Approximation Algorithms for the Euclidean Traveling Salesman Problem with Discrete and Continuous Neighborhoods*

Khaled Elbassioni† Aleksei V. Fishkin‡ René Sitters§

Abstract

In the Euclidean traveling salesman problem with discrete neighborhoods, we are given a set of points P in the plane and a set of n connected regions (neighborhoods), each containing at least one point of P. We seek to find a tour of minimum length which visits at least one point in each region. We give (i) an $O(\alpha)$-approximation algorithm for the case when the regions are disjoint and α-fat, with possibly varying size; (ii) an $O(\alpha^3)$-approximation algorithm for intersecting α-fat regions with comparable diameters. These results also apply to the case with continuous neighborhoods, where the sought TSP tour can hit each region at any point. We also give (iii) a simple $O(\log n)$-approximation algorithm for continuous non-fat neighborhoods. The most distinguishing features of these algorithms are their simplicity and low running-time complexities.

1 Introduction

In the *TSP with neighborhoods problem* we are given n *connected* subsets (*regions* or *neighborhoods*) of the Euclidean plane, and a set of specified points P. The objective is to find a tour of minimum length which visits each region in one of the specified points[1]. This generalizes the classical Euclidean TSP with applications in VLSI-design, and other routing-related applications (see e.g. [14, 16]). The problem is also a special case of the more general *Euclidean Group-TSP* problem, in which the specified regions are not necessarily connected, and hence can be arbitrary subsets of points in the plane. In this paper, we will impose the

*Preliminary versions of this paper containing some of the results appeared in Proceedings of ICALP 2005 [6] and ISAAC 2006 [7]

†Max-Planck-Institut für Informatik, Saarbrücken, Germany; elbassio@mpi-sb.mpg.de

‡University of Liverpool, UK; avf@csc.liv.ac.uk

§Department of Mathematics and Computer Science, Eindhoven University of Technology, the Netherlands; r.sitters@tue.nl

[1]For instance, a salesman wants to meet a set of potential buyers. Each buyer indicates a set of potential locations (neighborhoods) where he or she can meet the buyer. The salesman would like to minimize the total length of the tour required to meet all the potential buyers. How to construct such a tour?

connectivity requirement and consider two variants of the problem (see Figure 1): (i) the *continuous* case, in which P is the whole plane, i.e., the sought TSP tour can hit each region at any of its points, and (ii) the *discrete* case, in which P is a finite set of points intersecting all regions, and the tour is allowed to hit the region only in one of the specified points. Henceforth, we shall refer to the former problem as *Continuous TSPN* and for the latter as *Discrete TSPN*.

Although the problem has been extensively studied in the last decade after Arkin and Hassin [1] introduced it in 1994, still large discrepancies remain between known inapproximability and approximation ratios for various cases. For the most general case in which the subsets are unrestricted, and the metric is not necessarily Euclidean, the gap is almost closed: Garg et al. [9] gave a randomized $O((\log N)^2 \log k \log n)$-approximation algorithm for the group Steiner tree problem, the variant in which we are given a graph with N vertices and a set of n groups with at most k vertices per group, and seek a minimum cost Steiner tree connecting to at least one point from each group[2]. This approximation ratio can be improved to $O(\log N \log k \log n)$ using the results of Fakcharoenphol et al. [8]. Slavík [18] showed that the problem can be approximated within $O(k)$. On the negative side, Halperin and Krauthgamer [11] gave an inapproximability threshold of $\Omega(\log^{2-\epsilon} n)$ for any fixed $\epsilon > 0$.

With this being the situation for the general case, recent research has considered the cases where the given subsets are connected regions in the plane. Typically, the putting the discrete restriction on the problem makes it harder. Safra and Schwartz [17] showed that the discrete problem is NP-hard to approximate, in the general case, within any constant factor, and is APX-hard if each subset forms a connected region in the plane. For connected polygonal regions, Mata and Mitchell [12] gave an $O(\log n)$-approximation in $O(N^5)$-time based on"guillotine rectangular subdivision", where N is the total number of vertices of the polygons. Gudmundsson and Levcopoulos [10] reduced the running time to $O(N^2 \log N)$.

Figure 1: (a) TSP with discrete neighborhoods (b) TSP with continuous neighborhoods

Previous results also distinguish between *fat* regions, such as disks, and *non-fat* regions, such as line-segments, and between instances with *disjoint* regions and *intersecting* regions. Non-fatness and intersections seem to make the problem much harder and, in fact, no constant factor approximation algorithm is known for the general case of intersecting non-fat regions.

For the continuous case when the regions are translates of disjoint convex polygons, and for disjoint unit disks, Arkin and Hassin [1] presented constant-factor approximations. Dumitrescu and Mitchell [5] gave an $O(1)$-approximation

[2]Note that, in the metric case, the minimum cost Steiner tree and the minimum cost group-TSP tour are within a constant factor of each other.

algorithm for arbitrary connected regions of comparable diameters. For disjoint varying-sized convex fat regions in the continuous case, de Berg et al. [4] presented an $O(\alpha^3)$-approximation algorithm, with a very big hidden constant, where α is a measure of fatness of the regions. In a recent paper, Mitchell [15] gave a PTAS for disjoint varying-sized connected fat regions which works both in the continuous and discrete cases, even under a weaker notion of fatness than the one used in [4]. As the first result in this paper, we give a very simple algorithm with approximation factor $9.1\alpha + 1$. This algorithm also works for the continuous and discrete cases of disjoint varying-sized connected fat neighborhoods. Although this is only a constant factor approximation if α is a constant, the simplicity and the low running-time complexity is an advantage over PTAS's. Furthermore, it is not clear how to generalize the PTAS of [15] to higher dimensions, while the generalization of our $O(1)$-approximation algorithm to any fixed dimension is straightforward.

Perhaps the two most natural extensions for which no constant factor algorithm is known, are that of non-fat disjoint objects, and that of fat intersecting objects. In the second part of the paper, we consider the discrete version of the latter problem and give an $O(\alpha^3)$-approximation algorithm for intersecting convex α-fat objects of comparable size. The proof follows from an interesting observation about the relation between the length of the optimal TSP tour inside a square and the distribution of the points. This algorithm also works under the same assumptions in the continuous case. In the last part of the paper, we give a simple $O(\log n)$-approximation algorithm for the general continuous case of connected regions, which does not require the regions to be simple polygons. This algorithm is also an $O(1)$-approximation in the case when the neighborhoods have comparable diameters. A summary of the currently best known approximation factors for the different variants of the problem is given in Table 1.

	Discrete		Continuous	
General	$O(\log N \log k \log n)$ [9] $O(k)$ [18]		–	
Connected Intersecting	?		polygonal regions $O(\log n)$ [12, 10]	general $O(\log n)$ [*]
Intersecting fat, convex	comparable size $O(1)$ [*]	general ?	comparable size $O(1)$ [5], [*]	general ?
Disjoint fat	PTAS [15]		PTAS [15]	

Table 1: The Euclidean TSPN problem; the state of the art. n is the number of regions, $N = |P|$ is the total number of specified points, and k is the maximum number of points per region. (?) means nothing specific known for this particular case. [*] refers to the present article. The entries in the table give the approximation ratios.

Paper outline. In the next section, we recall a few basic definitions and facts that will be used throughout the algorithms and their analysis. We begin in Section 3 by observing that, even though it is not known if the TSP with neighborhoods problem is APX-hard for intersecting fat objects with the same size, the situation changes if we do not insist on fatness, but assume even that the objects are line segments of almost the same length. Section 4 presents an $O(\alpha)$-approximation algorithm for the discrete case with α-fat disjoint neighborhoods. In Section 4.2, we analyze the approximation factor of the algorithm, and in Sections 4.3 and 4.4, we respectively comment on how to extend it to the continuous case and to higher dimensions. Section 5 presents our $O(\alpha^3)$-approximation algorithm for intersecting convex α-fat objects, with comparable diameters. The analysis of the approximation factor follows from two lemmas presented in Section 5.1. Finally, Section 6 gives a simple $O(\log n)$-approximation algorithm for the general case of connected neighborhoods.

2 Preliminaries

An instance of the Euclidean TSP with neighborhoods (TSPN) problem is given by a subset P of the Euclidean plane $\mathbb{R}^2$, and a set of n objects (regions, or neighborhoods) $\mathcal{O} = \{O_1, \ldots, O_n\}$ in $\mathbb{R}^2$. The objective is to find a minimum length TSP tour that hits each object O_i at some point in $S_i \stackrel{\text{def}}{=} P \cap O_i$. Without loss of generality we assume that $S_i \neq \emptyset$ for all $i \in [n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$ and $P = S_1 \cup \ldots \cup S_n$. We denote the optimal tour by Opt, and its length by $|\text{Opt}|$, and use $\delta_1, \ldots, \delta_n$ to denote the diameters of the objects. Note that our definition of the TSPN problem also includes the TSP with continuous neighborhoods if we set $P = \mathbb{R}^2$.

There are several definitions of fatness in the literature and the following is commonly used for the problem we consider [4, 6, 19].

Definition 1 *An object $O \subseteq \mathbb{R}^2$ is said to be α-fat if for any disk D which does not fully contain O and whose center lies in O, the area of the intersection of O and D is at least $1/\alpha$ times the area of D.*

Notice for example that the plane $\mathbb{R}^2$ has fatness 1, a halfspace has fatness 2, a disk has fatness 4, while a line segment has fatness ∞.

The following packing lemma will be used in our analysis of the approximation ratios of various algorithms.

Lemma 1 *The length of the shortest path connecting k disjoint α-fat objects in $\mathbb{R}^2$ is at least $(k/\alpha - 1)\pi\delta/4$, where δ is the diameter of the smallest object.*

Proof. Consider a disk with diameter δ following the shortest path connecting the k objects (see Figure 2). At the point where the path touches a certain object, the disk intersects its boundary and hence at least an $1/\alpha$ fraction of the disk at that point intersects the object. The total area covered by the moving disk must be at least k/α times the area $\pi\delta^2/4$ of the disk. On the

Figure 2: Illustrating Lemma 1.

other hand, it is known that the total area covered by a disk that follows a continuous path of length l in $\mathbb{R}^2$ is at most $\pi\delta^2/4 + \delta l$. Combining the upper and lower bounds on the area we get,

$$k\pi\delta^2/(4\alpha) \leq \pi\delta^2/4 + \delta l \;\Rightarrow\; l \geq (k/\alpha - 1)\pi\delta/4.$$

□

Remark 1 *Mitchell uses weaker definition of fatness in [15]: An object of diameter δ is fat if it contains a disk of diameter $\Omega(\delta)$ inside it. Under this definition, there is only a constant factor change in the bound of Lemma 1 and, consequently, a constant factor change in the performance of the algorithms in Section 4. However, for intersecting regions, discussed in Section 5, we do need the stronger definition of fatness.*

Definition 2 *A subset $P' \subseteq P$ is called a* hitting pointset *for $\mathcal{O}$ if $P' \cap O_i \neq \emptyset$ for $i = 1, \ldots, n$, and is called furthermore a* minimal hitting pointset *if for every $x \in P'$ there exists an $i \in [n]$ such that $(P' \setminus \{x\}) \cap O_i = \emptyset$.*

A minimal hitting pointset for $\mathcal{O}$ can be found by the natural greedy algorithm: set $P' = P$, and keep deleting points from P' as long as it is still a hitting set.

Definition 3 *An axis-aligned square B is called* a covering box *for the set of objects $\mathcal{O}$ if B contains a hitting pointset for $\mathcal{O}$, and is called a* minimum covering box *if it is of smallest size among all such covering boxes.*

Since a minimum covering box is determined by at most three points of P on its boundary, there are only $O(|P|^3)$ such candidates. Thus, if P is finite, by enumerating over all such boxes, and verifying if they contain a hitting set, one can compute a minimum covering box. For our purposes, it is enough to work with a $(1+\epsilon)$-approximation, which can be obtained for any fixed $\epsilon > 0$, even if P is infinite, by the following "standard" technique: Let B_0 be a box of minimum size L_0 that contains all the objects in $\mathcal{O}$. Clearly, there is a minimum covering box for $\mathcal{O}$ is contained in B_0 and $L_0 \leq \sum_{i=1}^{n} \delta_i$.

For $i = 1, 2, \ldots,$ iterate:

1. Consider the four smaller squares of size $L_{i-1}/2 + L_{i-1}/(2(1+\epsilon))$, each touching one of the four corner points of B_{i-1} and completely contained inside B_{i-1};

2. If none of the four squares is a covering box for $\mathcal{O}$ then return B_{i-1}.

If B_i contains a covering box of size $L_{i-1}/(1+\epsilon)$ then it must be completely contained in one of the four candidate squares. Thus, a $(1+\epsilon)$-approximation of the minimum covering box for $\mathcal{O}$ can be computed[3] in time $O(n \log(L_0/L))$, where L is the size of the minimum covering box.

Remark 2 *From now on, we shall ignore factors of* $(1+\epsilon)$*, with the understanding that the stated bounds on the approximation ratios hold with an additional multiplicative factor of* $(1+\epsilon)$*, for a sufficiently small* $\epsilon > 0$*.*

The following lower bound will be used frequently in our analysis.

Lemma 2 *([1]) If* L *is the size of the minimum covering box for a given instance of the TSPN problem, then* $2L \leq |\text{OPT}|$*.*

3 APX-hardness

De Berg et al. [4] prove that the TSP with connected neighborhoods problem is APX-hard. The constant was raised to $2-\varepsilon$ by Safra and Schwartz [17]. Both reductions use curved objects of varying sizes. We can show that the problem is even APX-hard for the very restricted case where all objects are line-segments of almost equal length.

Theorem 1 *The TSP with neighborhoods problem is APX-hard, even if all objects are line segments of approximately the same length, i.e, the lengths differ by an arbitrarily small constant factor.*

Proof. Clementie et al. [3] show that vertex cover for 3-partite graphs cannot be approximated within a factor 34/33, unless $P = NP$. The reduction to the TSPN problem is pretty straightforward. Given a 3-partite graph G we plot the graph corresponding to the 3-partition as shown in Figure 3. The vertices of the graph correspond to end-points of line-segments, and the edges correspond to the objects (line-segments) of the TSPN instance in the obvious way: a line segment connects two points if and only if the corresponding vertices in the graph are adjacent. Further, we define a large number of point-objects which together form a polygon with perimeter L. The small equilateral triangle in the closeup has side-length d. If d is small enough, then an optimal tour follows the polygon and jumps up and down to some of the vertices. The extra cost of the detour for each such vertex is $2d-d = d$. Consider an optimal tour and let S be the set of vertices of G that are visited, then $\text{OPT} = L + |S|d$. In the reduction of Clementie et al., the size of the optimal vertex cover is $\Omega(n)$. More precisely, it is NP-hard to decide if the minimum vertex cover has size at most $n/2$ or at least $n/2 \cdot 34/33$. Now we let $d = 1/n$ and choose the distance between any two vertices substantially larger, say $4/n$. We let the perimeter of the polygon be sufficiently large, say 10. If there is a vertex cover of size $n/2$, then there exists a tour of length $L + nd/2 = 10.5$. On the other hand, if there exists a tour of

[3]We assume that testing if an axis-parallel rectangle intersects an object can be done in $O(1)$ time; if each object O_i is described by a finite set of points S_i, then the bound should be multiplied by the maximum number of points in an object.

Figure 3: The reduction for a 3-partite graph on 12 vertices.

length at most $10 + \beta$, then there must be a vertex cover of size at most βn. Taking $\beta = 34/66$ shows that TSPN cannot be approximated within a factor $(10 + 34/66)/10.5 \approx 1.0014$.

Notice that we can replace the points on the polygon by segments directed outwards and of arbitrary length without breaking the reduction. Note also that the reduction works for both the continuous and discrete variants of the problem. □

4 Varying-sized disjoint fat objects

In this section, we consider the case when the objects $O_1, \ldots, O_n$ are pairwise disjoint α-fat, and the required TSP tour must hit object O_i in a point of the set $S_i = P \cap O_i$, for $i = 1, \ldots, n$. We present the algorithm in Section 4.1, analyze it in Section 4.2, and show how to extend it to the continuous case, and to higher dimensions in Sections 4.3 and 4.4 respectively.

4.1 The algorithm

There exists a simple $(n-1)$-approximation algorithm that we denote by Dtspn-Greedy (for *discrete* TSPN). We define the *distance* between a point p and a set X as $d(p, X) = \min_{x \in X} d(p, x)$, where $d(p, q)$ denotes the Euclidean distance between two points $p, q \in \mathbb{R}^2$.

Algorithm Dtspn-Greedy:

(1) Pick the points $p_i \in S_i$ $(i = 1, \ldots, n)$ that minimize $\sum_{j=2}^{n} d(p_1, p_j)$.

(2) For all $j \geq 2$ add two copies of the edge (p_1, p_j) and construct a tour by short-cutting the edges.

We shall use the above greedy algorithm only when the number of objects is smaller than some constant. In general, we use the following algorithm.

Algorithm DTSPN-DISJOINT:

(1) Order the point sets by their diameter $\delta_1 \le \delta_2 \le \ldots \le \delta_n$.

(2) Pick any $p_1 \in S_1$. For $i = 2, \ldots, n$, pick the point $p_i \in S_i$ that minimizes $d(p_i, \{p_1, \ldots, p_{i-1}\})$, i.e., pick the point that is closest to the already chosen points..

(3) Construct a $(1+\epsilon)$-approximate TSP tour T on this set of n points.

(4) Output the shorter of T and the tour constructed by algorithm DTSPN-GREEDY.

The third step can be done efficiently for any $\epsilon > 0$ using techniques from [2] and [13].

Remark 3 *Algorithm* DTSPN-DISJOINT *can be implemented in time*[4] $\tilde{O}(n|P|)$.

4.2 The analysis

Lemma 3 *Algorithm* DTSPN-GREEDY *gives an* $(n-1)$*-approximate solution for Discrete TSPN.*

Proof. Any TSP tour contains two vertex disjoint paths from S_1 to S_i for all $i \in \{2, \ldots, n\}$. Therefore, $(n-1)|\text{OPT}| \ge 2\sum_{i=2}^{n} d(p_1, p_i)$, which is at most the length of the tour constructed by the algorithm. □

Theorem 2 *Algorithm* DTSPN-DISJOINT *gives a* $(9.1\alpha+1)$*-approximate solution for the TSP with non-intersecting* α*-fat neighborhoods.*

Proof. By Lemma 3 we only need to prove the the tour T in step (3) is $(9.1\alpha+1)$-approximate for large values of n, i.e., for $n-1 \ge 9.1\alpha+1$.

Denote the set of chosen points by $P' = \{p_1, \ldots, p_n\}$ and denote by $p_i^* \in \{p_1, \ldots, p_{i-1}\}$ the point at minimum distance from p_i and denote the distance by x_i.

Consider some optimal solution OPT and fix an orientation of this tour. Let $k = \lceil \alpha(4/\pi+1) \rceil$ and notice that $k \in \{1, \ldots, n\}$ is satisfied by the assumption $n \ge 9.1\alpha+2$. We define T_i as the part of this directed tour that connects exactly k sets and starts from the point in S_i. Let t_i be the length of path T_i.

Consider some $i \in \{1, \ldots, n\}$ and let S_h be a set with smallest diameter among those from the k sets on the path T_i. Then, by Lemma 1 and the choice of k we have

$$t_i \ge (k/\alpha - 1)\pi\delta_h/4 \ge \delta_h. \qquad (1)$$

[4] where $\tilde{O}(\cdot)$ hides some polylogarithmic factors in n and $|P|$ and $|\text{OPT}|$

Since S_i is on this path T_i and we ordered the sets by their diameter we have $h \in \{1, 2, \ldots, i\}$. We distinguish two cases.

If $h = i$, meaning that S_i has smallest diameter, then by (1) we have

$$t_i \geq \delta_i. \tag{2}$$

If $h \in \{1, 2, \ldots, i-1\}$, then we argue as follows. Since the algorithm picked point p_i we know that the distance from any point in S_i to the point p_h (which is chosen before p_i) is at least x_i. Hence, the distance from any point in S_i to any point in S_h is at least $x_i - \delta_h$, implying $t_i \geq x_i - \delta_h$. Together with (1) this yields

$$t_i \geq \max\{\delta_h, x_i - \delta_h\} \geq x_i/2. \tag{3}$$

We will construct a TSP tour on the set of points P' chosen by the algorithm, and bound its length using the bounds (2) and (3). Let H be the set of all indices i for which $t_i \geq \delta_i$, and let $\textsc{Opt}_H$ be a shortest TSP tour on the points $\{p_i | i \in H\}$. Clearly,

$$|\textsc{Opt}_H| \leq |\textsc{Opt}| + 2\sum_{i \in H} \delta_i \leq |\textsc{Opt}| + 2\sum_{i \in H} t_i.$$

Let $\bar{H} = [n] \setminus H$. Then by (3), we know that for any $i \in \bar{H}$, the length of the edge $d(p_i, p_i^*)$ equals $x_i \leq 2t_i$. We add this edge twice, for every $i \in \bar{H}$, to the tour $\textsc{Opt}_H$. Clearly, the resulting graph is Eulerian. Moreover, it is connected since $1 \in H$, and for any i we have $p_i^* = p_j$ for some $j < i$. The total length of this Eulerian graph is

$$|\textsc{Opt}_H| + 2\sum_{i \in \bar{H}} d(p_i, p_i^*) \leq |\textsc{Opt}| + 2\sum_{i \in H} t_i + 2\sum_{i \in \bar{H}} 2t_i \leq |\textsc{Opt}| + 4\sum_{i=1}^{n} t_i,$$

and hence there exists a TSP tour on P' of at most this length. When we take the sum over all t_i, every edge is counted $k-1$ times, implying $(k-1)|\textsc{Opt}| = \sum_{i=1}^{m} t_i$. Substituting the value of k we conclude that the tour given by our algorithm has length at most

$$(1 + 4(k-1))|\textsc{Opt}| < (9.1\alpha + 1)|\textsc{Opt}|.$$

□

4.3 Extension to the continuous case

It is easy to see that the algorithm and proof of the previous section apply directly to the Euclidean TSPN problem under the weaker assumption that, given the points $\{p_1, \ldots, p_{i-1}\}$, we can efficiently find the point p_i in the infinite set of points S_i that minimizes $d(p_i, \{p_1, \ldots, p_{i-1}\})$.

4.4 Higher dimensions

The generalization of the definitions and lemmas of the previous subsections to higher dimensions is straightforward.

Definition 4 *An object $O \subseteq \mathbb{R}^d$ is said to be α-fat if for any d-dimensional sphere D which does not fully contain O and whose center lies in O, the volume of the intersection of O and D is at least $1/\alpha$ times the volume of D.*

We denote the volume of a d-dimensional sphere with radius r by $V_d(r)$.

Lemma 4 *If the center of a d-dimensional sphere with radius r follows a path P in $\mathbb{R}^d$, then the volume covered by the sphere is at most $|P|V_{d-1}(r) + V_d(r)$.*

Lemma 5 *The length of the shortest path connecting k disjoint α-fat objects is at least $(k/\alpha - 1)V_d(r)/V_{d-1}(r)$, where r is half the diameter of the smallest object.*

The volume of a d-dimensional sphere with radius r is

$$V_d(r) = \frac{\pi^{d/2} r^d}{\Gamma(\frac{d+2}{2})},$$

where Γ is the well-known gamma function. For $d \geq 3$ we get

$$\frac{V_d(r)}{V_{d-1}(r)} = r\sqrt{\pi}\frac{\Gamma((d+2)/2)}{\Gamma((d+1)/2)} > r\sqrt{\pi d/2}.$$

For small values of n we can simply get an $(n-1)$-approximation as described in the previous section. If we choose k such that $(k/\alpha - 1)\sqrt{\pi d/2} \geq 1$, then the proof of Theorem 2 applies here without any adjustment.

Theorem 3 *Algorithm* DTSPN-DISJOINT *is an $O(\alpha/\sqrt{d})$-approximation algorithm for TSPN in $\mathbb{R}^d$ with disjoint α-fat neighborhoods.*

Notice that the approximation factor decreases in the dimension for constant α. However, for bounded objects, α grows exponentially in d. For example, $\alpha = 2^d$ for a d-dimensional sphere.

Safra and Schwartz [17] showed that TSP with neighborhoods in $\mathbb{R}^3$ is unlikely to be approximable within $O(\log^{1/2} n)$. Hence, there is little hope to improve our result for $d \geq 3$ to a ratio independent of the fatness α.

5 Intersecting convex fat objects

In this section, we assume again that each object $O_i \in \mathcal{O}$ can be hit only at specified points, i.e., we are given a set of points P and the required TSP tour must hit O_i at some point in $S_i = P \cap O_i$. We consider the case when $O_1, \ldots, O_n$ are intersecting convex α-fat objects of the same (or comparable) diameter δ.

Figure 4: (a) Partitioning the covering box. (b) A $(\beta L, \theta)$-sector $S(p) \in \mathcal{S}_{i,j}$.

Before we present our algorithm for intersecting convex fat neighborhoods, we give two lemmas which we use in the analysis and which may be of independent interest. The first lemma relates the length of a TSP tour on a set of points in the plane to the distribution of the points. Then, we use this in the second lemma to show that, even though a collection of convex fat objects with exactly one point in each might be intersecting, it still exhibits a *packing property* that admits a "short tour" visiting all the points.

5.1 Two packing lemmas

Call a circular sector with head angle $\theta \leq \pi$, and radius γ a (γ, θ)-sector.

Lemma 6 *Let $P \subseteq \mathbb{R}^2$ be a a set of points with covering box of size L, and $\beta > 0$ and $0 < \theta < \pi/2$ be two constants. Then there exists an absolute constant $c = c(\beta, \theta)$ such that the following holds:*

If for every point $p \in P$ there is a $(\beta L, \theta)$-sector centered at p which contains no other point from P, then the optimum TSP tour on P has length $|\textsc{Opt}| \leq cL$.

Proof. Let P be a point set satisfying the conditions of the lemma, i.e., for every point $p \in P$ there is a $(\beta L, \theta)$-sector $S(p)$, centered at p, which contains no other point from P. We begin by partitioning the set of sectors $\mathcal{S} = \{S(p) : p \in P\}$ into $k = hg$ groups, depending on their orientations and locations, where $h = \lceil 2\pi/\theta \rceil$ and $g = \lceil \sqrt{2}/(\beta \cos \theta) \rceil$. The precise partitioning is done as follows. Fix h directions, $\xi_1, \ldots, \xi_h$, where ξ_i, for $i \in [h]$, makes an angle of $(i-1)\theta$ with the horizontal direction. For each direction ξ_i we partition the covering box into g parallel slabs $\rho_{i,j}$ $(j = 1, \ldots, g)$ of equal width along the direction $\xi_i^{\perp}$ orthogonal to ξ_i. See Figure 4-(a) for an example. For $i \in [h]$ and $j \in [g]$, let $\mathcal{S}_{i,j} \subseteq \mathcal{S}$ be the set of sectors $S(p)$ with the following two properties (see Figure 4-(b)): (i) $p \in \rho_{i,j}$, and (ii) the line through p with direction ξ_i intersects the circular arc part of $S(p)$.

Since $h = \lceil 2\pi/\theta \rceil$ we can find for each $p \in P$ a direction ξ_i such that (ii) is satisfied. Clearly (i) is satisfied for some value j given the direction ξ_i for p. Hence, $\bigcup_{i \in [h],\ j \in [g]} \mathcal{S}_{i,j} = \mathcal{S}$.

Claim 1 *For any $i \in [h]$ and $j \in [g]$, the set of points $\{p \in P : S(p) \in \mathcal{S}_{i,j}\}$ can be covered by at most two paths T' and T'', of length at most $2(\sqrt{2}+\beta)L/\sin(\theta/2)$ each.*

Proof. By performing the appropriate rotation, we may assume without loss of generality that ξ_i is the vertical direction, and thus the slab $\rho_{i,j}$ is horizontal. Since the diameter of the covering box is $\sqrt{2}L$, the width of such a slab is at most $\sqrt{2}L/g \leq L\beta \cos\theta$. In particular, if we consider any point $p \in P$ such that $S(p) \in \mathcal{S}_{i,j}$, then the circular arc of $S(p)$ lies completely outside $\rho_{i,j}$ (see Figure

4-(b)), and thus the boundary of the intersection of $S(p)$ and $\rho_{i,j}$ is a triangle $\Delta(p)$, with head angle θ. A line passing through p parallel to the direction ξ_i divides this triangle into two, one on the left $\Delta'(p)$ and one on the right $\Delta''(p)$ of the line (see Figure 4-(b)). Clearly, the angle with head p in one of these triangles is at least $\theta/2$. Now we partition $\mathcal{S}_{i,j}$ further into two groups of sectors: $\mathcal{S}'_{i,j}$ is the set of sectors $S(p)$ whose left triangle $\Delta'(p)$ makes an angle of at least $\theta/2$ with the vertical direction, and $\mathcal{S}''_{i,j} = \mathcal{S} \setminus \mathcal{S}'_{i,j}$.

We claim that there is a path T' connecting all the points in $P' = \{p \in P \,:\, S(p) \in \mathcal{S}'_{i,j}\}$, with total length

$$|T'| \leq (\sqrt{2}+\beta)L\left(\frac{1+\cos(\theta/2)}{\sin(\theta/2)}\right) \leq 2(\sqrt{2}+\beta)L/\sin(\theta/2). \tag{4}$$

To see this, we may assume without loss of generality that each triangle $\Delta'(p)$, for $p \in P'$ makes an angle of exactly $\theta/2$ with the vertical direction. The path T' is obtained by traversing the boundary of these triangles from left to right as shown in Figure 5-(a). By projecting the sides of each such triangle on the big dotted triangle Δ_0 containing all of them (see Figure 5-(a)), we observe that the sum of all these lengths is at most the sum of the two non-horizontal sides of Δ_0, which in turn implies (4). Applying the same argument for $\mathcal{S}''_{i,j}$ and connecting both paths by a segment of length at most $\sqrt{2}L$ implies Claim 1. □

Now construct a connected Eulerian graph by taking the minimum covering box, together with two copies of each of the at most $2hg$ paths defined in the claim above, but extended to start and end at the covering box, which adds at most L for each path. The total length is at most

$$4L + 4\left(\frac{2(\sqrt{2}+\beta)}{\sin(\theta/2)} + 1\right)Lhg \stackrel{\text{def}}{=} c(\beta,\theta)L, \tag{5}$$

where $h = \lceil 2\pi/\theta \rceil$ and $g = \lceil \sqrt{2}/(\beta\cos\theta) \rceil$. □

Notice that the upper bound in the previous lemma does not depend on the number of points. An infinite set of points could still satisfy the condition of the lemma. The next lemma is an analogue for the TSP with intersecting neighborhoods. As we shall see later, conditions (i) and (ii) of the lemma can be easily satisfied if we select any minimal hitting set for the given set of objects $\mathcal{O}$.

Lemma 7 *Let B be a box of size L containing a set of points $P = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^2$. Assume that there is a collection of n convex α-fat objects $\mathcal{O} = \{O_1, \ldots, O_n\}$, each of diameter δ, such that: (i) each point $p \in P$ is contained in exactly one object $O(p) \in \mathcal{O}$, (ii) each object O contains exactly one point $p(O) \in P$. Then there exists a tour T on P with length $O(L^2\alpha^2/\delta + \alpha^2 L)$.*

Figure 5: (a) Bounding the tour length in the proof of Lemma 6. (b) Illustration for the proof of Lemma 7.

Proof. Consider an object O with its unique point $p = p(O) \in P$. We will prove that there is $(\beta L, \theta)$-sector with center p that lies completely inside O, with $\theta = 2\pi/(3\alpha)$ and $\beta = \delta/(4L)$.

Let p' be a point in O at maximum distance, say R, from p (See Figure 5-b). Obviously, $R \geq \delta/2$. Let s be the point in the middle of line segment pp' and consider a disk D with center p' and radius $R/2$. Let u and v be points in O on the circumference of D such that the angle $\phi = \langle upv \rangle$ is maximum. Finally, denote by S the (R, ϕ)-sector passing through u and v, with center p, radius R, and head angle ϕ.

We use $A(U)$ to denote the area of a given region U of the plane. By definition of fatness we have $A(O \cap D) \geq \pi(R/2)^2/\alpha$. Further, $A(S \cap D) \leq 3/4 \cdot A(S) = 3/4 \cdot R^2\phi/2$. It follows from the definition of u, v, and p', and the convexity of O that $O \cap D \subseteq S \cap D$. Therefore, $\pi(R/2)^2/\alpha \leq A(O \cap D) \leq A(S \cap D) \leq 3R^2\phi/8$. Hence, $\pi/\alpha \leq 3\phi/2$, implying $\phi \geq 2\pi/(3\alpha)$. The sector with center p, radius $R/2$ and head angle upv is contained in O.

Now we apply bound (5) from Lemma 6 with $\theta = 2\pi/(3\alpha)$ and $\beta = \delta/(4L)$. Note that as $\alpha \to \infty$, $\cos(\theta/2) \to 1$ and $\sin(\theta/2) = O(1/\alpha)$. We conclude that the length of the optimum tour is $O(L^2\alpha^2/\delta + \alpha^2 L)$. □

5.2 The algorithm

Now we give the algorithm for constructing an $O(1)$-approximate TSPN tour for intersecting convex α-fat objects with comparable sizes. We note that this algorithm also works for the continuous case, with comparable size neighborhoods. Thus our result in this section extends the $O(1)$-approximation for the continuous case of intersecting unit disks, given in [5].

Algorithm Dtspn-Intersecting:

(1) Compute a minimum covering box B of $\mathcal{O}$.

(2) Find a minimal hitting pointset $P' \subseteq P$ for $\mathcal{O}$ inside B.

(3) Compute a $(1+\epsilon)$-approximate TSP tour on P'.

Remark 4 *Algorithm* Dtspn-Intersecting *can be implemented in time* $\tilde{O}(n|P|)$, *assuming we use the simple* $(1+\epsilon)$*-approximation of the minimum covering box, as explained in Section 2.*

5.3 The analysis

Theorem 4 *Algorithm* DTSPN-INTERSECTING *is an $O(\alpha^3)$-approximation algorithm for the Discrete TSPN problem, with convex and α-fat neighborhoods of the same diameter.*

Proof. The algorithm constructs a $(1+\epsilon)$-approximate TSP-tour on a hitting pointset P'. It will be enough to show that there exists a tour T on P' whose total cost is within $O(\alpha^3)$ of the optimum for $\mathcal{O}$.

To P' we can associate a subset of the objects $\mathcal{O}' \subseteq \mathcal{O}$ with the property that $|P' \cap O| = 1$ for all $O \in \mathcal{O}'$ and $|\{O \in \mathcal{O}' \ : \ p \in O\}| = 1$ for all $p \in P'$. The set $\mathcal{O}'$ can be found as follows. By the minimality of P', for every point $p \in P'$ there exists an object $O(p) \in \mathcal{O}$ such that $O \cap P' = \{p\}$. Let $\mathcal{O}' = \{O(p) \ : \ p \in P'\}$. For an object $O \in \mathcal{O}'$, let us denote by $p(O)$ the (unique) point of P' contained in O. The sets P' and $\mathcal{O}'$ satisfy the conditions of Lemma 7. Denote by L the size of the minimum covering box B. Recall by Lemma 2 that $2L \leq |\text{OPT}|$.

If $L \leq \delta$ then the tour guaranteed by Lemma 7 has length $f(\alpha) \cdot L \leq f(\alpha)|\text{OPT}|$, where $f(\alpha) = O(\alpha^2)$.

Now assume $L \geq \delta$. Let $\mathcal{I} \subseteq \mathcal{O}'$ be a *maximal independent set* of objects in $\mathcal{O}'$, i.e., a maximal collection of pairwise disjoint objects. By the maximality of $\mathcal{I}$, every object in $\mathcal{O}'$ must intersect some object in $\mathcal{I}$. Let $\{\mathcal{O}_I \subseteq \mathcal{O}' \ : \ I \in \mathcal{I}\}$ be a partition of $\mathcal{O}'$, such that any $\mathcal{O}_I$ contains only objects intersecting I. For an arbitrary set of objects $U \subseteq \mathcal{O}'$ let us denote $p(U) = \{p(O) \ : \ O \in U\}$.

For any $I \in \mathcal{I}$ all objects in $\mathcal{O}_I$ lie in a square of size at most 3δ. By Lemma 7 there exists a tour T_I on $p(\mathcal{O}_I)$ of length $f(\alpha) \cdot \delta$, for some $f(\alpha) = O(\alpha^2)$.

Let $T_{\mathcal{I}}$ be an optimal TSP tour on $p(\mathcal{I})$. Adding $T_{\mathcal{I}}$ to the tours T_I $(I \in \mathcal{I})$ gives, after shortcutting, a tour (see Figure 6) on $p(\mathcal{O}')$ of length at most

$$|T_{\mathcal{I}}| + f(\alpha) \cdot \delta|\mathcal{I}|.$$

Let $\text{OPT}_{\mathcal{I}}$ be the optimal tour on $\mathcal{I}$ which can use any point from P to connect the objects, then, clearly, $|T_{\mathcal{I}}| \leq |\text{OPT}_{\mathcal{I}}| + 2|\mathcal{I}|\delta$. Hence, the tour T constructed by the algorithm has length

$$|T| \leq |\text{OPT}_{\mathcal{I}}| + (f(\alpha) + 2) \cdot \delta|\mathcal{I}|.$$

On the other hand, by Lemma 1, we know that $|\text{OPT}_{\mathcal{I}}| \geq (|\mathcal{I}|/\alpha - 1)\pi\delta/4$, and thus $\delta|\mathcal{I}| \leq 4|\text{OPT}_{\mathcal{I}}|\alpha/\pi + \delta\alpha$. Combining these together, and using $\delta \leq L \leq |\text{OPT}|$ and $|\text{OPT}_{\mathcal{I}}| \leq |\text{OPT}|$, we get

$$\begin{aligned} |T| &\leq |\text{OPT}_{\mathcal{I}}| + (f(\alpha)+2)\delta|\mathcal{I}| \\ &\leq \left(\alpha(f(\alpha)+2)(\frac{4}{\pi}+1)+1\right)|\text{OPT}| \\ &= O(\alpha^3)|\text{OPT}|. \end{aligned}$$

□

Figure 6: Constructing an approximate tour as in the proof of Theorem 4.

6 General objects - the continuous case

In this section we consider the case when we have arbitrary connected neighborhoods and are allowed to hit each neighborhood at any point inside it. If we do not restrict the shape of the objects then no better approximation algorithm than $O(\log n)$ is known. Gudmundsson and Levcopoulos [10] used a guillotine subdivision of the plane to obtain an algorithm which runs in time $O(N^2 \log N)$, where N is the total number of vertices of the polygonal objects. Here, we give a very simple approximation algorithm, with the same approximation guarantee, which does not require the objects to be simple polygons.

6.1 The algorithm

As before, we denote the objects by $O_1, \ldots, O_n$ and their respective diameters by $\delta_1, \ldots, \delta_n$.

Lemma 8 *Let $\mathcal{O}$ be a set of connected objects with diameter at least δ and a minimum covering box of size at most L. Then there exists a TSP tour T of length $|T| \leq 4\left[\frac{L}{\delta}\sqrt{2}+2\right]L$, connecting all the objects in $\mathcal{O}$.*

Proof. Consider a covering box B of size at most L. A grid G of granularity $\delta/\sqrt{2}$ in $\mathcal{B}$ covers all the objects in $\mathcal{O}$, and the total length of all the lines in G is at most $2(L\sqrt{2}/\delta+2)L$. By doubling each line of G we can build a TSP tour with length as stated in the lemma. □

For two objects $O_i, O_j \in \mathcal{O}$, define their distance $d(O_i, O_j) \stackrel{\text{def}}{=} \min\{d(x,y) : x \in O_i,\ y \in O_j\}$, and for an object $O_i \in \mathcal{O}$ and $r \in \mathbb{R}_+$, define the *r-neighborhood* of O_i to be $N(O_i, r) \stackrel{\text{def}}{=} \{O_j \in \mathcal{O} : d(O_i, O_j) \leq r, \delta_i \leq \delta_j\}$. Finally, we fix a constant c, and for $i = 1, \ldots, n$, define the neighborhood of O_i as $N(O_i) \stackrel{\text{def}}{=} N(O_i, c\delta_i)$.

Algorithm TSPN-GENERAL:

(1) For $i = 1, 2, \ldots$, let Q_i be the *smallest* diameter object in $\mathcal{O}$ not belonging to $N(Q_1) \cup N(Q_2) \cup \ldots \cup N(Q_{i-1})$. Let k be the largest value of i for which such Q_i exists.

(2) Let B be a minimum covering box for $\mathcal{O}$, and B_i be a minimum covering box for $\{O \cap B \ : \ O \in N(Q_i)\}$, for $i = 1, \ldots, k$.

(3) For $i = 1, \ldots, k$, construct the TSP tour T_i guaranteed by Lemma 8 on the set of objects $N(Q_i)$ with the covering box B_i.

(4) Construct a $(1 + \epsilon)$-approximate TSP tour T_0 on $\{p_1, \ldots p_k\}$, where $p_i \in Q_i$ is on tour T_i.

(5) Combine the tours $T_0, T_1, \ldots, T_k$ into a single TSP tour T (using shortcutting on the Eulerian graph defined by them).

(6) Output the shorter of T and the tour implied by Lemma 8 on the set $\mathcal{O}$ with covering box B.

Remark 5 *Algorithm* TSPN-GENERAL *can be implemented in time* $\tilde{O}(n^2)$*, assuming*[5] *we use the simple* $(1 + \epsilon)$*-approximations of the minimum covering boxes, as explained in Section 2.*

6.2 The analysis

Theorem 5 *Algorithm* TSPN-GENERAL *has approximation factor*

(i) $O(\log n)$ *if each neighborhood is a connected region, and*

(ii) $O(1)$ *if all the neighborhoods are connected and have the same (or comparable) diameters.*

Proof. Let OPT and OPT$'$ be respectively optimal TSPN tours on $\mathcal{O}$ and on $\mathcal{O}' = \{Q_1, \ldots, Q_k\}$. Note that $\delta_1 \leq \delta_2 \leq \ldots \leq \delta_k$. Let $L, L_1, \ldots, L_k$ be respectively the sizes of the minimum covering boxes $B, B_1, \ldots, B_k$.

To prove (i) we first establish the following claim.

Claim 2 *If* $k \geq 2$ *then* $|\text{OPT}'| > c \sum_{i=1}^{k-1} \delta_i / \log k$.

Proof. If $i < j$ then $\delta_i \leq \delta_j$ and $Q_j \notin N(Q_i)$. By the definition of the neighborhoods we have $d(Q_i, Q_j) > c\delta_i$. If $k = 2$ then $|\text{OPT}'| \geq 2d(Q_1, Q_2) > 2c\delta_1$. We continue the proof by induction on k. Fix an orientation of OPT$'$ and fix for each i a connection point $q_i \in Q_i$ on OPT$'$. Define W_i as the arc directed from q_i to the next connection point on the tour. For any $i = 1, \ldots, k$, let $Q_{h(i)}$

[5] we also assume that the intersection of an object with a disk can be done in $O(1)$ time

be an object with the smaller diameter among the two objects on the arc W_i. Then, by the same argument as in the first two lines of the proof, $|W_i| > c\delta_{h(i)}$ for $i = 1, \ldots, k$. Consequently, $|\text{OPT}'| = \sum_{i=1}^{k} |W_i| > c \sum_{i=1}^{k} \delta_{h(i)}$.

Let $\mathcal{O}'' = \mathcal{O}' \setminus \{Q_{h(i)} : i = 1, \ldots, k\}$ and OPT'' be an optimum TSPN tour connecting the objects in $\mathcal{O}''$. By induction we have $|\text{OPT}''| > c \sum_{i:Q_i \in \mathcal{O}'', i \neq k} \delta_i / \log(|\mathcal{O}''|)$ (we assume w.l.o.g. that the $Q_k \in \mathcal{O}''$), and thus

$$|\text{OPT}'| + \log(|\mathcal{O}''|)|\text{OPT}''| > c \sum_{i \in \{1,\ldots,k\}} \delta_{h(i)} + c \sum_{i:Q_i \in \mathcal{O}'', i \neq k} \delta_i.$$

Note that $|\mathcal{O}''| \leq |\mathcal{O}'|/2$ and $|\text{OPT}''| \leq |\text{OPT}'|$. Hence, the left side is at most $\log(|\mathcal{O}'|)|\text{OPT}'|$. Note further that $\sum_{i \in \{1,\ldots,k\}} \delta_{h(i)} = \sum_{i:Q_i \in \mathcal{O}' \setminus \mathcal{O}''} \delta_i$. This gives the claimed bound. □

Claim 3 $|T_0| \leq (1 + \sqrt{2})|\text{OPT}| + 2 \sum_{i=1}^{k-1} \delta_i$.

Proof. Let OPT'' be an optimum tour on $\{Q_1, \ldots, Q_{k-1}\}$ and q_i be a point in $Q_i \cap \text{OPT}''$, for $i = 1, \ldots, k-1$. Then the union of OPT'' with two copies of each of the segments $p_i q_i$, for $i = 1, \ldots, k-1$, and of $p_1 p_k$, forms a connected Eulerian graph that can be shortcut to a TSP tour T' on $\{p_1, \ldots, p_k\}$ of length

$$\begin{aligned} |T'| &\leq |\text{OPT}''| + 2 \sum_{i=1}^{k-1} d(p_i, q_i) + 2d(p_1, p_k) \\ &\leq (1 + \sqrt{2})|\text{OPT}| + 2 \sum_{i=1}^{k-1} \delta_i, \end{aligned}$$

since $|\text{OPT}''| \leq |\text{OPT}|$, $d(p_i, q_i) \leq \delta_i$, and $d(p_1, p_k) \leq \sqrt{2}L \leq |\text{OPT}|/\sqrt{2}$. □

All objects in the set $N(Q_i)$ have diameter at least δ_i and their minimum covering box has size $L_i \leq (2c+1)\delta_i$. Thus Lemma 8 gives, for $i = 1, \ldots, k$,

$$|T_i| \leq 4[\frac{L_i}{\delta_i}\sqrt{2} + 2]L_i \leq 4[(2c+1)\sqrt{2} + 2]L_i. \tag{6}$$

Now, the length of the tour T returned by the algorithm can be bounded as follows:

$$\begin{aligned} |T| &\leq |T_0| + \sum_{i=1}^{k} |T_i| \\ &\leq (1 + \sqrt{2})|\text{OPT}| + 2 \sum_{i=1}^{k-1} \delta_i + 4[(2c+1)\sqrt{2} + 2] \sum_{i=1}^{k} L_i \\ &\leq (1 + \sqrt{2})|\text{OPT}| + \Big[2 + 4(2c+1)[(2c+1)\sqrt{2} + 2]\Big] \sum_{i=1}^{k-1} \delta_i \\ &\quad +4[(2c+1)\sqrt{2} + 2]L_k, \end{aligned}$$

using $L_i \le (2c+1)\delta_i$ for the last inequality. Next we use $L_k \le L \le |\textsc{Opt}|$ and Claim 2 to get

$$|T| \le (f_1(c) + f_2(c)\log k)|\textsc{Opt}|,$$

where

$$\begin{aligned} f_1(c) &= 9 + 5\sqrt{2} + 8\sqrt{2}c, \\ f_2(c) &= \frac{1}{c}[16\sqrt{2}c^2 + (16\sqrt{2}+16)c + 10 + 4\sqrt{2}]. \end{aligned}$$

The value of $c = 0.832..$ minimizes $f_2(c)$ and gives the approximation ratio of $25.5 + 76.3 \log k$.

Now we prove (ii).

Let δ be the diameter of the smallest object, and assume that the diameters of all other objects are bounded by $\rho\delta$, for some constant $\rho \ge 1$. We require that the constant c used in the definition of the neighborhood is at least $\sqrt{2}(1+\rho)$. In case the objects have comparable diameters, we can strengthen Claim 2 as follows.

Claim 4 $|\textsc{Opt}'| \ge \left[\frac{k}{4} - 1\right] \frac{\pi\delta}{4}\sqrt{2}$.

Proof. Let D_i, for $i = 1, \ldots, k$, be a disk of diameter $\delta_i\sqrt{2}$, enclosing object Q_i. We observe that, since $c \ge \sqrt{2}(1+\rho)$, all these disks are disjoint. (Otherwise, there exist two indices $i < j$, such that $D_i \cap D_j \ne \emptyset$, implying that $d(Q_i, Q_j) \le \sqrt{2}(\delta_i + \delta_j) \le \sqrt{2}(1+\rho)\delta_i \le c\delta_i$, and contradicting the fact that $Q_j \notin N(Q_i)$.) Thus we can apply Lemma 1 to the disks $D_1, \ldots, D_k$, (using $\alpha = 4$ for disks) to conclude that any TSP tour connecting these disks, and hence connecting the objects inside them, must have length bounded as stated in the claim. □

To bound the tour T returned by the algorithm, we observe that $|T_0| \le |\textsc{Opt}| + 2\sum_{i=1}^{k} \delta_i$, and combine this with (6) and Claim 4 to get

$$\begin{aligned} |T| &\le |T_0| + \sum_{i=1}^{k} |T_i| \\ &\le |\textsc{Opt}| + 2\sum_{i=1}^{k} \delta_i + 4[(2c+1)\sqrt{2} + 2] \sum_{i=1}^{k} L_i \\ &\le |\textsc{Opt}| + \left[2 + 4(2c+1)[(2c+1)\sqrt{2} + 2]\right] \sum_{i=1}^{k} \delta_i \\ &\le |\textsc{Opt}| + \left[2 + 4(2c+1)[(2c+1)\sqrt{2} + 2]\right] \rho\delta k \\ &\le \left[1 + \left[2 + 4(2c+1)[(2c+1)\sqrt{2} + 2]\right] 2\rho \left(\frac{4\sqrt{2}}{\pi} + 1\right)\right] |\textsc{Opt}| \\ &= O(\rho^3)|\textsc{Opt}|, \end{aligned}$$

assuming that $L \geq \delta$ (otherwise, a tour of length at most $2(\sqrt{2}+2)|\textsc{Opt}|$ is guaranteed by Lemma 8 and Step (6) of the algorithm). $\square$

Remark 6 *The lower bound in Claim 2 is tight up to a constant factor, see [15].*

References

[1] E. M. Arkin and R. Hassin, *Approximation algorithms for the geometric covering salesman problem*, Discrete Applied Mathematics **55** (1994), no. 3, 197–218.

[2] S. Arora, *Nearly linear time approximation schemes for euclidean TSP and other geometric problems*, J. ACM **45** (1998), no. 5, 1–30.

[3] Andrea E. F. Clementi, Pierluigi Crescenzi, and Gianluca Rossi, *On the complexity of approximating colored-graph problems*, Proc. 5th Internat. Computing and Combinatorics Conference, Lecture Notes in Computer Science, vol. 1627, Springer, 1999, pp. 281–290.

[4] 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*, J. of Algorithms **57** (2005), 22–36.

[5] A. Dumitrescu and J.S.B. Mitchell, *Approximation algorithms for TSP with neighborhoods in the plane*, J. Algorithms **48** (2003), no. 1, 135–159.

[6] K. Elbassioni, A.V. Fishkin, N. Mustafa, and R. Sitters, *Approximation algorithms for Euclidean group TSP*, Proc. 15th Internat. Coll. on Automata, Languages and Programming, Lecture Notes in Computer Science, vol. 3580, Springer, 2005, pp. 1115–1126.

[7] K. Elbassioni, A.V. Fishkin, and R. Sitters, *On approximating the tsp with intersecting neighborhoods*, ISAAC, 2006, pp. x–y, to appear.

[8] J. Fakcharoenphol, S. Rao, and K. Talwar, *A tight bound on approximating arbitrary metrics by tree metrics*, STOC, 2003, pp. 448–455.

[9] N. Garg, G. Konjevod, and R. Ravi, *A polylogarithmic approximation algorithm for the group Steiner tree problem*, J. Algorithms **37** (2000), no. 1, 66–84.

[10] J. Gudmundsson and C. Levcopoulos, *A fast approximation algorithm for TSP with neighborhoods*, Nordic J. Computing **6** (1999), no. 4, 469–488.

[11] E. Halperin and R. Krauthgamer, *Polylogarithmic inapproximability*, Proc. 35th Annual ACM Symposium on Theory of Computing, 2003, pp. 585–594.

[12] C.S. Mata and J.S.B. Mitchell, *Approximation algorithms for geometric tour and network design problems (extended abstract)*, Proc. 11th Annual ACM Symposium on Computational Geometry, 1995, pp. 360–369.

[13] J.S.B. Mitchell, *Guillotine subdivions approximate polygonal subdivisons: A simple polynomial-time approximation scheme for geometric TSP, k-MST and related problems*, SIAM J. Computing **28** (1999), no. 4, 1298–1309.

[14] ______, *Handbook of computational geometry*, ch. Geometric shortest paths and network optimization, pp. 633–701, Elsevier, North-Holland, Amsterdam, 2000.

[15] ______, *A PTAS for TSP with neighborhoods among fat regions in the plane*, Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms, 2007, To appear.

[16] G. Reich and P. Widmayer, *Beyond Steiner's problem: a VLSI oriented generalization*, Proc. 15th Int. Workshop on Graph-theoretic Concepts in Computer Science, Springer, 1990, pp. 196–210.

[17] S. Safra and O. Schwartz, *On the complexity of approximating TSP with Neighborhoods and related problems*, Proc. 11th Annual European Symposium on Algorithms, Lecture Notes in Computer Science, vol. 2832, Springer, 2003, pp. 446–458.

[18] P. Slavik, *The errand scheduling problem*, Tech. report, SUNY, Buffalo, USA, 1997.

[19] A. F. van der Stappen, *Motion planning amidst fat obstacles*, Ph.d. dissertation, Utrecht University, Utrecht, the Netherlands, 1994.