

# TSP with Neighborhoods of Varying Size

Mark de Berg

*Department of Computer Science, TU Eindhoven, the Netherlands.*

Joachim Gudmundsson

*NICTA, Sydney, Australia.*<sup>1</sup>

Matthew J. Katz<sup>2</sup>

*Department of Computer Science, Ben-Gurion University, Israel.*

Christos Levcopoulos

*Department of Computer Science, Lund University, Sweden.*

Mark H. Overmars

*Department of Information and Computing Sciences, Utrecht University, the Netherlands.*

A. Frank van der Stappen

*Department of Information and Computing Sciences, Utrecht University, the Netherlands.*

---

## Abstract

In TSP with neighborhoods (TSPN) we are given a collection  $\mathcal{S}$  of regions in the plane, called *neighborhoods*, and we seek the shortest tour that visits all neighborhoods. Until now constant-factor approximation algorithms have been known only for cases where the neighborhoods are of approximately the same size. In this paper we present the first polynomial time constant-factor approximation algorithm for disjoint convex fat neighborhoods of arbitrary size. We also show that in the general case, where the neighborhoods can overlap and are not required to be convex or fat, TSPN is APX-hard and cannot be approximated within a factor of  $391/390$  in polynomial-time, unless  $P=NP$ .

*Key words:* Computational geometry, approximation algorithms, TSP with neighborhoods

---

## 1 Introduction

The Traveling Salesman Problem (TSP) is one of the best known and most widely studied optimization problems. In the Euclidean setting, of this problem, one is given a set of  $n$  points in the plane (or a higher-dimensional Euclidean space), and one wants to find the tour—that is, closed curve—of shortest length that visits all points from the set. A natural generalization of the Euclidean TSP, first studied by Arkin and Hassin [3], is the TSP-with-neighborhoods problem (TSPN): given a collection of  $n$  regions in the plane, called *neighborhoods*, find a shortest tour that visits all neighborhoods. In principle, a neighborhood can be any subset of the plane, although in the papers dealing with TSPN the neighborhoods are often restricted to simple shapes like disks or squares—see Figure 1 for an example. (Stated in the traditional TSP setting, one can phrase the problem as follows. A salesman wants to meet some potential buyers, each of whom specifies a region in the plane where he is willing to meet the salesman. The salesman now wants to find a shortest tour such that he can meet with all potential buyers.) Since TSPN is a generalization of the Euclidean TSP, it is NP-hard [8,16].




Fig. 1. An example of TSPN when the input neighborhoods are axis-parallel squares.

In the case when the neighborhoods may be disconnected then the problem is called the One-of-a-Set TSP [14], and the tour is required to visit at least one point from each set. It has also been referred to as the Errand Scheduling Problem [19], since it models the problem of finding the best order in which to perform a set of errands, each of which can be performed at some subset of the nodes of a graph. This problem has applications in communication-network

---

*Email addresses:* [m.t.d.berg@TUE.nl](mailto:m.t.d.berg@TUE.nl) (Mark de Berg),  
[joachim.gudmundsson@nicta.com.au](mailto:joachim.gudmundsson@nicta.com.au) (Joachim Gudmundsson),  
[matya@cs.bgu.ac.il](mailto:matya@cs.bgu.ac.il) (Matthew J. Katz), [christos@cs.lth.se](mailto:christos@cs.lth.se) (Christos  
Levcopoulos), [markov@cs.uu.nl](mailto:markov@cs.uu.nl) (Mark H. Overmars), [frankst@cs.uu.nl](mailto:frankst@cs.uu.nl) (A.  
Frank van der Stappen).

<sup>1</sup> National ICT Australia is funded by the Australian Government’s Backing Australia’s Ability initiative, in part through the Australian Research Council.

<sup>2</sup> Supported by grant no. 2000160 from the U.S.-Israel Binational Science Foundation.

design [10], VLSI routing [17], and in the manufacturing industry [18]. For more information see the survey by Mitchell [14].

Mata and Mitchell [12] provided a general framework that gives an  $O(\log k)$ -approximation algorithm for TSPN with time complexity  $O(n^5)$ , where  $k$  is the number of neighborhoods and  $n$  is the total complexity of the input. The time complexity was later improved by Gudmundsson and Levcopoulos [9] to  $O(n^2 \log n)$ . Although no algorithms with a better approximation factor are known for the general TSPN problem, they are known for various special cases. For example, Arkin and Hassin [3] gave  $O(1)$ -approximation algorithms for the following cases: parallel unit-length segments, translates of a convex polygonal neighborhoods and, more generally, for neighborhoods with diameter segments that are parallel to a common direction and a ratio between the longest and the shortest diameter that is bounded by a constant. Later, Dumitrescu and Mitchell [7] extended and improved these results to include neighborhoods with comparable diameter, unit disks, and infinite lines. Gudmundsson et al. [9] gave a polynomial-time approximation scheme (PTAS) for the special case when the tour is short compared to the size of the neighborhoods.

The results quoted above leave the main question in this area still open: is there a constant-factor approximation, or perhaps even a polynomial-time approximation scheme, for TSPN? In this paper we present two results, which give a partial answer to this question.

First of all, in Section 2 we use a reduction from the Vertex-Cover Problem to show that TSPN is APX-hard and cannot be approximated within a factor of  $391/390$  in polynomial time, unless  $P=NP$ . Thus, a polynomial-time approximation scheme does not exist for TSPN, unless  $P=NP$ . This is in contrast to the standard Euclidean TSP, for which there is a polynomial-time approximation scheme allowing one to compute, for any given  $\varepsilon > 0$ , a  $(1 + \varepsilon)$ -approximation in  $O(n^{O(1/\varepsilon)})$  time [4,15].

Second, in Section 3 we give a new constant-factor approximation algorithm for TSPN in the case when the regions are disjoint convex fat neighborhoods of constant complexity. Although we are not able to solve the general case, our algorithm is able to handle cases that could not be handled before. In particular, all special cases that were successfully solved until now have one common property: it is required that the input neighborhoods have roughly the same size. The main contribution of our paper is a polynomial-time constant-factor approximation algorithm for a class of arbitrary-size input neighborhoods, namely for input sets  $\mathcal{S}$  consisting of  $n$  disjoint convex fat neighborhoods. We also assume that the neighborhoods have constant complexity—their boundaries consists of a constant number of algebraic arcs of bounded maximum degree—so that basic operations such as computing tangent points comput-

ing intersections with a line, and computing distances can be done in  $O(1)$  time.

Before we proceed we introduce some notation. We use  $|\tau|$  to denote the length of a curve  $\tau$ . In particular, for two points  $p$  and  $q$ , we use  $|pq|$  to denote the length of the line segment  $pq$ . Furthermore,  $\text{dist}(\sigma, \sigma')$  denotes the distance between neighborhoods  $\sigma$  and  $\sigma'$ , that is,  $\text{dist}(\sigma, \sigma') := \min_{p \in \sigma, q \in \sigma'} |pq|$ . Finally, for a set  $\mathcal{S}$  of neighborhoods we use  $\mathcal{T}_{opt}(\mathcal{S})$  to denote an optimal (that is, shortest) tour visiting each neighborhood in  $\mathcal{S}$ .

## 2 TSPN is APX-hard

In this section we show that TSPN is APX-hard in the case when the neighborhoods are non-convex and allowed to overlap. We use a reduction from VERTEXCOVER, which is defined as follows. Let  $\mathcal{G} = (V, A)$  be an undirected graph. A vertex cover for  $\mathcal{G}$  is a subset  $V' \subset V$  of vertices such that for any arc  $(u, v) \in A$  we have  $\{u, v\} \cap V' \neq \emptyset$ —that is, any arc has at least one incident vertex in  $V'$ . VERTEXCOVER asks for a vertex cover of minimum cardinality. Berman and Karpinski [5] have shown that VERTEXCOVER is hard to approximate:

**Fact 1** (*Theorem 1 (iv) in [5]*) VERTEXCOVER is APX-hard, and cannot be approximated within a factor of  $79/78$  in polynomial time unless  $P=NP$ , even when the degree of each vertex is bounded by four.

Next we perform the reduction from VERTEXCOVER. Let  $\mathcal{G} = (V, A)$  be an undirected graph with  $n$  vertices of maximum degree at most four. We can restrict ourselves to considering connected graphs with at least  $n$  arcs since the cases when this is not true easily can be handled separately: if the graph is not connected then we can consider each connected component separately, and in the case when the graph is a tree an optimal vertex cover can be found by dynamic programming.

We will construct a set  $\mathcal{S}(\mathcal{G})$  of  $n + |A|$  neighborhoods such that an approximation of  $\mathcal{T}_{opt}(\mathcal{S}(\mathcal{G}))$  can be transformed into an approximation of a minimum vertex cover for  $\mathcal{G}$ . The construction uses two types of neighborhoods: points and objects consisting of a line segment followed by a circular arc followed by another line segment. (The circular arc can easily be replaced by a few line segments.) The second type of neighborhoods will correspond to the arcs of  $\mathcal{G}$ , hence they are called *arc objects*. The points are called *helper objects*.

**The helper objects.** The helper objects are the points located at the vertices of a regular  $n$ -gon of unit edge length.

**The arc objects.** To define the arc objects, we first assign a location to each node of  $\mathcal{G}$ . Note that these locations are not neighborhoods in  $\mathcal{S}(\mathcal{G})$ . The  $n$  node locations are defined as follows. For each edge  $pq$  of the  $n$ -gon defining the helper objects, the point  $x$  outside the  $n$ -gon such that  $|px| = |qx| = 1$  is a node location. We assign each node in  $V$  a unique node location in an arbitrary manner. Let  $C$  be the circle with the same center as the  $n$ -gon and at distance  $2n$  from the node locations. We call the line segment of length  $2n$  connecting a node location to  $C$  the *spoke* of the node. Now the arc object for an arc  $(u, v) \in A$  consists of the spokes of  $u$  and  $v$  and one of the circular arcs (it doesn't matter which one) along  $C$  connecting the spokes.

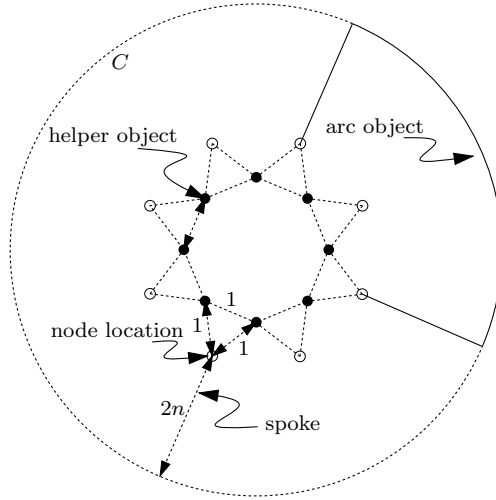


Fig. 2. The construction used in the reduction from VERTEXCOVER to TSPN.

Figure 2 illustrates the construction. The set of helper objects is denoted by  $\mathcal{S}_H$  and the set of arc objects by  $\mathcal{S}_A$ . We let  $\mathcal{S}(\mathcal{G}) := \mathcal{S}_H \cup \mathcal{S}_A$ . In the remainder of this section we prove that if we can approximate  $\mathcal{T}_{opt}(\mathcal{S}(\mathcal{G}))$ , then we can approximate a minimum vertex cover for  $\mathcal{G}$ . From this it will follow that TSPN is APX-hard.

We start by proving a bound on the length of any tour visiting a given subset of the spokes.

**Lemma 2** *Let  $S$  be an arbitrary subset of the spokes in the construction above. Then  $|\mathcal{T}_{opt}(\mathcal{S}_H \cup S)| = n + |S|$ .*

**PROOF.** An optimal tour for  $\mathcal{S}_H \cup S$  has  $n + |S|$  links, each connecting a pair of neighborhoods in  $\mathcal{S}_H \cup S$ . The distance between any two neighborhoods is at least 1, so  $|\mathcal{T}_{opt}(\mathcal{S}_H \cup S)| \geq n + |S|$ . Obviously, a tour of this length also exists: connect (the node location of) each spoke to the two closest helper points and complete the tour by adding the appropriate subset of edges of the  $n$ -gon.

Next we show that an approximation for  $\mathcal{T}_{opt}(\mathcal{S}(\mathcal{G}))$  induces an approximation of a minimum vertex cover for  $\mathcal{G}$ .

**Lemma 3** *Let  $\mathcal{T}$  be a tour visiting each neighborhoods in  $\mathcal{S}(\mathcal{G})$ , such that  $|\mathcal{T}| = c \cdot |\mathcal{T}_{opt}(\mathcal{S}(\mathcal{G}))|$  for some constant  $1 \leq c < 2$ . Then we can construct in polynomial time a vertex cover for  $\mathcal{G}$  whose cardinality is at most  $5c - 4$  times the minimum cardinality of a vertex cover for  $\mathcal{G}$ .*

**PROOF.** Let  $m$  be the minimum cardinality of a vertex cover for  $\mathcal{G}$ . Note that  $m \geq n/4$ , since the degree of the nodes in  $\mathcal{G}$  is bounded by four and since the number of arcs is at least  $n$ . By Lemma 2 this implies that there is a tour for  $\mathcal{S}(\mathcal{G})$  of length  $n+m$ . Since a tour of length  $n+m$  or less can never reach the circular-arc parts of the arc objects, we even have that  $|\mathcal{T}_{opt}(\mathcal{S}(\mathcal{G}))| = n + m$ . Using Lemma 2 again, and the fact that  $\mathcal{T}$  cannot visit the circular-arc parts of the arc objects (because  $c < 2$ ), we see that  $\mathcal{T}$  can visit at most

$$c(n + m) - n = ((c - 1)n/m + c)m \leq (5c - 4)m$$

spokes. The nodes corresponding to these spokes form a vertex cover, which can trivially be found in polynomial time.

The main result of this section now follows easily.

**Theorem 4** *TSP with neighborhoods is APX-hard, and cannot be approximated within a factor of  $391/390$  in polynomial time unless  $P=NP$ .*

**PROOF.** Lemma 3 implies that if we can approximate TSPN within a factor  $391/390$ , then we can approximate VERTEXCOVER (for graphs of maximum degree four) within a factor  $5 \cdot (391/390) - 4 = 79/78$ . Lemma 1, however, states that this is impossible in polynomial time unless  $P=NP$ .

### 3 A constant-approximation algorithm for fat neighborhoods

We now turn our attention to computing an approximation of the optimal tour through a given set of  $\alpha$ -fat disjoint bounded convex neighborhoods of constant complexity. The definition of fatness we use was introduced by Van der Stappen [20]. For convex neighborhoods it is basically equivalent to other definitions [1,2,11,13].

**Definition 5** [20] Let  $\alpha$  be a parameter with  $\alpha \geq 1$ . A neighborhood  $\sigma$  is  $\alpha$ -fat if for any disk  $D$  whose center lies in  $\sigma$  and whose boundary intersects  $\sigma$ , we have  $\text{area}(D) \leq \alpha \cdot \text{area}(\sigma \cap D)$ .

In our setting of bounded planar neighborhoods it turns out that  $\alpha \geq 4$  [20].

We shall need the following property of convex fat neighborhoods, which was proved by Chew *et al.* [6]. Below we will use the notation  $\partial\sigma$  to denote the boundary of  $\sigma$ .

**Lemma 6** [6] Let  $\sigma$  be a convex  $\alpha$ -fat neighborhood, and let  $p, q$  be two points on  $\partial\sigma$ . Then the length of the shorter path from  $p$  to  $q$  along  $\partial\sigma$  is at most  $(1 + 4\alpha/\pi) \cdot |pq|$ .

We also need the following lemma.

**Lemma 7** Let  $\ell$  and  $\ell'$  be two parallel lines intersecting a convex  $\alpha$ -fat neighborhood  $\sigma$ . Let  $p$  and  $q$  be two points between  $\ell$  and  $\ell'$  and on opposite sides of  $\sigma$ . Then

$$\min\{\text{dist}(q, \ell), \text{dist}(q, \ell')\} \leq (1 + 4\alpha/\pi) \cdot |pq|.$$

**PROOF.** Define  $p'$  and  $q'$  to be the two intersection points of  $pq$  with  $\partial\sigma$  closest to  $p$  and  $q$ , respectively—see Figure 3. (Possibly  $p' = p$  and/or  $q' = q$ .) We have




Fig. 3. Illustration for Lemma 7.

$$\begin{aligned} \min\{\text{dist}(q, \ell), \text{dist}(q, \ell')\} &\leq \min\{\text{dist}(q', \ell), \text{dist}(q', \ell')\} + |qq'| \\ &\leq |\text{shorter path along } \partial\sigma \text{ from } p' \text{ to } q'| + |qq'| \\ &\leq (1 + 4\alpha/\pi) \cdot |p'q'| + |qq'| \\ &\leq (1 + 4\alpha/\pi) \cdot (|p'q'| + |qq'|) \\ &\leq (1 + 4\alpha/\pi) \cdot |pq|. \end{aligned}$$

### 3.1 The algorithm

Let  $\mathcal{S}$  be a set of  $n$  disjoint bounded convex  $\alpha$ -fat neighborhoods, for some constant  $\alpha \geq 4$ . The overall approach we take is quite simple: for each neighborhood  $\sigma \in \mathcal{S}$  we compute a collection  $\mathcal{R}_\sigma$  of representative points, and then we compute a short tour that visits (that is, passes through at least one point of) each set  $\mathcal{R}_\sigma$ . The latter step is done using an algorithm by Slavik [19]. His algorithm computes, for a collection of  $n$  planar point sets, in polynomial time a tour that visits each set and whose length is at most  $3k/2$  times the optimal length of such a tour, where  $k$  is the maximum cardinality of the sets. Our goal is thus to construct the sets  $\mathcal{R}_\sigma$  in such a way that (i) their maximum cardinality is bounded by a constant, and (ii) the optimal tour visiting each set  $\mathcal{R}_\sigma$  is not much longer than the optimal tour visiting each neighborhood  $\sigma$ .

#### *The representative points.*

Fix an neighborhood  $\sigma \in \mathcal{S}$ . The set  $\mathcal{R}_\sigma$  of representative points of  $\sigma$  contains at most three types of points: several *corners*, at most one *anchor*, and possibly several *guards*. Next we define these points, and we prove some basic properties of them.

**3.1.0.1 The corners.** Let  $\gamma$  be a small constant such that  $\pi/2 \geq \gamma > 0$  and  $t = 2\pi/\gamma$  is an integer; later we shall see that we need to choose  $\gamma$  so that  $\tan(\gamma/2) \leq \pi/(21\alpha)$ . For some  $i$  with  $0 \leq i < t$ , we let  $p_i$  be a point on  $\partial\sigma$  that has a tangent making an angle  $i\gamma$  with the positive  $x$ -axis. Here the tangent is a directed line, whose direction is consistent with a counterclockwise orientation of  $\partial\sigma$ . For convenience, we assume that the points  $p_i$  are unique; since the neighborhoods have constant complexity—in particular, each neighborhood has only a constant number of straight segments on its boundary—we can easily enforce this by a suitable rotation of the coordinate axes. We call the points  $p_i$  the *corners* of  $\sigma$ . The portions of  $\partial\sigma$  connecting consecutive corners are called *sides* of  $\sigma$ . The side connecting corners  $p_i$  and  $p_{i+1}$  is denoted by  $e_i$ . It can happen that  $p_i = p_{i+1}$ , in which case the side  $e_i$  can be ignored. When necessary, we shall write  $p_i(\sigma)$ ,  $e_i(\sigma)$ , and so on, to indicate to which neighborhood the corner or side belongs.

Before we proceed, we need to introduce one more concept. Consider a side  $e_i$  of  $\sigma$ , connecting the corners  $p_i$  and  $p_{i+1}$ . For a point  $p$  and an angle  $\phi$ , let  $\ell(p, \phi)$  be the line through  $p$  that makes an angle  $\phi$  with the positive  $x$ -axis. Define  $\phi_i := (i + 1/2)\gamma - \pi/2$ . The *slab of  $e_i$* , denoted  $\text{slab}(e_i)$ , is the region bounded by the lines  $\ell(p_i, \phi_i)$  and  $\ell(p_{i+1}, \phi_i)$ . The part of this slab bounded



by  $e_i$ ,  $\ell(p_i, \phi_i)$ , and  $\ell(p_{i+1}, \phi_i)$ , and not containing  $\sigma$ , is called the *semi-slab* of  $e_i$ ; it is denoted  $\text{slab}(e_i)^+$ . Fig. 4 illustrates these definitions.




Fig. 4. Corners, sides and (semi-)slabs.

Next we prove two properties of the partitioning into sides. The first lemma is basically a consequence of the fact that a side is ‘almost straight’—the tangent does not vary too much along a side—and that the slab boundaries are ‘almost perpendicular’ to the side.

**Lemma 8** *Let  $\sigma$  be a convex neighborhood, and  $e_i$  be one of its sides. Let  $\ell$  be any line parallel to and inside  $\text{slab}(e_i)$ , and let  $p := \ell \cap e_i$ . Then for any point  $q \in e_i$  we have*

$$|qp| \leq (1/\cos(\gamma/2)) \cdot \text{dist}(q, \ell).$$

**PROOF.** Observe that there must be a point on  $e_i$  in between  $q$  and  $p$  whose tangent is parallel to  $pq$ . Hence, the line containing  $pq$  makes an angle with  $\ell$  that lies in the range  $[\pi/2 - \gamma/2, \pi/2 + \gamma/2]$ . It follows that  $|qp| \leq (1/\cos(\gamma/2)) \cdot \text{dist}(q, \ell)$ , as claimed.

The next lemma uses the fact that the sides are almost straight, as well as the fact that we deal with fat neighborhoods.

**Lemma 9** *Let  $\sigma$  and  $\sigma'$  be two convex  $\alpha$ -fat neighborhoods such that  $e_i(\sigma')$  lies at least partially inside  $\text{slab}(e_i(\sigma))^+$ . Let  $p$  be any point on  $e_i(\sigma)$  such that the line through  $p$  parallel to  $\text{slab}(e_i(\sigma))$  intersects  $\sigma'$ , and let  $q$  be any point in  $\sigma' \cap \text{slab}(e_i(\sigma))^+$ . If  $\gamma$  is such that  $\tan(\gamma/2) \leq \pi/(21\alpha)$ , then*

$$|pq| \leq c \cdot \text{dist}(e_i(\sigma), e_i(\sigma')),$$

for  $c = 2 + 10\alpha/\pi$ .

**PROOF.** Let  $B$  be the bounding box of  $\sigma' \cap \text{slab}(e_i(\sigma))^+$  with edges parallel and orthogonal to  $\text{slab}(e_i(\sigma))$ . Let  $w_B$  be the width and  $h_B$  be the height of  $B$ . See Figure 5.




Fig. 5. The bounding box  $B$ , the segments  $s$ ,  $s'$ , and  $s''$ , and the lines  $l$  and  $l'$ , and the region  $R$ .

The neighborhood  $\sigma'$  intersects all four edges of  $B$ . Consider the vertical line that cuts  $B$  into two equal halves; this line clearly intersects  $\sigma' \cap \text{slab}(e_i(\sigma))^+$ . Let  $m \in \sigma' \cap \text{slab}(e_i(\sigma))^+$  be a point on this vertical line. Consider the disk  $D$  with diameter  $w_B$  centered at  $m$ ;  $D$  touches both vertical edges of  $B$ . The boundary of  $D$  is intersected by  $\sigma'$  and thus it follows that  $\text{area}(D \cap \sigma') \geq \text{area}(D)/\alpha \geq \pi w_B^2/4\alpha$ . Since  $D \cap \sigma'$  lies completely inside  $B$  we also have  $\text{area}(B) \geq \pi w_B^2/4\alpha$ . As a result we obtain

$$h_B \geq (\pi/4\alpha)w_B. \quad (1)$$

A point  $p$  on  $e_i(\sigma)$  such that the line through  $p$  parallel to  $\text{slab}(e_i(\sigma))^+$  intersects  $\sigma'$  lies inside the imaginary subslab of  $\text{slab}(e_i(\sigma))^+$  bounded by the supporting lines of the vertical edges of  $B$ . We denote by  $h$  the distance from  $p$  to the bottom edge of the bounding box  $B$ . The distance from  $p$  to any point  $q \in \sigma' \cap \text{slab}(e_i(\sigma))^+$  is bounded by the distance from  $p$  to the farthest corner of the bounding box  $B$  of  $\sigma' \cap \text{slab}(e_i(\sigma))^+$ , which in turn is bounded by  $w_B + h$ , so

$$|pq| \leq h + w_B. \quad (2)$$

We shall see that  $h$  cannot be arbitrarily small.

We define three segments  $s$ ,  $s'$ , and  $s''$  orthogonal to and connecting the supporting lines of the vertical edges of  $B$ . The segment  $s$  runs through  $p$  and its left and right endpoints are denoted by  $v$  and  $v'$  respectively;  $s'$  is located at a distance  $w_B \tan(\gamma/2)$  below the top edge of  $B$ ;  $s''$  lies at a distance  $w_B \tan(\gamma/2)$  above the bottom edge of  $B$ . Finally, we let  $l$  be the line through  $v$  defining a clockwise angle of  $\gamma/2$  with  $s$ , and  $l'$  be the line through  $v'$  defining a counterclockwise angle of  $\gamma/2$  with  $s$ .

The next two inequalities follow from a series of similar observations involving the angle of  $\gamma/2$ . As  $\sigma' \cap \text{slab}(e_i(\sigma))^+$  intersects the top edge of  $B$  (implying that there is a point on  $e_i(\sigma)$  that is above the top edge of  $B$ ),  $e_i(\sigma)$  cannot reach the part of  $B$  below  $s'$ , as this would require a too steep descent of  $e_i(\sigma)$ . As a result, the point  $p$  cannot lie below  $s'$  and thus

$$h \geq h_B - w_B \tan(\gamma/2). \quad (3)$$

By definition, the side  $e_i(\sigma')$  touches the bottom edge of  $B$ . The bounds on its tangents dictate that  $e_i(\sigma') \cap \text{slab}(e_i(\sigma))^+$  remains completely inside the portion of  $B$  below  $s''$ . Similarly, the bounds on the tangents of  $e_i(\sigma)$ , which contains the point  $p$ , imply that  $e_i(\sigma)$  remains completely inside the region  $R$  consisting of points that are above  $l$  or above  $l'$ . The distance  $\text{dist}(e_i(\sigma), e_i(\sigma'))$  is therefore bounded from below by the distance between  $R$  (the upper shaded region in Figure 5) and the portion of  $B$  below  $s''$  (the lower shaded rectangle in Figure 5), so

$$\text{dist}(e_i(\sigma), e_i(\sigma')) \geq h \cos(\gamma/2) - 2w_B \sin(\gamma/2). \quad (4)$$

The inequality  $\tan(\gamma/2) \leq \pi/21\alpha$  implies that  $\sin(\gamma/2) \leq \pi/21\alpha$ , and that  $\cos(\gamma/2) > 0.98$ . The ratio

$$\frac{h + w_B}{h \cos(\gamma/2) - 2w_B \sin(\gamma/2)}$$

is decreasing for increasing  $h$  in the range given by inequality (3). As a consequence it attains its largest value when  $h$  is as small as possible. By applying inequalities (1) and (3), and the bounds on  $\sin(\gamma/2)$ ,  $\cos(\gamma/2)$ , and  $\tan(\gamma/2)$  we get

$$\begin{aligned} \frac{h + w_B}{h \cos(\gamma/2) - 2w_B \sin(\gamma/2)} &\leq \frac{((\pi/4\alpha)w_B - (\pi/21\alpha)w_B) + w_B}{0.98((\pi/4\alpha)w_B - (\pi/21\alpha)w_B) - 2(\pi/21\alpha)w_B} \\ &= \frac{17\pi + 84\alpha}{0.98 \cdot 17\pi - 8\pi} \\ &\leq 2 + 10\alpha/\pi \end{aligned}$$

Combination of this bound with inequalities (2) and (4) gives the final result.

**3.1.0.2 The anchor.** Consider a side  $e_i(\sigma)$ . We say that  $e_i(\sigma)$  *dominates* a side  $e_{i'}(\sigma')$  of another neighborhood  $\sigma'$  if  $i = i'$  and  $\sigma' \subset \text{slab}(e_i(\sigma))^+$ . We define an anchor for  $\sigma$  if  $\sigma$  has a side that dominates a side of another neighborhood, as follows. Of all the sides of other neighborhoods that are dominated by any of the sides of  $\sigma$ , take the one whose distance to its dominating side is minimal. Suppose this ‘closest dominated side’ is  $e_j(\sigma^*)$ . Let  $\ell := \ell(p_j(\sigma^*), \phi_j)$ . Note that  $\ell$  is parallel to  $\text{slab}(e_j(\sigma))$ . Then  $\text{anchor}(\sigma)$ , the anchor of  $\sigma$ , is the point  $\ell \cap e_j(\sigma)$ —see Figure 6.




Fig. 6. The anchor of a neighborhood  $\sigma$ .

**3.1.0.3 The guards.** Consider a side  $e_i$  of  $\sigma$ . We call  $e_i$  *extreme* if all other neighborhoods  $\sigma'$  intersect the semi-slab  $\text{slab}(e_i)^+$ . For each extreme side of  $\sigma$ , we add one or two more representative points, called guards. They are defined as follows.

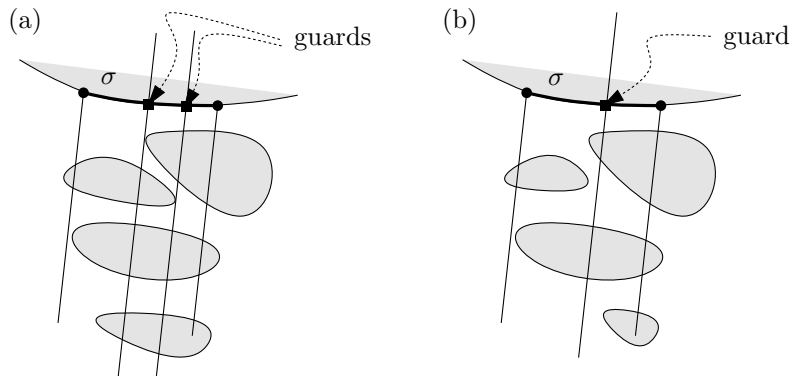


Fig. 7. The guards of a neighborhood  $\sigma$ .

Suppose there are lines parallel to and inside  $\text{slab}(e_i)$  that stab all the neighborhoods in  $\mathcal{S}$ . The collection of all such stabbing lines forms a subslab of

slab( $e_i$ ). We add two guards on  $e_i$  to  $\mathcal{R}_\sigma$ , namely the intersection points of the lines bounding this slab and  $e_i$ —see Figure 7(a).

If there are no stabbing lines, then there must be a line  $\ell$  parallel to and inside slab( $e_i$ ) such that both half-planes defined by  $\ell$  fully contain at least one neighborhood of  $\mathcal{S}$ . (Thus the line  $\ell$  separates at least two neighborhoods.) In this case we add one guard to  $\mathcal{R}_\sigma$ , namely the intersection point of  $\ell$  and  $e_i$ —see Figure 7(b).

### 3.2 The analysis

To show that our algorithm achieves a good approximation factor, we have to show that the optimal tour visiting each set  $\mathcal{R}_\sigma$  is not much longer than the optimal tour visiting each neighborhood  $\sigma$ .

**Theorem 10** *Let  $\mathcal{S}$  be a collection of  $n$  disjoint convex  $\alpha$ -fat neighborhoods in the plane, and let  $\mathcal{R}(\mathcal{S}) := \{\mathcal{R}_\sigma : \sigma \in \mathcal{S}\}$  be the collection of sets of representative points, as defined above. Let  $\mathcal{T}_{opt}(\mathcal{R}(\mathcal{S}))$  be an optimal tour visiting each set  $\mathcal{R}_\sigma \in \mathcal{R}(\mathcal{S})$ , and let  $\mathcal{T}_{opt}(\mathcal{S})$  be an optimal tour visiting each neighborhood  $\sigma \in \mathcal{S}$ . Then*

$$|\mathcal{T}_{opt}(\mathcal{R}(\mathcal{S}))| \leq (1 + c) \cdot |\mathcal{T}_{opt}(\mathcal{S})|,$$

for  $c = (4 + 20\alpha/\pi)t$  (where  $t = 2\pi/\gamma$  depends on  $\alpha$ ).

**PROOF.** The strategy of our proof is to add detours to  $\mathcal{T}_{opt} := \mathcal{T}_{opt}(\mathcal{S})$ , of total length  $c \cdot |\mathcal{T}_{opt}|$ , such that the extended tour visits each set  $\mathcal{R}_\sigma$ .

Let  $p_{start}$  be an arbitrary point on  $\mathcal{T}_{opt}$  outside all the neighborhoods. Imagine traversing the tour (in some arbitrary direction), starting at  $p_{start}$ . We define the *entry* of a neighborhood  $\sigma \in \mathcal{S}$  to be the point on  $\sigma$  where the tour first meets  $\sigma$  during the traversal. If the entry of  $\sigma$  happens to be one of its corners, then the set  $\mathcal{R}_\sigma$  is already visited by  $\mathcal{T}_{opt}$  and no detour has to be added for  $\sigma$ . The neighborhoods for which this is not the case have their entry in the interior of one of their sides. We partition the collection of these neighborhoods into  $t$  subsets,  $\mathcal{S}_0, \dots, \mathcal{S}_{t-1}$ , such that  $\mathcal{S}_i$  contains the neighborhoods  $\sigma$  whose entry lies in the interior of  $e_i(\sigma)$ .

Fix a non-empty subset  $\mathcal{S}_i$ . We claim that we can add detours of total length  $(4 + 20\alpha/\pi)|\mathcal{T}_{opt}|$  to  $\mathcal{T}_{opt}$ , such that each set  $\mathcal{R}_\sigma$  with  $\sigma \in \mathcal{S}_i$  is visited. Proving this claim will establish the theorem. We distinguish two cases.

**Case (i):**  $|\mathcal{S}_i| > 1$ .

Number the neighborhoods in  $\mathcal{S}_i$  as  $\sigma_1, \sigma_2, \dots$  in the order as they are encountered during the traversal of  $\mathcal{T}_{opt}$ . Consider a neighborhood  $\sigma_j \in \mathcal{S}_i$ , and let  $p_i := p_i(\sigma_j)$  and  $p_{i+1} := p_{i+1}(\sigma_j)$ . Denote the entry of  $\sigma_j$  by  $q_j$ , and let  $\mathcal{T}_{opt}(q_j, q_{j+1})$  be the portion of  $\mathcal{T}_{opt}$  from  $q_j$  to  $q_{j+1}$ , with indices taken modulo  $(|\mathcal{S}_i| + 1)$ . We have three subcases.

- If  $q_{j+1} \notin \text{slab}(e_i(\sigma_j))^+$ , then we add the shorter of the two detours  $q_j p_i q_j$  and  $q_j p_{i+1} q_j$  to  $\mathcal{T}_{opt}$ .

If  $q_{j+1} \notin \text{slab}(e_i(\sigma_j))$  then  $\mathcal{T}_{opt}(q_j, q_{j+1})$  crosses one of the bounding lines of  $\text{slab}(e_i(\sigma_j))$  and we can use Lemma 8 to bound the length of the detour by  $2 \cdot (1/\cos(\gamma/2)) \cdot |\mathcal{T}_{opt}(q_j, q_{j+1})|$ . Otherwise  $q_{j+1}$  lies inside  $\text{slab}(e_i(\sigma_j))$  but on the opposite side of  $\sigma_j$  as seen from  $q_j$ . In this case we can use Lemma 6 to bound the length of the detour as follows. Let  $q' \neq q_j$  be the intersection point of  $q_j q_{j+1}$  with  $\partial\sigma_j$ .

$$\begin{aligned} \min\{|q_j p_i q_j|, |q_j p_{i+1} q_j|\} &\leq 2 \cdot |\text{shortest path along } \partial\sigma_j \text{ from } q_j \text{ to } q'| \\ &\leq 2 \cdot (1 + 4\alpha/\pi) \cdot |q_j q'| \\ &\leq 2 \cdot (1 + 4\alpha/\pi) \cdot |q_j q_{j+1}| \\ &\leq 2 \cdot (1 + 4\alpha/\pi) \cdot |\mathcal{T}_{opt}(q_j, q_{j+1})|. \end{aligned}$$

- If  $q_{j+1} \in \text{slab}(e_i(\sigma_j))^+$  and  $\sigma_{j+1}$  lies only partly in  $\text{slab}(e_i(\sigma_j))^+$ , then we add one of the two detours  $q_{j+1} p_i q_{j+1}$  and  $q_{j+1} p_{i+1} q_{j+1}$  to  $\mathcal{T}_{opt}$  as follows. Let  $p \in \{p_i, p_{i+1}\}$  be a corner of  $e_i(\sigma_j)$  such that the bounding line of  $\text{slab}(e_i(\sigma_j))$  that passes through  $p$  intersects  $\sigma_{j+1}$ . We add the detour  $q_{j+1} p q_{j+1}$ , and use Lemma 9 to bound the length of the detour by  $2 \cdot (2 + 10\alpha/\pi) \cdot |\mathcal{T}_{opt}(q_j, q_{j+1})|$ .
- If  $q_{j+1} \in \text{slab}(e_i(\sigma_j))^+$  and  $\sigma_{j+1}$  lies entirely in  $\text{slab}(e_i(\sigma_j))^+$ , then we add a detour to  $\text{anchor}(\sigma_j)$ . (Note that  $\text{anchor}(\sigma_j)$  must exist in this case.) More precisely, if  $\sigma^*$  denotes the neighborhood that determines  $\text{anchor}(\sigma_j)$ , then we add a detour from the entry of  $\sigma^*$  to  $\text{anchor}(\sigma_j)$  and back. (That  $\sigma^*$  need not be in  $\mathcal{S}_i$  does not matter.)

Let  $e_k(\sigma_j)$  be the side containing  $\text{anchor}(\sigma_j)$ . Then, using Lemma 9 and the definition of the anchor, we have

$$\begin{aligned} |\text{length of detour}| &\leq 2 \cdot (2 + 10\alpha/\pi) \cdot \text{dist}(e_k(\sigma_j), e_k(\sigma^*)) \\ &\leq 2 \cdot (2 + 10\alpha/\pi) \cdot \text{dist}(e_i(\sigma_j), e_i(\sigma_{j+1})) \\ &\leq 2 \cdot (2 + 10\alpha/\pi) \cdot |\mathcal{T}_{opt}(q_j, q_{j+1})|. \end{aligned}$$

The total length of the detours we add in Case (i) is bounded by

$$\sum_{\sigma_j \in \mathcal{S}_i} 2 \cdot \max\left\{\frac{1}{\cos(\gamma/2)}, \left(2 + \frac{10\alpha}{\pi}\right)\right\} \cdot |\mathcal{T}_{opt}(q_j, q_{j+1})| \leq \max\left\{\frac{2}{\cos(\gamma/2)}, \left(4 + \frac{20\alpha}{\pi}\right)\right\} \cdot |\mathcal{T}_{opt}|.$$

**Case (ii):**  $|\mathcal{S}_i| = 1$ .

Let  $\sigma$  be the neighborhood in  $\mathcal{S}_i$  and let  $q$  be its entry. There are two subcases.

- If  $\mathcal{T}_{opt}$  does not stay inside  $\text{slab}(e_i(\sigma))^+$ , then we can follow the proof of the first subcase of Case (i) to show that we can add a detour for  $\sigma$  of length at most  $\max\{1/\cos(\gamma/2), (1 + 4\alpha/\pi)\} \cdot |\mathcal{T}_{opt}|$ .
- Otherwise  $\mathcal{T}_{opt}$  lies completely inside  $\text{slab}(e_i(\sigma))^+$ , so  $e_i(\sigma)$  is extreme. In this case we add the shortest detour from  $q$  to a guard on  $e_i(\sigma)$  and back.

To bound the length of the detour, first suppose that there is no stabbing line for  $\mathcal{S}$  parallel to and inside  $\text{slab}(e_i(\sigma))$ . In this case we added a single guard on  $e_i(\sigma)$ , which lies on a line  $\ell$  that has neighborhoods on both sides of it. But then  $\mathcal{T}_{opt}$  must cross  $\ell$ , so we can use Lemma 8 to bound the length of the detour by  $(1/\cos(\gamma/2)) \cdot |\mathcal{T}_{opt}|$ .

Next suppose that there is a collection of lines inside and parallel to  $\text{slab}(e_i(\sigma))$  that stab all neighborhoods. Let  $\ell$  and  $\ell'$  be the two lines bounding the subslab of all such stabbing lines. The guards we added in this case were the points  $\ell \cap e_i(\sigma)$  and  $\ell' \cap e_i(\sigma)$ . If  $\mathcal{T}_{opt}$  does not lie entirely inside the subslab, then we can use Lemma 8 again to bound the length of the detour by  $(1/\cos(\gamma/2)) \cdot |\mathcal{T}_{opt}|$ . If, on the other hand,  $\mathcal{T}_{opt}$  stays inside the subslab, then we can bound the length of the detour as follows. Let  $\sigma'$  be the neighborhood that lies directly below  $\sigma$  inside the subslab. Clearly  $\mathcal{T}_{opt}$  has to cross  $\sigma'$  to reach a point  $p$  below  $\sigma'$  (assuming there are at least three neighborhoods). Lemma 7 tells us that  $\min\{\text{dist}(q, \ell), \text{dist}(q, \ell')\} \leq (1 + 4\alpha/\pi) \cdot |pq|$ . If we combine this with Lemma 8 and use the fact that  $\mathcal{T}_{opt} \geq 2|pq|$ , then we see that we can bound the length of the detour by  $((1 + 4\alpha/\pi)/\cos(\gamma/2)) \cdot |\mathcal{T}_{opt}|$ .

We conclude that in Case (ii) we can bound the length of the detour by  $(\frac{1+4\alpha/\pi}{\cos(\gamma/2)}) \cdot |\mathcal{T}_{opt}|$ .

The condition  $\tan(\gamma/2) \leq \pi/(21\alpha)$  needed for Lemma 9 implies that  $0.98 < \cos(\gamma/2)$ , so the length of the detour in both cases can be bounded by  $(4 + 20\alpha/\pi)|\mathcal{T}_{opt}|$ . Since there are at most  $t$  non-empty subsets  $\mathcal{S}_i$ , we obtain that  $c = (4 + 20\alpha/\pi)t$ . Notice that  $t = 2\pi/\gamma$  depends on  $\alpha$  because of the condition on  $\gamma$  above.

Our algorithm computes an approximation of the optimal tour visiting each set  $\mathcal{R}_\sigma$ , using Slavik's algorithm. The latter algorithm has an approximation factor  $3k/2$ , where  $k$  is the maximum cardinality of any of the sets. In our case, the maximum size of  $\mathcal{R}_\sigma$  is  $3t + 1$  ( $t$  corners,  $2t$  guards and 1 anchor), where  $t = t(\alpha)$  is the number of sides of a neighborhood in  $\mathcal{S}$ . Combining this with the theorem above, we get the following result.

**Corollary 11** *Let  $\mathcal{S}$  be a collection of  $n$  disjoint convex  $\alpha$ -fat neighborhoods of constant complexity in the plane. We can compute in polynomial time a tour visiting each neighborhood in  $\sigma$ , whose length is at most  $3(3t + 1)/2 \cdot (1 + c)$  times the length of an optimal tour, where  $t = \pi/\arctan(\pi/21\alpha)$  and  $c = (4 + 20\alpha/\pi)t$ .*

**Remark 12** *A very rough estimate of the approximation factor would be  $12000\alpha^3$ . A more careful analysis would give a better value, although still large. For example, in the case when the input is a set of disjoint axis-aligned squares the approximation factor can be improved to roughly 93. There are two main reasons for this: (1) if the slabs are chosen to be axis-aligned then a square can have at most one extreme edge, and (2) there are only four squares with extreme edges and each of these only need two corner points. Hence, the number of representative points is at most five per square. Summing up the lengths of the detours we get  $5 \cdot \frac{3}{2}(4 \cdot 2\sqrt{2} + 1) < 93$ .*

*In our analysis we assume that a neighborhood can have  $2t$  guards which is a very conservative bound. If a neighborhood has many extreme edges then most guards will coincide with the corners, hence we expect the number of guards to be much smaller than the number of corners.*

*To improve the approximation factor one might be tempted to increase the value of  $t$ , i.e. the number of representative points, but due to the limitations of our proof technique this will only increase the value of the approximation factor.*

## 4 Conclusions and open problems

In this paper we have made two steps towards a better understanding of the approximability of TSPN. We have shown that in the most general setting—intersecting and non-convex neighborhoods—TSPN is APX-hard, and cannot be approximated within a factor  $391/390$ , unless  $P=NP$ . We also presented a constant-factor approximation algorithm for disjoint neighborhoods that are convex and fat. This is the first constant-factor approximation algorithm that does not require the neighborhoods to have roughly the same size.

Despite this progress, the problem is far from resolved. Probably the most interesting case to focus on is that of disjoint convex neighborhoods. Here the best known approximation factor is  $O(\log k)$  [9,12], where  $k$  is the number of neighborhoods. In fact, even for horizontal line segments this is the best known bound. We have tried to apply our techniques to obtain a constant-factor approximation for this case, but without success. There might even be a polynomial-time approximation scheme for disjoint convex neighborhoods; our lower bound result uses intersecting non-convex neighborhoods, and it seems hard to get around this.



## Acknowledgement

We thank V. Kann for suggesting to use the Vertex-Cover problem in the reduction proving the APX-hardness of TSPN. Many thanks also go to the anonymous referees who gave many helpful comments and pointed out some errors in a preliminary version.

## References

- [1] P. Agarwal, M. Katz and M. Sharir. Computing depth orders for fat objects and related problems. *Computational Geometry – Theory & Applications*, 5:187–206, 1995.
- [2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra and C. Uhrig. Approximate motion planning and the complexity of the boundary of the union of simple geometric figures. *Algorithmica*, 8(5 & 6):391–406, 1992.
- [3] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55:197–218, 1994.
- [4] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
- [5] P. Berman and M. Karpinski. On some tighter inapproximability results. In *Proc. 26th International Colloquium on Automata, Languages and Programming*, volume 1644 of *Lecture Notes in Computer Science*, pages 200–209, Springer-Verlag, 1999.
- [6] L. P. Chew, H. David, M. J. Katz and K. Kedem. Walking around fat obstacles. *Information Processing Letters*, 83(3):135–140, 2002.
- [7] A. Dumitrescu and J. S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. In *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 38–46, 2001.
- [8] M. R. Garey, R. L. Graham and D. S. Johnson. Some NP-complete geometric problems. In *Proc. 8th Annual ACM Symposium on Theory of Computing*, pages 10–22, 1976.
- [9] J. Gudmundsson and C. Levcopoulos. A fast approximation algorithm for TSP with neighborhoods. *Nordic Journal of Computing*, 6(4):469–488, 1999.
- [10] J. Hershberger and Subhash Suri. A pedestrian approach to ray shooting: Shoot a ray, take a walk. *Journal of Algorithms*, 18(3):403–431, 1995.
- [11] M. van Kreveld. On fat partitioning, fat covering, and the union size of polygons. In *Proc. 3rd Workshop on Algorithms and Data Structures*, volume 709 of *Lecture Notes in Computer Science*, pages 452–463. Springer-Verlag, 1993.

- [12] C. Mata and J. S. B. Mitchell. Approximation algorithms for geometric tour and network design problems. In *Proc. 11th Annual ACM Symposium on Computational Geometry*, pages 360–369, 1995.
- [13] J. Matoušek, N. Miller, J. Pach, M. Sharir, S. Sifrony and E. Welzl. Fat triangles determine linearly many holes. In *Proc. 32nd Annual IEEE Symposium on Foundations of Computer Science*, pages 49–58, 1991.
- [14] J. S. B. Mitchell. Geometric shortest paths and network optimization. In J.-R. Sack and J. Urrutia (eds.), *Handbook of Computational Geometry*, Elsevier Science.
- [15] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal of Computing*, 28(4):1298–1309, 1999.
- [16] C. H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoretical Computer Science*, 4(3):237–244, 1977.
- [17] G. Reich and P. Widmayer. Beyond Steiner’s problem: A VLSI oriented generalization. In *Proc. 15th International Workshop Graph-Theoretic Concepts in Computer Science*, volume 411 of Lecture Notes in Computer Science, pages 196–210, Springer-Verlag, 1989.
- [18] B. Shaleooi. Algoritmer för plåtskärning (*Eng. transl.* Algorithms for cutting sheets of metal). LUNDFD6/NFCS-5189/1–44/2001, Master thesis, Department of Computer Science, Lund University, 2001.
- [19] P. Slavík. The Errand Scheduling Problem. Technical report 97-02, Department of Computer Science and Engineering, SUNY Buffalo, 1997.
- [20] A. F. van der Stappen. Motion Planning amidst Fat Obstacles, Ph.D. Dissertation, Department of Computer Science, Utrecht University, Utrecht, Netherlands, 1994.