

GEOMETRIC SPANNERS FOR COMPACT REGIONS

by

FRANÇOIS-XAVIER CHAPLAIN-CORRIVEAU

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfillment
of the requirements for the degree of

Master

in

Computer Science

Carleton University

Ottawa, Ontario

© 2021

François-Xavier Chaplain-Corriveau

ABSTRACT

We generalize the spanner problem for additively weighted point sets introduced by Bose, Carmi and Couture (Spanners of additively weighted point sets, *J. Discrete Algorithms*, 2011) to finite sets of compact and path connected sets in \mathbb{R}^2 . Namely, let \mathcal{R} be a set of n pairwise disjoint compact path connected subsets of the plane. We refer to the elements of \mathcal{R} as regions. We suppose that travelling inside each of these regions is free, whereas the cost of travelling outside these regions corresponds to the travelled distance. This gives rise to a metric space (\mathcal{R}, δ_T) where T corresponds to the most cost effective graph connecting the regions of \mathcal{R} and δ_T corresponds to the length of a shortest path in T . In this frame of work, the sparse spanner problem translates to approximating this graph T , with not too many edges, in such a way that the cost of travelling between any pair of regions is at most some multiplicative factor of the optimal cost.

We show that T can be approximated by two different cone based algorithms: the first algorithm, based on *Yao* graphs, yields a graph whose spanning ratio is $1/(1 - 2 \sin(\theta/2))$ and the second, based on θ graphs, yields a graph whose spanning ratio is $1/(\cos \theta - \sin \theta)$.

We also introduce a notion of "fatness" and show that, when the regions are "fat" and convex, a simple greedy algorithm builds a t approximation of T with $\mathcal{O}((f + \frac{1}{t-1})^3 n)$ edges where $f \geq 1$ corresponds to the fatness parameter and $t > 1$ is a real number sufficiently close to 1. We also show with the notion of fatness introduced by van der Stappen, Halperin and Overmars (The complexity of the free space for a robot moving amidst fat obstacles, *CGTA*, 1993), the same greedy algorithm outputs a t approximation of T with $\mathcal{O}((k/(t-1))^2 n)$ edges, where k is the fatness parameter and $t > 1$ is a real number sufficiently close to 1.

ACKNOWLEDGMENTS

I would like thank both my supervisors Dr. Prosenjit Bose and Dr. Michiel Smid. They introduced me to the world of research and provided me with advice and support as I needed it. They and the many other members of the Computational Geometry Lab form a very warm group of researchers and I could not have found a better place to develop my interest for research.

I would also like to thank all the COMP2804 students I have had the pleasure to meet through out my thesis. Our many interactions provided me with a great way to relax and take my mind off things when I needed it the most.

CONTENTS

1	INTRODUCTION	1
2	PRELIMINARIES	5
3	REGION GRAPHS	9
4	EXISTENCE OF SPARSE SPANNERS OF T	15
5	POLYGONAL SPANNERS	19
6	REGION SPANNERS FOR FAT CONVEX REGIONS	31
6.1	Introduction	31
6.2	Bounding the number of good edges in T	33
6.3	Bounding the number of nasty edges in T	34
6.4	The Greedy Spanner for Fat Convex Regions	42
6.5	Other notions of fatness	45
7	OPEN AND RELATED PROBLEMS	51
A	APPENDIX	53
A.1	Proof of 5.1.2	53
A.2	Defining the event queue when several events can be hit by the sweep line	54
	BIBLIOGRAPHY	57

LIST OF FIGURES

Figure 1.1	A network N on D_1, D_2 and D_3	2
Figure 2.1	The selection process in the construction of a θ_k -graph. We add an edge between p and r since its projection on $bis(C_{i,p})$ is closest to p .	6
Figure 3.1	A set of pairwise disjoint regions in \mathbb{R}^2	9
Figure 3.2	The shortest line segment S joining R and R' is intersected by R'' . The portion of S in red that goes inside R'' has no cost.	10
Figure 3.3	The graph T_0 induced by a set of pairwise disjoint regions. The dotted edges are useless.	11
Figure 3.4	The transportation graph T induced by a set of pairwise disjoint regions.	12
Figure 3.5	A region graph for the set of regions $\mathcal{R} = \{R_1, R_2, R_3, R_4, R_5\}$	13
Figure 4.1	The selection process in the construction of a Yao graph Y_k . The yellow edge is added to Y_k since it is the shortest edge with p as an endpoint contained in $C_{i,p}$	15
Figure 4.2	Each edge incident to the center of D in the complete bipartite graph induces an edge of the transportation graph T on the corresponding set of disks. This is true for the center of each disk, resulting in $\mathcal{O}(n^2)$ edges.	18
Figure 5.1	The selection process in the construction of θ_k .	20
Figure 5.2		22
Figure 5.3	In red we have the point q and in purple we have the line orthogonal to $bis(C_{i,p})$. In blue we have x and the portion of S corresponding to L_{xq} .	23
Figure 5.4	In orange we have r_6 for $\theta = \pi/8$. For each vertex p we want to report the first intersection on $r_6 + p$	24
Figure 5.5	In orange we have a sweep line oriented with respect to the ray it corresponds to. At that moment the leaves in the BST storing the segment are from left to right $S_1, S_2, S_3, S_4, S_5, S_6$ which corresponds to the order of their intersection as we are walking in the direction of L_i .	25

- Figure 6.1 Two regions R and R' that are 4-fat and 2-fat respectively. In blue we have ρ_R^- and $\rho_{R'}^-$. In orange we have ρ_R^+ and $\rho_{R'}^+$. 31
- Figure 6.2 The oriented transportation graph T on a set of 2-fat regions. 32
- Figure 6.3 A long edge between two regions. The regions are far apart compared to ρ_R^- . 32
- Figure 6.4 A nasty edge between two regions. Although the regions are close, their inner disks are far apart 32
- Figure 6.5 A good edge between two regions. The two inner disks are close. 33
- Figure 6.6 In orange we have the point u . D'_0 and c'_0 are shown in blue and purple respectively. 34
- Figure 6.7 36
- Figure 6.8 In orange and purple we have the lines L_1 and L_2 respectively. In blue and green we have the line V_2 and V_1 respectively. The point u is in red. 37
- Figure 6.9 In orange we have the cone C_B and in blue we have the cone C_A . In green we have L_q^- and in purple we have L_q^+ 38
- Figure 6.10 39
- Figure 6.11 39
- Figure 6.12 In blue we have the intersection point c' . In green we have the point u and in pink we have the point b . 40
- Figure 6.13 In light blue we have the disk D'_0 41
- Figure 6.14 The disks D_1 and D_2 are in U_R , whereas D_3 and D_4 are not. 46
- Figure 6.15 The degree of the big disk is linear in terms of the total number disks: there is an edge between each tiny disk and the big disk. 50
- Figure A.1 53
- Figure A.2 55

INTRODUCTION

Networks and graphs are prime objects for data visualization and data handling in several fields including mathematics and computer science. A graph N on a set of points in the plane is *geometric* if each edge of N between two points p, q corresponds to a line segment between p and q and is assigned a weight equal to the length of that segment. One of the most extensively studied problem in Computational Geometry consists of designing algorithms that produce geometric networks with not too many edges such that between any two points p and q there exists a path whose length is within a small multiplicative factor of the Euclidean distance between p and q . This is typically referred to as the *spanner problem*.

As far as we know, *Bose et al.* [4] were the first to introduce the spanner problem on additively weighted point sets: Let S be a set of points in the plane, and let each point p of S be assigned some weight $\omega(p) \geq 0$. The distance between two weighted points p and q is defined as

$$|pq| - \omega(p) - \omega(q)$$

where $|pq|$ denotes the Euclidean distance between p and q . Under the assumption that $|pq| > \omega(p) + \omega(q)$ for all $p, q \in S$, in which case S can be thought of as a set of disjoint disks, they were able to show that it is possible to design a graph N with not too many edges such that for any pair of disks D, D' of S there is a path in N whose length is arbitrarily close to the Euclidean distance between D and D' . Recently Smid [15] improved this result by reducing the number of edges needed by a factor of 4.

More specifically, as done in [15], we can redefine the additively weighted point set spanner problem as follows; Let \mathcal{D} be a set of n pairwise disjoint disks in the plane. Suppose that inside each disk we can travel between a pair of points at no cost whereas outside all disks the travel cost between two points is their Euclidean distance. For any two disks D, D' of \mathcal{D} , a point x in D , and a point y in D' , let $\delta(x, y)$ be the minimum cost required to travel from x to y . Note that by our assumption, for any point x' of D and y' of D'

$$\delta(x, y) = \delta(x', y')$$

Let $\delta(D, D')$ be that common value; it is the minimum cost needed to travel from any point of D to any point of D' . Solving the spanner problem for this set \mathcal{D} entails producing a network N with edges

connecting the disks of \mathcal{D} such that for any disks $D, D' \in \mathcal{D}$ the cost required to travel from D to D' using the edges of N , call it $\delta_N(D, D')$, is within some multiplicative factor of $\delta(D, D')$. For simplicity of notation we will write $\text{dist}(p, q)$ as $|pq|$ in the case of points and $\text{dist}(P, Q)$ as $|PQ|$ in the case of sets.

As we can see in the case of disks in Figure 1.1 below, one of the difficulties of this problem lies in the fact that, although D_1, D_2 are far apart, the cost for travelling from D_1 to D_2 is small relative to $|D_1D_2|$. Namely, the triangle inequality does not hold: $|D_1D_2| > |D_1D_3| + |D_3D_2|$.

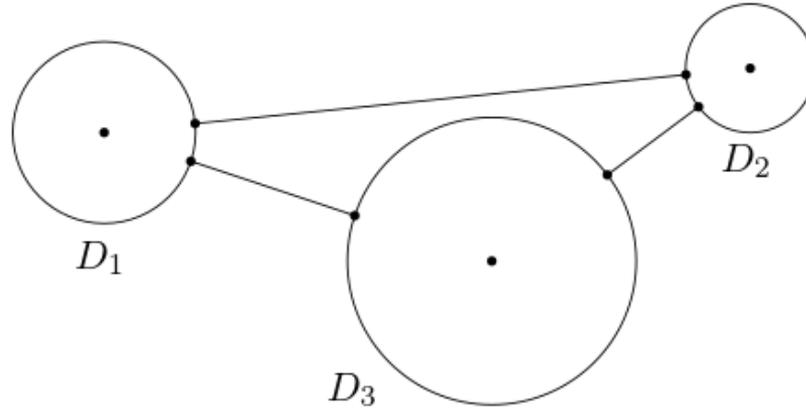


Figure 1.1: A network N on D_1, D_2 and D_3

Smid [15] shows the following:

Theorem 1.1 *Let \mathcal{D} be a set of n pairwise disjoint disks in plane. Let $0.389\text{rad} > \theta > 0$ such that $2\pi/\theta$ is an integer. There exists a network N on \mathcal{D} with at most $(2\pi/\theta)n$ edges such that for any disks D, D' of \mathcal{D}*

$$\delta_N(D, D') \leq t \cdot \delta(D, D')$$

where

$$t = \frac{1-2\sin\theta}{1-2\sin\theta-2(1-\sin\theta)\sin(\theta/2)}$$

In this thesis we tackle the problem of generating sparse spanners on a collection $\mathcal{R} = \{R_1, \dots, R_n\}$ of n pairwise disjoint closed, bounded, and path connected sets. We suppose that the cost of segments inside one of these sets R_i is zero. The cost of a segment that does not intersect any of the interiors of the sets R_i is its length. We show that analogous results to Theorem 1.1 hold and explicitly describe algorithms for constructing spanner networks.

In Chapter 2 we briefly review some of the known results for generating sparse spanners on sets of points.

In Chapter 3 we formally define the notions of regions and a transportation graph T for sets of regions \mathcal{R} , and give a few preliminary

results.

In Chapter 4 we prove the existence of sparse (region) spanners of the transportation graph induced by a finite set of regions in \mathbb{R}^2 by using a Yao graph construction on the edges of the transportation graph.

In Chapter 5, we show that a Θ -graph construction can be used to produce sparse region spanners for a finite set of pairwise disjoint polygons whose size depends on the number of vertices. The method described in this chapter does not require the precomputation of the transportation graph.

In Chapter 6, we introduce the following notion of fatness: A region R is f -fat for $f \geq 1$ if there exists two concentric disks D_R, D_R^+ such that D_R is contained in R , D_R^+ contains R , and the radius of D_R^+ is at most f times the radius of D_R . We show that for finite sets of "fat" convex regions sparse spanners of the transportation graph exists, regardless of the number of endpoints of the edges of the transportation graph. Then, using an alternate definition of fatness, we show the same result without the requirement for the regions to be convex.

PRELIMINARIES

Let S be a set of points in \mathbb{R}^2 . A *network connecting the points of S* is a connected graph $N = (S, E)$ with vertex set S and edge set $E \subseteq S \times S$. N is said to be *geometric* if each edge of N is weighted by the Euclidean distance between its endpoints. The following Theorem is a well-known result from graph theory.

Theorem 2.1.1. *Let S be a set of n points. Any connected network on the points of S must have at least $n - 1$ edges.*

We shall say that a network N on a set of n points is *sparse* if $|E| \in \mathcal{O}(n)$. Let $t \geq 1$. A network N is a *t -spanner* for S if for any points p, q of S there is a path between p and q in N of length at most $t|pq|$. The smallest t such that N is a t -spanner of S is called the *spanning ratio* of N . One technique to construct a t -spanner is based on the idea of using cones.

Let $k \geq 2$ be some integer and let $\theta = 2\pi/k$. We define k cones as follows: for $0 \leq i \leq k - 1$ the boundary of C_i corresponds to the rays obtained by rotating the positive x -axis about the origin by an angle of $i\theta$ and $(i + 1)\theta$ in counter clockwise direction. We shall sometimes refer to the rays defining the boundary of a cone C as the *bounding rays* of C . We shall denote by \mathcal{C}_k the collection of these k cones and for each cone $C_i \in \mathcal{C}_k$ we shall denote its bisector by $bis(C_i)$.

Let $C_i \in \mathcal{C}_k$ and let p be a point in the plane. We will denote by $C_{i,p}$ the set $\{x + p : x \in C_i\}$. In other words $C_{i,p}$ is the cone obtained by translating C_i so that its apex is p .

We define the graph θ_k on S as follows:

- The vertex set of θ_k is S .
- For each cone $C_i \in \mathcal{C}_k$ and for each point p of S such that $C_{i,p}$ contains at least one point of $S \setminus \{p\}$, θ_k contains the edge $\{p, r\}$ where r is the point in $C_{i,p} \cap S \setminus \{p\}$ whose orthogonal projection onto $bis(C_{i,p})$ is closest to p .

Theorem 2.1.2. [6] *Let S be a set of n points in the plane. Let $k \geq 1$ be an integer. The graph θ_{4k+2} is a sparse t -spanner for S , where $t = (1 + 2 \sin(\theta/2))$.*

Theorem 2.1.3. [6] *Let S be a set of n points in the plane. Let $k \geq 1$ be an integer. The graphs θ_{4k+3} and θ_{4k+5} are sparse t -spanners for S , where $t = \frac{\cos(\theta/4)}{(\cos(\theta/2) - \sin(3\theta/4))}$.*

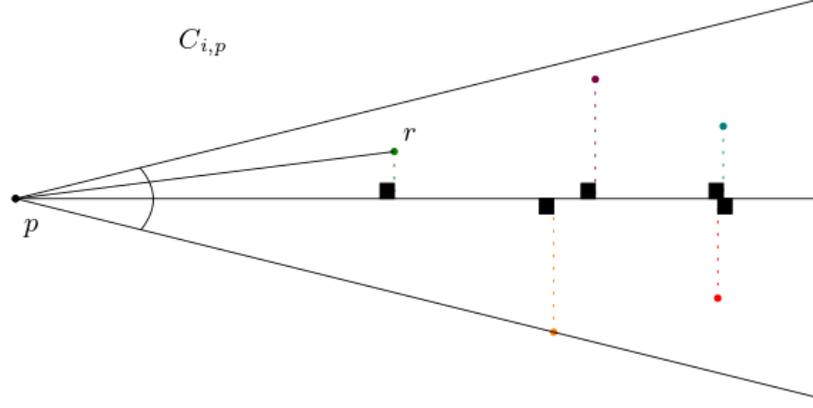


Figure 2.1: The selection process in the construction of a θ_k -graph. We add an edge between p and r since its projection on $\text{bis}(C_{i,p})$ is closest to p .

Theorem 2.1.4. [6] Let S be a set of n points in the plane. Let $k \geq 1$ be an integer. The graph θ_{4k+4} is a sparse t -spanner for S , where $t = 1 + \frac{2 \sin(\theta/2)}{\cos(\theta/2) - \sin(\theta/2)}$.

It is worth mentioning that Bose *et al.* [6] also show that the bound on the spanning ratio of θ_{4k+2} obtained from Theorem 2.1.2 is tight and give lower bounds on the spanning ratio of θ_{4k+3} , θ_{4k+4} , and θ_{4k+5} . A closely related graph introduced independently by Yao [18] and Flinchbaugh and Jones [12] commonly referred as Yao graph denoted as Y_k is defined as follows:

- The vertex set of Y_k is S .
- For each cone $C_i \in \mathcal{C}_k$ and for each point p of S such that $C_{i,p}$ contains at least one point of $S \setminus \{p\}$, Y_k contains the edge $\{p, r\}$ where r is the point in $C_{i,p} \cap S \setminus \{p\}$ that is closest to p with respect to the Euclidean distance.

Theorem 2.1.5. [3] For all $k \geq 1$, the graph Y_{4k+3} has spanning ratio at least $1 + 2 \sin(3\theta/8) + 4 \frac{(\sin(13\theta/18) + \sin(19\theta/16)) \sin(\theta/16) \sin(3\theta/8)}{\sin(2\theta)}$, where $\theta = 2\pi/(4k+3)$.

Theorem 2.1.6. [3] For all $k \geq 1$, the graph Y_{4k+4} has spanning ratio at least $1 + 2 \sin(\theta/2)(1 + \tan(\theta/2))$, where $\theta = 2\pi/(4k+4)$.

Theorem 2.1.7. [3] For all $k \geq 1$, the graph Y_{4k+5} has spanning ratio at least $1 + 2 \sin(3\theta/8) + 4 \sin(5\theta/16) \sin(3\theta/8)$, where $\theta = 2\pi/(4k+5)$.

The notions of network and spanners have natural extensions in \mathbb{R}^d . In fact the Yao graph was introduced for any dimension $d \geq 2$ [18] and the family of θ graphs were defined in \mathbb{R}^d by Ruppert and Seidel [14].

Although both constructions are fairly similar, θ graphs are sometimes favored over Yao graphs because there exists a simpler algorithm for building the edge set of θ_k that naturally extends to higher

dimensions (see [13]). In particular, the θ_k graph on a set S of n points in \mathbb{R}^d can be constructed in $\mathcal{O}((n \log^{d-1} n)/\theta^{d-1})$ and using $\mathcal{O}(n/\theta^{d-1} + n \log^{d-2} n)$ space. It is not known whether or not Y_k can be built within that time frame for $d > 2$. In Chapter 5 we present an analogous algorithm for building θ graphs to the one described in [13].

Another technique for building spanners on sets of points is based on the notion of *well-separated pair decomposition* (WSPD) introduced by Callahan and Kosaraju [7]. Let A and B be two finite sets of points in \mathbb{R}^d . We say that A and B form a *well-separated pair* with respect to some real number s if there are two hyperdisks D_A, D_B of equal radius ρ such that

1. D_A contains the bounding box of A and D_B contains the bounding box of B .¹
2. $|D_A D_B| \geq s \cdot \rho$

A *well-separated pair decomposition* (WSPD) with respect to some positive real number s of a set S of n points is a sequence

$$\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}$$

of pairs of nonempty subsets of S , for some integer m , such that:

- (1) A_i and B_i are well-separated with respect to s , for each $1 \leq i \leq m$
- (2) For any two points p, q of S , there is exactly one index i such that $p \in A_i$ and $q \in B_i$, or $p \in B_i$ and $q \in A_i$

Theorem 2.1.8. [7] *Let S be a set of points in \mathbb{R}^d . Let $t > 1$ be a real number and let $s = 4(t + 1)/(t - 1)$. Given an arbitrary WSPD for S , let a_i be an arbitrary point of A_i and b_i be an arbitrary point of B_i for any for each i such that $A_i B_i$ form a well-separated pair. The graph $G = \{\{a_i, b_i\} \mid 1 \leq i \leq m\}$ is a t -spanner of S .*

Theorem 2.1.9. [7] *Let S be a set of points in \mathbb{R}^d and let $t > 1$ be a real number. In $\mathcal{O}(n \log n + n/(t - 1)^d)$ time, we can construct a t -spanner for S having $\mathcal{O}(n/(t - 1)^d)$ edges.*

We refer the reader to Narasimham and Smid [13] for a good survey of existing results for WSPD and their applications.

We present one last approach for building spanners on sets of points. Consider the following algorithm:

¹ The bounding box of a set of points A in \mathbb{R}^d is the smallest hyperrectangle containing A

Algorithm PATHGREEDY(S, t)

```

1 Input. A set  $S$  of  $n$  points, a real number  $t > 1$ .
2 Output. A  $t$ -region spanner of  $T$ 
3 Sort the  $\binom{n}{2}$  pairs of distinct points in non-decreasing order of their
  distances (breaking ties arbitrarily), and store them in a list  $L$ ;
4  $E := \emptyset$ ;
5  $G := (\mathcal{R}, E)$ ;
6 for each  $\{p, q\} \in L$  (consider pairs in sorted order)
7 do  $\delta_G(p, q) :=$  length of a shortest path in  $G$  between  $p$  and  $q$ ;
8   if  $\delta_G(p, q) > t|pq|$ 
9   then
10      $E := E \cup \{\{p, q\}\}$ ;
11      $G := (S, E)$ 
12   endif
13 endfor;
14 Output the graph  $G$ 

```

Note that PATHGREEDY outputs a t -spanner of S by construction. *Chandra et al.* showed that the graph constructed by PATHGREEDY is actually sparse.

Theorem 2.1.10. [9], [8] *Let S be a set of n points in \mathbb{R}^d , let $t > 1$ be a real number, and let $G = (S, E)$ be the t -spanner computed by Algorithm PATHGREEDY. For any two real numbers θ and w with $0 < \theta < \pi/4$, $0 \leq w < (\cos \theta - \sin \theta)/2$, and $t \geq 1/(\cos \theta - \sin \theta - 2w)$.*

- (1) *the degree in G of each vertex of S is $\mathcal{O}(1/\theta^{d-1})$, and*
- (2) *if $w > 0$, the weight of G is*

$$\mathcal{O}\left(\frac{1}{w\theta^{d-1}} \log n\right)$$

where the weight of a graph N corresponds to the sum of the weight of the edges N^2

A complete overview of PATHGREEDY for point sets is presented in [13]. For more results on the spanner problem for points and many of its variants we refer the reader to both [5] and [13].

² In our case the weight of an edge corresponds to the distance between its two endpoints.

REGION GRAPHS

Let R be a subset of \mathbb{R}^d . The *boundary* of R , denoted by ∂R , is the set of points p of R such that any disk D centered on p contains a point of R and a point not in R . R is said to be *closed* if $\partial R \subseteq R$. We say that R is *bounded* if for any points p, q of R , $|pq| < m$ for some positive real number m . A set that is both closed and bounded is said to be *compact*. R is said to be *path connected* if for any two points of R there is a path contained in R from p to q .¹

Let \mathcal{R} be collection of pairwise disjoint compact and path connected sets in the plane. We refer to the elements of \mathcal{R} as *regions*.

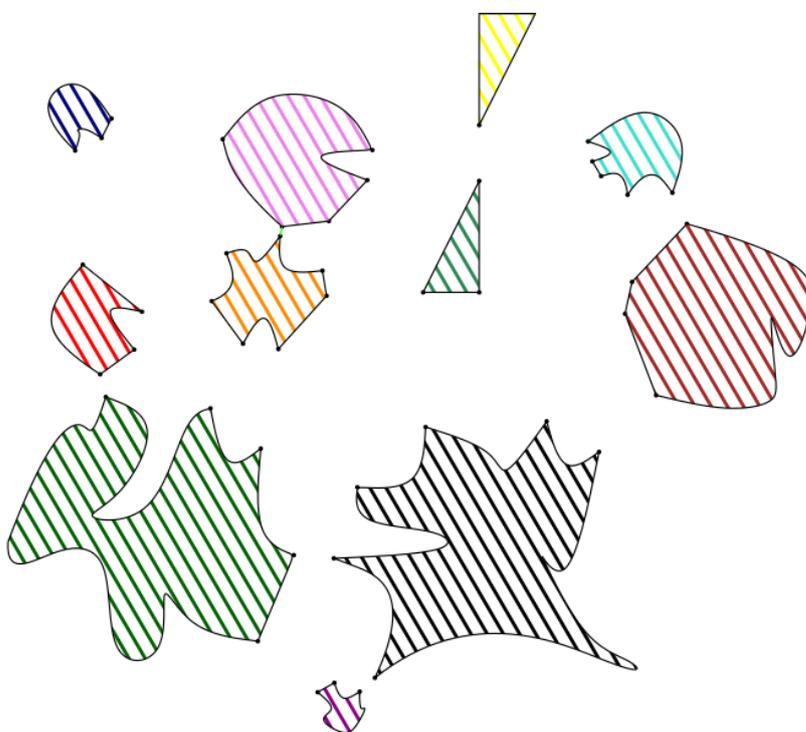


Figure 3.1: A set of pairwise disjoint regions in \mathbb{R}^2

We suppose that inside each region we can travel between a pair of points at no cost whereas outside all regions the travel cost between two points is their Euclidean distance. For any two regions R, R' of \mathcal{R} , a point x in R , and a point y in R' , let $\delta(x, y)$ be the minimum cost required to travel from x to y . Note that by our assumption, for any point x' of R and y' of R' $\delta(x, y) = \delta(x', y')$. Let $\delta(R, R')$ be that

¹ A path from p to q in R is a continuous function f from the interval $[0, 1]$ to R such that $f(0) = p$ and $f(1) = q$.

common value; it is the minimum cost needed to travel from any point of R to any point of R' . Solving the spanner problem for this set \mathcal{R} entails producing a graph N with edges connecting the regions of \mathcal{R} such that for any regions $R, R' \in \mathcal{R}$ the cost required to travel from R to R' using the edges of N , call it $\delta_N(R, R')$, is within some multiplicative factor of $\delta(R, R')$.

Lemma 3.1.1. *Let S be a shortest line segment joining two disjoint regions R, R' . S does not go through the interior² of R or R' .*

Proof. By definition, $\text{dist}(R, R') = \inf\{|pq| : p \in R, q \in R'\}$. Since both R and R' are closed and bounded, by the Extreme Value theorem $\text{dist}(R, R')$ achieves a minimum at some point p of R and q of R' [1]. Suppose p is not on the boundary of R . Since R is closed, the line segment from p to q intersects the boundary of R in a point $x \neq p, x \in R$. But then $\text{dist}(p, q) = \text{dist}(p, x) + \text{dist}(x, q)$. That is $\text{dist}(p, q) - \text{dist}(p, x) = \text{dist}(x, q)$. Since $x \neq p, \text{dist}(x, p) > 0$. Hence $\text{dist}(x, q) < \text{dist}(p, q)$ which contradicts the minimality of $\text{dist}(p, q)$. The argument for q is analogous.

Now suppose that S goes inside R . Since R is closed S must intersect the boundary of R in at least some other point. Let u be the point on the boundary of R that is hit last by walking on S starting from p . Since $u \neq p$ and $u \in S$, the line segment from q to u is strictly shorter than S . Indeed, we have that $\text{dist}(p, q) = \text{dist}(p, u) + \text{dist}(u, q)$ and so $\text{dist}(u, q) = \text{dist}(p, q) - \text{dist}(p, u) < \text{dist}(p, q)$ since $u \neq p$. This contradicts the minimality of S . \square

Lemma 3.1.1 tells us that a cost efficient edge between two regions R and R' corresponds to a shortest a line segment between one point on ∂R and one point on $\partial R'$. However, as we can see in the figure below if a line segment S between R and R' intersected some other region R'' , then S can be broken down into two segments whose combined length is at most the length of S .

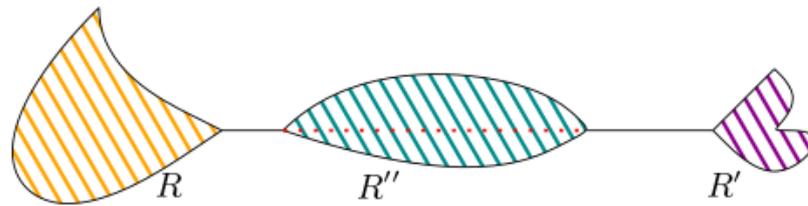


Figure 3.2: The shortest line segment S joining R and R' is intersected by R'' . The portion of S in red that goes inside R'' has no cost.

This inspires us to look at the following graph T_0 . Let \mathcal{R} be a set of pairwise disjoint regions, the graph T_0 is the graph whose vertex set is \mathcal{R} and, for each pair of regions R, R' of \mathcal{R} , if the shortest (in the Euclidean sense) line segment between R and R' does not intersect

² The interior of a set $R \subseteq \mathbb{R}^d$ corresponds to $R \setminus \partial R$

any other region of $\mathcal{R} \setminus \{R, R'\}$ then it is an edge of T_0 . We denote the shortest line segment between two regions R and R' by $L_{RR'}$ and its weight by $|RR'|$. Notice that in Figure 3.3 the dotted edges do not belong to any shortest path of T_0 . Such edges $\{R, R'\}$ are called *useless*. Otherwise $\{R, R'\}$ is called *useful*.

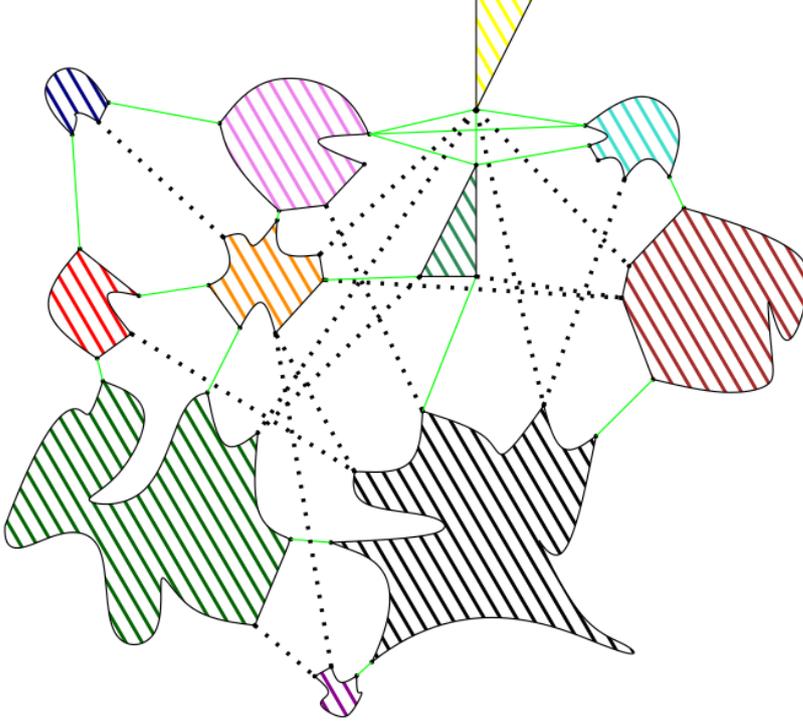


Figure 3.3: The graph T_0 induced by a set of pairwise disjoint regions. The dotted edges are useless.

The *transportation graph* T is the graph obtained from T_0 by removing all useless edges of the complete region graph T_0 . In particular, T corresponds to the most cost efficient graph on \mathcal{R} . The cost of travelling in T_0 (or in T) between two regions R and R' corresponds to the sum of the weights of the edges defining a shortest path from R to R' in T_0 (in T). We denote this cost by $\delta_{T_0}(R, R')$ ($\delta_T(R, R')$). The definition of T entails that for any regions R, R'

$$\delta_{T_0}(R, R') = \delta_T(R, R')$$

We now give a criterion to determine whether or not an edge of T_0 is an edge of the transportation graph T .

Lemma 3.1.2. *Let $\{R, R'\}$ be an edge of the graph T_0 . $\{R, R'\}$ is useless if and only if $\delta_T(R, R') < |RR'|$, where T is the transportation graph.*

Proof. T is a subgraph of T_0 , and $\{R, R'\}$ is an edge of T_0 , so it follows that $\delta_T(R, R') \leq |RR'|$. If $\delta_T(R, R') = |RR'|$ then the edge $\{R, R'\}$ corresponds to a shortest path from R to R' in T_0 thereby making $\{R, R'\}$ useful. If $\delta_T(R, R') < |RR'|$ then $\{R, R'\}$ is useless since any

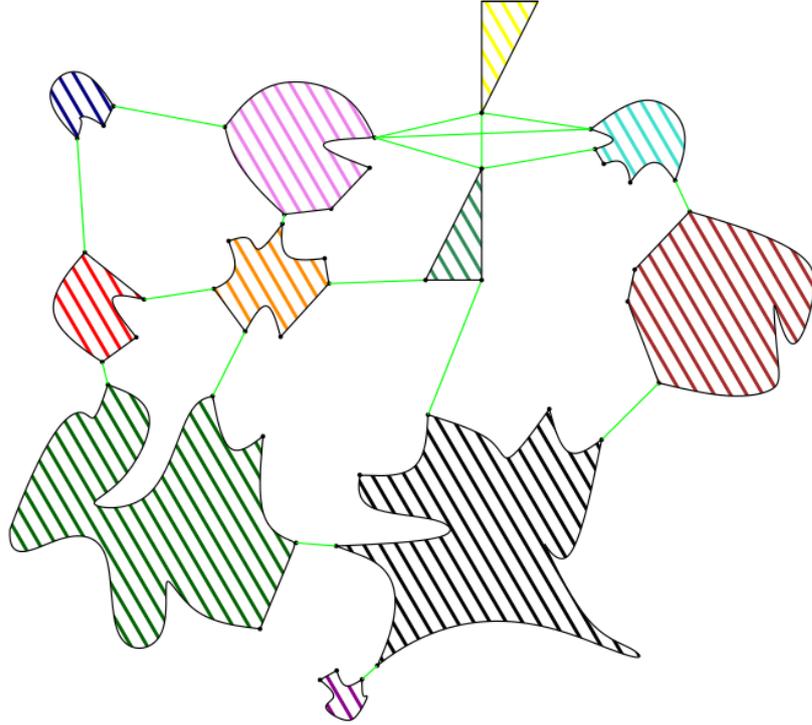


Figure 3.4: The transportation graph T induced by a set of pairwise disjoint regions.

shortest path in T_0 using $\{R, R'\}$ can be made shorter by replacing $\{R, R'\}$ with a shortest path in T (made of edges of T_0), \square

The following lemma states that the weight of a shortest path in T cannot exceed the Euclidean distance between R and R' .

Lemma 3.1.3. *For any two regions $R, R' \in \mathcal{R}$, $\delta_T(R, R') \leq |RR'|$.*

Proof. For sake of a contradiction, let R, R' be regions with shortest $|RR'|$ such that $\delta_T(R, R') > |RR'|$. $\{R, R'\}$ is not an edge of T_0 ; otherwise $\{R, R'\}$ would be a path from R to R' of weight $|RR'|$ which contradicts our assumption $\delta_T(R, R') > |RR'|$. Therefore, we must have that $L_{RR'}$ intersects at least 1 region R'' other than R and R' in some point p . Notice that

$$|RR''| + |R''R'| \leq |Rp| + |pR'| = |RR'|.$$

And so

$$\begin{aligned} \delta_T(R, R') &\leq \delta_T(R, R'') + \delta_T(R'', R') \quad \text{by def. of shortest path} \\ &\leq |RR''| + |R''R'| \quad |RR'| \text{ is the shortest counter example} \\ &\leq |RR'| \text{ which contradicts our initial assumption} \end{aligned}$$

\square

In the previous chapter we introduced the notion of a spanner for points. We now generalize this notion for regions.

A *region graph* G on \mathcal{R} is a graph whose vertex set is \mathcal{R} and a set of edges $E \subseteq \mathcal{R} \times \mathcal{R}$ such that each edge $\{R, R'\}$ has one endpoint on ∂R , one endpoint on $\partial R'$, and does not intersect any other regions in $\mathcal{R} \setminus \{R, R'\}$. In this thesis we assume that the edges corresponds to line segments since we are approximating shortest distances (in the Euclidean sense) between regions. Notice that we also allow for several edges between two regions. However, we also require the endpoints of each edge $\{R, R'\}$ of a region graph to be given as part of the description of the graph and that a region graph can be uniquely embedded in the plane. Notice that by definition both T_0 and T are region graphs.

Every edge of G is given a weight corresponding to the Euclidean

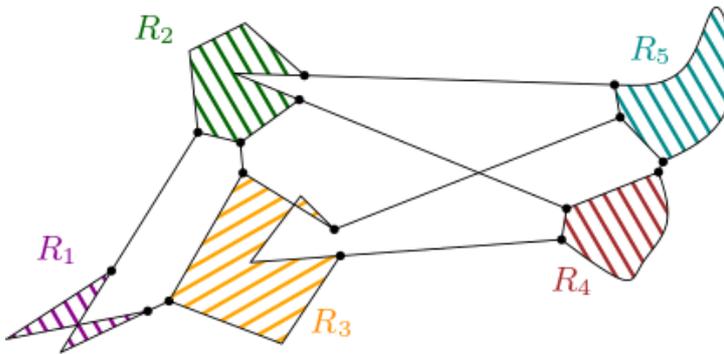


Figure 3.5: A region graph for the set of regions $\mathcal{R} = \{R_1, R_2, R_3, R_4, R_5\}$

distance between its endpoints.³ Given a region graph G , the cost of travelling in G between two regions R and R' corresponds to the sum of the weights of the edges defining a shortest path from R to R' . We denote this cost by $\delta_G(R, R')$. If there are no paths in G between two regions R and R' then we set $\delta_G(R, R') = \infty$.

Let $t \geq 1$ and let G be a region graph on a set of regions \mathcal{R} . G is a t -spanner of the transportation graph T if for all $R, R' \in \mathcal{R}$

$$\delta_G(R, R') \leq t \cdot \delta_T(R, R')$$

Importantly, notice that we do not require for a spanner of T to be a subgraph of T . Lemma 3.1.3 entails that T is 1-spanner of the complete region graph T_0 . Lemma 3.1.4 states a sufficient condition for a region graph to be a spanner of the transportation graph T .

Lemma 3.1.4. *Let G be a region graph on a set of disjoint regions \mathcal{R} . If for every edge $\{R, R'\}$ of T we have $\delta_G(R, R') \leq t \cdot |RR'|$ for some $t \geq 1$, then G is a t -spanner of T .*

³ With respect to how we introduced the problem in Chapter 1, the weight of an edge of G corresponds to the cost of travelling on that edge.

Proof. Suppose that for every edge $\{R, R'\}$ in T we have $\delta_G(R, R') \leq t \cdot |RR'|$. Let R and R' be two distinct regions of \mathcal{R} and consider a sequence of regions $R = R_1, R_2, R_3, \dots, R_k = R'$ which describes a shortest path in T from R to R' . By definition of T , every edge $\{R_i, R_j\}$ is useful which is to say that $|R_i R_j| = \delta_T(R_i, R_j)$ by Lemma 3.1.2. And so we have

$$\delta_G(R, R') \leq \sum_{i=1}^{k-1} \delta_G(R_i, R_{i+1}) \leq t \cdot \sum_{i=1}^{k-1} |R_i R_{i+1}| = t \cdot \delta_T(R, R')$$

as desired. □

We shall say that a region graph G on a set of n disjoint regions is sparse if the total number of edges of G is $\mathcal{O}(n)$. Our focus throughout this thesis consists of understanding how we can obtain sparse region spanners.

EXISTENCE OF SPARSE SPANNERS OF T

In this chapter we show that there exists families of sparse t -spanners of the transportation graph induced by a collection of n disjoint regions \mathcal{R} for $t > 1$

As mentioned in Chapter 2, sparseness results for both the Yao and θ graphs have been proven for sets of points. We generalize these methods for sets of disjoint regions.

One question arises naturally: where should we put our cones? By Lemma 3.1.1 we know that the shortest line segment $L_{RR'}$ between two regions R and R' has two endpoints: one on the boundary of R and the other on the boundary of R' . In particular, we can associate to each region R a set of at most $n - 1$ points on the boundary of R where each point corresponds to an endpoint of $L_{RR'}$ for some region $R' \neq R$. Therefore, by partitioning the plane with cones around each of these *boundary points* for every region R in \mathcal{R} we guarantee that any line segment $L_{RR'}$, which is to say every edge of the transportation graph T , is contained in at least 1 cone.

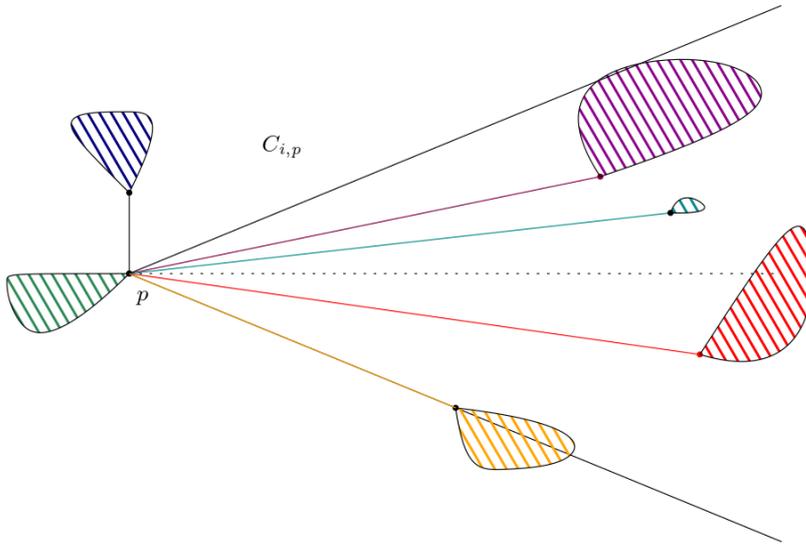


Figure 4.1: The selection process in the construction of a Yao graph Y_k . The yellow edge is added to Y_k since it is the shortest edge with p as an endpoint contained in $C_{i,p}$

More formally, let $\theta > 0$ be an angle such that $2\pi/\theta = k$ is an integer. We define the region graph Y_k as follows: for all points p corresponding to the endpoint of some line segment $L_{RR'}$, we cover the plane with k cones of angle θ with apex centered on p .

For each cone C with apex p we select, among all edges of T contained in C with p as an endpoint, the one whose weight is minimum.

In the upcoming Theorem we will make use of the following Lemma, whose proof can be found in [3]:

Lemma 4.1.1. *Assume that $\theta < \pi/3$ and let a, b, c be three distinct points in the plane. If $|ca| \leq |cb|$, then $|ab| \leq |cb| - (1 - 2 \sin(\theta/2))|ca|$.*

Theorem 4.1.2. *Let \mathcal{R} be a set of n disjoint regions in the plane and let T be the transportation graph corresponding to \mathcal{R} . For any $0 < \theta < \pi/3$, Y_k is a spanner of T for $t = 1/(1 - 2 \sin(\theta/2))$.*

Proof. By definition Y_k is indeed a region graph since Y_k is a subgraph of T . By Lemma 3.1.4 we only need to show that for every edge $\{R, R'\}$ of T

$$\delta_{Y_k}(R, R') \leq t \cdot |RR'|$$

To achieve this we proceed by induction on the rank of the length $|RR'|$ in the sequence of all edges of T , sorted by non-decreasing order.

Base Case: $\{R, R'\}$ is an edge of T of minimum rank. It follows that $\{R, R'\}$ is the edge of minimum weight $|RR'|$ in at least 1 cone. That is to say that $\{R, R'\}$ is an edge of Y_k . And so $\delta_{Y_k}(R, R') = |RR'|$ as desired.

Inductive Step: Let $\{R, R'\}$ be an edge of T of rank at least 2. We suppose that

$$\delta_{Y_k}(R_1, R_2) \leq t \cdot |R_1R_2|$$

for every edge $\{R_1, R_2\}$ in T with $|R_1R_2| < |RR'|$.

If $\{R, R'\}$ is an edge of Y_k then $\delta_{Y_k}(R, R') = |RR'|$ and the Theorem holds. Assume that $\{R, R'\}$ is not an edge of Y_k . Let C be the cone with apex centered on the boundary point of R , call it r , that contains $\{R, R'\}$.

Since $\{R, R'\}$ is not an edge of Y_k the cone C must contain another edge of T of smaller weight. Thus there is at least 1 region R'' such that the edge $\{R, R''\}$ is contained in C and $|RR''| < |RR'|$. Let us denote by r'' (respectively r') the endpoint of $L_{RR''}$ ($L_{RR'}$) on the boundary of R'' (R'). Since $|rr''| < |rr'|$, By Lemma 4.1.1 this implies:

$$|r'r''| \leq |rr'| - (1 - 2 \sin(\theta/2))|rr''|$$

Noting that $|R'R''| \leq |r'r''|$, $|RR''| = |rr''|$ and that $|rr'| = |RR'|$ the above can be rewritten as

$$\begin{aligned} |R'R''| &\leq |r'r''| \\ &\leq |rr'| - (1 - 2 \sin(\theta/2))|rr''| \\ &= |RR'| - (1 - 2 \sin(\theta/2))|RR''| \quad (1) \end{aligned}$$

Combining (1) and Lemma 3.1.3 yields

$$\delta_T(R', R'') \leq |R'R''| < |RR'|$$

In particular, each edge of a shortest R'' - R' path $R'' = R_1, \dots, R_j = R'$ in T from R'' to R' is of smaller rank than $\{R, R'\}$.¹ And so, by our induction hypothesis, we have that

$$\delta_{Y_k}(R_i, R_{i+1}) \leq t \cdot \delta_T(R_i, R_{i+1}) \quad 1 \leq i \leq j-1$$

Just as in the proof of Lemma 3.1.4, summing over all i yields:

$$\delta_{Y_k}(R'', R') \leq t \cdot \delta_T(R'', R') \quad (2)$$

Now consider the path in Y_k that consists of first travelling from R to R'' via $\{R, R''\}$ and then from R'' to R' by taking a shortest path from R'' to R' . Setting $f(\theta) = (1 - 2 \sin(\theta/2))$, the following must hold:

$$\begin{aligned} \delta_{Y_k}(R, R') &\leq |RR''| + \delta_{Y_k}(R'', R') \quad \text{by def. of shortest path} \\ &\leq |RR''| + t \cdot \delta_T(R'', R') \quad \text{by (2)} \\ &\leq |RR''| + t \cdot |R'R''| \quad \text{by Lemma 3.1.3} \\ &\leq |RR''| + t(|RR'| - f(\theta)|RR''|) \quad \text{by (1)} \\ &= t \cdot |RR'| + (1 - t \cdot f(\theta))|RR''| \end{aligned}$$

Thus, when $t = 1/f(\theta)$, we have:

$$\delta_{Y_k}(R, R') \leq t \cdot |RR'|$$

as desired. \square

How good is the spanner Y_k ? We know that any spanner on a set of n points requires at least $n - 1 \in \Omega(n)$ edges by Theorem 2.1.1. In the construction of Y_k we are picking at most 1 edge of T for each cone. The total number of cones depends on the total number of distinct boundary points: if every edge of T yields two unique boundary points, then the total number of edges is $\mathcal{O}(n^2)$. As shown in Figure 4.2, consider a set S of n points where $n/2$ of these points are spaced by distance ϵ on a horizontal line H_1 and the remaining $n/2$ points are spaced by distance ϵ on another horizontal line H_2 at distance ϵ from H_1 for some $\epsilon > 0$.

The transportation graph on this set of points contains the complete bipartite graph where one part corresponds to the points on H_1 and the other corresponds to the points of H_2 , which has $\mathcal{O}(n^2)$ edges. We remark that there exists a sufficiently small positive real number ρ such that the transportation graph induced by forming a disk of radius ρ around each point in S contains the complete bipartite graph on the resulting set of disks. This results in Y_k having $\mathcal{O}(n^2)$ edges for

¹ Notice that we are implicitly showing that Y_k is connected simultaneously. A more formal proof could be done independently. The argument would be the same and for that reason we omit it. Consequently at least 1 path from R'' to R' does exist in Y_k .

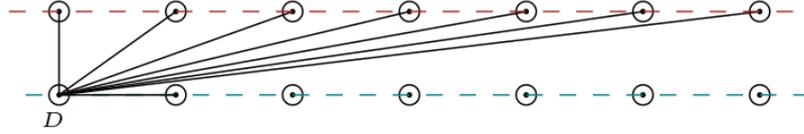


Figure 4.2: Each edge incident to the center of D in the complete bipartite graph induces an edge of the transportation graph T on the corresponding set of disks. This is true for the center of each disk, resulting in $\mathcal{O}(n^2)$ edges.

any k on this set of disks.

Nonetheless, we remark that if there are enough line segments $L_{RR'}$ sharing endpoints then the number of edges in Y_k can be of smaller order than n^2 . Namely we have the following:

Corollary 4.1.3. *Let \mathcal{R} be a set of n disjoint closed and bounded regions in the plane and let T be the transportation graph induced by \mathcal{R} . Let M be the total number of distinct endpoints of the line segments $L_{RR'}$, $R, R' \in \mathcal{R}$. The number of edges is $\mathcal{O}(M)$. For any $0 < \theta < \pi/3$, if $M = \mathcal{O}(n)$, then Y_k is a sparse t -spanner of T for $t = 1/(1 - 2 \sin(\theta/2))$.*

For instance if each region is a polygon with at most a constant number of vertices, then the graph Y_k is sparse. Similarly, if each region is a polygon and the total number of vertices of all n regions is $\mathcal{O}(n)$, then the graph Y_k is sparse.

Our approach from Chapter 4 has the following two problems:

- 1) The data structures for computing Yao graphs on sets of points are more difficult to implement than the ones used for computing θ_k graphs.
- 2) Y_k requires to already have computed all the edges of T since we are building a subgraph of T . To the best of our knowledge it is unclear whether fast algorithms can perform this task. However, when the regions are polygons we will see that this is no longer an issue.

In this chapter we construct spanners of the transportation graph when the set of regions corresponds to a collection of non-intersecting polygons. We first introduce a collection of region graphs θ_k and show that it yields spanners of the transportation graph T . We then design an algorithm that generates sparse spanners of T for sets of disjoint polygons if these polygons do not have too many vertices.

Let \mathcal{R} be a set of regions and let θ be a cone angle such that $2\pi/\theta = k$ is an integer. We formally define the region graph θ_k as follows:

- The vertex set of θ_k is \mathcal{R} .
- For each cone $C_i \in \mathcal{C}_k$, for each region R , and for each point p of R corresponding to an endpoint of one of the edges of the transportation graph T ; if $C_{i,p}$ is intersected by at least one region different than R , then θ_k contains the edge $\{p, q\}$ where q is the point in $C_{i,p} \cap \mathcal{R} \setminus \{R\}$ whose orthogonal projection onto $bis(C_{i,p})$ is closest to p .

Informally, as shown in Figure 5.1 below, for each cone $C_{i,p}$ a line orthogonal to $bis(C_{i,p})$ is swept starting at p . It stops as soon as it hits some point of $C_{i,p} \cap \mathcal{R} \setminus \{R\}$ and an edge is added between p and (one of) the first point that was hit.

The proof that this definition of θ_k gives rise to a family of spanner of T is analogous to the one exhibited in Theorem 4.1.2. More specifically we proceed by induction on the rank of the length $|RR'|$ in the sequence of all edges of T sorted in non-decreasing order. However, Lemma 4.1.1 cannot be applied in the induction step because the weight of an edge is not necessarily reflected by its orthogonal

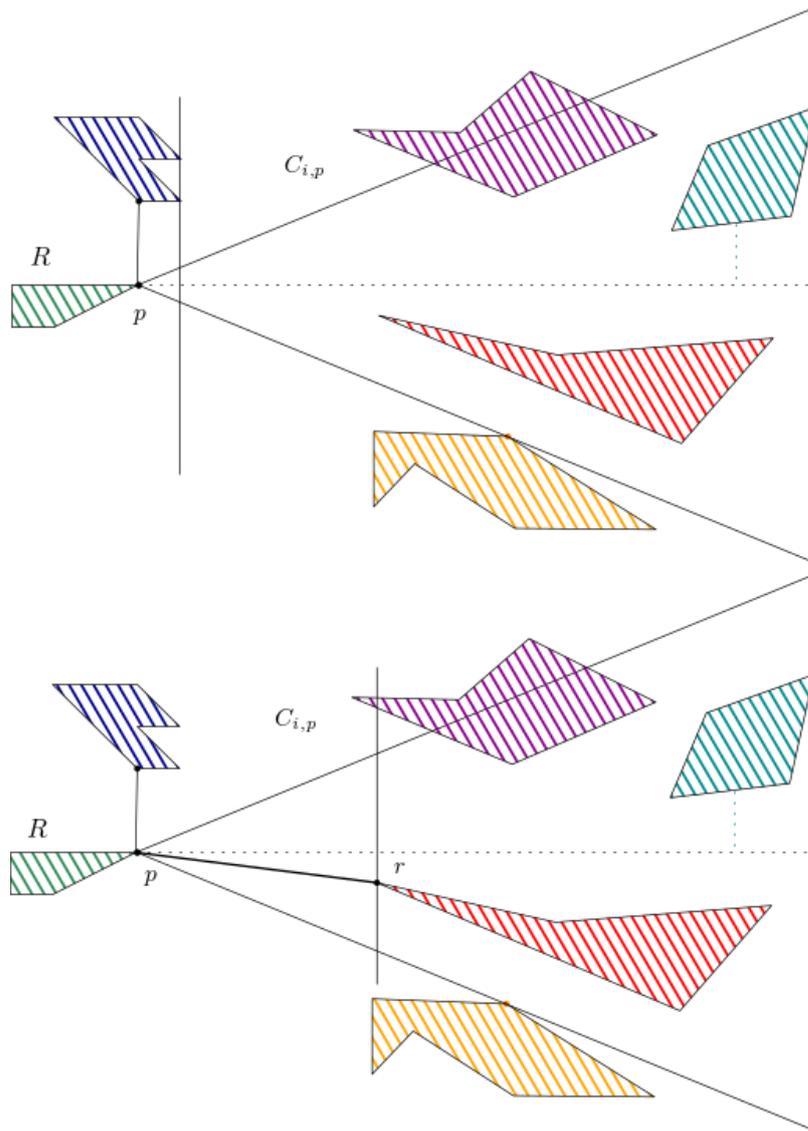


Figure 5.1: The selection process in the construction of θ_k .

projection. We resort instead to the following lemma whose proof can be found in [2]

Lemma 5.1.1. *Let $k > 8$ be an integer, let $\theta = 2\pi/k$ and let p, q be distinct points of the plane. Denote by C_p the cone with apex centered on p of angle θ containing q . Let r be any point in C_p such that the orthogonal projection of r onto the bisector of the cone C_p is at least as close as the projection of q . Then:*

1. $|pr| \cos \theta \leq |pq|$, and
2. $|rq| \leq |pq| - (\cos \theta - \sin \theta)|pr|$

The full details of the proof of Theorem 5.1.2 can be found in the appendix.

Theorem 5.1.2. *For any $0 < \theta < \pi/4$, θ_k is a t -spanner of the transportation graph T for $t = 1/(\cos \theta - \sin \theta)$.*

Corollary 5.1.3. *Let \mathcal{R} be a set of n pairwise disjoint regions in the plane and let T be the transportation graph corresponding to \mathcal{R} . Let M be the total number of distinct endpoints of all the line segments $L_{RR'}$. For any $0 < \theta < \pi/4$, θ_k is a spanner of T for $t = 1/(\cos \theta - \sin \theta)$ and with at most $\mathcal{O}(M)$ edges.*

We are still facing the obstacle outlined earlier in this chapter: constructing this θ_k depends on having prior knowledge of the edges of T . We now proceed to resolve this issue.

So far we have restrained ourselves from specifying any constraint on the regions. Having some extra information on the shape of the regions is relevant. For instance when \mathcal{R} is a set of points, the endpoints of the edges of T are part of the input: they correspond exactly to the points in \mathcal{R} . Hence we can build a θ_k spanner of T by building a θ graph on the points of \mathcal{R} as described in Chapter 2 without the need of generating T . We now proceed to show that, under the assumption that \mathcal{R} is a set of disjoint polygons, we can again use existing algorithms to produce spanners of T without having to compute T beforehand.

By Lemma 3.1.1 we know that the shortest line segment $L_{RR'}$ connecting two regions has one endpoint on the boundary of R and the other one the boundary of R' . When the regions are polygons we can view the boundary of a polygon P as a cycle of line segments. For a polygon P , we denote this cycle by $Seg(P)$. Hence for two disjoint polygons P and P' , the line segment $L_{PP'}$ connects a line segment $S \in Seg(P)$ to another line segment $S' \in Seg(P')$. We denote by $Seg(\mathcal{R})$ the set of line segments describing the boundaries of the polygons in \mathcal{R} .

In the upcoming proofs we will make use of the notion of *convexity*: a set $S \subset \mathbb{R}^d$ is convex if for any pair of points p, q in S the line segment L_{pq} joining p and q is contained in S .

Lemma 5.1.4. *Let S and S' be two non-crossing non-parallel line segments in the plane. At least 1 endpoint of $L_{SS'}$ corresponds to some endpoint of either S or S' .¹*

Proof. Let S, S' be non-crossing non-parallel line segments such that both endpoints of $L_{SS'}$ do not correspond to any of the endpoints of S and S' . Let s and s' denote respectively the points of S and S' for which $|SS'|$ is minimum. The ball $B_{s'}$ centered on s' of radius $|SS'|$ only intersects S on the point s : Suppose there was another point $x \neq s$

¹ Note that when two segments S and S' are parallel, we enforce $L_{SS'}$ to emanate from at least one endpoint of either S or S' .

of S intersecting this ball. Then S would contain the line segment L_{sx} between s and x . But since $B_{s'}$ is convex, L_{sx} would also be contained in $B_{s'}$. In particular, there would be a point y on L_{sx} intersecting the interior of $B_{s'}$. This would entail that $|ys'| < |ss'|$ which would contradict the minimality of $|ss'|$. Thus the line segment S is tangent to $B_{s'}$ and so $L_{SS'}$ is orthogonal to S . The same argument can be made for S' i.e. $L_{SS'}$ is orthogonal to S' . This entails that S and S' are parallel which contradicts our initial assumption. \square

As a result of Lemma 5.1.4, by placing the C_i cones on each endpoint of the line segments of $Seg(\mathcal{R})$, each line segments $L_{PP'}$ is contained in (at least) one of these cones for all polygons P, P' in \mathcal{R} . Moreover these endpoints are part of the input: they are the vertices of the polygons in \mathcal{R} . Although this is a good step towards not having to compute T it is not enough. Namely, we need to describe how to identify, out of all line segments in $Seg(\mathcal{R})$ intersecting a cone C , the one whose orthogonal projection onto $bis(C)$ is closest to the apex of C .

Let p be a vertex of a polygon P in \mathcal{R} . Let $C_{i,p}$ be one of the k cones of angle θ with apex p . Let S be one of the line segments of $Seg(P)$ with p as an endpoint. As shown in Figure 5.2, let $S' \neq S$ be a line segment of $Seg(\mathcal{R})$ intersecting $C_{i,p}$ and let q be a point of $S' \cap C_{i,p}$ whose orthogonal projection onto $bis(C_{i,p})$ is closest to p . If $bis(C_{i,p})$ is orthogonal to the line induced by S' then any point of S' yields the same orthogonal projection. In particular, q can be any point of $S' \cap C_{i,p}$.

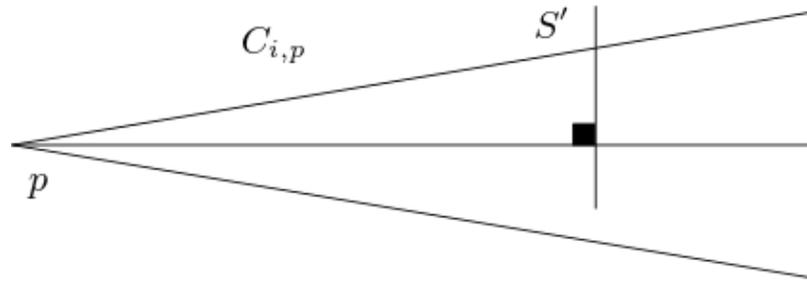


Figure 5.2

Otherwise we have the following:

Lemma 5.1.5. *If $bis(C_{i,p})$ is not orthogonal to the line induced by S' , then q corresponds to either one of the endpoints of S' or the intersection point of S' and one of the rays defining the boundary of $C_{i,p}$.*

Proof. Suppose that q is not an endpoint of S' . As shown in Figure 5.3 consider the line ℓ passing through q that is orthogonal to $bis(C_{i,p})$.

Since $bis(C_{i,p})$ is not orthogonal to the line induced by S' , q is the

only point of $S' \cap \ell$. In particular, one of the endpoints of S' , call it x , lies to the left of ℓ .² Now consider the line segment $L_{qx} \subset S'$. Notice that $L_{qx} \cap C_{i,p} = \{q\}$: if this was not the case, then there would be a point of L_{qx} that is to the left of ℓ and contained in $C_{i,p}$ and this would contradict the minimality of q . By convexity of $C_{i,p}$, it follows that q must be on the boundary of $C_{i,p}$. In other words, q is the intersection point of one of the rays defining the boundary of C with S' . \square

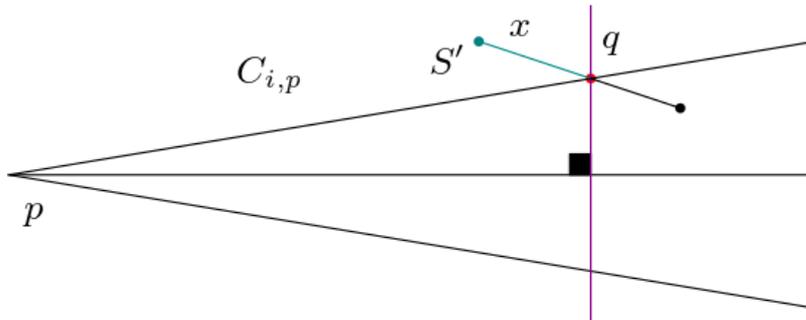


Figure 5.3: In red we have the point q and in purple we have the line orthogonal to $bis(C_{i,p})$. In blue we have x and the portion of S corresponding to L_{xq} .

It follows that in order to compute a point q for one of the cones $C_{i,p}$ it is sufficient to proceed as follows:

1. Seek, among all the endpoints contained in $C_{i,p}$ of the line segments of $Seg(\mathcal{R})$, the one whose projection onto $bis(C_{i,p})$ is closest to p .
2. Out of all the points of intersection (if any) of the boundary of $C_{i,p}$ with the line segments of $Seg(\mathcal{R})$, report the two (one for each ray defining $C_{i,p}$) closest to p .
3. Out of the at most 3 reported points in the first two steps, select the one whose projection onto $bis(C_{i,p})$ is closest to p .

In Chapter 2 we have seen that the θ graph of the endpoints of $Seg(\mathcal{R})$ is constructed using the first step. Moreover, the endpoints of the segments of $Seg(\mathcal{R})$ are given as part of the input: they are the vertices of the polygons in \mathcal{R} .

The second step reports for each ray defining the boundary of one

² The notion of being to the left of a line L when looking at $C_{i,p}$ refers to the coordinate system obtained by translating the origin to p and rotating the positive x -axis to $bis(C_{i,p})$

of these cones $C_{i,p}$ its point of intersection (if there is one) with the segments in $Seg(\mathcal{R})$ that is closest to the apex of the cone that ray defines. This can be computed in $\mathcal{O}(n \log n + k \log n)$ time and $\mathcal{O}(n)$ space where k is the number of intersection points using plane sweep (see [17]). In Figure 5.4 below we extended one of the legs of C into a line L . If we were to sweep this line through the plane our problem now becomes the following: upon reaching p we want to report the point q on L that lies "above" p .

Recall from Chapter 2 that each cone $C_{i,p}$ has a fixed orientation. Let us denote by r_0, \dots, r_{k-1} the k rays emanating from the origin and outlining the boundary of the cones C_i . More specifically, the boundary of C_i is given by the rays r_i and r_{i+1} (note that $r_k = r_0$). For every vertex p of the polygons in \mathcal{R} output the first segment of $Seg(\mathcal{R})$ hit by the ray $r_i + p$ for every $0 \leq i \leq k - 1$. We denote this segment by $S(p, i)$. Since the slope of each ray r_i is predetermined, for each r_i a line L_i parallel to r_i will serve as our sweep line. Our task is to maintain through out each sweep a list of all the segments of $Seg(\mathcal{R})$ intersecting the sweep line L_i in sorted order of their intersection along the sweep line with respect to the orientation of r_i .

In [17] they outline how to maintain such a list by storing the segments in a balanced binary search tree (BST). This BST is commonly referred to as the *sweep status*. Let L_i be one of the sweep lines. As L_i sweeps the plane, every time L_i reaches a vertex p of some polygon in \mathcal{R} we want to compute $S(p, i)$. In particular, each vertex p defines an *event* during the sweep. The order in which each endpoint of $Seg(\mathcal{R})$ is hit by a sweep line L_i is stored in the list $EventQueue_i(\mathcal{R})$. $S(p, i)$ is obtained by looking at the successor of p in the BST storing the segments of $Seg(\mathcal{R})$ at the moment L_i reaches p during the sweep. For the sake of convenience we will assume that no two events are hit at the same

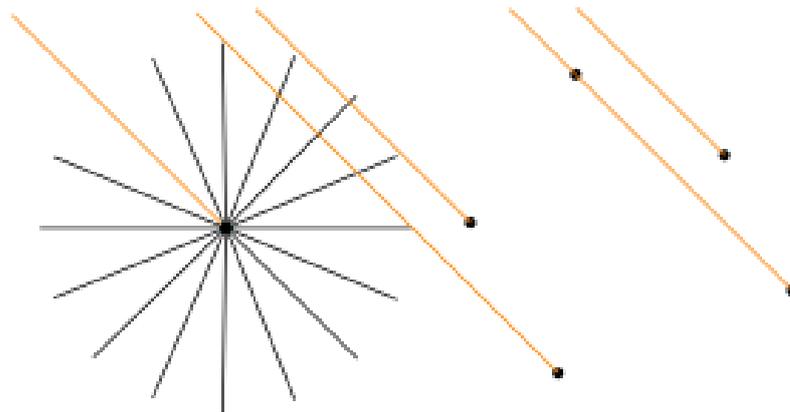


Figure 5.4: In orange we have r_6 for $\theta = \pi/8$. For each vertex p we want to report the first intersection on $r_6 + p$

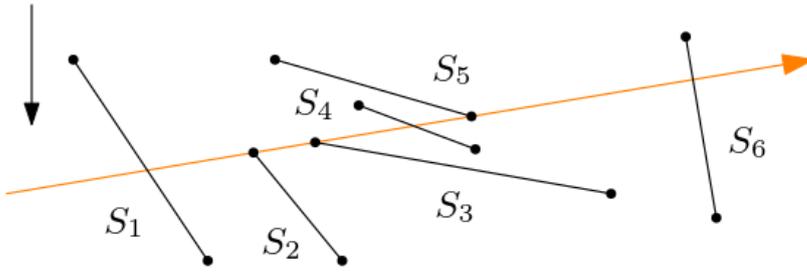


Figure 5.5: In orange we have a sweep line oriented with respect to the ray it corresponds to. At that moment the leaves in the BST storing the segment are from left to right $S_1, S_2, S_3, S_4, S_5, S_6$ which corresponds to the order of their intersection as we are walking in the direction of L_i .

time by one of the sweep lines. The details for when this assumption is not made can be found in the appendix A.2.

For each segment we also determine which of its endpoint is hit first by the sweep line. We store and sort with respect to the same ordering in $EventQueue_i(\mathcal{R})$ all of these endpoint in a list $First_i(\mathcal{R})$. Similarly, we define the list $Second_i(\mathcal{R})$: if some vertex p is in $Second_i(\mathcal{R})$ then p is the endpoint of some segment S of $Seg(\mathcal{R})$ that is hit second (compared to the other endpoint of S). Notice that each vertex p is the endpoint of exactly two segments of $Seg(\mathcal{R})$.

We shall refer to the routine described above by `FINDINTERSECTION`.

Lemma 5.1.6. *For any vertex p and for any ray r_i , `FINDINTERSECTION` reports a successor to p when L_i reaches p if and only if there is an intersection to be detected on the ray $r_i + p$.*

Proof. Let r_i be one of the k rays. We proceed by induction on the rank of the vertices in the list $EventQueue_i(\mathcal{R})$.

Base Case: The vertex p of rank 1 in the list $EventQueue_i(\mathcal{R})$ is the point whose scalar component on the line $L_i + \pi/2$ is maximal. In particular, this point has no successor $S(p, i)$ in the BST and the ray $r_i + p$ has no intersection.

Inductive Step: Let p be a vertex of rank $n \geq 2$ in $EventQueue_i(\mathcal{R})$. (I.H.) We suppose that all vertices of rank $k < n$ in $EventQueue_i(\mathcal{R})$ were treated correctly. We show that vertex p is also treated correctly which is to say that `FINDINTERSECTION` reports a successor to p when L_i reaches p if and only if there is an intersection on the ray $r_i + p$.

Suppose that `FINDINTERSECTION` reports a successor to p when L_i reaches p . By definition of `FINDINTERSECTION`, more precisely of how the line segments are maintained in the sweep status, this successor lies on the portion of the line L_i corresponding to the ray $r_i + p$. Con-

versely, suppose that there is an intersection point q to be detected on the ray $r_i + p$. By the assumption that no two vertices are hit at the same time during the sweep, q is not an endpoint of any segment of $\text{Seg}(\mathcal{R})$. In particular, the segment S_q it corresponds to was correctly added to the BST prior to reaching p by our induction hypothesis (I.H) since the "first" endpoint of S_q was of rank lower than p in $\text{EventQueue}_i(\mathcal{R})$ (one of the endpoints of S_q is "above" the line going through p and q). \square

We can now describe our algorithm

Algorithm POLYGONSPANNER

Input. A set \mathcal{R} of n disjoint polygons, an angle $\theta = 2\pi/k$

Output. T_θ a spanner of the transportation graph T of the set \mathcal{R}

1. Build the θ_k graph on the set of vertices of the polygons in \mathcal{R} : For each vertex p and for $0 \leq i \leq k-1$ store as $x_{p,i}$ the point of $C_{i,p}$ that is selected (if any) in the construction of the θ_k graph.
 2. Run FINDINTERSECTION on the set of vertices of \mathcal{R}
 3. For all vertex p and for $0 \leq i \leq k-1$, compute and store the intersection of $S(p, i)$ with r_i as $r_{i,p}$
 4. For all vertex p and for all $0 \leq i \leq k-1$ select out of $x_{p,i}, r_{i,p}$ and $r_{i+1 \pmod k, p}$ the point whose orthogonal projection onto $\text{bis}(C_{i,p})$ is minimal and store it as p_i . In case of ties select the point that was hit first in the line sweep.
 5. if p and p_i belong to the same polygon, then do nothing
 6. otherwise add the edge $\{P_p, P_{p_i}\}$ to the set of edges of T_θ with endpoints p and p_i between the polygons P_p, P_{p_i} of \mathcal{R} containing p and p_i respectively.
-

Lemma 5.1.7. *The graph T_θ generated by POLYGONSPANNER is a region graph.*

Proof. Let $\{P, P'\}$ be one of the edges of T_θ with endpoints p and p' . By construction p and p' lie on the boundaries of P and P' . Without loss of generality suppose that p' was the point whose orthogonal projection onto $\text{bis}(C_{i,p})$ was closest to p for some cone $C_i \in \mathcal{C}_\parallel$. Note that the line segment $L_{pp'}$ is fully contained in $C_{i,p}$. Suppose that $L_{pp'}$ intersects some other polygon P'' . Then we know that $L_{pp'}$ intersects $\text{Seg}(P'')$ in some point p'' : otherwise $L_{p,p'}$ is contained in P'' which contradicts our assumption that each polygon is pairwise disjoint. Therefore, p'' is in $C_{i,p}$. In particular, the projection of p'' onto $\text{bis}(C_{i,p})$ is closer to p than the one of p' . Thus, POLYGONSPANNER would not add the edge $L_{pp'}$ to T_θ which contradicts our initial assumption and so $\{P, P'\}$ does not intersect any other polygon of \mathcal{R} . \square

Theorem 5.1.8. *Given a set \mathcal{R} of n disjoint polygons `POLYGONSPANNER` generates a spanner of the transportation graph T induced by \mathcal{R} for $t = 1/(\cos \theta - \sin \theta)$, $0 < \theta < \pi/4$.*

Proof. The correctness of `POLYGONSPANNER` follows from the correctness of the algorithm building the θ_k graph on the vertices of the polygons in \mathcal{R} together with Lemma 5.1.6. By Lemma 5.1.7 we know that T_θ is a region graph. Consequently to show that T_θ is a spanner of T by Lemma 3.1.4 it is sufficient to show that for every edge $\{P, P'\}$ of T

$$\delta_{T_\theta}(P, P') \leq t \cdot |PP'|$$

We once again proceed by induction on the rank of the length $|PP'|$ in the sequence of all edges of T , sorted by non-decreasing order.

Base Case: $\{P, P'\}$ is an edge in T having minimum length $|PP'|$. By Lemma 5.1.4, $L_{PP'}$ emanates from at least one of the endpoints of a segment from either $Seg(P)$ or $Seg(P')$, in other words one of the vertex of P or P' . Without loss of generality, we can assume that P is the polygon with vertex p which $L_{PP'}$ emanates from. Let C be the cone with apex centered on p containing $L_{PP'}$. We denote by p' the point of (one of the segments of) P' in C whose projection onto $bis(C)$ is closest to p . Although $|PP'|$ is minimal, this does not necessarily entail that $\{P, P'\}$ is an edge of T_θ (see Figure A.1). However, if θ is sufficiently small it does. More specifically, if $\{P, P'\}$ is not an edge of T_θ there would be a point p'' on the boundary of some polygon P'' that lies in C whose projection onto $bis(C)$ was closer to p . And so by Lemma 5.1.1

$$\begin{aligned} |P'P''| &\leq |p'p''| \\ &\leq |pp'| - (\cos \theta - \sin \theta)|pp''| \end{aligned}$$

Noticing that $|PP'| = |pp'|$ and $|PP''| \leq |pp''|$ it follows that

$$|P'P''| \leq |PP'| - (\cos \theta - \sin \theta)|PP''| < |PP'|$$

since $\theta < \pi/4$. But this contradicts the minimality of $|PP'|$.

Inductive Step: Let $\{P, P'\}$ be an edge of T of rank $n \geq 2$. (I.H.) We suppose that $\delta_{T_\theta}(P_1, P_2) \leq t \cdot |P_1P_2|$ for every edge $\{P_1, P_2\}$ of T with rank $k < n$

Notice that if $\{P, P'\}$ is a edge of T_θ then $\delta_{T_\theta}(P, P') \leq |PP'|$ and the Theorem holds. So suppose that $\{P, P'\}$ is not a edge of T_θ . As we have done in the base case, we can assume that P is the polygon with vertex p which $L_{PP'}$ emanates from. Let C be the cone with apex centered on p that contains $L_{PP'}$.

Since $\{P, P'\}$ is not a edge of T_θ , there must be some polygon P'' which yields a point p'' contained in C whose projection onto $bis(C)$

was closer than the point p' of P' . Repeating the same argument from the base case, by Lemma 5.1.1 this entails that $|P'P''| < |PP'|$. And so $\delta_T(P', P'') \leq |P'P''| < |PP'|$. In particular, every edge on a shortest path $P'' = P_1, \dots, P_j = P'$ in T from P'' to P' has smaller weight than $|PP'|$ and thus is of smaller rank than $\{P, P'\}$. Hence by the induction hypothesis (I.H.):

$$\delta_{T_\theta}(P_i, P_{i+1}) \leq t \cdot |P_i P_{i+1}| \quad 1 \leq i \leq j-1$$

Summing over all j yields:

$$\delta_{T_\theta}(P'', P') \leq t \cdot \delta_T(P'', P') \quad (1)$$

Now consider the path in T_θ that consists of first travelling from P to P'' via the edge $L_{pp''}$ and then from P'' to P' . The following must hold:

$$\begin{aligned} \delta_{T_\theta}(P, P') &\leq |pp''| + \delta_{T_\theta}(P', P'') \quad \text{by def. of shortest path} \\ &\leq |pp''| + t \cdot \delta_T(P', P'') \quad \text{by (1)} \\ &\leq |pp''| + t \cdot |P'P''| \quad \text{by Lemma 3.1.3} \\ &\leq |pp''| + t \cdot |p'p''| \quad (2) \end{aligned}$$

But notice that by Lemma 5.1.1

$$|p''p'| + (1/t)|pp''| \leq |pp'|$$

In particular,

$$t|p''p'| + |pp''| \leq t|pp'| = t \cdot |PP'|$$

Combining the above with (2) we get:

$$\delta_{T_\theta}(P, P') \leq t \cdot |PP'|$$

as desired □

We know that building the θ_k graph on the set of M vertices of the polygons in \mathcal{R} requires $\mathcal{O}(M \log M)$ time and $\mathcal{O}(M)$ storage. What about FINDINTERSECTION?

For each sweep with respect to one of the lines L_i we first sort in the list $EventQueue_i(\mathcal{R})$ the vertices in non-decreasing order of their scalar component on $L_i + \pi/2$. Computing the scalar component of each vertex takes $\mathcal{O}(1)$ time and thus this sorting requires $\mathcal{R}(M \log M)$. The lists $First_i(\mathcal{R})$ and $Second_i(\mathcal{R})$ can then be built in $\mathcal{O}(M)$ time. Every time a vertex p is reached the time required for reporting $S(p, i)$ and updating the sweep status is $\mathcal{O}(\log M)$: this is insertion and deletion in a balanced binary search tree. Thus each sweep takes $\mathcal{O}(M \log M)$ time and FINDINTERSECTION takes $\mathcal{O}((1/\theta) \cdot M \log M)$ time where $\theta = 2\pi/k$. Moreover, FINDINTERSECTION requires $\mathcal{O}((1/\theta) \cdot M)$ storage; we store k successors for each vertex and the lists $EventQueue_i(\mathcal{R})$, $First_i(\mathcal{R})$, $Second_i(\mathcal{R})$ together with the sweep status require $\mathcal{O}(M)$ space. Summarising:

Corollary 5.1.9. *Given a set \mathcal{R} of n disjoint polygons `POLYGONSPANNER` generates a spanner of T for $t = 1/(\cos \theta - \sin \theta)$, $0 < \theta < \pi/4$, with $\mathcal{O}(M)$ edges in $\mathcal{O}(1/\theta \cdot M \log M)$ time and $\mathcal{O}(M)$ space where M is the total number of distinct endpoints of all the line segments $L_{pp'}$.*

Some final comments for this chapter:

We first remark that `POLYGONSPANNER` might add several edges between two polygons. In practice we would only like keep the shortest of these edges. This can be done after `POLYGONSPANNER` by traversing the resulting list of edges of T_θ in sorted non-decreasing order of their weight.

Through out this chapter we made the assumption that the polygons are *simple*: each polygon does not self intersect and do not contain any holes. However this is not strictly required: if a polygon does contain holes then one can think of its boundary as a union of disjoint cycles of line segments (as opposed to only one cycle). In particular, as long as the boundary of each polygon is given as part of the input, `POLYGONSPANNER` still generates spanners correctly.

As we have already commented at the end Chapter 4, the graph T_θ is, under the assumption that the total number of vertices M is linear in terms of the number of polygons, a "good" spanner: the number of edges of T_θ is linear in n . However we still have not found sparse spanners in the case M is not $\mathcal{O}(n)$. In other words, can we show the existence of sparse spanners whose number of edges does not depend on the number of boundary point? This is the problem tackled in Chapter 6.

6.1 INTRODUCTION

Previous work in motion planning problems ([16], [11]) demonstrates that knowing whether or not the geometric objects of interest are "fat" leads to faster algorithms. In this chapter we employ such a notion in order to improve Theorem 4.1.2 and Corollary 4.1.3: we show that sparse spanners on finite sets of pairwise disjoint convex regions exist regardless of the total number of boundary points, given that these regions satisfy some notion of fatness.

Let R be a region in the plane. We say R is f -fat for some $f \geq 1$ if there exist two concentric disks D_R^- and D_R^+ such that D_R^- is contained in R , D_R^+ contains R , and $\text{radius}(D_R^+) \leq f \cdot \text{radius}(D_R^-)$. We will denote by ρ_R^-, ρ_R^+ the radii of D_R^- and D_R^+ respectively. We will also denote by c_R the center of D_R^- (and D_R^+). Thus, R is f -fat for some $f \geq 1$ if $\rho_R^+ \leq f \cdot \rho_R^-$. We assume that the interior of R is non-empty. This implies that $\rho_R^- > 0$.

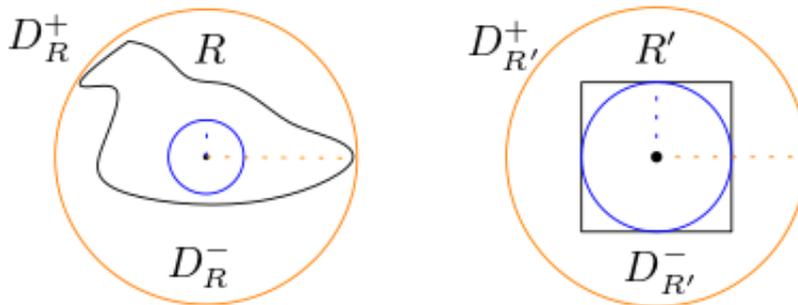


Figure 6.1: Two regions R and R' that are 4-fat and 2-fat respectively. In blue we have ρ_R^- and $\rho_{R'}^-$. In orange we have ρ_R^+ and $\rho_{R'}^+$.

For a given region R , we shall refer to D_R^- as the *inner disk* of R and D_R^+ as the *outer disk* of R .

Throughout sections 6.2 to 6.4, \mathcal{R} denotes a set of n pairwise disjoint f -fat regions for some $f \geq 1$. Let T be the transportation graph induced by \mathcal{R} . We direct the edges of T in the following way: edge $\{R, R'\}$ is oriented as (R, R') if $\rho_R^- \leq \rho_{R'}^-$.

Let λ be a real positive number. Each oriented edge (R, R') of the transportation graph is classified as follows:

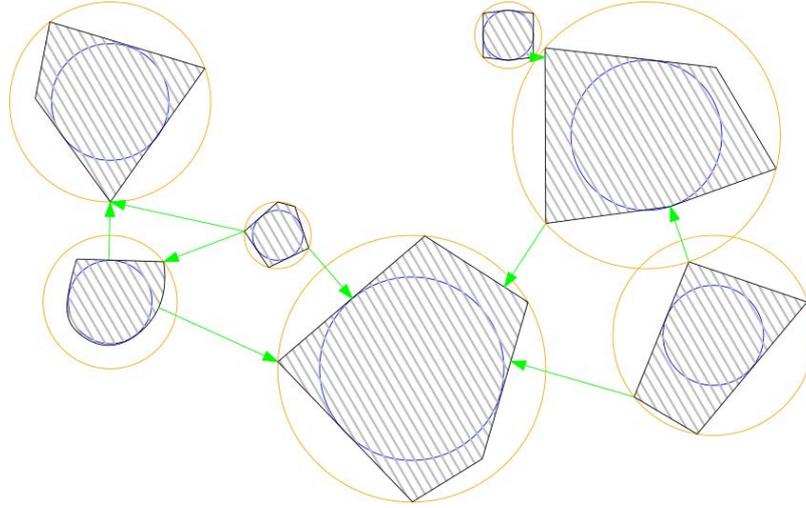


Figure 6.2: The oriented transportation graph T on a set of 2-fat regions.

- (R, R') is called *long* if $|RR'| \geq \lambda \cdot \rho_R^-$. Otherwise (R, R') is *short*.



Figure 6.3: A long edge between two regions. The regions are far apart compared to ρ_R^- .

- (R, R') is called *nasty* if (R, R') is short and $|D_R^- D_{R'}^-| > \lambda \cdot \rho_R^-$.

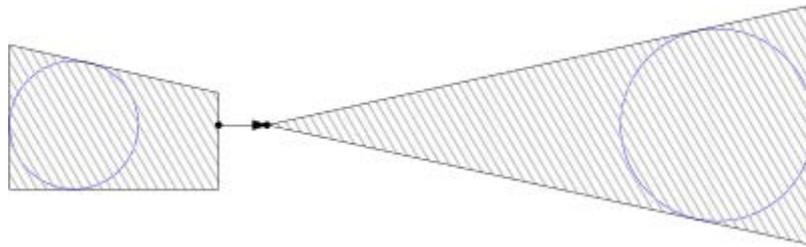


Figure 6.4: A nasty edge between two regions. Although the regions are close, their inner disks are far apart

- (R, R') is called *good* if (R, R') is short and $|D_R^- D_{R'}^-| \leq \lambda \cdot \rho_R^-$

In sections 6.2, we show an upper bound on the number of good edges in the transportation graph T that depends only on λ , regardless of whether or not the regions are convex. In Section 6.3, we show that the

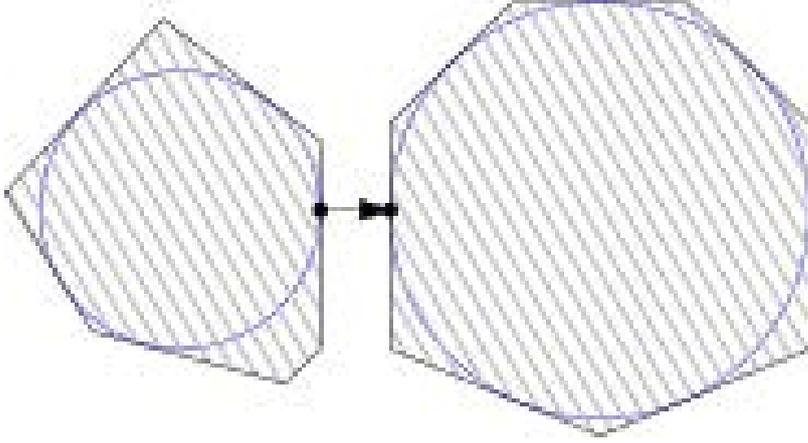


Figure 6.5: A good edge between two regions. The two inner disks are close.

number of nasty edges in T is bounded from above by a function of f and λ , under the assumption that the regions are convex. In Section 6.4, we introduce algorithm `PATHGREEDY` for regions. We show an upper bound on the number of edges of the t -spanner produced by this algorithm that depends only on t , f , and λ . Combining this bound and the ones obtained in section 6.2 and 6.3 entails that `PATHGREEDY` produces a sparse spanner for the case when the regions are convex and fat. Finally, in section 6.5 we present another definition of fatness introduced by *van der Stappen et al.* [16]. With respect to that definition, we show how `PATHGREEDY` generates a sparse spanners on sets of fat regions, regardless of the regions being convex.

6.2 BOUNDING THE NUMBER OF GOOD EDGES IN T

In this section we show that for a given region R , the number of good edges (R, R') in the transportation graph T is bounded from above by a function of λ .

Lemma 6.2.1. *Let R be a region of \mathcal{R} . The number of edges (R, R') of the transportation graph T such that $|D_R^- D_{R'}^-| \leq \lambda \cdot \rho_R^-$ is at most $(\lambda + 3)^2$.*

Proof. We know that $\rho_R^- \leq \rho_{R'}^-$. As shown in Figure 6.6, let u be the intersection of the boundary of $D_{R'}^-$ with the line segment $L_{c_R c_{R'}}$.

Let D'_0 be the disk of radius ρ_R^- whose center is the point c'_0 on $L_{c_R c_{R'}}$ at distance $\rho_R^- + |c_R u|$ from c_R . Note that this disk is contained in $D_{R'}^-$: for all $x \in D'_0$, $|c_{R'} x| \leq |c_{R'} c'_0| + |c'_0 x| \leq (\rho_{R'}^- - \rho_R^-) + \rho_R^-$. We have that

$$|c_R c'_0| = |c_R u| + |u c'_0| = 2 \cdot \rho_R^- + |D_R^- D_{R'}^-| \leq (\lambda + 2) \rho_R^-$$

Therefore, D'_0 is contained in the disk D'' of radius $(\lambda + 3) \rho_R^-$ centered on c_R . Thus, for each (R, R') of T such that $|D_R^- D_{R'}^-| \leq \lambda \cdot \rho_R^-$ we get such a disk D'_0 of radius ρ_R^- contained in D'' . Since the inner disks are pairwise disjoint, so are these disks D'_0 . Moreover all the

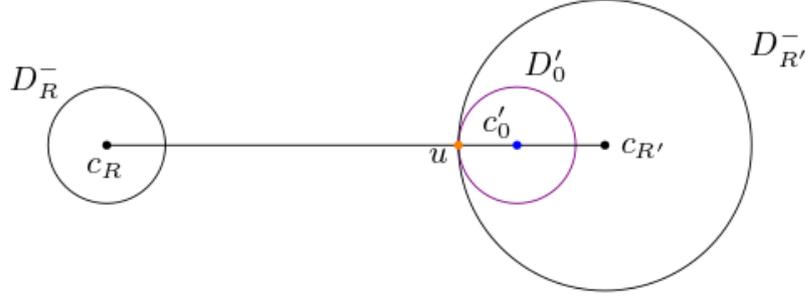


Figure 6.6: In orange we have the point u . D'_0 and c'_0 are shown in blue and purple respectively.

disks are contained in D'' . Thus the number of such disks is at most $\pi((\lambda + 3)\rho_{\bar{R}})^2 / (\pi\rho_{\bar{R}})^2 = (\lambda + 3)^2$ \square

Corollary 6.2.2. *There are at most $(\lambda + 3)^2 \cdot n$ good edges in the transportation graph T .*

6.3 BOUNDING THE NUMBER OF NASTY EDGES IN T

In this section, we assume that the f -fat regions in \mathcal{R} are convex. We will show that, for any region R in \mathcal{R} , the number of nasty edges (R, R') in the transportation graph T is bounded from above by a function of f and λ .

Lemma 6.3.1. *Let (R, R') be a short edge of the transportation graph T between two regions R and R' . Then*

$$\rho_{R'}^- \geq |c_R c_{R'}| / (2f + \lambda)$$

Proof. Let r be the point on the boundary of R corresponding to the endpoint of (R, R') and denote the other endpoint of (R, R') by s . We have:

$$|rs| \leq \lambda \cdot \rho_{\bar{R}} \quad (1)$$

Note that

$$\begin{aligned} |c_R c_{R'}| &\leq |c_R r| + |rs| + |s c_{R'}| \\ &\leq f \cdot \rho_{\bar{R}} + |rs| + f \cdot \rho_{R'}^- \quad \text{by } f\text{-fatness} \\ &\leq f \cdot \rho_{\bar{R}} + \lambda \cdot \rho_{\bar{R}} + f \cdot \rho_{R'}^- \quad \text{by (1)} \\ &\leq (2f + \lambda)\rho_{R'}^- \quad \text{since } \rho_{\bar{R}} \leq \rho_{R'}^- \end{aligned}$$

It follows that $|c_R c_{R'}| / (2f + \lambda) \leq \rho_{R'}^-$. \square

Recall that we say that a region R is *convex* if for any pair of points p, q in R the line segment L_{pq} joining p and q is contained in R .

We now show that there cannot be arbitrarily long chains of f -fat

convex regions R_i with centers c_{R_i} aligned on a ray emanating from c_R such that that (R, R_i) defines a nasty edge of the transportation graph T .

Lemma 6.3.2. *Let R be a convex region and let r be a ray emanating from the center c_R of the inner disk of R . Let k be the number of regions convex R' such that (R, R') is a nasty edge of T and the center of the inner disk of R' is on the ray r . Then $k = \mathcal{O}((f + \lambda)^2)$.*

Proof. If $k \leq 1$, then the claim is true. Assume that $k \geq 2$. Let R be a region of \mathcal{R} . We consider a sequence $R_1, R_2, \dots, R_{k-1} = R_A, R_k = R_B$ of regions of \mathcal{R} such that:

- (R, R_i) is a nasty edge of T for $1 \leq i \leq k$
- $c_R, c_{R_1}, c_{R_2}, \dots, c_{R_k}$ are, in this order, on a ray r emanating from c_R

For convenience we will assume that the line V_{c_R} induced by the ray r is vertical. Since the inner disks $D_{R_i}^-, D_{R_j}^-$ are pairwise disjoint for all $i \neq j$, we have that

$$2(\rho_{R_1}^- + \rho_{R_2}^- + \dots + \rho_{R_{i-1}}^-) < |D_R^- D_{R_i}^-|$$

for $2 \leq i \leq k$. In particular,

$$2(\rho_{R_1}^- + \rho_{R_2}^- + \dots + \rho_{R_{k-2}}^-) < |D_R^- D_{R_A}^-|$$

Since $\rho_{R_i}^- \geq \rho_R^-$ for all i we get

$$2(k-2)\rho_R^- < |D_R^- D_{R_A}^-| \quad (1)$$

For $i \leq k$, let x_i and y_i denote the two points on the boundary of $D_{R_i}^-$ such that $\angle c_R c_{R_i} x_i = \angle c_R c_{R_i} y_i = \pi/2$. We distinguish between x_i and y_i as follows: x_i is the point to the left of V_{c_R} . Note that V_{c_R} defines the unique line equidistant to x_i and y_i .

Denote by p the point of R corresponding to the endpoint of (R, R_B) . We denote the other endpoint of (R, R_B) by q . Let L_A denote the horizontal line tangent to $D_{R_A}^-$ that is closest to c_R .

Suppose that the q lies above L_A . The distance from c_R to L_A is $|c_R c_{R_A}| - \rho_{R_A}^-$, and so $|qc_R|$ is at least $|c_R c_{R_A}| - \rho_{R_A}^-$. Since p is the other endpoint of the nasty edge (R, R_B) we have $|pq| \leq \lambda \cdot \rho_R^-$.

If $|c_R p| < |c_R c_{R_A}| - \rho_{R_A}^- - \lambda \cdot \rho_{R_A}^-$, then

$$\begin{aligned} |c_R p| + |pq| &< |c_R c_{R_A}| - \rho_{R_A}^- - \lambda \cdot \rho_R^- + |pq| \\ &\leq |c_R c_{R_A}| - \rho_{R_A}^- - \lambda \cdot \rho_R^- + \lambda \cdot \rho_R^- \\ &= |c_R c_{R_A}| - \rho_{R_A}^- \\ &\leq |qc_R| \end{aligned}$$

which contradicts the triangle inequality. This entails that $|c_R p|$ must be at least $|c_R c_{R_A}| - \rho_{R_A}^- - \lambda \cdot \rho_R^-$. We note that if

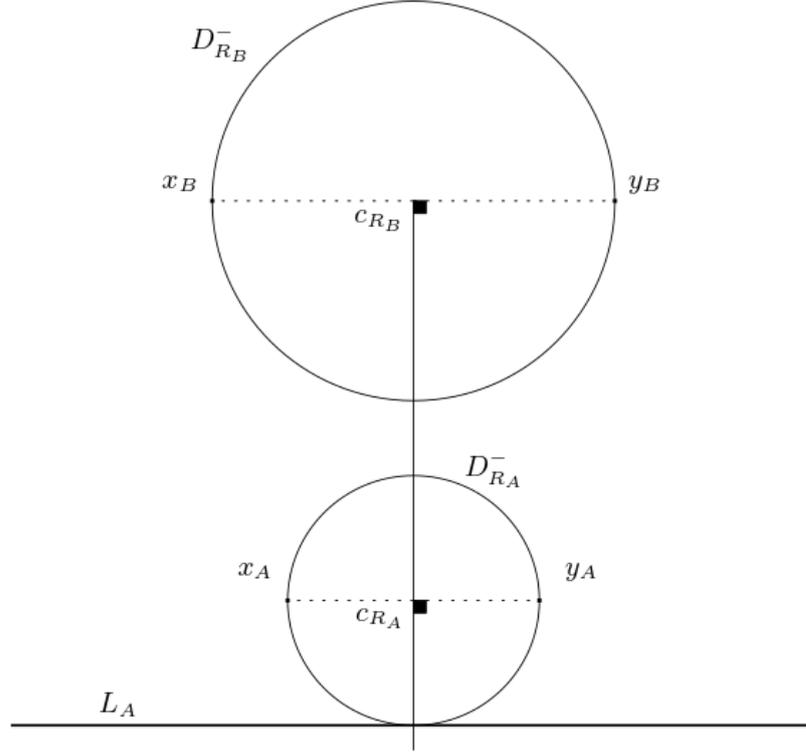


Figure 6.7

$$|D_{R_B}^- D_{R_A}^-| > (f + \lambda) \cdot \rho_{R_B}^-$$

then

$$\begin{aligned} f \cdot \rho_{R_B}^- &< |D_{R_B}^- D_{R_A}^-| - \lambda \cdot \rho_{R_B}^- \\ &= |c_{R_B} c_{R_A}| - \rho_{R_A}^- - \rho_{R_B}^- - \lambda \cdot \rho_{R_B}^- \\ &< |c_{R_B} c_{R_A}| - \rho_{R_A}^- - \lambda \cdot \rho_{R_B}^- \\ &\leq |c_{R_B} p| \end{aligned}$$

which contradicts the f -fatness of R . Hence we must have $|D_{R_B}^- D_{R_A}^-| \leq (f + \lambda) \cdot \rho_{R_B}^-$. And so by (1) we have $2(k - 2)\rho_{R_B}^- \leq (f + \lambda) \cdot \rho_{R_B}^-$. Thus $(k - 2) \leq (f + \lambda)/2$ and $k \in \mathcal{O}(f + \lambda)$.

Now suppose that q lies below L_A . Let L_1 be the line passing through x_B and y_A , and L_2 be the line passing through y_B and x_A . Let H_A (H_B) be the horizontal line passing through x_A and y_A (x_B and y_B). Let V_1 , V_2 be the vertical lines passing through y_A and x_A respectively. Let u be the intersection point of L_1 and L_2 . Let V_u be the vertical line perpendicular to L_A passing through u , and let H_u be the horizontal line parallel to L_A containing u .

Consider the isosceles trapezoid Q defined by x_A, y_A, x_B , and y_B . Note that u corresponds to the intersection point of the diagonals of Q . Since Q is convex, u must be inside Q . Moreover, u must be equidistant

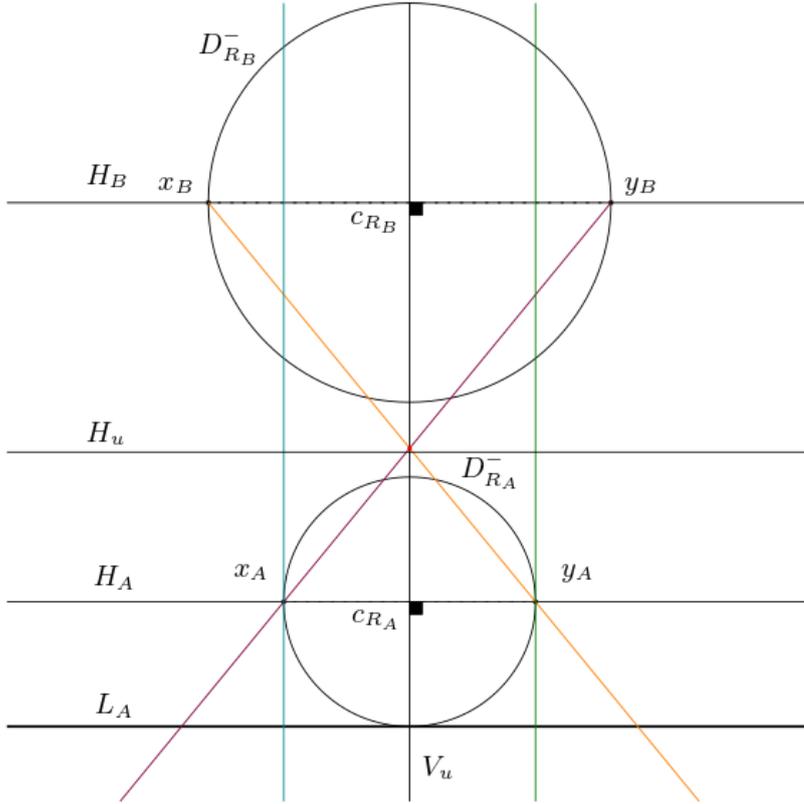


Figure 6.8: In orange and purple we have the lines L_1 and L_2 respectively. In blue and green we have the line V_2 and V_1 respectively. The point u is in red.

to x_A and y_A . In particular, u must lie on the line equidistant from x_A and y_A . Thus u must lie on V_{c_R} .

Let C_A be the cone with apex u defined by the rays going from u to x_A and from u to y_A . Similarly, let C_B be the cone with apex u defined by the rays going from u to x_B and u to y_B . Let $\gamma = \angle x_B u y_B = \angle x_A u y_A$. We observe that C_A can be defined as the set containing all points v below H_u such that $\angle v u c_{R_A} \leq \gamma/2$. Similarly, C_B can be defined as the set containing all points v above H_u such that $\angle v u c_{R_B} \leq \gamma/2$. Note that V_u corresponds to the union of $\text{bis}(C_A)$ and $\text{bis}(C_B)$.

Suppose q lies below both L_1 and L_2 . Then q lies in C_A .

Let L_q be the line passing through u and q . We denote by L_q^+ and L_q^- the portion of the line of L_q that lies above and below H_u respectively¹. Let q' be some point of L_q^+ . Notice that the angles $\angle q u c_{R_A}$ and $\angle q' u c_{R_B}$ are opposites. Since q is in C_A , we have $\gamma/2 \geq \angle q u c_{R_A} = \angle q' u c_{R_B}$. Thus q' is in C_B and so L_q^+ is contained in C_B . Since $L_{x_B y_B}$ intersects both legs of C_B , we know it intersects L_q^+ in some point q_B . Similarly,

¹ L_q^+ and L_q^- are not empty: q lies below L_A which is below H_u .

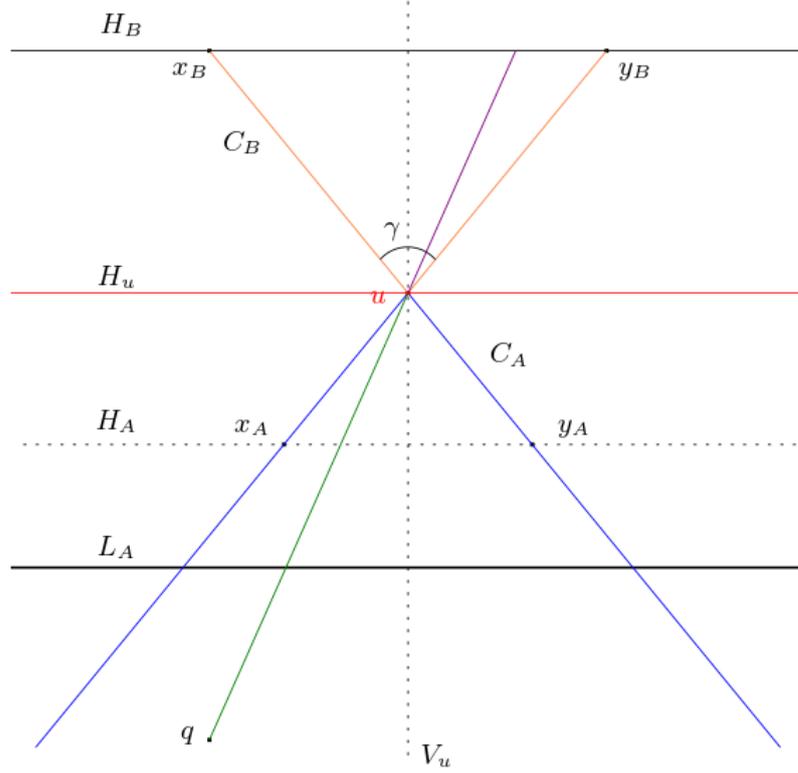


Figure 6.9: In orange we have the cone C_B and in blue we have the cone C_A . In green we have L_q^- and in purple we have L_q^+

$L_{x_A y_A}$ intersects both legs of C_A and so it intersects L_q^- in some point q_A . Note that q_A lies on H_A which is above L_A . Thus, q_A lies on the line segment L_{qq_B} . Moreover q_A is in $D_{R_A}^-$ and q_B is in D_B^- . However, L_{qq_B} is contained in R_B by convexity. In particular, q_A is in R_B which contradicts the fact that the regions are disjoint. Thus q must lie above either L_1 or L_2 .

Let $X = |L_1 c_R| = |L_2 c_R|$. Since q is above either L_1 or L_2 , p must be at least $X - \lambda \cdot \rho_R^-$ away from c_R .

As shown in Figure 6.10, c_R lies on V_u , below c_{R_A} and at distance at least $\rho_{R_A}^-$ from c_{R_A} . Note that V_1 is closer to c_R than L_1 and V_2 is closer to c_R than L_2 . It follows that if $\rho_{R_A}^- > (f + \lambda) \cdot \rho_R^-$, then

$$\begin{aligned} f \cdot \rho_R^- &< \rho_{R_A}^- - \lambda \cdot \rho_R^- \\ &< X - \lambda \cdot \rho_R^- \\ &\leq |c_R p| \end{aligned}$$

which contradicts the f -fatness of R . Hence, $\rho_{R_A}^- \leq (f + \lambda) \cdot \rho_R^-$. Combining this with Lemma 6.3.1 entails that $|c_R c_{R_A}| \leq (2f + \lambda)(f + \lambda) \cdot \rho_R^-$. Since $|D_R^- D_{R_A}^-| < |c_R c_{R_A}|$, (1) implies

$$2(k - 2) \cdot \rho_R^- < (f + \lambda)(2f + \lambda) \cdot \rho_R^-$$

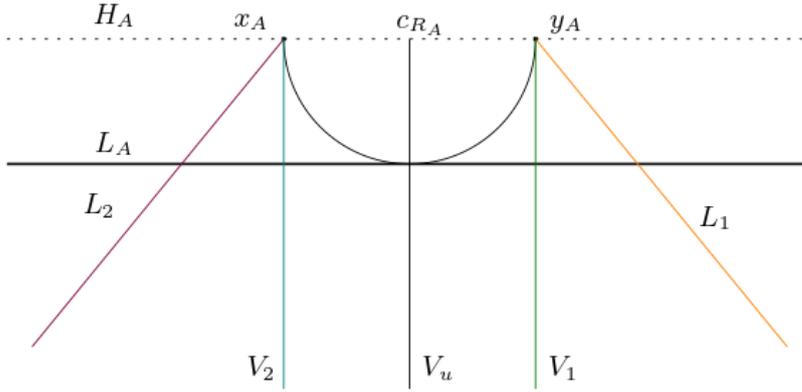


Figure 6.10

In other words, $k \in \mathcal{O}((f + \lambda)^2)$.

□

Let R be a region in \mathcal{R} and consider a cone with apex c_R whose angle is sufficiently small. We now proceed to show that there cannot be an arbitrarily large number of regions R' such that (R, R') is a nasty edge in T and the center of the inner disk of R' is in this cone; in Lemma 6.3.3 we find an upper bound on this angle θ . Then in Lemma 6.3.4 we show that each such regions R' is $(f + 1)/(1 - \sin(\theta/3))$ -fat with respect to inner disks with centers aligned on a ray emanating from c_R .

Lemma 6.3.3. *Let (R, R') be a short edge of T . Let x' and y' correspond to the two points on the boundary of $D_{R'}^-$ such that $\angle c_R c_{R'} x' = \angle c_R c_{R'} y' = \pi/2$. Then $\angle x' c_R c_{R'} = \angle y' c_R c_{R'} \geq \theta$ where $\theta = \arctan(1/(2f + \lambda))$*

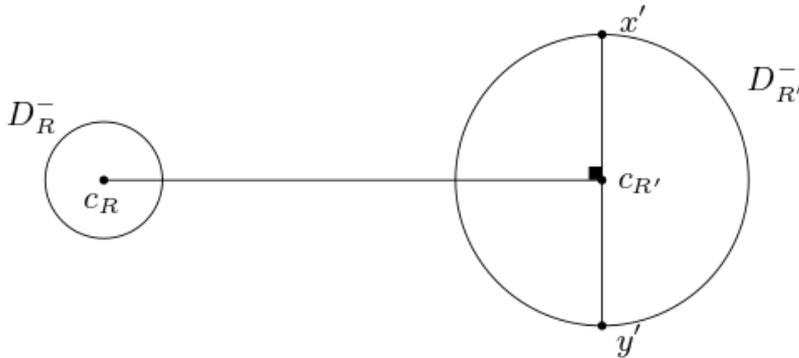


Figure 6.11

Proof. As illustrated in Figure 6.11, by Lemma 6.3.1, $\tan(\angle x' c_R c_{R'}) = \rho_{R'}^- / |c_R c_{R'}| \geq 1/(2f + \lambda)$ by Lemma 6.3.1. Thus $\angle x' c_R c_{R'} = \angle y' c_R c_{R'} \geq \arctan(1/(2f + \lambda)) = \theta$. □

Lemma 6.3.4. *Let (R, R') be a nasty edges of T , and let D_R^- and $D_{R'}^-$ denote the inner disks of R and R' respectively. Let $\theta = \arctan(1/(2f + \lambda))$.*

Consider a cone C with apex c_R of angle $2\theta/3$ and suppose that $c_{R'}$ is in C . Then:

1. $D_{R'}^-$ intersects both legs of C
2. There is a disk D'_0 contained in $D_{R'}^-$ with center c'_0 on $\text{bis}(C)$ and with radius ρ'_0 such that $\rho'_0 > (1 - \sin(\theta/3)) \cdot \rho_{R'}^-$.

Proof. For convenience we will assume that the $\text{bis}(C)$ is horizontal. Without loss of generality, suppose that $c_{R'}$ is to the below of $\text{bis}(C)$. Let x' and y' be the two points of $D_{R'}^-$ defined in Lemma 6.3.3 such that $\angle x'c_{R'}c_R = \angle y'c_{R'}c_R = \pi/2$. Let u denote the point on $\text{bis}(C)$ obtained by walking from $c_{R'}$ in the direction of x' and let $\alpha = \angle uc_Rc_{R'}$. Since $\alpha + \theta/3 \leq 2\theta/3 < \theta$, x' must lie outside of C by Lemma 6.3.3. In particular, the line segment $L_{x'c_{R'}}$ and the upper leg of C intersect in some point w of $D_{R'}^-$.

Now let v denote the projection of $c_{R'}$ onto the upper leg of C . The points $wvc_{R'}$ form a right triangle. Therefore, $|vc_{R'}| \leq |wc_{R'}| < \rho_{R'}^-$ and so v lies in $D_{R'}^-$.

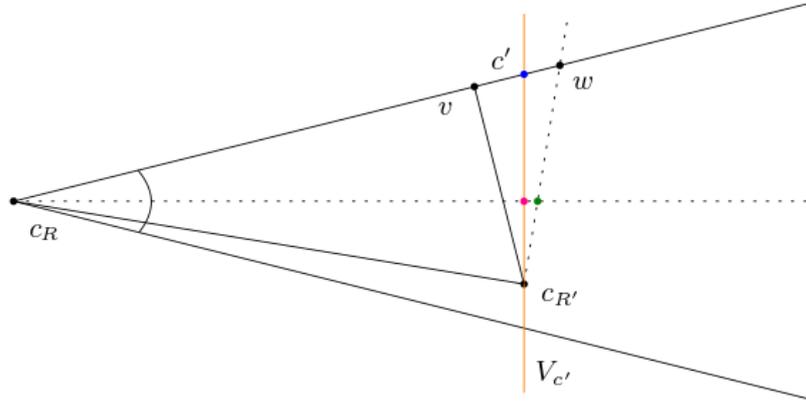


Figure 6.12: In blue we have the intersection point c' . In green we have the point u and in pink we have the point b .

Let $V_{c'}$ denote the vertical line passing through $c_{R'}$ and let c' denote the point of intersection of the upper leg of C and $V_{c'}$. Notice that v must be to the left of $V_{c'}$ since the angle $\angle c_Rc_{R'}c'$ is acute. Similarly, w must be to the right of $V_{c'}$ since $c_{R'}$ is not on the bisector of C .

Thus, c' lies on the line segment L_{vw} which is contained in $D_{R'}^-$ by convexity of $D_{R'}^-$. Now consider the diameter of $D_{R'}^-$ passing through c' . It intersects $\text{bis}(C)$ in some point b . Observe that the disk D'_0 with center b of radius $\rho_{R'}^- - |bc_{R'}|$ is fully contained in $D_{R'}^-$.

Note that $\angle bc_Rc_{R'} = \angle uc_Rc_{R'} = \alpha$. It follows that $|bc_{R'}| = |c_Rc_{R'}| \cdot \sin(\alpha) \leq |c_Rc_{R'}| \cdot \sin(\theta/3)$ and thus

$$\rho_{R'}^- - |bc_{R'}| \geq \rho_{R'}^- - \sin(\theta/3) \cdot |c_Rc_{R'}| > \rho_{R'}^- - \sin(\theta/3) \cdot \rho_{R'}^-$$

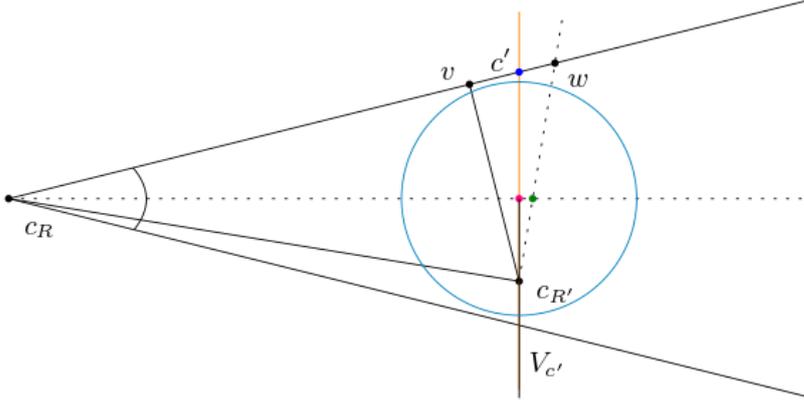


Figure 6.13: In light blue we have the disk D'_0

since $\rho_{R'}^- < |c_R c_{R'}|$. Finally, notice that since $c_{R'}$ is below $\text{bis}(C)$ and that $D_{R'}^-$ intersects the upper leg of C , $D_{R'}^-$ also intersects the lower leg of C . \square

Corollary 6.3.5. *With respect to D'_0 , and $D_{R'}^+$ the region R' is f_0 -fat, where $f_0 = (f + 1)/(1 - \sin(\theta/3))$.*

Proof. Since R' is f -fat we know that R' is contained in a disk $D_{R'}^+$ centered at c_R of radius at most $f \cdot \rho_{R'}^-$. Since the center b of D'_0 is at distance at most $\sin(\theta/3) \cdot |c_R c_{R'}|$ of $c_{R'}$, b is as far as $f \cdot \rho_{R'}^- + \sin(\theta/3) \cdot |c_R c_{R'}|$ away from the boundary of $D_{R'}^+$. Therefore, we have that:

$$\begin{aligned} f \cdot \rho_{R'}^- + \sin(\theta/3) \cdot |c_R c_{R'}| &\leq f \cdot \rho_{R'}^- + \sin(\theta) \cdot |c_R c_{R'}| \\ &\leq f \cdot \rho_{R'}^- + \tan(\theta) \cdot |c_R c_{R'}| \\ &\leq f \cdot \rho_{R'}^- + \rho_{R'}^- \quad \text{by def. of } \theta \\ &= f_0 \cdot (1 - \sin(\theta/3)) \rho_{R'}^- \end{aligned}$$

Thus R' is f_0 -fat. \square

Lemma 6.3.6. *Let R be a region, let $\theta = \arctan(1/(2f + \lambda))$, and let C be a cone with apex c_R and angle $2\theta/3$. Let k be the number of regions R' such that (R, R') is a nasty edge of T and the center of the inner disk of R' is in the cone C . Then $k = \mathcal{O}((f + \lambda)^2)$.*

Proof. By Lemma 6.3.4, the inner disk $D_{R'}^-$ of such a region R' can be replaced by a disk $D'_{R'}$ whose center $c'_{R'}$ lies on $\text{bis}(C)$. By Corollary 6.3.5, such a region R' is f_0 -fat where $f_0 = (f + 1)/(1 - \sin(\theta/3))$. Notice that $f_0 > f$ and so R is also f_0 -fat. The number of such regions R' is therefore $\mathcal{O}((f_0 + \lambda)^2)$ by Lemma 6.3.2. For large values of f and λ , we can observe that θ is close to $1/(2f + \lambda)$ since $\arctan x$ is close to x when x is close to 0. It follows that $f_0 \in \mathcal{O}(f)$. Thus $k \in \mathcal{O}((f + \lambda)^2)$ \square

Corollary 6.3.7. *Let \mathcal{R} be a set of n pairwise disjoint convex f -fat regions for some $f \geq 1$. The number of nasty edges in the transportation graph T is $\mathcal{O}((f + \lambda)^3 n)$.*

Proof. Let R be some region of \mathcal{R} , let $\theta = \arctan(1/(2f + \lambda))$. Consider $3\pi/\theta$ cones of angle $2\theta/3$ and apex at c_R that cover the plane. By Lemma 6.3.6, for each cone, there are $\mathcal{O}((f + \lambda)^2)$, nasty edges (R, R') such that the center of the inner disk of R' is in this cone. We also noted in the proof of 6.3.6 that for large values of f and λ , θ is about $1/(2f + \lambda)$. In other words $1/\theta = \mathcal{O}(f + \lambda)$. Therefore, the total number of nasty edges (R, R') is

$$\mathcal{O}(3\pi/\theta \cdot (f + \lambda)^2 n) = \mathcal{O}((f + \lambda)^3 n).$$

□

6.4 THE GREEDY SPANNER FOR FAT CONVEX REGIONS

We begin by describing algorithm PATHGREEDY for sets of regions.

Algorithm PATHGREEDY(\mathcal{R}, t)

1 *Input.* A set \mathcal{R} of regions of n regions, the edge set E_T of the transportation graph T induced by \mathcal{R} , a real number $t > 1$.
2 *Output.* A t -region spanner of T
3 Sort the at most $\binom{n}{2}$ undirected edges of T in non-decreasing order of their weights (breaking ties arbitrarily), and store them in a list L ;
4 $E := \emptyset$;
5 $G := (\mathcal{R}, E)$;
6 **for each** $\{R, R'\} \in L$ (consider pairs in sorted order)
7 **do** $\delta_G(R, R') :=$ length of a shortest path in G between R and R' ;
8 **if** $\delta_G(R, R') > t \cdot |RR'|$
9 **then**
10 $E := E \cup \{\{R, R'\}\}$;
11 $G := (\mathcal{R}, E)$
12 **endif**
13 **endfor**;
14 Output the graph G

Let \mathcal{R} be a set of pairwise disjoint f -fat convex regions for some $f \geq 1$. Let $G = (\mathcal{R}, E)$ be the graph computed by PATHGREEDY. We know that, by construction and Lemma 3.1.4, PATHGREEDY outputs a region t -spanner of T . In sections 6.2 and 6.3 we have obtained upper bounds on the number of good and nasty edges in the transportation T . Since G is a subgraph of T , these bounds also applies to G . We now prove an upper bound on the number of long edges in G .

Lemma 6.4.1. *Let $1.99 > t \geq 1$ be a real number sufficiently close to 1. Let $\theta = (t - 1)/2$ and $w = 2\theta/5$. Then*

(a) $0 \leq w < (\cos \theta - \sin \theta)/2$

(b) $t \geq 1/(\cos \theta - \sin \theta - 2w)$

Proof. The equation $2x/5 < (\cos x - \sin x)/2$ is true whenever $x < 0.49916$. Since $\theta = (t - 1)/2 < 0.99/2 = 0.495$ we have $w < (\cos \theta - \sin \theta)/2$ and (a) is satisfied. (b) is equivalent to

$$1/t + 2w \leq \cos \theta - \sin \theta$$

which is equivalent to

$$1/(1 + 2\theta) + 4\theta/5 - \cos \theta + \sin \theta \leq 0 \quad (1)$$

By using Taylor series we note that

$$\cos \theta - \sin \theta = 1 - \theta + \mathcal{O}(\theta^2)$$

and that

$$1/(1 + 2\theta) = 1 - 2\theta + \mathcal{O}(\theta^2)$$

so the left-hand side is

$$1 - 2\theta + 4\theta/5 - 1 + \theta + \mathcal{O}(\theta^2) = -\theta/5 + \mathcal{O}(\theta^2) \quad (2)$$

Remark that (2) is negative when θ is sufficiently small. \square

We state the following lemma without its proof. The proof can be found in [13].

Lemma 6.4.2. *Let t, θ and w be real numbers such that $0 < \theta < \pi/4$, $0 \leq w < (\cos \theta - \sin \theta)/2$, and $t \geq 1/(\cos \theta - \sin \theta - 2w)$. Let p, q, r, s be 4 points in the plane such that:*

- $p \neq q, r \neq s$
- $\text{angle}(rs, pq) \leq \theta$
- $|rs| \leq |pq|/\cos \theta$
- $|pr| \leq w|rs|$

Then $|pr| < |pq|$, $|sq| < |pq|$, and $t|pr| + |rs| + t|sq| \leq t|pq|$

Recall that an edge (R, R') of T is called long if $|RR'| \geq \lambda \cdot \rho_R^-$, where λ is a sufficiently large number. In the next lemma, we choose λ as a function of f and t . We now prove the upper bound on the number of long edges in the graph G computed by PATHGREEDY.

Lemma 6.4.3. *Let $1.99 > t > 1$ for some t sufficiently close to 1. Let $\lambda = 10f/(t - 1)$. Let G be the graph computed by PATHGREEDY. For a given region R of \mathcal{R} , there are $\mathcal{O}(1/(t - 1))$ long edges (R, R') in G .*

Proof. Let R, R' and R'' be regions of \mathcal{R} such that (R, R') and (R, R'') are long edges of the greedy spanner G . Let r be the point on the boundary of R corresponding to the endpoint of (R, R') and denote the other endpoint of (R, R') by s . Similarly let p be the point on the boundary of R corresponding to the endpoint of (R, R'') and denote the other endpoint of (R, R'') by q . We have:

$$|rs| \geq \lambda \cdot \rho_{\bar{R}} \text{ and } |pq| \geq \lambda \cdot \rho_{\bar{R}} \quad (1)$$

Without loss of generality, we can assume that $|rs| \leq |pq|$. Let $\text{angle}(rs, pq)$ denote the angle formed by translating both segments rs and pq to the origin. Let $\theta = (t-1)/2$ and $w = 2\theta/5$. Suppose that $\text{angle}(rs, pq) < \theta$. Then we have

$$\begin{aligned} |pr| &\leq 2f \cdot \rho_{\bar{R}} \quad \text{since } R \text{ is } f\text{-fat} \\ &\leq 2f \cdot |rs|/\lambda \quad \text{by (1)} \\ &\leq w \cdot |rs| \quad \text{since } \lambda = 2f/w \end{aligned}$$

Moreover, notice that since $t < 1 + \pi/2$ then $\theta < \pi/4$. By Lemma 6.4.1, all conditions for Lemma 6.4.2 are satisfied.

Thus $|R'R''| \leq |sq| < |pq|$. Now consider the moment at which `PATHGREEDY` reaches the edge (R, R'') . Let the sequence of regions $R' = R_0, R_1, R_2, \dots, R_k = R''$ define a shortest path between R' and R'' in T . Since $\delta_T(R', R'') \leq |R'R''| < |pq|$ by Lemma 3.1.3, at that moment all pairs of regions (R_i, R_{i+1}) are connected by a t -spanner path P_i in G . Let $|P_i|$ denote the length of the path P_i . We have:

$$\begin{aligned} \sum_{i=0}^{k-1} |P_i| &\leq \sum_{i=0}^{k-1} t \cdot |R_i R_{i+1}| \\ &= t \cdot \sum_{i=0}^{k-1} \delta_T(R_i, R_{i+1}) \\ &= t \cdot \delta_T(R', R'') \end{aligned}$$

Hence there is a t -spanner path in G between R' and R'' . It follows that there is a path from R to R'' of length at most

$$|rs| + t \cdot |R'R''| \leq |rs| + t|sq| \leq t|pq| = t \cdot \delta_T(R, R'') \text{ by Lemma 6.4.3}$$

This implies that `PATHGREEDY` does not add the edge (R, R'') to G which contradicts our initial assumption. Therefore $\text{angle}(rs, pq) \geq \theta$ and so there cannot be more than $2\pi/\theta \in \mathcal{O}(1/(t-1))$ long edges with R as a source. \square

Corollary 6.4.4. *Let $1.99 > t > 1$ for some t sufficiently close to 1. Let $\lambda = 10f/(t-1)$. Let G be the graph computed by `PATHGREEDY`. G has $\mathcal{O}(n/(t-1))$ long edges.*

We now possess all the tools to prove the following Theorem:

Theorem 6.4.5. *Let \mathcal{R} be a set of n pairwise disjoint convex f -fat regions for some $f \geq 1$ and let $t > 1$ be a real number sufficiently close to 1. Running `PATHGREEDY` on the edges of transportation graph T yields a t -spanner G of T with $\mathcal{O}(((f + \frac{1}{t-1}))^3 n)$ edges. In particular, if f and t are constant, then G is sparse.*

Proof. We already know that the graph G constructed by `PATHGREEDY` is a t -region spanner. By Corollary 6.2.2, Corollary 6.3.7, and Corollary 6.4.4 respectively there are at most $(\lambda + 3)^2 n$ good edges, $\mathcal{O}((f + \lambda)^3 n)$ nasty edges, and $\mathcal{O}(n/(t - 1))$ long edges in G . Since $\lambda = 10f/(t - 1)$, it follows that G has at most

$$\mathcal{O}((f + \lambda)^3 + (1/(t - 1))n) = \mathcal{O}\left(\left(f + \frac{1}{t-1}\right)^3 n\right)$$

edges. Note that both f and t do not depend on n . Thus, G is indeed sparse. \square

6.5 OTHER NOTIONS OF FATNESS

In the previous section we introduced a notion of fatness. The requirement for the two disks to be concentric is a bit superfluous: it only made the proof of Lemmas 6.3.2, 6.3.4, and 6.3.6 a bit easier to write. Rather than looking at any two such disks we can require for D_R^- to be the biggest disk contained, call it $MAX(R)$, in R and for D_R^+ to be the smallest disk containing R , call it $MED(R)$. A region would then be f -fat for some $f \geq 1$ if $radius(MED(R)) \leq f \cdot radius(MAX(R))$.

Lemma 6.5.1. *Let R be some region. If $radius(MED(R)) \leq f \cdot radius(MAX(R))$ for some $f \geq 1$, then there exists two concentric disks D_R^-, D_R^+ such that D_R^- is contained in R , D_R^+ contains R and $radius(D_R^+) \leq f' \cdot radius(D_R^-)$, where $f' = 2f$.*

Proof. Let $c_{MAX(R)}$ be the center of $MAX(R)$ and let $c_{MED(R)}$ be the center of $MED(R)$. Note that $|c_{MAX(R)}c_{MED(R)}| \leq radius(MED(R))$. Therefore, the disk D_R^+ of radius $2 \cdot radius(MED(R))$ centered on $c_{MAX(R)}$ contains R and we have

$$radius(D_R^+) = 2 \cdot radius(MED(R)) \leq 2f \cdot radius(MAX(R))$$

Thus, R is f' -fat with respect to D_R^+ above and $D_R^- = MAX(R)$ for $f' = 2f$. \square

Note that if R is f -fat with respect to some concentric disks D_R^- and D_R^+ respectively contained in R and containing R , then R is f -fat with respect to $MAX(R)$ and $MED(R)$. The conjunction of this together with Lemma 6.5.1 implies that the two notions of fatness are equivalent (up to a factor of 2). We now present a different definition of fatness that is introduced in by *van der Stappen et al.* [16].

Let p be a point in \mathbb{R}^2 . We denote by $D_{p,r}$ the disk centered on p of radius r . That is

$$D_{p,r} = \{x \in \mathbb{R}^2 \mid |xp| \leq r\}$$

We denote the boundary of $D_{p,r}$ by $\partial D_{p,r}$. The interest of studying "fat" shapes comes in part from the following remark: the number of fat shapes surrounding or intersecting a region should not be too

big. With that idea in mind, let us consider the two following sets: Let R be a region and let $p \in \mathbb{R}^2$ be a point, the sets $U_{p,R}$ and U_R are respectively defined as

$$U_{p,R} = \{D_{p,r} \subseteq \mathbb{R}^2 \mid \partial D_{p,r} \cap R \neq \emptyset\}$$

and

$$U_R = \bigcup_{p \in R} U_{p,R}$$

U_R is the set of all disks with center inside R that do not fully contain R . We shall say that R is k -fat if the volume² of the part covered by R of any disk D whose boundary intersects R and whose center lies inside R is at least $1/k$ the volume of D . More formally, let R be a region in the plane and let k be a positive number. We shall say that R is k -fat if for all $D \in U_R$

$$k \cdot \text{volume}(R \cap D) \geq \text{volume}(D)$$

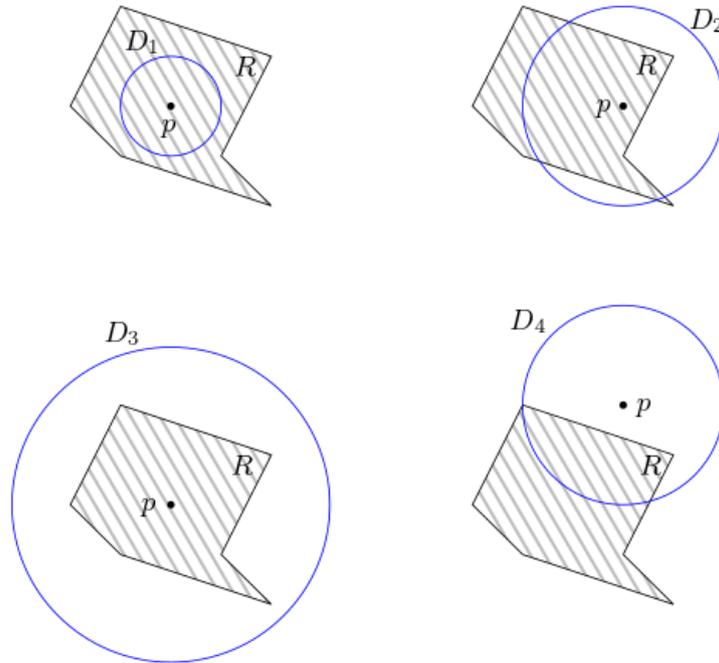


Figure 6.14: The disks D_1 and D_2 are in U_R , whereas D_3 and D_4 are not.

Let R be an k -fat region. The *diameter* δ of R corresponds to the longest distance between two points of R . In particular, there is two points a, b of R such that $|ab| = \delta$. By Jung's Theorem [10], the volume of R is at most the volume of a disk of radius $\delta/2$. Moreover, the disk $D_{a,\delta}$ is an element of U_R . Therefore we have:

$$\text{volume}(R \cap D_{a,\delta}) \leq \text{volume}(R) \leq \pi \cdot (\delta/2)^2$$

² Technically we are still only talking about regions in \mathbb{R}^2 so perhaps the word area would be less confusing. However, these results extend to \mathbb{R}^d for any $d \geq 2$ and so we stick with the word volume.

and

$$\text{volume}(D_{a,\delta}) = \pi \cdot \delta^2$$

It follows that

$$1/k \cdot \pi \cdot \delta^2 \leq \pi \cdot (\delta/2)^2$$

Thus $k \geq 2^2$.

Lemma 6.5.2. *Let \mathcal{R} be a set of k -fat regions. Let $R \in \mathcal{R}$ and let $\delta = 2 \cdot \text{radius}(\text{MED}(R))$ be the diameter of $\text{MED}(R)$. Then the number of regions $R' \in \mathcal{R}$ with larger minimal enclosing disk that lies within a distance of $\lambda \cdot \delta$ from R is at most $4 \cdot k \cdot (\lambda + 1)^2$, where λ is some positive real number.*

Proof. Let d be the center of $\text{MED}(R)$. Let R' be a region of \mathcal{R} such that $\text{radius}(\text{MED}(R)) < \text{radius}(\text{MED}(R'))$ and that is less than $\lambda \cdot \delta$ away from R . $D_{d, \frac{\delta}{2} + \delta \cdot \lambda}$ contains all points at distance $\lambda \cdot \delta$ of R .

Therefore, $D_{d, \frac{\delta}{2} + \delta \cdot \lambda} \cap R'$ is not empty. Let x be a point in that intersection.

Now consider $D_{x, \frac{\delta}{2}}$. We have that $x \in D_{x, \frac{\delta}{2}} \cap R'$. Moreover, R' is not fully contained in $D_{x, \frac{\delta}{2}}$; Otherwise $\text{radius}(\text{MED}(R')) \leq \delta/2 = \text{radius}(\text{MED}(R))$ which contradicts our selection of R' . This implies that $\partial D_{x, \frac{\delta}{2}} \cap R'$ is not empty thereby making $D_{x, \frac{\delta}{2}}$ an element of $U_{R'}$.

We notice that $D_{x, \frac{\delta}{2}} \subset D_{d, \delta + \lambda \cdot \delta}$: for all $y \in D_{d, \delta + \lambda \cdot \delta}$,

$$|dy| \leq |dx| + |xy| \leq \delta/2 + \lambda \cdot \delta + \delta/2$$

By k -fatness of R' ,

$$\text{volume}(R \cap D_{x, \frac{\delta}{2}}) \geq (\text{volume}(D_{x, \frac{\delta}{2}})/k) = (4\pi(\delta/2)^2/k)$$

In particular, each region R' at distance less than $\lambda \cdot \delta$ of R with larger enclosing disk must take at least $(4\pi(\delta/2)^2/k)$ of the volume of $D_{d, \delta + \lambda \cdot \delta}$ which is $4\pi((1 + \lambda) \cdot \delta)^2$.

Thus the number of such region R' is at most

$$(4\pi((1 + \lambda) \cdot \delta)^2)/(4\pi(\delta/2)^2/k) = 2^2 \cdot k \cdot (\lambda + 1)^2$$

□

Let \mathcal{R} be a set of n pairwise disjoint k -fat regions for some $k \geq 4$. In the previous section we oriented the edges of T with respect to the radii of some inner disks of each region. We now do something similar with respect to the smallest outer disk of each region $\text{MED}(R)$: each edge $\{R, R'\}$ of the transportation graph is oriented from the smaller minimal enclosing disk to the larger. Let λ be some positive real number. Each edge (R, R') can be classified as follows: (R, R')

is called long if $|RR'| > 2\lambda \cdot \text{radius}(\text{MED}(R))$. Otherwise (R, R') is called short. The following corollary follows immediately from Lemma 6.5.2.

Corollary 6.5.3. *Let \mathcal{R} be a set of n pairwise disjoint k -fat regions for some $k \geq 4$. The number of short edges of the transportation graph T is at most $4k \cdot (\lambda + 1)^2$.*

Lemma 6.5.4. *Let \mathcal{R} be a set of n pairwise disjoint k -fat regions for some $k \geq 4$. Let G be the graph computed by `PATHGREEDY` for some $1.99 > t > 1$ sufficiently close to 1. For a given region R of \mathcal{R} , there is $\mathcal{O}(1/(t-1))$ long edges (R, R') in G if $\lambda = 5/(t-1)$.*

Proof. Let R, R' and R'' be regions of \mathcal{R} such that (R, R') and (R, R'') are long edges of the greedy spanner G . Let r be the point on the boundary of R corresponding to the endpoint of (R, R') and denote the other endpoint of (R, R') by s . Similarly let p be the point on the boundary of R corresponding to the endpoint of (R, R'') and denote the other endpoint of (R, R'') by q . Without loss of generality, we can assume that $|rs| \leq |pq|$. Let $\theta = (t-1)/2$ and $w = 2\theta/5$. Suppose that $\text{angle}(rs, pq) < \theta$. We have

$$\begin{aligned} |pr| &\leq 2 \cdot \text{radius}(\text{MED}(R)) \\ &\leq 2 \cdot |rs|/2\lambda \\ &\leq w \cdot |rs| \quad \text{since } \lambda = 1/w \end{aligned}$$

By Lemma 6.4.1, all conditions for Lemma 6.4.2 are met. Hence, $|R'R''| \leq |sq| < |pq|$. The rest of the proof is identical to the proof of Lemma 6.4.3. Namely, we can conclude that `PATHGREEDY` does not add the edge (R, R'') to G which contradicts our initial assumption. Thus $\text{angle}(rs, pq) \geq \theta$ and so there cannot be more than $2\pi/\theta \in \mathcal{O}(1/(t-1))$ long edges with R as a source. \square

Theorem 6.5.5. *Let \mathcal{R} be a set of n pairwise disjoint k -fat regions for some $k \geq 4$. Let G be the graph computed by `PATHGREEDY` for some $t > 1$ sufficiently close to 1. G is a t -spanner of T with $\mathcal{O}(((k/(t-1))^2)n)$ edges.*

Proof. Corollary 6.5.3 and Lemma 6.5.4 together imply that for some region R , there is at most $\mathcal{O}((k/t^2 + 1/(t-1)))$ edges (R, R') . The theorem therefore follows immediately. \square

In stark difference of Theorem 6.4.5, the regions are no longer required to be convex. In section 6.3 we only made use of the convexity of the regions in the proof of Lemma 6.3.2. It is not clear whether or not this assumption can be removed. Moreover, we do not know whether or not the notions of f -fatness and k -fatness are equivalent. We first observe that if a region R is f -fat, then

$$\text{volume}(\text{MED}(R)) \leq f \cdot \text{volume}(\text{MAX}(R)) \leq f \cdot \text{volume}(R)$$

We shall say that a region R is k -thick if $\text{volume}(\text{MED}(R)) \leq k \cdot \text{volume}(R)$.

Lemma 6.5.6. *Let R be a convex region. Then*

- (a) *If R is k -fat for some $k \geq 4$, then R is k' -thick with $k' \in \mathcal{O}(k)$.*
- (b) *If R is k -thick for some $k \geq 1$, then R is k' -fat for some $k' \in \mathcal{O}(k)$.*

The proof of Lemma 6.5.6 is given in [16].

Corollary 6.5.7. *Let R be a convex f -fat regions for some $f \geq 1$. Then R is k -fat for some $k \in \mathcal{O}(f)$.*

Proof. We noted that any f -fat regions is f -thick. Since R is convex the claim holds by Lemma 6.5.6 \square

We do not know whether or not either k -fatness or k -thickness implies f -fatness for convex regions. In other words both following questions remain open:

- Let R be a convex region that is k -fat for some $k \geq 4$. Is there some $f \geq 1$ (possibly dependent on k) such that R is f -fat?
- Let R be a region that is k -thick for some $k \geq 1$. Is there some $f \geq 1$ (possibly dependent on k) such that R is f -fat?

There exists many other definitions of fatness, most of which were shown to be equivalent to k -fatness defined in [16] when the objects are convex. It is worth noting that Lemma 6.5.4 does not depend on the definition of fatness chosen. As long as analogous results to Lemma 6.2.1 and Corollary 6.3.7 or Lemma 6.5.2 hold for the chosen definition of fatness, the graph computed by `PATHGREEDY` will be a sparse spanner.

Finally, in this chapter we resorted for the first time to algorithm `PATHGREEDY` in order to produce a sparse spanner. Fatness allowed us to show that the degree of each region in the t region spanner generated by `PATHGREEDY` is some constant which only depends on the fatness parameter and t . A few questions arise naturally:

- Is fatness required in order to obtain such results? Notice how in the proof of Lemma 6.5.4 fatness is not invoked once. Therefore, if we were to orient the edges of the transportation graph T induced by a set of regions \mathcal{R} , can we show that the outdegree of each region is some constant which does not depends on n ?
- The way in which we oriented the edges of T played an important role in obtaining a bound on the degree of every region. Is this strictly required? Do we need to orient the edges of T of a set \mathcal{R} of f -fat (or k -fat) regions in order to show that the number of edges, either in T or in the graph computed by `PATHGREEDY`, incident to some region R does not depend on n ?

As we can see in Figure 6.15 below, the answer to the second item above is yes; some orientation is required to get a bound independent of n . The big disk has degree which depends on the number of tiny disks surrounding it. Therefore, the number of edges of T incident to that big disk depends on n .

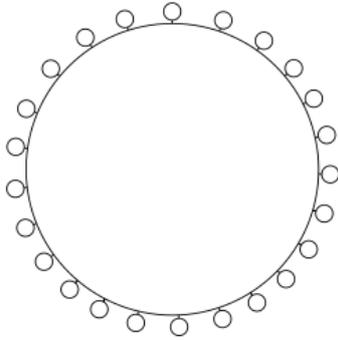


Figure 6.15: The degree of the big disk is linear in terms of the total number disks: there is an edge between each tiny disk and the big disk.

Throughout this thesis we have scrutinized a few techniques for building sparse spanners for arbitrary sets of regions. We insisted on the regions to be pairwise disjoint. However, if two regions R and R' had a nonempty intersection, the time travelling from R to R' would be zero. In particular, intersecting regions could be merged into one. Moreover, we also required for regions to be path-connected sets of \mathbb{R}^2 . We never made use of this assumption. In fact, we observe that it is somewhat redundant with respect to the assumption that the travel time (distance) between two points of a region R is zero: from a theoretical point view, there is no need to "move" from p to q for distinct locations in a region.

In applications however one would still expect to have to move from p to q albeit it being "free". Therefore, the path-connected constraint for regions is still one that is worth keeping. For example, if we had an entity located inside a region that wished to travel to some other region of a set \mathcal{R} , then it should be able to travel (at no cost) from its location to the boundary of R while staying inside R . Theorem 6.4.5 states that `PATHGREEDY` produces a sparse graph that can be as cost efficient as desired compared to the optimal travelling graph T for f -fat convex regions. We obtain a similar result for k -fat regions (with respect to the definition from *van der Stappen et al.* [16]). This gives rise to the following question. Denote by δ the cost, independent of the travelled distance ¹, of travelling inside any region of a set of regions \mathcal{R} . If $\delta \neq 0$, does `PATHGREEDY` still generate a cost efficient region spanner G ?

There are several other possible variants of the problem described above. For instance, we might consider the scenario in which travelling inside each region is charged a flat rate per unit of distance travelled. We could also assume that each region possesses its own flat rate, dependent or not, on the travelled distance. These problems remain open.

We also observed that several known results for spanners on sets of points can be generalized for region spanners. This is far from an exhaustive list. In particular, our main focus was to show the existence of sparse region spanners. Spanners have many other desirable properties that were not studied or mentioned here.

¹ In this scenarios, the cost can be conceptualized as a flat entering fee for all regions.

In Chapter 2 we briefly introduced the notion of WSPD. As we know, the standard Euclidean metric fails to encapsulate that two subsets of regions contained in two disks that are far apart in the Euclidean sense might be close in the sense of the metric δ_T . This inspires the following: instead of looking at the Euclidean metric we look at the δ_T metric.

Let \mathcal{R} be some finite set of regions, and let A be some subset of \mathcal{R} . The diameter of A denoted by δ_A corresponds to the biggest distance with respect to δ_T between the regions contained in A .

$$\delta_A = \max\{\delta_T(P, Q) \mid P, Q \in A\}$$

Two finite sets of regions A and B form a well-separated pair with respect to some real number $s > 0$ if

$$\delta_T(A, B) = \min\{\delta_T(P, Q) \mid P \in A, Q \in B\} \geq s \cdot \max\{\delta_A, \delta_B\}$$

A *well-separated pair decomposition* of a set \mathcal{R} of n regions with respect to some positive real number s is a sequence

$$\{A_1, B_1\}, \dots, \{A_m, B_m\}$$

of pairs of nonempty subsets of \mathcal{R} , for some integer m , such that

- (1) A_i and B_i are well-separated with respect to s for each $1 \leq i \leq m$
- (2) For any two regions P, Q of \mathcal{R} , there is exactly one index i such that $P \in A_i$ and $Q \in B_i$, or $P \in B_i$ and $Q \in A_i$

The construction outlined in Theorem 2.1.8 also extends to arbitrary metric space. Therefore, WSPDs yield spanners in the case of regions². In particular, if a WSPD of m pairs exists, then we obtain a t -spanner with m edges. We do not know whether or not WSPDs of size $m = \mathcal{O}(n)$ on a set of n regions exists. We also must not forget that, besides from the spanner construction presented in Chapter 5, the methods that described in other chapters required the transportation graph T to precomputed. We do not know if efficient algorithms exists that could perform this task.

² If we remove the requirement in the definition of region graph for the edges to not intersect other regions

APPENDIX

A.1 PROOF OF 5.1.2

Proof. By Lemma 3.1.4 it is sufficient to show that for any edge $\{R, R'\}$ of T

$$\delta_{T_\theta}(R, R') \leq t \cdot |RR'|$$

We proceed by induction on the rank of the length $|RR'|$ in the sequence of all edges of T .

Base Case: $\{R, R'\}$ is an edge in T having minimum length $|RR'|$. In particular, R, R' is the pair of region that is closest to each other. Let p and p' be the endpoints of $L_{RR'}$. Let C be the cone with apex centered on p containing $L_{RR'}$. Notice that, although $|RR'|$ is minimal, this does not necessarily entail that $\{R, R'\}$ is an edge of T_θ

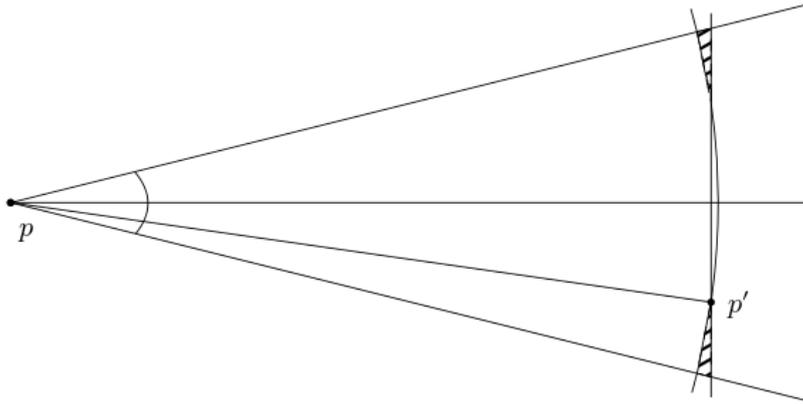


Figure A.1

However, if θ is sufficiently small, then it does by Lemma 5.1.1. More specifically if $\{R, R'\}$ is not an edge of T_θ there would be a point p'' corresponding to the endpoint of $L_{RR''}$ for some region R'' inside one of the highlighted regions. And so by Lemma 5.1.1

$$\begin{aligned} |R'R''| &\leq |p'p''| \\ &\leq |pp'| - (\cos \theta - \sin \theta)|pp''| \end{aligned}$$

Noticing that $|RR'| = |pp'|$ and $|RR''| \leq |pp''|$ we get

$$|R'R''| \leq |RR'| - (\cos \theta - \sin \theta)|RR''| < |RR'|$$

since $\theta < \pi/4$. But this contradicts the minimality of $|RR'|$.

Inductive Step: Let $\{R, R'\}$ be an edge of T that does not have minimum length. (I.H.) We suppose that $\delta_{T_\theta}(R_1, R_2) \leq t \cdot |R_1R_2|$ for every edge $\{R_1, R_2\}$ of T with $|R_1R_2| < |RR'|$

We first notice that if $\{R, R'\}$ is an edge of T_θ then $\delta_{T_\theta}(R, R') \leq |RR'|$ and so we are done. So suppose that $\{R, R'\}$ is not an edge of T_θ . Let p and p' be the endpoints of $L_{RR'}$. Let C be the cone with apex centered on p that contains $L_{RR'}$.

Since $\{R, R'\}$ is not an edge of T_θ , there must be another edge $\{R, R''\}$ of T incident to p such that R'' yielded a point whose projection on the bisector of C was closer than p' . Repeating what we did for the base case, Lemma 5.1.1 entails that $|R'R''| < |RR'|$. And so we have $\delta_T(R', R'') \leq |R'R''| < |RR'|$. In particular, every edge on a shortest path $R'' = R_1, \dots, R_j = R'$ in T from R'' to R' has smaller weight than $|RR'|$. By our induction hypothesis (I.H.) we have that

$$\delta_{T_\theta}(R_i, R_{i+1}) \leq t \cdot |R_iR_{i+1}| \quad 1 \leq i \leq j-1$$

Summing of all i yields:

$$\delta_{T_\theta}(R'', R') \leq t \cdot \delta_T(R'', R') \quad (1)$$

Now consider the path in T_θ that consists of first travelling from R to R'' via the edge $\{R, R''\}$ and then from R'' to R' by some shortest path. Setting $f(\theta) = (\cos \theta - \sin \theta)$, the following must hold:

$$\begin{aligned} \delta_{T_\theta}(R, R') &\leq |pp''| + \delta_{T_\theta}(R', R'') \quad \text{by def. of shortest path} \\ &\leq |RR''| + t \cdot \delta_T(R', R'') \quad \text{by (1)} \\ &\leq |RR''| + t \cdot |R'R''| \quad \text{by Lemma 3.1.3} \\ &\leq |RR''| + t \cdot (|RR'| - f(\theta)|RR''|) \quad \text{by Lemma 5.1.1} \end{aligned}$$

Setting $t = 1/f(\theta)$, we have

$$\delta_{T_\theta}(R, R') \leq t \cdot |RR'|$$

as desired. □

A.2 DEFINING THE EVENT QUEUE WHEN SEVERAL EVENTS CAN BE HIT BY THE SWEEP LINE

Recall that we define a priority queue or *event queue* on the vertices of the polygons of \mathcal{R} based the following observations:

- We want to insert each segment S only once in the BST. This should occur precisely the first time the sweep line reaches S . In particular, for each segment we need to be able to tell which of its endpoint is hit first by L_i . Similarly we want to delete each segment S only once. This should happen only when reaching the endpoint of S that is hit second by L_i .

- We want to ensure that all insertions are performed before the deletions in the BST. Moreover, we also need to worry about the order in which we perform the insertions and the deletions.

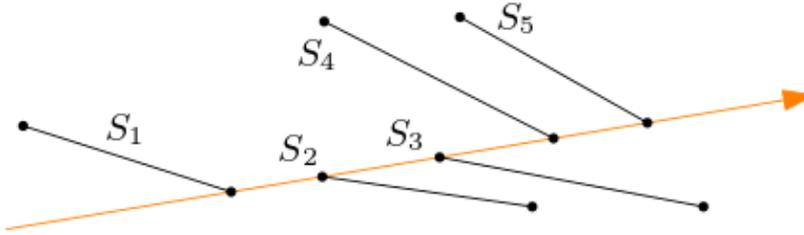


Figure A.2

As shown in Figure A.2, if S_4 was removed before S_3 was added then the successor of S_3 would falsely be S_5 . If S_5 was removed before S_4 , then upon removing S_4 we say falsely report that S_4 has no successor. If S_2 is added before S_3 then S_2 would have S_4 as its successor.

Let L_i be one of the sweep lines. We sort the points with respect to their scalar component on the line $L_{i+\pi/2}$ in increasing order oriented in the direction of $r_{i+\pi/2}$. In order to break ties we also compute the scalar component of each point on L_i oriented in the direction of r_i . It follows that the vertex whose scalar component(s) was maximal is hit first and so on. This order is stored in a list $EventQueue_i(\mathcal{R})$. In the same manner, for each segment we can determine which of its endpoint is hit first. We store and sort with respect to the same ordering in $EventQueue_i(\mathcal{R})$ all of these endpoint in a list $First_i(\mathcal{R})$. Similarly, we define the list $Second_i(\mathcal{R})$. Again these lists are not necessarily disjoint. Once those lists are built, we begin the sweep following the ordering given by $EventQueue_i(\mathcal{R})$. Upon reaching one or more vertices v_1, \dots, v_j

1. We first treat the vertices that are in $First_i(\mathcal{R})$ in the same order prescribed by $First_i(\mathcal{R})$. In particular, we search for $S(v_k, i)$ in the BST and then we insert the line segment(s) with v_k as their "first" endpoint.
2. We then treat the vertices that in $Second_i(\mathcal{R})$ with respect to the reversed order given by $Second_i(\mathcal{R})$. In particular, we search for $S(v_k, i)$ in the BST and then we remove from BST the line segment(s) with v_k as their "second" endpoint.

Lemma A.2.1 For any vertex p and for any ray r_i , `FINDINTERSECTION` reports a successor to p when L_i reaches p if and only if there is an intersection on the ray $r_i + p$

Proof. Let r_i be one of the k rays. We proceed by induction on the rank of the vertices in the list $EventQueue_i(\mathcal{R})$.

Base Case: The vertex p of priority 1 in $EventQueue_i(\mathcal{R})$ is the point in $First(\mathcal{R})$ whose scalar components was maximal. In particular, this point has no successor in the BST and the ray $r_i + p$ has no intersection.

Induction Step: Let p be a vertex whose rank in $EventQueue_i(\mathcal{R})$ is not minimal. We assume that all vertices prior to p in $EventQueue_i(\mathcal{R})$ were treated correctly. We show that vertex p is also treated correctly which is to say that $FINDINTERSECTION$ reports a successor to p when L_i reaches p if and only if there is an intersection on the ray $r_i + p$.

Suppose that $FINDINTERSECTION$ reports a successor to p when L_i reaches p . By definition of $FINDINTERSECTION$ this successor lies on the portion of the line L_i corresponding to the ray $r_i + p$. Conversely, suppose that there is an intersection to be detected on the ray $r_i + p$. We denote this intersection by q . If q is not an endpoint of any segment of $Seg(\mathcal{R})$ then the segment it corresponds to was correctly added to the BST prior to reaching p . The same can be said if q is a "second" endpoint. In particular, if both p and q are in $Second_i(\mathcal{R})$, then p is of lower priority and $FINDINTERSECTION$ correctly treated q before p . If both p and q are "first" endpoints then q is still of higher priority than p since the y coordinate of q is bigger. The cases where p, q are not the same type of endpoint follows immediately from the ordering of (1) and (2) in $FINDINTERSECTION$. \square

BIBLIOGRAPHY

- [1] Colin Conrad Adams and Robert David Franzosa. *Introduction to Topology: Pure and Applied*. Sirsi.
- [2] Sunil Arya and Michiel Smid. "Efficient construction of a bounded-degree spanner with low weight." In: *Algorithmica* 17.1 (1997), pp. 33–54.
- [3] Luis Barba, Prosenjit Bose, Mirela Damian, Rolf Fagerberg, Wah Loon Keng, Joseph O'Rourke, André Van Renssen, Perouz Taslakian, Sander Verdonschot, and Ge Xia. "New and improved spanning ratios for Yao graphs." In: *Proceedings of the thirtieth annual symposium on Computational geometry*. 2014, pp. 30–39.
- [4] Prosenjit Bose, Paz Carmi, and Mathieu Couture. "Spanners of additively weighted point sets." In: *J. Discrete Algorithms* 9.3 (2011), pp. 287–298.
- [5] Prosenjit Bose and Michiel Smid. "On plane geometric spanners: A survey and open problems." In: *Computational Geometry* 46.7 (2013), pp. 818–830.
- [6] Prosenjit Bose, Jean-Lou De Carufel, Pat Morin, André Van Renssen, and Sander Verdonschot. "Towards tight bounds on theta-graphs: More is not always better." In: *Theoretical Computer Science* 616 (2016), pp. 70–93.
- [7] Paul B Callahan and S Rao Kosaraju. "A decomposition of multi-dimensional point sets with applications to k-nearest-neighbors and n-body potential fields." In: *Journal of the ACM (JACM)* 42.1 (1995), pp. 67–90.
- [8] Barun Chandra. "Constructing sparse spanners for most graphs in higher dimensions." In: *Information Processing Letters* 51.6 (1994), pp. 289–294.
- [9] Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. "New sparseness results on graph spanners." In: *Proceedings of the eighth annual symposium on Computational geometry*. 1992, pp. 192–201.
- [10] Boris V Dekster. "The Jung theorem for spherical and hyperbolic spaces." In: *Acta Mathematica Hungarica* 67.4 (1995), pp. 315–331.
- [11] Alon Efrat, Matthew J Katz, Frank Nielsen, and Micha Sharir. "Dynamic data structures for fat objects and their applications." In: *Computational Geometry* 15.4 (2000), pp. 215–227.
- [12] BE Flinchbaugh and LK Jones. "Strong connectivity in directional nearest-neighbor graphs." In: *SIAM Journal on Algebraic Discrete Methods* 2.4 (1981), pp. 461–463.

- [13] Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
- [14] Jim Ruppert and Raimund Seidel. "Approximating the d-dimensional complete Euclidean graph." In: *Proceedings of the 3rd Canadian Conference on Computational Geometry (CCCG 1991)*. 1991, pp. 207–210.
- [15] Michiel Smid. "An improved construction for spanners of disks." In: *Computational Geometry* 92 (2021), p. 101682.
- [16] A Frank van der Stappen, Dan Halperin, and Mark H Overmars. "The complexity of the free space for a robot moving amidst fat obstacles." In: *Computational Geometry* 3.6 (1993), pp. 353–373.
- [17] Marc Van Kreveld, Otfried Schwarzkopf, Mark de Berg, and Mark Overmars. *Computational geometry algorithms and applications*. Springer, 2000.
- [18] Andrew Chi-Chih Yao. "On constructing minimum spanning trees in k-dimensional spaces and related problems." In: *SIAM Journal on Computing* 11.4 (1982), pp. 721–736.