On Generalized Diamond Spanners*

Prosenjit Bose Aaron Lee Michiel Smid

Abstract

Given a set P of points in the plane and a set L of non-crossing line segments whose endpoints are in P, a *constrained* plane geometric graph is a plane graph whose vertex set is P and whose edge set contains L. An edge e has the α-visible diamond property if one of the two isosceles triangles with base e and base angle α does not contain any points of P visible to both endpoints of e. A constrained plane geometric graph has the d-good polygon property provided that for every pair x, y of visible vertices on a face f, the shorter of the two paths from x to y around the boundary has length at most $d \cdot |xy|$. If a constrained plane geometric graph has the α-visible diamond property for each of its edges and the d-good polygon property, we show it is a $\frac{8d(\pi-\alpha)^2}{\alpha^2 \sin^2(\alpha/4)}$-spanner of the visibility graph of P and L. This is a generalization of the result in Das and Joseph [3] to the constrained setting as well as a slight improvement on their spanning ratio of $\frac{8d\pi^2}{\alpha^2 \sin^2(\alpha/4)}$. We then show that several well-known constrained triangulations (namely the constrained Delaunay triangulation, constrained greedy triangulation and constrained minimum weight triangulation) have the α-visible diamond property for some constant α. In particular, we show that the greedy triangulation has the $\pi/6$-visible diamond property, which is an improvement over previous results.

1 Introduction

A graph G whose vertices are points in the plane and edges are segments weighted by their length is a t-spanner (for $t \geq 1$) provided that the shortest path in G between any two vertices x, y does not exceed $t|xy|$ where $|xy|$ is the Euclidean distance between x and y. The value t is the *spanning ratio* or *stretch factor* of the graph. The spanning properties of various geometric graphs has been studied extensively in the literature (see Eppstein [7], Knauer and Gudmundsson [10], Narasimhan and Smid [13], Smid [14] for several surveys on the topic). Our work is a generalization of the result by Das and Joseph [3] to the constrained setting. Das and Joseph [3] showed that any graph possessing the *diamond* property and the *good polygon* property is a t-spanner where the constant t depends on parameters of each of the two properties.

Before we can state our results precisely, we outline what these properties are, what we mean by the constrained setting and how the spanning ratio of a geometric graph is measured in this setting. Throughout this paper, a graph will refer to a *geometric 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 by their length. Let P denote a set of points in the plane and L be a set of non-crossing line segments whose endpoints are in P. Two points p and q of P are *visible* with respect to L provided the segment pq does not properly intersect any segment of L. Two line segments intersect

*Research supported in part by NSERC. School of Computer Science, Carleton University, Ottawa, Canada. Email: {jit, michiel }@scs.carleton.ca, aaron@aaronlee.ca

properly if they share a common interior point. The visibility graph of P constrained to L, denoted $Vis(P, L)$, is a geometric graph whose vertex set is P and whose edge set contains L as well as one edge for each visible pair of vertices (See Figure 1). A spanning subgraph of $Vis(P, L)$ whose edge set contains L is a geometric graph *constrained to* L. In such a graph, the set L is referred to as the *constrained edges* and all other edges are referred to as *unconstrained edges* or *visibility edges*.

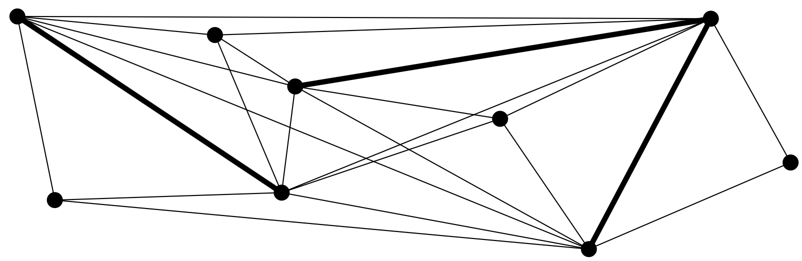

Figure 1: The visibility graph $Vis(P, L)$ where segments of L are shown in bold.

Definition 1. Let $t \geq 1$ be a real number. A constrained geometric graph $G(P, L)$ is a *constrained t-spanner* provided that for every visibility edge $[pq]$ in $Vis(P, L)$, the length of the shortest path between p and q in $G(P, L)$ is at most t times the Euclidean distance between p and q. We refer to t as the *spanning ratio* or the *stretch factor* of $G(P, L)$.

Note that if $G(P, L)$ is a constrained t-spanner, then for every pair of points p, q in P (not just visible edges), the shortest path from p to q in $G(P, L)$ is at most t times the shortest path from p to q in $Vis(P, L)$. We now define the two essential properties.

Definition 2. Refer to Figure 2. Fix $\alpha \in (0, \pi/2)$. A constrained graph $G(P, L)$ is said to have the *visible α-diamond property* if, for every unconstrained edge e in the graph, at least one of the two isosceles triangles, with e as the base and base angle α, does not contain any points of P visible to the endpoints of e. We label this empty triangle as $\triangle(e)$, and the apex of $\triangle(e)$ as $a(e)$.

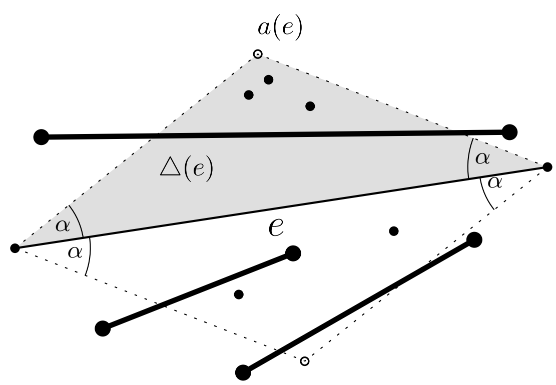

Figure 2: The edge e has the visible α-diamond property

Definition 3. A constrained plane graph $G(P, L)$ has the *d-good polygon property* if for every visible pair of vertices a and b on a face 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.

Our main results are the following:

Theorem 1. *Given fixed $\alpha \in (0, \pi/2)$ and $d \geq 1$, if a constrained plane graph $G(P, L)$ has both the visible α-diamond property and the d-good polygon property, then its stretch factor is at most $\frac{8(\pi-\alpha)^2 d}{\alpha^2 \sin^2(\alpha/4)}$.*

This is a generalization of the result in Das and Joseph [3] to the constrained setting as well as a slight improvement on the spanning ratio from $\frac{8d\pi^2}{\alpha^2 \sin^2(\alpha/4)}$ to $\frac{8d(\pi-\alpha)^2}{\alpha^2 \sin^2(\alpha/4)}$.

Theorem 2. *The Constrained Greedy Triangulation has the visible $\frac{\pi}{6}$-diamond property.*

This is an improvement over the results of Das and Joseph [3], Dickerson *et al.* [4] and Drysdale *et al.* [6] on this problem. In Das and Joseph [3], they showed that the Greedy Triangulation has the $\pi/8$-diamond property. The results in Drysdale *et al.* [6], which is an extension of the results in Dickerson *et al.* [4], imply that the Greedy triangulation has the $\arctan(1/\sqrt{5})$-diamond property. Note that $\arctan\left(1/\sqrt{5}\right) \approx 24.1°$.

2 Constructing Spanner Paths

The proof of the main result is constructive. Consider a constrained plane graph $G(P, L)$ that has the visible α-diamond property and the d-good polygon property. Given a pair of points $a, b \in P$ that are visible with respect to L, we show how to construct a path from a to b in $G(P, L)$ whose length is at most $\frac{8d(\pi-\alpha)^2}{\alpha^2 \sin^2(\alpha/4)}$ times the Euclidean distance between a and b.

If $[ab]$ is an edge of $G(P, L)$ then such a path trivially exists. Therefore, assume that $[ab]$ is not an edge of the graph. In this case, either $[ab]$ intersects some edges of the graph or intersects no edges of the graph. In the latter case, this means that the segment $[ab]$ is a chord in a face of $G(P, L)$. The d-good polygon property ensures that the required path exists in this case. In the remainder of this section, we show that when $[ab]$ intersects some edges of $G(P, L)$, we can construct a spanner path from a to b.

Re-orient the coordinate system such that $[ab]$ lies on the x-axis. Let $e_1, e_2, \ldots, e_k$ be the edges of $G(P, L)$ that cross $[ab]$ in order from a to b. For simplicity of exposition, assume that none of these edges share a common endpoint. Sharing endpoints is a degenerate situation that only makes the proof simpler. Label the endpoint of e_i above $[ab]$ as u_i and the endpoint below $[ab]$ as l_i. The fact that a and b are visible with respect to L ensures that each of these edges is an unconstrained edge. Moreover, the visible α-diamond property implies that each of these edges is the base of a visibly empty triangle $\triangle(e_i)$ with apex $a(i)$, $1 \leq i \leq k$. Since $G(P, L)$ is a plane graph, e_i and e_{i+1} lie on a common face f. Let U_i be the shortest path from u_i to u_{i+1} in the face f. Since a and b are visible, this implies that U_i is a convex path. Similarly, let L_i be the shortest path from l_i to l_{i+1} in f. Note that the d-good polygon property ensures that there is path in $G(P, L)$ from u_i to u_{i+1} whose length does not exceed d times the length of U_i. Define U_0 (resp. U_k) to be the shortest path from a to u_1 (resp. u_k to b). L_0 and L_k are defined symmetrically. See Figure 3.

We have two different construction methods depending on where the apices of the empty triangles are with respect to $[ab]$. If all of the apices lie on the same side of the line through $[ab]$, we

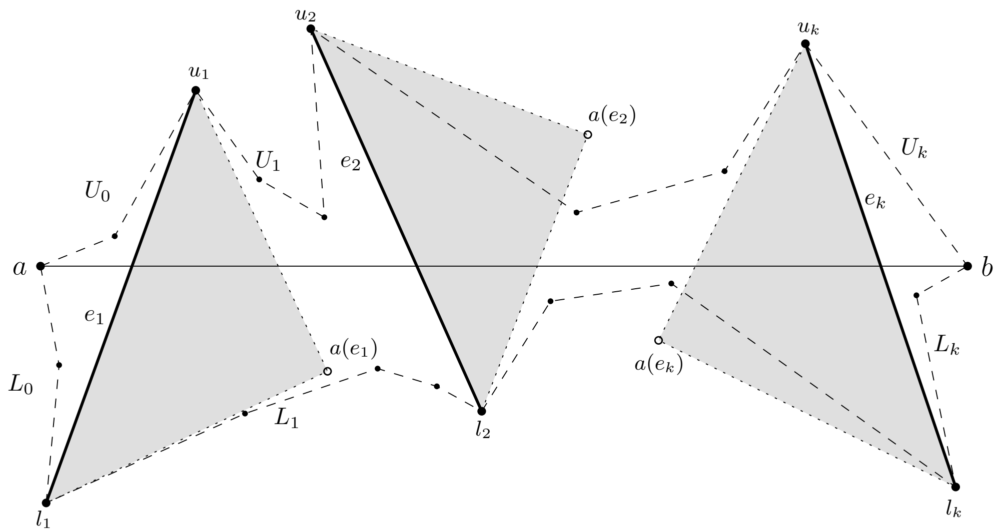

Figure 3: Illustration of Structures.

construct a path called a *one-sided* path, otherwise, we construct a *two-sided* path. We first show the construction of one-sided paths and bound their length.

2.1 One-sided paths

In this case, we assume that all the apices of the empty triangles lie below $[ab]$. The construction of the one-sided path starts with the union of the U_i. Now, each edge e in U_i can be approximated by a path in $G(P, L)$ whose length is at most $d \cdot e$. Therefore, the length of the one-sided path is at most $d(|U_0| + |U_1| + \ldots + |U_k|)$. What remains to be shown is that this is a good approximation of $|ab|$ since a and b are visible. Let h be the line through $[ab]$ and h^- be the closed half-plane below h. To obtain a bound on $|U_0| + |U_1| + \ldots + |U_k|$, we consider the following structure $T = h^- \cup \bigcup_{i=1}^{k} \triangle(e_i)$ (See Figure 4). Denote by $T(a, b)$ the portion of the boundary of T between a and b. The fact that each of the triangles $\triangle(e_i)$ is empty of points visible to both endpoints of the edge e_i, all the apices of the empty triangles are below $[ab]$ and each of the U_i is formed by a shortest path imply that no edge of $T(a, b)$ can intersect any of the upper chains U_i. Since each of the upper chains U_i is a shortest path, we conclude the following:

Lemma 1. $\sum_{i=0}^{k} |U_i| \leq |T\,(a, b)\,|$

Before proceeding, we need a simple property about triangles that essentially follows from the sine law.

Lemma 2. *Given a triangle $\triangle(u, v, w)$ such that angle at vertex v is α, we have that $|uv| + |vw| \leq |uw| / \sin(\alpha/2)$.*

We now show how to bound the length of the one-sided path in terms of $|ab|$.

Lemma 3. *The length of a one-sided path from a to b in $G(P, L)$ is at most $|ab| \frac{2(\pi-\alpha)d}{\alpha \cdot \sin(\alpha/4)}$.*

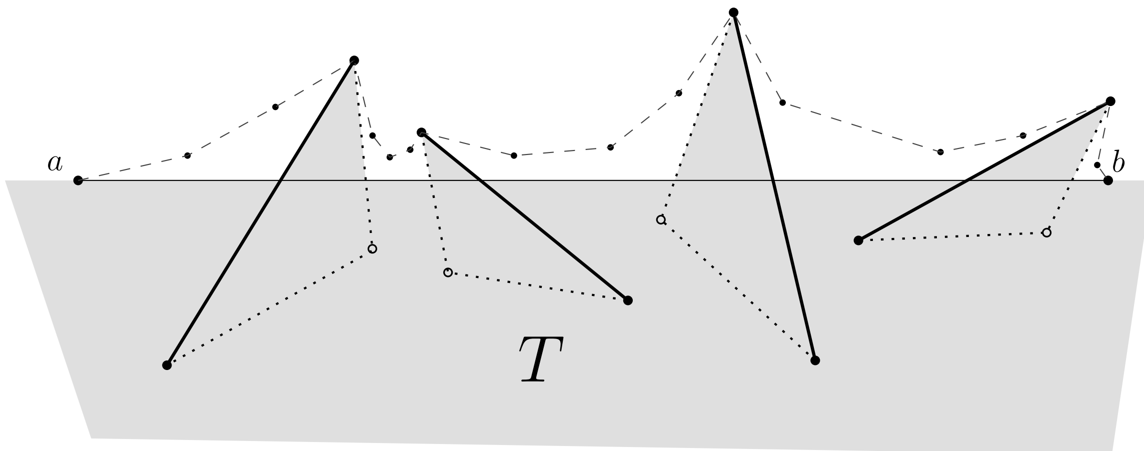

Figure 4: Using the shaded region T to approximate the length of the upper chains.

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

To simplify the exposition, we assume that π is a multiple of α (a condition that can be easily removed). Partition the empty triangles $\triangle(e_1), \triangle(e_2), \ldots, \triangle(e_k)$ into $\frac{2(\pi-\alpha)}{\alpha}$ groups labelled $G_0, G_{\alpha/2}, \cdots, G_\theta, \cdots, G_{\pi-3\alpha/2}$, such that the left edges of the empty triangles in group G_θ make an angle in the range $\left[\theta, \theta+\frac{\alpha}{2}\right]$ with the x-axis. Since the base angle of the empty triangles is α, we see that the right edges will be in the range $\left[\theta+\alpha, \theta+\frac{3\alpha}{2}\right]$ with the x-axis. This is why the last group ends at $\pi - 3\alpha/2$.

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 $T = T_0 \cup T_{\alpha/2} \cup \cdots \cup T_{\pi-3\alpha/2}$. Hence, it follows that the length of the boundary of T from a to b is bounded by $|T_0(a,b)| + \left|T_{\alpha/2}(a,b)\right| + \cdots + \left|T_{\pi-3\alpha/2}(a,b)\right|$

Consider the group G_θ, as shown in Figure 5. The edges of the boundary $T_\theta(a,b)$ are shown in bold. Let p be the point such that $\measuredangle apb = \frac{\alpha}{2}$, $\measuredangle pab = \theta + \frac{\alpha}{2}$, and $p \in h^+$. By construction, the portion of the triangles in G_θ that lie above the x-axis (and thus $T_\theta(a,b)$) are completely contained inside $\triangle pab$.

$T_\theta(a,b)$ is a polygonal chain consisting of portions of left edges of empty triangles, portions of right edges of empty triangles, and portions of $[ab]$. Note that the angle restriction implies that the chain is monotone both in the direction pa and the direction pb. To bound the length of an edge xy of $T_\theta(a,b)$, project x and y onto pa by translating in a direction parallel to pb and denote the projected vertices on pa by x_l and y_l, respectively. Similarly, project xy onto pb by translating in a direction parallel to pa resulting in projected vertices x_r and y_r, respectively. The triangle inequality guarantees that $|xy| \leq |x_l y_l| + |x_r y_r|$. Monotonicity guarantees that none of the projected edges of $T_\theta(a,b)$ overlap. Therefore, we have that $|T_\theta(a,b)| \leq |pa| + |pb|$.

Using Lemma 2, it follows that $|pa| + |pb| \leq \frac{|ab|}{\sin(\alpha/4)}$, regardless of the angle θ. Therefore, since there are $\frac{2(\pi-\alpha)}{\alpha}$ many groups, $|T(a,b)| \leq \frac{2(\pi-\alpha)|ab|}{\alpha\cdot\sin(\alpha/4)}$. □

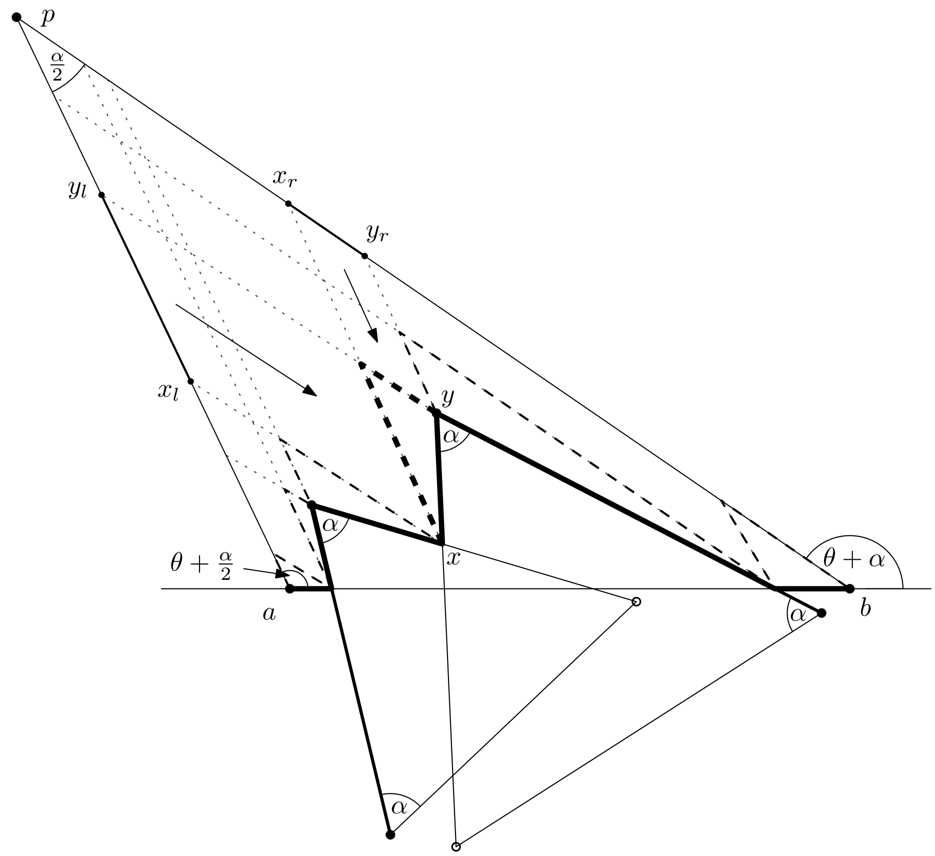

Figure 5: The triangles of G_θ

2.2 Two-Sided Paths

We have seen that if every apex $a(e_1), a(e_2), \ldots, a(e_k)$ 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 $|ab|$. We now outline the process of constructing a short path from a to b in the case where some apices $a(e_i)$ lie above the x-axis, while others lie below. The one-sided paths either followed the upper chain or the lower chain. In two-sided paths, we may need to cross over from upper chains to lower chains. Recall that U_i is the shortest path from u_i to u_{i+1} and L_i is the shortest path from l_i to l_{i+1}. For each pair of upper and lower chains, U_i and L_i, respectively, we add the unique edge on the shortest path from u_i to l_{i+1} that is not on either chain and the unique edge on the shortest path from l_i to u_{i+1} that is not on either chain. We refer to these two edges as *tangents* between the upper and lower chains.

Divide the set of edges $e_1, e_2, \ldots, e_k$ 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 $T_U = h^- \cup \bigcup_{e \in U} \triangle(e)$. Correspondingly, define the region $T_L = h^+ \cup \bigcup_{e \in L} \triangle(e)$. Let $T_U\,(a,b)$ denote the upper boundary of T_U between a and b and 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 $|ab|\frac{2(\pi-\alpha)d}{\alpha \cdot \sin(\alpha/4)}$ as shown in Lemma 3. The two-sided path from a to b is constructed using disjoint portions of T_U, T_L and tangents. Since $|T_U(a,b)| + |T_L(a,b)| \le 2|ab|\frac{2(\pi-\alpha)d}{\alpha \cdot \sin(\alpha/4)}$, we only need to bound the length of the tangents used.

The two-sided path from a to b is constructed as follows. Without loss of generality, assume that e_1 has its apex below $[ab]$. Let e_{i+1} be the first edge whose apex is above $[ab]$. If no subsequent edge has an apex below $[ab]$, then the path follows the upper chains from a to u_i, follows the path from u_i to l_{i+1} and follows the lower chains to b. This path has length at most $|T_U(a,b)| + |T_L(a,b)|$ since each of the two paths is one-sided and the length of the tangent is subsumed by the unused portions of the upper and lower chain.

The situation where a decision needs to be made on how to proceed is when there are two edges e_i and e_j (with $j > i+1$) having apices $a(e_i)$ and $a(e_j)$ below the x-axis and at least one edge e_k (with $i < k < j$) with apex $a(e_k)$ above the x-axis. We show how to construct a short path from u_i to u_j in this case. There are two possibilities in this case, either the path from u_i to u_j follows the upper chain, or it follows the path from u_i to l_{i+1}, continues on the lower chain until l_{j-1} and follows the path from l_{j-1} to u_j.

The decision whether or not to cross over from the upper chain to the lower chain depends on the tangents. Let $t_a = [u_a, l_a]$ be the tangent on the path from u_i to l_{i+1} and $t_b = [u_b, l_b]$ be the tangent on the path from l_{j-1} to u_j. 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 $\theta = 0$. There are two cases to consider depending on the angle θ.

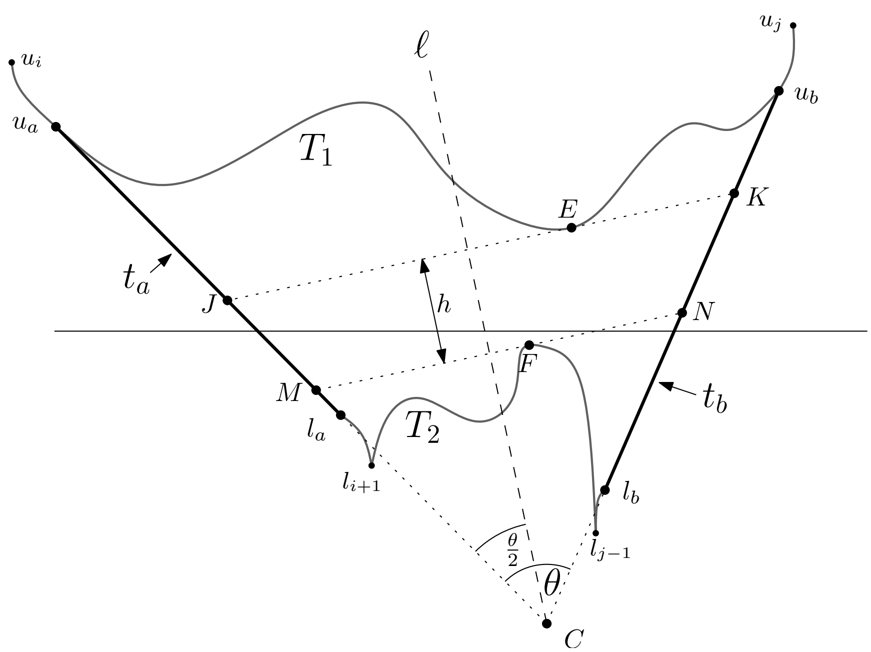

Figure 6: Case 2: Construct an empty region

Case 1: $\theta \ge \alpha$. In this case, the angle between the tangents is relatively wide. Follow the path from u_i to l_{i+1} using the tangent t_a. Continue on the lower convex chains $L_{i+1} \cdots L_{j-1}$ from l_a until l_b, and cross back up along t_b to u_j.

Case 2: $\theta < \alpha$. Refer to Figure 6 for this case. 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 $h \leq \frac{w}{2\tan(\alpha/2)}$, then follow the same path from u_i to u_j as in Case 1.

b) If $h > \frac{w}{2\tan(\alpha/2)}$, then simply follow the upper chains $U_i, U_{i+1}, \ldots, U_{j-1}$ from u_i and u_j.

We now bound the lengths of the paths constructed in the context of the different cases stated above. Note that the main difficulty is in bounding the length of the tangents since we have a bound on the length of the upper and lower chains.

Lemma 4. *If $\theta \geq \alpha$, the length of the portion of the two-sided path from a to b that is between u_a and u_b is at most $\frac{2}{\sin(\alpha/2)}\big(|T_1| + |T_2|\big)$.*

Proof. Consider the triangle $\triangle(u_a, C, u_b)$. The path follows the two tangents and the lower chain. The length of the lower chain is $|T_2|$. We want to bound the length of the two tangents in terms of $|T_1|$ and $|T_2|$. If C is below $[ab]$, by Lemma 2, we have that $|u_aC| + |Cu_b| \leq |T_1|/\sin(\alpha/2)$. If C is above $[ab]$, by Lemma 2, we have that $|u_aC| + |Cu_b| \leq |T_2|/\sin(\alpha/2)$. Note that these above bounds on the tangents only hold when C does not lie on one of the tangents. Should C land on one of the tangents then an extra $|T_1|$ or $|T_2|$ term can bound the portion of the tangent that lies outside the triangle by the triangle inequality. Putting all the inequalities together completes the proof. □

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

Proof. (Sketch): The path constructed in this case is identical to the path constructed in Case 1. The path follows the two tangents and the lower chains. The length of the lower chains is $|T_2|$. We need to bound the length of the tangents. Refer to Figure 6.

Note that the tangent t_a is decomposed into three segments: $[u_aJ], [JM], [Ml_a]$. Similarly t_b is decomposed into three segments: $[u_bK], [KN], [Nl_b]$. Since the angle at J in triangle $\triangle(u_aJE)$ is obtuse, we have that $|u_aJ|$ is shorter than the portion of T_1 from u_a to E. By this argument, we have that $|u_aJ| + |u_bK| < |T_1| + |T_2|$. To bound $|JM|$ and $|KN|$, we use the fact that $h \leq \frac{w}{2\tan(\alpha/2)}$. This allows us to show that $|JK| + |KN| \leq 2h/\cos(\alpha/2) \leq |T_1|/\sin(\alpha/2)$. Finally, by elementary trigonometry, we have that $|l_aM| + |l_bN| \leq (|l_aF| + |l_bF|)/\sin(\alpha/2) \leq |T_2|/\sin(\alpha/2)$. Combining the inequalities, we have that the length of the two-sided path is at most $2(|T_1| + |T_2|)/\sin(\alpha/2)$.

Note that this bound only holds when C does not lie on one of the tangents. However, C may lie on one of the two tangents. In this case, an extra $|T_1|/\sin(\alpha/2)$ or $|T_2|/\sin(\alpha/2)$ term needs to be added giving the stated bound*. □

*This is one of the cases that was omitted from the original proof by Das and Joseph [3].

A fairly lengthy argument along the same lines allows us to bound the path when $h > \frac{w}{2\tan(\alpha/2)}$ †.

Lemma 6. *If $\theta < \alpha$, and $h > \frac{w}{2\tan(\alpha/2)}$, then the length of the portion of the two-sided path from a to b that is between u_a and u_b is at most $\frac{2(\pi-\alpha)}{\alpha\sin(\alpha/4)}\big(|T_1| + |T_2|\big)$.*

2.3 The Final Spanning Ratio

We now have all the pieces to prove Theorem 1. From the above lemmas, we have that the maximum length of the path between u_a and u_b is at most:

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

Since $|T_U(a,b)| + |T_L(a,b)| \le 2|ab|\frac{2(\pi-\alpha)d}{\alpha\cdot\sin(\alpha/4)}$, we have that the path from a to b has length at most:

$$\begin{aligned} & \frac{2(\pi-\alpha)}{\alpha\sin\left(\frac{\alpha}{4}\right)} \cdot \Big(\big|T_U(a,b)\big| + \big|T_L(a,b)\big|\Big) \\ \le & \frac{2(\pi-\alpha)}{\alpha\sin\left(\frac{\alpha}{4}\right)} \cdot 2\left(\frac{2(\pi-\alpha)d|ab|}{\alpha\sin\left(\frac{\alpha}{4}\right)}\right) \\ = & \frac{8(\pi-\alpha)^2 d|ab|}{\alpha^2\sin^2\left(\frac{\alpha}{4}\right)} \end{aligned}$$

proving Theorem 1.

3 Triangulations that have the Diamond Property

In this section, we note that three constrained versions of classical triangulations have the visible α-diamond property, namely the constrained Delaunay triangulation, the constrained minimum weight triangulation and the constrained greedy triangulation. Since they are all triangulations, they satisfy the requirements for the d-good polygon property, for $d = 1$. Therefore, all three triangulations are constant spanners where the constant depends on α.

The first constrained triangulation considered is the constrained Delaunay triangulation (CDT), also called the Generalized Delaunay triangulation [11], and the Obstacle triangulation [2]. 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 both endpoints of the edge e [11] [2]. For such an edge e, we refer to $\mathcal{C}_e$ as its *visibly empty circle*. The existence of the visibly empty circle for each unconstrained edge implies that the edge has the visible $\frac{\pi}{4}$-diamond property.

† Full details of this proof are made available to the Program Committee in the version of this paper with appendix which can be found at http://cg.scs.carleton.ca/~jit/CDiamond.pdf

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

We note that techniques exploiting additional properties of the Delaunay triangulation have been used to reduce the spanning ratio from the one implied by the visible $\frac{\pi}{4}$-diamond property (see Dobkin *et al.* [5], Keil and Gutwin [9] for the unconstrained setting, and Karavelas [8], Bose and Keil [1] for the constrained setting). Currently, the best known spanning ratio for the Constrained Delaunay triangulation is $\frac{4\pi\sqrt{3}}{9}$ as shown by Bose and Keil [1]. It is conjectured that the spanning ratio for the Delaunay and Constrained Delaunay triangulation is $\pi/2$.

The second constrained triangulation that we consider is the constrained Greedy triangulation (CGT), which is a generalization of the standard Greedy triangulation [12]. The algorithm for computing such a triangulation is as follows: Sort the edges of $Vis(P, L)$ by length. First insert all the constrained edges to CGT. Next insert the unconstrained edges in sorted order into CGT as long as they do not introduce a crossing. In order to prove the result for the Constrained Greedy Triangulation (CGT), we make extensive use of the following.

Lemma 7. *Let x and y be points of P such that $xy \in Vis(P, L)$, but xy is not an edge of $CGT(P, L)$. Let e be the edge of $CGT(P, L)$ 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 $Vis(P, L)$ 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 CGT, which is in contradiction to the assumption that xy is not in the triangulation. □

Using this simple lemma, we can show that the Constrained Greedy Triangulation has the visible $\frac{\pi}{6}$-diamond property. The main approach to proving this theorem is by contradiction. If an edge xy of CGT does not have the visible $\frac{\pi}{6}$-diamond property, then both visible triangles adjacent to xy contain at least one point visible to both x and y. We use those points to show that the greedy process would have inserted a shorter edge intersecting xy, thereby contradicting that xy is part of the greedy triangulation. The analysis is similar in approach to the one presented in Das and Joseph [3], however, by carefully reviewing each of the 6 cases in their analysis, we are able increase the angle from $\pi/8$ to $\pi/6$[‡].

Theorem 4. *The Constrained Greedy Triangulation (CGT) has the visible $\frac{\pi}{6}$-diamond property.*

The size of the diamond in the above proof is an improvement over the original value of $\frac{\pi}{8}$ shown by Das and Joseph [3]. It is also an improvement over the values shown by Dickerson *et al.* [4] and Drysdale *et al.* [6]. Dickerson *et al.* [4], 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 Drysdale *et al.* [6] to include the tangents to the region. By basic trigonometry, it can be shown that the largest visible diamond inscribed in this region has $\alpha = \arctan\left(1/\sqrt{5}\right)$. Therefore, $\frac{\pi}{6}$ is currently the largest diamond for the Greedy and constrained Greedy triangulations.

[‡]Full details of this proof are made available to the Program Committee in the version of this paper with appendix which can be found at http://cg.scs.carleton.ca/∼jit/CDiamond.pdf

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.

An argument similar to the one for Theorem 4 shows that the constrained minimum weight triangulation has the visible $\pi/8$-diamond property.

Theorem 5. *The Constrained Minimum Weight Triangulation has the visible $\frac{\pi}{8}$-diamond property.*

References

[1] P. Bose and M. Keil, On the stretch factor of the constrained delaunay triangulation. In *Proceedings of the 3rd Symposium on Voronoi Diagrams in Science and Engineering*, pp. 25–31, 2006.

[2] L. P. Chew, There is a planar graph almost as good as the complete graph. In *Proc. 2nd Annu. ACM Sympos. Comput. Geom.*, pp. 169–177, 1986.

[3] G. Das and D. Joseph, Which triangulations approximate the complete graph. In *Proc. International Symposium on Optimal Algorithms*, vol. 401 of *Lecture Notes Comput. Sci.*, pp. 168–192, Springer-Verlag, 1989.

[4] 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.*, pp. 211–220, 1994.

[5] 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.*, pp. 20–26, 1987.

[6] R. L. Drysdale, G. Rote, and O. Aichholzer, A simple linear time greedy triangulation algorithm for uniformly distributed points. Tech. Rep. IIG-408, Institutes for Information Processing, Technische Universität Graz, 1995.

[7] D. Eppstein, Spanning trees and spanners. In J.-R. Sack and J. Urrutia, eds., *Handbook of Computational Geometry*, pp. 425–461, Elsevier Science, Amsterdam, 2000.

[8] M. I. Karavelas, *Proximity Structures for Moving Objects in Constrained and Unconstrained Environments*. Ph.D. thesis, Dept. Comput. Sci., Stanford University, 2001.

[9] J. M. Keil and C. A. Gutwin, Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, **7**:13–28, 1992.

[10] C. Knauer and J. Gudmundsson, Dilation and detour in geometric networks. In *Handbook of approximation algorithms and metaheuristics*, CRC Press, 2007.

[11] D. T. Lee, Proximity and reachability in the plane. Report R-831, Dept. Elect. Engrg., Univ. Illinois, Urbana, IL, 1978.

[12] 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.

[13] G. Narasimhan and M. Smid, *Geometric Spanner Networks.* Cambridge University Press, 2007.

[14] M. Smid, Closest-point problems in computational geometry. In J.-R. Sack and J. Urrutia, eds., *Handbook of Computational Geometry*, pp. 877–935, Elsevier Science, Amsterdam, 2000.