Diamonds are a Plane Graph's Best Friend

A Corrected, Improved, and Generalized Result for α -Diamond Spanners

By
Aaron Wojcik Lee

A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfilment of
the requirements for the degree of
Master of Computer Science

Ottawa-Carleton Institute for Computer Science
School of Computer Science
Carleton University
Ottawa, Ontario

September 2004

© Copyright
2004, Aaron Wojcik Lee

Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 0-612-97456-1

Our file Notre référence

ISBN: 0-612-97456-1

The author has granted a non-exclusive license allowing the Library and Archives Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

Canada

PAGINATION ERROR.

TEXT COMPLETE.

ERREUR DE PAGINATION.

LE TEXTE EST COMPLET.

Abstract

Let P be a set of points in the plane, and let p and q be two points of P . Denote the Euclidean distance between p and q as $|pq|$. Let $G = (V, E)$ be a geometric Euclidean graph defined on the set of points P . Denote the shortest path in G between p and q as $d_G(p, q)$. G is said to be a t -spanner for fixed $t \geq 1$ provided that $\frac{d_G(p, q)}{|pq|} \leq t$ for all points $p, q \in P$. Das and Joseph [7] show if a plane graph G has the α -Diamond property and the d -Good Polygon property, then G is a $\frac{8\pi^2 d}{\alpha^2 \sin^2(\alpha/4)}$ -spanner.

In this thesis, we provide a complete and thorough proof of this result [7], as well as correct errors and omissions in [7]. Further, we improve the spanning ratio of α -diamond spanners to $\frac{8(\pi-\alpha)^2 d}{\alpha^2 \sin^2(\alpha/4)}$. These results are extended to the constrained setting by modifying the definition of the α -diamond for visibility.

We show that the constrained Delaunay and constrained Greedy triangulations both satisfy the visible α -diamond property. An improved α value of $\frac{\pi}{6}$ for the Greedy triangulation is shown.

Acknowledgements

After all this time, I am probably not the only one who thought I might never finish this degree. I would quickly like to acknowledge a few people who gave me the strength to persevere.

First and foremost, my deepest thanks to my two supervisors, Jit Bose and Michiel Smid. Thank you for your guidance and support, and for seeing this project through to the end with me — it has been a life changing experience.

To my wife Karen, for giving me strength when I needed it most. We have been through a lot during the course of this degree, and I will never forget the journey.

To my family, both blood and by law, for supporting me (and feeding me) when I needed it. For being concerned and interested in my work, even at those times when I was not.

To the many friendly faces around the department, both faculty and staff, thank you for making my days at school better. And to all the colleagues and friends I've made during my time at Carleton, I'm glad we've been able to help each other out down this long road.

The research in this thesis was supported in part by NSERC and Nortel Networks. I also recommend to any person interested in geometry to look into the software *Cinderella*; without it, many concepts within this thesis might well have eluded me.

I would like to dedicate this thesis to the memory of my father, Robert Lee.

Contents

1	Introduction	1
1.1	Previous Work	3
1.2	Preview of Things to Come	6
2	Definitions and Preliminaries	7
2.1	Basic Graph Definitions	7
2.2	t -Spanners	8
2.3	Triangulations	9
2.4	Constrained Graphs and Visibility Graphs	10
2.5	Constrained t -spanners	13
3	The Spanning Ratio of α-Diamond Spanners	14
3.1	One-Sided Paths	18
3.2	Two-Sided Paths	26
3.2.1	Tangents between the Convex Chains	26
3.2.2	A Second Empty Region	27
3.2.3	Upper and Lower Boundaries	28
3.2.4	Constructing the Final Path	30
3.2.5	Case 1	32
3.2.6	Case 2 <i>a</i>)	35
3.2.7	Case 2 <i>b</i>)	43
3.3	The Final Spanning Ratio	48

4	The Visible α-Diamond Property	49
4.1	One-Sided Paths	51
4.2	Two-Sided Paths	52
5	Graphs that have the Diamond Property	54
5.1	Constrained Delaunay Triangulation	54
5.2	Constrained Greedy Triangulation	55
5.2.1	Trigonometry is Fun	58
5.2.2	The Required Construction	59
5.2.3	Proof by Cases	62
6	Conclusion	70
6.1	Future Work	71
A	Maple Output for <i>CGT</i> Proof	73
	Bibliography	80

List of Figures

1.1	The shortest route from City E to City F along the highway	1
1.2	The route around a lake is expected to be long	2
1.3	The α -Diamond property, and d -Good Polygon property	4
2.1	The visibility graph $vis(S)$ of $S = P \cup L$	11
2.2	Variations on how the visibility graph is defined	12
3.1	Every edge in the graph has an empty α -triangle	14
3.2	Walk along AB , labelling the edges you encounter	15
3.3	Constructing an upper convex chain between u_i and u_{i+1}	17
3.4	Augment the faces from A to B by adding convex chains	18
3.5	Construction for an upper one-sided path between p and q	19
3.6	Using shaded region T to approximate the length of the convex chains	19
3.7	Passing through an empty triangle $\Delta(e_i)$	20
3.8	Bounding the length of the upper boundary of an α -triangle	21
3.9	The triangles of G_θ	25
3.10	Building the tangents $t_{a,i}$ and $t_{b,i}$ between UC_i and LC_i	27
3.11	Every edge e_i has a 2 nd empty (shaded) region associated with it: $\Delta'(e_i)$	28
3.12	The boundary of T_U will approximate the length of the upper path .	29
3.13	Case 2: Construct an empty region	31
3.14	Extend the tangents t_a and t_b until they intersect	33
3.15	Tangents can stick out underneath each other	34
3.16	Diagram for The small Error	36
3.17	Diagram for The Critical Flaw	37

3.18	Bounding the length of the tangents for Lemma 5	39
3.19	It is possible for h to be negative	42
3.20	Violation of a one-sided path requires angle $\angle ulz_{i+1} > \alpha$	45
3.21	The empty trapezoid is too narrow for v to be above JK	46
4.1	Every edge e has at least one <i>visibly</i> empty α -triangle	49
4.2	A constrained edge must pass through an empty triangle $\Delta(e_i)$	51
5.1	Every edge e of $CDT(S)$ has a visibly empty circle	55
5.2	Every edge of the CGT has a $\frac{\pi}{6}$ empty triangle	57
5.3	The required construction for Theorem 4	59
5.4	The eight different combinations of cases	63
5.5	The location of X is restricted by the Greedy construction	65
5.6	The location of X can be determined by σ and ρ	66
5.7	Point Y depends on the location of X	67
5.8	The function $f_{\sigma,\rho}$ over the ranges $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$	68
5.9	New and previous exclusion regions for the Greedy triangulation . . .	69
A.1	Co-ordinates of points on the Cartesian plane	73
A.2	The function f over the ranges $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$	78
A.3	The function f over the range $\rho \in [0, \frac{\pi}{30}]$, when $\sigma = 0$ and $\frac{\pi}{10}$	79

Chapter 1

Introduction

The study of graphs and networks is interesting and important, both from a practical and theoretical perspective. Many real-world problems can be expressed as graph problems. For example, a national highway system of a country can be represented as a graph, where cities are represented by points, and highways are connections between these points. Obviously not all distances between cities are the same. We can therefore attach a weight to every highway, corresponding to the length of the highways between the cities, or perhaps the time it takes to travel from one to another. If the graph is simplified even further as to assume that highways between cities are straight lines, and the weight of each edge is the straight-line distance between cities, then this would be an example of a geometric graph, which will be defined more formally in the next chapter.

Figure 1.1: The shortest route from City E to City F along the highway

Consider the graph of a fictitious road network between eight cities in Figure 1.1. One might use such a map to plan a trip between two cities, determining which route would be the shortest. If a traveller uses the highways to get from City E to City F, he/she would have to travel quite far in comparison to the distance “as the crow flies,” as indicated by a dashed line in the figure. Obviously, this is not desirable: if two cities are not that far apart, it makes sense to want the distance travelled in the road network to not be much greater. But how much greater would be acceptable? If the citizens of these cities simply wanted to travel the least distance possible, they could build a highway between every pair of cities. However, such a road network would require $\binom{8}{2} = 28$ highways. In a country with 100 cities, it would mean almost 5,000 highways! This would be too costly, and require too much maintenance. Is it possible to have a road network where the number of highways is linearly proportional to the number of cities, and the shortest distance between any two cities is not more than, say, twice the straight-line distance? This is the informal definition of a spanner, which will be properly defined in Chapter 2.

Figure 1.2: The route around a lake is expected to be long

Yet what if the cities in consideration were situated around a large bay, or other impassable physical object, as in Figure 1.2. In such a situation, the shortest possible distance between any two cities might not necessarily be a straight line anymore. Consider a path between City A and City I. Here, the shortest possible distance

between them is to follow along the shore of the bay, as driving across the bay in a car is particularly difficult. The *constraint* of having to go around the bay modifies the definition of shortest distance. Therefore, the highway route from City A to City I, via B, C, E, and H, doesn't seem as out of the way as first thought, as it closely approximates the shortest distance around the bay (again, shown by the dashed line). This concept of having constraints in a graph which modify the shortest feasible distance between points, and still finding short paths through the network with these constraints in mind is part of the topic of constrained graphs and constrained spanners, and will be elaborated upon later.

This thesis examines the results of a conference paper by Das and Joseph [7], and the thesis by Das [6]. Since only small portions of the thesis are available, we will concentrate on the results as reported in the conference paper [7]. The paper defines a pair of properties: the α -diamond property and the d -good polygon property, both of which will be precisely defined in Section 1.1. Informally, the diamond property ensures that every edge within a graph has a certain size diamond-shaped region around it such that vertices of the graph can exist only on one side of the diamond. The good polygon property ensures that the distance around the boundary of every face in the graph is proportional to the distance across it. Das and Joseph [7] show that any graph which satisfies these two properties is well-connected in the sense described in this introduction. We correct some errors in [7], prove that this result extends into the constrained setting by modifying one of the properties, and show constrained versions of two well known graphs that have these properties.

1.1 Previous Work

Most of the results discussed in this section relate to graph spanners. Informally, a spanner is a graph which has the property that every pair of points in the graph is connected by a path, whose length is not longer than a constant times the actual distance between the points. We formalize this definition in Chapter 2.

Chew [4] was the first to show how to construct a planar graph with a constant spanning ratio (defined in Chapter 2). He proved that the Delaunay triangulation in

the L_1 norm is a $\sqrt{10}$ -spanner of the complete graph in the Euclidean norm.

Dobkin et al. [11] then derived a spanning ratio for the Delaunay triangulation in the usual L_2 norm, a value of $\pi(1 + \sqrt{5})/2 \approx 5.08$. Using an incremental construction, they showed how to construct a short path between every pair of vertices in the graph by exploiting properties of the Delaunay triangulation, and its dual, the Voronoi diagram.

A few years later, Keil and Gutwin [16] improved on the result in [11] by bounding the spanning ratio of the Delaunay triangulation with a value of $2\pi/(3\cos(\frac{\pi}{6})) \approx 2.42$. They used induction to show how to build a short path between any pair of vertices in the Delaunay triangulation.

Shortly after the paper by Dobkin et al. [11], Das and Joseph [7] proved that any graph which has the two properties mentioned in the previous section (and defined below) is a spanner, whose spanning ratio depends on these properties. The techniques used in the paper were similar to those employed by Dobkin et al. [11]. Das and Joseph [7] went on to prove that the Delaunay, Greedy, and Minimum Weight triangulations all had these particular properties, which are now formally defined:

Figure 1.3:

Definition 1. Refer to Figure 1.3(a). Fix $\alpha \in (0, \pi/2)$. A plane graph G is said to have the α -diamond property if, for every edge e in the graph, the two isosceles triangles, having e as the base and base angle α , has vertices of G in at most one of the two triangles. That is, at least one of the triangles (not including the boundary of the region) must be empty of vertices of G . Throughout this thesis, we label the empty triangle as $\Delta(e)$, and the apex of $\Delta(e)$ as $v(e)$.

Definition 2. Refer to Figure 1.3(b). Fix $d \geq 1$. A plane graph G is said to have the d -good polygon property if for every face f in G , and for every pair of vertices a and b of f whose segment is a chord of f , the shortest distance from a to b around the boundary of f is at most d times the Euclidean distance between a and b .

Das and Joseph [7] showed that if a graph G has the diamond property for some angle α , and the good polygon property for some $d \geq 1$, then G is a t -spanner, for $t = \frac{8\pi^2 d}{\alpha^2 \sin^2(\alpha/4)}$. Note that for triangulations, the value of $d = 1$ since every face in the graph is a triangle. They also showed that the Delaunay triangulation has the $\frac{\pi}{4}$ -diamond property, while the Greedy and Minimum Weight triangulations both have the $\frac{\pi}{8}$ -diamond property.

Dickerson, Drysdale, McElfresh, and Welzl [10] proved that the Greedy triangulation (GT) has a *disc* shaped type of empty region, or exclusion region. They showed that for any edge pq of the GT , a circle of diameter $|pq|/\sqrt{5}$ centered at the midpoint of pq has no vertices of the graph in at least one half of the circle, either above or below pq . Drysdale, Rote, and Aichholzer [13] later showed that this exclusion region can be enlarged by including the areas under the tangents to the circle from p and q .

Drysdale, McElfresh, and Snoeyink [12] improved upon the results in [7], proving that the Minimum Weight triangulation (MWT) has a larger diamond shaped empty region, with base angle $\pi/4.6$. They also showed that the MWT has a disc shaped empty region of radius $|pq|/\sqrt{2}$ centered about the midpoint of pq .

Knowing that a graph has a particular empty region associated with it is helpful. For example, Beirouti and Snoeyink [1] use the diamond shaped region of the Minimum Weight triangulation as an edge test to improve the performance of their algorithm which calculates the MWT of a set of uniformly-distributed points.

1.2 Preview of Things to Come

The remainder of the thesis is summarized as follows:

Chapter 2: This chapter introduces some graph theory and computational geometry definitions. It covers the nomenclature used in the discussion of spanners and constrained graphs, and provides a framework and motivation for this thesis.

Chapter 3: In this chapter, we examine the previous work by Das and Joseph [7]. The entire proof of the spanning ratio being derived from the diamond property and the good polygon property is examined, and every detail of the proof is fully explained. An error and a few small omissions in the original paper are noted, and corrected. Additionally, an improvement to the function bounding the spanning ratio is provided.

Chapter 4: The α -diamond property is extended to the constrained setting by defining the visible α -diamond property. We show that a constrained graph with this property, and the d -good polygon property, is a constrained spanner whose spanning ratio is based on α and d .

Chapter 5: Here we consider two constrained graphs: the constrained Delaunay triangulation, and the constrained Greedy triangulation. We show how both of these graphs have the visible α -diamond property, for values of α equal to $\frac{\pi}{4} = 45^\circ$ and $\frac{\pi}{6} = 30^\circ$ respectively. The α value for the Greedy triangulation is an improvement over the currently best known value for diamond-shaped exclusion regions, which is $\arctan(\frac{1}{\sqrt{5}}) \approx 24.1^\circ$ by [10] and [13].

Chapter 6: This chapter summarizes the results of the previous chapters, and provides some interesting problems that could be considered for future work.

Appendix A: Part of the proof for the constrained Greedy triangulation in Chapter 5 is shown here in its entirety. The proof required the use of the mathematical software Maple to aid in the derivation of some functions.

Chapter 2

Definitions and Preliminaries

Now that the context and motivation for this thesis has been established in Chapter 1, more precise definitions of previously mentioned constructions are presented. Most of the definitions and concepts in this chapter are well known in the literature (see [14] and [20]).

2.1 Basic Graph Definitions

Definition 3. A *graph* $G = (V, E)$ consists of a finite set V of vertices, and a set E of unordered pairs of distinct elements of V , called edges.

Definition 4. A *geometric Euclidean graph* is a graph whose vertex set is a set of points in the plane, and whose edge set is a set of line segments joining pairs of vertices. The edges are weighted, where the weight of each edge is the Euclidean length of the segment.

Definition 5. A graph is called *planar* if it can be drawn in the 2-dimensional plane as a set of points and line segments such that no two edges cross, except at a common endpoint. Such a drawing is called an *embedding*.

Definition 6. A *plane graph* is a planar graph that has been embedded in the plane such that no two edges cross.

Although an edge in a plane graph may be a curve, throughout this thesis we will assume that all edges in plane graphs are straight line segments. We will also assume that when discussing geometric graphs, the Euclidean metric will be implied, unless otherwise stated. Lastly, we assume that the embedded vertices of the graph are in general position in the plane, which means that no three vertices are co-linear, and no four vertices are co-circular.

2.2 t -Spanners

A spanner graph is a particular type of geometric graph. If a graph is a spanner, then the shortest distance between any two vertices in the graph is proportional to the Euclidean distance between these points. This is formalized in the following:

Definition 7. Let $t \geq 1$ be a real number. A geometric graph G is said to be a t -spanner if, for every pair of vertices p and q in G , the length of the shortest path between them in the graph, $d_G(p, q)$, is at most t times the Euclidean distance between the points, or:

$$\frac{d_G(p, q)}{|pq|} \leq t$$

The smallest such value t is called the *spanning ratio* of the spanner.

When the vertices of the graph are in general position, as the spanning ratio approaches 1, the required number of edges must approach $\binom{n}{2}$, where n is the number of vertices in the graph. Spanners have obvious applications to any situation where generating short paths within sets of points is required. Examples include communication routing, shortest path planning in maps, etc. Refer to [20] for a discussion on spanners.

The road network example at the beginning of Chapter 1 (Figure 1.1) can now be expressed in terms of spanners. If the highway engineers were successful in designing a system of roads such that the distance between any two cities is not more than twice the distance it would take to fly, then the highway network would be a 2-spanner of the cities.

2.3 Triangulations

Triangulations of planar point sets are well known and well studied structures, both in a geometric sense and topological sense. The definition commonly used in geometry for a triangulation is:

Definition 8. A *triangulation* of a point set P is a plane graph, with vertex set P , that cannot have another edge inserted between vertices of P without inducing an edge crossing. In a triangulation, every face of the graph is a triangle except possibly the outer exterior face.

Since triangulations are planar, they have a linear number of edges. There has been significant work done to prove that various types of triangulations are spanners (see for example: [7] [11] [4] [15] [16]). We review a few of the results pertinent to this thesis.

One of the most famous triangulations is the Delaunay triangulation (DT), named after the Russian mathematician Boris Nikolaevich Delone (Delaunay is the French transliteration) who discovered the structure [9]. The Delaunay triangulation of a point set P in general position can be defined in several equivalent ways:

- Consider every triple of points in P , and construct the unique circle through these three points. If this circle contains no points of P in its interior, then these three points will have edges between them forming one of the triangles.
- Consider every pair of points in P . If it is possible to construct a circle with these two points on its boundary such that the circle contains no points of P in its interior, then there is an edge between these two vertices in the Delaunay triangulation.
- It is the geometric dual of the Voronoi diagram of P [9].
- It is the angle-optimal triangulation of P (Theorem 9.9 of [8]).

Another well-known triangulation is the Greedy triangulation (GT) [18]. A simplified algorithm for constructing it on a point set P is:

- Consider every pair of points in P , and sort them into a non-decreasing order by the lengths of the segments connecting each pair.
- For each pair of points in this sorted order, insert the corresponding edge into the partially completed triangulation if and only if the edge does not properly intersect any edge of the triangulation that has been inserted prior to it.

As the complete graph has $\binom{n}{2}$ edges, where $n = |P|$, the greedy algorithm described above runs in $O(n^2 \log n)$ time. This is, however, a brute-force type construction, and there are faster ways of computing the Greedy triangulation [21].

One of the most enigmatic triangulations to compute is the Minimum Weight triangulation (*MWT*). Quite simply, it is defined as the triangulation, out of every possible triangulation of a point set, that has the least total edge weight [18]. It is still not known whether computing the *MWT* is NP-complete.

2.4 Constrained Graphs and Visibility Graphs

The results discussed in Section 1.1 were proven for graphs that are unconstrained — graphs that had no restrictions on where edges could be placed. However, consider our second example at the beginning of Chapter 1 (Figure 1.2). In that situation, the bay acted as a barrier, or constraint, which prevented highways from crossing it. In this thesis, we shall use the following definition of a constrained graph. Let P be a set of points, and L a set of non-crossing line segments whose endpoints are in P . Let $S = P \cup L$.

Definition 9. A *constrained graph* $G(S) = (V, E)$ is a plane graph which consists of a finite set V of vertices such that $V = P$, and a set E of segments such that $L \subseteq E$, and no edge in $E \setminus L$ properly intersects any edge of L (i.e. only intersect at endpoints).

Definition 10. The edges of a constrained graph $G(S) = (V, E)$ in set L are referred to as *constrained edges*, or *constraints*. Similarly, the edges in $E \setminus L$ are called *unconstrained edges*.

Hence, the edges in L act as barriers, blocking visibility between the vertices of P . We formalize the concept of visibility in the following:

Definition 11. Two points p and q of S are *visible* if and only if the segment pq does not intersect any constrained edge of L .

Definition 12. A point p and segment e of S are visible if and only if p is visible to at least one point on e .

From these definitions, it is clear that the endpoints of every unconstrained edge in E of a constrained graph $G(S)$ are visible to each other. The edge set E of $G(S)$ is a subset of another graph, which contains the constrained edges L , and all possible visible edges of the graph. This graph is called the visibility graph of S :

Figure 2.1: The visibility graph $vis(S)$ of $S = P \cup L$

Definition 13. The *visibility graph* of S , denoted $vis(S)$, is a graph containing all possible unconstrained edges between the vertices of S that are visible.

Constrained graphs are sometimes referred to as *generalized* graphs, simply because if a constrained graph had no constraints ($L = \emptyset$), then it would be unconstrained, and its visibility graph would be the complete graph.

There are several ways a visibility graph can be defined, as in Figure 2.2. In 2.2(a), the shortest path between p and q goes through the constrained segments, using one of their vertices along the way. In Figure 2.2(b), the polygonal chain of constrained edges acts as a wall, forcing the path from p to q to go completely around the constraints.

Figure 2.2: Variations on how the visibility graph is defined

Both visibility constructions have useful applications. For example, the constraints in 2.2(a) simply act as placeholder edges, ensuring those edges get used in the final graph. They really are not barriers as the ones in Figure 2.2(b) are. This second example of constraints could be used to represent high terrain in line-of-sight communication, or a lake which must be navigated around.

Regardless of which type of visibility graph is used, it will not affect the results discussed in this thesis. It simply changes how the shortest constrained distance between vertices is defined.

Several interesting results have been shown for constrained graphs. For example, the constrained Delaunay triangulation (*CDT*) was first called generalized Delaunay triangulation by Lee [17], and later an obstacle triangulation by Chew [4]. It is a constrained triangulation as close to the regular Delaunay triangulation as possible. That is, for every edge e in the *CDT*, there exists a circle C_e through the endpoints of e such that C_e has no vertices in its interior that are visible to e . This differs from the regular Delaunay triangulation simply because C_e would not be allowed to have any vertices in its interior. Chew [5], and Wang and Schubert [22] later proved that the *CDT* can be built in $O(n \log n)$ time.

2.5 Constrained t -spanners

Extending the concept of a t -spanner into the domain of the constrained graph results in the following definition:

Definition 14. Let $t \geq 1$ be a real number. A constrained geometric graph $G(S)$ is said to be a *constrained t -spanner* if, for every pair of vertices p and q in $G(S)$, the length of the shortest path between them in the graph, $d_{G(S)}(p, q)$, is at most t times the shortest distance between p and q in the visibility graph $vis(S)$, or:

$$\frac{d_{G(S)}(p, q)}{d_{vis(S)}(p, q)} \leq t$$

Recently, Guibas and Karavelas [15] show that the technique of Dobkin et al. [11] used to prove a spanning ratio of the Delaunay triangulation can be extended to the constrained setting. That is, for every pair of vertices p and q of the constrained Delaunay triangulation, a path can be constructed between p and q whose length is at most $\frac{\pi(1+\sqrt{5})}{2} \cdot d_{vis(S)}(p, q)$.

Chapter 3

The Spanning Ratio of α -Diamond Spanners

Figure 3.1: Every edge in the graph has an empty α -triangle

This chapter examines the central result from the paper by Das and Joseph [7], which is the basis of this thesis. We give a complete proof of the entire result, and explain details that are left out of the original paper. An error in [7] which breaks the result is discussed, and a correction for it is made. Further, an improvement to the function bounding the spanning ratio of α -diamond spanners is provided.

Recall the definitions of the α -diamond property, and the d -good polygon property from Chapter 1 (Page 5). See Figure 3.1 for an example. Intuitively, both of these properties together seem to produce graphs with low spanning ratios: the diamond property ensures that no edge of the graph can act as a very large obstacle between any pair of vertices, which would cause a major detour on the path between these vertices. The good polygon property guarantees that we can always get from one side of a face to the other without travelling “too far.” We formalize this intuition regarding these properties in the following theorem:

Theorem 1. Fix $\alpha \in (0, \pi/2)$ and $d \geq 1$. If a plane graph G has both the α -diamond property and the d -good polygon property, then for every pair of vertices A and B in G , there exists a path in G between A and B whose length is at most $c_{\alpha,d}|AB|$, where $c_{\alpha,d} = \frac{8(\pi-\alpha)^2 d}{\alpha^2 \sin^2(\alpha/4)}$.

We assume that π is a multiple of α ; this simplifies later calculations, without affecting the final result. The proof of Theorem 1 will comprise the remainder of this chapter.

Figure 3.2: Walk along AB , labelling the edges you encounter

Select two vertices A and B in the graph G . Reorient the graph such that the segment AB is on the x -axis, and A lies to the left of B , as in Figure 3.2. If no edges of G intersect AB , then A and B must lie on the same face, and by the good polygon property, the shortest path between them has length at most $d|AB|$. Since $d < c_{\alpha,d}$, the theorem holds. Otherwise, label the edges of G that intersect AB in order from A to B as $e_1, e_2, e_3, \dots, e_{k-1}, e_k$.

Since each edge e_i is an edge in G , then by the α -diamond property, at least one of its adjacent α -triangles must be empty. If both triangles are empty, then arbitrarily choose one of them. Label the empty triangles associated with each edge as $\Delta(e_1), \Delta(e_2), \dots, \Delta(e_k)$ respectively. Now label the apex of each of these empty triangles $v(e_1), v(e_2), \dots, v(e_k)$ accordingly. Note that these apices are not actual vertices in G — they are simply apices of the empty triangles.

For a given edge e_i , by definition, it crosses AB . Therefore, it has one endpoint above AB , and one below AB ; label its upper vertex u_i , and its lower vertex l_i . We will assume that no adjacent edges e_i and e_{i+1} share a common vertex. It will be clear that when edges share common vertices, the proof is simpler.

The goal is to construct a path from A to B whose length is at most $c_{\alpha,d}|AB|$. The path from A to B will be constructed by walking around the boundaries of the faces that are adjacent to the aforementioned k edges. By planarity, adjacent edges in the ordering, e_i and e_{i+1} , lie on the same face in the graph. Similarly, A and e_1 line on the same face, as well as B and e_k lie on the same face.

During the construction of the path from A to B , there must be a way of getting from each u_i to the next u_{i+1} , and similarly for the lower vertices. Since each u_i and u_{i+1} lie on the same face, by employing the good polygon property we can ensure that the length of the path between these adjacent pairs of upper (and lower) vertices is bounded by d times the shortest path between them inside the face.

Consider the face between e_i and e_{i+1} as in Figure 3.3 — for ease of reference, label this face f . We restrict our analysis to the upper boundary of f between u_i and u_{i+1} . This upper portion of the boundary, which is a chain of edges, does not intersect AB , since e_i and e_{i+1} are adjacent in the ordering of edges from A to B .

Build a convex chain from u_i to u_{i+1} by taking the *lower convex hull* of the upper

Figure 3.3: Constructing an upper convex chain between u_i and u_{i+1}

boundary of f . The entire upper boundary of f between u_i and u_{i+1} is thus contained above this convex chain. Hence, each segment of the convex chain must either be an edge of the boundary of f , or must lie completely inside the face, and is therefore a chord of f . Therefore, by the good polygon property, we know that the length of each segment of the convex chain is a d -approximation to the shortest path along the boundary of f between the endpoints of this segment. That is, if the total length of the convex chain from u_i to u_{i+1} is ψ , then there exists a path around the boundary of f from u_i to u_{i+1} whose length is at most $d\psi$. We will refer to this chain as UC_i , and denote its length by $|UC_i|$.

In a symmetric way, construct an *upper* convex hull of the lower boundary of f between l_i and l_{i+1} , and label it LC_i . By symmetry, we know that the length of LC_i d -approximates the shortest path from l_i to l_{i+1} around f .

Repeat this construction process as shown in Figure 3.4 for all upper and lower vertices of adjacent pairs of edges. Additionally, construct *lower* convex hulls UC_0 between A and u_1 , as well as UC_k between u_k and B . Similarly, build *upper* convex hulls LC_0 between A and l_1 , and LC_k between l_k and B .

Figure 3.4: Augment the faces from A to B by adding convex chains

By using these convex chains, $UC_0, UC_1, \dots, UC_k$ and $LC_0, LC_1, \dots, LC_k$ to construct the path between A and B , any such path will be a d -approximation to a corresponding path using existing edges in the original graph G . We can therefore shift our attention to finding a path of length at most $\frac{8(\pi-\alpha)^2}{\alpha^2 \sin^2(\alpha/4)} |AB|$ from A to B that traverses the edges of these convex chains.

3.1 One-Sided Paths

We introduce the notion of one-sided paths as defined by Das and Joseph [7]. For the remainder of this section, we assume that each of the apices between A and B , $v(e_1), v(e_2), \dots, v(e_k)$, is located on the same side of AB . Further, without loss of generality, assume the apices lie in the half-plane h^- below the x -axis. In this case, the path from A to B is constructed by walking along the upper convex chains $UC_0, UC_1, \dots, UC_k$ above the x -axis, from $A \rightarrow u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_k \rightarrow B$. We call this construction a *one-sided path*:

Figure 3.5: Construction for an upper one-sided path between p and q

Definition 15. Refer to Figure 3.5. Let p and q be vertices of G , and let $e_1, e_2, \dots, e_n$ be the edges of G that intersect pq . Assume that the apices $v(e_1), v(e_2), \dots, v(e_n)$ are all located on the same side of pq . A *one-sided path* from p to q is constructed by following either the upper convex chains $UC_0, UC_1, \dots, UC_n$ if the apices are located below pq , or the lower convex chains $LC_0, LC_1, \dots, LC_n$ if the apices are above pq .

We prove that one-sided paths have a bounded spanning ratio by bounding the length of the convex chains which make up the one-sided path. This will be accomplished by using the empty triangles $\Delta(e_1), \Delta(e_2), \dots, \Delta(e_k)$ to approximate the length of the convex chains.

Figure 3.6: Using shaded region T to approximate the length of the convex chains

Lemma 1. *Refer to Figure 3.6. Let $T = \bigcup_{i=1}^k \Delta(e_i) \cup h^-$. That is, T is the union of all the empty triangles together with the half-plane below AB . Denote the upper boundary of T between A and B as $T(A, B)$. The length of $T(A, B)$ is at least the length of the upper convex chains from A to B :*

$$|UC_0| + |UC_1| + \cdots + |UC_k| \leq |T(A, B)|$$

Proof. Note that since every triangle $\Delta(e_1), \Delta(e_2), \dots, \Delta(e_k)$ is an empty region, no vertices of any convex chain $UC_0, UC_1, \dots, UC_k$ can lie inside any of the triangles. On the other hand, it seems possible for a segment of a convex chain to pass *through* a triangle (as seen in Figure 3.4 on page 18 between u_2 and u_k). However, also note the location of the apex of e_2 : $v(e_2)$ is above the x -axis. When all of the $v(e_i)$'s are below AB , no segment of any upper convex chain UC_i can pass through any of the empty triangles.

Figure 3.7: Passing through an empty triangle $\Delta(e_i)$

To see this, consider the contrary: some segment zz' of one of the convex chains above AB passes through the empty triangle $\Delta(e_i)$. Since there are no vertices of G in $\Delta(e_i)$, the endpoints z and z' are outside the triangle, and the segment must enter and exit the empty triangle, piercing it twice. The convex chains exist strictly inside the faces of G and therefore cannot intersect any edges of G , so zz' cannot cross e_i . It must therefore pierce $\Delta(e_i)$ through its other two sides, as in Figure 3.7. However,

we know that $v(e_i)$ lies below the x -axis, as does l_i . Therefore the lower side of the triangle, $l_i v(e_i)$, lies completely below the x -axis as well. So if zz' exits the empty triangle by crossing $l_i v(e_i)$, it must also cross AB . However, this contradicts the fact that the convex chain lies completely above AB .

Hence, between each u_i and u_{i+1} , we have a convex shape UC_i that is entirely contained above the boundary of T between u_i and u_{i+1} , which we denote $T(u_i, u_{i+1})$. Moreover, because UC_i is convex, its length *must* be at most $T(u_i, u_{i+1})$ (page 42 of [2]). Therefore, the total length of all the convex chains $|UC_0| + |UC_1| + \dots + |UC_k|$ must be at most $|T(A, B)|$. $\square$

Lemma 1 implies that $d \cdot |T(A, B)|$ is an upper bound on the length of the actual one-sided path in G from A to B . Before proving a numeric bound on the length of $T(A, B)$, a useful result must first be shown.

Figure 3.8: Bounding the length of the upper boundary of an α -triangle

Lemma 2. *The ratio of the length of the upper boundary of a triangle, with upper angle α , compared to its horizontal base is at most $\frac{1}{\sin(\alpha/2)}$.*

Proof. Refer to Figure 3.8. Let the length of the base segment equal 1. The unknown lengths of the other two sides are labelled X and Y . Hence, the ratio of the upper boundary to the base is simply $X+Y$. By sine law:

$$\frac{1}{\sin \alpha} = \frac{X}{\sin(\alpha + \beta)} = \frac{Y}{\sin \beta}$$

Note that $\sin(\pi - (\alpha + \beta)) = \sin(\alpha + \beta)$. To obtain $X + Y$, rearrange the above equations to get:

$$X + Y = \frac{\sin(\alpha + \beta) + \sin \beta}{\sin \alpha} \quad (3.1)$$

To maximize the value of this function for a given α , differentiate with respect to β :

$$\frac{\partial(X + Y)}{\partial \beta} = \frac{\cos(\alpha + \beta) + \cos \beta}{\sin \alpha} \quad (3.2)$$

and solve for β with the above equation set to 0:

$$\begin{aligned} \frac{\cos(\alpha + \beta) + \cos \beta}{\sin \alpha} &= 0 \\ \cos(\alpha + \beta) + \cos \beta &= 0 \\ \cos(\alpha + \beta) &= -\cos \beta \\ \cos(\alpha + \beta) &= \cos(\pi - \beta) \\ \alpha + \beta &= \pi - \beta + 2\pi m, \text{ for some } m \in \mathbb{Z} \\ 2\beta &= \pi - \alpha + 2\pi m \\ \beta &= \frac{\pi}{2} - \frac{\alpha}{2} + \pi m \end{aligned}$$

Since $\alpha \in (0, \frac{\pi}{2})$, and $\beta \in [0, \pi - \alpha)$:

$$\begin{aligned} \pi m &= \beta - \frac{\pi}{2} + \frac{\alpha}{2} > 0 - \frac{\pi}{2} + 0 = -\frac{\pi}{2} \\ \pi m &= \beta - \frac{\pi}{2} + \frac{\alpha}{2} < \pi + \frac{\pi}{2} + \frac{\pi}{4} = \frac{3\pi}{4} \\ \therefore -\frac{\pi}{2} &< \pi m < \frac{3\pi}{4} \\ \therefore -\frac{1}{2} &< m < \frac{3}{4} \\ \therefore m &= 0 \\ \therefore \beta &= \frac{\pi - \alpha}{2} \end{aligned} \quad (3.3)$$

This proves that there is only one local extremum. To show that Equation (3.1) is maximized at $\beta = \frac{\pi - \alpha}{2}$, we must show that Equation (3.2) is positive for $\beta < \frac{\pi - \alpha}{2}$, and negative for $\beta > \frac{\pi - \alpha}{2}$ (i.e. the slope changes from positive to negative).

When $\beta = 0$:

$$\frac{\cos(\alpha + 0) + \cos(0)}{\sin \alpha} = \frac{\cos(\alpha) + 1}{\sin(\alpha)}$$

Since each of these three values is positive when $\alpha \in (0, \frac{\pi}{2})$, this function is also positive.

When $\beta = \pi - \alpha$:

$$\frac{\cos(\alpha + (\pi - \alpha)) + \cos(\pi - \alpha)}{\sin \alpha} = \frac{-1 + \cos(\pi - \alpha)}{\sin(\alpha)}$$

With $\alpha \in (0, \frac{\pi}{2})$, both terms in the numerator are negative. Since the denominator is positive, this function is always negative. Therefore, Equation (3.1) is maximized when $\beta = \frac{\pi - \alpha}{2}$.

Continuing, substitute Equation (3.3) into Equation (3.1), to arrive at:

$$\begin{aligned} X + Y_{\max} &= \frac{\sin(\alpha + \frac{\pi - \alpha}{2}) + \sin(\frac{\pi - \alpha}{2})}{\sin \alpha} \\ &= \frac{\sin(\frac{\pi}{2} + \frac{\alpha}{2}) + \sin(\frac{\pi}{2} - \frac{\alpha}{2})}{\sin \alpha} \\ &= \frac{\cos(\frac{\alpha}{2}) + \cos(\frac{\alpha}{2})}{\sin \alpha} \\ &= \frac{2 \cos(\frac{\alpha}{2})}{2 \sin(\frac{\alpha}{2}) \cos(\frac{\alpha}{2})} \\ &= \frac{1}{\sin(\frac{\alpha}{2})} \end{aligned}$$

as required. □

The results of Lemma 1 and Lemma 2 can now be combined to show a bound on the length of the one-sided path.

Lemma 3. *The length of a one-sided path from A to B in G is at most $c_{\text{OSP}}|AB|$, where $c_{\text{OSP}} = \frac{2(\pi - \alpha)d}{\alpha \sin(\alpha/4)}$.*

Proof. By Lemma 1, and the d -good polygon property, it suffices to show that $|T(A, B)| \leq \frac{2(\pi - \alpha)|AB|}{\alpha \sin(\alpha/4)}$. Note that since the apex $v(e_i)$ of every empty triangle $\Delta(e_i)$ is below AB , only two sides of $\Delta(e_i)$ lie in the upper half-plane h^+ . Hence, there is a well-defined left edge and right edge for the portion of $\Delta(e_i)$ that lies above AB .

Since π is a multiple of α , partition the triangles $\Delta(e_1), \Delta(e_2), \dots, \Delta(e_k)$ into $\frac{2(\pi-\alpha)}{\alpha}$ groups labelled $G_0, G_{\alpha/2}, \dots, G_\theta, \dots, G_{\pi-3\alpha/2}$, such that the left edges of the empty triangles in group G_θ make an angle in the range $[\theta, \theta + \frac{\alpha}{2}]$ with the x -axis. Note that since the apex of every empty triangle is below AB , there is a well-defined left edge and right edge above AB . Thus their right edges will be in the range $[\theta + \alpha, \theta + \frac{3\alpha}{2}]$ with the x -axis.

We now point out a small detail from the original paper [7] that leads to an improvement in the spanning ratio of the one-sided path. Das and Joseph [7] divide the empty triangles into $\frac{2\pi}{\alpha}$ groups, rather than $\frac{2(\pi-\alpha)}{\alpha}$ groups. The reason that fewer groups are needed becomes evident when considering the angles that the right edges make with the x -axis. With $\frac{2\pi}{\alpha}$ groups, the last group is $G_{\pi-\alpha/2}$. Therefore, the range of angles that the right edges of this group would make with the x -axis are:

$$\left[\left(\pi - \frac{\alpha}{2} \right) + \alpha, \left(\pi - \frac{\alpha}{2} \right) + \frac{3\alpha}{2} \right] = \left[\pi + \frac{\alpha}{2}, \pi + \alpha \right]$$

Clearly, the largest angle that a right edge can make with the x -axis is π . For this reason, the last two groups described in [7] are unnecessary, and the upper bound of the spanning ratio is improved slightly from this observation.

Continuing, let T_θ be the union of all the triangles in G_θ with the half-plane h^- below the x -axis. Recall T from Lemma 1. Note that:

$$T = T_0 \cup T_{\alpha/2} \cup \dots \cup T_{\pi-\alpha/2}$$

Hence, it follows that the length of the boundary of T from A to B is bounded by the sum:

$$|T(A, B)| \leq |T_0(A, B)| + |T_{\alpha/2}(A, B)| + \dots + |T_{\pi-\alpha/2}(A, B)|$$

Consider the group G_θ , as shown in Figure 3.9. Also note the boundary $T_\theta(A, B)$ shown as a thick black line between A and B . Define point P such that $\angle APB = \frac{\alpha}{2}$, $\angle PAB = \theta + \frac{\alpha}{2}$, and $P \in h^+$. Note that the portion of the triangles in G_θ that lie above the x -axis are completely contained inside $\triangle PAB$, and as such $T_\theta(A, B)$ is also inside $\triangle PAB$.

Figure 3.9: The triangles of G_θ

$T_\theta(A, B)$ is a chain of segments: segments of left edges, segments of right edges, and segments of AB . We will pay for each of these segments using portions of PA and PB . Draw lines parallel to the left edge PA at every vertex of the chain $T_\theta(A, B)$, until intersecting the right edge PB . Translate each piece of the segmented PB downwards, parallel to PA , such that every corresponding segment of $T_\theta(A, B)$ has a piece of PB touching it, as shown by in Figure 3.9. Similarly, draw lines parallel to the right edge PB at every vertex of the boundary, and translate the pieces of PA parallel to PB onto $T_\theta(A, B)$. Note that the segmented PA and PB are now triangles sitting on each segment of the boundary $T_\theta(A, B)$. Due to the geometric construction of $\triangle PAB$ and the angles of the edges in $T_\theta(A, B)$, by the triangle inequality, each segment of $T_\theta(A, B)$ can be paid for by a disjoint pair of segments from PA and PB .

Hence,

$$|T_\theta(A, B)| \leq |AP| + |PB|$$

Referring to Lemma 2, since the upper angle of $\triangle APB$ is $\frac{\alpha}{2}$, it follows that $|AP| + |BP| \leq \frac{|AB|}{\sin(\alpha/4)}$, regardless of the angle θ . Therefore, as there are $\frac{2(\pi-\alpha)}{\alpha}$ many groups, $|T(A, B)| \leq \frac{2(\pi-\alpha)|AB|}{\alpha \sin(\alpha/4)}$. $\square$

3.2 Two-Sided Paths

We have seen that if every apex $v(e_1), v(e_2), \dots, v(e_k)$ lies on the same side of the x -axis, then a one-sided path from A to B can be constructed, whose length is $\frac{2(\pi-\alpha)d}{\alpha \sin(\alpha/4)}$ times the length $|AB|$. But what if some $v(e_i)$'s lie above the x -axis, while others lie below? We now outline the process of constructing a short path from A to B in this case. The path will consist of disjoint pieces of upper UC_i and lower LC_i chains. For this, a way of crossing between the upper and lower chains within each face is needed.

3.2.1 Tangents between the Convex Chains

We introduce the concept of *tangents*. Between every pair of upper and lower chains UC_i and LC_i , except those incident to A or B (namely UC_0, LC_0, UC_k , and LC_k), two tangents will be constructed.

Consider the face f bounded by e_i and e_{i+1} , as in Figure 3.10, and suppose that both UC_i and LC_i are closed convex polygons, by adding the segments $u_i u_{i+1}$ and $l_i l_{i+1}$ respectively. Construct two *inner tangents* (or supporting lines) between UC_i and LC_i as detailed on pages 119-123 of [19]. Note that these tangents must exist strictly within f . Label the two tangents $t_{a,i}$ and $t_{b,i}$ such that the line through $t_{a,i}$ has vertex A on one side of it, and UC_i on the other. The other inner tangent is $t_{b,i}$.

Every face between e_1 and e_k is now augmented with two tangents to cross between the upper and lower set of convex chains. Note that each of these inner tangents lie completely inside their respective faces, and therefore must have their endpoints visible inside the face. Hence, the good polygon property applies, and the distance

Figure 3.10: Building the tangents $t_{a,i}$ and $t_{b,i}$ between UC_i and LC_i

around the boundary of the face between the upper and lower vertices of every tangent cannot be longer than d times the length of the tangent.

3.2.2 A Second Empty Region

Next, consider an edge e_i . We define an additional empty region $\Delta'(e_i)$ with respect to this edge. If the apex $v(e_i)$ is situated below AB , as in Figure 3.11, then $\Delta'(e_i)$ is defined as the region above AB , and below $t_{b,i-1}$, UC_{i-1} , UC_i , and $t_{a,i}$. Clearly, each $\Delta'(e_i)$ region is empty of vertices of G — this can be seen since the area between the upper and lower convex chains is empty of vertices of G (by definition), and $\Delta'(e_i)$ is contained strictly within this area. Note that the portion of the edge's empty triangle $\Delta(e_i)$ above AB need not be completely contained inside $\Delta'(e_i)$, unlike the illustration in Figure 3.11: it is possible for the tangent $t_{a,i}$ to extend from u_i down through the empty triangle.

If $v(e_i)$ is above AB , then $\Delta'(e_i)$ is symmetrically defined, and the region is bounded by AB , $t_{a,i-1}$, LC_{i-1} , LC_i , and $t_{b,i}$.

Figure 3.11: Every edge e_i has a 2nd empty (shaded) region associated with it: $\Delta'(e_i)$

3.2.3 Upper and Lower Boundaries

Divide the set of edges $e_1, e_2, \dots, e_k$ into two disjoint groups, U and L . U contains the edges that have their apex below AB , and L contains the edges whose apex is above AB . For the first group, define the region $T_U = h^- \cup \bigcup_{e \in U} (\Delta(e) \cup \Delta'(e))$, i.e., the union of every edge's empty triangle $\Delta(e)$, $e \in U$, and additional empty region $\Delta'(e)$, $e \in U$, with the half-plane h^- below AB , as illustrated in Figure 3.12. This region is similar to region T of Lemma 1. Correspondingly, define the region $T_L = h^+ \cup \bigcup_{e \in L} (\Delta(e) \cup \Delta'(e))$, i.e. the union of all empty triangles that have their apex above AB , their respective empty regions $\Delta'(e)$, $e \in L$, as well as the half-plane h^+ above AB . Let $T_U(A, B)$ denote the upper boundary of T_U between A and B ; similarly $T_L(A, B)$ for the lower boundary. Note that the length of $T_U(A, B)$ and $T_L(A, B)$ are each less than $c_{\text{OSP}}|AB|$, where c_{OSP} is as in Lemma 3.

To see why this is, the consider the region $\overline{T_U} = h^- \cup \bigcup_{e \in U} \Delta(e)$, i.e. just the

Figure 3.12: The boundary of T_U will approximate the length of the upper path

empty triangles. The length of the boundary of this region between A and B is bounded by $c_{\text{OSP}}|AB|$, as all of the empty triangles have their apices below AB , and is therefore analogous to the region T in Lemma 1. Consider the boundary of T_U between adjacent upper vertices u_p and u_q . Adding the secondary empty regions Δ' into the union of triangles cannot make the boundary of T_U between u_p and u_q any longer. Since each Δ' is composed of convex chains, the boundary of T_U between vertices u_p and u_q can only get shorter [2] by adding the convex shape of the second empty region Δ' . Therefore:

$$|T_U(A, B)| \leq |\overline{T_U}(A, B)| \leq c_{\text{OSP}} |AB|$$

The final path from A to B will be constructed and bounded using alternating, disjoint sections of the boundaries of T_U and T_L , connected by tangents. Ultimately, $|AB|$ times the constant $c_{\alpha, d}$ from Theorem 1, which bounds the length of the final path, will pay for all of $T_U(A, B)$ and $T_L(A, B)$. Hence, the portions of the final path that lie on $T_U(A, B)$ and $T_L(A, B)$ can be charged appropriately. However, the sections of the final path which cross between the upper and lower boundaries along the tangents are not yet paid for, and it is calculating the length of these tangents that comprises the majority of the remainder of the proof of Theorem 1.

3.2.4 Constructing the Final Path

Everything required to formalize the construction process of the path from A to B has now been defined. We shall use an incremental construction, walking around the boundaries of the faces of G that intersect AB . Initially, $s = i = 0$, and $u_s = A$. Assume that the construction process of the path has been one-sided above the x -axis from u_s until some u_i , with $i \geq s$. Assume the next edge e_{i+1} has its apex $v(e_{i+1})$ above the x -axis. The path will cross AB to the lower boundary. This will be accomplished by continuing along UC_i until the upper vertex of the tangent $t_{a,i}$ is reached. The path will then cross from the upper chain to the lower chain along $t_{a,i}$, and from the tangent's lower vertex, continue along LC_i until the path reaches l_{i+1} .

Looking ahead along the sequence of edges, either there exists another edge e_j (with $j > i$) whose apex is below the x -axis, or the path terminates at B . If the path ends at B , then the remainder of the path is one-sided, and is charged to the boundary of T_L . Otherwise, the path returns above AB as follows: e_{j-1} is the last edge whose apex is above AB . The path goes from l_{j-1} to u_j by following LC_{j-1} until the lower vertex of the tangent $t_{b,j-1}$, crossing back up above AB along $t_{b,j-1}$, and then continuing to u_j by way of UC_{j-1} . This construction is referred to in [7] as the *direct path*.

For ease of reference, the remainder of the proof of Theorem 1 will refer to tangent $t_{a,i}$ simply as t_a , and similarly refer to $t_{b,j-1}$ as t_b . Additionally, refer to the upper and lower vertices of the tangents as u_a , l_a , u_b , and l_b , respectively.

As stated at the end of Section 3.2.3, the parts of the path that follow any chain UC_i or LC_i can be paid for by $T_U(A, B)$ or $T_L(A, B)$. However, the lengths of the tangents between T_U and T_L are not “charged” yet. It is necessary to ensure that when the path from u_i to u_j goes below the x -axis, as described in the direct path construction of the preceding paragraph, that the lengths of the tangents t_a and t_b are not too great relative to the distance travelled along AB . In other words, the direct path should not go “too far down” between the tangents’ upper vertices u_a and u_b . If the direct path is too long, a *shortcut* (defined later) must be constructed between u_a and u_b . We now present criteria to determine how the path from u_a to u_b will be constructed, either along the direct path, or using a shortcut, and provide a bound

for the path's length.

Extend the tangents t_a and t_b until they intersect. Label their intersection point as C . Note that C may be below or above the x -axis. Label the smaller of the two angles between the two extended tangents as θ . If t_a and t_b are parallel, then C is a point at infinity, and $\theta = 0$.

Figure 3.13: Case 2: Construct an empty region

Case 1: $\theta \geq \alpha$. In this case, the angle between the tangents is relatively wide.

Continue along the direct path: cross tangent t_a , follow the lower convex chains $LC_i \cdots LC_{j-1}$ from l_a until l_b , and cross back up along t_b .

Case 2: $\theta < \alpha$. As the tangents become more parallel, a second criterion is needed to determine if the direct path should be maintained, or if the shortcut is needed. Refer to Figure 3.13. Label the portion of $T_U(A, B)$ between the upper vertices u_a and u_b of the tangents as T_1 . Similarly, label the portion of $T_L(A, B)$ between l_a and l_b as T_2 . Consider the region bounded by the tangents t_a and t_b and the

boundaries T_1 and T_2 . By definition, this region is empty of vertices of G . Bisect the angle θ at C with a line labelled ℓ . (If t_a and t_b happen to be parallel, then ℓ is defined as the line parallel to the tangents, and halfway between them.) Let E be the point on T_1 whose orthogonal projection onto ℓ is lowest, and let F be the point on T_2 whose orthogonal projection onto ℓ is highest. Define h to be the distance between the orthogonal projections of E and F onto ℓ . Define points J and M on t_a , and K and N on t_b , with JEK and MFN perpendicular to ℓ . Let $w = \max(|JK|, |MN|)$.

- a) If $h \leq \frac{w}{2 \tan(\alpha/2)}$, the shortcut is not necessary, and we continue along the direct path from u_a to u_b as described in Case 1.
- b) If $h > \frac{w}{2 \tan(\alpha/2)}$, then a shortcut above the x -axis is constructed. The shortcut path will traverse the upper boundaries of the faces between u_i and u_j .

We now bound the lengths of the paths constructed in the context of the different cases stated above.

3.2.5 Case 1

Lemma 4. *If $\theta \geq \alpha$, the length of the portion of the direct path from A to B that is between u_a and u_b is at most $c_1(|T_1| + |T_2|)$, where $c_1 = \frac{2}{\sin(\alpha/2)}$.*

The proof of this lemma is the remainder of this section. As shown in Section 3.2.3, the length of the section of the direct path between l_a and l_b is bounded by $|T_2|$. It remains to calculate the length of the tangents.

Refer to Figure 3.14. Assume *for now* that C , the intersection of the two extended tangents, is below AB . The goal is to show that the length of each tangent can be paid for by the length of T_1 . Since T_1 joins u_a and u_b , the straight-line segment $u_a u_b$ must have length at most $|T_1|$. Considering the triangle $\triangle u_a u_b C$ as in Figure 3.14, it must be shown that the bottom two sides of the triangle cannot be “too much” longer than the length of the top side.

Figure 3.14: Extend the tangents t_a and t_b until they intersect

Label the extended tangent $u_a C$ as X , and the other segment $u_b C$ as Y . Label the angle opposite Y as β , with $\beta \in [0, \pi - \theta]$. The problem of bounding the lengths of the two tangents has been reduced to determining the sum of the lengths $|X|$ and $|Y|$ in relation to $|u_a u_b|$.

Refer to Lemma 2 on Page 21, which bounded the length of two sides of a triangle relative to its base. Observe that this is the exact same problem, albeit upside-down, with the exception that the upper angle is called θ instead of α . Hence, the result of that lemma can be directly applied here: $|X| + |Y| \leq \frac{1}{\sin(\theta/2)} \cdot |u_a u_b|$. Since $\theta \in [\alpha, \pi)$, $\frac{1}{\sin(\theta/2)}$ is maximized when $\sin(\theta/2)$ is minimized — when θ is as small as possible. That occurs at the boundary condition of this case: $\theta = \alpha$. Hence, the sum of the lengths of the two extended tangents is at most $\frac{|u_a u_b|}{\sin(\alpha/2)} \leq \frac{|T_1|}{\sin(\alpha/2)}$.

Symmetrically, if C were located above the x -axis, then the sum of the lengths of the tangents must be at most $\frac{|T_2|}{\sin(\alpha/2)}$.

Problems with Extended Tangents

In the original paper [7], the authors failed to consider the case where one of the tangents lies “underneath” the other. That is, when one tangent was extended, it would intersect the other, and C would exist on a tangent rather than below them both. Referring to Figure 3.15, the section of the tangent t_b between C and l_b would

Figure 3.15: Tangents can stick out underneath each other

not be paid for by the previous argument. There is a simple solution for this case:

- using the technique described above, the entire segment $X = u_a C$ gets paid for, but the section from l_a to C is unnecessary for bounding the length of t_a , so it is essentially extra
- the straight-line distance between l_a and l_b can not be greater than $|T_2|$. This follows since $|T_2|$ bounds the length of the path connecting l_a and l_b .

By the triangle inequality, $|l_b C| \leq |l_b l_a| + |l_a C| \leq |T_2| + |l_a C|$. Since $|l_a C|$ has already been paid for, an extra $|T_2|$ term is all that is needed. Note that this would be the case regardless of which tangent extended underneath the other. By a symmetric argument, if C lies above AB , then an extra $|T_1|$ term would be needed instead.

Summing up, the length of the direct path between the upper vertices u_a and u_b , which includes the lower piece of the path between l_a and l_b , the two tangents, and the extra term to pay for “over-extended” tangent, is at most:

- $\overbrace{|T_2|}^{\text{lower path}} + \overbrace{\frac{|T_1|}{\sin(\alpha/2)}}^{\text{tangents}} + \overbrace{|T_2|}^{\text{extra}} = 2|T_2| + \frac{|T_1|}{\sin(\alpha/2)} \leq \frac{2|T_2|}{\sin(\alpha/2)} + \frac{|T_1|}{\sin(\alpha/2)}$, if C is *below* AB .
- $|T_2| + \frac{|T_2|}{\sin(\alpha/2)} + |T_1| \leq \frac{2|T_2|}{\sin(\alpha/2)} + |T_1| \leq \frac{2|T_2|}{\sin(\alpha/2)} + \frac{|T_1|}{\sin(\alpha/2)}$, if C is *above* AB .

Both of these are at most $\frac{2}{\sin(\alpha/2)} (|T_1| + |T_2|)$, completing the proof of Lemma 4.

3.2.6 Case 2a)

The original paper [7] contains both a small error, as well as a major critical flaw, from which the smaller error derives. Both cannot be fixed given their original context. However, before continuing, we shall now examine the lemma for Case 2a) in its original state.

The Original Lemma

Below is the Lemma that is incorrect, as it appeared in [7]. For reference, the preceding case definitions which define the context have also been included. For simplicity, the terminology of the original paper has been translated to coincide with our framework.

DJ Case 1: $\theta > \frac{\pi}{3}$. The shortcut is not necessary.

DJ Case 2: $\theta \leq \frac{\pi}{3}$. This intuitively means that the tangents are almost parallel.

The region bounded by the tangents t_a and t_b and the paths $T_U(A, B)$, and $T_L(A, B)$, is empty of sites. Let the upper (lower) boundary of this region be T_1 (T_2). E (F) is the point on T_1 (T_2) with the minimum (maximum) y coordinate, and $h = y(E) - y(F)$. JEK and MFN are parallel to the x -axis. Let $w = \max(|JK|, |MN|)$. If $\frac{h}{w} \leq \frac{1}{2 \tan(\alpha/2)}$, the shortcut is not necessary. Otherwise, the shortcut is taken.

You may notice that the original paper divided its two cases based on θ and $\frac{\pi}{3}$, rather than if θ is greater or less than α . We outline the errors in Das and Joseph's approach [7].

DJ Lemma 1. *If the shortcut is not taken due to DJ Case 2, then the length of the portion of the path between u_a and u_b is smaller than $a_\alpha(|T_1| + |T_2|)$, where $a_\alpha = \frac{5}{\sin(\alpha/2)}$.*

The small Error

Only the first small section of the proof where the smaller error lies will be included, as taken from Das' thesis [6].

Proof. As before, the lower boundary of [the direct path] between l_a and l_b is smaller than $|T_2|$. We now compute the total length of the two tangents. Since $\theta \leq \pi/3$, we see that irrespective of where C lies, the following conditions hold:

$$\begin{aligned} |u_a J| &\leq \frac{2}{\sqrt{3}} |u_a E|, \\ |u_b K| &\leq \frac{2}{\sqrt{3}} |u_b E|, \\ |l_a M| &\leq \frac{2}{\sqrt{3}} |l_a F|, \\ |l_b N| &\leq \frac{2}{\sqrt{3}} |l_b F|. \end{aligned}$$

... (rest of the original proof omitted).

Figure 3.16: Diagram for The small Error

However, consider Figure 3.16. Note that we are certainly within this case: $\theta = \pi/3$, and w is much larger than h , as required. But it is quite easy to see in this configuration that $|l_b N| > \frac{2}{\sqrt{3}}|l_b F| \approx 1.15|l_b F|$. In fact, in the figure, $l_b N$ is more than 6 times longer than $l_b F$, far exceeding the allowed $\frac{2}{\sqrt{3}}$. Perhaps it could be fixed, using $|T_1|$ to help bound $|l_b N|$. However, it seems that as the tangents approach being parallel with the x -axis, they appear to be able to grow unbounded. This observation is what was used to break the original proof, and what the following section is based on.

the reason for the error. It seemed like a good choice based on Dobkin's work [11], and also because there is a natural separation of the faces of the graph above and below AB . However, since the proof attempts to bound the length of the tangents using only $|T_1|$ and $|T_2|$, whose length have nothing to do with the x -axis, the choice of empty region should also have nothing to do with the x -axis.

It might be helpful to the reader to now refer back to the modified case definitions and Figure 3.13 on page 31. Bisecting the angle between the two tangents with ℓ , and aligning the construction to ℓ , changes the co-ordinate system so that the very sloped, parallel configuration that broke the original lemma (Figure 3.17) cannot happen: it forces w to be as small as possible since w is perpendicular to the bisector of the tangents. The construction in Figure 3.17 now becomes a tall, narrow column, which is very appropriate to constructing a shortcut, and will be discussed at length in Section 3.2.7. Before getting too far ahead, the corrected lemma for Case 2a) must first be proven.

Lemma 5 (Case 2a). *If $\theta < \alpha$, and $h \leq \frac{w}{2 \tan(\alpha/2)}$, then the length of the portion of the direct path from A to B that is between u_a and u_b is at most $c_{2a}(|T_1| + |T_2|)$, where $c_{2a} = \frac{3}{\sin(\alpha/2)}$.*

Proof. Orient the construction such that line ℓ is vertical, as in Figure 3.18. Assume that C is located below AB , and hence $w = |JK| \geq |MN|$. Also, assume that $h > 0$ — that is, the segment JK is “above” segment MN when the construction is reoriented. It is possible for h to be negative; however, this particular case will be dealt with afterwards, as well as when C is above AB .

As in Lemma 4, the section of the direct path between l_a and l_b in the graph is at most $|T_2|$. It remains to bound the length of the tangents t_a and t_b .

Consider the parts of t_a and t_b above the segment JK : u_aJ and u_bK respectively. By the triangle inequality, the segment u_aE must have length less than or equal to the portion of T_1 that runs between u_a and E . Symmetrically, $|u_bE|$ is less than or equal to the part of T_1 which runs from E to u_b . Therefore, $|u_aE| + |u_bE| \leq |T_1|$. Note that the angle $\angle u_aJE$ is always obtuse ($\geq \pi/2$) when C is below AB . This can be seen since $\angle CJE = \frac{\pi - \theta}{2}$, which means the other side of the straight angle

Figure 3.18: Bounding the length of the tangents for Lemma 5

$\angle u_a J E = \pi - \frac{\pi - \theta}{2} = \frac{\pi + \theta}{2} \geq \frac{\pi}{2}$. Hence, $|u_a J| \leq |u_a E|$. Similarly, since $\angle u_b K E = \angle u_a J E$, $|u_b K| \leq |u_b E|$. Combining the above inequalities:

$$|u_a J| + |u_b K| \leq |u_a E| + |u_b E| \leq |T_1|$$

Next, consider the sections of the tangents JM and KN . Since h is measured parallel to ℓ , which bisects the angle between t_a and t_b , the lengths $|JM|$ and $|KN|$ can be expressed as a function of α and h . Referring to Figure 3.18, consider the segment KN and its adjacent right triangle. By elementary trigonometry:

$$\begin{aligned}
& \cos\left(\frac{\theta}{2}\right) = \frac{h}{|KN|} \\
\stackrel{\text{rearrange}}{\Rightarrow} \quad & |KN| = \frac{h}{\cos\left(\frac{\theta}{2}\right)} \\
& \leq \frac{h}{\cos\left(\frac{\alpha}{2}\right)}
\end{aligned}$$

Similarly, $|JM| \leq \frac{h}{\cos(\alpha/2)}$. Recall that due to Case 2a), $h \leq \frac{w}{2 \tan(\alpha/2)}$. Therefore:

$$\begin{aligned}
|JM| + |KN| & \leq \frac{2h}{\cos\left(\frac{\alpha}{2}\right)} \\
& \leq \frac{2w}{\cos\left(\frac{\alpha}{2}\right) 2 \tan\left(\frac{\alpha}{2}\right)} \\
& = \frac{w}{\cos\left(\frac{\alpha}{2}\right) \tan\left(\frac{\alpha}{2}\right)} \\
& = \frac{w}{\sin\left(\frac{\alpha}{2}\right)} \\
& = \frac{|JK|}{\sin\left(\frac{\alpha}{2}\right)} \\
& \leq \frac{|T_1|}{\sin\left(\frac{\alpha}{2}\right)}
\end{aligned}$$

Finally, consider the sections of the tangents below MN : $l_a M$ and $l_b N$. Like the first part of this proof, we have $|l_a F| + |l_b F| \leq |T_2|$. Unlike the top sections of the tangents, the angle $\angle l_a M F$ is acute ($\leq \pi/2$), which means that $|l_a M|$ is not necessarily less than $|l_a F|$.

To maximize the length of $|l_a M|$ with regards to $|l_a F|$ is to locate F directly above l_a (or conversely, have l_a situated vertically below F). This is because the shortest distance between a point and a line is to travel perpendicular to that line — if $l_a F$ was perpendicular to MN (and parallel to ℓ), its length would be minimized relative to $l_a M$. Of course, this same logic holds true for $|l_b N|$ and $|l_b F|$.

With F located directly above l_a , and the angle $\angle F l_a M = \frac{\theta}{2}$, then by basic trigonometry:

$$\begin{aligned}
\cos\left(\frac{\theta}{2}\right) &= \frac{|l_a F|}{|l_a M|} \\
\stackrel{\text{rearrange}}{\Rightarrow} |l_a M| &= \frac{|l_a F|}{\cos\left(\frac{\theta}{2}\right)} \\
&\leq \frac{|l_a F|}{\cos\left(\frac{\alpha}{2}\right)}
\end{aligned}$$

Similarly, $|l_b N| \leq \frac{|l_b F|}{\cos(\alpha/2)}$. Therefore:

$$\begin{aligned}
|l_a M| + |l_b N| &\leq \frac{|l_a F| + |l_b F|}{\cos\left(\frac{\alpha}{2}\right)} \\
&\leq \frac{|T_2|}{\cos\left(\frac{\alpha}{2}\right)}
\end{aligned}$$

The lengths of each section of the tangents have now been successfully bounded to a constant (based on α) times either $|T_1|$ or $|T_2|$. Combining all of the above steps, and summing the lengths of the tangents with the length of the path from l_a to l_b :

$$\begin{aligned}
|T_2| + |t_a| + |t_b| &= |T_2| + (|u_a J| + |JM| + |l_a M|) + (|u_b K| + |KN| + |l_b N|) \\
&= |T_2| + (|u_a J| + |u_b K|) + (|JM| + |KN|) + (|l_a M| + |l_b N|) \\
&\leq |T_2| + |T_1| + \frac{|T_1|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{|T_2|}{\cos\left(\frac{\alpha}{2}\right)} \quad \text{Note: } \frac{1}{\cos\beta} < \frac{1}{\sin\beta}, \beta \in [0, \frac{\pi}{4}) \\
&\leq |T_2| + |T_1| + \frac{|T_1|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{|T_2|}{\sin\left(\frac{\alpha}{2}\right)} \\
&\leq \frac{|T_2|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{|T_1|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{|T_1|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{|T_2|}{\sin\left(\frac{\alpha}{2}\right)} \\
&= \frac{2|T_1|}{\sin\left(\frac{\alpha}{2}\right)} + \frac{2|T_2|}{\sin\left(\frac{\alpha}{2}\right)} \\
&= \frac{2}{\sin\left(\frac{\alpha}{2}\right)} (|T_1| + |T_2|)
\end{aligned}$$

We are not done yet. There are several special cases to consider. The first is similar to the special case in Lemma 4 (Case 1) where C actually lies on one of the

tangents because one of the tangents is extending “underneath” the other. Clearly, the ratio between $|l_a M|$ and $|l_a F|$ is still maximized as long as $l_a F$ is parallel to ℓ . Similarly for $|l_b N|$ and $|l_b F|$. If one tangent was underneath the other, then it would force either $l_a F$ or $l_b F$ away from the parallel with ℓ , and numerically away from the upper bound. Hence, the worst case is still realized with F directly over top of both l_a and l_b . This ensures that the bound still holds regardless of how the tangents are situated relative to each other.

Figure 3.19: It is possible for h to be negative

Next, consider what happens if $h \leq 0$, or in other words, if JK was “below” MN relative to ℓ . Refer to Figure 3.19. Rather than be a special case that requires a lot of attention, when h is less than 0, it means the sections of T_1 and T_2 that we are using to bound the tangents are quite long. In fact, all of the above work still applies, with the only difference that the sections of the tangents JM and KN do not need to be included. The other sections of the tangents either overlap (and essentially count the section MJ and NK twice!), or they are just touching (when $h = 0$).

Finally, what if C resides above the x -axis? With the exception of the bound on the length of the path between l_a and l_b (which is still $|T_2|$), all of the other terms simply switch T_1 ’s for T_2 ’s, and vice versa. Hence, when C is above AB , the length of the path from u_a to u_b is bounded by:

$$\begin{aligned}
|T_2| + |t_a| + |t_b| &= |T_2| + (|u_a J| + |JM| + |l_a M|) + (|u_b K| + |KN| + |l_b N|) \\
&= |T_2| + (|u_a J| + |u_b K|) + (|JM| + |KN|) + (|l_a M| + |l_b N|) \\
&\leq |T_2| + \frac{|T_1|}{\sin(\frac{\alpha}{2})} + \frac{|T_2|}{\cos(\frac{\alpha}{2})} + |T_2| \\
&\leq |T_2| + \frac{|T_1|}{\sin(\frac{\alpha}{2})} + \frac{|T_2|}{\sin(\frac{\alpha}{2})} + |T_2| \\
&\leq \frac{|T_2|}{\sin(\frac{\alpha}{2})} + \frac{|T_1|}{\sin(\frac{\alpha}{2})} + \frac{|T_2|}{\sin(\frac{\alpha}{2})} + \frac{|T_2|}{\sin(\frac{\alpha}{2})} \\
&= \frac{|T_1|}{\sin(\frac{\alpha}{2})} + \frac{3|T_2|}{\sin(\frac{\alpha}{2})} \\
&\leq \frac{3}{\sin(\frac{\alpha}{2})} (|T_1| + |T_2|)
\end{aligned}$$

Hence, for Case 2a), we have shown that a short path between u_a and u_b can be constructed, regardless of where C is located. If one of the tangents t_a and t_b stuck out underneath the other, it does not affect the result of this proof. And finally, if $h < 0$, the tangents are still paid for by the same arguments. This summarizes all of the possible configurations for this case, thus completing the proof of Lemma 5. $\square$

3.2.7 Case 2b)

Finally, if constructing a shortcut between u_a and u_b is necessary, we can show that the total length of said shortcut is not too long. The shortcut will not follow the direct path below the x -axis. Instead, the shortcut path will traverse the upper boundaries of the faces along AB , from u_a to u_b . We formalize this construction, and bound its length in the following lemma:

Lemma 6. *On the direct path from A to B , if the shortcut is constructed from u_a to u_b because $h > \frac{w}{2 \tan(\alpha/2)}$, then the length of the shortcut is at most $c_{2b}(|T_1| + |T_2|)$, where $c_{2b} = \frac{2(\pi-\alpha)}{\alpha \sin(\alpha/4)}$.*

Proof. The goal here is to construct a shortcut path between u_a and u_b that stays

completely above AB . Consider each of the convex chains between u_a and u_b , including the chains they are on: $UC_i, UC_{i+1}, \dots, UC_{j-1}$. Let UC_Z be the lower convex hull of $UC_i, \dots, UC_{j-1}$. Label the vertices of UC_Z as $u_a = z_0, z_1, \dots, z_{n-1}, z_n = u_b$. This “super-hull” UC_Z is below all of the upper boundaries of the faces between e_i and e_j , and hence is completely contained inside these faces. This means that the only edges of G that can possibly cross any given segment $z_i z_{i+1}$ of UC_Z are $e_{i+1}, e_{i+2}, \dots, e_{j-1}$.

Note that $|UC_Z| \leq |T_1|$. To see this, consider the edges that intersect AB between u_a and u_b : $e_{i+1}, e_{i+2}, \dots, e_{j-1}$. Each of these has their apex above AB , and therefore do not contribute anything to the boundary $T_U(A, B)$, i.e. T_1 , between u_a and u_b . Starting from u_a , T_1 follows the tangent t_a until either intersecting t_b and returning to u_b , or intersects AB , in which case it follows AB until intersecting t_a . In either case, T_1 is completely below the vertices of the upper boundaries of the faces between u_a and u_b . Since UC_Z is the convex hull of these vertices between u_a and u_b , and completely above T_1 , the length of the UC_Z is at most $|T_1|$ [2].

The shortcut path will be constructed by traversing the direct path between vertices of UC_Z . We will show that for each segment $z_i z_{i+1}$ of UC_Z , the direct path between z_i and z_{i+1} in G must be a one-sided path, with respect to z_i and z_{i+1} . Due to Lemma 3, this immediately provides a bound on the length of the shortcut path from u_a to u_b . Further, we show that the direct path of each segment of UC_Z is an *upper* one-sided path: that is, every edge of G crossing a segment of the super-hull must have its apex below said segment. However, the edges that cross UC_Z — $e_{i+1}, e_{i+2}, \dots, e_{j-1}$ — all have their apices at least above AB . But how far above? Why can none of $v(e_{i+1}), v(e_{i+2}), \dots, v(e_{j-1})$ be above UC_Z ?

Assume some segment $z_i z_{i+1}$ has some edge e intersecting it such that e ’s apex $v(e)$ lies above $z_i z_{i+1}$, and hence violates the segment’s upper one-sided path property. Refer to Figure 3.20. Assume that e ’s empty triangle is rightward facing. For $v(e)$ to lie above $z_i z_{i+1}$, then $\angle ulz_{i+1} > \alpha$. Otherwise, if $\angle ulz_{i+1} \leq \alpha$, z_{i+1} would either be inside the empty triangle $\Delta(e)$ which violates the diamond property, or above $\Delta(e)$ (inside the shaded region in Figure 3.20) such that $v(e)$ lies below the segment $z_i z_{i+1}$, which contradicts the assumption that e violates the upper one-sided path for $z_i z_{i+1}$.

Refer to Figure 3.21. Assume the construction has been oriented, as in Case 2a),

Figure 3.20: Violation of a one-sided path requires angle $\angle ulz_{i+1} > \alpha$

such that JK and MN are horizontal. Note that every segment $z_i z_{i+1}$ of UC_Z must exist between the two tangents t_a and t_b , and above or on JK . This follows since UC_Z is the convex hull of the upper convex chains, and starts and ends at the tangents' upper vertices u_a and u_b , therefore ensuring all of its segments are between t_a and t_b . No segment lies below JK simply because J and K correspond to a point E (Figure 3.13) which is selected as the lowest point along T_1 . If a vertex of UC_Z was below JK , then E would not be the lowest possible point on T_1 .

Hence, as the vertex z_{i+1} from Figure 3.20 is a vertex of UC_Z , it must lie above JK , and to the left of t_b . However, z_{i+1} it must also lie to the right of the line through $lv(e)$, so that it is outside of the empty triangle $\Delta(e)$, thus ensuring the violation of the one-sided path. We shall demonstrate that this particular configuration is impossible, thereby disproving the premise of a non-one-sided path by contradiction.

Figure 3.21: The empty trapezoid is too narrow for v to be above JK

From Figure 3.21, observe that:

$$\begin{aligned}
 \tan(\sigma) + \tan(\rho) &= \frac{|u'd|}{|ld|} + \frac{|v'd|}{|ld|} \\
 &= \frac{|u'd| + |v'd|}{|ld|} \\
 &\leq \frac{|u'd| + |v'd|}{h} \\
 &\leq \frac{w}{h} \\
 &< 2 \tan\left(\frac{\alpha}{2}\right)
 \end{aligned} \tag{3.4}$$

Since $\angle ulv(e) = \alpha$, then $\sigma + \rho = \alpha \Rightarrow \rho = \alpha - \sigma$. Hence:

$$\tan(\sigma) + \tan(\rho) = \tan(\sigma) + \tan(\alpha - \sigma) = f(\alpha, \sigma)$$

Due to equation (3.4), $f(\alpha, \sigma) < 2 \tan\left(\frac{\alpha}{2}\right)$ for some σ . To verify this is true, differentiate with respect to σ , and solve for α where the minimum occurs:

$$\begin{aligned}
\frac{\partial f}{\partial \sigma} &= (1 + \tan^2(\sigma)) + (-1 - \tan^2(\alpha - \sigma)) = 0 \\
\tan^2(\sigma) - \tan^2(\alpha - \sigma) &= 0 \\
\tan^2(\sigma) &= \tan^2(\alpha - \sigma)
\end{aligned}$$

Since $\tan(x)$ is a bijection, and entirely positive in the range $x \in [0, \pi/2)$, it follows that $\tan^2(x)$ is as well. Hence, because $\alpha \in (0, \pi/2)$, and $\sigma \in [0, \alpha]$:

$$\begin{aligned}
\tan^2(\sigma) &= \tan^2(\alpha - \sigma) \\
\text{implies that } \Rightarrow \quad \sigma &= \alpha - \sigma \\
\sigma &= \frac{\alpha}{2}
\end{aligned}$$

Therefore, the minimum value that $f(\alpha, \sigma)$ can possibly achieve for any given α is when σ is equal to $\frac{\alpha}{2}$:

$$f_{\min}(\alpha, \sigma = \frac{\alpha}{2}) = \tan\left(\frac{\alpha}{2}\right) + \tan\left(\alpha - \left(\frac{\alpha}{2}\right)\right) = \tan\left(\frac{\alpha}{2}\right) + \tan\left(\frac{\alpha}{2}\right) = 2 \tan\left(\frac{\alpha}{2}\right)$$

This however is in contradiction to equation (3.4), which states that $f(\alpha, \sigma)$ must be *less* than $2 \tan(\frac{\alpha}{2})$. Hence, the assumption that such an edge e exists is invalid, and it is impossible for the direct path of any segment of UC_Z to not be upper one-sided.

Thus, for every segment $z_i z_{i+1}$ of the super-hull UC_Z , the direct path between its endpoints z_i and z_{i+1} is a one-sided path in G . As such, the result of Lemma 3 implies that the length of this path from z_i to z_{i+1} is at most $\frac{2(\pi-\alpha)}{\alpha \cdot \sin(\frac{\alpha}{4})} |z_i z_{i+1}|$. Therefore, the entire length of the shortcut is:

$$\begin{aligned}
& \frac{2(\pi-\alpha)}{\alpha \cdot \sin(\frac{\alpha}{4})} \cdot (|z_0 z_1| + |z_1 z_2| + \dots + |z_{n-1} z_n|) \\
&= \frac{2(\pi-\alpha)}{\alpha \cdot \sin(\frac{\alpha}{4})} \cdot |UC_Z| \\
&\leq \frac{2(\pi-\alpha)}{\alpha \cdot \sin(\frac{\alpha}{4})} \cdot |T_1| \\
&\leq \frac{2(\pi-\alpha)}{\alpha \cdot \sin(\frac{\alpha}{4})} \cdot (|T_1| + |T_2|)
\end{aligned}$$

This completes the proof of Lemma 6. □

3.3 The Final Spanning Ratio

Regardless of which case constructs the path between the upper vertices of the two tangents t_a and t_b , the maximum length of the path between u_a and u_b is at most:

$$\begin{aligned}
 & \max(c_1, c_{2a}, c_{2b}) \cdot (|T_1| + |T_2|) \\
 = & \max\left(\frac{2}{\sin\left(\frac{\alpha}{2}\right)}, \frac{3}{\sin\left(\frac{\alpha}{2}\right)}, \frac{2(\pi - \alpha)}{\alpha \sin\left(\frac{\alpha}{4}\right)}\right) \cdot (|T_1| + |T_2|) \\
 = & \frac{2(\pi - \alpha)}{\alpha \sin\left(\frac{\alpha}{4}\right)} \cdot (|T_1| + |T_2|)
 \end{aligned}$$

Recall the incremental construction from Section 3.2.4 on Page 30. Let $u_s = u_j$, and $u_i = u_j$. Continue along the direct path from u_s , modifying for shortcuts if necessary, until reaching B . At each stage, bound the construction of the path with whichever case properly applies, charging T_1 and T_2 appropriately to pieces of $T_U(A, B)$ and $T_L(A, B)$.

Recall from Section 3.2.3 that $|T_U(A, B)|$ and $|T_L(A, B)|$ are bounded by the one-sided path constant c_{OSP} . Further, since $|T_U(A, B)|$ and $|T_L(A, B)|$ pay for the portions of the path along the upper and lower convex chains, respectively, the constant d must be included because the convex chains are merely d -approximations to actual paths around the boundaries of the faces in G . Therefore, the constructed path between A and B has length at most:

$$\begin{aligned}
 & \frac{2(\pi - \alpha)}{\alpha \sin\left(\frac{\alpha}{4}\right)} \cdot (|T_U(A, B)| + |T_L(A, B)|) \\
 \leq & \frac{2(\pi - \alpha)}{\alpha \sin\left(\frac{\alpha}{4}\right)} \cdot 2 \left(\frac{2(\pi - \alpha)d|AB|}{\alpha \sin\left(\frac{\alpha}{4}\right)} \right) \\
 = & \frac{8(\pi - \alpha)^2 d |AB|}{\alpha^2 \sin^2\left(\frac{\alpha}{4}\right)}
 \end{aligned}$$

proving Theorem 1. ■

Chapter 4

Generalization to the Visible α -diamond Property

Recall the definitions relating to constrained graphs from Chapter 2. This chapter demonstrates how the concept of α -diamond spanners translates to the constrained setting. The following definition shows the modification to the diamond property:

Figure 4.1: Every edge e has at least one *visibly* empty α -triangle

Definition 16. Refer to Figure 4.1. Fix $\alpha \in (0, \pi/2)$. A constrained plane graph $G(S)$ is said to have the *visible α -diamond property* if, for every unconstrained edge e in the graph, a pair of isosceles triangles, with e as the base and base angle α , has vertices of $G(S)$ visible to e in at most one of the two triangles. That is, at least one of the triangles (not including the boundary of the region) must not have any vertices of $G(S)$ visible to e . Label the visibly empty triangle as $\Delta(e)$, and the apex of $\Delta(e)$ as $v(e)$.

We show that a constrained graph, built on a set S of points and non-crossing constrained edges, $G(S)$ with the visible α -diamond property, and the d -good polygon property, is a spanner of the visibility graph $vis(S)$.

Theorem 2. Fix $\alpha \in (0, \pi/2)$ and $d \geq 1$. If a constrained plane graph $G(S)$ has both the visible α -diamond property and the d -good polygon property, then for every pair of vertices A and B in $G(S)$, there exists a path in $G(S)$ between A and B whose length is at most $c_{\alpha,d} \cdot d_{vis(S)}(A, B)$, where $c_{\alpha,d} = \frac{8(\pi-\alpha)^2 d}{\alpha^2 \sin^2(\alpha/4)}$.

Proof. We first assume that A and B are visible in $G(S)$. Hence, all of the edges $e_1, e_2, \dots, e_k$ that cross AB must be unconstrained edges, and as such are all incident to visibly empty triangles. As in Figure 4.1, these triangles $\Delta(e_1), \Delta(e_2), \dots, \Delta(e_k)$ may or may not have any constrained edges passing through them. Additionally, the faces along AB adjacent to the k edges may have any number of constrained edges on their boundaries. $G(S)$ is still a planar graph however, and the boundaries of each face between adjacent edges e_i and e_{i+1} can again be divided into an upper boundary and a lower boundary. Further, upper and lower convex hulls of these boundaries, $UC_0, \dots, UC_k$ and $LC_0, \dots, LC_k$ can still be constructed in the exact same manner as the non-constrained proof of Chapter 3.

The tangents $t_{a,1}, \dots, t_{a,k-1}$ and $t_{b,1}, \dots, t_{b,k-1}$ can also be constructed between their respective upper and lower convex hulls, simply because (by definition) the area between the hulls inside the faces are completely empty of all vertices and edges of $G(S)$, regardless of constraints or visibility.

4.1 One-Sided Paths

Assume the direct path from A to B is one-sided, and without loss of generality, assume all of the apices lie in the half-plane h^- below AB . As before, let $T = h^- \cup \Delta(e_1) \cup \Delta(e_2) \cup \dots \cup \Delta(e_k)$, and let $T(A, B)$ denote the boundary of this region between A and B .

From the good polygon property, the distance along the convex chains $|UC_0| + |UC_1| + \dots + |UC_k|$ from A to B is still a d -approximation to an actual path in $G(S)$. However, due to the constrained nature of $G(S)$, it is unclear whether the length of the boundary of T between adjacent upper vertices u_i and u_{i+1} , i.e. $T(u_i, u_{i+1})$, is still an appropriate bound on the length of their respective convex chain UC_i . However, since UC_i is convex, as long as it remains completely above the boundary of T , $|UC_i|$ will be at most $|T(u_i, u_{i+1})|$ [2].

Figure 4.2: A constrained edge must pass through an empty triangle $\Delta(e_i)$

The reason no convex chain UC_i can enter any empty triangle $\Delta(e)$, regardless of whether its visible or not, follows from the same reasoning as in Lemma 1 in Chapter 3: since UC_i cannot cross AB , it would have to enter and leave the triangle along the same side $uv(e)$, as in Figure 4.2. Hence, some vertex of UC_i would lie inside the empty triangle. As $\Delta(e)$ is *visibly* empty, there must exist some constrained edge xx'

blocking visibility between e and the vertex inside $\Delta(e)$.

As A and B are visible to each other, xx' cannot cross AB . However, the vertex x' is visible to e , thereby violating the visibly-empty diamond property. One could then argue that another constrained edge yy' could block visibility between e and x' , however this new constrained edge also cannot cross AB , and must therefore have one of its endpoints y' inside $\Delta(e)$. This endpoint would have to be visible to e since xx' is the only other edge that could block visibility between e and y' . Since constrained edges of S do not cross, and we know yy' is between e and x' , the endpoint y' is visible to e . This cyclical argument shows that a total ordering of constrained edges can be determined, with the constrained edge closest to e having an endpoint inside $\Delta(e)$ being visible to e .

Hence, for any edge e , whichever side of the x -axis $v(e)$ lies on, the portion of triangle $\Delta(e)$ on the other side of AB is *completely* empty, not just visibly empty. Therefore, when the direct path from A to B is one-sided above the x -axis, the sum of the lengths of all the convex chains $UC_0, UC_1, \dots, UC_k$ from A to B is still bounded by $|T(A, B)|$.

4.2 Two-Sided Paths

When the direct path from A to B is not one-sided, the three cases which determine the construction of the final path must be analyzed to ensure the methods for bounding the length of the final path translate into the constrained setting.

In chapter 3, sections of the boundaries $T_U(A, B)$ and $T_L(A, B)$ were used to bound the length of the final path. Consider region T_U . Recall that it is the union of all empty triangles whose apices lie below the x -axis, as well as the secondary empty regions Δ' , with the half plane h^- below AB . In the previous section, it was shown why no constrained edge could lie inside the triangles of T_U above AB . Additionally, no constrained edge may lie inside any region Δ' , as the secondary empty region is contained entirely within the empty space between convex chains. Hence, both of $T_U(A, B)$ and $T_L(A, B)$ can be constructed as per the usual non-constrained setting.

Therefore, the constrained edges do not impact the convex chains along the direct

path, which tries to be one-sided. The convex chains in turn are bounded by their respective boundaries of $T_U(A, B)$ and $T_L(A, B)$. Therefore, all of the pieces of the direct path in Case 1 (Section 3.2.5) and Case 2a (Section 3.2.6) that do not lie on the tangents t_a and t_b can be charged as appropriate to $T_U(A, B)$ and $T_L(A, B)$.

As the tangents between upper and lower chains in all faces of $G(S)$ along AB can be constructed despite the constrained setting, the lengths of the tangents t_a and t_b in Cases 1 and 2a can also be charged to sections of $T_U(A, B)$ and $T_L(A, B)$ as well.

Finally, in Case 2b, if the direct path is too long, and a shortcut around the upper boundaries of the faces must be constructed, recall that the super-hull UC_Z is entirely contained beneath the convex hulls $UC_i, UC_{i+1}, \dots, UC_{j-1}$ that it spans. As such, it can only intersect the unconstrained edges $e_{i+1}, \dots, e_{j-1}$. Hence, UC_Z can be constructed as in the non-constrained case, and for every segment $z_i z_{i+1}$ of UC_Z , z_i and z_{i+1} are visible in $G(S)$. Due to the shortcut condition, which specifies that $h > \frac{w}{2 \tan(\alpha/2)}$, the direct path between the endpoints of each segment of UC_Z is one-sided. Therefore, the result of visible one-sided paths of Section 4.1 applies.

The conclusion is that when A and B are visible in $G(S)$, the final path from A to B can be constructed in the same incremental process as in Chapter 3, following the direct path and modifying it with shortcuts as necessary, notwithstanding the constrained edges of $G(S)$.

Next, we must consider when A and B are not visible in $G(S)$. Let $d_{vis(S)}(A, B)$ denote the shortest path between A and B in the visibility graph of $G(S)$. Label the points along this shortest path as: $A = p_0, p_1, \dots, p_{n-1}, p_n = B$, for some $n > 1$. For each segment of the shortest path $p_i p_{i+1}$, the endpoints p_i and p_{i+1} are visible, ergo Theorem 2 states the length of the path between them is no more than $c_{\alpha,d} |p_i p_{i+1}|$. It follows that a path of length $c_{\alpha,d} |p_0 p_1| + c_{\alpha,d} |p_1 p_2| + \dots + c_{\alpha,d} |p_{n-1} p_n|$ can be constructed between A and B . Therefore, if the shortest distance between A and B in the visibility graph is $d_{vis(S)}(A, B)$, our result provides a bound on the length of the path that can be constructed from A to B , namely: $\frac{8(\pi-\alpha)^2 d}{\alpha^2 \sin^2(\alpha/4)} \cdot d_{vis(S)}(A, B)$. $\square$

Chapter 5

Triangulations that have the Diamond Property

In this chapter, we show two constrained versions of classical triangulations that both have the visible α -diamond property. Note that as they are both triangulations, they satisfy the requirements for the d -good polygon property, for $d = 1$.

5.1 Constrained Delaunay Triangulation

The first constrained triangulation considered is the constrained Delaunay triangulation (CDT), also called the Generalized Delaunay triangulation [17], and the Obstacle triangulation [4]. Let P be a set of points, and L a set of non-crossing line segments with endpoints in P . Let $S = P \cup L$. Recall that the points of P are assumed to be in general position. Given the set S , with $|S| = n$, we denote the constrained Delaunay triangulation of the set S by $CDT(S)$. Both Chew [5], and Wang and Schubert [22] independently show an optimal efficiency algorithm of $O(n \log n)$ for constructing $CDT(S)$.

One of the important properties of the constrained Delaunay triangulation is that for every unconstrained edge e in the graph, there exists a circle $\mathcal{C}_e$ such that the endpoints of e lie on the boundary of this circle, and there are no vertices of S that are visible to the edge e [17] [4]. For such an edge e , we refer to $\mathcal{C}_e$ as its *visibly empty*

Figure 5.1: Every edge e of $CDT(S)$ has a visibly empty circle

circle. With the use of this empty circle, the proof of the subsequent theorem follows:

Theorem 3. *The Constrained Delaunay Triangulation (CDT) has the visible $\frac{\pi}{4}$ -diamond property.*

Proof. Refer to Figure 5.1. As in the definitions for a constrained Delaunay triangulation in both [17] and [4], it is known that every unconstrained edge in the CDT is incident to a visibly empty circle $\mathcal{C}$. It follows from simple geometry that two isosceles triangles, with base angles $\frac{\pi}{4}$ can be inscribed on the edge, and one of these triangles will lie completely inside $\mathcal{C}$. Since the circle is visibly empty, the triangle which is entirely contained inside the circle must also be visibly empty. If the edge happens to be the diameter of the empty circle, then both triangles are empty, and either can be picked arbitrarily. $\square$

5.2 Constrained Greedy Triangulation

The second constrained triangulation that we consider is the constrained Greedy triangulation (CGT), which is a generalization of the standard Greedy triangulation

[18]. As before, let $S = P \cup L$: a set P of points in the plane, and a set L of non-crossing line segments joining pairs of points in P . Given the set S , we denote the constrained Greedy triangulation of the set S by $CGT(S)$.

Recall the construction process for the standard Greedy triangulation, outlined in Section 2.3 on Page 9. We generalize the construction by modifying the first bullet in the algorithm description:

- Consider every pair of points in S whose corresponding line segment is in the visibility graph $vis(S)$, and sort them into a non-decreasing order by the lengths of the segments connecting each pair.

Thus, only edges of the visibility graph of S will be considered for insertion. It is clear that this is a generalization of the standard Greedy triangulation, since if there were no constrained edges in S , then the visibility graph of S would equal the complete graph.

In order to prove the modified diamond result for the Constrained Greedy Triangulation (CGT), the following lemma must first be shown:

Lemma 7. *Let X and Y be points of S such that the segment XY is in the visibility graph $vis(S)$, but XY is not an edge of $CGT(S)$. Let e be the edge of $CGT(S)$ of shortest length that properly intersects the segment XY . Then $|e| \leq |XY|$.*

Proof. Recall that the CGT is constructed by considering all possible edges of $vis(S)$ in non-descending sorted order; an edge is inserted only if it does not intersect any previously inserted edge. If the lemma were false, then at the point in the algorithm when XY would be considered for insertion, none of the edges that intersect it would have been considered yet since they are all longer than XY . Hence, with no edges yet crossing XY , the segment joining X and Y would be inserted, which is in contradiction to the assumption that XY is not in the triangulation. $\square$

The use of this small but powerful lemma will facilitate the proof of the following theorem:

In the remainder of this chapter, we will prove Theorem 4.

Figure 5.2: Every edge of the CGT has a $\frac{\pi}{6}$ empty triangle

Consider Figure 5.2. Let AB be an unconstrained edge of the CGT . Let C and D be the apices of two isosceles triangles, with AB as the base and base angle $\frac{\pi}{6}$. Construct two circles of radius $|AB|$, one centered at A , the other at B . Extend the segment BC (BD) upwards until it intersects the circle centered at B , and label the intersection E (F). Similarly, extend the segment AC (AD) downwards until it intersects the circle centered at A , and label the intersection K (L). Construct a

perpendicular from segment BE (BF) that goes through A , and label the intersection of BE (BF) with the perpendicular as G (H). Repeat for segments AK and AL , labelling the intersection points I , and J as shown in the figure.

If AB is an unconstrained edge of $CGT(S)$, then in order for the visible $\frac{\pi}{6}$ -diamond property to hold, only one of the triangles $\triangle ABC$ and $\triangle ABD$ can contain vertices of S visible to the edge AB . We will show by contradiction that it is impossible for there to be vertices of S visible to AB in both triangles.

5.2.1 Trigonometry is Fun

First, it is necessary to calculate the lengths of some segments. Assume for the remainder of this proof that $|AB| = 1$. Therefore, segment HI is also of length 1. Note that by basic trigonometry:

$ AI = \cos\left(\frac{\pi}{6}\right)$ $= \frac{\sqrt{3}}{2} \approx 0.866$	$ Ig_1 = \cos\left(\frac{\pi}{6}\right) \cdot BI $ $= \frac{\sqrt{3}}{4} \approx 0.433$	$ CI = \tan\left(\frac{\pi}{6}\right) \cdot BI $ $= \frac{\sqrt{3}}{6}$
$ IK = 1 - AI $ $= 1 - \frac{\sqrt{3}}{2}$	$\therefore IJ = \frac{\sqrt{3}}{2}$	$ AC = AI - CI $ $= \frac{\sqrt{3}}{3}$
$ BI = \sin\left(\frac{\pi}{6}\right)$ $= \frac{1}{2}$	$ CO = \tan\left(\frac{\pi}{6}\right) \cdot AO $ $= \frac{\sqrt{3}}{6} \approx 0.289$ $\therefore CD = \frac{\sqrt{3}}{3} \approx 0.577$	$ Cg_3 = AC /2 \cos\left(\frac{\pi}{6}\right)$ $= \frac{1}{3}$ $\therefore g_2g_3 = \frac{2}{3}$

Next, by Law of Cosines:

$$\begin{aligned}
 |ID|^2 &= |CI|^2 + |CD|^2 - 2|CI||CD|\cos(\angle ICD) \\
 &= \left(\frac{\sqrt{3}}{6}\right)^2 + \left(\frac{\sqrt{3}}{3}\right)^2 - 2 \cdot \frac{\sqrt{3}}{6} \cdot \frac{\sqrt{3}}{3} \cdot \cos\left(\frac{2\pi}{3}\right) \\
 &= \frac{1}{12} + \frac{1}{3} + \frac{1}{6} \\
 |ID| &= \sqrt{\frac{7}{12}} \approx 0.764
 \end{aligned}$$

$$\begin{aligned}
 |ED|^2 &= |CE|^2 + |CD|^2 - 2|CE||CD|\cos(\angle ECD) \\
 &= \left(1 - \frac{\sqrt{3}}{3}\right)^2 + \left(\frac{\sqrt{3}}{3}\right)^2 - 2 \cdot \left(1 - \frac{\sqrt{3}}{3}\right) \cdot \frac{\sqrt{3}}{3} \cdot \cos\left(\frac{2\pi}{3}\right) \\
 |ED| &= \frac{\sqrt{12-3\sqrt{3}}}{3} \approx 0.869
 \end{aligned}$$

5.2.2 The Required Construction

Refer to Figure 5.3. Consider all points of S inside $\triangle ABC$ that are visible to AB , and let M be the point of these that is closest to AB . We can show that M is visible

to the entire segment AB . By contradiction, if this were not so, there must at least one constrained edge that partially blocks visibility between M and AB . Consider the triangle $\triangle AMB$. Since M is at least partially visible to AB , no constrained edge can pass completely through $\triangle AMB$, as that would totally block visibility between AB and M . Hence, every constrained edge between AB and M must have at least one endpoint inside $\triangle AMB$. However, each of these endpoints would lie closer to AB than M is, and at least one of them must be visible to AB , contradicting the fact that M is the closest visible vertex to AB .

Let N be the point on AB which is the perpendicular projection of M onto AB . Walk along the segment NM from N towards M , and label each triangle of the CGT that is passed through in order: t_1, t_2 , and so on. Let t_i be the first triangle that has one of its vertices strictly inside the region $ABICG$. We know that at least one triangle has a vertex inside this region since vertex M , which is definitely inside $ABICG$, is incident to the last triangle. Let V be the vertex of triangle t_i that is inside the region. Without loss of generality, it can be assumed that V lies in the lower half of $ABICG$, i.e., in the region $OBIC$. Note that V is visible to the entire segment AB . This follows from the fact that M is also visible to all of AB , which means that M is visible from N . Hence, all of the edges of $CGT(S)$ that intersect segment NM must be unconstrained. By construction, the only edges of the graph between V and AB are the edges of the triangles $t_1, \dots, t_i$ that intersect NM . As each of these edges is unconstrained, V is also completely visible to AB .

Let polygon $P = t_1 \cup t_2 \cup \dots \cup t_i$. Let the upper part of the boundary of this polygon between V and A be referred to as the upper chain, and the lower part of the boundary between V and B be the lower chain. Note the lower chain is a single edge VB . If this were not the case, then when segment VB is considered by the greedy algorithm, there must be an edge e already intersecting it of shorter length. However, since V is the only vertex of the first i triangles to be inside the region $ABICG$, the endpoints of e must lie outside this region: above AG and below BI . Further, the edge e must lie to the right of C , simply because e must intersect segment NM . With these restrictions, the shortest e can possibly be is when its endpoints are located at g_2 and g_3 in Figure 5.2, and $|g_2g_3| = \frac{2}{3}$. However, since V is inside the region $OBIC$,

the longest VB can possibly be is $|BC| = \frac{\sqrt{3}}{3} \approx 0.577$. Therefore, this contradicts Lemma 7, and implies that VB must be an edge of the CGT .

However, the upper chain can either be a single edge VA , or be comprised of a number of segments. In the case of the latter, let $V = u_0, u_1, \dots, u_m, u_{m+1} = A$ denote the vertices along the upper chain. Note that at least one of the u_i 's between V and A must lie inside the semi-curved region AEG . If this were not the case, then every segment that crosses VA would be longer than AB , since we assume $u_1, \dots, u_m$ lie outside of the circle centered at B . Regardless of where V is located in the lower region $OBIC$, VA cannot be longer than AB . Hence, this is in contradiction to Lemma 7, and one of $u_1, u_2, \dots, u_m$ must lie inside AEG . In fact, this can be further refined by saying that the last vertex u_m must lie inside AEG . If u_j , with $j < m$, was the last vertex inside AEG , then every edge of the CGT that crosses the segment u_jA will have length greater than $|AB|$, which is longer than u_jA can possibly be.

Symmetrically, let Q be the visible (to AB) vertex inside triangle $\triangle ABD$ that is closest to the edge AB . All of the previous reasoning directly translates to this side of the edge, resulting in a vertex X (which duplicates the “role” of vertex V), and polygon P' . Due to loss of generality, it cannot be assumed that X resides in the bottom half of the region $ABJDH$. However, similar to V , if X lies in $OBJD$, then the lower chain of its polygon P' degenerates into a single edge XB , and the upper chain of P' will either be a single edge XA , or a chain of segments denoted $X = u'_0, u'_1, \dots, u'_m, u'_{n+1} = A$. Likewise, and if X lies in $OAHD$, then the upper chain of P' is an edge XA , and the lower chain will either be a single edge XB , or a chain denoted $X = l'_0, \dots, l'_{n+1} = B$.

The combination of the various constructions can be summarized as:

Case 1: V is located in $OBIC$, and the upper chain of P is a single edge VA .

Case 2: V is located in $OBIC$, and the upper chain of P is a chain of segments between vertices $u_1, u_2, \dots, u_m$.

Case 3: X is located in the bottom half $OBJD$, and the upper chain of polygon P' is a single edge XA .

Case 4: X is located in $OBJD$, and the upper chain of polygon P' a chain of segments between vertices $u'_1, u'_2, \dots, u'_n$.

Case 5: X is located in the upper half $OAHD$, and the lower chain of polygon P' is a single edge XB .

Case 6: X is located in $OAHD$, and the lower chain of polygon P' a chain of segments between vertices $l'_1, l'_2, \dots, l'_n$.

The contradictions arising from the eight possible combinations of these cases results in proof of Theorem 4.

5.2.3 Proof by Cases

Case 1 and Case 3:

Refer to Figure 5.4(a) on page 63. In this case, V and X are both located in the bottom half of the figure, and each upper chain of their respective polygons is simply a single edge. Therefore, the only edge of the greedy triangulation intersecting segment VX is AB . Note that no matter where V and X lie inside their respective regions, the length of VX is less than $\frac{\sqrt{3}}{2} \approx 0.866$, which is when V is located at I , and X is at J . Therefore, AB is the only edge intersecting VX , and it is longer than VX . Hence, by Lemma 7, AB cannot be an edge of the greedy triangulation, a contradiction.

Case 1 and Case 4:

Refer to Figure 5.4(b). As in the previous section, $|VX| < \frac{\sqrt{3}}{2}$. Note that the inequality is a “strictly less than”, due to V and X being defined as residing strictly inside their respective regions. The only edges of the greedy triangulation that intersect VX are AB , and any of $u'_1B, u'_2B, \dots, u'_nB$. Since each u'_i must lie outside of region $ABJDH$, and between segments BF and AB , the shortest any segment u'_iB can be is $|BH| = \cos(\frac{\pi}{6}) = \frac{\sqrt{3}}{2}$. Hence, by Lemma 7, VX must be an edge of the greedy triangulation.

Figure 5.4: The eight different combinations of cases

Case 1 and Case 5:

Refer to Figure 5.4(c). X is now located in the upper half of the figure, inside region $OAHD$, and the lower boundary of polygon P' is a single edge XB . The only edge of the greedy triangulation intersecting segment VX is AB . Note that the region $AIBH$ is a rectangle, with AB and HI as diagonals. $|VX|$ is maximized as $V \rightarrow B$ and $X \rightarrow A$, or as $V \rightarrow I$ and $X \rightarrow H$. However, since V and X are strictly within the boundary of this rectangle, $|VX| < 1$, which violates Lemma 7.

Case 1 and Case 6:

Refer to Figure 5.4(d). X lies in $OAHD$, and the lower boundary of polygon P' is a chain of segments between the lower vertices $l'_1, l'_2, \dots, l'_n$. The only edge of the greedy triangulation that intersects Vl'_n is AB . Since l'_n must lie within the region BJL , the length of Vl'_n is maximized when V is at I , and l'_n is at L ; in this configuration, $|Vl'_n| < |IL| = |ED| \approx 0.94$. (The other option for placement of V is at C , but $|CL| \approx 0.869 > |IL|$) Hence, $|Vl'_n| \leq |AB|$, which violates Lemma 7.

Case 2 and Case 3:

Refer to Figure 5.4(e). In this particular construction, both V and X are located in the lower half of the figure, and the upper boundary of P is a chain of segments. As in Case 1 and 4, the length of VX is less than $\frac{\sqrt{3}}{2}$. The only edges of the greedy triangulation that intersect VX are $u_1B, u_2B, \dots, u_mB$ and AB . The shortest any segment u_iB can be is when u_i is located at G , which means that $|u_iB| \geq |BG| = \frac{\sqrt{3}}{2}$, for all $i \in [1, m]$. Therefore, since $|VX|$ is strictly less than $|BG|$, Lemma 7 breaks this case.

Case 2 and Case 4:

Refer to Figure 5.4(f). This combination of cases is an extension of the previous case (2 & 3). $|VX|$ is still less than $\frac{\sqrt{3}}{2}$, and every intersecting edge $u_1B, \dots, u_mB, AB, u'_nB, \dots, u'_1B$ has length greater or equal to $\frac{\sqrt{3}}{2}$. Hence, the contradiction.

Case 2 and Case 5:

Refer to Figure 5.4(g). This configuration is symmetric to Case 1 and Case 6. Here, $|u_m X| < |EH| \approx 0.94 < |AB|$.

Case 2 and Case 6:

Refer to Figure 5.4(h). This configuration of V in the lower half, and X in the upper half, with both of their polygons having multi-edge upper and lower chains, respectively, results in the most difficult analysis of this proof. Consider segment $u_m X$. Note that the edges of $CGT(S)$ that intersect this segment are $AB, Al'_n, \dots, Al'_1$. Without much consideration, it appears that $|u_m X| < |EH| \approx 0.94$, while the segment Al'_1 could have length as little as $|AJ| \approx 0.866$, thereby breaking the proof by contradiction. A more sophisticated approach is needed.

(a) X must lie outside C_1 (b) and also outside of C_2 Figure 5.5: The location of X is restricted by the Greedy construction

First of all, recall that every edge $Al'_1, \dots, Al'_n$ must have length at least $\frac{\sqrt{3}}{2}$. Hence, if the distance between X and u_m was less than this, it would violate Lemma 7, and we would be done. Therefore, the distance between X and u_m must be at least $\frac{\sqrt{3}}{2}$. Regardless of where X is located in the region $OAHD$, the furthest u_m can possibly be from X is when it is located at point E (recall that u_m must be inside AEG).

Consider a circle C_1 centered at E of radius $\frac{\sqrt{3}}{2}$, as in Figure 5.5(a): X must be located outside, or on the boundary, of this circle.

In addition to $|u_m X|$ being at least $\frac{\sqrt{3}}{2}$, the distance between V and X must also be at least this much, otherwise Lemma 7 breaks this case. Therefore, regardless of where X is located outside C_1 , the furthest that V can be from X is when V is located at point I . Again, consider the circle C_2 centered at I with radius $\frac{\sqrt{3}}{2}$, as in Figure 5.5(b). X must also lie outside of C_2 .

(a) Determine the intersection points p_1 and p_2

(b) Calculate the location of X using angles σ and ρ

Figure 5.6: The location of X can be determined by σ and ρ

Hence, X must lie inside the region $R = OAH D \cap \overline{C_1} \cap \overline{C_2}$. Label the intersection of C_1 with AH as p_1 , and the intersection of C_2 with BH as p_2 . Omitting some rather ugly algebra (refer to Appendix A), it can be shown that the angle $\angle p_1 B H < 6^\circ = \frac{\pi}{30}$. Additionally, the angle $\angle p_2 A H < 18^\circ = \frac{\pi}{10}$. Since region R is a subset of the shaded quadrilateral in Figure 5.6(a), it is enough to prove this case when X is within the shaded region.

The location of X can now be expressed using the angles it makes with segments AH and BH — label these two angles $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$ respectively, as in Figure 5.6(b).

One final observation is required to complete the proof of this case. Recall Figure 5.3 on page 59. Just as V is defined as being the first vertex inside region $ABICG$ along the triangles $t_1, t_2, \dots$ from N to M , X is defined as being the first vertex inside region $ABJDH$ belonging to the triangles from AB towards Q . Due to Case 6, X is joined by an edge to the lower vertex l'_1 , which lies somewhere on or below BJ . Regardless of where X lies in the upper region, note that the segment XL'_1 must lie to the left of point D . If it did not, Q could not exist inside triangle $\triangle ABD$. Hence, every l'_i must lie to the left of the line through X and D . Label the point where this line crosses segment BJ as Y . Therefore, the shortest possible length of any edge Al'_j would be $|AY|$, which depends on the location of X , as illustrated in Figure 5.7.

Everything needed has now been defined for the proof of this final case. It is required to show that regardless of where X lies inside the shaded region, that $|EX| < |AY|$. If this were not the case, then it would be possible for some edge Al'_j to be inserted before $u_m X$ is considered, thereby allowing this case to be true, and breaking the proof by contradiction. Refer to Appendix A for the following analysis using Maple:

- The cartesian coordinates of X are computed by calculating the intersection of two lines: one through A and making an angle σ with AH , the other through

B making an angle ρ with BH .

- The equation of the line through X and D is computed by determining the slope of the line (rise over run), and substituting the coordinates of X into the equation to determine the y intercept.
- The location of Y on segment BJ is calculated by solving the intersection of the segment BJ with the line through XD .
- The distances $|EX|$ and $|AY|$ are calculated using the standard Euclidean method $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$.
- The function $f_{\sigma,\rho}$ is defined as $|AY| - |EX|$, and shown to have a value greater than 0 in the ranges $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$.
- The lowest value $f_{\sigma,\rho}$ obtains is approximately 0.3824, when $\sigma = 0, \rho \approx 0.05476$.

Figure 5.8: The function $f_{\sigma,\rho}$ over the ranges $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$

Therefore, regardless of the location of X , the segment $u_m X$ is shorter than any edge of the CGT that crosses it, thereby contradicting Lemma 7.

In Conclusion

It has been shown that regardless of the configuration of the vertices V and X , a valid construction of the constrained Greedy triangulation is impossible. Therefore, we can conclude that the original hypothesis of the existence of visible vertices on both sides of AB inside triangles $\triangle ABC$ and $\triangle ABD$ was false, thus proving that the CGT has the visible $\frac{\pi}{6}$ -diamond property. This concludes the proof of Theorem 4.

The size of the diamond in the above proof is an improvement over the original value of $\frac{\pi}{8} = 22.5^\circ$ shown by Das and Joseph [7], although the same technique was used. The result of Theorem 4 is also an improvement over the values shown by Dickerson et al. [10] and Drysdale et al. [13]. In [10], they prove that every edge e of the Greedy triangulation has a disc-shaped exclusion region, centered at the midpoint of e , of radius $\frac{|e|}{\sqrt{5}} \approx 0.447|e|$. The size of this region is extended in [13] to include the tangents to the region. By basic trigonometry, it can be shown that the largest diamond inscribed in this region has $\alpha = \arctan(1/\sqrt{5}) \approx 24.1^\circ$.

Therefore, $\frac{\pi}{6} = 30^\circ$ is now the current largest diamond for the Greedy and constrained Greedy triangulations. Further, if a circle is inscribed inside the diamond of an edge e , its radius is $\frac{|e|}{2} = 0.5|e|$, an improvement on the size of a disc-shaped exclusion region for the Greedy triangulation. These new regions are shown by the heavy dashed line in Figure 5.9 below.

Figure 5.9: New and previous exclusion regions for the Greedy triangulation

Chapter 6

Conclusion

In this thesis, we have expanded upon the conference paper [7] by Das and Joseph. The motivation of [7] was to show that the class of graphs which satisfy the α -diamond property and the d -good polygon property, are spanner graphs in the Euclidean metric. As a full journal version of the proof was never produced, this thesis proved the result in its entirety. The original proof of Theorem 1 in [7] was shown to be flawed, and corrections for it are provided here. An improvement was also made to the function bounding the spanning ratio for the class of α -diamond spanners.

The result of Theorem 1 was then generalized into the constrained setting, defined by the *visible* α -diamond property. Informally, it states that any constrained graph with this property (and the d -good polygon property) is a spanner of its visibility graph, meaning that a path can be constructed between any two vertices of such a constrained graph whose length is no more than a constant (based on α and d) times the shortest distance between the two vertices in the visibility graph.

It was then shown that two important graphs, the constrained Delaunay and constrained Greedy triangulations, both have the visible α -diamond property, for α equal to $\frac{\pi}{4}$ and $\frac{\pi}{6}$ respectively. The α value for the Greedy triangulation is an improvement over the previously best known value of $\tan^{-1}(1/\sqrt{5}) \approx 24^\circ$ by Dickerson et al. [10] and Drysdale et al. [13]. Two other types of graphs which were not considered in this thesis are the Minimum Weight triangulation (*MWT*), and the class of Locally Optimal triangulations (*LOTs*). A *LOT* has the property that if every quadrilateral

formed by two adjacent triangles is considered, each quadrilateral will either have a reflex vertex, or the common edge between the triangles will be the shorter of the two diagonals in the quadrilateral. Currently, the largest α value proven for the *MWT* is $\frac{\pi}{4.6}$, shown by Drysdale et al. [12], who also showed that the general class of Locally Optimal triangulations does not have the diamond property.

The use of the empty diamond, also referred to as *exclusion regions* in [12], has direct application to any algorithm attempting to build a sparse, spanning network. That is, it can be used as a “black box” to ensure that any graph with the diamond property is also a spanner. For example, Bose and Morin [3] show that with any triangulation satisfying the diamond property, their online competitive routing algorithm bounds the distance a packet must travel in the network. It is the hope of the author that results such as this one can be extended into the realm of the constrained graph, using the results of this thesis.

6.1 Future Work

Future work might consider:

- Improving the spanning ratio $c_{\alpha,d}$ from Theorem 1. Since the value of $c_{\alpha,d}$ is equal to 2 times the square of the one-sided path spanning ratio, c_{OSP} , any improvement to the bound of the one-side path will have a quadratic effect on the bound of the final path. Dobkin et al. [11] proved that a one-sided path in the Delaunay triangulation has a spanning ratio of $\frac{\pi}{2} \approx 1.57$. However, the spanning ratio for one-sided paths in the Delaunay triangulation using the diamond property is $\frac{2(3\pi/4)}{(\pi/4) \cdot \sin(\pi/16)} \approx 31$. Surely a tighter bound for α -diamond one-sided paths can be found. Otherwise, perhaps the use of induction, as in [16], can be used to construct the path from A to B .
- What other, possibly non-triangular, planar graphs have the diamond property and the good polygon property? In [7], the authors propose an algorithm to generate a sparse spanner graph, whose total edge weight is $O(|MST(P)|)$, where *MST* is the Minimum Spanning Tree of the point set P . Are there other

graph generating algorithms, or path generating algorithms, that can use the diamond property to ensure a bounded path length?

- Improving the size of the α -diamonds, or exclusion regions, for the graphs that are known to satisfy the diamond property. For example, in [12], the authors state that they suspect their analysis of the $\frac{\pi}{4.6}$ -diamond for the Minimum Weight triangulation is not tight. Additionally, the author of this thesis believes that the proven α value of $\frac{\pi}{6}$ for the Greedy triangulation is not tight either. Careful exploitation of the properties of these graphs would result in larger diamond regions.
- What different shaped exclusion regions exist for known graphs, besides the diamond-shaped, and disc-shaped? Perhaps elliptical-shaped or lune-shaped exclusion regions could be proven for the previously discussed triangulations.
- Finally, of all the previously known results using the α -diamond property, which of these can be extended into the constrained setting, using the visible diamond property? And can it be shown that the Constrained Minimum Weight triangulation has an exclusion region as well?

Appendix A

Maple Output for *CGT* Proof

This appendix contains a transcription of the output generated by the mathematical software Maple, used in the proof of Case 2 and Case 6 of Theorem 4.

Figure A.1: Co-ordinates of points on the Cartesian plane

Consider an edge AB of the constrained Greedy triangulation. Assume that the length of AB is 1, that the mid-point of this edge lies at the origin of the graph — at co-ordinates $(0,0)$ — and that the edge is aligned with the y -axis. With such a configuration, the co-ordinates of the points shown in Figure A.1 can be deduced using trigonometry. We now present the Maple code used in the proof.

The equation of the line through the segment BH is:

```
> l_BH := sqrt(3)*x - 1/2;
```

$$\ell_{BH} := \sqrt{3}x - \frac{1}{2}$$

And the equation of the line through A and H is:

```
> l_AH := -x/sqrt(3)+1/2;
```

$$\ell_{AH} := -\frac{\sqrt{3}}{3}x + \frac{1}{2}$$

Next, consider the general equation for a circle:

```
> c_general := (x-a)^2 + (y-b)^2 = r^2;
```

$$C_{general} := (x - a)^2 + (y - b)^2 = r^2$$

Using the co-ordinates for point E , compute the equation for the circle centered at E with a radius of $\frac{\sqrt{3}}{2}$, and solve for y :

```
> c_E := solve(subs(a=-1/2,b=(sqrt(3)-1)/2,r=sqrt(3)/2,c_general),y);
```

$$C_E := \frac{\sqrt{3}}{2} - \frac{1}{2} + \frac{\sqrt{2-4x-4x^2}}{2}, \frac{\sqrt{3}}{2} - \frac{1}{2} - \frac{\sqrt{2-4x-4x^2}}{2}$$

Repeat for the circle centered at I :

```
> c_I := solve(subs(a=-sqrt(3)/4,b=-1/4,r=sqrt(3)/2,c_general),y);
```

$$C_I := -\frac{1}{4} + \frac{\sqrt{9-16x^2-8\sqrt{3}x}}{4}, -\frac{1}{4} - \frac{\sqrt{9-16x^2-8\sqrt{3}x}}{4}$$

Determine the x -coordinate of p_1 by solving the intersection of circle C_E with the line through AH :

```
> p1_x := simplify(solve(c_E[2]=l_AH,x));
```

$$p1_x := -\frac{3}{4} + \frac{\sqrt{3}}{4} + \frac{\sqrt{-3+6\sqrt{3}}}{4}$$

Substitute this value back into the equation for the line through AH to determine the y -coordinate of p_1 :

```
> p1_y := simplify(subs(x=p1_x, l_AH));
```

$$p1_y := \frac{\sqrt{3}}{4} + \frac{1}{4} - \frac{\sqrt{3}\sqrt{-3+6\sqrt{3}}}{12}$$

Repeat this process for p_2 , intersecting circle C_I with line BH :

```
> p2_x := simplify(solve(c_I[1]=l_BH, x));
```

$$p2_x := \frac{\sqrt{2}}{4}$$

```
> p2_y := simplify(subs(x=p2_x, l_BH));
```

$$p1_y := \frac{\sqrt{3}\sqrt{2}}{4} - \frac{1}{2}$$

Load the geometry package, with the useful distance function. Define points p_1 , p_2 , and H .

```
> with(geometry):
> point(p_1, p1_x, p1_y):
> point(p_2, p2_x, p2_y):
> point(p_H, sqrt(3)/4, 1/4):
```

Using trigonometry, determine the angle $\angle p_1BH$, and show that it is less than $\frac{\pi}{30}$:

```
> max_rho := arctan(distance(p_1, p_H)/(sqrt(3)/2));
```

$$max_\rho := \arctan \left(\frac{2\sqrt{3} \sqrt{\left(-\frac{3}{4} + \frac{\sqrt{-3+6\sqrt{3}}}{4}\right)^2 + \left(\frac{\sqrt{3}}{4} - \frac{\sqrt{3}\sqrt{-3+6\sqrt{3}}}{12}\right)^2}}{3} \right)$$

```
> evalf(max_rho);
```

.09343404434

```
> evalf(Pi/30);
```

.1047197551

Repeat for $\angle p_2AH$:

```
> max_sigma := arctan(distance(p_2,p_H)/(1/2));
```

$$max_\sigma := \arctan \left(2 \sqrt{\left(\frac{\sqrt{2}}{4} - \frac{\sqrt{3}}{4} \right)^2 + \left(\frac{\sqrt{3}\sqrt{2}}{4} - \frac{3}{4} \right)^2} \right)$$

```
> evalf(max_sigma);
```

.3077398543

```
> evalf(Pi/10);
```

.3141592654

Next, we want to express the co-ordinates of X using σ and ρ . The equations of the lines through BX and AX are:

```
> l_X_rho := tan(Pi/3 + rho)*x -1/2;
```

$$\ell_{X\rho} := \tan\left(\frac{\pi}{3} + \rho\right)x - \frac{1}{2}$$

```
> l_X_sigma := -tan(Pi/6 +sigma)*x +1/2;
```

$$\ell_{X\sigma} := -\tan\left(\frac{\pi}{6} + \sigma\right)x + \frac{1}{2}$$

Determine the location X in terms of ρ and σ by solving the intersection of these two lines:

```
> x1 := solve(l_X_rho=l_X_sigma,x);
```

$$x1 := \frac{1}{\tan\left(\frac{\pi}{3} + \rho\right) + \tan\left(\frac{\pi}{6} + \sigma\right)}$$

```
> x2 := subs(x=x1,l_X_rho);
```

$$x2 := \frac{\tan\left(\frac{\pi}{3} + \rho\right)}{\tan\left(\frac{\pi}{3} + \rho\right) + \tan\left(\frac{\pi}{6} + \sigma\right)} - \frac{1}{2}$$

```
> point(p_X,x1,x2):
```

```
> point(p_D,sqrt(3)/6,0):
```

Consider the line through XD . Determine the slope of this line using rise-over-run:

```
> l_XD_slope := (coordinates(p_X)[2]-coordinates(p_D)[2])/
                 (coordinates(p_X)[1]-coordinates(p_D)[1]);
```

$$\ell_{XD_slope} := \frac{\frac{\tan(\frac{\pi}{3}+\rho)}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{1}{2}}{\frac{1}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{\sqrt{3}}{6}}$$

Then compute the y -intercept by substituting the slope and x -coordinate of D into the equation for the line:

```
> l_XD_yint := -l_XD_slope*coordinates(p_D)[1];
```

$$\ell_{XD_yint} := -\frac{\sqrt{3}}{6} \left(\frac{\frac{\tan(\frac{\pi}{3}+\rho)}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{1}{2}}{\frac{1}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{\sqrt{3}}{6}} \right)$$

The equation of the line through XD is now:

```
> l_XD := l_XD_slope*x + l_XD_yint;
```

$$\ell_{XD} := \frac{\left(\frac{\tan(\frac{\pi}{3}+\rho)}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{1}{2} \right) x - \frac{\sqrt{3}}{6} \left(\frac{\frac{\tan(\frac{\pi}{3}+\rho)}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{1}{2}}{\frac{1}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{\sqrt{3}}{6}} \right)}{\frac{1}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)} - \frac{\sqrt{3}}{6}}$$

Next, the location of point Y on segment BJ must be computed. Determine the coordinates of Y by solving the intersection of the lines through XD and BJ (note that the line through BJ is the inverse of the line through AH):

```
> y1 := solve(l_XD=-l_AH,x);
```

$$y1 := \frac{1}{2} \left(\frac{-3 + \sqrt{3} \tan(\frac{\pi}{3} + \rho)}{2 \tan(\frac{\pi}{3} + \rho) - \tan(\frac{\pi}{6} + \sigma) - \sqrt{3}} \right)$$

```
> y2 := subs(x=y1,-l_AH);
```

$$y2 := \frac{\sqrt{3}}{6} \left(\frac{-3 + \sqrt{3} \tan(\frac{\pi}{3} + \rho)}{2 \tan(\frac{\pi}{3} + \rho) - \tan(\frac{\pi}{6} + \sigma) - \sqrt{3}} \right) - \frac{1}{2}$$

```
> point(p_Y,y1,y2):
```

```
> point(p_A,0,1/2):
```

```
> point(p_E,-1/2,(sqrt(3)-1)/2):
```

The distances $|AY|$ and $|EX|$ can now be computed:

```
> d_EX := distance(p_E,p_X):
```

```
> d_AY := distance(p_Y,p_A):
```

The complete equation of the function $f_{\sigma,\rho}$ is now known:

```
> f := d_AY - d_EX;
```

$$f := \sqrt{\frac{(-3+\sqrt{3}\tan(\frac{\pi}{3}+\rho))^2}{4(2\tan(\frac{\pi}{3}+\rho)-\tan(\frac{\pi}{6}+\sigma)-\sqrt{3})^2} + \left(\frac{1}{6}\frac{\sqrt{3}(-3+\sqrt{3}\tan(\frac{\pi}{3}+\rho))}{2\tan(\frac{\pi}{3}+\rho)-\tan(\frac{\pi}{6}+\sigma)-\sqrt{3}} - 1\right)^2} - \sqrt{\left(-\frac{1}{2} - \frac{1}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)}\right)^2 + \left(\frac{\sqrt{3}}{2} - \frac{\tan(\frac{\pi}{3}+\rho)}{\tan(\frac{\pi}{3}+\rho)+\tan(\frac{\pi}{6}+\sigma)}\right)^2}$$

The value of f can now be computed over the ranges of $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$:

```
> plot3d(f,rho=0..Pi/30,sigma=0..Pi/10,axes=frame,filled=true,
      shading=zgreyscale,numpoints=225,style=patch,projection=0.95,
      orientation=[-55,70]);
```

Figure A.2: The function f over the ranges $\sigma \in [0, \frac{\pi}{10}]$ and $\rho \in [0, \frac{\pi}{30}]$

It is clear that the value of f within the ranges of σ and ρ is positive. However, to verify it completely, one could take the partial derivative of f with respect to σ . This

function does have values equal to 0, however f achieves a maximum at these values. Therefore, f is minimized at either of the two boundary cases ($\sigma = 0$ and $\sigma = \frac{\pi}{10}$). Due to the periodic nature of the tangent function, solving for where these minimums occur produce some numerically ugly solutions. As such, a plot of the values of f at these boundary cases follows:

```
> plot([subs(sigma=0,f),subs(sigma=Pi/10,f)],rho=0..Pi/30);
```

Figure A.3: The function f over the range $\rho \in [0, \frac{\pi}{30}]$, when $\sigma = 0$ and $\frac{\pi}{10}$

The lower value is when $\sigma = 0$. To compute where the minimum occurs, let $\sigma = 0$ in f , and differentiate f with respect to ρ and solve where it is equal to 0

```
> sol := solve(diff(subs(sigma=0,f),rho)=0,rho):
> evalf(sol);
```

$sol = .054757673$

Finally, substitute this value of ρ back into f to determine the minimum value f obtains:

```
> subs(rho=sol,sigma=0,f):
> evalf(%);
```

$.0382437819$

Bibliography

- [1] R. Beirouti and J. Snoeyink. Implementations of the lmt heuristic for minimum weight triangulations. In *Proceedings of the Fourteenth Annual ACM Symposium on Computational Geometry*, pages 96–105, 1998.
- [2] R. V. Benson. *Euclidean Geometry and Convexity*. McGraw-Hill, New York, NY, 1966.
- [3] P. Bose and P. Morin. Competitive online routing in geometric graphs. *Theoretical Computer Science*, 324:273–288, 2004.
- [4] L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proc. 2nd Annu. ACM Sympos. Comput. Geom.*, pages 169–177, 1986.
- [5] L. P. Chew. Constrained Delaunay triangulations. In *Proc. 3rd Annu. ACM Sympos. Comput. Geom.*, pages 215–222, 1987.
- [6] G. Das. *Approximation Schemes in Computational Geometry*. Ph.D. thesis, Dept. Comput. Sci., University of Wisconsin, Madison, Wisconsin, 1990.
- [7] G. Das and D. Joseph. Which triangulations approximate the complete graph. In *Proc. International Symposium on Optimal Algorithms*, volume 401 of *Lecture Notes Comput. Sci.*, pages 168–192. Springer-Verlag, 1989.
- [8] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.
- [9] B. Delaunay. Sur la sphère vide. A la memoire de Georges Voronoi. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennyh Nauk*, 7:793–800, 1934.

- [10] M. T. Dickerson, R. L. S. Drysdale, S. A. McElfresh, and E. Welzl. Fast greedy triangulation algorithms. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 211–220, 1994.
- [11] D. P. Dobkin, S. J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. In *Proc. 28th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 20–26, 1987.
- [12] R. L. Drysdale, S. McElfresh, and J. S. Snoeyink. An improved diamond property for minimum-weight triangulations. In *Abstracts 14th European Workshop Comput. Geom.*, pages 65–67, 1998.
- [13] R. L. Drysdale, G. Rote, and O. Aichholzer. A simple linear time greedy triangulation algorithm for uniformly distributed points. Technical Report IIG-408, Institutes for Information Processing, Technische Universität Graz, February 1995.
- [14] F. Harary. *Graph Theory*. Addison-Wesley, Reading, Mass., 1969.
- [15] M. I. Karavelas. *Proximity Structures for Moving Objects in Constrained and Unconstrained Environments*. Ph.D. thesis, Dept. Comput. Sci., Stanford University, 2001.
- [16] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, 7:13–28, 1992.
- [17] D. T. Lee. Proximity and reachability in the plane. Report R-831, Dept. Elect. Engrg., Univ. Illinois, Urbana, IL, 1978.
- [18] G. K. Manacher and A. L. Zobrist. Neither the greedy nor the Delaunay triangulation of a planar point set approximates the optimal triangulation. *Inform. Process. Lett.*, 9:31–34, 1979.
- [19] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, NY, 1985.

- [20] J.-R. Sack and J. Urrutia, editors. *Handbook of Computational Geometry*. Elsevier Science, 2000.
- [21] C. A. Wang. An optimal algorithm for greedy triangulation of a set of points. In *Proc. 6th Canad. Conf. Comput. Geom.*, pages 332–338, 1994.
- [22] C. A. Wang and L. Schubert. An optimal algorithm for constructing the Delaunay triangulation of a set of line segments. In *Proc. 3rd Annu. ACM Sympos. Comput. Geom.*, pages 223–232, 1987.