Constant-Factor Approximation for TSP with Disks

Adrian Dumitrescu[*]   Csaba D. Tóth[†]

Thursday 15[th] October, 2015

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 with high probability 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 $\mathbb{R}^d$, $d \geq 3$, one seeks a shortest closed curve (a.k.a. *tour*) that visits each point. In the TSP *with neighborhoods* (TSPN), first studied by Arkin and Hassin [2], each point is replaced by a (possibly disconnected) region. The TSP tour must visit at least one point in each of the given regions, i.e., it must intersect each region. Since the Euclidean TSP is known to be NP-hard in $\mathbb{R}^d$ for every $d \geq 2$ [21, 22, 36], TSPN is also NP-hard for every $d \geq 2$. TSP is recognized as one of the corner-stone problems in combinatorial optimization. For related problems in geometric network optimization, the reader is referred to the two surveys [30, 31].

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 [29]. Subsequent running time improvements have been obtained by Rao and Smith [37]; 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 regions with pairwise *parallel* diameters and with the ratio between the longest and the shortest diameter bounded by a constant. Dumitrescu and Mitchell [16] extended the above result to arbitrary connected neighborhoods with comparable diameters. Bodlaender et al. [8] describe a PTAS for TSPN with disjoint fat regions of about the same size (this includes the case of disjoint unit disks), in $\mathbb{R}^d$, where d is constant. Earlier Dumitrescu and Mitchell [16] proposed a PTAS for TSPN with fat regions of about the same size and bounded depth in the plane (where Spirkl [40] recently found and filled a gap).

Mata and Mitchell [28] 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 [23] improved the running time of the algorithm. The $O(\log n)$-approximation relies on the following early result by Levcopoulos and Lingas [27]: 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 [39] 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 fat and arbitrarily intersecting convex regions of comparable size. Preliminary work by Mitchell [32] 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 dimensions (combining an earlier QPTAS by Chan and Elbassioni [11] with a PTAS for TSP in doubling metrics by Bartal et al. [4]).

The disks are undoubtedly among the simplest neighborhood types [2, 16, 24]. 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 where 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 many new ideas.

Degree of approximation. Regarding the degree of approximation achievable, TSPN with arbitrary neighborhoods is generally APX-hard [6, 38], and it remains so even for segments of nearly the same length [19]. The problem is also APX-hard for disconnected regions [38], such as point-pair regions [15]. Further, approximating TSPN for (arbitrary) connected regions in the plane within a factor smaller than 2 is NP-hard [38]. Delineating the class of neighborhoods for which constant 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 [38]. 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 [38] are intractable.

Our results. In this paper we present a polynomial-time randomized 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 with high probability.

Theorem 1. *Given a set of n disks in the plane, a TSP tour whose length is at most $O(1)$ times the optimal with high probability 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 an $O(1)$-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 randomized $O(1)$-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, 35].

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 $w : V \to \mathbb{R}^+$.

Brönnimann and Goodrich [9] gave a $O(\log \mathsf{OPT})$-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 [34] 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 [41] 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 $\mathbb{R}^2$ [25]. Specifically, we use the following result due to Chan et al. [13].

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

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

The graph G is the union of three plane straight-line graphs, G_1, G_2 and G_3. The graph G_1 is an $O(1)$-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 $\text{len}(\mathsf{OPT})/n$ from G_1. The graph G_3 connects any remaining disks to G_1; this step is based on recent results on weighted hitting sets due to Chan et al. [13] (Section 5).

The interface between TSP and the hitting set problem is established by a quadtree [5, Ch. 14]. Previously, Arora [3] and Mitchell [33] 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 [33] for certain orthogonal polygons. We define stratified grids in general for arbitrary planar straight-line graphs (Section 4).

Definitions. Let $\mathcal{R}$ be a set of regions in $\mathbb{R}^d$, $d \geq 2$. A set $\Xi \subset \mathbb{R}^d$ *intersects* $\mathcal{R}$ if Ξ intersects each region in $\mathcal{R}$, that is, $\Xi \cap r \neq \emptyset$, $\forall r \in \mathcal{R}$. An optimal TSP tour for a set $\mathcal{R}$ of regions (neighborhoods), denoted by $\mathsf{OPT}(\mathcal{R})$, is a shortest closed curve in the ambient space that intersects $\mathcal{R}$.

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 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. For a point set $Q \subset \mathbb{R}^2$, we denote by $\text{width}(Q)$ and $\text{height}(Q)$, respectively, the horizontal and vertical extents of the minimum axis-parallel rectangle containing Q.

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).
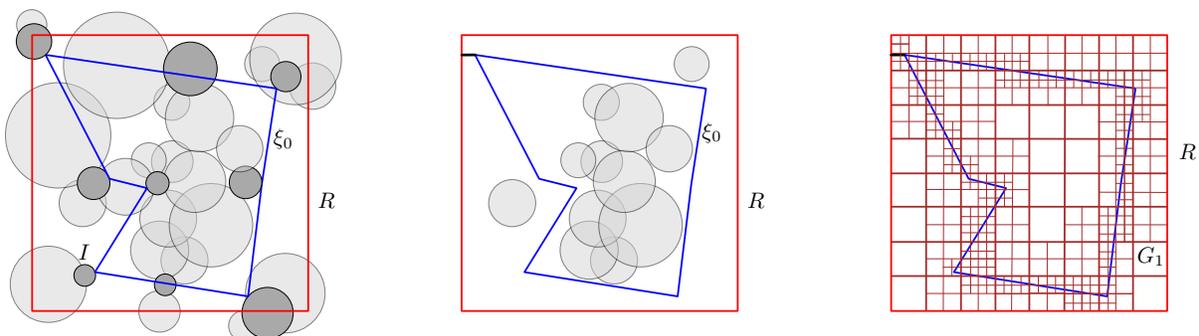

Figure 1: Left: A set of disks in the plane. The independent set I selected by a greedy algorithm is highlighted. A TSP tour ξ_0 for I, and a minimum square R intersecting all disks. Middle: G_1 is the union of ξ_0, R, and a shortest line segment connecting them. Only the disks in $S_2 \cup S_3$ that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

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 sorted by increasing radii (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, 40] but not needed for our purposes). It is clear that $\text{len}(\xi_0) = O(\text{len}(\text{OPT}))$.

3. Let R be a minimum axis-parallel square that intersects every disk in S (i.e., every disk intersects its boundary or is contained in its interior). R can be computed in $O(n)$ time by linear programming. Let r denote its side length; obviously, we have $r \le \text{len}(\text{OPT})$.

4. Let G_1 be the union of ξ_0, the boundary of R, and (if ξ_0 and R are disjoint) a shortest line segment connecting ξ_0 and R.

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.

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

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, $\text{dist}(D, G_1) \le \text{diam}(D)$. $\square$

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 $\text{dist}(D, G_1) \le \frac{r}{n}$; 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 $\text{len}(G_2) = \sum_{D \in S_2} \text{dist}(D, G_1) \le |S_2| \cdot \frac{r}{n} \le r \le \text{len}(\text{OPT})$. By construction, we have

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

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 planar straight line 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 planar straight-line 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, $r \le \text{len}(\text{OPT})$, satisfying (1). Let $\text{OPT}(S_3, G_1)$ denote a planar straight-line graph Γ of minimum length such that $G_1 \cup \Gamma$ is connected and intersects every disk in S_3. In Sections 5 and 6, we use hitting sets to compute a $O(1)$-approximation for $\text{OPT}(S_3, G_1)$. Similarly to a quadtree decomposition, we recursively construct a subdivision of R into squares of side lengths $r/2^i$, for $i = 0, 1, \ldots, \lceil \log n \rceil$. Refer to Fig. 1 (right).

Previously, Mitchell [33] used a similar quadtree decomposition for TSPN with disjoint regions in the plane, coined the name "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 [33]. However, Mitchell used stratified grids only for special types of orthogonal polygons, called *histograms* [26]: we generalize this tool to arbitrary planar straight-line graphs.

A square Q in the process gets subdivided unless it is too small ($\mathrm{diam}(Q) < \frac{r}{2n}$) or it is relatively far from G_1 (i.e., $\mathrm{diam}(Q) < \mathrm{dist}(Q, G_1)$).

> **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 $Q \leftarrow \mathtt{dequeue}(L)$. If $\mathrm{diam}(Q) \geq \max(\frac{r}{2n}, \mathrm{dist}(Q, G_1))$, 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 $\mathrm{dist}(Q, G_1)$ 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 $\mathrm{dist}(Q, G_1) \leq 3\,\mathrm{diam}(Q)$.*

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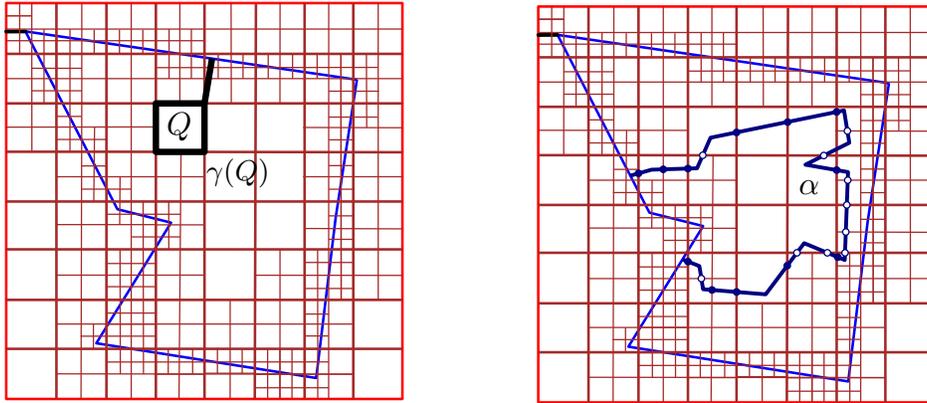

Figure 2: Left: A square Q of the stratified grid, and the graph $\gamma(Q)$. Right: a polygonal curve α. The intersections of α with horizontal (resp., vertical) edges of the stratified grid are marked with empty (resp., full) dots.

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 $\mathrm{len}(\gamma(Q)) \leq (3 + 2\sqrt{2})q$; on the other hand, $\mathrm{len}(\gamma(Q)) \geq \mathrm{per}(Q) = 2\sqrt{2}\,q$, and so we have the following.

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

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

Lemma 3. *If a square $Q \in \mathcal{Q}$ intersects a disk $D \in S_3$, then the boundary of Q also intersects D, and so the graph $\gamma(Q)$ intersects D, as well.*

6

Proof. 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 $Q \in \mathcal{Q}$, it is *not* subdivided by Algorithm **Stratify**(R, G_1), 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 is a contradiction. $\qquad\square$

Recall that the squares in $\mathcal{Q}$ can only intersect at common boundary points.

Lemma 4. *If two squares $Q_1, Q_2 \in \mathcal{Q}$ intersect, then $\frac{1}{2}\text{diam}(Q_1) \leq \text{diam}(Q_2) \leq 2\,\text{diam}(Q_1)$.*

Proof. Assume, without loss of generality, that $\text{diam}(Q_1) < \text{diam}(Q_2) = q$. Then Algorithm **Stratify**(R, G_1) subdivided a square Q_1' such that $Q_1 \subset Q_1'$, $\text{diam}(Q_1') = q$, and $Q_1' \cap Q_2 \neq \emptyset$.

The algorithm subdivided Q_1' but did not subdivide Q_2. This implies $q \geq \max(\frac{r}{2n}, \text{dist}(Q_1', G_1))$ and $q < \max(\frac{r}{2n}, \text{dist}(Q_2, G_1))$. 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, we have $|\text{dist}(Q_1', G_1) - \text{dist}(Q_2, G_1)| \leq q$, and so $q < \text{dist}(Q_2, G_1) \leq 2q$. Similarly, since Q_1 and Q_2 intersect, $\text{dist}(Q_1, G_1) \geq \text{dist}(Q_2, G_1) - \text{diam}(Q_1) > q - \text{diam}(Q_1)$. Combined with Lemma 2, we have $3\,\text{diam}(Q_1) > q - \text{diam}(Q_1)$, 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. $\qquad\square$

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 straight-line 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 an $O(1)$-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 that starts and ends at some points on G_1 and 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 $(Q_0, Q_1, Q_2, \ldots, Q_m)$ 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 $\text{diam}(Q_0) < \max(\frac{r}{2n}, 0) = \frac{r}{2n}$; since $\text{len}(\alpha) \geq \frac{2r}{n}$, we have $\text{len}(\alpha) \geq 4\,\text{diam}(Q_0)$.

For $i = 1, \ldots, m$, let p_i be the first intersection point of α with Q_i, and let e_i be an edge of Q_i that contains p_i. We partition the sequence $P = (p_1, \ldots, p_m)$ into two subsequences as follows. Let $P_{\text{horizontal}} = (p_{\sigma(1)}, \ldots, p_{\sigma(a)})$ be the subsequence of points p_i for which e_i is horizontal, and $P_{\text{vertical}} = (p_{\tau(1)}, \ldots, p_{\tau(b)})$ be the subsequence of remaining points p_i, where e_i is vertical. For two points $p, q \in \alpha$, denote by $\alpha(p, q)$ the portion of α between p and q. With this notation, we have

$$\text{len}(\alpha) \geq \sum_{i=1}^{a-1} \text{len}(\alpha(p_{\sigma(i)}, p_{\sigma(i+1)})) \geq \sum_{i=1}^{a-1} \text{dist}(p_{\sigma(i)}, p_{\sigma(i+1)}), \tag{3}$$

$$\text{len}(\alpha) \geq \sum_{i=1}^{b-1} \text{len}(\alpha(p_{\tau(i)}, p_{\tau(i+1)})) \geq \sum_{i=1}^{b-1} \text{dist}(p_{\tau(i)}, p_{\tau(i+1)}). \tag{4}$$

Since $p_i \in e_i$ for all $1 \leq i \leq m$, we have an obvious bound $\text{dist}(p_i, p_j) \geq \text{dist}(e_i, e_j)$. For any three points, p_i, p_j, p_k, with $i < j < k$, this implies

$$\text{dist}(p_i, p_j) + \text{dist}(p_j, p_k) \geq \max\left(\text{dist}(p_i, p_j), \text{dist}(p_j, p_k), \text{dist}(p_i, p_k)\right),$$
$$\geq \max\left(\text{dist}(e_i, e_j), \text{dist}(e_j, e_k), \text{dist}(e_i, e_k)\right).$$

A horizontal edge e_i of a square Q_i overlaps with at most two horizontal edges of other squares in $\mathcal{Q}$ by Lemma 4; and e_i intersects (at its endpoints) at most two additional horizontal edges of other squares in $\mathcal{Q}$, and the distance between those two edges is precisely $\text{len}(e_i) = \text{diam}(Q_i)/\sqrt{2}$. By Lemma 4, the distance between e_i and any disjoint horizontal edge of other squares in $\mathcal{Q}$ is at least $\text{len}(e_i)/2 = \text{diam}(Q_i)/(2\sqrt{2})$. Consequently,

$$\text{dist}(p_{\sigma(i)}, p_{\sigma(i+1)}) + \text{dist}(p_{\sigma(i+1)}, p_{\sigma(i+2)}) \geq \frac{\max\left(\text{diam}(Q_{\sigma(i)}), \text{diam}(Q_{\sigma(i+1)}), \text{diam}(Q_{\sigma(i+2)})\right)}{2\sqrt{2}}, \tag{5}$$

$$\text{dist}(p_{\tau(i)}, p_{\tau(i+1)}) + \text{dist}(p_{\tau(i+1)}, p_{\tau(i+2)}) \geq \frac{\max\left(\text{diam}(Q_{\tau(i)}), \text{diam}(Q_{\tau(i+1)}), \text{diam}(Q_{\tau(i+2)})\right)}{2\sqrt{2}}. \tag{6}$$

Combining (3) and (5) yields

$$\text{len}(\alpha) = \Omega\left(\sum_{i=1}^{a-1} \text{diam}(Q_{\sigma(i)})\right) \text{ if } a \geq 3. \tag{7}$$

Similarly, Combining (4) and (6) yields

$$\text{len}(\alpha) = \Omega\left(\sum_{i=1}^{b-1} \text{diam}(Q_{\tau(i)})\right) \text{ if } b \geq 3. \tag{8}$$

Observe that $\mathcal{Q}_{\text{horizontal}} = (Q_{\sigma(1)}, \ldots, Q_{\sigma(a)})$ and $\mathcal{Q}_{\text{vertical}} = (Q_{\tau(1)}, \ldots, Q_{\tau(b)})$ is a partition of the sequence $Q_1, \ldots, Q_m$. Recall also that $\text{len}(\alpha) \geq 4\,\text{diam}(Q_0)$ and $\text{len}(\alpha) \geq \frac{2r}{n}$.

If $a \geq 3$ and $b \geq 3$, adding (7) and (8) yields $\text{len}(\alpha) = \Omega\left(\sum_{i=1}^{m} \text{diam}(Q_i)\right)$; since $\text{len}(\alpha) \geq 4\,\text{diam}(Q_0)$, it follows that $\text{len}(\alpha) = \Omega\left(\sum_{i=0}^{m} \text{diam}(Q_i)\right)$, as required. It remains to consider the cases when a or b is smaller than 3.

If $a, b \in \{0, 1, 2\}$, then α intersects at most 7 squares, and thus $m \leq 6$. The first square Q_0 intersects G_1 and has $\text{diam}(Q_0) \leq \frac{r}{2n}$. By Lemma 4, the diameter of subsequent squares can increase at most by factors of 2, hence $\sum_{i=0}^{m} \text{diam}(Q_i) \leq \frac{r}{2n} \sum_{i=0}^{m} 2^i = O(r/n)$. Hence $\text{len}(\alpha) = \Omega(r/n) = \Omega(\sum_{i=0}^{m} \text{diam}(Q_i))$, as required.

8

Assume that $a < 3 \leq b$ (the case that $b < 3 \leq a$ is analogous). If $a = 0$, then (8) completes the proof. If $a \in \{1, 2\}$, then each square $Q_{\sigma(i)}$, for $i = 1, \ldots, a$, is adjacent to some square $Q_{j(i)}$ where the edge $e_{j(i)}$ is vertical. By Lemma 4, $\mathrm{diam}(Q_{\sigma(i)}) \geq \frac{1}{2}\mathrm{diam}(Q_{j(i)})$. Consequently, (8) implies $\mathrm{len}(\alpha) = \Omega(\sum_{i=0}^{m} \mathrm{diam}(Q_i))$, as required. $\qquad\square$

Lemma 6. *If Γ is a planar straight-line 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*

$$\mathrm{len}(\Gamma) = \Omega\left(\sum_{Q \in \mathcal{H}} \mathrm{diam}(Q)\right). \tag{9}$$

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 an at most twice. Then $\mathrm{len}(C) \geq \frac{1}{2}\mathrm{len}(\alpha)$, and $\mathrm{len}(\alpha) = \Omega(\sum_{Q \in \mathcal{Q}(\alpha)} \mathrm{diam}(Q))$ by Lemma 5. Summation over all the components of Γ yields (9). $\qquad\square$

Recall that $\mathsf{OPT}(S_3, G_1)$ is a planar straight-line 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 $W_0 = O(\mathrm{len}(\mathsf{OPT}(S_3, G_1)))$.*

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

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 each hyperedge corresponds to a 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 (Corollary 4) 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, this is the case as shown in the following lemma.

Lemma 7. *If a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$, then $\mathrm{diam}(Q) \leq 2\,\mathrm{diam}(D)$.*

Proof. Since $Q \in \mathcal{Q}$, Algorithm **Stratify**(R, G_1) did not subdivide Q, and so $\mathrm{diam}(Q) < \frac{r}{2n}$ or $\mathrm{diam}(Q) < \mathrm{dist}(Q, G_1)$. We distinguish two cases.
Case 1: $\mathrm{diam}(Q) < \frac{r}{2n}$. In this case, Corollary 1 yields $\mathrm{diam}(Q) < \frac{r}{2n} < \mathrm{diam}(D) < 2\,\mathrm{diam}(D)$.
Case 2: $\mathrm{diam}(Q) < \mathrm{dist}(Q, G_1)$. We distinguish two subcases.
Case 2a: $\mathrm{dist}(Q, G_1) \leq 2\,\mathrm{dist}(D, G_1)$. Then the claim follows from Lemma 1:

$$\mathrm{diam}(Q) < \mathrm{dist}(Q, G_1) \leq 2\,\mathrm{dist}(D, G_1) \leq 2\,\mathrm{diam}(D).$$

9

Case 2b: $\mathrm{dist}(Q, G_1) > 2\,\mathrm{dist}(D, G_1)$. In this case, we use the fact that Q and D intersect. The triangle inequality yields

$$\mathrm{dist}(Q, G_1) \leq \mathrm{dist}(D, G_1) + \mathrm{diam}(D) < \frac{1}{2}\,\mathrm{dist}(Q, G_1) + \mathrm{diam}(D).$$

Rearranging this inequality, we have $\mathrm{dist}(Q, G_1) \leq 2\,\mathrm{diam}(D)$, and further by Lemma 1 that

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

as required. □

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).

Corollary 4. *If a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$, then D contains one of the sentinel points of Q.*

Proof. Assume that a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$ of side length h and refer to Fig. 3 (right). By Lemma 7, $\mathrm{diam}(Q) \leq 2\,\mathrm{diam}(D)$. By scaling down D from a center in $D \cap Q$, we find a disk D' intersecting Q with $\mathrm{diam}(D') = \frac{1}{2}\,\mathrm{diam}(Q)$. The inscribed axis-aligned square Q' of D' has $\mathrm{diam}(Q') = \frac{1}{2}\,\mathrm{diam}(Q)$. That is, the side length of Q' is $h/2$, and $\mathrm{dist}(Q', Q) \leq (\sqrt{2} - 1)h$. 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. □
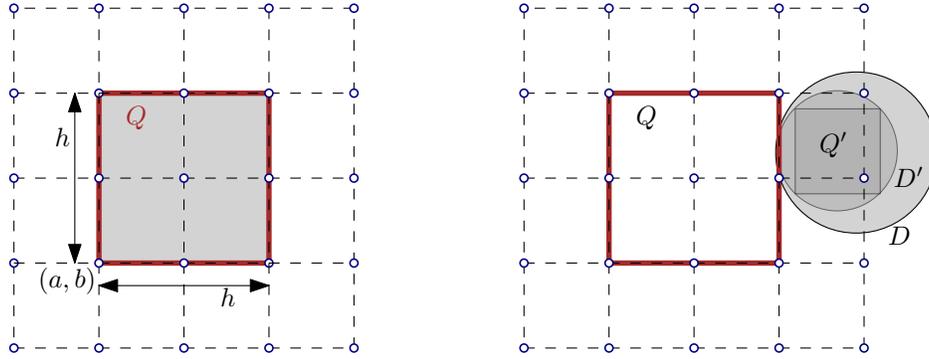

Figure 3: Left: the set of 25 sentinel points for a square Q. Right: a disk $D \in S_3$ intersects a square $Q \in \mathcal{Q}$.

We define a new weighted hypergraph $\mathcal{G}' = (V', E')$, where V' is the union of sentinel sets for all $Q \in \mathcal{Q}$, and each hyperedge in E' is the set of sentinels that lie in a disk $D \in S_3$. Let the weight of a sentinel $s \in \mathbb{R}^2$ 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 $25 \sum_{Q \in \mathcal{Q}} \mathrm{diam}(Q)$. Note that not all sentinels of Q are contained in Q. Since the side lengths of adjacent squares of the stratified grid differ by at most a factor of 2 (cf. Lemma 4), every sentinel of Q is contained in Q or in an adjacent square of $\mathcal{Q}$. Conversely, a sentinel lying in a square $Q \in \mathcal{Q}$ corresponds to Q or to an adjacent square in $\mathcal{Q}$; thus every sentinel corresponds to at most 9 squares in $\mathcal{Q}$, and so the weight of a sentinel corresponding to $Q \in \mathcal{Q}$ is $O(\mathrm{diam}(Q))$.

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

Lemma 8.   *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 $O(\sum_{Q \in \mathcal{H}} \operatorname{diam}(Q))$.*

   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 Corollary 4, D contains one of the sentinels of Q. Consequently, every disk $D \in S_3$ contains a sentinel in H'. By definition, the weight of H' is $O(\sum_{Q \in \mathcal{H}} \operatorname{diam}(Q))$.

(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 $Q \in \mathcal{H}$. Every sentinel $s \in V'$ corresponding to a square $Q \in \mathcal{Q}$ lies in Q or in an adjacent square by Lemma 4. The diameters of adjacent squares differ by at most a factor of 2. Therefore, $s \in V'$ lies in a square of diameter at most $2\operatorname{diam}(Q) = 2w(s)$. It follows that the weight of $\mathcal{H}$ is at most $2\sum_{s \in H'} w(s)$, 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 $\operatorname{len}(\xi_0) = O(\operatorname{len}(\mathsf{OPT}))$ by [6, 19]. Compute the graph G_1 with $\operatorname{len}(G_1) = O(\operatorname{len}(\mathsf{OPT}))$, 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 $\operatorname{len}(G_2) = \sum_{D \in S_2} \operatorname{dist}(D, G_1) \leq |S_2| \cdot \frac{r}{n} \leq r \leq \operatorname{len}(\mathsf{OPT})$.

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 randomized algorithm by Chan et al. [13] to compute a hitting set H' for $\mathcal{G}'$ whose weight is $O(1)$ times the minimum with high probability. Let $\mathcal{H}$ be the set of squares in $\mathcal{Q}$ containing the sentinels in H'. By Lemma 8, $\mathcal{H}$ is a hitting set for the hypergraph $\mathcal{G} = (\mathcal{Q}, E)$ whose weight is at most $O(1)$ times the minimum W_0 with high probability. Put $G_3 = \cup_{Q \in \mathcal{H}} \gamma(Q)$. Then G_3 intersects every disk in S_3 by Lemma 3, and $\operatorname{len}(G_3) = \Theta(\sum_{Q \in \mathcal{H}} \operatorname{diam}(Q)) = \Theta(W_0)$ by Corollary 2. Finally, Corollary 3 yields $\operatorname{len}(G_3) = O(\operatorname{len}(\mathsf{OPT}(S_3, G_1))) = O(\operatorname{len}(\mathsf{OPT}))$.

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 $\operatorname{len}(G) = \operatorname{len}(G_1) + \operatorname{len}(G_2) + \operatorname{len}(G_3) = O(\operatorname{len}(\mathsf{OPT}))$ with high probability. Consequently, an Eulerian tour of the multi-graph containing each edge of G twice visits each disk and has length $2\operatorname{len}(G) = O(\operatorname{len}(\mathsf{OPT}))$, with high probability, as required. $\square$

7   Conclusion

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 and arbitrarily intersections. In light of the connection we established between TSPN and MWHS in geometric hypergraphs, the following question is relevant:

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

---

[1]A 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.

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 approximation algorithm for TSP with a set of objects of linear union complexity, e.g., pseudo-disks?

3. Is there a constant approximation algorithm for TSP with convex bodies in the plane?

4. Is TSP with disks in the plane APX-hard? Is TSP with convex bodies in the plane APX-hard?

5. Is there a constant approximation algorithm for TSP with balls (of arbitrary radii) 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 Appl. Math.* **55** (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 Sympos. 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 Verlag, 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. Penninkx, R. Sitters, and T. Wolle, On the minimum corridor connection problem and other generalized geometric problems, *Computational Geometry* **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 International Symposium on Theoretical Aspects of Computer Science (STACS)*, LIPIcs, 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 Sympos. 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 J. Discrete Math.* **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, preprint, October 2014, `arXiv:1303.6659v2`. A preliminary version in *Proc. 24th ACM-SIAM Sympos. on Discrete Algorithms (SODA)*, SIAM, 2013.

[19] K. M. Elbassioni, A. V. Fishkin, and R. Sitters, Approximation algorithms for the Euclidean traveling salesman problem with discrete and continuous neighborhoods, *Int. J. Comput. Geometry Appl.* **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, *Proc. 8th Annual 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] J. Gudmundsson and C. Levcopoulos, A fast approximation algorithm for TSP with neighborhoods, *Nordic J. of Comput.* **6** (1999), 469–488.

[24] P. Kamousi and S. Suri, Euclidean traveling salesman tours through stochastic neighborhoods, *Proc. 24th International Symposium on Algorithms and Computation (ISAAC)* LNCS 8283, Springer, 2013, pp. 644–654.

[25] 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** (1986), 59–70.

[26] C. Levcopoulos, Heuristics for Minimum Decompositions of Polygons, PhD thesis, Linkping Studies in Science and Technology No. 74, 1987.

[27] 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.

[28] C. Mata and J. S. B. Mitchell, Approximation algorithms for geometric tour and network design problems, *Proc. 11th ACM Sympos. Comput. Geom. (SOCG)*, ACM Press, 1995, pp. 360–369.

[29] J. S. B. Mitchell, Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems, *SIAM J. on Comput.* **28(4)** (1999), 1298–1309.

[30] 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.

[31] 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.

[32] J. S. B. Mitchell, A PTAS for TSP with neighborhoods among fat regions in the plane, *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, ACM Press, 2007, pp. 11–18.

[33] J. S. B. Mitchell, A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane, in *Proc. 26th Annual Symposium on Computational Geometry (SOCG)*, ACM Press, 2010, pp. 183–191.

[34] N. H. Mustafa and S. Ray, Improved results on geometric hitting set problems, *Discrete & Computational Geometry* **44(4)** (2010), 883–895.

[35] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, John Wiley, New York, 1995.

[36] C. H. Papadimitriou, Euclidean TSP is NP-complete, *Theor. Comp. Sci.* **4** (1977), 237–244.

[37] S. B. Rao and W. D. Smith, Approximating geometrical graphs via "spanners" and "banyans", *Proc. 30th ACM Sympos. Theory Comput. (STOC)* ACM Press, 1998, pp. 540–550.

[38] S. Safra and O. Schwartz, On the complexity of approximating TSP with neighborhoods and related problems, *Computational Complexity* **14(4)** (2005), 281–307.

[39] P. Slavik, The errand scheduling problem, CSE Technical Report 97-02, University of Buffalo, Buffalo, NY1, 1997.

[40] S. Spirkl, The guillotine subdivision approach for TSP with neighborhoods revisited, preprint, April 2014, `arXiv:1312.0378v2`.

[41] K. R. Varadarajan, Epsilon nets and union complexity, in *Proc. 25th Annual Sympos. on Computational Geometry (SOCG)*, ACM Press, 2009, pp. 11–16.