

TSP With Locational Uncertainty: The Adversarial Model^{*†}

Gui Citovsky¹, Tyler Mayer², and Joseph S. B. Mitchell³

- 1 Google Manhattan, New York, NY, USA
gcitovsky@gmail.com
- 2 Dept. of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY, USA
tyler.mayer@stonybrook.edu
- 3 Dept. of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY, USA
joseph.mitchell@stonybrook.edu

Abstract

In this paper we study a natural special case of the Traveling Salesman Problem (TSP) with point-locational-uncertainty which we will call the *adversarial TSP* problem (ATSP). Given a metric space (X, d) and a set of subsets $R = \{R_1, R_2, \dots, R_n\} : R_i \subseteq X$, the goal is to devise an ordering of the regions, σ_R , that the tour will visit such that when a single point is chosen from each region, the induced tour over those points in the ordering prescribed by σ_R is as short as possible. Unlike the classical locational-uncertainty-TSP problem, which focuses on minimizing the expected length of such a tour when the point within each region is chosen according to some probability distribution, here, we focus on the *adversarial model* in which once the choice of σ_R is announced, an adversary selects a point from each region in order to make the resulting tour as long as possible. In other words, we consider an offline problem in which the goal is to determine an ordering of the regions R that is optimal with respect to the “worst” point possible within each region being chosen by an adversary, who knows the chosen ordering. We give a 3-approximation when R is a set of arbitrary regions/sets of points in a metric space. We show how geometry leads to improved constant factor approximations when regions are parallel line segments of the same lengths, and a polynomial-time approximation scheme (PTAS) for the important special case in which R is a set of disjoint unit disks in the plane.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases traveling salesperson problem, TSP with neighborhoods, approximation algorithms, uncertainty

Digital Object Identifier 10.4230/LIPIcs.SoCG.2017.32

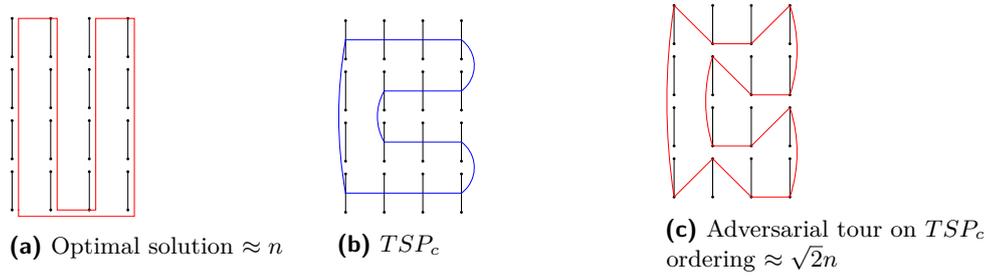
1 Introduction

We consider the travelling salesperson problem (TSP) on uncertain sites. We are given as input a set of n uncertainty regions $R = \{R_1, R_2, \dots, R_n\}$, each of which is known to contain exactly one site that must be visited by the tour. In the standard TSP, the regions R_i are singleton points. In the *TSP with neighborhoods* (TSPN), or *one-of-a-set TSP*, model, the

* A full version of the paper is available at <https://arxiv.org/abs/1705.06180>.

† This research was partially supported by the National Science Foundation (CCF-1526406).





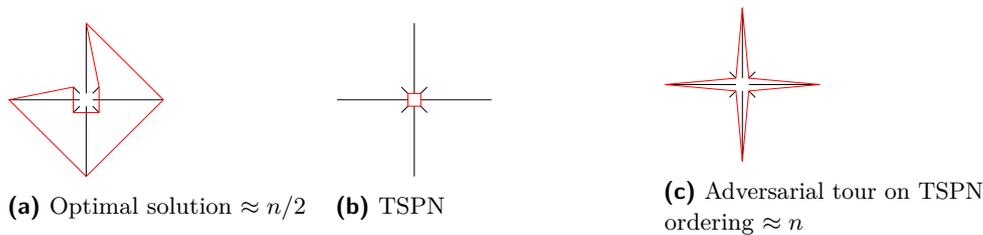
■ **Figure 1** TSP on center points ordering does not always provide an optimal solution to ATSP.

goal is to compute an optimal tour that visits some point of each region R_i , and we are allowed to pick any point $p_i \in R_i$ to visit, making this choice in the most advantageous way possible, to minimize the length of the resulting tour that we compute. In models of TSP with locational uncertainty, the regions R_i model the support sets of probability distributions for the uncertain locations of the (random variable) sites p_i . The objective, then, may be to optimize some statistic of the tour length; e.g., we may wish to minimize the expected tour length, or minimize the probability that the tour length is greater than some threshold, etc. In this paper, we study the version of the stochastic TSP model in which our goal is to optimize for the *worst case* choice of p_i within each R_i . We call this problem the *adversarial TSP*, or ATSP, as one can think of the choice of p_i within each R_i as being made by an adversary. Our goal is to compute a permutation σ_R on the regions R_i so that we minimize the length of the resulting tour on the points p_i , assuming that an adversary makes the choice of $p_i \in R_i$, given our announced permutation σ_R on the regions. While the TSPN seeks an optimal tour for the *best* choices of $p_i \in R_i$, the ATSP seeks an optimal tour for the *worst* choices of $p_i \in R_i$.

Another motivation for the ATSP solution is that one may seek a single permutation of the set of input sets R_i so that the permutation is “good” (controls the worst-case choices of $p_i \in R_i$) for any of the numerous ($|R_1| \cdot |R_2| \cdot |R_3| \cdots |R_n|$) instances of TSP associated with the sets R_i , thereby avoiding repeated computations of TSP tours. In certain vehicle routing applications, it may also be beneficial to establish a fixed ordering of visits to clients, even if the specific locations of these visits may vary in the sets R_i . Further, in locationally uncertain TSP one may expect that probability distributions over the regions R_i are imperfect and not known precisely, and that customer locations are known imprecisely (possibly for privacy concerns, with deliberate noise added for protecting the identity/privacy of users), making it important to optimize over all possible choices of site locations.

In Figure 1 we give a simple example showing that the ordering given by a TSP on center points (TSP_c) can be suboptimal, by at least a factor of $\sqrt{2}$. The input R is a $\sqrt{n} \times \sqrt{n}$ grid of vertical unit-length line segments with distance 1 between midpoints of horizontally adjacent segments and with distance $1 + \epsilon$ between midpoints of vertically adjacent segments. In Figure 2 we show that the ordering prescribed by a TSPN over the input regions can be at least a factor 2 away from optimal. The input is a set of n segments, $n/2$ of which have length 1, and the remainder have length ϵ ; they are arranged in alternating order radially around a point or the boundary of a small circle.

In this paper, we initiate the study of the ATSP. We give a 3-approximation when R is a set of arbitrary regions/sets of points in a metric space. We exploit geometry to give an improved approximation bound for the case of regions that are unit line segments of the same orientation in the plane; we compute a permutation with adversarial tour length at



■ **Figure 2** TSPN ordering does not always provide an optimal solution to ATSP.

most $(7/3 + \epsilon)|OPT| + 1$, where $|OPT|$ is the length of an optimal solution. We further exploit geometry to give a polynomial time approximation scheme (PTAS) for the important special case when R is a set of disjoint unit disks in the plane.

Related Work

Geometric problems on imprecise points have been the subject of many recent investigations. Löffler et al. [11] study, given a set of n uncertainty regions in the plane, the problem of *selecting* a single point within each region so that the area of the resulting convex hull is as large/small as possible. They show a number of results, including an $O(n^3)$ -time and an $O(n^7)$ -time exact algorithm for maximizing the area of the convex hull of selected points when the uncertainty regions are parallel line segments and disjoint axis aligned squares respectively. They show that this problem is NP-Hard when the regions are line segments with arbitrary orientations. In the same paper, Löffler et al. show that the problem of selecting a point within each region so that the resulting minimum spanning tree over those points is as small as possible is NP-Hard when the uncertainty regions are overlapping disks and squares. In his thesis [6], Fraser extends the prior minimum spanning tree result to show that the problem is still NP-Hard even when the regions are pairwise disjoint. He provides several constant factor approximation algorithms for the special case of disjoint disks in the plane. Dorrigiv et al. [4] show that neither the minimization nor the maximization version of this problem admit an FPTAS when the regions are disjoint disks. Yang et al. [16] give a PTAS for the minimization version. In a thesis by Montanari [15], it is shown that the minimization version when the input regions are vertically or horizontally aligned segments is NP-Hard and that this problem does not admit a FPTAS. Interestingly, in another paper by Liu and Montanari [10] it is shown that selecting a point from each region so that *diameter* of a minimum spanning tree on the selected points is minimized is polynomially solvable when the regions are arbitrary sized (possibly overlapping) disks in the plane.

Montanari [15] also studies the problem of placing a single point within each region so that the resulting shortest s, t path is either maximized or minimized. They show that the minimization version of this problem can be solved in polynomial time in the L_1 metric when the polygons are rectilinear (not necessarily disjoint, or convex). They also show that the maximization version of the problem is NP-Hard to approximate to any factor $(1 - \epsilon) : \epsilon < 1/4$ even in the case where the polygons are vertically aligned segments.

There has been a considerable amount of work done on studying TSP variants with point-existential uncertainty. Two main models in the literature are the *a priori model* proposed by Bertsimas et al. [2] and Jaillet [7], in which each point x_i (with a known, fixed, location) is independently present with probability p_i , and the *universal model* [8], which asks for a tour over the entire data set such that for *any* subset of active requests, the master tour restricted to this active subset is *close* to optimal.

The TSP with neighborhoods (TSPN) problem was introduced by Arkin and Hassin [1] and has been studied extensively from the perspective of approximation algorithms, particularly in geometric domains (see, e.g., [13]). Kamousi et al [9] study a stochastic TSPN model where each client lies within a region, a disk with a fixed center and stochastic radius.

Preliminaries

We are given regions $R = \{R_1, R_2, \dots, R_n\}$, with each R_i a subset of a metric space (X, d) . We seek a cyclic permutation $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ (an *ordering*) of the regions R , in order to minimize the length, $\max_{p_i \in R_i} [d(p_{\sigma_1}, p_{\sigma_2}) + d(p_{\sigma_2}, p_{\sigma_3}) + \dots + d(p_{\sigma_{n-1}}, p_{\sigma_n}) + d(p_{\sigma_n}, p_{\sigma_1})]$, of a cycle on adversarial choices of the points in the respective regions. We let σ_R^* denote an optimal ordering for R , and we let $|OPT|$ denote the length of the corresponding cycle, OPT , that is based on the optimal adversarial choices of the points $p_i \in R_i$, for the ordering σ_R^* . The following lemmas are shown in the full paper [3].

► **Lemma 1.** *The length, $|OPT|$, of OPT satisfies $TSPN^* \leq |OPT| \leq TSPN^* + \sum_{R_i \in R} 2 \cdot \text{diam}(R_i)$, where $TSPN^*$ is the length of an optimal TSPN tour on the regions R , and $\text{diam}(R_i)$ denotes the diameter of region $R_i \in R$.*

► **Lemma 2.** *For a set R of convex regions in the Euclidean plane, and any ordering σ of the regions R , any longest cycle corresponding to an adversarial choice of points $p_i \in R_i$ is a polygonal cycle, with edges $(p_{\sigma_i}, p_{\sigma_{i+1}})$ and with each point p_{σ_i} an extreme point of its corresponding region, R_{σ_i} .*

2 3-Approximation for Arbitrary Regions in a Metric Space

We begin by giving a 3-approximation to the ATSP problem when R is a set of arbitrary regions in a metric space.

Consider the complete graph \hat{G} whose nodes are the regions R and whose edges join every pair of regions with an edge, (R_i, R_j) , whose weight is defined to be $w(R_i, R_j) = \max_{s \in R_i, t \in R_j} \{d(s, t)\}$, the maximum distance between a point $s \in R_i$ and a point $t \in R_j$. For distinction, we will speak of edge “weights” in the graph \hat{G} and of edge “lengths” in the original metric space (X, d) . It is not hard to see that the edge-weighted graph \hat{G} defines a metric (see the full paper [3]).

► **Lemma 3.** *An optimal TSP tour in \hat{G} yields a 2-approximation to the ATSP on R .*

Proof. Let $\sigma_R^* = \langle R_1^*, R_2^*, \dots, R_n^* \rangle$ be an optimal (cyclic) permutation of the regions R for the adversarial TSP on R , and let $p_i^* \in R_i^*$ be the adversary’s choice of points corresponding to σ_R^* . Then, $|OPT| = d(p_1^*, p_2^*) + d(p_2^*, p_3^*) + \dots + d(p_n^*, p_1^*)$ is the length of the cycle $C = \langle p_1^*, p_2^*, \dots, p_n^* \rangle$, an optimal adversarial TSP solution.

Let $w_{\sigma_R^*} = w(R_1^*, R_2^*) + w(R_2^*, R_3^*) + \dots + w(R_n^*, R_1^*)$ be the total weight of the cycle σ_R^* in \hat{G} . Let w_{TSP}^* be the total weight of a minimum-weight Hamiltonian cycle, given by (cycle) permutation σ_{TSP} , in \hat{G} ; then, $w_{TSP}^* \leq w_{\sigma_R^*}$.

Our goal is to show that the permutation σ_{TSP} yields a 2-approximation for the adversarial TSP on R . Since the length of the adversarial cycle corresponding to σ_{TSP} is at most w_{TSP} , and since $w_{TSP}^* \leq w_{\sigma_R^*}$, it suffices to show that $w_{\sigma_R^*} \leq 2|OPT|$.

Consider the cycle $C = \langle p_1^*, p_2^*, \dots, p_n^* \rangle$ whose length is $|OPT|$. If we modify C by choosing points within each region R_i^* differently from $p_i^* \in R_i^*$, the length of C can only go down, since the points p_i^* were chosen adversarially to make the cycle C as long as possible (for the given permutation σ_R^*). Consider two copies of C (of total length $2|OPT|$); we will

modify these two cycles into two (possibly shorter) cycles, C_1 and C_2 , by making different choices for the points in each region R_i^* .

Consider first the case that n is even. Then, we define C_1 to be the modification of cycle C in which the points are chosen in regions R_i^* in order to maximize the lengths of the “odd” edges, $(R_1^*, R_2^*), (R_3^*, R_4^*), \dots, (R_{n-1}^*, R_n^*)$, and we define C_2 to be the modification of cycle C in which the points are chosen in regions R_i^* in order to maximize the lengths of the “even” edges, $(R_2^*, R_3^*), (R_4^*, R_5^*), \dots, (R_n^*, R_1^*)$. The cycle C_1 , then, has length at least $w(R_1^*, R_2^*) + w(R_3^*, R_4^*) + \dots + w(R_{n-1}^*, R_n^*)$, the total weights of the odd edges in the cycle in \hat{G} corresponding to σ_R^* . Similarly, the cycle C_2 has length at least $w(R_2^*, R_3^*) + w(R_4^*, R_5^*) + \dots + w(R_n^*, R_1^*)$, the total weights of the even edges in the cycle in \hat{G} corresponding to σ_R^* . Together, then, the lengths of the two cycles C_1 and C_2 total at least the weight, $w_{\sigma_R^*}$, of the cycle σ_R^* in the graph \hat{G} . Since each of the weights of C_1 and C_2 are at most $|OPT|$ (the weight of C), we conclude that $w_{\sigma_R^*} \leq 2|OPT|$, as claimed.

The case in which n is odd is handled similarly; details appear in the full paper [3]. ◀

► **Theorem 4.** *The permutation σ_R corresponding to a Christofides 3/2-approximate TSP tour in \hat{G} yields a 3-approximation to the adversarial TSP on R .*

3 Unit Line Segments of the Same Orientation in the Plane

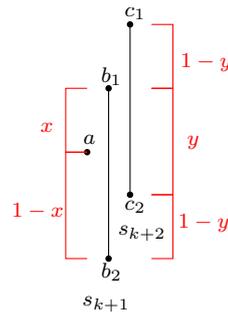
In this section, we assume that R consists of a set of n unit-length segments of the same orientation; without loss of generality, we assume the segments are vertical. We show that the ordering, TSP_c , given by an optimal TSP tour on the segment center points yields an adversarial tour of length at most $(7/3)|OPT| + 1$; thus, a PTAS to approximate TSP_c yields an algorithm with adversarial tour length at most $(7/3 + \epsilon)|OPT| + 1$, for any fixed ϵ .

► **Lemma 5.** *For the ATSP on a set R of unit vertical segments in the plane, $|OPT| \geq TSP_c^*$, where TSP_c^* is the length of an optimal TSP tour on the segment center points.*

Proof. Consider an ATSP optimal ordering σ_R^* of the vertical segments R . The cycle γ_c that visits the center points of segments R in the order σ_R^* has length at least TSP_c^* , and the length, $|OPT|$, of an adversarial cycle for σ_R^* is at least the length of γ_c . ◀

► **Lemma 6.** *$|OPT| \geq \frac{3}{4}(n - 1)$ when n is odd and $|OPT| \geq \frac{3}{4}n$ when n is even.*

Proof. It suffices to show the claim for ATSP paths; an ATSP cycle is at least as long. The proof is by induction on $n = |R|$. First, suppose that n is odd. The base case is trivially true. Assume that the claim holds for $n \leq k$, for k odd. Next, consider an instance S' with $k + 2$ segments. We know that for the first k segments in an optimal permutation for S' , that $|OPT| \geq \frac{3}{4}(k - 1)$. Next, we show that regardless of the placement of the next two unit segments in S' , s_{k+1} and s_{k+2} , an adversary can make us pay at least 3/2 units for every independent pair of consecutive segments in σ_S^* . We can assume that (vertical) segments s_{k+1} and s_{k+2} are vertically collinear. Next we assume, without loss of generality, that s_{k+2} is above s_{k+1} . Let a be the point on s_k that the adversary chose; refer to Figure 3. Let b_1 (resp., c_1) be the top endpoint of s_{k+1} (resp., s_{k+2}). Let b_2 (resp. c_2) be the bottom endpoint of s_{k+1} (resp., s_{k+2}). Now, let $|b_1a| = x$ and $|c_2b_1| = y$. This implies that $|ab_2| = 1 - x$, $|c_2b_2| = 1 - y$ and $|c_1b_1| = 1 - y$. The three candidate routes for the adversary to take are (a, b_2, c_1) or (a, b_1, c_2) or (a, b_1, c_1) . These paths have lengths $3 - x - y$, $x + y$, $x + 1 - y$, respectively. Thus, we solve $\min - \max_{x,y} \{ \{3 - x - y, x + y, x + 1 - y\} : 0 \leq x \leq 1, 0 \leq y \leq 1 \}$ to find the minimum possible length of the adversarial route; the solution is 3/2. Thus, the adversary can make us pay 3/2 for each pair of segments; thus, $|OPT| \geq \frac{3}{4}(n - 1)$ for n odd.



■ **Figure 3** Illustration of the induction step in the proof of Lemma 6.

In the case that n is even, the adversary can make us pay at least $3/2$ between every consecutive pair of segments in the optimal ordering; thus, $|OPT| \geq \frac{3}{4}n$. ◀

► **Theorem 7.** *For the ATSP on a set R of unit-length vertical segments, the ordering given by an optimal TSP on the segment center points yields an adversarial tour of length at most $(7/3)|OPT| + 1$. Thus, a PTAS for the TSP on center points yields an approximation algorithm for ATSP, with tour length at most $(7/3 + \epsilon)|OPT| + 1$.*

Proof. Let APX_c be the ordering in which the segments are visited by a $(1 + \epsilon)$ -approximate TSP tour on their center points, and let $|APX_c|$ be the cost of the resulting adversarial tour for this ordering. We know that $|APX_c| \leq |TSP_c| + n$, where $|TSP_c|$ is the length of an optimal TSP on center points, since a tour on the center points can be made to detour to either endpoint and back, for each segment, at a total increase in length of n . Since $|TSP_c| \leq |OPT|$ (by Lemma 5) and $n \leq 4/3|OPT| + 1$ (by Lemma 6), we have that $|APX| \leq (7/3 + \epsilon)|OPT| + 1$. ◀

4 PTAS for Disjoint Unit Disks in the Plane

In this section we give a PTAS for the adversarial TSP problem when the regions $R = \{d_1, \dots, d_n\}$ are n disjoint unit-diameter disks in the plane. We employ the m -guillotine method [12], which has been applied to give approximation schemes for a wide variety of geometric network optimization problems, including the Euclidean TSP and the *TSP with Neighborhoods* (TSPN) when the regions are disjoint disks or fat regions in the plane [5, 14].

The challenge in applying known PTAS techniques is being able to handle the adversarial nature of the tour. For the TSPN problem, one computes (using dynamic programming) a shortest connected m -guillotine, Eulerian, spanning subgraph of the regions; a tour visiting each region can then be extracted from this network. A structure lemma shows that an optimal TSPN solution can be converted to an m -guillotine solution whose weight is at most $(1 + \epsilon)|OPT|$. Since m -guillotine networks have a recursive structure, we can apply dynamic programming in order to find the cheapest such structure over the input. Then, by extracting a tour from the optimal m -guillotine network, we obtain a permutation of the input disks, as well as a *particular* point within each region that the tour visits.

For the ATSP problem, we require new ideas and a new structure theorem to account for the fact that our algorithm must search for a permutation of the input disks that is good with respect to an adversarial path through the ordered disks. We seek to optimize a network that has a recursive structure (to allow dynamic programming to be applied) and that yields an ordering of the disks so that the length of the adversary’s tour is “very

close” to optimal among all possible permutations. We do this by searching for a shortest (embedded) network having an m -guillotine structure that has additional properties that guarantee that the adversary’s path through the sequence of regions we compute is not much longer than that of the network we compute. To accomplish this, we will require several structural results about an optimal solution to ATSP.

4.1 Discretization and a Structural Theorem

In order to make our problem and our algorithm discrete, for a fixed integer $m = O(1/\epsilon)$, we place m *sample points* evenly spaced around the boundaries of each of the n disks $d_i \in R$. Let \mathcal{G} be the set of all nm sample points. Let $E_{\mathcal{G}}$ denote the set of edges (line segments) between two sample points of \mathcal{G} that lie on the boundaries of different disks of R . The following lemma shows that for any adversarial (polygonal) tour T associated with σ_R there is a polygonal tour T' visiting the sequence σ_R whose vertices are among the sample points \mathcal{G} and whose length is at least $(1 - O(1/m))|T|$.

► **Lemma 8.** *Given an adversarial (polygonal) path/cycle, T , associated with a sequence σ_R of input disks, there is a polygonal path/cycle T' that visits sample points \mathcal{G} , exactly one per disk, in the order σ_R , such that $|T| \leq (1 + O(1/m))|T'|$.*

Proof. We let T' be the path/cycle obtained from T by rounding each of its vertices to the closest sample point of the associated (unit-diameter) disk. This rounding results in each edge decreasing in length by at most $2 \cdot \frac{\pi}{m}$, since the sample points are spaced on the disk boundary at distance (along the boundary) of $\frac{2\pi(1/2)}{m}$. Thus, $|T| \leq |T'| + \frac{2\pi}{m}n$. We obtain a lower bound on $|T'|$, in terms of n , using an area argument (as done in [5], but included here for completeness). Let $A(T')$ be the area swept by a disk of radius 1 whose center traverses T' ; it is well known that the area swept by a disk of radius δ whose center moves on a curve of length λ is at most $2\delta\lambda + \pi\delta^2$, implying that $A(T') \leq 2|T'| + \pi$. Since T' meets all n of the unit-diameter disks d_i , we know that $A(T') \geq n \cdot \pi(1/2)^2$. Thus, $n \leq (8/\pi)|T'| + 4 \leq O(|T'|)$ (assuming that $|T'| \geq c$, for some constant c , which holds if $n \geq 2$). Since $n \leq O(|T'|)$, the inequality $|T| \leq |T'| + \frac{2\pi}{m}n$ implies that $|T| \leq (1 + O(1/m))|T'|$, as desired. ◀

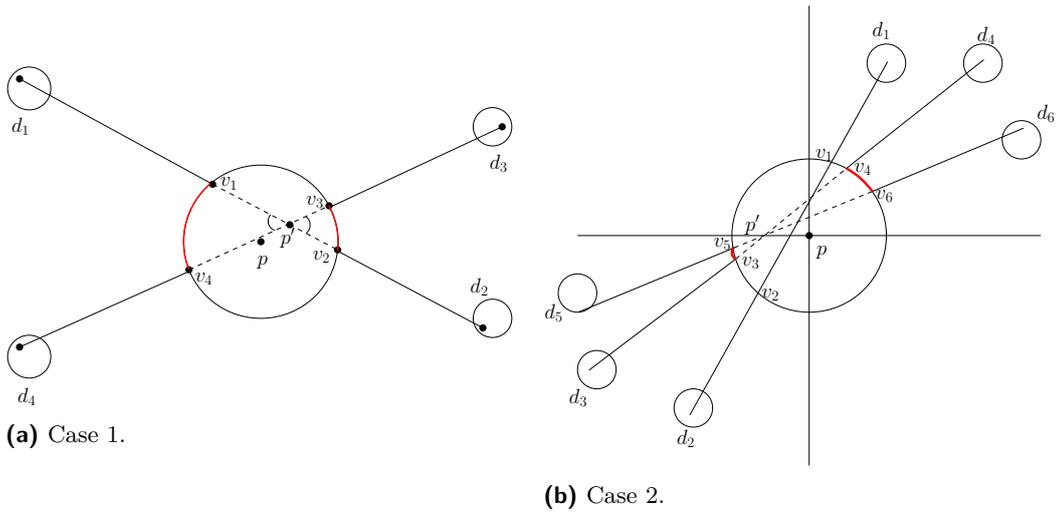
A corollary of Lemma 8 is that, for purposes of obtaining a PTAS, it suffices to search for an optimal adversarial tour in the discrete graph of edges $E_{\mathcal{G}}$ on sample points.

For two consecutive disks, d_i and d_{i+1} in an ordering σ_R , we refer to the convex hull of d_i and d_{i+1} as the *fat edge* associated with (d_i, d_{i+1}) . The collection of such fat edges will be called the *convex hull tour* associated with σ_R .

► **Theorem 9.** *No point $p \in \mathbb{R}^2$ in the plane lies within more than a constant number of fat edges of the convex hull tour, OPT , associated with an optimal ordering σ_R^* .*

Proof. Consider an arbitrary point $p \in \mathbb{R}^2$ and consider its intersection with the convex hull tour of OPT . Center a disk D_p , of radius K centered at p , with $K = O(1)$ a constant to be determined later. Since the disks d_i are disjoint, there are only a constant number ($O(K^2)$) that intersect D_p . We remove from R these disks, as well as the (at most two) disks adjacent to them in the tour OPT . Let R' be the remaining set of disks after these (constant number of) disks are removed from R .

We claim that p is contained in no more than a constant number of the fat edges of OPT joining two disks of R' . Assume to the contrary that more than a constant number of remaining fat edges of the convex hull tour of OPT connecting disks of R' contain p . Consider two such fat edges, (d_1, d_2) and (d_3, d_4) , containing p in the region where they



■ **Figure 4** Case analysis for fat edge swapping.

properly cross. Each of these fat edges must pass “nearly” diametrically across D_p . That is they must cross D_p in such a way that they contain its center point p . We will show that by uncrossing these two fat edges we obtain a strictly shorter adversarial tour, thereby contradicting the optimality assumption. Suppose, without loss of generality, that, in order to preserve connectivity, the uncrossing replaces (d_1, d_2) and (d_3, d_4) with (d_1, d_4) and (d_3, d_2) . Let v_i be the point of intersection closest to d_i where the adversarial edge incident on d_i crosses the boundary of D_p . There are two cases.

Case 1: First suppose that $\angle v_1 p' v_4 = \angle v_3 p' v_2 \leq \pi/2$, where p' is the point where the adversarial edges correspond to (d_1, d_2) and (d_3, d_4) cross; refer to Figure 4a. Note that we could delete the portions of adversarial edges (v_1, v_2) , and (v_3, v_4) crossing D_p and replace these with the two portions of the circumference of D_p connecting points v_1, v_4 and v_2, v_3 (see Figure 4a). In deleting the portions of the adversarial edges which intersect the interior of D_p , we saved at least $4\sqrt{K^2 - 1}$. This value comes from the fact that the adversarial edge is contained within the fat edge connecting these two disks, which needs only “nearly” pass diametrically across D_p ; it could be the case that p is contained within a fat edge on its boundary. In replacing the deleted portions of adversarial edges with two arcs comprising at most half of the circumference of D_p (with arc length at most πK) we still get an overall savings of at least $4\sqrt{K^2 - 1} - \pi K$. Thus, we need to choose K so that $4\sqrt{K^2 - 1} - \pi K \geq 9$ implying $K \geq 11 = O(1)$. We will show later that this savings of 9 units of tour length is more than enough to compensate for the adversarial increase in the new proposed ordering.

Case 2: Next, suppose that $\angle v_1 p' v_4 = \angle v_3 p' v_2 > \pi/2$ for any pair of fat edges still containing p . We will begin by breaking the plane into quadrants whose origin is p and now consider triples of fat edges that contain p . We will only consider those triples of fat edges whose disk endpoints lie in quadrants I, and III, as we can repeat this process a finite number of times, each with a new perturbed (rotated) set of quadrants whose origin is p so that eventually all remaining fat edges containing p have this property.

Let (d_1, d_2) be some remaining fat edge containing p whose disk endpoints are in quadrants I, and III. Let (d_3, d_4) , (d_5, d_6) be the second and third fat edge respectively that contain p ,

and have an endpoint in each of quadrants I and III, found in order by walking along the optimal tour from d_2 away from d_1 . As in case 1, let v_i be the point of intersection of the adversarial edge emanating from disk d_i and the boundary of D_p . We have that all of the v_i are in quadrants I or III as well (see Figure 4b). Given that v_1, v_2 are in opposite quadrants, as well as points v_3, v_4 , and v_5, v_6 , a simple case analysis will show that we can delete two edges (v_i, v_{i+1}) , (v_j, v_{j+1}) that cross the interior of D_p , and replace them with two arcs of D_p , lying strictly within quadrants I and III, which make up at most half the circumference of D_p , while preserving connectivity of the tour. This case analysis is independent of the specifics of which quadrant contains disk d_i , and only requires that each triple of edges we try to uncross go between opposite quadrants.

Thus, as in Case 1, we can argue that in replacing two edges crossing D_p (saving at least $4\sqrt{K^2 - 1}$ in length) and replacing these with the two arcs of D_p (which comprise at most half the circumference of D_p) we have a net savings of at least $4\sqrt{K^2 - 1} - \pi K$, which is at least 9 when $K \geq 11$.

Each round of uncrossing (Case 1 or Case 2) reduces the tour length by a positive amount and reduces the depth of p by at least one. Therefore, this process will terminate in a finite number of rounds. The number of fat edges containing p remaining after the process (Case 2) terminates will be at most (another) constant.

Finally, we argue that the constant 9 we save in tour length in each local uncrossing is enough to compensate for whatever global increase in adversarial tour length may occur due to the new proposed ordering (since the adversary gets to re-optimize his selection of points).

Again, consider an uncrossing of the original, hypothesized optimal tour, replacing (d_1, d_2) and (d_3, d_4) with (d_1, d_4) and (d_2, d_3) . Let x, y, u, v be the (original) points adversarially chosen in disks d_1, d_2, d_3, d_4 . After performing the uncrossing, we get a new tour, and thus the adversary gets to re-optimize by choosing a different set of points. From the adversarial property of the initial solution, we have that the initial paths from x to u and from v to y were as long as possible over the intermediate choice of disks if we fix points x, y, u, v . The new path chosen between disks d_1 and d_3 is at most that of the original path between x and u , plus two diameters, one per disk. That is, suppose the adversary chose new points x', u' in disks d_1, d_3 respectively. We can model the new path as traveling from x' to x in d_1 following the original path from x to u and then traveling from u to u' in d_3 costing at most two diameters. Similarly for the path between v , and y . Finally, in arguing about the additional length reconnecting the tour after the swap, we can upper bound, by triangle inequality, the length of the edge (x', v') and (u', y') as at most four diameters, one per disk d_1, d_2, \dots, d_4 the portions of the edges (x, y) (u, v) strictly exterior to D_p as well as at most half the circumference of D_p . However, we have the savings of removing those portions of edges (x, y) and (u, v) that were strictly interior to D_p . Recall that the diameter of D_p was chosen such that removing two edges that pass “nearly” diametrically across D_p and replacing them with two arcs comprising at most half of its circumference results in a net savings of 9 units. Therefore in adding at most 8 diameters (or 8 units) upper bounding the adversarial increase, we still have a net savings of at least 1 unit. Thus, we have a strictly shorter adversarial tour after performing the uncrossing, thereby contradicting the optimality assumption of the original tour. ◀

4.2 The m -Guillotine Structure Theorem

We begin with some notation largely following [5, 12]. Let G be an embedded planar straight line graph (PSLG) with edge set E of total length L , and let $R = \{d_1, \dots, d_n\}$ be a set of

disjoint unit-diameter disks d_i in the plane. (In our setting, there will be exactly one vertex of G within each disk $d_i \in R$.) Let \mathcal{B} be an axis-aligned bounding square of R . We refer to an axis-aligned box $W \subset \mathcal{B}$ as a *window*, which will correspond to a particular subproblem of our dynamic program. We refer to an axis-parallel line ℓ that intersects window W as a *cut* of window W .

Consider a cut ℓ for window W ; assume, without loss of generality, that ℓ is vertical. The intersection $\ell \cap (E \cap W)$ of ℓ with the edge set contained in W consists of a, possibly empty, set of subsegments (which include, as a degenerate case, singleton points) along ℓ . We let ξ be the number of endpoints of subsegments along ℓ , and let these endpoints along ℓ be denoted by $\beta_1, \beta_2, \dots, \beta_\xi$ ordered by decreasing y coordinate. For a positive integer m we define the m -span $\sigma_m(\ell)$ of ℓ to be \emptyset if $\xi \leq 2(m-1)$, and the possibly zero length segment $\beta_m, \beta_{\xi-m+1}$, joining the m th and the m th from the last endpoints along ℓ otherwise.

The intersection of $\ell \cap R \cap W$ consists of a possibly empty set of $\xi_R \leq |R \cap W|$ subsegments of ℓ , one subsegment for each disk (bounding box) intersected by $\ell \cap W$. Let these disk/boxes be $d_1, d_2, \dots, d_{\xi_R}$ in order of decreasing y coordinate. For a positive integer m we define the m -disk-span $\sigma_{m,R}(\ell)$ of ℓ to be the (possibly empty) line segment joining the bottom endpoint of $d_m \cap \ell$ to the top endpoint of $d_{\xi_R-m+1} \cap \ell$. In fact, as observed in [14], it suffices to consider the m -disk-span of the set of axis-aligned bounding squares of the input disks, since the charging scheme charges the perimeters of the regions, which are, within a constant factor, the same whether we deal with circular disks or square (L_∞) disks.

As in [5] we define a line (cut) ℓ to be an m -good cut with respect to W if $\sigma_m(\ell) \subseteq E$ and $\sigma_{m,R} \subseteq E$. Finally, we say that E satisfies the m -guillotine property with respect to W if 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 satisfies the m -guillotine property with respect to W_1 , and W_2 . The following is shown in [5], using a variant of the charging scheme of [12]:

► **Theorem 10** ([5]). *Let G be an embedded connected planar graph with edge set E of total length L , and let R be a given set of pairwise-disjoint equal-radius disks (of radius δ) each of which intersects E . Assume that E and R are contained in the square \mathcal{B} . Then for any positive integer m there exists a connected planar graph G' that satisfies the m -guillotine property with respect to \mathcal{B} and has edge set $E' \supseteq E$ of length $L' \leq (1 + O(1/m))L + O(\delta/m)$.*

In the constructive proof of Theorem 10, m -spans are added to E , whose lengths are charged off to a small fraction ($O(1/m)$) of the length L of E . Consider the edges of E that cross an m -span, ab that is added: By Theorem 9 we know that the associated fat edges (of width 1) have constant depth. This implies that the number of edges of E that cross an m -span, ab , that arises in the constructive proof of Theorem 10 is $O(|ab|)$.

► **Theorem 11.** *In the graph G' that is obtained from G according to Theorem 10, the segments of E' that arise as m -span edges for the input edges E are such that the number of edges of E intersecting an m -span edge ab is at most $O(|ab|)$.*

Provided that the input R is nontrivial ($n \geq 2$), the length L^* of an optimal solution OPT (path or cycle) to ATSP is at least 2; thus, Theorem 10 shows that there exists an m -guillotine supergraph of OPT of length $L' \leq (1 + O(1/m))L^*$. Further, as shown in [5, 12, 14], one can make the m -guillotine conversion using cuts whose coordinates are from among a discrete set of $O(n)$ candidate x - and y -coordinates, for fixed m . We will show how to use this fact, along with the structure of an optimal adversarial solution, to construct via dynamic programming an m -guillotine structure from which we can extract an approximation to OPT , with approximation factor $(1 + \epsilon)$, for any $\epsilon > 0$. (Here, $m = O(1/\epsilon)$.)

4.3 The Dynamic Program

A subproblem of our dynamic program (DP) is responsible for computing a shortest total length connected network that spans the input set R of disks (at their sample points) while satisfying a constant-size, $O(m)$, set of boundary conditions. The boundary conditions specify $O(m)$ disks that the subproblem is responsible for interconnecting, as well as conditions on how the computed network within this subproblem should interact with optimal solutions computed within abutting subproblems. As we cannot afford to keep track of all (potentially $\Omega(n)$) interconnections of the optimal ATSP solution, OPT , between two rectangles that bound subproblems, the m -guillotine structure theorem, together with our additional structural results, allow us to compactly summarize the interconnection information well enough to ensure approximation within factor $(1 + \epsilon)$ of optimal.

Unlike the PTAS for TSPN, where the DP can choose any point within each region of R , in computing a minimum-weight connected, Eulerian, m -guillotine spanning subgraph over \mathcal{G} , in the ATSP we have no control over the point being spanned within each region: Once we produce an ordering σ_R , the adversary gets to solve an offline longest path problem to choose the (“worst possible”) point $p_i \in d_i$ within each region $d_i \in R$ our tour must visit. Thus, we need to create a minimum weight connected spanning Eulerian subgraph over \mathcal{G} that satisfies the m -guillotine property *and* satisfies a certain *adversarial subpath property*, which allows us to show that in the resulting network computed by DP, we can extract a polygonal tour of R that satisfies the adversarial property. In essence, we need the DP subproblems to be able to estimate (approximately) what the cost of an *adversarial* solution will be, if we extract from the optimized m -guillotine network a tour through R .

In particular, each DP subproblem is specified by a window $W \subseteq \mathcal{B}$, along with the following additional information:

1. An m -span (possibly empty) on each of the 4 sides of W , each with a parity bit indicating whether the number of edges incident to the m -span from outside of W is even or odd;
2. $O(m)$ *specified edges*, which are the network edges crossing the boundary of W that are not crossing one of the (up to 4) m -spans;
3. An m -disk span (possibly empty) on each of the 4 sides of W , with a specified sample point given for the first and for the last disk along the m -disk span;
4. $O(m)$ *specified input disks* (i.e., disks of R not intersecting an m -disk span) intersecting the boundary of W ;
5. A specified sample point of \mathcal{G} on the boundary of each of the $O(m)$ specified input disks, where the network is required to visit the associated disks (these are the “guessed” positions of the adversarial visitation points for the specified disks);
6. For each of the $O(m)$ specified input disks, we indicate whether the specified sample point of the disk is visited by the network being computed for the subproblem, and, if so, whether its degree in that network is 1 or 2. (The total degree of the sample point, using edges associated with subproblems on both sides of the cut, will be 2.)
7. An interconnection pattern specifying the subsets of the $O(m)$ boundary elements (specified input disks, specified edges, m -spans, and m -disk spans) that form connected components within W .

There are only $O(n^4)$ choices for W , $n^{O(m)}$ choices for the specified edges/disks, and a constant ($O(g(m))$, for some function g) number of choices of the $O(m)$ bits and the interconnection patterns. Thus, there are a polynomial number of subproblems for the DP.

A subproblem in the dynamic program requires one to compute a minimum-length m -guillotine network satisfying the following constraints:

- (i) The network is comprised of edges of the following types: (a). edges from the set $E_{\mathcal{G}}$ of edges linking a sample point of \mathcal{G} on one disk to a sample point of \mathcal{G} on another disk; (b). edges of type (a), $E_{\mathcal{G}}$, truncated at a (Steiner) attachment point on an m -span where the edge crosses the m -span or passes through an endpoint of the m -span; and (c) m -spans and m -disk spans that lie along cuts in the decomposition (recall that cuts lie along $O(n)$ discrete horizontal/vertical lines). The attachment points and the endpoints of m -spans and m -disk spans constitute a set, \mathcal{H} , of *Steiner points*, distinct from the sample points \mathcal{G} on the boundaries of the disks.
- (ii) Each sample point of \mathcal{G} within W that is visited by the network has degree 2.
- (iii) The number of edges of type (b) (i.e., edges of $E_{\mathcal{G}}$ truncated at an m -span) incident on an m -span segment ab is even or odd, according to whether the parity bit of the m -span is even or odd, so that the total sum of the degrees of the Steiner points \mathcal{H} along an m -span is even. Further, the number of edges of type (b) incident on an m -span segment ab is bounded by $c_0 \cdot |ab|$, where c_0 is a constant arising from the structure Theorem 9.
- (iv) The network must be m -guillotine with respect to W , and, for each cut in the recursive partitioning of W , in the total length of the network we count each m -span twice; these doubled m -spans allow us to augment the resulting network to be Eulerian [12], and thereby to extract a tour (see below). Further, we count the length of each m -disk span a constant ($O(1)$) times as well; this will allow the m -disk spans to be converted into adversarial subpaths visiting the set of disks that are spanned.
- (v) The network must utilize the specified edges (which straddle the boundary of W).
- (vi) The network must visit, at a sample point, each of the input disks interior to W .
- (vii) The network must visit each specified disk whose bit indicates it should be visited by the subproblem, at the specified sample point for that disk. Further, the network must visit the specified sample points for the first and last disk associated with each nonempty m -disk span.
- (viii) The network must obey the interconnection pattern of the subproblem.
- (ix) The network obeys the *adversarial subpath property*: Any maximal path, endpoints non-inclusive, within the network that goes through only sample points \mathcal{G} is a longest path through the sequence of disks on which the sample points lie (one per disk).

► **Lemma 12.** *When an optimal tour OPT is rounded to the grid \mathcal{G} and then converted to become m -guillotine in the process that proves Theorem 10, the network that results from the augmentation of OPT satisfies conditions (i)-(viii) at every level of the recursive process, for appropriate choices of the specified edges, disks, and interconnection patterns.*

Proof. During the process that converts OPT to be m -guillotine, according to the constructive proof of Theorem 10, most of the conditions hold automatically, by construction. Edges of OPT that cross an m -span, ab , do so at a point of \mathcal{H} that has degree 4 (since the crossing edge is partitioned at the crossing point, becoming two truncated type-(b) edges, and the m -span is partitioned at the crossing point as well). An m -span edge ab , by construction, extends between two points (a and b , each a Steiner point) on edges of OPT (each of which is thereby partitioned into two truncated type-(b) edges). Theorem 11 implies that the fat edges associated with the edges of OPT have bounded depth, implying condition (iii) holds. Condition (viii) holds for the choice of interconnection pattern that is implied by OPT . The adversarial subpath property holds because of the adversarial path property of OPT itself. ◀

We now discuss the enforcement of condition (ix), the adversarial subpath property, which is key to our being able to account for the adversary's choices during our optimization of the

network length, assuring that, in the end, we can extract from the computed network a tour that is adversarial and not much longer than the overall network.

Let (W, Σ) denote a subproblem associated with window W , where Σ is a specification of the boundary constraints information (1)-(7). The dynamic programming recursion optimizes the partition of the subproblem (W, Σ) into two subproblems, (W_1, Σ_1) and (W_2, Σ_2) , by a horizontal or vertical cut line ℓ (intersecting W and passing through one of the $O(n)$ discrete values of x, y -coordinates that define windows). Crucial to the correctness of the algorithm is that this recursion preserves the properties specified by the conditions (i)-(ix).

The objective function, $f(W, \Sigma)$, measures the total length of the network restricted to the window W ; in particular, edges of E_G that are specified in the boundary constraints Σ have their length partitioned and assigned to subproblems through which they pass.

The DP recursion optimizes over the choice of the cut line ℓ that partitions W into W_1 and W_2 , as well as the boundary conditions, Σ_ℓ , along the cut, which will be part of the specifications Σ_1 and Σ_2 . The conditions Σ_1 and Σ_2 must be compatible with each other and with the choice of boundary conditions, Σ_ℓ , across the cut ℓ . In particular, in order for Σ_1 and Σ_2 to be compatible with each other and with Σ , the specified edges of E_G across ℓ must match, as well as the m -span and m -disk span along the cut ℓ . Further, the interconnection pattern of Σ must specify subsets of boundary elements for W that are yielded by taking the union of interconnection patterns for (W_1, Σ_1) and (W_2, Σ_2) .

We let $\Sigma_\ell^{(R)}$ denote the partial specification of the boundary conditions Σ_ℓ , in which we specify which pairs of disks from R constitute the specified edges crossing ℓ , but do not specify the actual sample points on the boundaries of these disks that define the endpoints of the edges from E_G being specified. (In other words, $\Sigma_\ell^{(R)}$ specifies only the equivalence classes of the full set of conditions, Σ_ℓ ; the refinement of these equivalence classes will be specified in the optimization within the “max” term of the recursion below.) The DP recursion is

$$f(W, \Sigma) = \min_{\ell, \Sigma_\ell^{(R)}} \left\{ \max_{\Sigma_\ell \in X(\Sigma_\ell^{(R)})} (f(W_1(\ell), \Sigma_1(\Sigma_\ell)) + f(W_2(\ell), \Sigma_2(\Sigma_\ell))) \right\}$$

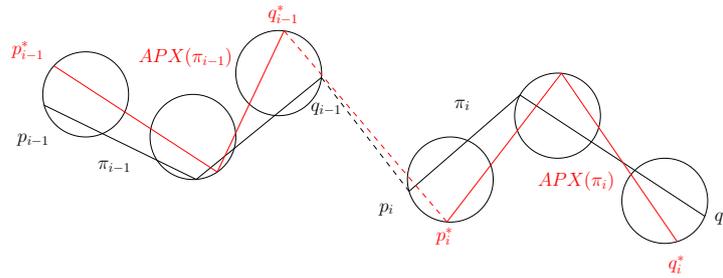
where the outer minimization is over choices of the cut ℓ and the cross-cut boundary conditions $\Sigma_\ell^{(R)}$, and the inner maximization is over choices of Σ_ℓ that are in the set $X(\Sigma_\ell^{(R)})$ of all boundary conditions across the cut ℓ that are refinements of the choice $\Sigma_\ell^{(R)}$, specifying precisely which sample points are utilized for each of the disks of R that are involved in the specification $\Sigma_\ell^{(R)}$ (and not already specified by the “parent” choice, Σ , in cases in which edges crossing ℓ also extend outside of W and have their sample points specified within Σ). In the expression above, $W_1(\ell)$ and $W_2(\ell)$ are the subwindows of W on either side of the cut ℓ , and $\Sigma_1(\Sigma_\ell)$ and $\Sigma_2(\Sigma_\ell)$ are the corresponding boundary conditions on either side of ℓ that are inherited from Σ and consistent with the conditions Σ_ℓ . The fact that we maximize over the choices that the adversary can make, in all choices that cross the cut ℓ , implies that we preserve the adversarial subpath property:

► **Lemma 13.** *The DP yields a network satisfying (ix), the adversarial subpath property.*

4.4 Extracting an Approximating ATSP Tour

The output of the DP algorithm is an m -guillotine network G of minimum cost, where cost is total length, taking into account that m -spans are counted twice, and m -disk spans are counted $O(1)$ times (and are each of length at least 1). From the structure Theorem 10, we know that the total length of edges of G is at most $|OPT|(1 + O(1/m))$.

The fact that we accounted for the doubling of the m -spans in the optimization implies that we can afford to augment the edges along each m -span, in order that every Steiner



■ **Figure 5** Bounding the adversarial increase over extracted approximate tour T' .

point along an m -span has degree 4: Initially, the points \mathcal{H} along an m -span have degree 3, being either endpoints of the m -span (having a T-junction with an edge between two sample points of \mathcal{G}), or being a T-junction where an edge between two sample points of \mathcal{G} is truncated, terminating on the m -span. By the parity condition at the m -span, we know that there are an even number of T-junctions along the m -span, implying that we can add a perfect matching of segments along the m -span, joining consecutive pairs of T-junctions. The total length of this matching is less than the length of the m -span, and is “paid for” by the doubling of the m -span lengths in the DP optimization.

The fact that we accounted for $O(1)$ copies of the m -disk spans in the optimization implies that we can afford to augment the edges of G with an adversarial path through the sequence of disks stabbed by the m -disk span; such a path has length proportional to the length of the m -disk span, assuming the m -disk span is nontrivial in length.

By the above discussion, the result of our algorithm is, then, a connected Eulerian network of length at most $|OPT|(1 + O(1/m))$. From this network, we extract a tour T . The tour T is a cycle consisting of straight line segments joining points that are either sample points, \mathcal{G} , or Steiner points, \mathcal{H} . Let $\pi_1, \pi_2, \dots, \pi_k$ be the maximal subpaths along T whose vertices are all sample points \mathcal{G} ; i.e., each path π_i has only vertices of \mathcal{G} (no Steiner points \mathcal{H}), and every sample point of \mathcal{G} that is a vertex of T lies in exactly one path π .

Now, the number, k , of subpaths is at most the number of Steiner points \mathcal{H} along the tour T , and this number is upper bounded by the number of Steiner points along m -spans in the entire network. But, the total length of all m -spans is at most $O(1/m) \cdot |OPT|$, by the proof of Theorem 10. This implies that $k \leq O(1/m) \cdot |OPT|$.

The adversarial subpath property that was enforced in the dynamic programming algorithm implies that the subpaths π_i are each adversarial – their lengths are longest possible, for the given sequence of disks through which it passes (given that the path π_i begins at sample point p_i and ends at sample point q_i as chosen by our dynamic program). We obtain a new tour, T' , by chaining together the subpaths π_i , omitting any Steiner points that were along T . The resulting tour T' is not necessarily adversarial, but the following lemma shows that it is close to being so.

► **Lemma 14.** *Let σ'_R be the order in which the disks R are visited by the tour T' . Then, the adversarial tour, APX , associated with σ'_R has length at most $|T'| + O(1/m) \cdot |OPT|$.*

Proof. For each subpath π_i in our approximate tour T' let $APX(\pi_i)$ be the adversarial path computed over the sequence of disks associated with π_i in the final adversarial tour associated with σ'_R . Similarly let p_i^* (resp., q_i^*) be the point chosen in the first (resp., last) disk along π_i in APX . Assume that we have chosen the points p_i (resp., q_i) in the first (resp., last) disk along π_i in our extracted tour T' ; see Figure 5 for an illustration.

From the adversarial property of our computed solution we know that the length of the path $|\pi_i|$ computed from p_i to q_i is as long as possible over choices in intermediate disks. Therefore we can over estimate the length of APX by walking around T' adding at most two unit diameter detours to the first and last disk in each sub-path π_i . That is we can begin at p_i detour to p_i^* and back, follow π_i until we reach q_i and then detour from q_i to q_i^* and back and follow T' to p_{i+1} and so on. By triangle inequality and the fact that π_i is a longest path for fixed choices of p_i , and q_i this over estimates the length of APX .

We have $|APX| \leq |T'| + 4k$, and therefore $|APX| \leq |T'| + O(1/m)|OPT|$, because, as previously stated $k \leq O(1/m)|OPT|$. ◀

Since we know that the computed T , and thus T' , has length at most $|OPT|(1 + O(1/m))$, Lemma 14 implies that the overall solution extracted from our computed tour yields a PTAS.

References

- 1 Esther M. Arkin and Refael Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55(3):197–218, 1994.
- 2 Dimitris J. Bertsimas, Patrick Jaillet, and Amedeo R. Odoni. A priori optimization. *Operations Research*, 38(6):1019–1033, 1990.
- 3 Gui Citovsky, Tyler Mayer, and Joseph S. B. Mitchell. TSP With Locational Uncertainty: The Adversarial Model, March 2017. arXiv:1705.06180 [cs.CG]. URL: <https://arxiv.org/abs/1705.06180>.
- 4 Reza Dorrigiv, Robert Fraser, Meng He, Shahin Kamali, Akitoshi Kawamura, Alejandro López-Ortiz, and Diego Seco. On minimum-and maximum-weight minimum spanning trees with neighborhoods. *Theory of Computing Systems*, 56(1):220–250, 2015.
- 5 Adrian Dumitrescu and Joseph S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *Journal of Algorithms*, 48:135–159, 2003. Special issue devoted to 12th ACM-SIAM Symposium on Discrete Algorithms, Washington, DC, January, 2001.
- 6 Robert Fraser. *Algorithms for geometric covering and piercing problems*. PhD thesis, University of Waterloo, 2012.
- 7 Patrick Jaillet. A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Operations Research*, 36(6):929–936, 1988.
- 8 Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for TSP, Steiner tree, and set cover. In *Proc. 37th ACM Symposium on Theory of Computing*, pages 386–395. ACM, 2005.
- 9 Pegah Kamousi and Subhash Suri. Euclidean traveling salesman tours through stochastic neighborhoods. In *International Symposium on Algorithms and Computation*, pages 644–654. Springer, 2013.
- 10 Chih-Hung Liu and Sandro Montanari. Minimizing the diameter of a spanning tree for imprecise points. In *International Symposium on Algorithms and Computation*, pages 381–392. Springer, 2015.
- 11 Maarten Löffler and Marc van Kreveld. Largest and smallest convex hulls for imprecise points. *Algorithmica*, 56(2):235–269, 2010.
- 12 Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.
- 13 Joseph S. B. Mitchell. Shortest paths and networks. In Jacob E. Goodman and Joseph O’Rourke, editors, *Handbook of Discrete and Computational Geometry (2nd Edition)*, chapter 27, pages 607–641. Chapman & Hall/CRC, Boca Raton, FL, 2004.

32:16 TSP With Locational Uncertainty: The Adversarial Model

- 14 Joseph 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. Society for Industrial and Applied Mathematics, 2007. URL: <http://www.ams.sunysb.edu/~jsbm/papers/tspn-soda07-rev.pdf>.
- 15 Sandro Montanari. *Computing routes and trees under uncertainty*. PhD thesis, Dissertation, ETH-Zürich, 2015, No. 23042, 2015.
- 16 Yang Yang, Mingen Lin, Jinhui Xu, and Yulai Xie. Minimum spanning tree with neighborhoods. In *International Conference on Algorithmic Applications in Management*, pages 306–316. Springer, 2007.