Approximation algorithms for TSP with neighborhoods in the plane

Adrian Dumitrescu* Joseph S. B. Mitchell†

Abstract

In the Euclidean TSP with neighborhoods (TSPN), we are given a collection of n regions (*neighborhoods*) and we seek a shortest tour that visits each region. As a generalization of the classical Euclidean TSP, TSPN is also NP-hard. In this paper, we present new approximation results for the TSPN, including (1) a constant-factor approximation algorithm for the case of arbitrary connected neighborhoods having comparable diameters; and (2) a PTAS for the important special case of disjoint unit disk neighborhoods (or nearly disjoint, nearly-unit disks). Our methods also yield improved approximation ratios for various special classes of neighborhoods, which have previously been studied. Further, we give a linear-time $O(1)$-approximation algorithm for the case of neighborhoods that are (infinite) straight lines.

1 Introduction

A salesman wants to meet a set of potential buyers. Each buyer specifies a (connected) region in the plane, his *neighborhood*, within which he is willing to meet the salesman. For example, the neighborhoods may be disks centered at buyers' locations, and the radius of each disk specifies the maximum distance a buyer is willing to travel to the meeting place. The salesman wants to find a tour of shortest length that visits all of buyers' neighborhoods and finally returns to its initial departure point. A variant of the problem, which we will address in this paper, is that in which no departure point is specified, and only a tour of the neighborhoods is to be found. This problem, which is known as the *TSP with neighborhoods* (TSPN), is a generalization of the classic Euclidean Traveling Salesman Problem (TSP), in which the regions (neighborhoods) are single points, and consequently is NP-hard [6, 16].

Related Work. The TSP has a long and rich history of research in combinatorial optimization. It has been studied extensively in many forms, including geometric instances; see [3, 9, 10, 14, 18]. The problem is known to be NP-hard, even for points in the Euclidean plane [6, 16]. It has recently been shown that the geometric instances of TSP (e.g., Euclidean TSP) have a polynomial-time approximation scheme, as developed by Arora [2] and Mitchell [13], and later improved by Rao and Smith [17].

Arkin and Hassin [1] were the first to study approximation algorithms for the geometric TSPN. They gave $O(1)$-approximation algorithms for several special cases, including parallel segments of equal length, translates of a convex region, translates of a connected region, and more generally, for regions which have diameter segments that are *parallel* to a common direction, and the ratio between the longest and the shortest diameter is bounded by a constant.

For the general case of connected polygonal regions, Mata and Mitchell [11] obtained an $O(\log n)$-approximation algorithm, based on "guillotine rectangular subdivisions", with time bound $O(m^5)$, where m is the total complexity of the n regions. Gudmundson and Levcopoulos [7] have recently obtained a faster method, which, for any fixed $\epsilon > 0$, is guaranteed to perform at least one of the following tasks (although one does not know in advance which one will be accomplished): (1) it outputs a tour of length at most $O(\log n)$ times optimum in time $O(n \log n + m)$; (2) it outputs a tour of length at most $(1+\epsilon)$ times optimum in time $O(m^3)$. So far, no polynomial-time approximation algorithm is known for general connected regions.

Recently it was shown that TSPN is APX-hard and cannot be approximated within a factor of 1.000374 unless P=NP [8].

Summary of Our Results. In this paper, we obtain several approximation results on the geometric TSPN, including:

(1) We extend the approaches initiated in [1] and obtain the first $O(1)$-approximation algorithm for the TSPN having connected regions of the same (or similar) diameter. This solves among others, the open problem posed in [1], to provide a constant-

*`dumitres@ams.sunysb.edu`. This work was done while the author was a visiting faculty member at the University at Stony Brook.

†`jsbm@ams.sunysb.edu`. Dept. of Applied Mathematics and Statistics, University at Stony Brook, Stony Brook, NY 11794-3600. Partially supported by grants from HRL Laboratories, NASA Ames, the National Science Foundation (CCR-9732220), Northrop-Grumman Corporation, Sandia National Labs, and Sun Microsystems.

factor approximation algorithm for TSPN on segments of the same length and *arbitrary* orientation.

(2) We give a PTAS for the case of disjoint unit disks (or nearly disjoint disks of nearly the same size). The algorithm is based on applying the m-guillotine method with a new area-based charging scheme. The fact that there is a PTAS for the case in which the neighborhoods are "nice" and not severely overlapping should be contrasted with the fact that the TSPN on arbitrary regions is APX-hard. (The construction in the proof of [8] utilizes very "skinny" neighborhoods, which intersect each other extensively.)

(3) We also give modest improvements on earlier approximation bounds in [1] for the cases of parallel segments of equal length, translates of a convex region, and translates of a connected region.

(4) We present simple algorithms which achieve a constant-factor guarantee for the case of equal disks and for the case of infinite straight lines.

2 Equal Disks

We begin by giving some simple arguments and corresponding algorithms which achieve a constant ratio for TSPN on disks of the same size. (Our results carry over naturally to disks of nearly the same size, with corresponding changes in the approximation factor. We defer details to the full paper.) Without loss of generality, assume that all disks have unit radius.

First, we consider the case of disjoint unit disks. The algorithm is simple and natural: compute a $(1+\epsilon)$-approximate tour, T_C, of the center points (the *center tour*) of the collection of n disks. Clearly T_C is a valid region tour. Denote by T_R an optimal region (disk) tour. For any tour T, let $|T|$ be its length. We claim that

$$|T_C| \leq ((1+\frac{8}{\pi})|T_R|+8)(1+\epsilon).$$

Put $l = |T_R|$. Since T_R visits all disks, the area swept by a disk of radius 2, whose center moves along T_R, covers all unit disks. This area A_l is bounded as follows

$$\pi n \leq A_l \leq 4l + 4\pi,$$

from which we get $n \leq 4 + \frac{4l}{\pi}$. A center tour of length

$$\leq l + 2n \leq l + 2(4+\frac{4l}{\pi}) = (1+\frac{8}{\pi})l + 8$$

can be obtained by going along T_R and making a detour of length at most 2 to visit the center of each disk, when T_R visits the disk. Hence the length of the computed tour is bounded as claimed. For n large, the approximation ratio is $(1+\frac{8}{\pi})(1+\epsilon) \leq 3.55$. (With a somewhat more delicate analysis, we are able to improve this factor by using the packing density of disks, and using a more refined method of taking detours to center points.) We note that for the algorithm that computes an approximate TSP tour, T_C, on the center points, we cannot expect a ratio smaller than 2. To see this, consider a large square, and place almost touching unit disks along its perimeter, both on its inside and on its outside. All the disks touch the perimeter which is also an optimal disk tour except at the four corners of the square. The disk center tour has length roughly two times that of the perimeter.

Next we consider the case when the disks can overlap. First, we compute a maximal independent (pairwise-disjoint) set I of disks. Next, we compute C_I, a $(1+\epsilon)$-approximate tour of the center points of disks in I. Finally, we output the tour R obtained by going along C_I and along the boundaries of each of the disks in I, as illustrated in Figure 1. Clearly, R is a valid region tour.

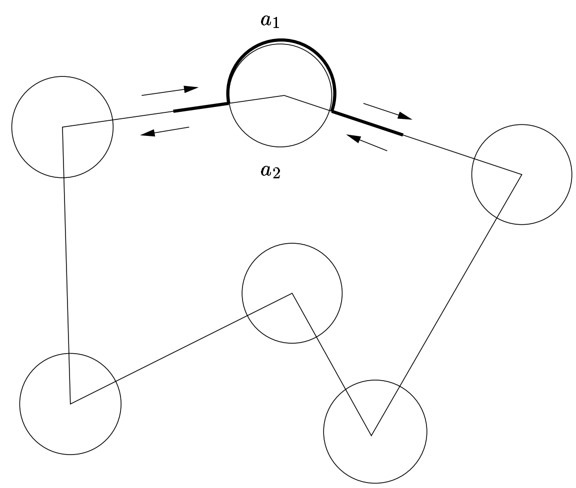

Figure 1: Disjoint unit disks.

Denote by OPT an optimal disk tour and by OPT_I an optimal disk tour of I. A constant approximation ratio can be derived from the following three inequalities

$$|R| \leq \pi|C_I| + 2\pi, \tag{1}$$

$$|C_I| \leq ((1+\frac{8}{\pi})|OPT_I|+8)(1+\epsilon), \tag{2}$$

$$|OPT_I| \leq |OPT|. \tag{3}$$

The third inequality follows from the fact that $I \subseteq \mathcal{R}$, and the second from the case of disjoint unit disks studied previously. To check the first inequality, decompose C_I into $|I|$ parts (assume $|I| \geq 2$), one for each disk in I, by cutting each segment between two consecutive centers in the middle. Let $x + y = d \geq 2$ be the length of one of these parts in the tour, corresponding to a disk D, where x, y are the lengths of the two segments of C_I adjacent to the center of D. Write a_1, a_2 for the arc lengths of D when its boundary is traversed by R (we have $|a_1| + |a_2| = 2\pi$). Comparing the lengths of this part in C_I and R, we get

$$\frac{(x - 1 + a_1 + y - 1) + (x - 1 + a_2 + y - 1)}{d} =$$

$$\frac{2(d + \pi - 2)}{d} \geq \pi,$$

the minimum being attained when $d = 2$. When $|I| = 1$, $|C_I| = 0, |R| = 2\pi$. Putting (1), (2), (3) together, we get

$$|R| \leq ((\pi + 8)|OPT| + 8\pi)(1 + \epsilon) + 2\pi.$$

For n large, the approximation ratio is $(\pi + 8)(1 + \epsilon) \leq 11.15$.

We note that the approximation ratio we have obtained with this approach for disjoint unit disks, 3.55, (resp. 11.15 for unit disks) is better (resp. weaker) than $\sqrt{3^2 + 7^2} \approx 7.62$, the approximation ratio that will be given at the end of Section 4 for translates of a convex region (which applies for unit disks).

3 A PTAS for Disjoint Equal Disks

Given the powerful methods that have been developed to obtain PTAS's for various geometric optimization problems, such as the Euclidean TSP, it is natural to suspect that these same techniques may apply to the TSPN. Indeed, one may expect that TSPN should have a PTAS based on applying existing methods. However, we know now, from the recent APX-hardness result of [8], that this cannot be. What goes wrong?

The basic issue we must address in order to apply these techniques is to be able to write a recursion to solve an appropriate "succinct" subproblem with dynamic programming. What is the subproblem "responsible" to solve? For problems involving *points*, the subproblem can be made responsible for constructing some kind of inexpensive network on the points *inside* the subproblem (rectangle), and to interconnect this network with the boundary in some nicely controlled way (e.g., with only a constant complexity of connection, in the case of m-guillotine methods). The problem with *regions* is that they can cross subproblem boundaries. Then, we do not know if the subproblem is responsible to visit the region, or if the region is visited *outside* the subproblem. We cannot afford to enumerate the subset of regions that cross the boundary for which the subproblem is responsible – there are too many such subsets, leading to too many subproblems. Thus, we need a new idea.

Our approach is to employ a new type of structural result, based on the general method of m-guillotine subdivisions. In particular, we show how to transform an optimal tour into one of a special class of tours that recursively has a special m-guillotine structure, permitting us to have a succinct (constant-size) specification of the subset of regions, crossing the subproblem boundary, for which the subproblem is "responsible" (must visit on its interior). In order to bound the increase in tour length in performing this transformation, we must "charge off" the added tour length to some small fraction of the length of the optimal tour (just as is done in proving the bounds for the m-guillotine PTAS method for TSP). In order to do this charging, we must assume some special structure to the class of neighborhoods in the TSPN, e.g., that the regions $\mathcal{D} = \{D_1, \ldots, D_n\}$ are pairwise disjoint equal-size (without loss of generality, unit) disks (or have a similar structure allowing us to relate tour length to area).

Here, we outline how the approach applies to disjoint unit disks; generalizations will appear in the full paper. We begin with some definitions, largely following the notation of [13]. Let G be an embedding of a planar graph, and let L denote the total Euclidean length of its edges, E. We can assume (without loss of generality) that G is restricted to the unit square, B (i.e., $E \subset int(B)$).

Consider an axis-aligned rectangle W (a *window*) with $W \subseteq B$. (W will correspond to a subproblem.) Let s_1, s_2, s_3, s_4 be the four segments that form the sides of W, and let $\mathcal{D}_i \subseteq \mathcal{D}$ be the set of disks intersecting s_i. Since the disks are disjoint, they can be ordered along side s_i; thus, we think of $\mathcal{D}_i$ as an ordered sequence, whose subsequences of consecutive disks (consecutive along edges in E) are called *disk-intervals*. Since our goal is to have a succinct representation of $\mathcal{D}_i$, for each side s_i of W, we are motivated to look at situations in which $\mathcal{D}_i$ can be represented as a small ($O(m)$-size) set of disk-intervals. Given a positive integer m, we say that E is *m-good with respect to W* if, for each $i = 1, 2, 3, 4$, the set $\mathcal{D}_i$ consists of at most $2m$ disk-intervals.

An axis-parallel line ℓ is an *m-good cut with respect to W* if it splits W into two subwindows, W_1 and W_2, such that E is m-good with respect to both W_1 and W_2.

We now say that G (or E) satisfies the *m-guillotine property with respect to W* if E is m-good with respect to W and either (1) W does not fully contain any disk; or (2) there exists an m-good cut, ℓ, that splits W into W_1 and W_2, and, recursively, E is m-good with respect to both W_1 and W_2.

Theorem 3.1. *Let G be an embedded planar graph, with edge set E, of total length L, and let $\mathcal{D}$ be a given set of pairwise-disjoint equal-radius disks (radius δ). Assume that $E, \mathcal{D} \subset B$. Then, for any positive integer m, there exists a planar graph G' that satisfies the m-guillotine property with respect to B and has an edge set $E' \supseteq E$ of length $L' \leq (1 + \frac{32(1+(1/\pi))}{m})L$.*

Proof. (sketch) We convert G into a new graph G' by adding to E a new set of horizontal/vertical edges whose total length is at most $\frac{32(1+(1/\pi))}{m}L$. The construction is recursive: at each stage, we show that there exists a cut, ℓ, with respect to the current window W (which initially is the box B), such that we can afford to add the following "m-bridging" segment to E: the segment joining the bottom endpoint of $\ell \cap D_T$ with the top endpoint of $\ell \cap D_B$, where D_T (resp., D_B) is the mth disk encountered along $\ell \cap W$ when walking down from the top of W (resp., up from the bottom of W).

In order to make the charging scheme work, we have to define a notion of "favorable cut", as was done in [13]. Here, a cut is *favorable* if the length of its "chargeable" portion within W is *at least as long as* the length of the m-bridging segment. (The chargeable portion consists of those points $p \in \ell \cap W$ for which both the leftward and rightward rays from p intersect at least m disks before leaving W.) The existence of such a favorable cut is guaranteed by the following key lemma, whose proof is similar to that of the key lemma in [13]:

Lemma 3.1. *For any G and any window W, there is a favorable cut.*

The rest of the charging scheme is based on charging off the lengths of the m-bridges to $(1/m)$th of the projected perimeter lengths of the disks (which is at most 8δ, for a δ-radius disk, since we project up, down, left, and right). We then use the lower bound, from the previous section, which shows that the optimal length of the TSPN tour is $L^* \geq (n\delta - 4)\pi/4$, from which we get that the total charge is at most $8\delta n/m \leq \frac{8}{m}(4 + \frac{4}{\pi})L^*$.

Corollary 3.1. *The TSPN on a set of disjoint equal-size disks has a PTAS.*

Proof. (sketch) Once we have the structure Theorem 3.1, we proceed to write a dynamic programming algorithm, running in $O(n^{O(m)})$ time, to compute a minimum-length planar graph having a prescribed set of properties: (1) it satisfies the m-guillotine property (necessary for the dynamic program to have the claimed efficiency); (2) it visits at least one point of each region D_i, at one of a small constant number $(O(m^2))$ of grid points in D_i; and (3) it is Eulerian (a condition that is necessary in order to extract a tour in the end). A subproblem is defined by a rectangle W, together with a constant amount of boundary information, specifying $O(m)$ disk-intervals, a grid point inside each disk that is an "endpoint" of each disk-interval, and a "connection pattern" that is required within W (e.g., by giving a partition of the $O(m)$ disk-intervals into one of the $2^{O(m)}$ possible subsets).

The result is that in polynomial time $(O(n^{O(m)}))$ one can compute a shortest possible Eulerian graph, from a special class of such graphs, and this graph spans the regions $\mathcal{D}$. Theorem 3.1 guarantees that the length of the resulting graph is very close (within factor $1 + O(1/m)$) to the length of an optimal solution to the TSPN. Thus, once we extract a tour from the Eulerian graph, we have the desired $(1+\epsilon)$-approximation solution, where $\epsilon = O(1/m)$.

4 Connected Regions of the Same Diameter

The *diameter* of a region, δ, is the distance between two points in the region that are farthest apart. Without loss of generality, we assume that all regions have unit diameter, $\delta = 1$. The general method we use is to carefully select a representative point in each region, and then compute an almost optimal tour on these representative points. This approach was initiated in [1]. We also employ the following tool:

Lemma 4.1. (Combination Lemma) [1] *Given regions that can be partitioned into two types, and constants c_1, c_2 bounding the error ratios with which we can approximate optimal tours on regions of types 1 and 2, then we can approximate the optimal tour on all regions with an error ratio bounded by $c_1 + c_2 + 2$.*

As noted in [1], repeated application of this lemma may be useful: if there are k different region types, with individual approximation bounds of $c_1, \ldots, c_k$, the resulting ratio is bounded by $c_1 + \ldots + c_k + 2(k-1)$. Now we describe the algorithm.

Select (compute) arbitrarily in each region a unit diameter segment. Classify the regions into two types: (1) those for which the selected diameter is *almost horizontal*, by which we mean its slope is between $-45°$ and $45°$; (2) those for which the selected diameter is *almost vertical*, by which we mean all the others. (Computing a diameter segment of a region can be done efficiently, in time linear in the complexity (e.g., number of vertices) of the region.) Use algorithm **A** (below) for each of these two region types. We will prove that a constant ratio is achievable for each class. We then apply the Combination Lemma to obtain a constant-factor approximation for all regions.

Algorithm A.

1. Construct a greedy cover of the regions by a minimum number of vertical lines. (A set of lines is a *cover* of a set of regions if each region is intersected by at least one line from the covering set. We refer to this set of lines as *covering lines.*) This procedure works in a greedy fashion, namely the leftmost line is as far right as possible, so that it is a right tangent of some region. To obtain this cover, the intervals of projection on the x-axis of all regions are computed and a greedy cover of this set of intervals is found. After removing all segments covered by a previous line, another covering line is repeatedly added to the cover; at the same time, representative points of each region are arbitrarily selected on the corresponding covering lines.
2. Proceed according to the following three cases.

Case 1 The greedy cover contains one covering line. Compute a smallest perimeter isothetic rectangle Q (of width w and height h), that touches (intersects) all regions (Q is considered a two-dimensional region). Add twice the two vertical segments of height h which divide its width in three equal parts, to get a tour R. Output R.

Case 2 The greedy cover contains two covering lines. Move (if possible) the rightmost vertical covering line to the left as much as possible (while still covering all regions). Recompute the representative points obtained in this way. Set D to be the distance between the two covering lines (clearly $D > 0$).

Case 2.1 $D \geq 3$. Construct (compute) rectangle Q of width $w = D$, with its vertical sides along the two covering lines, and of minimal height h, which includes all representative points (on the two covering lines). Output the tour R that is the perimeter of Q.

Case 2.2 $D \leq 3$. Compute a smallest perimeter isothetic rectangle Q (with width w and height h), that touches (intersects) all regions. Add twice the seven vertical segments of height h which divide its width into eight equal parts, to get a tour R. Output R.

Case 3 The greedy cover contains at least three covering lines. Construct R, a $(1+\epsilon)$-approximation tour of the representative points as the output tour.

We prove that algorithm **A** gives a constant-factor approximation of the optimal tour for regions of type (1). (Regions of type (2) are readily handled by rotating them by 90° to obtain type (1) regions.) (In the full paper, we give also tradeoffs in the approximation bound when the diameters are not identical, but are "similar" in that the ratio of the largest to the smallest is bounded.)

THEOREM 4.1. *In polynomial time, one can compute an $O(1)$-approximation of an optimal tour on a set of n connected regions having the same diameter.*

Proof. Let OPT be an optimal region tour. We prove individually each of the cases we distinguished in the previous algorithm. We will use repeatedly the following simple fact (see [1]): For positive a, b, w, h the following inequality holds

$$aw + bh \leq \sqrt{a^2+b^2}\sqrt{w^2+h^2}. \tag{4}$$

Case 1. Write $diag(Q)$ for the diagonal of the rectangle Q. We first argue that R visits all regions. Since all regions are covered by a unique covering line, they lie in a vertical strip of width $\leq 2\delta = 2$. So $w \leq 2$. The horizontal projection of each region (on the x-axis) is at least $\frac{1}{\sqrt{2}}$. Hence, all regions are intersected either by the boundary (perimeter) of Q or by the two vertical segments inside Q (since $3\frac{1}{\sqrt{2}} > 2$). Thus, R is a valid region tour. $|OPT|$ and $|R|$ are bounded as follows.

$$|OPT| \geq 2diag(Q) = 2\sqrt{w^2+h^2},$$

$$|R| = 2w + 6h = 2(w+3h) \leq 2\sqrt{10}\sqrt{w^2+h^2} \leq \sqrt{10}|OPT|.$$

Case 2.1. $D \geq 3$. We distinguish two sub-cases.

Case 2.1.a $h \leq 2$. Recall that $w = D \geq 3$.

$$|OPT| \geq 2(w-2),$$

$$|R| = 2w + 2h \leq 2w + 4 \leq 10(w-2) \leq 5|OPT|.$$

Case 2.1.b $h \geq 2$.

$$|OPT| \geq 2\sqrt{(w-2)^2+(h-2)^2},$$

$$|R| = 2w + 2h \leq 10((w-2)+(h-2)) \leq 10\sqrt{2}\sqrt{(w-2)^2+(h-2)^2} \leq 5\sqrt{2}|OPT|.$$

Case 2.2. $D \leq 3$. Since the horizontal projection of each region is at least $\frac{1}{\sqrt{2}}$ and $8\frac{1}{\sqrt{2}} > 5$, R visits all regions.

$$|OPT| \geq 2diag(Q) = 2\sqrt{w^2+h^2}.$$

$$|R| = 2w + 16h \leq 2\sqrt{1^2+8^2}\sqrt{w^2+h^2} \leq 2(8.1)\sqrt{w^2+h^2} \leq 8.1|OPT|.$$

Case 3. Partition the optimal tour OPT into blocks OPT_i, with $i \geq 1$. OPT_1 starts at an arbitrary point of intersection of OPT with the leftmost covering line (notice that OPT does not cross to the left of this line), and ends at the last intersection of OPT with the second left covering line, before OPT intersects a different covering line. In general, the blocks OPT_i are determined by the last point of intersection of OPT with a covering line, before OPT crosses a different covering line. Consider the bounding box Q of OPT_i, the smallest perimeter aligned rectangle which includes OPT_i. Write w for its width and h for its height. There are two cases to consider.

Case 3.1 OPT_i intersects regions stabbed by two consecutive covering lines only, l_1, l_2 say, at distance w_1; this implies that OPT_i lies between these two covering lines, so $w = w_1$. Without loss of generality, OPT_i touches the lower side of Q (at some point C) before it touches the upper side of Q. Since the horizontal projection of each region is at least $\frac{1}{\sqrt{2}}$, $w_1 \geq \frac{1}{\sqrt{2}}$. Let $0 \leq a, b \leq h$ specify the starting and ending points A and B of OPT_i (see Figure 2). By considering the reflections of OPT_i with respect to the two horizontal sides of Q, we get

$$|OPT_i| \geq \sqrt{(h+a+b)^2 + {w_1}^2}. \qquad (5)$$

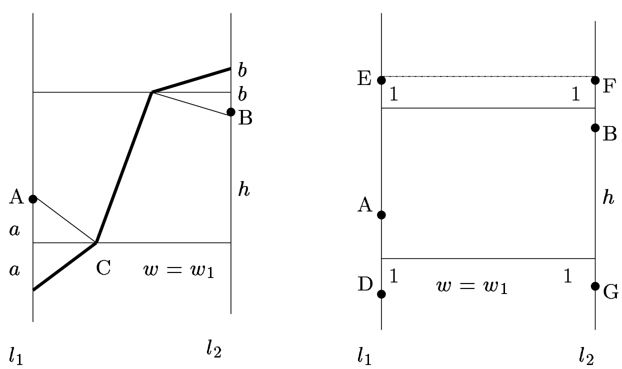

Figure 2: Case 3.1: 2 lines.

We show that there exists a partial tour R_i (a path) of the representative points of all the regions OPT_i visits, of length bounded by $c|OPT_i|$ for some positive constant c, where R_i starts at A and ends at B. Take $R_i = ADEFGB$, where the points D, E, F, G are on the lines l_1, l_2 at unit distance from the corners of Q (see Figure 2).

$$|R_i| \leq (a+1)+(1+h+1)+(w_1)+(1+h+1)+(1+h-b)$$

$$\leq 3(h+a+b) + (w_1+6) \leq 3(h+a+b) + (6\sqrt{2}+1)w_1$$

$$\leq \sqrt{3^2 + (6\sqrt{2}+1)^2}\sqrt{(h+a+b)^2 + {w_1}^2}$$

$$< 10|OPT_i|. \qquad (6)$$

Put R_i together to get a tour R' of length smaller than $10|OPT|$. Then the $(1+\epsilon)$ approximate tour of representatives has length at most $10|OPT|$.

Case 3.2 OPT_i intersects regions stabbed by three consecutive covering lines only, l_1, l_2, l_3 say at distances w_1, w_2; this implies that OPT_i lies between these three covering lines, but it does not touch the rightmost one l_3. We can assume that OPT_i starts at A and ends at B, and as in the previous case, that it touches the lower side of Q (at some point C) before it touches the upper side of Q. Let l' be the supporting right line of the right side of Q. We also have $w_1, w_2 \geq \frac{1}{\sqrt{2}}$ and $w' > \max(0, w_2 - 1)$. The case when OPT_i touches the upper side of Q before it touches l' is shown is Figure 3; the other case is similar and we get the same bound on OPT_i. Let a, b, A, B be as before. We distinguish two sub-cases:

Case 3.2.a $w_2 \leq 1$. The lower bound on OPT_i we have used earlier is still valid.

$$|OPT_i| \geq \sqrt{(h+a+b)^2 + {w_1}^2}.$$

We show a partial tour R_i of the representative points of all the regions OPT_i visits. Take $R_i = ADEFBGHIFB$ (see Figure 3).

$$|R_i| \leq (a+1) + (1+h+1) + w_1 + (1+h+1) +$$

$$w_2 + (1+h+1) + w_2 + (1+b)$$

$$\leq 3(h+a+b) + (w_1 + 2w_2 + 8)$$

$$\leq 3(h+a+b) + (10\sqrt{2}+1)w_1$$

$$\leq \sqrt{3^2 + (10\sqrt{2}+1)^2}\sqrt{(h+a+b)^2 + {w_1}^2}$$

$$< 15.5|OPT_i|. \qquad (7)$$

Then the $(1+\epsilon)$-approximate tour of representatives has length at most $15.5|OPT|$.

Case 3.2.b $w_2 \geq 1$. We use a different lower bound on OPT_i.

$$|OPT_i| \geq \sqrt{(h+a+b)^2 + (w_1 + 2w_2 - 2)^2}, \qquad (8)$$

which we get by considering the reflections of OPT_i with respect to the two horizontal sides of Q and with respect to l' (see Figure 3). Take $R_i = ADEFBGHIFB$ as in Case 3.2.a.

$$|R_i| \leq 3(h+a+b) + (w_1 + 2w_2 + 8)$$

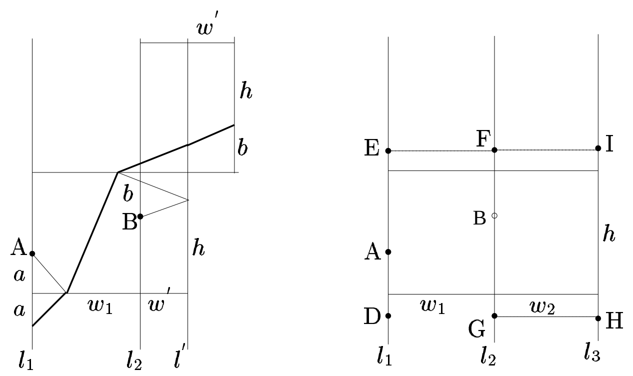

Figure 3: Case 3.2: 3 lines.

$$\leq 3(h+a+b) + (10\sqrt{2}+1)(w_1+2w_2-2)$$
$$\leq \sqrt{3^2+(10\sqrt{2}+1)^2}\sqrt{(h+a+b)^2+(w_1+2w_2-2)^2}$$
$$< 15.5|OPT_i|. \qquad (9)$$

Then the $(1+\epsilon)$-approximate tour of representatives has length at most $15.5|OPT|$. The overall approximation ratio of the algorithm, derived from the Combination Lemma is $15.5+15.5+2=33$.

Some Special Cases

We note that our calculations of the approximation ratio for connected regions of a same diameter, give improved bounds for three cases addressed in [1]:

1. parallel equal segments, from $3\sqrt{2}+1$, to $3\sqrt{2}$,
2. translates of a convex region, from $\sqrt{3^2+7^2}+1$, to $\sqrt{3^2+7^2}$,
3. translates of a connected region, from $\sqrt{3^2+11^2}+1$, to $\sqrt{3^2+11^2}$.

We exemplify here the case of parallel equal segments and omit details for the rest. The algorithm computes a greedy cover of the segments (assume of unit length) using vertical lines. then it proceeds according to the cardinality of the cover. Cases 1 and 2 are treated in [1], and the ratio is $\sqrt{2}$ (this is not the bottleneck case). In Case 1 (one covering line), an optimal tour is easy to obtain. In Case 2 (two covering lines), a smallest aligned rectangle which touches all segments is the output tour. In Case 3 (three or more covering lines) the algorithm computes an almost optimal tour of the representative points (as algorithm **A** does). Its analysis is divided into two sub-cases, as in the proof of Theorem 4.1. In the first sub-case (OPT_i intersects segments covered by two consecutive covering lines only), the lower bound in Equation (5) on $|OPT_i|$ is valid. The upper bound in Equation (6) on $|R_i|$ is adjusted by dropping the constant term equal to $+6$. Then

$$|R_i| \leq 3(h+a+b)+w_1$$
$$\leq \sqrt{3^2+1^2}|OPT_i| = \sqrt{10}|OPT_i|.$$

In the second sub-case (OPT_i intersects segments covered by three consecutive covering lines only), the lower bound in Equation (8) on $|OPT_i|$ is valid. The upper bound in Equation (9) on $|R_i|$ is adjusted by dropping the constant term equal to $+8$. We also have that $w_1, w_2 > 1$. Then

$$|R_i| < 3(h+a+b)+w_1+2w_2$$
$$\leq 3(h+a+b)+3(w_1+2w_2-2) \leq 3\sqrt{2}|OPT_i|.$$

The overall approximation ratio obtained for parallel equal segments is $3\sqrt{2}$.

5 Lines

Consider now the case in which the n regions are infinite straight lines in the plane. It is interesting that this case allows for an exact solution in polynomial time:

Proposition 5.1. *Given a set L of n infinite straight lines in the plane, a shortest tour that visits L can be computed in polynomial time.*

Proof. (sketch) We convert the problem to an instance of the *watchman route problem* in a simple polygon, which is known to have an $O(n^6)$ algorithm (see [19], as well as [5, 15]). Let B be a rectangle that contains all of the vertices of the arrangement of L. At each point p where a line $l_i \in L$ intersects the boundary of B, we extend a very narrow "spike" outward from p, along l_i, for some fixed distance. Let P be the simple polygon having $6n+4$ vertices that is the union of B and these $2n$ spikes (in fact, only one spike per line is enough). We claim that a tour T visits all the lines in L if and only if it sees all of the polygon P. Thus, we can solve the TSPN on L by solving the watchman route problem on the simple polygon P.

Given the high running time of the watchman route algorithms, it is of interest to consider more efficient algorithms that may approximate the optimal solution. To this end, we now present a *linear-time* constant-factor approximation algorithm.

Let $L = \{l_1, \ldots, l_n\}$ be the input set of n lines. A *minimum touching circle (disk)* is a circle of minimum radius which touches (intersects) all of the lines in L. The algorithm computes and outputs C_L, a minimum touching circle for L. We will show that this provides a tour of length at most $\frac{\pi}{2}|OPT|$, where, as usual, OPT

denotes an optimal tour. First we argue about the approximation ratio, and leave for later the presentation of the algorithm. For simplicity, we assume that no two lines are parallel, though this assumption can be removed.

OBSERVATION 5.1. *The optimal tour is a (possibly degenerate) convex polygon P.*

(It is easy to see that OPT consists of connected straight line segments, thus forming a polygon P; assuming P non-convex, its convex hull $CH(P)$ would still visit all the lines in L and have a shorter length, which is a contradiction.)

OBSERVATION 5.2. *C_L is determined by 3 lines in L, i.e. it is the inscribed circle in the triangle Δ formed by these 3 lines.*

We distinguish two cases:

Case 1. Δ is an acute triangle. It is well known that for an acute triangle, the minimum perimeter inscribed triangle (having a vertex on each side of the triangle) is its orthic triangle. So in this case the optimal tour OPT_Δ, which visits the 3 lines of Δ is its orthic triangle; we denote by y its perimeter. Clearly y is a lower bound on $|OPT|$: $|OPT| \geq |OPT_\Delta| = y$. Denote by s the semi-perimeter of Δ, by R the radius of its circumscribed circle, and by r the radius of its inscribed circle.

FACT 5.1. *For an acute triangle Δ, $s > 2R$.*

Proof. If A, B, C are the angles of Δ, this is equivalent to

$$R(\sin A + \sin B + \sin C) > 2R.$$

After simplification with R, this is a well known inequality in the geometry of an acute triangle ([4], page 18).

FACT 5.2. *For an acute triangle Δ, $y = \frac{2rs}{R}$.*

A proof of this equality can be found in [4], page 86.

CLAIM 1. *For an acute triangle Δ, $r < \frac{y}{4}$.*

Proof. Putting the above together, we get $r = \frac{Ry}{2s} < \frac{Ry}{4R} = \frac{y}{4}$.

The length of the output tour is bounded as follows

$$|C_L| = 2\pi r < 2\pi \frac{y}{4} = \frac{\pi}{2} y \leq \frac{\pi}{2}|OPT|.$$

Case 2. Δ is an obtuse triangle. In this case, $|OPT_\Delta| = 2h$, where h is the length of the altitude corresponding to the obtuse angle, say A. Clearly $2h$ is a lower bound on $|OPT|$: $|OPT| \geq |OPT_\Delta| = 2h$.

FACT 5.3. *For any triangle Δ, $h > 2r$.*

Proof. We express the area S of the triangle (with side lengths a, b, c) in two ways, from which we get the inequality,

$$2S = a \cdot h = (a + b + c)r,$$

which implies that

$$\frac{h}{r} = \frac{a+b+c}{a} > \frac{a+a}{a} = 2.$$

Using this, we have $|OPT_\Delta| = 2h > 4r$. The length of the output tour is bounded as follows

$$|C_L| = 2\pi r = \frac{\pi}{2}(4r) < \frac{\pi}{2}|OPT_\Delta| \leq \frac{\pi}{2}|OPT|.$$

Thus, in both cases, the approximation ratio is $\frac{\pi}{2} \leq 1.58$.

We now describe the algorithm for computing C_L. The distance $d(p, l)$ from a point p of coordinates (x_0, y_0) to a line l of equation $ax + by + c = 0$ is

$$d(p, l) = \frac{|ax_0 + by_0 + c|}{\sqrt{a^2 + b^2}}.$$

Let the lines in L have the equations

$$(l_i :) \qquad a_i x + b_i y + c_i = 0, \; i = 1, \ldots, n.$$

Finding a minimum touching circle amounts to finding (x, y), (the center coordinates) and a minimum radius z, s.t.

$$\frac{|a_i x + b_i y + c_i|}{\sqrt{a_i^2 + b_i^2}} \leq z, \quad i = 1, \ldots, n.$$

This is equivalent to solving the following 3-dimensional linear program

min z :
subject to

$$\begin{aligned} \frac{a_i x + b_i y + c_i}{\sqrt{a_i^2 + b_i^2}} &\leq z \\ \frac{-a_i x - b_i y - c_i}{\sqrt{a_i^2 + b_i^2}} &\leq z, \quad i = 1, \ldots, n, \end{aligned}$$

which takes $O(n)$ time [12].

6 Conclusion

Several open problem remain, including

(1) Is there a constant-factor approximation algorithm for arbitrary connected regions in the plane? What if the regions are disconnected? (giving us a geometric version of a "one-of-a-set TSP")

(2) What approximation bounds can be obtained in higher dimensions? Our packing arguments for disjoint disks lift to higher dimensions, but our other methods do not readily generalize.

A particularly intriguing special case is the generalization of the case of infinite straight lines: What can be said in 3-space for the TSPN on a set of lines or of planes?

(3) Is there a PTAS for general pairwise-disjoint regions in the plane?

References

[1] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Appl. Math.*, 55:197–218, 1994.

[2] S. Arora. Nearly linear time approximation schemes for Euclidean TSP and other geometric problems. *J. of the ACM*, 45(5):1–30, 1998.

[3] J. L. Bentley. Fast algorithms for geometric traveling salesman problems. *ORSA J. Comput.*, 4(4):387–411, 1992.

[4] O. Bottema, R. Ž. Djordjević, R. Janić, D. S. Mitrinović, and P. M. Vasić. *Geometric Inequalities.* Wolters-Noordhoff, Groningen, 1969.

[5] S. Carlsson, H. Jonsson, and B. J. Nilsson. Finding the shortest watchman route in a simple polygon. In *Proc. 4th Annu. Internat. Sympos. Algorithms Comput.*, volume 762 of *Lecture Notes Comput. Sci.*, pages 58–67. Springer-Verlag, 1993.

[6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman, New York, NY, 1979.

[7] J. Gudmundsson and C. Levcopoulos. A fast approximation algorithm for TSP with neighborhoods. *Nordic Journal of Computing*, 6:469–488, 1999.

[8] J. Gudmundsson and C. Levcopoulos. Hardness result for TSP with neighborhoods. Technical Report LU-CS-TR:2000-216, Department of Computer Science, Lund University, Sweden, 2000.

[9] M. Jünger, G. Reinelt, and G. Rinaldi. The traveling salesman problem. In M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, editors, *Network Models*, Handbook of Operations Research/Management Science, pages 225–330. Elsevier Science, Amsterdam, 1995.

[10] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys, editors. *The Traveling Salesman Problem.* John Wiley & Sons, New York, NY, 1985.

[11] C. Mata and J. S. B. Mitchell. Approximation algorithms for geometric tour and network design problems. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 360–369, 1995.

[12] N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31:114–127, 1984.

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

[14] J. S. B. Mitchell. Geometric shortest paths and network optimization. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 633–701. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[15] B. J. Nilsson. *Guarding Art Galleries — Methods for Mobile Guards.* PhD thesis, Lund University, 1995.

[16] C. H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoret. Comput. Sci.*, 4:237–244, 1977.

[17] S. B. Rao and W. D. Smith. Improved approximation schemes for traveling salesman tours. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, 1998.

[18] G. Reinelt. Fast heuristics for large geometric traveling salesman problems. *ORSA J. Comput.*, 4:206–217, 1992.

[19] X. Tan, T. Hirata, and Y. Inagaki. Corrigendum to 'An incremental algorithm for constructing shortest watchman routes'. Manuscript (submitted to internat. j. comput. geom. appl.), Tokai University, Japan, 1998.