D \in S$, we have $\text{dist}(D, G_1) \leq \text{diam}(D)$.*

Proof. Let $D \in S$. If G_1 intersects D , then $\text{dist}(D, G_1) = 0$, and the claim is trivial. Assume that D is disjoint from G_1 . Since ξ_0 intersects every disk in I and $\xi_0 \subseteq G_1$, we have $D \in S \setminus I$. By the greedy choice of $I \subseteq S$, the disk D intersects some disk $D' \in I$ of equal or smaller radius, where G_1 intersects D' . Consequently, $\text{dist}(D, G_1) \leq \text{diam}(D') \leq \text{diam}(D)$. \square

Connecting nearby disks to G_1 . We partition S into three subsets: let S_1 be the set of disks in S that intersect G_1 ; let S_2 be the set of disks $D \in S \setminus S_1$ such that $\text{dist}(D, G_1) \leq \frac{r}{n}$; and let $S_3 = S \setminus (S_1 \cup S_2)$. Let G_2 be a graph that consists of $|S_2|$ line segments: specifically for every $D \in S_2$, G_2 contains a shortest segment connecting D and G_1 . Then $\text{len}(G_2) = \sum_{D \in S_2} \text{dist}(D, G_1) \leq |S_2| \cdot \frac{r}{n} \leq r \leq \text{len}(\text{OPT})$. By construction, we have

$$\text{dist}(D, G_1) > \frac{r}{n} \text{ for every } D \in S_3. \quad (1)$$

By Lemma 1 and inequality (1) we have

Corollary 1. *For every disk $D \in S_3$, we have $\text{diam}(D) > \frac{r}{n}$.*

In the next section, we show how to find a geometric graph G_3 such that G_3 intersects every disk in S_3 and $G_1 \cup G_3$ is connected (note, however, that G_3 need not be connected).

4 Stratified Grids

Recall that we have a geometric graph G_1 , and a set S_3 of at most n disks in the interior of an axis-aligned square R of side-length r , $r \leq \text{len}(\text{OPT})$, satisfying (1). Let $\text{OPT}(S_3, G_1)$ denote a geometric graph Γ of minimum length such that $G_1 \cup \Gamma$ is connected and intersects every disk in S_3 . Note that $\text{len}(\text{OPT}(S_3, G_1)) \leq \text{len}(\text{OPT}(S_3)) \leq \text{len}(\text{OPT})$, for every G_1 .

In Sections 5 and 6, we use hitting sets to compute a $O(1)$ -approximation of $\text{OPT}(S_3, G_1)$. Similarly to a quadtree decomposition, we recursively construct a subdivision of R into squares of side-lengths $r/2^i$, for $i = 0, 1, \dots, \lceil \log n \rceil$. Refer to Fig. 1 (right).

Previously, Mitchell [35] used a similar quadtree decomposition for TSPN with disjoint regions in the plane, coined the term “stratified grid,” and derived several basic properties of quadtrees that we rederive here. Specifically, he proved analogues of Lemmas 2 and 5 for the problem studied in [35]. However, Mitchell used stratified grids only for special types of orthogonal polygons, called *histograms* [27]; here we generalize this tool to arbitrary geometric graphs.

The following algorithm subdivides a square Q unless it is too small (i.e., $\text{diam}(Q) < \frac{r}{2n}$) or it is relatively far from G_1 (i.e., $\text{diam}(Q) < \text{dist}(Q, G_1)$).

Stratify(R, G_1). Let L be a FIFO queue and \mathcal{Q} be a set of axis-aligned squares. Set $L = (R)$ and $\mathcal{Q} = \emptyset$. Repeat the following while L is nonempty. Set $Q \leftarrow \text{dequeue}(L)$. If $\text{diam}(Q) \geq \max(\frac{r}{2n}, \text{dist}(Q, G_1))$, then subdivide Q into four congruent axis-aligned squares, and enqueue them onto L . Otherwise, let $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{Q\}$. Return \mathcal{Q} .

It is worth noting that \mathcal{Q} does not directly depend on the disks in S_3 , but only indirectly, via G_1 . By construction, the squares in \mathcal{Q} are interior-disjoint, and every square in \mathcal{Q} has diameter at least $r/(4n)$. Consequently, the number of squares in \mathcal{Q} is $O(n^2)$. Thus the algorithm **Stratify**(R, G_1) runs in polynomial time in n , since $O(n^2)$ squares are enqueued onto L , and $\text{dist}(Q, G_1)$ can be computed in polynomial time for all $Q \in L$. We show that the squares in \mathcal{Q} have a property similar to the disks in S_3 (cf. Lemma 1): only larger squares can be farther from G_1 .

Lemma 2. *For every square $Q \in \mathcal{Q}$, we have $\text{dist}(Q, G_1) \leq 3 \text{diam}(Q)$.*

Proof. Put $q = \text{diam}(Q)$. Recall that Q is obtained by subdividing a square Q' , $Q \subset Q'$, with $\text{diam}(Q') = 2q$. Since Q' is subdivided by the algorithm, we have $\text{diam}(Q') \geq \max(\frac{r}{2n}, \text{dist}(Q', G_1))$. Since $\text{dist}(p', Q) \leq q$ for every point $p' \in Q'$, the triangle inequality yields $\text{dist}(Q, G_1) \leq 3q$. \square




Figure 2: Left: a square Q of the stratified grid, and the graph $\gamma(Q)$. Right: a polygonal curve α ; the intersections of α with horizontal (resp., vertical) edges of the stratified grid are marked with empty (resp., full) dots.

For every square $Q \in \mathcal{Q}$, we define a graph $\gamma(Q)$ that consists of the boundary of Q and a shortest line segment from Q to G_1 ; see Fig. 2 (left). By Lemma 2, we have $\text{len}(\gamma(Q)) \leq (3 + 2\sqrt{2}) \text{diam}(Q)$; on the other hand, $\text{len}(\gamma(Q)) \geq \text{per}(Q) = 2\sqrt{2} \text{diam}(Q)$, and so we have the following.

Corollary 2. *For every $Q \in \mathcal{Q}$, we have $\text{len}(\gamma(Q)) = \Theta(\text{diam}(Q))$.*

The following observation is crucial for reducing the problem of approximating $\text{OPT}(S_3, G_1)$ to a minimum weight hitting set problem.

Lemma 3. *If a square $Q \in \mathcal{Q}$ intersects a disk $D \in S_3$, then*

- (i) $\text{diam}(Q) \leq 2 \text{diam}(D)$, and
- (ii) D intersects the boundary of Q (and the graph $\gamma(Q)$ in particular).

Proof. (i) Since $Q \in \mathcal{Q}$, Algorithm **Stratify** (R, G_1) did not subdivide Q , and so we have $\text{diam}(Q) < \frac{r}{2n}$ or $\text{diam}(Q) < \text{dist}(Q, G_1)$. If $\text{diam}(Q) < \frac{r}{2n}$, then Corollary 1 yields

$$\text{diam}(Q) < \frac{r}{2n} < \frac{r}{n} < \text{diam}(D) < 2 \text{diam}(D).$$

If $\text{diam}(Q) < \text{dist}(Q, G_1)$, then $\text{dist}(Q, G_1) \leq \text{dist}(D, G_1) + \text{diam}(D)$ follows from the intersection condition and the triangle inequality. Consequently,

$$\text{diam}(Q) < \text{dist}(Q, G_1) \leq \text{dist}(D, G_1) + \text{diam}(D) \leq 2 \text{diam}(D),$$

where the last inequality holds by Lemma 1.

(ii) Suppose, to the contrary, that the boundary of Q is disjoint from D , hence D lies in the interior of Q . This immediately implies

$$\text{dist}(Q, G_1) \leq \text{dist}(D, G_1). \tag{2}$$

Since $Q \in \mathcal{Q}$, Algorithm **Stratify** (R, G_1) did not subdivide Q , and so we have $\text{diam}(Q) < \frac{r}{2n}$ or $\text{diam}(Q) < \text{dist}(Q, G_1)$. If $\text{diam}(Q) < \frac{r}{2n}$, Corollary 1 yields $\text{diam}(Q) < \frac{r}{2n} < \frac{r}{n} < \text{diam}(D)$. If $\text{diam}(Q) < \text{dist}(Q, G_1)$, then the combination of (2) and Lemma 1 yield

$$\text{diam}(Q) < \text{dist}(Q, G_1) \leq \text{dist}(D, G_1) < \text{diam}(D).$$

In both cases, we have shown that $\text{diam}(Q) < \text{diam}(D)$. Therefore D cannot lie in the interior of Q , which contradicts the assumption. \square

Recall that the squares in \mathcal{Q} can only intersect at common boundary points; we call such squares *adjacent*.

Lemma 4. *If two squares $Q_1, Q_2 \in \mathcal{Q}$ are adjacent and $\text{diam}(Q_1) \leq \text{diam}(Q_2)$, then*

$$\frac{1}{2} \text{diam}(Q_2) \leq \text{diam}(Q_1) \leq \text{diam}(Q_2) \leq 2 \text{diam}(Q_1).$$

Proof. If $\text{diam}(Q_1) = \text{diam}(Q_2)$, the inequalities are satisfied. We may thus assume that $\text{diam}(Q_1) < \text{diam}(Q_2) = q$. Then Algorithm **Stratify** (R, G_1) subdivided a square Q'_1 such that $Q_1 \subset Q'_1$, $\text{diam}(Q'_1) \leq q$, and $Q'_1 \cap Q_2 \neq \emptyset$. The algorithm subdivided Q'_1 but did not subdivide Q_2 . This implies

$$q \geq \max\left(\frac{r}{2n}, \text{dist}(Q'_1, G_1)\right) \text{ and } q < \max\left(\frac{r}{2n}, \text{dist}(Q_2, G_1)\right).$$

The first inequality yields $q \geq \max(\frac{r}{2n}, \text{dist}(Q'_1, G_1)) \geq \frac{r}{2n}$, and then the second inequality yields $q < \max(\frac{r}{2n}, \text{dist}(Q_2, G_1)) = \text{dist}(Q_2, G_1)$. Consequently, $\text{dist}(Q'_1, G_1) \leq q < \text{dist}(Q_2, G_1)$.

Since Q'_1 and Q_2 intersect, and their diameters are at most q , the triangle inequality yields $|\text{dist}(Q'_1, G_1) - \text{dist}(Q_2, G_1)| \leq q$. It follows that

$$q < \text{dist}(Q_2, G_1) \leq \text{dist}(Q'_1, G_1) + q \leq 2q.$$

Similarly, since Q_1 and Q_2 intersect, $\text{dist}(Q_1, G_1) \geq \text{dist}(Q_2, G_1) - \text{diam}(Q_1) > q - \text{diam}(Q_1)$. Combining with Lemma 2, we get

$$3 \text{diam}(Q_1) \geq \text{dist}(Q_1, G_1) > q - \text{diam}(Q_1),$$

that is, $\text{diam}(Q_1) > q/4$. Finally, recall that the ratio between the diameters of any two squares in \mathcal{Q} is a power of 2. Therefore $q/4 < \text{diam}(Q_1) < q$ yields $\text{diam}(Q_1) = q/2$, as required. \square

5 Hitting Sets for Squares and Disks

For the graph G_1 and the set of disks S_3 , we define a hypergraph $\mathcal{G} = (\mathcal{Q}, E)$, where the vertex set is the set \mathcal{Q} of squares in the stratified grid; and for every disk $D \in S_3$, the set of squares in \mathcal{Q} that intersect D forms a hyperedge in E . Thus, a subset $\mathcal{H} \subseteq \mathcal{Q}$ of squares is a *hitting set* in the hypergraph \mathcal{G} if and only if every disk in S_3 intersects some square in \mathcal{H} .

For every hitting set \mathcal{H} , the geometric graph $\Gamma = \cup_{Q \in \mathcal{H}} \gamma(Q)$ intersects every disk in S_3 by Lemma 3, and $G_1 \cup \Gamma$ is connected by construction. Let the *weight* of a square $Q \in \mathcal{Q}$ be $w(Q) = \text{diam}(Q)$. In this section, we show that the minimum-weight hitting set for $\mathcal{G} = (\mathcal{Q}, E)$ is a $O(1)$ -approximation for $\text{OPT}(S_3, G_1)$. The following technical lemma considers a single curve (i.e., a Jordan arc). For a curve α , let $\mathcal{Q}(\alpha)$ denote the set of squares in \mathcal{Q} that intersect α . Refer to Fig. 2 (right).

Lemma 5. *Let α be a directed polygonal curve whose start and end points lie on G_1 . If α intersects at least one disk in S_3 , then $\text{len}(\alpha) = \Omega(\sum_{Q \in \mathcal{Q}(\alpha)} \text{diam}(Q))$.*

Proof. By (1), we have $\text{dist}(D, G_1) > \frac{r}{n}$ for every disk $D \in S_3$. Since α intersects at least one disk in S_3 , we have $\text{len}(\alpha) \geq \frac{2r}{n}$.

Let $A = (Q_0, Q_1, \dots, Q_m)$ be the sequence of distinct squares that intersect α in the order in which they are first encountered by α (with no repetitions and ties broken arbitrarily). Since Q_0 intersects G_1 , we have $\text{diam}(Q_0) < \max(\frac{r}{2n}, 0) = \frac{r}{2n}$, and consequently

$$\text{len}(\alpha) \geq \frac{2r}{n} \geq 4 \text{diam}(Q_0). \quad (3)$$

Let $B = (Q_{\sigma(0)}, Q_{\sigma(1)}, \dots, Q_{\sigma(\ell)})$ be the subsequence of A such that $\sigma(0) = 0$ and a square Q_i , $1 \leq i \leq m$, is added to B if it is disjoint from Q_j for all $0 \leq j < i$. By construction, B consists of pairwise disjoint squares, and every square in A is either in B or adjacent to some square in B . By Lemma 4, the sizes of adjacent squares in \mathcal{Q} differ by a factor of at most 2. Consequently, each square in \mathcal{Q} is adjacent to at most 12 squares in \mathcal{Q} (at most two along each side and at most one at each corner). It follows that

$$\sum_{i=0}^m \text{diam}(Q_i) = \Theta \left(\sum_{j=0}^{\ell} \text{diam}(Q_{\sigma(j)}) \right). \quad (4)$$

For $j = 0, \dots, \ell$, let $p_{\sigma(j)}$ be the first intersection point of α with $Q_{\sigma(j)}$. For two points $p, q \in \alpha$, denote by $\alpha(p, q)$ the portion of α between p and q . Since the squares in B are pairwise disjoint, and the sizes of adjacent squares differ by at most a factor of 2 (Lemma 4), we have

$$\text{len}(\alpha(p_{\sigma(j)}, p_{\sigma(j+1)})) \geq |p_{\sigma(j)} p_{\sigma(j+1)}| \geq \frac{1}{2\sqrt{2}} \max(\text{diam}(Q_{\sigma(j)}), \text{diam}(Q_{\sigma(j+1)}))$$

for $j = 0, \dots, \ell - 1$. Consequently, if $\ell \geq 1$, we have

$$\text{len}(\alpha) = \sum_{j=0}^{\ell-1} \text{len}(\alpha(p_{\sigma(j)}, p_{\sigma(j+1)})) = \Omega \left(\sum_{j=0}^{\ell} \text{diam}(Q_{\sigma(j)}) \right). \quad (5)$$

The combination of (3), (4), and (5) yields $\text{len}(\alpha) = \Omega(\sum_{i=0}^m \text{diam}(Q_i))$, as required. \square

Lemma 6. *If Γ is a geometric graph such that Γ intersects every disk in S_3 and $G_1 \cup \Gamma$ is connected, then there is a hitting set $\mathcal{H} \subseteq \mathcal{Q}$ for $\mathcal{G} = (\mathcal{Q}, E)$ such that*

$$\text{len}(\Gamma) = \Omega \left(\sum_{Q \in \mathcal{H}} \text{diam}(Q) \right). \quad (6)$$

Proof. Let \mathcal{H} be the set of squares in \mathcal{Q} that intersect Γ , and observe that \mathcal{H} is a hitting set for \mathcal{G} . For each connected component C of Γ , let α be a directed polygonal curve that starts and ends at some points in $G_1 \cap C$ and traverses every edge of C at least once and at most twice. Then $\text{len}(C) \geq \frac{1}{2} \text{len}(\alpha)$, and $\text{len}(\alpha) = \Omega(\sum_{Q \in \mathcal{Q}(\alpha)} \text{diam}(Q))$ by Lemma 5. Summation over all the components of Γ yields (6). \square

Recall that $\text{OPT}(S_3, G_1)$ is a geometric graph Γ of minimum length such that $G_1 \cup \Gamma$ is connected and intersects every disk in S_3 . Let W_0 denote the minimum weight of a hitting set in the hypergraph \mathcal{G} . The main result of this section is the following.

Corollary 3. *We have $W_0 = O(\text{len}(\text{OPT}(S_3, G_1)))$.*

Proof. Invoke Lemma 6 with $\Gamma = \text{OPT}(S_3, G_1)$. Then \mathcal{G} has a hitting set $\mathcal{H} \subset \mathcal{Q}$ of weight $\sum_{Q \in \mathcal{H}} \text{diam}(Q) = O(\text{len}(\Gamma)) = O(\text{len}(\text{OPT}(S_3, G_1)))$. This is clearly an upper bound on the minimum weight W_0 of a hitting set in \mathcal{G} . \square

6 Hitting Sets for Points and Disks

In Section 5 we defined a hypergraph $\mathcal{G} = (\mathcal{Q}, E)$ for squares and disks; that is, the vertices are squares in \mathcal{Q} and the hyperedges are the squares intersecting a disk in S_3 . In order to apply Theorem 2 by Chan et al. [13], we reduce the problem to a traditional geometric hypergraph problem, where the vertices are points in \mathbb{R}^2 and a hyperedge corresponds to the set of points contained in a disk $D \in S_3$.

For each square $Q \in \mathcal{Q}$, we define a set of 25 *sentinel* points, and show (Lemma 7) that if a disk $D \in S_3$ intersects Q , then D contains one of the sentinel points of Q . A constant number of sentinels suffice if none of the disks intersecting Q is too small, and indeed, Lemma 3 has shown that this is the case.

For a square $Q \in \mathcal{Q}$, where $Q = [a, a+h] \times [b, b+h]$, let the 25 sentinel points be $(a+ih/2, b+jh/2)$ for all $i, j \in \{-1, 0, 1, 2, 3\}$; see Fig. 3 (left).




Figure 3: Left: the set of 25 sentinel points for a square Q . Right: a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$.

Lemma 7. *If a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$, then D contains a sentinel point corresponding to Q .*

Proof. Assume that a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$ of side length h ; refer to Fig. 3 (right). By Lemma 3, $\text{diam}(Q) \leq 2 \text{diam}(D)$. By scaling down D from an arbitrary center in $D \cap Q$, we find a disk D' intersecting Q with $\text{diam}(D') = \frac{1}{2} \text{diam}(Q)$. The inscribed axis-aligned square Q' of D' has $\text{diam}(Q') = \frac{1}{2} \text{diam}(Q)$. That is, the side-length of Q' is $h/2$, and $\text{dist}(Q', Q) \leq (\sqrt{2} - 1)h$. Since the sentinels of Q form a section of a square lattice of (the same) side-length $h/2$, within distance $\frac{\sqrt{2}}{2}h$ from Q , some sentinel of Q lies in Q' , and hence in $D' \subseteq D$, as claimed. \square

We define a new weighted hypergraph $\mathcal{G}' = (V', E')$, where V' is the union of sentinel point sets for all $Q \in \mathcal{Q}$ that lie in R (sentinels in the exterior of R are discarded); and each hyperedge in E' is the set of sentinels in V' contained in a disk $D \in S_3$. Note that a sentinel $s \in V'$ may correspond to several squares in \mathcal{Q} . Let the *weight* of a sentinel $s \in V'$ be the sum of the diameters of the squares $Q \in \mathcal{Q}$ that correspond to s . Hence the total weight of all sentinels is at most $25 \sum_{Q \in \mathcal{Q}} \text{diam}(Q)$. We next derive a bound on the weight of each sentinel.

Lemma 8. *For every $Q \in \mathcal{Q}$, the weight of every sentinel corresponding to Q is $O(\text{diam}(Q))$.*

Proof. By Lemma 4, the side-lengths of adjacent squares of the stratified grid differ by at most a factor of 2. Consequently, every sentinel in V' corresponding to a square $Q \in \mathcal{Q}$ is contained in Q or in a square of \mathcal{Q} adjacent to Q .

Let $s \in V'$ be a sentinel. Then s may correspond to all squares in \mathcal{Q} that contain s , and to adjacent squares in \mathcal{Q} . Every point is contained in at most 4 squares of \mathcal{Q} , whose side-lengths differ by a factor of at most 2; and they are each adjacent to $O(1)$ additional squares whose side-lengths differ by another factor of at most 2. Overall, s corresponds to $O(1)$ squares in \mathcal{Q} whose side-lengths differ by a factor of $\Theta(1)$. Therefore, the weight of s is $O(\text{diam}(Q))$ for every square $Q \in \mathcal{Q}$ corresponding to s . \square

By Theorem 2, there is a polynomial-time $O(1)$ -approximation algorithm for MWHS on $\mathcal{G}' = (V', E')$. It remains to show that a $O(1)$ -approximation for MWHS on the hypergraph $\mathcal{G}' = (V', E')$ provides a $O(1)$ -approximation for MWHS on the hypergraph $\mathcal{G} = (\mathcal{Q}, E)$.

Lemma 9. *1. For every hitting set $\mathcal{H} \subseteq \mathcal{Q}$ for \mathcal{G} , the set H' of sentinels in V' corresponding to the squares $Q \in \mathcal{H}$ is a hitting set for \mathcal{G}' of weight $O(\sum_{Q \in \mathcal{H}} \text{diam}(Q))$.*

2. For every hitting set $H' \subseteq V'$ for \mathcal{G}' , the set \mathcal{H} of squares $Q \in \mathcal{Q}$ that contain the sentinel points in H' is a hitting set for \mathcal{G} of weight $O(\sum_{s \in H'} w(s))$.

Proof. (1) If \mathcal{H} is a hitting set for \mathcal{G} , then every disk $D \in S_3$ intersects some square $Q \in \mathcal{H}$. By Lemma 7, D contains one of the sentinels of Q . Consequently, every disk $D \in S_3$ contains a sentinel in H' . Every square $Q \in \mathcal{H}$ corresponds to 25 sentinels, each of weight $O(\text{diam}(Q))$ by Lemma 8. The weight of H' is $O(\sum_{Q \in \mathcal{H}} \text{diam}(Q))$.

(2) If H' is a hitting set for \mathcal{G}' , then every disk $D \in S_3$ contains some point $s \in H'$. The point s lies in a square $Q \in \mathcal{Q}$ of the stratified grid, which is in \mathcal{H} . Consequently, every disk $D \in S_3$ intersects some square $Q \in \mathcal{H}$. By construction, the weight of each sentinel s is the sum of weights of the corresponding squares in \mathcal{Q} , including all squares in \mathcal{Q} that contain s . Therefore, the weight of \mathcal{H} is at most $\sum_{s \in H'} w(s)$, as required. \square

We are now ready to prove Theorem 1 by analyzing the constructed graph $G = G_1 \cup G_2 \cup G_3$.

Proof of Theorem 1. Let S be a set of n disks in \mathbb{R}^2 . Compute an independent set $I \subset S$ as described in Section 3, and a TSP tour ξ_0 for I with $\text{len}(\xi_0) = O(\text{len}(\text{OPT}))$ (as in [6] or [19]). Compute the graph G_1 with $\text{len}(G_1) = O(\text{len}(\text{OPT}))$, and the partition $S = S_1 \cup S_2 \cup S_3$ as described in Section 3. The graph G_1 intersects the disks in S_1 .

Construct the graph G_2 , which contains a shortest segment between G_1 and every disk $D \in S_2$. The length of this graph is $\text{len}(G_2) = \sum_{D \in S_2} \text{dist}(D, G_1) \leq |S_2| \cdot \frac{r}{n} \leq r \leq \text{len}(\text{OPT})$.

Compute the stratified grid \mathcal{Q} , and construct the weighted hypergraph $\mathcal{G}' = (V', E')$, where V' is the set of sentinel points for all squares $Q \in \mathcal{Q}$, the weight of a sentinel s is the sum of diameters of the corresponding squares $Q \in \mathcal{Q}$, and for every disk $D \in S_3$, the set of sentinels lying in D forms a hyperedge in E' . Use the algorithm by Chan et al. [13] to compute a hitting set H' for \mathcal{G}' whose weight is $O(1)$ times the minimum. Let \mathcal{H} be the set of squares in \mathcal{Q} containing the sentinels in H' . By Lemma 9, \mathcal{H} is a hitting set for the hypergraph $\mathcal{G} = (\mathcal{Q}, E)$ whose weight is at most $O(1)$ times the minimum W_0 . Put $G_3 = \cup_{Q \in \mathcal{H}} \gamma(Q)$. Then G_3 intersects every disk in S_3 by Lemma 3, and $\text{len}(G_3) = \Theta(\sum_{Q \in \mathcal{H}} \text{diam}(Q)) = \Theta(W_0)$ by Corollary 2. Finally, Corollary 3 yields $\text{len}(G_3) = O(\text{len}(\text{OPT}(S_3, G_1)))$. By the definition of $\text{OPT}(S_3, G_1)$, we have $\text{len}(\text{OPT}(S_3, G_1)) \leq \text{len}(\text{OPT}(S_3)) \leq \text{len}(\text{OPT})$, and consequently $\text{len}(G_3) = O(\text{len}(\text{OPT}))$.

Note that the graph $G = G_1 \cup G_2 \cup G_3$ is connected by construction, it intersects every disk in $S = S_1 \cup S_2 \cup S_3$, and $\text{len}(G) = \text{len}(G_1) + \text{len}(G_2) + \text{len}(G_3) = O(\text{len}(\text{OPT}))$. Consequently, an Eulerian tour of the multi-graph containing each edge of G twice visits each disk and its length is $2\text{len}(G) = O(\text{len}(\text{OPT}))$, as required.

Since the above steps as well as algorithm **Stratify** (R, G_1) all run in time that is polynomial in n , the constant-factor approximation algorithm for TSP with disks runs in polynomial time. \square

7 Conclusions

In this paper, we obtained the first constant-ratio approximation for TSP with disks in the plane. This is the first result of this kind for a class of planar convex bodies of arbitrary size that can intersect in an arbitrary fashion. In light of the connection we established between TSPN and MWHS in geometric hypergraphs, the following question emerges:

1. Besides regions of linear union complexity (e.g., disks and pseudo-disks¹), what other types of regions admit a constant-factor approximation for the minimum weight hitting set problem?

Obviously, a constant-factor approximation for MWHS with a certain type of neighborhoods does not automatically imply a constant-factor approximation for TSPN with the same type of neighborhoods. We conclude with a few, perhaps the simplest still unsolved questions on TSPN that we could identify:

2. Is there a constant-factor approximation algorithm for TSP with a set of objects of linear union complexity, e.g., pseudo-disks?
3. Is there a constant-factor approximation algorithm for TSP with convex bodies in the plane?²
4. Is TSP with disks in the plane APX-hard? Is TSP with convex bodies in the plane APX-hard?

¹A set of regions \mathcal{R} consists of *pseudo-disks*, if every pair of regions $\omega_1, \omega_2 \in \mathcal{R}$ satisfies the *pseudo-disk property*: the sets $\omega_1 \setminus \omega_2$ and $\omega_2 \setminus \omega_1$ are connected [5, p. 293]. Equivalently, the boundaries $\partial\omega_1$ and $\partial\omega_2$ have at most two proper intersection points.

²Very recently, Mitchell [37] proposed a constant-factor approximation algorithm for this variant. However, no complete proof is available at the time of this writing. In fact, we believe that TSP with planar convex bodies is much harder to approximate than TSP with disks.

5. Is there a constant-factor approximation algorithm for TSP with balls (with arbitrary radii and intersections) in \mathbb{R}^d , in fixed dimension $d \geq 3$?

References

- [1] P. K. Agarwal, E. Ezra, and M. Sharir, Near-linear approximation algorithms for geometric hitting sets, *Algorithmica* **63**(1-2) (2012), 1–25.
- [2] E. M. Arkin and R. Hassin, Approximation algorithms for the geometric covering salesman problem, *Discrete Applied Mathematics* **55**(3) (1994), 197–218.
- [3] S. Arora, Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems, *Journal of the ACM* **45**(5) (1998), 753–782.
- [4] Y. Bartal, L.-A. Gottlieb, and R. Krauthgamer, The traveling salesman problem: low-dimensionality implies a polynomial time approximation scheme, in *Proc. 44th Symposium on Theory of Computing (STOC)*, ACM Press, 2012, pp. 663–672.
- [5] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry*, 3rd edition, Springer, Heidelberg, 2008.
- [6] M. de Berg, J. Gudmundsson, M. J. Katz, C. Levcopoulos, M. H. Overmars, and A. F. van der Stappen, TSP with neighborhoods of varying size, *Journal of Algorithms* **57**(1) (2005), 22–36.
- [7] M. Bern and D. Eppstein, Approximation algorithms for geometric problems, in *Approximation Algorithms for NP-hard Problems (D. S. Hochbaum, ed.)*, PWS Publishing Company, Boston, MA, 1997, pp. 296–345.
- [8] H. L. Bodlaender, C. Feremans, A. Grigoriev, E. Penninx, R. Sitters, and T. Wolle, On the minimum corridor connection problem and other generalized geometric problems, *Computational Geometry: Theory and Applications* **42**(9) (2009), 939–951.
- [9] H. Brönnimann and M. T. Goodrich, Almost optimal set covers in finite VC-dimension, *Discrete & Computational Geometry* **14**(4) (1995), 463–479.
- [10] N. Bus, S. Garg, N. H. Mustafa, and S. Ray, Improved local search for geometric hitting set, in *Proc. 32nd Symposium on Theoretical Aspects of Computer Science (STACS)*, vol. 30 of LIPIcs, Schloss Dagstuhl, 2015, pp. 184–196.
- [11] T.-H. H. Chan and K. Elbassioni, A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics, *Discrete & Computational Geometry* **46**(4) (2011), 704–723.
- [12] T.-H. H. Chan and S. H.-C. Jiang, A PTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics, technical report, e-prints for the Optimization Community, http://www.optimization-online.org/DB_HTML/2014/11/4649.html.
- [13] T.M. Chan, E. Grant, J. Könemann, and M. Sharpe, Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling, in *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 2012, pp. 1576–1585.

- [14] K. L. Clarkson and K. R. Varadarajan, Improved approximation algorithms for geometric set cover, *Discrete & Computational Geometry* **37(1)** (2007), 43–58.
- [15] M. Dror and J. B. Orlin, Combinatorial optimization with explicit delineation of the ground set by a collection of subsets, *SIAM Journal on Discrete Mathematics* **21(4)** (2008), 1019–1034.
- [16] A. Dumitrescu and J. S. B. Mitchell, Approximation algorithms for TSP with neighborhoods in the plane, *Journal of Algorithms* **48(1)** (2003), 135–159.
- [17] A. Dumitrescu and Cs. D. Tóth, On the total perimeter of homothetic convex bodies in a convex container, *Beiträge zur Algebra und Geometrie* **56(2)** (2015), 515–532.
- [18] A. Dumitrescu and Cs. D. Tóth, The traveling salesman problem for lines, balls and planes, *ACM Transactions on Algorithms* **12(3)** (2016), article 43.
- [19] K. M. Elbassioni, A. V. Fishkin, and R. Sitters, Approximation algorithms for the Euclidean traveling salesman problem with discrete and continuous neighborhoods, *International Journal on Computational Geometry & Applications* **19(2)** (2009), 173–193.
- [20] G. Even, D. Rawitz, and S. Shahar, Hitting sets when the VC-dimension is small, *Information Processing Letters* **95(2)** (2005), 358–362.
- [21] M. R. Garey, R. Graham, and D. S. Johnson, Some NP-complete geometric problems, in *Proc. 8th ACM Symposium on Theory of Computing (STOC)*, ACM Press, 1976, pp. 10–22.
- [22] M. R. Garey and D. S. Johnson: *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, New York, 1979.
- [23] B. Gärtner, J. Matoušek, L. Rüst, and P. Škovroň, Violator spaces: Structure and algorithms, *Discrete Appl. Math.* **156(11)** (2008), 2124–2141.
- [24] J. Gudmundsson and C. Levcopoulos, A fast approximation algorithm for TSP with neighborhoods, *Nordic Journal of Computing* **6(4)** (1999), 469–488.
- [25] P. Kamousi and S. Suri, Euclidean traveling salesman tours through stochastic neighborhoods, in *Proc. 24th International Symposium on Algorithms and Computation (ISAAC)*, LNCS 8283, Springer, 2013, pp. 644–654.
- [26] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete & Computational Geometry* **1(1)** (1986), 59–70.
- [27] C. Levcopoulos, Heuristics for Minimum Decompositions of Polygons, PhD thesis, Linköping Studies in Science and Technology, No. 74, 1987.
- [28] C. Levcopoulos and A. Lingas, Bounds on the length of convex partitions of polygons, in *Proc. 4th Conf. on Foundations of Software Technology and Theoretical Computer Science*, LNCS 181, Springer, 1984, pp. 279–295.
- [29] C. Mata and J. S. B. Mitchell, Approximation algorithms for geometric tour and network design problems, in *Proc. 11th ACM Symposium on Computational Geometry (SOCG)*, ACM Press, 1995, pp. 360–369.

- [30] J. Matoušek, M. Sharir, and E. Welzl, A subexponential bound for linear programming, *Algorithmica* **16(4)** (1996), 498–516.
- [31] J. S. B. 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)** (1999), 1298–1309.
- [32] J. S. B. Mitchell, Geometric shortest paths and network optimization, in *Handbook of Computational Geometry (J.-R. Sack, J. Urrutia, eds.)*, Elsevier, 2000, pp. 633–701.
- [33] J. S. B. Mitchell, Shortest paths and networks, in *Handbook of Computational Geometry (J. E. Goodman and J. O’Rourke, eds.)*, Chapman & Hall/CRC, 2004, pp. 607–641.
- [34] J. S. B. Mitchell, A PTAS for TSP with neighborhoods among fat regions in the plane, in *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 2007, pp. 11–18.
- [35] J. S. B. Mitchell, A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane, in *Proc. 26th Symposium on Computational Geometry (SOCG)*, ACM Press, 2010, pp. 183–191.
- [36] J. S. B. Mitchell, Clarifying remarks added to “Approximation algorithms for TSP with neighborhoods in the plane,” manuscript, <http://www.ams.sunysb.edu/~jsbm/papers/tspn-jalgorithms-updated.pdf>, accessed in Feb. 2016.
- [37] J. S. B. Mitchell, Updated version of “A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane,” manuscript, <http://www.ams.sunysb.edu/~jsbm/papers/tspn-socg10-updated.pdf>, accessed in Feb. 2016.
- [38] N. H. Mustafa and S. Ray, Improved results on geometric hitting set problems, *Discrete & Computational Geometry* **44(4)** (2010), 883–895.
- [39] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, John Wiley, New York, 1995.
- [40] C. H. Papadimitriou, Euclidean TSP is NP-complete, *Theoretical Computer Science* **4(3)** (1977), 237–244.
- [41] S. B. Rao and W. D. Smith, Approximating geometrical graphs via “spanners” and “banyans,” in *Proc. 30th ACM Symposium on Theory of Computing (STOC)*, ACM Press, 1998, pp. 540–550.
- [42] S. Safra and O. Schwartz, On the complexity of approximating TSP with neighborhoods and related problems, *Computational Complexity* **14(4)** (2005), 281–307.
- [43] P. Slavik, The errand scheduling problem, CSE Technical Report 97-02, University of Buffalo, Buffalo, NY, 1997.
- [44] S. Spirkl, The guillotine subdivision approach for TSP with neighborhoods revisited, preprint, [arXiv:1312.0378v2](https://arxiv.org/abs/1312.0378v2), 2014.
- [45] K. R. Varadarajan, Epsilon nets and union complexity, in *Proc. 25th Symposium on Computational Geometry (SOCG)*, ACM Press, 2009, pp. 11–16.