TSP with Neighborhoods of Varying Size

Mark de Berg

Joachim Gudmundsson

Matthew J. Katz

Christos Levcopoulos

Mark H. Overmars

A. Frank van der Stappen

TSP with Neighborhoods of Varying Size

Mark de Berg[*]   Joachim Gudmundsson[*]   Matthew J. Katz[†]

Christos Levcopoulos[‡]   Mark H. Overmars[*]   A. Frank van der Stappen[*]

3 September 2002

Abstract

In TSP with neighborhoods (TSPN) we are given a collection $\mathcal{S}$ of objects 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 objects are of approximately the same size. In this paper we present the first polynomial time constant-factor approximation algorithm for disjoint convex fat objects of arbitrary size. We also show that TSPN is APX-hard and cannot be approximated within a factor of 391/390 in polynomial-time, unless P=NP.

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 set of objects in the plane, called *neighborhoods*, find a shortest tour that visits all neighborhoods—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].

One can think of TSPN as a generalization of the Errand Scheduling Problem [20], sometimes also called the One-of-a-Set TSP [14]. In the One-of-a-Set TSP one is given a collection of sets of points and the aim is to visit at least one point from each set. This problem has applications in communication-network design [10], VLSI routing [17], and in the manufacturing industry [19]. The difference between TSPN and One-of-a-Set TSP is that in the
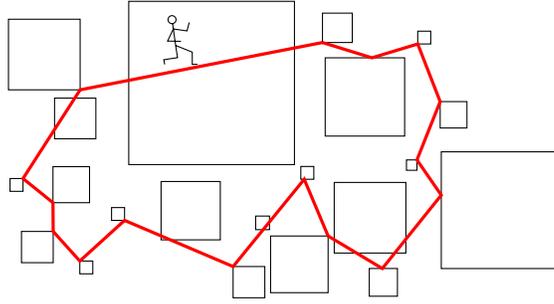
1

Figure 1: An example of TSPN when the input objects are axis-parallel squares.

latter problem each set to be visited has a finite number of points, unlike the neighborhoods in TSPN.

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 object and, more generally, for objects 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.[1] 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-approximation algorithm for TSPN. 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 objects 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 objects, namely for input sets $\mathcal{S}$ consisting of n disjoint convex fat objects.

---

[1]This was recently improved by Schwartz and Safra [18] who showed that TSPN cannot be approximated within a factor of $(2 - \varepsilon)$, unless $NP \subseteq TIME(n^{O(\log \log n)})$.

2

Before we proceed we introduce some notation. We use $|\tau|$ to denote the length of a curve τ. 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 that distance between objects σ and σ', that is, $\text{dist}(\sigma, \sigma') := \min_{p \in \sigma, q \in \sigma'} |pq|$. Finally, for a set $\mathcal{S}$ of objects we use $\mathcal{T}_{opt}(\mathcal{S})$ to denote an optimal (that is, shortest) tour visiting each object in $\mathcal{S}$.

2 TSPN is APX-hard

In this section we show that TSPN is APX-hard. 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:

Lemma 2.1 *[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.*

Let $\mathcal{G} = (V, A)$ be an undirected graph with n vertices of maximum degree at most four. We can restrict ourselves to consider 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 $n + |A|$ objects 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 objects: 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 objects 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 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 objects 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.

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.

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

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

Lemma 2.2 *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 objects in $\mathcal{S}_H \cup S$. The distance between any two objects 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. $\square$

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 2.3 *Let $\mathcal{T}$ be a tour visiting each object 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 2.2 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 2.2 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. $\square$

The main result of this section now follows easily.

4

Theorem 2.4 *TSP with neighborhoods is APX-hard, and cannot be approximated within a factor of 391/390 in polynomial time unless P=NP.*

Proof: Lemma 2.3 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 2.1, however, states that this is impossible in polynomial time unless P=NP. □

3   A constant-approximation algorithm for fat objects

We now turn our attention to computing an approximation of the optimal tour through a given set of α-fat disjoint bounded convex objects. The definition of fatness we use was introduced by Van der Stappen [21]. For convex objects it is basically equivalent to other definitions [1, 2, 11, 13].

Definition 3.1 *[21] Let α be a parameter with $\alpha \geq 1$. An object σ is α-fat if for any disk D whose center lies in σ and whose boundary intersects σ, we have $\mathrm{area}(D) \leq \alpha \cdot \mathrm{area}(\sigma \cap D)$.*

In our setting of bounded planar objects it turns out that $\alpha \geq 4$ [21].

We shall need the following property of convex fat objects, which was proved by Chew *et al.* [6].

Lemma 3.2 *[6] Let σ be a convex α-fat object, 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 3.3 *Let ℓ and ℓ' be two parallel lines intersecting a convex α-fat object σ. Let p and q be two points between ℓ and ℓ' and on opposite sides of σ. Then*

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

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

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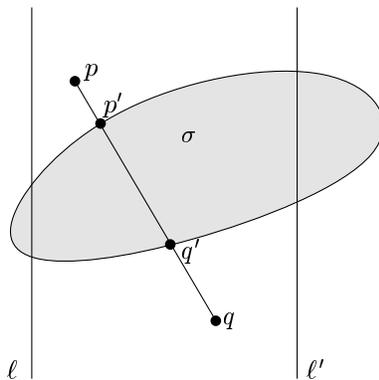&\leq (1 + 4\alpha/\pi) \cdot |pq|.
\end{aligned}
$$

□

Figure 3: Illustration for Lemma 3.3.

3.1 The algorithm

Let $\mathcal{S}$ be a set of n disjoint bounded convex α-fat objects, for some constant α (≥ 4). The overall approach we take is quite simple: for each object $\sigma \in \mathcal{S}$ 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 [20]. 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 $\mathcal{R}_s$ 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 object σ.

The representative points.

Fix an object $\sigma \in \mathcal{S}$. The set $\mathcal{R}_\sigma$ of representative points of σ 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.

The corners. Let γ 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 γ 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 (and in the sequel) 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 objects have constant complexity—in particular, each object 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 σ. The portions of $\partial\sigma$ connecting consecutive corners are called *sides* of σ. 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 object the corner or side belongs.

Before we proceed, we need to introduce one more concept. Consider a side e_i of σ, connecting the corners p_i and p_{i+1}. For a point p and an angle ϕ, let $\ell(p, \phi)$ be the line

through p that makes an angle ϕ with the positive x-axis. Define $\phi_i := (i + 1/2)\gamma - \pi/2$. The *slab of* e_i, denoted slab(e_i), is the slab 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 σ, is called the *semi-slab of* e_i; it is denoted slab$(e_i)^+$. Fig. 4 illustrates these definitions.
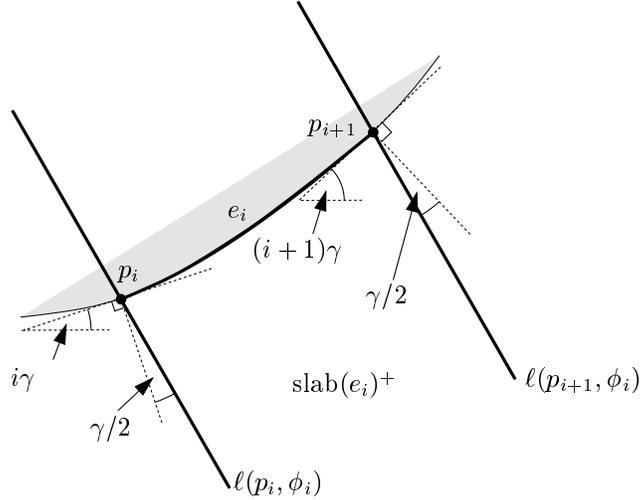

Figure 4: Corners, sides and (semi-)slabs.

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 3.4 *Let σ be a convex object, and e_i be one of its sides. Let ℓ be any line parallel to and inside* 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 ℓ 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. $\qquad\square$

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

Lemma 3.5 *Let σ and σ' be two convex α-fat objects such that $e_i(\sigma')$ lies at least partially inside* slab$(e_i(\sigma))^+$. *Let p be any point on $e_i(\sigma)$ such that the line through p parallel to* slab$(e_i(\sigma))$ *intersects σ', and let q be any point in $\sigma' \cap$ slab$(e_i(\sigma))^+$. If γ 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 $\sigma' \cap$ slab$(e_i(\sigma))^+$ with edges parallel and orthogonal to slab$(e_i(\sigma))$. Let w_B be the width and h_B be the height of B. See Figure 5.

The object σ' intersects all four edges of B. Consider the vertical line that cuts B into two equal halves; this line clearly intersects $\sigma' \cap$ slab$(e_i(\sigma))^+$. Let $m \in \sigma' \cap$ slab$(e_i(\sigma))^+$ be a
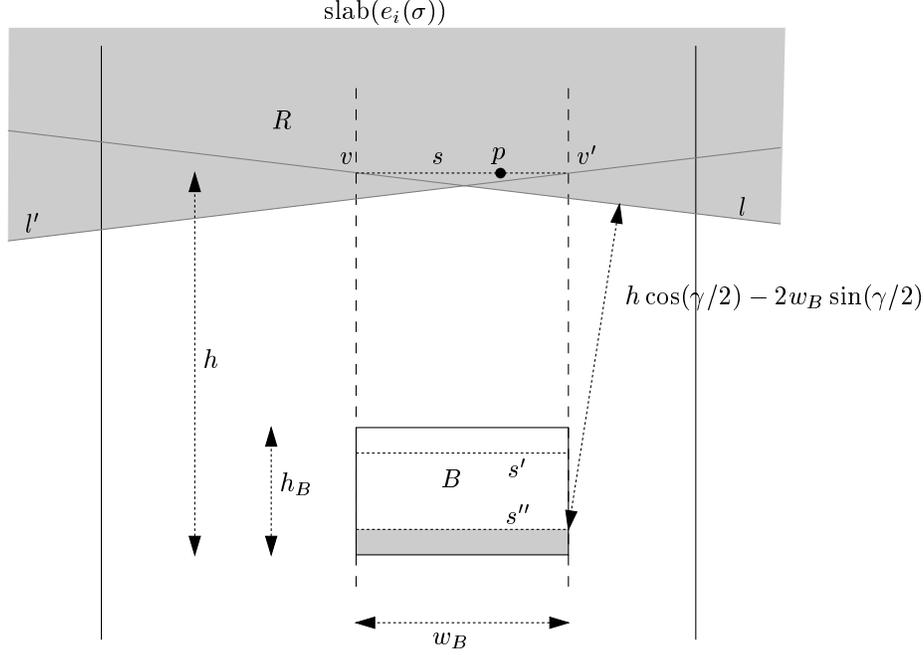
7

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

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 σ' and thus it follows that $\mathrm{area}(D \cap \sigma') \geq \mathrm{area}(D)/\alpha \geq \pi w_B^2/4\alpha$. Since $D \cap \sigma'$ lies completely inside B we also have $\mathrm{area}(B) \geq \pi w_B^2/4\alpha$. As a result we obtain

$$h_B \geq (\pi/4\alpha)w_B. \tag{1}$$

A point p on $e_i(\sigma)$ such that the line through p parallel to $\mathrm{slab}(e_i(\sigma))^+$ intersects σ' lies inside the imaginary subslab of $\mathrm{slab}(e_i(\sigma))^+$ 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 $q \in \sigma' \cap \mathrm{slab}(e_i(\sigma))^+$ is bounded by the distance from p to the farthest corner of the bounding box B of $\sigma' \cap \mathrm{slab}(e_i(\sigma))^+$, which in turn is bounded by $w_B + h$, so

$$|pq| \leq h + w_B. \tag{2}$$

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 $\sigma' \cap \mathrm{slab}(e_i(\sigma))^+$ 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

8

would require a too steep descent of $e_i(\sigma)$. As a result, the point p cannot lie below s' and thus

$$h \geq h_B - w_B \tan(\gamma/2). \tag{3}$$

By definition, the side $e_i(\sigma')$ touches the bottom edge of B. The bounds on its tangents dictate that $e_i(\sigma') \cap \mathrm{slab}(e_i(\sigma))^+$ 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 $\mathrm{dist}(e_i(\sigma), e_i(\sigma'))$ 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

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

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

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

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

The anchor. Consider a side $e_i(\sigma)$. We say that $e_i(\sigma)$ *dominates* a side $e_{i'}(\sigma')$ of another object σ' if $i = i'$ and $\sigma' \subset \mathrm{slab}(e_i(\sigma))^+$. We define an anchor for σ if σ has a side that dominates a side of another object, as follows. Of all the sides of other objects that are dominated by any of the sides of σ, 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 ℓ is parallel to $\mathrm{slab}(e_j(\sigma))$. Then $\mathrm{anchor}(\sigma)$, the anchor of σ, is the point $\ell \cap e_j(\sigma)$—see Figure 6.

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

Suppose there are lines parallel to and inside $\mathrm{slab}(e_i)$ that stab all the objects in $\mathcal{S}$. The collection of all such stabbing lines forms a subslab of $\mathrm{slab}(e_i)$. We add two guards on e_i to $\mathcal{R}_\sigma$, namely the intersection points of the lines bounding this subslab and e_i—see Figure 7(a).

If there are no stabbing lines, then there must be a line ℓ parallel to and inside $\mathrm{slab}(e_i)$ such that both half-planes defined by ℓ fully contain at least one object of $\mathcal{S}$. (Thus the line ℓ separates at least two objects.) In this case we add one guard to $\mathcal{R}_\sigma$, namely the intersection point of ℓ and e_i—see Figure 7(b).
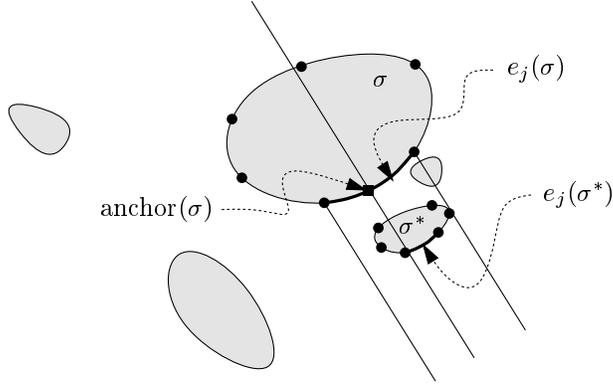
Figure 6: The anchor of an object σ.

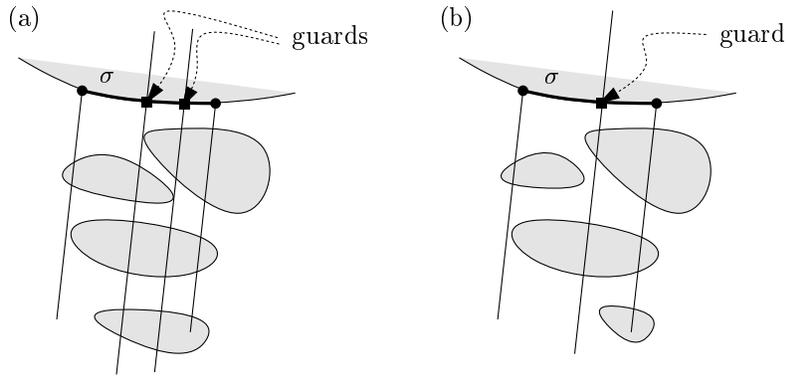
Figure 7: The guards of an object σ.

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 object σ.

Theorem 3.6 *Let $\mathcal{S}$ be a collection of n disjoint convex α-fat objects 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 object $\sigma \in \mathcal{S}$. 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 α).

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 objects. Imagine traversing the tour (in some arbitrary direction), starting at p_{start}. We define the *entry* of an object $\sigma \in \mathcal{S}$ to be

10

the point on σ where the tour first meets σ during the traversal. If the entry of σ 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 σ. The objects for which this is not the case have their entry in the interior of one of their sides. We partition the collection of these objects into t subsets, $\mathcal{S}_0, \ldots, \mathcal{S}_{t-1}$, such that $\mathcal{S}_i$ contains the objects σ 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 objects in $\mathcal{S}_i$ as $\sigma_1, \sigma_2, \ldots$ in the order as they are encountered during the traversal of $\mathcal{T}_{opt}$. Consider an object $\sigma_j \in \mathcal{S}_i$, and let $p_i := p_i(\sigma_j)$ and $p_{i+1} := p_{i+1}(\sigma_j)$. Denote the entry of σ_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 3.4 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 σ_j as seen from q_j. In this case we can use Lemma 3.2 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 σ_{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 σ_{j+1}. We add the detour $q_{j+1} p q_{j+1}$, and use Lemma 3.5 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 σ_{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 σ^* denotes the object that determines $\text{anchor}(\sigma_j)$, then we add a detour from the entry of σ^* to $\text{anchor}(\sigma_j)$ and back. (That σ^* 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 3.5 and the definition of the anchor, we have

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

11

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

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

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

Let σ be the object 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 σ 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 ℓ that has objects on both sides of it. But then $\mathcal{T}_{opt}$ must cross ℓ, so we can use Lemma 3.4 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 objects. Let ℓ and ℓ' 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 3.4 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 σ' be the object that lies directly below σ inside the subslab. Clearly $\mathcal{T}_{opt}$ has to cross σ' to reach a point p below σ' (assuming there are at least three objects). Lemma 3.3 tells us that $\min\{\text{dist}(q, \ell), \text{dist}(q, \ell')\} \le (1 + 4\alpha/\pi) \cdot |pq|$. If we combine this with Lemma 3.4 and use the fact that $\mathcal{T}_{opt} \ge 2|pq|$, 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 $\tan(\gamma/2) \le \pi/(21\alpha)$ needed for Lemma 3.5 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 α because of the condition on γ above.

$\square$

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 an object in $\mathcal{S}$. Combining this with the theorem above, we get the following result.

Corollary 3.7 *Let $\mathcal{S}$ be a collection of n disjoint convex α-fat objects in the plane. We can compute in polynomial time a tour visiting each object in σ, whose length is at most $3(3t + 1)/2 \cdot (1 + c)$ times the length of an optimal tour, where $t = \pi/\arctan(\pi/21\alpha)$ is the number of sides of an object in $\mathcal{S}$ and $c = (4 + 20\alpha/\pi)t$.*

Remark: 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 an object can have $2t$ guards which is a very conservative bound. If an object 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 objects. 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.

Recently Schwartz and Safra [18] showed that TSPN cannot be efficiently approximated to within a factor of $(2 - \varepsilon)$, unless $NP \subseteq TIME(n^{O(\log \log n)})$, by using a similar construction as in Section 2. In the same paper the authors also show that TSPN cannot be approximated within a factor of $\Omega(\log \log^{1/6} n)$ in 3D and $\Omega(\log \log^{\frac{d-2}{3(d-1)}} n)$ in d dimensions in polynomial time.

Acknowledgement

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. Technical Report 98-029, ECCC, 1998.

[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 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.

14

[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] O. Schwartz and S. Safra. On the complexity of approximating TSP with neighborhoods and related problems. Manuscript, 2002.

[19] 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.

[20] P. Slavík. The Errand Scheduling Problem. Technical report 97-02, Department of Computer Science and Engineering, SUNY Buffalo, 1997.

[21] A. F. van der Stappen. Motion Planning amidst Fat Obstacles, Ph.D. Dissertation, Department of Computer Science, Utrecht University, Utrecht, Netherlands, 1994.