A Constant-Factor Approximation Algorithm for TSP with Pairwise-Disjoint Connected Neighborhoods in the Plane

Joseph S. B. Mitchell*
Stony Brook University
Stony Brook, NY 11794-3600
joseph.mitchell@stonybrook.edu

ABSTRACT

In the Euclidean TSP with neighborhoods (TSPN) problem we seek a shortest tour that visits a given set of n *neighborhoods*. The Euclidean TSPN generalizes the standard TSP on points.

We present the first constant-factor approximation algorithm for planar TSPN with pairwise-disjoint connected neighborhoods of any size or shape. Prior approximation bounds were $O(\log n)$, except in special cases. The methods also apply to the case of arbitrarily overlapping regions that are convex.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms

1. INTRODUCTION

The traveling salesperson problem with neighborhoods (or, TSPN) is a generalization of the classic TSP in which the traveler is to visit a of neighborhoods rather than a set of specific points. We focus primarily on the case in which the neighborhoods are *connected* regions in the plane.

The TSPN was introduced by Arkin and Hassin [1], who gave the first algorithmic study of the problem and provided the first approximation algorithms for special cases of the problem. The problem arises naturally in in various optimal covering tour problems, network design problems, relay placement in sensor networks, robotics mission planning, and robot localization.

*Partially supported by grants from the National Science Foundation (CCF-0729019), Metron Aviation, and NASA Ames.

SCG'10, June 13–16, 2010, Snowbird, Utah, USA.

Our main result is a polynomial-time algorithm that gives the first constant-factor approximation for the TSPN for the general case of disjoint connected neighborhoods in the plane. The regions can be of arbitrary size and shape. In contrast, the best approximation ratio previously known for the problem is $O(\log n)$. Only special cases in which the regions are fat and/or of nearly the same size (diameter) were known previously to have constant-factor approximation algorithms or a polynomial-time approximation scheme (PTAS). Our method also gives a constant-factor approximation for the case of arbitrarily overlapping convex regions.

1.1 Related Work

The TSP with neighborhoods has been actively studied since its first formalization and study by Arkin and Hassin [1].

The TSPN is a generalization of the classic geometric TSP on points and is therefore NP-hard. While the TSP on point sets in geometric domains admits a PTAS by the results of Arora [2], Mitchell [15], and Rao and Smith [19], the TSP with neighborhoods seems to be a considerably harder optimization problem to approximate.

Dumitrescu and Mitchell [7] give a PTAS, based on the m-guillotine method, if the regions are all about the same size (the ratio of largest to smallest diameter is bounded), have bounded depth (no point lies in more than a constant number of regions), and are *fat*. Feremans and Grigoriev [11] give a PTAS, based on the method of Arora [2], if the regions are disjoint fat polygons of comparable size in the plane; their algorithm applies also in higher dimensions.

By mapping to the "one-of-a-set" TSP, de Berg et al. [5] gave an $O(1)$-approximation for disjoint fat convex regions. Elbassioni et al. [8, 10] improved the running time and generalized to the case in which the neighborhoods to be visited may be arbitrary sets of points, with each set lying within a (disjoint) fat region that is not necessarily convex. Elbassioni, Fishkin, and Sitters [9, 10] consider the case in which the corresponding regions are intersecting, convex, and fat, of comparable size; they give an $O(1)$-approximation.

Mitchell [18] obtained a PTAS for bounded depth fat regions in the plane, allowing them to be of arbitrary sizes. His definition of "fat" is very weak; he requires only that the area is at least a constant times the square of the diameter. His results apply also to the case in which the regions to be visited may be disconnected sets of points that each lie within one of the (nearly) disjoint fat polygons.

Very recently, Chan and Elbassioni [4] have given a quasi-polynomial-time approximation scheme (QPTAS) for geometric instances of TSPN in any fixed dimension for the case of *fat, weakly disjoint* regions; their results apply also to metric spaces of bounded doubling dimension.

If the regions are allowed to overlap, the problem is APX-hard [5, 20], even if the regions are segments all of very nearly the same length [9]. The problem is also APX-hard if the regions are disconnected, with each being a pair of points in the plane [6]. Thus, we do not expect to find a PTAS for the general case. There might be a PTAS in the case of disjoint regions (the case for which we give the first constant-factor approximation); this is a fascinating open problem.

For connected sets in the plane, possibly overlapping, the best known approximation ratio is $O(\log n)$, given in an algorithm discovered more than 15 years ago [14]; the running time has been improved more recently [9, 12].

For further background about the geometric TSP, the TSP with neighborhoods, and related problems, see the surveys [3, 16, 17].

1.2 Our Contribution

All prior constant-factor approximation algorithms (and PTAS's) relied on exploiting special structure that comes from assuming that regions are either all about the same size or all are fat in some sense. Fatness was crucial in order to apply a packing argument to give effective lower bounds on the length of an optimal tour. Even the special case of horizontal line segments has had no constant-factor approximation prior to this work.

Several new ideas are needed in order to address the general case in which the given regions, $\{P_1, P_2, \dots, P_n\}$, can be of arbitrary sizes and are not fat. We outline our approach:

1. We address the "skinniness" of general regions by surrounding each region by its four "directional hulls", which we prove are fat, for an appropriate choice of one of *two* coordinate systems. This allows us to partition the regions into two sets ("blue" regions and "red" regions); we separately find approximating tours of each set, and then appeal to the Combination Lemma of [1] to combine them into a tour of the entire set. Thus, it suffices to consider only the set of blue regions, which have all four of their directional hulls in our (x, y)-coordinate system fat.

2. Let T be a tour that visits all of the fat blue directional hulls. We convert T into a polygonal subdivision, G, each of whose faces are *histograms*; this increases the total length of T by a constant factor.

3. We argue that if all four of the (fat) directional hulls of the original (possibly skinny) input region P_i is visited by G, then G also visits P_i. This follows from the fact that the faces of G are histograms. (Thus, if G happens to visit all of the directional hulls of all of the (blue) regions, then G already visits all of the regions themselves, and we are done.)

   This allows us to conclude with our first result (Proposition 1), a problem reduction: A c-approximation for the case of computing a tour T for arbitrary *fat* connected regions implies an $O(c)$-approximation for arbitrary connected regions, *even if the input regions are not disjoint*.

4. Since we do not yet know a method to compute an $O(1)$-approximation for arbitrary fat connected regions in the plane, we need some further ideas to obtain our second (and main) result, a constant-factor approximation for the case of *disjoint* connected regions (Theorem 1).

   We assume now that our original regions are disjoint. Their minimum enclosing balls (MEB's) are fat, of course, but they overlap. We select a disjoint subset, $\mathcal{E}_0$, of the MEB's; we do this greedily, in order of increasing size. We use the algorithm of [18] to compute an approximately optimal tour, T, of the disjoint subset $\mathcal{E}_0$ of MEB's. We convert T into a polygonal subdivision, G, each of whose faces are *histograms*; this increases the total length of T by a constant factor.

5. The fact that the regions $\mathcal{E}_0$ were selected greedily by size allows us to argue that any region P_i that is *not* visited already by T must be *close* to the tour T (and therefore close to the boundaries of faces of G), where "close" means within distance that is $O(diam(P_i))$. (I.e., there can be no "missed" regions P_i that are substantially interior to some face of G.) This motivates the definition of a *stratified grid* within a histogram face, H; this grid has small squares near the boundary of H, then progressively larger squares as the distance to the boundary increases. Since regions P_i that are "far" from the boundary of H must be "large", we know that we are able to visit all regions with a slight refinement of the squares of the stratified grid. Also, each square of the stratified grid can be attached to the boundary of H by a connection whose length is proportional to the size of the square.

6. We then focus on a single histogram face, H, and its stratified grid. We first argue a lower bound on the total length of any forest F that visits all regions P_i interior to H and for which the union $F \cup \partial H$ is connected: Such a forest must have length at least a constant fraction of the weight of a minimum-weight cover of the regions P_i by squares of the stratified grid, where "weight" refers to the sum of the sizes of the squares. Thus, our goal is to solve a weighted set cover problem that has special structure. We are able to exploit the disjointness of the regions P_i to obtain a dynamic programming algorithm to compute a minimum-weight cover. Alternatively, dynamic programming can be applied to solve the problem in which the regions are *not* disjoint, provided the regions are *convex*.

2. PRELIMINARIES

Let $\mathcal{R} = \{P_1, P_2, \dots, P_n\}$ be the input set of n disjoint, connected regions in the plane.

We say that a tree or a tour T *visits* region P_i if $T \cap P_i \neq \emptyset$. We say that T visits $\mathcal{R}$ if T visits each of the regions P_i, for $i = 1, \dots, n$. Let L^* denote the length of an optimal tour.

We utilize some basic notation and observations from [18]. Let R_0 be a minimum-diameter (blue) axis-aligned rectangle that intersects or contains all regions P_i, and let D be the diameter of R_0. (Note that R_0 is easily computed in polynomial time by critical placement arguments; for our purposes, a constant-factor approximation suffices, and this is even more readily computed.) Then, bounds on the length, L^*, of an optimal tour can be written in terms of D; specifically, as shown in [18], Lemma 2.2, $2D \le L^* \le nD$. Let $\mathcal{G}$ denote the regular grid (lattice) of points $(i\delta, j\delta)$, for integers i and j, where $\delta = D/n$. We let Γ_i denote the subset of grid points $\mathcal{G}$ at distance at most $\delta/\sqrt{2}$ from region P_i. Then, $\Gamma_i \neq \emptyset$. Note that it may be that $\Gamma_i = \Gamma_j$ for distinct regions P_i and P_j, $i \neq j$. (In particular, many "tiny" regions may all map to the same singleton grid point.) By Lemma 2.3 of [18] (with $\epsilon = 1$), we know that any tour T, of length L, that visits $\mathcal{R} = \{P_1, P_2, \ldots, P_n\}$ can be modified (by adding n detours, each of length at most $2\delta/\sqrt{2} = \sqrt{2}D/n \le L^*/(\sqrt{2}n)$) to be a tour $T_{\mathcal{G}}$, of length $O(L)$, that visits $\{\Gamma_1, \Gamma_2, \ldots, \Gamma_n\}$. Similarly, any tour $T_{\mathcal{G}}$, of length $L_{\mathcal{G}}$, that visits $\{\Gamma_1, \Gamma_2, \ldots, \Gamma_n\}$ can be modified to be a tour T, of length $O(L_{\mathcal{G}})$, that visits $\mathcal{R} = \{P_1, P_2, \ldots, P_n\}$. Further, by Lemma 2.4 of [18], it suffices to search for optimal tours within the ball, $B(c_0, 2D)$, of radius $2D$ centered at the center point, c_0, of R_0; this *localizes* our problem.

Thus, for our approximation purposes, it suffices to consider polygonal tours whose vertices lie on the grid $\mathcal{G}$ (and within distance $O(D)$ of the point c_0); in fact, giving up at most a factor $\sqrt{2}$, we can assume, for simplicity, that we are searching over the set of rectilinear tours on $\mathcal{G}$, with all edges axis-parallel, on the grid. By rescaling, we can also assume that $\delta = 1$, without loss of generality.

Finally, it suffices to give an approximation algorithm for the *minimum spanning tree with neighborhoods* (MSTN) problem, since associated with a spanning tree T of length $|T|$ is a tour of length $2|T|$ obtained by walking around the tree T. Thus, we desire a tree, T, that visits all of the regions $\mathcal{R} = \{P_1, P_2, \ldots, P_n\}$ and has (Euclidean) length close to $|T^*|$, the length of an optimal spanning tree, T^*. We can assume that the tree is rectilinear, lies on the integer grid, and the goal is to find a short tree T that visits, or comes close (within distance $1/\sqrt{2}$) to, each region P_i.

3. REDUCTION TO CASE OF OVERLAPPING FAT REGIONS

We consider two Cartesian coordinate systems: (1) The "base" (or "blue") system has coordinates (x, y), and (2) the "rotated" (or "red") system is rotated by 45 degrees with respect to the blue system and has coordinates (x', y'). Each region P_i has an axis-aligned bounding box in each of the two coordinate systems; we let $BB(P_i)$ denote the blue bounding box and $BB'(P_i)$ denote the red bounding box.

We let $E^{+x}(P_i) \subseteq BB(P_i)$ denote the set of all points $p \in BB(P_i)$ for which the rightwards ray (in the $+x$ direction) from p intersects P_i. We similarly define $E^{-x}(P_i)$, $E^{+y}(P_i)$, and $E^{-y}(P_i)$, and refer to these as the *blue directional hulls* of P_i. Exactly analogously, we define the four red directional hulls of P_i, with respect to the red coordinates x', y' and the red bounding box, $BB'(P_i)$. Refer to Figure 1.

We use the same definition of fatness as in [18]: a connected planar region X is α-*fat* (or simply *fat*) if X contains a disk of radius at least $diam(X)/\alpha$, where $diam(X)$ denotes the diameter of X. Note that this definition of fat applies to convex as well as non-convex regions.

LEMMA 1. *Let X be a connected subset of the plane. Either the four blue directional hulls of X are fat or the four red directional hulls of X are fat.*

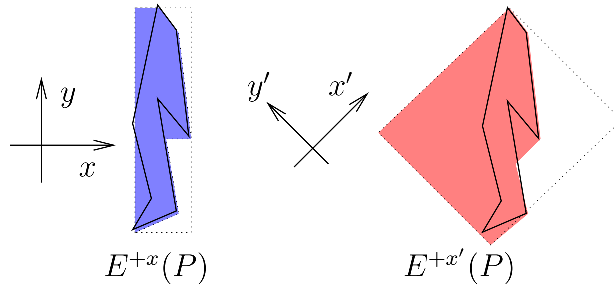

Figure 1: Directional hulls: blue $E^{+x}(P)$ and red $E^{+x'}(P)$ for polygon P.

PROOF. If the four blue directional hulls are fat, we are done. Thus, assume that at least one of the blue directional hulls, say $E^{-y}(X)$, is not fat.

We claim that all four red directional hulls are fat.

There are two cases: (1) The aspect ratio of $BB(X)$ is at least 10; and (2) The aspect ratio of $BB(X)$ is less that 10. (The number 10 is not specially significant or optimally chosen, but it suffices for our purposes.)

In case (1) we assume, without loss of generality, that $BB(X)$ has its longer sides horizontal (in the (x, y)-coordinate system), and we partition $BB(X)$ into a chain of 10 subrectangles by partitioning the longer dimension of $BB(X)$ into 10 equal-length pieces. By the fact that $BB(X)$ is a minimal bounding box, we know that there must be a point $a \in X$ within the first subrectangle and a point $b \in X$ within the last (10th) subrectangle of the chain. (In fact, by connectivity, each subrectangle has a connected portion of X within it.) We readily see that the red bounding box $BB'(\{a, b\})$ extends enough outside $BB(X)$ to contain the (x', y')-aligned square, A, of size $diam(X)/10$ in the upper corner of $BB'(\{a, b\})$; similarly, there is a square B of size $diam(X)/10$ in the lower corner of $BB'(\{a, b\})$. Refer to Figure 2. Since X is connected, there is a path, γ, within X (and therefore within $BB(X)$ and $BB'(X)$) connecting a and b. We see that $A \subset BB'(\{a, b\}) \subseteq BB'(X)$, and, by definition of directional hulls, $A \subset E^{-y'}(X)$ and $A \subset E^{-x'}(X)$; similarly, $B \subset BB'(\{a, b\}) \subseteq BB'(X)$, and $B \subset E^{+y'}(X)$ and $B \subset E^{+x'}(X)$. Thus, $BB'(X)$ has all four of its red directional hulls fat (containing A or B, each of which has a constant fraction fraction of the area of $BB'(X)$).

In case (2), we again assume, without loss of generality, that $BB(X)$ has its longer sides horizontal (in the (x, y)-coordinate system). We partition the longer dimension of $BB(X)$ into 10 equal-length pieces, and consider the 10 corresponding squares within $BB(X)$ along the bottom edge of $BB(X)$. Refer to Figure 2. Since $E^{-y}(X)$ is not fat, we know that none of these squares is disjoint from X (otherwise, such a square would lie inside $E^{-y}(X)$, contradicting the assumption that $E^{-y}(X)$ is not fat). Thus, we know that there must be a point $a \in X$ within the leftmost square and a point $b \in X$ within the rightmost square. While the red bounding box $BB'(\{a, b\})$ extends enough below $BB(X)$ to contain the (x', y')-aligned square, B, of size $diam(X)/10$ in the lower corner of $BB'(\{a, b\})$, the similar square, A, in the upper corner of $BB'(\{a, b\})$ may lie partly or entirely within $BB(X)$. We have $B \subset BB'(\{a, b\}) \subseteq BB'(X)$, $B \subset$

$E^{+y'}(X)$ and $B \subset E^{+x'}(X)$; thus, $E^{+y'}(X)$ and $E^{+x'}(X)$ are fat. We need a different argument to show that $E^{-y'}(X)$ and $E^{-x'}(X)$ are fat, since it may be that X enters square A, so it is not immediate that A is contained in $E^{-y'}(X)$ and $E^{-x'}(X)$. Consider a path γ within X from b to a. Such a path must at some point cross the polygonal path (p, q, r) shown in Figure 2. (Here, pq lies on the vertical (y-parallel) line supporting A on its right; qr lies on the y'-parallel line supporting A from below.) If the path γ hits segment pq before segment qr, then γ is a witness to the containment of square C within $E^{-y}(X)$, a contradiction to the fact that $E^{-y}(X)$ is assumed not to be fat. Otherwise, if γ hits segment qr before (or at the same point as) it hits segment pq, then γ is a witness to the containment of A in the red directional hull $E^{-y'}(X)$, showing that $E^{-y'}(X)$ is fat. A similar argument shows that $E^{-x'}(X)$ is fat. $\square$

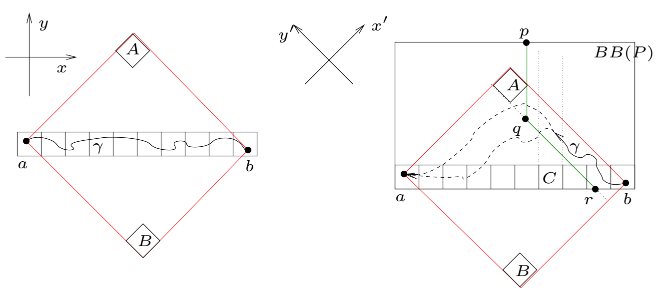

Figure 2: Proof of Lemma 1. Left: Case (1), aspect ratio is high; Right: Case (2), aspect ratio is low.

The next lemma shows that the problem of visiting a set of arbitrary connected regions can be reduced to the problem of visiting a set of fat regions (that may overlap, even if the original regions are disjoint).

LEMMA 2. *Let T be a tree that visits all four of the blue directional hulls of each of the regions, $\mathcal{R} = \{P_1, P_2, \ldots, P_n\}$. Then there exists a connected planar network, $G \supseteq T$, of length $|G| = O(|T|)$ that visits the regions in $\mathcal{R}$.*

PROOF. We can assume that T is rectilinear (with axis-parallel edges). First, we add to T the boundary of the (axis-aligned) bounding box, $BB(T)$, of tree T. The set $BB(T) \setminus T$ consists of a set of simple rectilinear polygonal *pockets*, each with at least one edge (the pocket "lid") on the boundary of $BB(T)$.

For each pocket, Q, we build its *histogram decomposition*, as follows. Let e be one edge of Q; for specificity, we can take e to be a lid of pocket Q, and assume that e is horizontal. We illuminate edge e within Q, using *rectilinear visibility*: the set $VP(e, Q)$ of all points in Q that are rectilinearly visible from e is the set of all points $q \in Q$ for which the axis-parallel segment, $qq_\perp$, joining q to the point $q_\perp$ that is the orthogonal projection of q onto the line through e, satisfies $q_\perp \in e$ and $qq_\perp \subset Q$. The set $VP(e, Q)$ is a (rectilinear) histogram with base e. (A *histogram* is a polygon H having an edge β, the *base* of the histogram, such that the segment $pp_\perp \subset H$ for every $p \in H$, where $p_\perp$ is the perpendicular projection of p onto the line through β.) The boundary of $VP(e, Q)$ consists of segments contained in the tree T and chord segments, known as *windows*, that pass through the interior of Q. Each window segment w is a chord separating $VP(e, Q)$ from a subpolygon ("subpocket") of Q. We then consider each window w, and the corresponding subpocket Q', and determine the histogram, $VP(w, Q')$, consisting of the points within Q' (and therefore within Q) that are rectilinearly visible from w. We continue this process until each subpocket has no windows (only edges contained in T). This process is a rectilinear version of the "window partition tree" used to compute link distances in polygons [21]. It was used also in [14], where it was argued that the total length of the resulting histogram decomposition is $O(|T|)$, where $|T|$ denotes the total length of tree T: The key observation is that any axis-parallel chord of Q crosses at most two windows; thus, the length of each window (the edges that were added to T to make the decomposition) can be charged to the length of T.

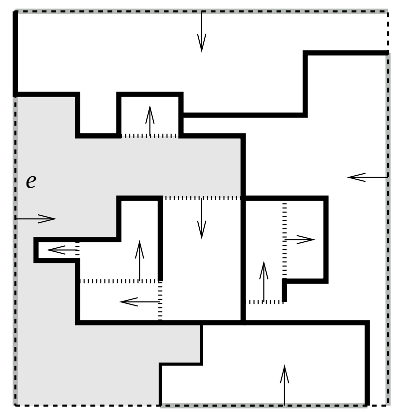

Figure 3: The histogram decomposition of tree T (shown bold black). $VP(e, Q)$ is shaded light gray for the edge e (shown in medium gray) on the boundary of the bounding box. Windows are shown dashed, with a short arrow indicating the direction of illumination into the subpocket.

Next we argue that the connected planar network, G, given by the union of the histogram decompositions of the pockets of T, visits all of the original regions, $\mathcal{R}$. Suppose to the contrary that P_i lies fully interior to a (histogram) face, H, of G. (We assume here that H is a bounded face, not the face at infinity; however, the argument below applies even more readily to the case that H is the face at infinity, i.e., the complement of the rectangle $BB(T)$.) Suppose that H is a *vertical histogram*, meaning its base is horizontal (x-parallel). Since T visits all four of the blue directional hulls of P_i, we know that G (whose edges are a superset of T) visits the hull $E^{+y}(P_i)$ in particular.

If some point $p \in E^{+y}(P_i)$ lies left of the leftmost (vertical) side of H, then, by definition of $E^{+y}(P_i)$, the upwards ray from p intersects P_i at some point $p' \in P_i$; thus, p' lies to the left of the bounding box of H, so p' is a point of P_i outside H, a contradiction. Similarly, we know that no part of $E^{+y}(P_i)$ lies to the right of the rightmost (vertical) side of H.

If the bottom edge (base) of $E^{+y}(P_i)$ lies below the bottom edge (base) of H, then some point of P_i that supports the base of $E^{+y}(P_i)$ lies outside H, a contradiction.

Finally, if some point $p \in E^{+y}(P_i)$ lies above the top (x-monotone) boundary of H, then, by definition of $E^{+y}(P_i)$, the upwards ray from p intersects P_i at some point $p' \in P_i$; thus, p' lies above the top boundary of H, so p' is a point of P_i outside H, a contradiction. Thus, $E^{+y}(P_i)$ lies completely below the top boundary of H. Since $E^{+y}(P_i)$ also lies completely above the base of H and to the right of the left side of H and to the left of the right side of H, we conclude that $E^{+y}(P_i)$ lies interior to H, a contradiction to the fact that G visits $E^{+y}(P_i)$. $\square$

Based on the above lemmas, we suggest an approach to computing a constant-factor approximation for an arbitrary set of connected regions $\mathcal{R}$ in the plane. First, we partition the regions $\mathcal{R}$ into two sets – the "blue" regions that have all four blue directional hulls fat, and the remaining "red" regions, which, by Lemma 1, must have all four red directional hulls fat. Then we compute a tree, T_{blue}, visiting all of the (fat) blue directional hulls of the blue regions and separately a tree, T_{red}, visiting all of the (fat) red directional hulls of the red regions. By Lemma 2, we know that T_{blue} visits all of the blue regions and T_{red} visits all of the red regions. We can then add a segment of length $O(D)$ (and recall that $L^* \geq 2D$) connecting T_{blue} and T_{red}, forming a single tree that visits all of the regions. If we have a constant-factor approximation algorithm for computing a short tree spanning a set of (overlapping) fat regions, then this approach yields an overall constant-factor approximation for MSTN (and therefore for TSPN). Note that this approach applies to *arbitrarily overlapping* regions $\mathcal{R}$; we made no assumption about disjointness. We summarize:

PROPOSITION 1. *A c-approximation for TSPN/MSTN for arbitrary connected fat regions in the plane yields an $O(c)$-approximation for TSPN/MSTN for arbitrary connected regions in the plane.*

The difficulty with the above approach is that we do not yet know how to give a constant-factor approximation for TSPN on a set of arbitrarily overlapping fat regions. Thus, we take a different approach. We *do* know how to compute a nearly optimal tree/tour of a set of *disjoint* fat regions (even if the regions are not disjoint, but have bounded depth); in fact, a PTAS is given in [18]. Thus, we will compute such an approximating tree/tour T of an appropriately chosen subset of *disjoint* fat regions, and then augment T so that it visits all of those original regions that T does not already visit. Since we can afford to add to T its histogram decomposition, and still have total length $O(|T|)$, it suffices to solve the augmentation problem for regions that lie within a single face (histogram), H, of the decomposition.

It is important for our approach that we have additional structure associated with the "floating" regions that lie interior to faces of the histogram decomposition. This structure comes from the method we use to select the set of disjoint fat regions. We now describe that method.

Let $\mathcal{E} = \{E_1, E_2, \ldots\}$ be the minimum enclosing balls (MEB's) of the input regions $\mathcal{R}$. We pick a subset $\mathcal{E}_0 \subseteq \mathcal{E}$ of *disjoint* MEB's by an iterative algorithm, examining the E_i's in increasing order of diameter, $diam(E_i)$. As we examine E_i, if it is disjoint from the MEB's already placed in the set $\mathcal{E}_0$, then we add it to $\mathcal{E}_0$; otherwise, we skip it and go to the next larger MEB. (In case of ties, we examine the MEB's of equal diameter in arbitrary order.)

Next, we compute a tree, $T_{\mathcal{E}_0}$, visiting the MEB's $\mathcal{E}_0$. For this, we can use the PTAS of [18]. We know that the length of an optimal tree/tour of $\mathcal{R}$ is at least the length of an optimal tree/tour of $\mathcal{E}_0$. We then extend $T_{\mathcal{E}_0}$ to a histogram decomposition, G, of the bounding box and pockets of $T_{\mathcal{E}_0}$.

Now, the remaining problem is that of augmenting G so that the augmented graph, $\overline{G}$, is connected and visits *all* of the regions $\mathcal{R}$. It suffices to consider one face, H, in the histogram decomposition G, and the set $\mathcal{R}_H$ of (input) regions that lie interior to H. We will exploit special structure of the set of regions $\mathcal{R}_H$ with respect to the boundary of H.

LEMMA 3. *Each $P \in \mathcal{R}_H$ lies within distance $O(diam(P))$ of the boundary, ∂H, of H.*

PROOF. Let $P \in \mathcal{R}_H$. If the MEB of P is in $\mathcal{E}_0$, then we know that $T_{\mathcal{E}_0}$ comes within distance $diam(P)$ of P. If the MEB of P is not in $\mathcal{E}_0$, then there must be another region $E' \in \mathcal{E}_0$ with diameter at most $diam(P) = diam(E)$ and with $E \cap E' \neq \emptyset$. Since $T_{\mathcal{E}_0}$ visits E', $T_{\mathcal{E}_0}$ comes within distance $diam(E') \leq diam(P)$ of E', and therefore within distance at most $2 \cdot diam(P)$ of P. Thus, P lies within distance $O(diam(P))$ of ∂H. $\square$

In the next section, we give an algorithm to solve (approximately) the *histogram problem*, which is the following: For a histogram H and a set $\mathcal{R}_H$ of connected regions interior to H, with each $P \in \mathcal{R}_H$ within distance $O(diam(P))$ of ∂H, construct a minimum-length network, F, such that $F \cup \partial H$ is connected and visits all of the regions $\mathcal{R}_H$.

4. SOLVING THE HISTOGRAM PROBLEM

Let $\mathcal{R}_H$ be a set of disjoint regions interior to a histogram H. As discussed earlier, we can assume that the vertices of H are integral (on the grid $\mathcal{G}$). Also, without loss of generality, H is a vertical histogram (i.e., H has a horizontal base). By Lemma 3, we know that each $P \in \mathcal{R}_H$ lies within distance $O(diam(P))$ of ∂H. Our goal is to construct a minimum-length network of edges, F, such that $F \cup \partial H$ is connected and visits all of the regions $\mathcal{R}_H$. If F is truly of minimum-length, it is a forest (since any cycle can be cut); we let F^* denote an optimal forest, with length $|F^*|$. The network we construct for our approximation will not, in general, be a forest, but it will have length $O(|F^*|)$.

4.1 The Stratified Grid

We now define the *stratified grid* decomposition of H. Informally, the decomposition is formed as follows: We place squares of size 1 (forming the first strata, S_1) inside H, around its boundary; this results in an "eroded" (shrunk by L_∞ distance 1) region. We then place squares of size 2 (forming the second strata, S_2) inside each component of the eroded region, around its boundary; this results in the eroded region getting smaller by L_∞ distance 2. This process continues until all of H has been packed (and covered) by squares of sizes 1, 2, 4, 8, etc., forming the strata $S_1, S_2, S_4, \ldots$. This description is informal since it does not take into account the situation in which the eroded region may contain narrow "necks" of insufficient width to accommodate squares of the appropriate size for the next strata of erosion. Thus, we make the process more formal with the following definitions. (There are multiple ways to establish stratified grids that work equally well to the method we describe here; we will point out below the basic properties that we need the stratification to have.)

The *standard squares* of size 2^i (or *standard* 2^i*-squares*) are those squares of size 2^i whose defining coordinates are integer multiples of 2^i ($i = 0, 1, 2, \ldots$).

The stratified grid decomposition of H is a polygonal subdivision of H into standard squares of sizes 1, 2, 4, 8, etc; we let S_i denote the set of squares of size i. We abuse notation slightly by using S_i also to denote the region within H that is the union of the squares of S_i.

Let H_{-1} denote the 1-offset of H, consisting of the set of all points within H that are at L_∞ distance at least 1 from the boundary of H. (Note that H_{-1} may consist of more than one connected component, each of which is a histogram with integral coordinates.) Let $H_{-1}^{(2)}$ denote the union of all standard 2-squares within H_{-1}. (Note that $H_{-1}^{(2)}$ is a union of histograms whose vertex coordinates are multiples of 2; thus, $H_{-1}^{(2)}$ can be exactly tiled by standard 2-squares.) We define the *1-strata*, S_1, to be the set of all (standard) 1-squares that lie within $H \setminus H_{-1}^{(2)}$. Note that each 1-square of S_1 is either in contact with the boundary of H or at L_∞ distance 1 from the boundary of H.

We define H_{-2} to be the 2-offset of $H_{-1}^{(2)}$, and $H_{-2}^{(4)}$ to be the union of all standard 4-squares within H_{-2}. The *2-strata*, S_2, is the set of all standard 2-squares that lie within $H_{-1}^{(2)} \setminus H_{-2}^{(4)}$. We similarly define H_{-2^i} and the regions $H_{-2^i}^{(2^{i+1})}$ for $i = 0, 1, 2, \ldots$; the 2^i-strata, S_{2^i}, is the set of all standard 2^i-squares that lie within $H_{-2^{i-1}}^{(2^i)} \setminus H_{-2^i}^{(2^{i+1})}$. (Note that both $H_{-2^{i-1}}^{(2^i)}$ and $H_{-2^i}^{(2^{i+1})}$ can be exactly tiled by standard 2^i-squares.) An example of a stratified grid of H is shown in Figure 4.

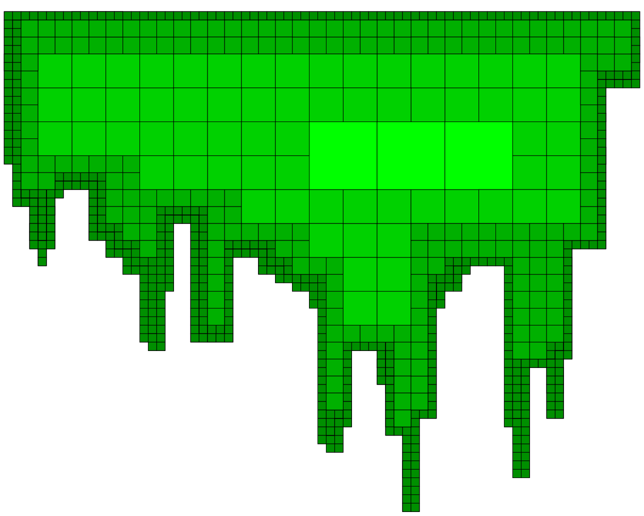

Figure 4: The stratified grid of a histogram H.

By the specification of the stratified grid, each grid square σ of size $|\sigma| = 2^i$ lies at L_∞ distance $\Theta(|\sigma|)$ from ∂H, (More precisely, the L_∞ distance is between $2^i - 1$ and $3 \cdot 2^i - 2$.)

We subdivide each grid square of the stratified grid into a constant number of subsquares. This is done so that no region $P \in \mathcal{R}_H$ can lie substantially inside any square (after subdivision), where "substantially" means at distance more than $1/\sqrt{2}$ (recall that we only need to get within distance $\delta/\sqrt{2}$ of each region, and we rescaled so that $\delta = 1$). There is no need to subdivide the finest squares (the squares of size 1 or 2). For all other squares σ, it suffices to partition σ with a "+" *mullion* into 4 equal-size "panes" (subsquares), each of half the size. We consider these mullions to be associated with each square σ having $|\sigma| > 1$. Then, since σ is at distance at least $|\sigma| - 1$ to ∂H, and we know that any $P \in \mathcal{R}_H$ is at distance $O(diam(P))$ from ∂H, we know that, while P may lie inside a square σ, it cannot avoid being intersected by the mullions of σ, since each subsquare has diameter $|\sigma|/\sqrt{2}$, and we have that $|\sigma|/\sqrt{2} < |\sigma| - 1$ for sizes $|\sigma| \geq 4$.

4.2 A Lower Bound on OPT

We establish a lower bound on the length, $|F^*|$, of a minimum-length forest, F^* whose addition to ∂H spans all of the regions, with $F^* \cup \partial H$ being connected. We argue that we can replace the forest F^* with the network consisting of the boundaries of the set S_{F^*} of all grid squares that F^* intersects, and that this network has length $O(|F^*|)$. This is a nontrivial statement and utilizes the special structure of the stratified grid, since, in a general tiling by squares of different sizes, a curve of length ℓ can intersect a set of n squares whose total perimeter is $\Omega(|\ell| \log n)$; see the example in [18].

LEMMA 4. *The total perimeter of the stratified grid squares in the set S_γ of squares intersected by a curve $\gamma \subset H$ of length $|\gamma|$ is $O(|\gamma|)$ plus the sizes of the squares containing γ's endpoints.*

PROOF. Suppose γ's starting endpoint, g_0, lies in a grid square of size 2^i (within strata $S^{(0)} = S_{2^i}$). As we traverse γ, starting at g_0, we either get to the ending point of γ without ever leaving the initial strata, $S^{(0)}$, or we enter a second, neighboring strata, $S^{(1)}$ (which is either $S_{2^{i-1}}$ or $S_{2^{i+1}}$). If we stay within a single strata, then our result follows immediately by standard packing arguments (a curve of length ℓ can intersect only $O(|\ell|)$ unit disks). Thus, assume we change strata. Then, γ stays within these two strata for some time, until we either reach the endpoint (in which case we are again done, since we have encountered only two different sizes of squares), or we enter a new strata, $S^{(3)}$. We mark this first entry point, g_3 along γ. Then, γ continues within the two strata $S^{(2)}$ and $S^{(3)}$, possibly zig-zagging back and forth, until it finally reaches a new strata $S^{(4)}$ (adjacent to $S^{(2)}$ or to $S^{(3)}$) for the first time at point g_4, etc. The standard packing argument applies to the portion of γ from g_3 to g_4, since the curve lies within squares of only two different sizes (one being twice the size of the other). Further, we know that γ had to traverse a distance of at least the size of one of the squares in the smaller strata. Thus, the sum of the sizes of the squares traversed by γ between point g_j and g_{j+1} is at most $O(|\gamma_{j,j+1}|)$ (the length of the subcurve), *plus* the sizes of the squares at g_j and at g_{j+1}, which we know is also $O(|\gamma_{j,j+1}|)$, since $\gamma_{j,j+1}$ had to cross one of the two strata. Thus, summing along γ, we get the claimed bound. □

LEMMA 5. *Let F^* be a minimum-length forest whose addition to ∂H spans the regions $\mathcal{R}_H$ inside histogram H, with $F^* \cup \partial H$ being connected. Then $|F^*|$ has length at least a* *constant times the sum of the sizes of the set S_{F^*} of stratified grid squares intersected by F^*.*

PROOF. The forest F^* can be traversed by a set of paths, one path walking around each tree (connected component of F^*), starting and ending at the point where the tree connects to ∂H. We then apply Lemma 4 to each such path. □

Our histogram problem is now reduced to the following minimum-weight cover problem: Find a minimum-weight set $\mathcal{S}$ of grid squares within the stratified grid of H such that each region of $\mathcal{R}_H$ is intersected by at least one square of $\mathcal{S}$. (The weight of the set of squares is the sum of the sizes of the squares in the set.) For a set $\mathcal{S}$ of squares covering $\mathcal{R}_H$, we can construct a network, F, of length proportional to the weight of $\mathcal{S}$, that visits all of the regions $\mathcal{R}_H$ and is connected to ∂H: The network consists of the boundaries of the squares in $\mathcal{S}$, together with the mullions of these squares, and together with a segment joining each square σ to ∂H (each segment is of length $O(|\sigma|)$).

The next subsection gives a constant-factor approximation algorithm to the minimum-weight cover problem.

4.3 The Minimum-Weight Cover Problem

We begin by defining a notion of "maximal" squares in the stratified grid. Consider the following embedded planar dual graph, $\mathcal{G}'$, of the stratified grid. The nodes of $\mathcal{G}'$ are the centerpoints of the grid squares; two nodes are joined by an edge (straight line segment) if the corresponding squares are edge-adjacent. (By construction, two edge-adjacent squares are either of the same size, or one square is twice the size of the other.) We augment $\mathcal{G}'$ with a node s_0, corresponding to ∂H, and a set of edges, each of length $1/2$, linking s_0 to each of the centerpoints of the 1-squares that are (edge) adjacent to the boundary, ∂H. Now consider the set of *all* shortest paths from s_0 to nodes of $\mathcal{G}'$. Those squares of the grid whose centerpoints are *not* internal to *any* shortest path from s_0 are said to be *maximal squares*; they are, in a sense, "locally furthest" from the boundary, ∂H. (An alternative definition that works for our purposes is to define a maximal square to be a square σ such that it and any of its neighboring grid squares of size $|\sigma|$ are not adjacent to any grid squares of size $2|\sigma|$.) We let $\mathcal{M}_H$ denote the set of maximal squares in H.

Consider the set of all maximal shortest paths from s_0 to the centerpoints of maximal squares. There is a subset, $\mathcal{P}$, of noncrossing maximal shortest paths, with each non-maximal square σ lying internal to at least one path in $\mathcal{P}$. Such a set $\mathcal{P}$ of noncrossing paths can be constructed incrementally, starting with any one path, π_0, and then adding maximal shortest paths one by one, maintaining the non-crossing property; for any non-maximal square σ, there is always a maximal shortest path through it that does not cross the paths already in the set $\mathcal{P}$. The paths of $\mathcal{P}$ have a natural ordering, starting at π_0 and proceeding counter-clockwise around the boundary of H. We say that $\pi \prec \pi'$, for $\pi, \pi' \in \mathcal{P}$, if π attaches to ∂H at a point between the attachment points of π_0 and π' (going counterclockwise around ∂H from π_0 to π') and, if π and π' share the same attachment point on ∂H, then at the point along the paths where the paths first diverge, π turns to the left from π'.

The following lemma justifies that, in solving the minimum-weight cover problem, we can afford to add to F the boundaries (and mullions) of the maximal squares of the stratified grid, since their length is $O(|\partial H|)$, and we know (Lemma 2) that the total length of the histogram decomposition, G, is $\sum_H |\partial H| = O(L^*)$.

LEMMA 6. *The sum of the sizes of all maximal squares within H is $O(|\partial H|)$.*

PROOF. This follows from a charging argument whose details are deferred to the full paper. We assume that H is a vertical histogram with a horizontal base across its top, as in Figure 4. The boundary of each maximal square is charged off to the boundary of H (such that no portion of ∂H is charged more than $O(1)$ times). We argue that each maximal square σ is either "horizontally chargeable", meaning that there is length $\Omega(|\sigma|)$ of ∂H that lies in the horizontal strip of height $2|\sigma|$ and width $O(|\sigma|)$ whose top edge is aligned with the top of σ, or "vertically" chargeable to the boundary of ∂H below σ or to the horizontal chords corresponding to horizontally charged squares. □

In addition to the boundaries of maximal squares (and their mullions), we add to our network F a shortest path (the *spoke* associated with σ) in $\mathcal{G}'$ linking each maximal square σ to ∂H. Since each square σ is within distance $O(|\sigma|)$ of ∂H, the total length of the spokes is also $O(|\partial H|)$, so we can afford it. We can assume that the spokes are noncrossing, since any pair of crossing shortest paths can be uncrossed.

Each grid square, $\sigma \in S_i$, of the stratified grid has an associated shortest path, $\pi_\sigma \in \mathcal{P}$ that passes through its center. We let L_σ denote the subpath of π_σ from s_0 to σ, and refer to L_σ as the *stem* of σ. The stems connect each non-maximal square to the boundary ∂H. We know that $|L_\sigma| = O(|\sigma|)$. Let M_σ denote the maximal square at the end of path π_σ, and let U_σ denote the subpath of π_σ joining σ to M_σ. (Thus, π_σ is the concatenation of L_σ and U_σ.)

The portion of H that is not covered by maximal squares, $H \setminus \mathcal{M}$, is partitioned by the spokes into *zones*. It suffices to solve our minimum-weight cover problem in each zone Z independently. For a given zone Z, our goal is to find a minimum-weight *covering* set of grid squares within Z that is covering in the sense that all regions P interior to Z are intersected by the square (and therefore by its boundary or mullions) or its stem that joins it to ∂H; the weight is the sum of the square sizes, which is proportional to the sum of the perimeters, lengths of mullions, and lengths of stems.

For this optimization problem, we use a dynamic programming algorithm. Our approach is similar to that of Katz et al. [13], who study the following covering problem: Given a set of downward-pointing rays and a set of disjoint line segments in the plane, each of which is stabbed by at least one ray, find a minimum-cardinality set of rays that stab all line segments. (In fact, our results here generalize and extend the results of [13] to the problem of covering a set of disjoint *polygons* with a set of downwards-point rays.) The stemmed squares (squares with their associated stems), serve as the set of "rays" that are "downward" in the sense that the stems point to points that get closer and closer to ∂H.

We need a bit more notation. Let $\mathcal{R}_\sigma$ denote the set of regions intersected by $\sigma \cup L_\sigma$. Let r_σ be the "rightmost" point of $\partial\sigma \cup \mathcal{R}_\sigma$, and let P_σ denote the region (or σ) containing r_σ, where we say that $p \in Z$ is to the *right* of $q \in Z$ if the non-maximal square of Z containing p (resp., q) lies on a path

$\pi_p \in \mathcal{P}$ (resp., $\pi_q \in \mathcal{P}$) with $\pi_q \prec \pi_p$ in the counterclockwise ordering of paths $\mathcal{P}$. (Possibly p and q lie within the same non-maximal square, in which case the points are considered to be "equal" in this ordering, since then $\pi_p = \pi_q$.) We let A_σ denote the subpath joining the non-maximal square containing r_σ to its terminus (a maximal square, one of the two that define the zone Z).

We define a partial order on stemmed squares within Z. Consider two distinct squares, σ and σ'. We say that σ *dominates* σ' if $\mathcal{R}_\sigma$ intersects $A_{\sigma'}$ and σ' lies to the left of the path, ν_σ, from ∂H to P_σ (along L_σ), then along the boundary of the region that contains r_σ, to r_σ, then to the boundary of Z along A_σ. (We think of ν_σ as the "vertical line" through σ.) The partial order is well defined. It is easy to verify transitivity. Also, if σ dominates σ', then σ' does not dominate σ. (This follows from the fact that if $\mathcal{R}_\sigma$ intersects $A_{\sigma'}$, then $\mathcal{R}_{\sigma'}$ does not intersect A_σ: If $P \in \mathcal{R}_\sigma$ intersects $A_{\sigma'}$, then the rightmost point of P must be to the right of the point $r_{\sigma'}$ that defines $A_{\sigma'}$, and we know that $r_{\sigma'}$ is rightmost among regions $\mathcal{R}_{\sigma'}$.)

If σ is undominated among squares in Z, then no region $P \in \mathcal{R}_H$ crosses A_σ, since any such region P has a rightmost point to the right of r_σ, and there is some other grid square σ' for which $P \in \mathcal{R}_{\sigma'}$ (e.g., the square σ' that intersects P works).

The above observation allows us to write a dynamic program as follows. The set of stemmed squares within Z are ordered, $(\sigma_1, \ldots, \sigma_{K-1}, \sigma_K)$, according to ruling given by the noncrossing paths $\mathcal{P}$ within Z. (Since the paths are noncrossing, they are ordered about ∂H.) Additionally, we let σ_0 (resp., σ_{K+1}) denote the spoke that defines the "left" (resp., "right") side of Z.

A *subproblem* within zone Z is specified by giving two stemmed squares, (σ_i, σ_j), with the property that no regions intersect A_{σ_i} or A_{σ_j}. The subproblem asks us to find a minimum-weight set of stemmed squares in the subzone between ν_{σ_i} and ν_{σ_j} that stab all of the regions P that lie *fully* between ν_{σ_i} and ν_{σ_j}.

The important property about the subproblem is that we are able to specify *exactly which regions are the responsibility of the subproblem* – there are no regions that "sneak across" the boundary and would need to be specified, potentially causing the state space to explode. (This is the typical difficulty with recursive methods of trying to solve TSPN; see [18].) This is where we critically use the disjointness of the regions: No regions can cross the left boundary of the subproblem, ν_{σ_i}, along a portion of the curve ν_{σ_i} that lies inside a region. (And we know separately, from undominance, that no region crosses A_{σ_i}, and any region crossing the stem of σ_i is already visited by the stem, so we need not worry further about them in the subproblem.)

Let $f(i,j)$ be the optimal value associated with the subproblem. Then, we get the dynamic programming recursion: $f(i,j) = \min_{k \in K(i,j)}(weight(\sigma_k) + f(i,k) + f(k,j))$, where $K(i,j)$ is the set of all $k \in (i,j)$ such that σ_k is undominated among stemmed squares that lie between ν_{σ_i} and ν_{σ_j}.

Correctness of the dynamic program follows from the observation that there *exists* an undominated stemmed square in any optimal solution, and the principle of optimality. The output is a set of stemmed squares that have a minimum total weight among sets of stemmed squares that cover all regions interior to the zone.

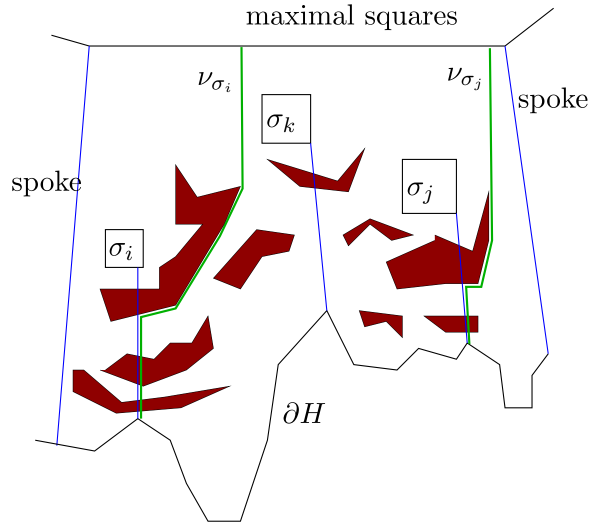

Figure 5: Illustration of the dynamic program. (The diagram is not intended to be an accurate depiction of the histogram, grid squares, etc.)

5. THE MAIN THEOREM

THEOREM 1. *There is a polynomial-time $O(1)$-approximation algorithm for the TSP with neighborhoods in the plane for a set $\mathcal{R}$ of disjoint, connected regions.*

PROOF. Our algorithm computes a tree $T_{\mathcal{E}_0}$ visiting a disjoint subset, $\mathcal{E}_0$, of the MEB's of the regions. It does so with a known algorithm ([18]) that yields an approximately optimal tree/tour that visits $\mathcal{E}_0$. We know that an optimal tree/tour T^* must visit all of $\mathcal{R}$ and, therefore, must visit all of the regions $\mathcal{E}_0$ (each of which contains a region of $\mathcal{R}$). Thus, the tree $T_{\mathcal{E}_0}$ that we produce has length $O(L^*)$. Then, using the fact that the histogram decomposition of $T_{\mathcal{E}_0}$ has length within a constant factor of the length of $T_{\mathcal{E}_0}$ (see the first part of the proof of Lemma 2 and reference [14]), the network G that we compute from $T_{\mathcal{E}_0}$ is also of length $O(L^*)$.

The output of our algorithm is the network G, augmented by a set of networks F, one per face H of G. We know that T^* restricted to any one face H of G is a forest of trees, each of which is connected to ∂H, spanning the regions that lie interior to H. Since we compute a set of networks, one per face H (spanning the regions within H, and connecting to ∂H), each of which is within a constant factor of optimal for the subproblem H, we know that the total length of all forests we compute is $O(L^*)$.

Since our solution that is output is the union of the edge set of G and the set of computed networks F for each face H of G, we conclude that the network we compute, which spans $\mathcal{R}$, is of length $O(L^*)$. □

The above result for the case of disjoint regions relied on disjointness in solving, via dynamic programming within a zone Z, the weighted set cover problem to which we reduced the problem. If the regions are allowed to overlap arbitrarily, but are *convex*, dynamic programming applies as well, with a minor modification. In particular, for a zone Z, with associated maximal square σ defining its "ceiling", we can solve the subproblem (σ_i, σ_j) between stemmed squares σ_i and σ_j easily in the case of convex regions. For a stemmed square σ_k within Z, let λ_k denote a minimum-length segment joining σ_k with σ; thus, λ_k is a line segment orthogonal to the ceiling of Z, joining the ceiling to the side of σ_k that is parallel to the ceiling. The main issue we face in trying to solve the weighted set cover problem with dynamic programming is that there may be regions within Z that cross "above" a stemmed square σ_k, passing through λ_k without intersecting σ_k, going between one subproblem and another; we cannot afford to keep track, for every such region, whether it is to be covered on one side or the other of σ_k, the boundary between, say, subproblem (σ_i, σ_k) and subproblem (σ_k, σ_j). However, if we assume that regions are *convex*, and that σ_k a "leftmost" stemmed square (closest to σ_i in terms of where the stems meet the boundary of the histogram) within subproblem (σ_i, σ_j) that minimizes its distance from the ceiling of Z (i.e., that is "highest" in the subproblem, minimizing the length of λ_k), then we can know definitively that any convex region P that is disjoint from σ_k and crosses λ_k (passing "above" σ_k) must be covered by a stemmed square on the "right" of σ_k, i.e., within the subproblem (σ_k, σ_j). This allows the dynamic program to solve the weighted set cover problem, recursively making the best choice of σ_k (the "leftmost" among the "highest" stemmed squares) to partition the subproblem (σ_i, σ_j). Because of convexity, it is well defined exactly which of the regions that intersect the subzone of Z between σ_i and σ_j, not intersecting stemmed squares σ_i or σ_j, are the responsibility of the subproblem (σ_i, σ_j); those that cross λ_j are not the responsibility of subproblem (σ_i, σ_j), while those that cross λ_i (and do not cross λ_j) are the responsibility of subproblem (σ_i, σ_j). Thus, the dynamic program solves the weighted set cover problem in polynomial time, and we have the following result.

THEOREM 2. *There is a polynomial-time $O(1)$-approximation algorithm for the TSP with neighborhoods in the plane for a set $\mathcal{R}$ of (not necessarily disjoint) convex regions.*

6. CONCLUSION

We have given the first constant-factor approximation algorithm for TSPN in the plane for a set of arbitrary disjoint connected regions. We have not yet attempted to optimize (or even to compute exactly) the constant factor in our approximation bounds for our algorithm; however, the factor is not particularly large, and we conjecture that with some care, it can be reduced to $(2 + \epsilon)$.

There are two main directions we are pursuing for future work on TSPN in the plane: (1) In the case of disjoint regions, it is not known if a PTAS may exist; (2) In the case of overlapping nonconvex regions of arbitrary depth, it is not known if a constant-factor approximation may exist. Our results point to a possible approach: It suffices, using our results, to obtain a constant-factor approximation to visit a set of *fat* (nonconvex) overlapping regions (the four directional hulls of the blue regions) that have the property that each region lies within a constant factor of its diameter to the boundary of a containing histogram (Proposition 1). (Using a modification of an argument in [10], we can show that the problem of visiting a set of overlapping fat convex regions is APX-hard, so we do not expect to obtain a PTAS.) Our dynamic programming algorithm exploits disjointness (or convexity), though did not require fatness.

Additionally, we are pursuing approximation algorithms for TSPN in higher dimensions for general regions (not necessarily fat, as required in [4]).

7. 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. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[3] S. Arora. Approximation schemes for NP-hard geometric optimization problems: A survey. *Mathematical Programming*, 97(1-2):43–69, 2003.

[4] T.-H. H. Chan and K. Elbassioni. A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics. In *Proc.21st ACM-SIAM Symposium on Discrete Algorithms*, pages 256–267, 2010.

[5] M. de Berg, J. Gudmundsson, M. J. Katz, C. Levcopoulos, M. H. Overmars, and A. F. van der Stappen. TSP with neighborhoods of varying size. *Journal of Algorithms*, 57:22–36, 2005.

[6] M. Dror and J. B. Orlin. Combinatorial optimization with explicit delineation of the ground set by a collection of subsets. Technical report, MIT, 2004.

[7] A. Dumitrescu and J. S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *Journal of Algorithms*, 48:135–159, 2003.

[8] K. Elbassioni, A. V. Fishkin, N. H. Mustafa, and R. Sitters. Approximation algorithms for Euclidean group TSP. In *Proc. 32nd Internat. Colloq. Automata Lang. Prog.*, volume 3580 of *Lecture Notes Comput. Sci.*, pages 1115–1126. Springer-Verlag, 2005.

[9] K. Elbassioni, A. V. Fishkin, and R. Sitters. On approximating the TSP with intersecting neighborhoods. In *Proc. 17th Internat. Sympos. Algorithms Comput.*, volume 4288 of *Lecture Notes Comput. Sci.*, pages 213–222. Springer-Verlag, 2006.

[10] K. M. Elbassioni, A. V. Fishkin, and R. Sitters. Approximation algorithms for the Euclidean traveling salesman problem with discrete and continuous neighborhoods. *Int. J. Comput. Geometry Appl.*, 19(2):173–193, 2009.

[11] C. Feremans and A. Grigoriev. Approximation schemes for the generalized geometric problems with geographic clustering. In *Abstracts 21st European Workshop Comput. Geom.*, pages 101–102, 2005.

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

[13] M. Katz, J. S. B. Mitchell, and Y. Nir. Orthogonal segment stabbing. *Comput. Geom. Theory Appl.*, 30(2):197–205, 2005.

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

[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 J. Comput.*, 28(4):1298–1309, 1999.

[16] 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.

[17] J. S. B. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry (2nd Edition)*, chapter 27, pages 607–641. Chapman & Hall/CRC, Boca Raton, FL, 2004.

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

[19] S. B. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *Proc. 30th ACM Sympos. Theory Comput.*, pages 540–550, 1998.

[20] S. Safra and O. Schwartz. On the complexity of approximating TSP with neighborhoods and related problems. In *Proc. 11th Annu. European Sympos. Algorithms*, volume 2832 of *Lecture Notes Comput. Sci.*, pages 446–458. Springer-Verlag, 2003.

[21] S. Suri. A linear time algorithm for minimum link paths inside a simple polygon. *Comput. Vision Graph. Image Process.*, 35:99–110, 1986.