

On the Stretch Factor of Convex Delaunay Graphs

Prosenjit Bose¹, Paz Carmi¹, Sébastien Collette^{2*}, and Michiel Smid¹

¹ School of Computer Science**, Carleton University, Ottawa, Ontario, K1S 5B6, Canada.

² Computer Science Department, Université Libre de Bruxelles, CP212, Bvd du Triomphe, 1050 Brussels, Belgium.

Abstract. Let C be a compact and convex set in the plane that contains the origin in its interior, and let S be a finite set of points in the plane. The Delaunay graph $DG_C(S)$ of S is defined to be the dual of the Voronoi diagram of S with respect to the convex distance function defined by C . We prove that $DG_C(S)$ is a t -spanner for S , for some constant t that depends only on the shape of the set C . Thus, for any two points p and q in S , the graph $DG_C(S)$ contains a path between p and q whose Euclidean length is at most t times the Euclidean distance between p and q .

1 Introduction

Let S be a finite set of points in the plane and let G be a graph with vertex set S , in which each edge (p, q) has a weight equal to the Euclidean distance $|pq|$ between p and q . For a real number $t \geq 1$, we say that G is a t -spanner for S , if for any two points p and q of S , there exists a path in G between p and q whose Euclidean length is at most $t|pq|$. The smallest such t is called the *stretch factor* of G . The problem of constructing spanners has received much attention; see Narasimhan and Smid [12] for an extensive overview.

Spanners were introduced in computational geometry by Chew [3,4], who proved the following two results. First, the L_1 -Delaunay graph, i.e., the dual of the Voronoi diagram for the Manhattan metric, is a $\sqrt{10}$ -spanner. Second, the Delaunay graph based on the convex distance function defined by an equilateral triangle, is a 2-spanner. We remark that in both these results, the stretch factor is measured in the Euclidean metric. Chew also conjectured that the Delaunay graph based on the Euclidean metric, is a t -spanner, for some constant t . (If not all points of S are on a line, and if no four points of S are cocircular, then the Delaunay graph is the well-known Delaunay *triangulation*.) This conjecture was proved by Dobkin *et al.* [8], who showed that $t \leq \pi(1 + \sqrt{5})/2$. The analysis was improved by Keil and Gutwin [9], who showed that $t \leq \frac{4\pi\sqrt{3}}{9}$.

* Chargé de recherches du F.R.S.-FNRS.

** Supported by NSERC.

In this paper, we unify these results by showing that the Delaunay graph based on any convex distance function has bounded stretch factor.

Throughout this paper, we fix a compact and convex set C in the plane. We assume that the origin is in the interior of C . A *homothet* of C is obtained by scaling C with respect to the origin, followed by a translation. Thus, a homothet of C can be written as

$$x + \lambda C = \{x + \lambda z : z \in C\},$$

for some point x in the plane and some real number $\lambda \geq 0$. We call x the *center* of the homothet $x + \lambda C$.

For two points x and y in the plane, we define

$$d_C(x, y) := \min\{\lambda \geq 0 : y \in x + \lambda C\}.$$

If $x \neq y$, then this definition is equivalent to the following: Consider the translate $x + C$ and the ray emanating from x that contains y . Let y' be the (unique) intersection between this ray and the boundary of $x + C$. Then

$$d_C(x, y) = |xy|/|xy'|.$$

The function d_C is called the *convex distance function* associated with C . Clearly, we have $d_C(x, x) = 0$ and $d_C(x, y) > 0$ for all points x and y with $x \neq y$. Chew and Drysdale [5] showed that the triangle inequality $d_C(x, z) \leq d_C(x, y) + d_C(y, z)$ holds. In general, the function d_C is not symmetric, i.e., $d_C(x, y)$ is not necessarily equal to $d_C(y, x)$. If C is symmetric with respect to the origin, however, then d_C is symmetric.

Let S be a finite set of points in the plane. For each point p in S , we define

$$V'_C(p) := \{x \in \mathbb{R}^2 : \text{for all } q \in S, d_C(x, p) \leq d_C(x, q)\}.$$

If C is not strictly convex, then the set $V'_C(p)$ may consist of a closed region of positive area with an infinite ray attached to it. For example, in figure 1, the set $V'_C(a)$ consists of the set of all points that are on or to the left of the leftmost zig-zag line, together with the infinite horizontal ray that is at the same height as the point a . Also, the intersection of two regions $V'_C(p)$ and $V'_C(q)$, where p and q are distinct points of S , may have a positive area. As a result, the collection $V'_C(p)$, where p ranges over all points of S , does not necessarily give a subdivision of the plane in which the interior of each cell is associated with a unique point of S . In order to obtain such a subdivision, we follow the approach of Klein and Wood [10] (see also Ma [11]): first, infinite rays attached to regions of positive area are not considered to be part of the region. Second, a point x in \mathbb{R}^2 that is in the interior of more than one region $V'_C(p)$ is assigned to the region of the lexicographically smallest point p in S for which $x \in V'_C(p)$.

To formally define Voronoi cells, let \prec denote the lexicographical ordering on the set of all points in the plane. Let $p_1 \prec p_2 \prec \dots \prec p_n$ be the points of S ,

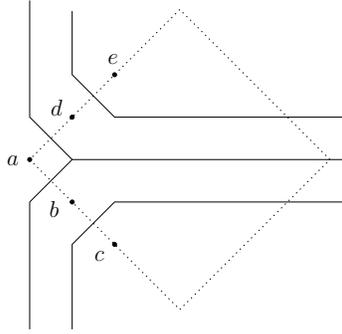


Fig. 1. The Voronoi diagram $VD_C(S)$ for the set $S = \{a, b, c, d, e\}$. The set C is the square as indicated by the dotted figure; the origin is at the center of C .

sorted according to this order. Then the Voronoi cells $V_C(p_i)$ of the points of S are defined as $V_C(p_1) := cl(int(V_C'(p_1)))$ and, for $1 < i \leq n$,

$$V_C(p_i) := cl \left(int \left(V_C'(p_i) \setminus \left(\bigcup_{j < i} V_C(p_j) \right) \right) \right),$$

where $cl(X)$ and $int(X)$ denote the closure and the interior of the set $X \subseteq \mathbb{R}^2$, respectively.

Thus, in figure 1, the Voronoi cell $V_C(a)$ consists only of the set of all points that are on or to the left of the leftmost zig-zag line; the infinite horizontal ray that is at the same height as the point a is not part of this cell.

The Voronoi diagram $VD_C(S)$ of S with respect to C is defined to be the collection of Voronoi cells $V_C(p)$, where p ranges over all points of S . An example is given in figure 1.

As for the Euclidean case, the Voronoi diagram $VD_C(S)$ induces Voronoi cells, Voronoi edges, and Voronoi vertices. Each point in the plane is either in the interior of a unique Voronoi cell, in the relative interior of a unique Voronoi edge, or a unique Voronoi vertex. Each Voronoi edge e belongs only to the two Voronoi cells that contain e on their boundaries. Observe that Voronoi cells are closed.

The Delaunay graph is defined to be the dual of the Voronoi diagram:

Definition 1. Let S be a finite set of points in the plane. The Delaunay graph $DG_C(S)$ of S with respect to C is defined to be the dual of the Voronoi diagram $VD_C(S)$. That is, the vertex set of $DG_C(S)$ is S and two distinct vertices p and q are connected by an edge in $DG_C(S)$ if and only if the Voronoi cells $V_C(p)$ and $V_C(q)$ share a Voronoi edge.

For example, the Delaunay graph $DG_C(S)$ for the point set in figure 1 consists of the five edges (a, b) , (a, d) , (b, c) , (b, d) , and (d, e) .

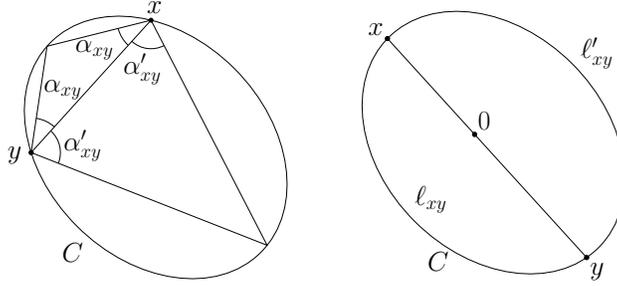


Fig. 2. The two parameters associated with C .

We consider the Delaunay graph $DG_C(S)$ to be a geometric graph, which means that each edge (p, q) is embedded as the closed line segment with endpoints p and q .

Before we can state the main result of this paper, we introduce two parameters whose values depend on the shape of the set C . Let x and y be two distinct points on the boundary ∂C of C . These points partition ∂C into two chains. For each of these chains, there is an isosceles triangle with base xy and whose third vertex is on the chain. Denote the base angles of these two triangles by α_{xy} and α'_{xy} ; see figure 2 (left). We define

$$\kappa_C := \min\{\max(\alpha_{xy}, \alpha'_{xy}) : x, y \in \partial C, x \neq y\}.$$

Consider again two distinct points x and y on ∂C , but now assume that x , y , and the origin are collinear. As before, x and y partition ∂C into two chains. Let ℓ_{xy} and ℓ'_{xy} denote the lengths of these chains; see figure 2 (right). We define

$$\kappa_{C,0} := \max\left\{\frac{\max(\ell_{xy}, \ell'_{xy})}{|xy|} : x, y \in \partial C, x \neq y, \text{ and } x, y, \text{ and } 0 \text{ are collinear}\right\}.$$

Clearly, the convex distance function d_C and, therefore, the Voronoi diagram $VD_C(S)$, depends on the location of the origin in the interior of C . Surprisingly, the Delaunay graph $DG_C(S)$ does not depend on this location; see Ma [11, Section 2.1.6]. We define

$$\kappa_C := \min\{\kappa_{C,0} : 0 \text{ is in the interior of } C\}.$$

In this paper, we will prove the following result:

Theorem 1. *Let C be a compact and convex set in the plane with a non-empty interior, and let S be a finite set of points in the plane. The stretch factor of the Delaunay graph $DG_C(S)$ is less than or equal to*

$$t_C := \begin{cases} 2\kappa_C \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right) & \text{if } DG_C(S) \text{ is a triangulation,} \\ 2\kappa_C^2 \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right) & \text{otherwise.} \end{cases}$$

Thus, for any two points p and q in S , the graph $DG_C(S)$ contains a path between p and q whose Euclidean length is at most t_C times the Euclidean distance between p and q .

We emphasize that we do *not* make any “general position” assumption; our proof of Theorem 1 is valid for *any* finite set of points in the plane.

Throughout the rest of this paper, we assume that the origin is chosen in the interior of C such that $\kappa_C = \kappa_{C,0}$.

The rest of this paper is organized as follows. In Section 2, we prove some basic properties of the Delaunay graph which are needed in the proof of Theorem 1. Section 3 contains a proof of Theorem 1. This proof is obtained by showing that the Delaunay graph satisfies the “diamond property” and a variant of the “good polygon property” of Das and Joseph [6]. The proof of the latter property is obtained by generalizing the analysis of Dobkin *et al.* [8] for the lengths of so-called one-sided paths.

Due to the lack of space, some proofs are omitted or are sketched.

2 Some properties of the Delaunay graph

Recall that in the Euclidean Delaunay graph, if two points p and q of S are connected by an edge, then there exists a disk having p and q on its boundary that does not contain any point of S in its interior. The next lemma generalizes this result to the Delaunay graph $DG_C(S)$.

Lemma 1. *Let p and q be two points of S and assume that (p, q) is an edge in the Delaunay graph $DG_C(S)$. Then, the following are true.*

1. *The line segment between p and q does not contain any point of $S \setminus \{p, q\}$.*
2. *For every point x in $V_C(p) \cap V_C(q)$, there exists a real $\lambda > 0$ such that*
 - (a) *the homothet $x + \lambda C$ contains p and q on its boundary, and*
 - (b) *the interior of $x + \lambda C$ does not contain any point of S .*

Due to space constraints, the proof is omitted.

As can be seen in figure 1, Voronoi cells are, in general, not convex. They are, however, star-shaped:

Lemma 2. *Let p be a point of S and let x be a point in the Voronoi cell $V_C(p)$. Then the line segment xp is completely contained in $V_C(p)$.*

Due to space constraints, the proof is omitted.

It is well known that the Euclidean Delaunay graph is a plane graph; see, for example, de Berg *et al.* [7, page 189]. The following lemma states that this is true for the Delaunay graph $DG_C(S)$ as well.

Lemma 3. *The Delaunay graph $DG_C(S)$ is a plane graph.*

Even though this fact seems to be well known, we have not been able to find a proof in the literature. The full proof is given in the full version of this paper.

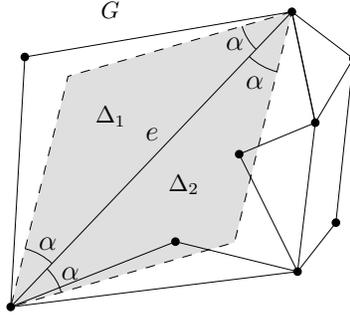


Fig. 3. The α -diamond property.

3 The stretch factor of Delaunay graphs

In this section, we will prove Theorem 1. First, we show that the Delaunay graph $DG_C(S)$ satisfies the diamond property and a variant of the good polygon property of Das and Joseph [6]. According to the results of Das and Joseph, this immediately implies that the stretch factor of $DG_C(S)$ is bounded. In fact, we will obtain an upper bound on the stretch factor which is better than the one that is implied by Das and Joseph's result.

3.1 The diamond property

Let G be a plane graph with vertex set S and let α be a real number with $0 < \alpha < \pi/2$. For any edge e of G , let Δ_1 and Δ_2 be the two isosceles triangles with base e and base angle α ; see figure 3. We say that e satisfies the α -diamond property, if at least one of the triangles Δ_1 and Δ_2 does not contain any point of S in its interior. The graph G is said to satisfy the α -diamond property, if every edge e of G satisfies this property.

Lemma 4. Consider the value α_C that was defined in Section 1. The Delaunay graph $DG_C(S)$ satisfies the α_C -diamond property.

Proof. Let (p, q) be an arbitrary edge of $DG_C(S)$ and let x be any point in the relative interior of $V_C(p) \cap V_C(q)$. By Lemma 1, there exists a real number $\lambda > 0$ such that p and q are on the boundary of the homothet $x + \lambda C$ and no point of S is in the interior of $x + \lambda C$. The points p and q partition $\partial(x + \lambda C)$ into two chains. For each of these chains, there is an isosceles triangle with base pq and whose third vertex is on the chain. We denote the base angles of these two triangles by β and γ ; see figure 4. We may assume without loss of generality that $\beta \geq \gamma$. Let a denote the third vertex of the triangle with base angle β . If we translate $x + \lambda C$ so that x coincides with the origin and scale the translated homothet by a factor of $1/\lambda$, then we obtain the set C . This translation and scaling does

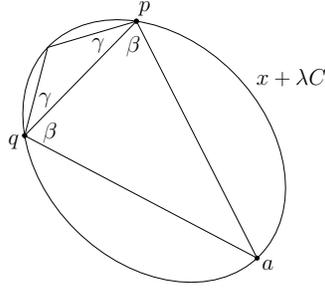


Fig. 4. Illustrating the proof of Lemma 4.

not change the angles β and γ . Thus, using the notation of Section 1 (see also figure 2), we have $\{\beta, \gamma\} = \{\alpha_{pq}, \alpha'_{pq}\}$. The definition of α_C then implies that

$$\alpha_C \leq \max(\alpha_{pq}, \alpha'_{pq}) = \beta.$$

Let Δ be the isosceles triangle with base pq and base angle α_C such that a and the third vertex of Δ are on the same side of pq . Then Δ is contained in the triangle with vertices p , q , and a . Since the latter triangle is contained in $x + \lambda C$, it does not contain any point of S in its interior. Thus, Δ does not contain any point of S in its interior. This proves that the edge (p, q) satisfies the α_C -diamond property. \square

3.2 The visible-pair spanner property

For a real number $\kappa \geq 1$, we say that the plane graph G satisfies the *strong visible-pair κ -spanner property*, if the following is true: For every face f of G , and for every two vertices p and q on the boundary of f , such that the open line segment joining p and q is completely in the interior of f , the graph G contains a path between p and q having length at most $\kappa|pq|$. If for every face f of G and for every two vertices p and q on the boundary of f , such that the line segment pq does not intersect the exterior of f , the graph G contains a path between p and q having length at most $\kappa|pq|$, then we say that G satisfies the *visible-pair κ -spanner property*. Observe that the former property implies the latter one. Also, observe that these properties are variants of the *κ -good polygon* property of Das and Joseph [6]: The κ -good polygon property requires that G contains a path between p and q that is along the boundary of f and whose length is at most $\kappa|pq|$; in the (strong) visible-pair spanner property, the path is not required to be along the boundary of f .

In this subsection, we will prove that the Delaunay graph $DG_C(S)$ satisfies the visible-pair κ_C -spanner property, where κ_C is as defined in Section 1. This claim will be proved by generalizing results of Dobkin *et al.* [8] on so-called one-sided paths.

Let p and q be two distinct points of S and assume that (p, q) is not an edge of the Delaunay graph $DG_C(S)$. Consider the Voronoi diagram $VD_C(S)$. We consider the sequence of points in S whose Voronoi cells are visited when the line segment pq is traversed from p to q . If pq does not contain any Voronoi vertex, then this sequence forms a path in $DG_C(S)$ between p and q . Since, in general, Voronoi cells are not convex, it may happen that this path contains duplicates. In order to avoid this, we define the sequence in the following way.

In the rest of this section, we will refer to the line through p and q as the X -axis, and we will say that p is to the *left* of q . This implies a *left-to-right* order on the X -axis, the notion of a point being *above* or *below* the X -axis, as well as the notions *horizontal* and *vertical*. (Thus, conceptually, we rotate and translate all points of S , the set C , the Voronoi diagram $VD_C(S)$, and the $DG_C(S)$, such that p and q are on a horizontal line and p is to the left of q . Observe that $VD_C(S)$ is still defined based on the lexicographical order of the points of S before this rotation and translation.) In the following, we consider the (horizontal) line segment pq . If this segment contains a Voronoi vertex, then we imagine moving pq vertically upwards by an infinitesimal amount. Thus, we may assume that pq does not contain any Voronoi vertex of the (rotated and translated) Voronoi diagram $VD_C(S)$.

The first point in the sequence is $p_0 := p$. We define $x_1 \in \mathbb{R}^2$ to be the point on the line segment pq such that $x_1 \in V_C(p_0)$ and x_1 is closest to q .

Let $i \geq 1$ and assume that the points p_0, p_1, \dots, p_{i-1} of S and the points x_1, \dots, x_i in \mathbb{R}^2 have already been defined, where x_i is the point on the line segment pq such that $x_i \in V_C(p_{i-1})$ and x_i is closest to q . If $p_{i-1} = q$, then the construction is completed. Otherwise, observe that x_i is in the relative interior of a Voronoi edge. We define p_i to be the point of $S \setminus \{p_{i-1}\}$ whose Voronoi cell contains x_i on its boundary, and define x_{i+1} to be the point on the line segment pq such that $x_{i+1} \in V_C(p_i)$ and x_{i+1} is closest to q .

Let $p = p_0, p_1, \dots, p_k = q$ be the sequence of points in S obtained in this way. By construction, these $k+1$ points are pairwise distinct and for each i with $1 \leq i \leq k$, the Voronoi cells $V_C(p_{i-1})$ and $V_C(p_i)$ share an edge. Therefore, by definition, (p_{i-1}, p_i) is an edge in $DG_C(S)$. Thus, $p = p_0, p_1, \dots, p_k = q$ defines a path in $DG_C(S)$ between p and q . We call this path the *direct path* between p and q . If all points p_1, p_2, \dots, p_{k-1} are strictly on one side of the line through p and q , then we say that the direct path is *one-sided*.

We will show in Lemma 6 that the length of a one-sided path is at most $\kappa_C |pq|$. The proof of this lemma uses a geometric property which we prove first.

Let C' be a homothet of C whose center is on the X -axis, and let x and y be two points on the boundary of C' that are on or above the X -axis. The points x and y partition the boundary of C' into two chains. One of these chains is completely on or above the X -axis; we denote this chain by $\text{arc}(x, y; C')$. The length of this chain is denoted by $|\text{arc}(x, y; C')|$.

For two points x and y on the X -axis, we write $x <_X y$ if x is strictly to the left of y , and we write $x \leq_X y$ if $x = y$ or $x <_X y$.

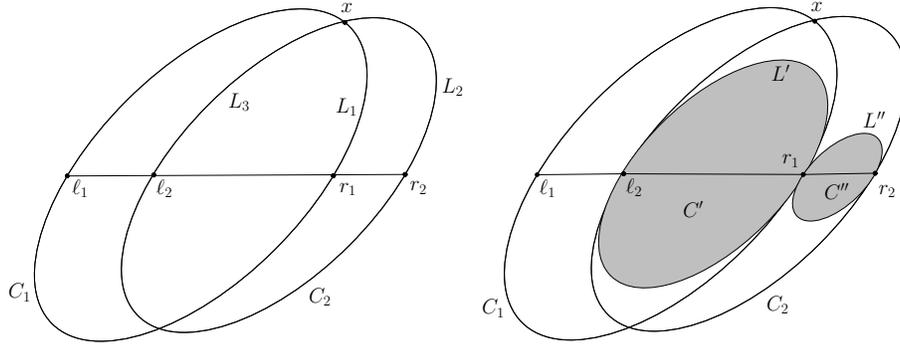


Fig. 5. Illustrating the proof of Lemma 5.

We now state the geometric property, which is illustrated in figure 5. Recall the value κ_C that was defined in Section 1.

Lemma 5. *Let $C_1 = y_1 + \lambda_1 C$ and $C_2 = y_2 + \lambda_2 C$ be two homothets of C whose centers y_1 and y_2 are on the X -axis. Assume that $\lambda_1 > 0$, $\lambda_2 > 0$, and $y_1 <_X y_2$. For $i = 1, 2$, let ℓ_i and r_i be the leftmost and rightmost points of C_i on the X -axis, respectively. Assume that $r_1 \leq_X r_2$ and $\ell_1 \leq_X \ell_2 <_X r_1$. Let x be a point that is on the boundaries of both C_1 and C_2 and on or above the X -axis. Let $L_1 = |\text{arc}(x, r_1; C_1)|$ and $L_2 = |\text{arc}(x, r_2; C_2)|$. Then*

$$L_2 \leq L_1 + \kappa_C |r_1 r_2|.$$

Proof. We define $L_3 = |\text{arc}(\ell_2, x; C_2)|$. Let C' be the homothet of C whose center is on the X -axis such that the intersection between C' and the X -axis is equal to the line segment $\ell_2 r_1$, and let $L' = |\text{arc}(\ell_2, r_1; C')|$; see Fig. 5. Observe that, for $\lambda := |\ell_2 r_1|/|\ell_2 r_2|$, C' is obtained from C_2 by a scaling by a factor of λ . Thus, since $|\text{arc}(\ell_2, r_2; C_2)| = L_2 + L_3$, we have $L' = \lambda(L_2 + L_3)$.

Let C'' be the homothet of C whose center is on the X -axis such that the intersection between C'' and the X -axis is equal to the line segment $r_1 r_2$, and let $L'' = |\text{arc}(r_1, r_2; C'')|$. Since C'' is obtained from C_2 by a scaling by a factor of $1 - \lambda$, we have $L'' = (1 - \lambda)(L_2 + L_3)$.

Thus, we have $L' + L'' = L_2 + L_3$. By convexity, we have $C' \subseteq C_1 \cap C_2$. Then it follows, again from convexity (see Benson [1, page 42]), that $L' \leq L_1 + L_3$. Thus, we have $L_2 + L_3 = L' + L'' \leq L_1 + L_3 + L''$, which implies that $L_2 \leq L_1 + L''$. Since, by the definition of κ_C , $L'' \leq \kappa_C |r_1 r_2|$, the proof is complete. \square

We are now ready to prove an upper bound on the length of a one-sided path.

Lemma 6. *If the direct path between p and q is one-sided, then its length is at most $\kappa_C |pq|$.*

Due to space constraints, we only give a proof sketch.

Proof (Sketch). As above, we assume that p and q are on the X -axis and that p is to the left of q . Consider the direct path $p = p_0, p_1, \dots, p_k = q$ in $DG_C(S)$ and the sequence x_1, x_2, \dots, x_k , as defined above. Since the direct path is one-sided, we may assume without loss of generality that the points p_1, p_2, \dots, p_{k-1} are strictly above the X -axis.

Recall that, for each i with $1 \leq i \leq k$, x_i is in the relative interior of $V_C(p_{i-1}) \cap V_C(p_i)$ and x_i is on the line segment pq . Therefore, by Lemma 1, if we define $\lambda_i := d_C(x_i, p_{i-1})$ (which is equal to $d_C(x_i, p_i)$), then the homothet $C_i := x_i + \lambda_i C$ contains p_{i-1} and p_i on its boundary and no point of S is in its interior.

Remark that the length of the upper boundary of the union of the C_i is an upper bound on the length of the one-sided path. If there are k edges, then the upper boundary is composed of portions of boundaries of k homothets. Now we proceed by induction on the number of edges (and thus of homothets):

For the base case, i.e., when there is one edge, the path is bounded by the upper boundary of a single homothet, and thus the fact that the length of the path between p and q is at most $\kappa_C |pq|$ follows from the definition of κ_C .

Assume it is true for all paths up to length k . By Lemma 5, a case analysis shows that when we add one more edge to the path, its length is not increasing by more than κ_C times the length of the projection of the added edge on the X -axis. This implies that the whole path has length less than $\kappa_C |pq|$. \square

We are now ready to prove that the Delaunay graph satisfies the visible-pair spanner property:

Lemma 7. *The Delaunay graph $DG_C(S)$ satisfies the visible-pair κ_C -spanner property.*

Proof. Recall from Lemma 3, that the graph $DG_C(S)$ is plane. It suffices to prove that $DG_C(S)$ satisfies the strong visible-pair κ_C -spanner property. Let f be a face of G and let p and q be two vertices on f such that the open line segment between p and q is contained in the interior of f . We have to show that there is a path in $DG_C(S)$ between p and q whose length is at most $\kappa_C |pq|$.

As before, we assume that p and q are on the X -axis and that p is to the left of q . Consider the direct path $p = p_0, p_1, \dots, p_k = q$ in $DG_C(S)$ and the sequence x_1, x_2, \dots, x_k , as defined in the beginning of this section. We will show that the direct path is one-sided. The lemma then follows from Lemma 6.

Since the open line segment between p and q is in the interior of f , none of the points p_1, \dots, p_{k-1} is on the closed line segment pq . Assume that for some i with $1 \leq i < k$, p_i is on the X -axis. Then p_i is either strictly to the left of p or strictly to the right of q . We may assume without loss of generality that p_i is strictly to the right of q . Consider the point x_i and the homothet $C_i = x_i + \lambda_i C$ as in the proof of Lemma 6. Since x_i is on pq and in the interior of C_i , and since p_i is on the boundary of C_i , it follows from convexity that q is in the interior of C_i , which is a contradiction. Thus we have shown that none of the points p_1, \dots, p_{k-1} is on the X -axis.

Assume that the direct path is not one-sided. Then there is an edge (p_{i-1}, p_i) on this path such that one of p_{i-1} and p_i is strictly below the X -axis and the other point is strictly above the X -axis. Let z be the intersection between $p_{i-1}p_i$ and the X -axis. By assumption, z is not on the open line segment joining p and q , and by Lemma 1, $z \neq p$ and $z \neq q$. Thus, z is either strictly to the left of p or strictly to the right of q . We may assume without loss of generality that z is strictly to the right of q . Consider again the point x_i and the homothet $C_i = x_i + \lambda_i C$ as in the proof of Lemma 6. This homothet contains the points x_i , p_{i-1} and p_i . Thus, by convexity, C_i contains the triangle with vertices x_i , p_{i-1} , and p_i . Since q is in the interior of this triangle, it follows that q is in the interior of C_i , which is a contradiction. \square

3.3 The proof of Theorem 1

Das and Joseph [6] have shown that any plane graph satisfying the diamond property and the good polygon property has a bounded stretch factor. The analysis of the stretch factor was slightly improved by Bose *et al.* [2]. A close inspection of the proof in [2] shows that the following holds: Let G be a geometric graph with the following four properties:

1. G is plane.
2. G satisfies the α -diamond property.
3. The stretch factor of any one-sided path in G is at most κ .
4. G satisfies the visible-pair κ' -spanner property.

Then, G is a t -spanner for

$$t = 2\kappa\kappa' \cdot \max\left(\frac{3}{\sin(\alpha/2)}, \kappa\right).$$

We have shown that the Delaunay graph $DG_C(S)$ satisfies all these properties: By Lemma 3, $DG_C(S)$ is plane. By Lemma 4, $DG_C(S)$ satisfies the α_C -diamond property. By Lemma 6, the stretch factor of any one-sided path in $DG_C(S)$ is at most κ_C . By Lemma 7, $DG_C(S)$ satisfies the visible-pair κ_C -spanner property. If $DG_C(S)$ is a triangulation, then obviously, $DG_C(S)$ satisfies the visible-pair 1-spanner property. Therefore, we have completed the proof of Theorem 1.

4 Concluding remarks

We have considered the Delaunay graph $DG_C(S)$, where C is a compact and convex set with a non-empty interior and S is a finite set of points in the plane. We have shown that the (Euclidean) stretch factor of $DG_C(S)$ is bounded from above by a function of two parameters α_C and κ_C that are determined only by the shape of C . Roughly speaking, these two parameters give a measure of the “fatness” of the set C .

Our analysis provides the first generic bound valid for any compact and convex set C . In all previous works, only special examples of such sets C were

considered. Furthermore, our approach does not make any “general position” assumption about the point set S .

Note that for the Euclidean Delaunay triangulation, we derive an upper bound on the stretch factor of ≈ 24.6 , which is worse than the currently best known upper bound ($4\pi\sqrt{3}/9 \approx 2.42$ as proved by Keil and Gutwin [9]). We leave open the problem of improving our upper bound. In particular, is it possible to generalize the techniques of Dobkin *et al.* [8] and Keil and Gutwin [9]?

References

1. R. V. Benson. *Euclidean Geometry and Convexity*. McGraw-Hill, New York, NY, 1966.
2. P. Bose, A. Lee, and M. Smid. On generalized diamond spanners. In *Proceedings of the 10th Workshop on Algorithms and Data Structures*, volume 4619 of *Lecture Notes in Computer Science*, pages 325–336, Berlin, 2007. Springer-Verlag.
3. L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the 2nd ACM Symposium on Computational Geometry*, pages 169–177, 1986.
4. L. P. Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39:205–219, 1989.
5. L. P. Chew and R. L. Drysdale. Voronoi diagrams based on convex distance functions. In *Proceedings of the 1st ACM Symposium on Computational Geometry*, pages 235–244, 1985.
6. G. Das and D. Joseph. Which triangulations approximate the complete graph? In *Proceedings of the International Symposium on Optimal Algorithms*, volume 401 of *Lecture Notes in Computer Science*, pages 168–192, Berlin, 1989. Springer-Verlag.
7. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 2nd edition, 2000.
8. D. P. Dobkin, S. J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5:399–407, 1990.
9. J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7:13–28, 1992.
10. R. Klein and D. Wood. Voronoi diagrams based on general metrics in the plane. In *Proceedings of the 5th Symposium on Theoretical Aspects of Computer Science*, volume 294 of *Lecture Notes in Computer Science*, pages 281–291, Berlin, 1988. Springer-Verlag.
11. L. Ma. *Bisectors and Voronoi Diagrams for Convex Distance Functions*. Ph.D. thesis, Department of Computer Science, FernUniversität Hagen, Germany, 2000.
12. G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, Cambridge, UK, 2007.