



Empty-ellipse graphs

Olivier Devillers, Jeff Erickson, Xavier Goaoc

► To cite this version:

Olivier Devillers, Jeff Erickson, Xavier Goaoc. Empty-ellipse graphs. 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08), 2008, San Francisco, United States. pp.1249–1256. inria-00176204

HAL Id: inria-00176204

<https://inria.hal.science/inria-00176204>

Submitted on 2 Oct 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Empty-Ellipse Graphs

Olivier Devillers*

Jeff Erickson[†]

Xavier Goaoc[‡]

Submitted to SODA 2008: July 5, 2007

Abstract

We define and study a geometric graph over points in the plane that captures the local behavior of Delaunay triangulations of points on smooth surfaces in \mathbb{R}^3 . Two points in a planar point set P are neighbors in the *empty-ellipse graph* if they lie on an axis-aligned ellipse with no point of P in its interior. The empty-ellipse graph can be a clique in the worst case, but it is usually much less dense. Specifically, the empty-ellipse graph of n points has complexity $\Theta(\Delta n)$ in the worst case, where Δ is the ratio between the largest and smallest pairwise distances. For points generated uniformly at random in a rectangle, the empty-ellipse graph has expected complexity $\Theta(n \log n)$. As an application of our proof techniques, we show that the Delaunay triangulation of n random points on a circular cylinder has expected complexity $\Theta(n \log n)$.

*INRIA, Sophia-Antipolis, France; <http://www-sop.inria.fr/geometrica/team/Olivier.Devillers>

[†]University of Illinois at Urbana-Champaign; <http://www.cs.uiuc.edu/~jeffe/>. Partially supported by NSF CAREER award CCR-0093348 and NSF grant DMS-0528086.

[‡]LORIA, INRIA-Lorraine, Nancy, France; <http://www.loria.fr/~goaoc/>

1 Introduction

Delaunay triangulations and Voronoï diagrams are widely used in application areas such as mesh generation [3], surface reconstruction [6] or molecular modeling [10]. The size of the Delaunay triangulation of points in \mathbb{R}^3 is quadratic in the worst-case and near-linear in most practical situation. To understand this discrepancy, the complexity of the Delaunay graph has been studied in many settings, including point sets generated by homogeneous Poisson processes over \mathbb{R}^d [18], point sets generated uniformly at random in a hypercube [9], well-spaced point sets [22, 23], point sets in \mathbb{R}^3 with bounded spread [11, 12], and several types of discrete samples of polyhedral and smooth surfaces in \mathbb{R}^3 [2, 11, 12, 13, 14] or \mathbb{R}^d [1].

The case where the point set P samples a 2-dimensional surface $\Sigma \subset \mathbb{R}^3$ is of particular interest as the size of the Delaunay triangulation of such point sets governs the complexity of typical surface reconstruction algorithms [6]. If P consists of n points and Σ is a smooth surface, this complexity is $O(n \log^3 n)$ [2] under the condition that the surface be generic, i.e. that it meets any medial ball in a constant number of points. Generic surfaces cannot contain, for instance, a portion of sphere or cylinder. Whether this bound generalizes to non-generic surfaces is a question both of theoretical interest and practical relevance: shapes such as portions of cylinders or of spheres are widely used in solid modelling. In this paper, we answer this question for the case of cylinders.

Results. We give a tight bound on the size of the Delaunay graph of n points uniformly distributed on a cylinder of constant height. These bounds are obtained by studying *empty-ellipse graphs*. Let P be a set of points in the plane. We say that two points in P are *empty-ellipse neighbors* if they lie on an axis-aligned ellipse¹ with no point of P in its interior. The resulting *empty-ellipse graph* is a supergraph of the Delaunay graph, in which two points are neighbors if and only if they lie on an empty circle. Specifically, we prove:

- The empty-ellipse graph of any n -point set with spread Δ has $\Theta(n\Delta)$ edges in the worst case, for any Δ between $\Omega(\sqrt{n})$ and $O(n)$ (Section 2). This bound matches the conjectured worst-case complexity of three-dimensional Delaunay graphs with spread in the same range [12]. The lower-bound examples for empty-ellipse graphs closely resemble lower-bound examples for three-dimensional Delaunay graphs [11].
- The empty-ellipse graph of a random set of points in any axis-aligned rectangle has $\Theta(n \log n)$ edges in expectation (Section 3). To prove the matching bounds, we sandwich the empty-ellipse graph between two different *empty-region graphs* [15, 16, 4, 5] that yield more easily to probabilistic analysis: the *empty-right-triangle* graph and the *empty-doubled-box* graph.
- The size of the Delaunay graph of n points uniformly distributed on a cylinder of constant height is $\Theta(n \log n)$ in expectation (Section 4). This significantly improves an earlier high-probability upper bound $O(n^{3/2} \log^{1/2} n)$, which follows from a much more complex analysis in terms of spread [12]. This result also shows that the worst-case bound $\Theta(n^{3/2})$ for the Delaunay complexity of uniform ε -samples of smooth surfaces, achieved by points on the cylinder [11, 12], is pathological.

In the appendix, we present several basic properties of empty-ellipse graphs, as well as an output-sensitive algorithm for their construction.

¹The restriction to axis-aligned ellipses is crucial; in any set of points in general position, *every* pair of points lies on an empty ellipse that approximates a line segment. We will omit the adjective ‘axis-aligned’ almost everywhere; all ellipses in this paper are axis-aligned.

2 Spread-Sensitive Bounds

The *spread* of a set of points is the diameter of the set divided by the distance between the closest pair of points. A set of n points in the plane is *dense* if its spread is $O(\sqrt{n})$. The second author recently proved that the Delaunay triangulation of any set of n points in \mathbb{R}^3 with spread Δ has complexity $O(\Delta^3)$ [12]; this bound is tight in the worst case for all $\Delta = O(\sqrt{n})$ [11] and improves the trivial $O(n^2)$ worst-case bound when $\Delta = o(n^{2/3})$. He also conjectured an upper bound of $O(\Delta n)$, which would be tight in the worst case for all Δ between \sqrt{n} and n . (For $\Delta = \Omega(n)$, the trivial $O(n^2)$ bound is the best possible.) The results in this section match these conjectured complexity bounds.

In this section, we prove that the empty-ellipse graph of n points in the plane with spread Δ has $O(\Delta n)$ edges and maximum degree $O(\Delta \log \Delta)$; we construct examples to show that both bounds are tight in the worst case. We derive both upper bounds by considering the related *empty-right-triangle graph*. Two points $p, q \in P$ are neighbors in this graph if and only if the segment \overline{pq} is the hypotenuse of an axis-aligned right triangle with no point of P in its interior. Any chord through an axis-aligned ellipse (in fact, through any axis-symmetric convex body) is the hypotenuse of an axis-aligned right triangle that lies entirely within that ellipse; see Figure 1. It follows that the empty-ellipse graph of any point set P is a subgraph of the empty-right-triangle graph of P .

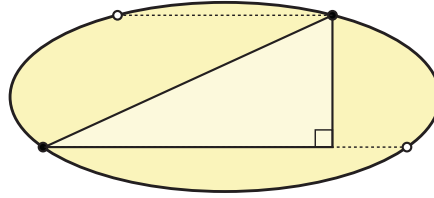


Figure 1. Any chord through an ellipse is the hypotenuse of an interior right triangle.

Any empty right triangle graph can be decomposed into four subgraphs, defined by restricting the orientation of the empty triangle. Two points (a, c) and (b, d) in P , where $a < b$ and $c < d$, are neighbors in the *southeast triangle graph* of P if the triangle with vertices (a, c) , (b, c) , and (b, d) has no point of P in its interior. More intuitively, every edge in the southeast triangle graph is the hypotenuse an empty axis-aligned right triangle whose right angle points southeast. Northeast, northwest, and southwest triangle graphs are defined analogously.

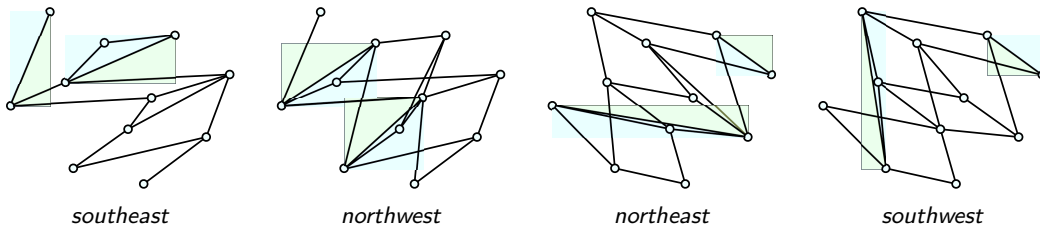


Figure 2. Four empty right triangle graphs for the same point set.

Lemma 2.1. *The southeast triangle graph of any set of n points in the plane with spread Δ has $O(\Delta n)$ edges.*

Proof: Fix a set P of n points in the plane with spread Δ . Without loss of generality, we assume that the closest pair of points in P is at unit distance, so that spread is synonymous with diameter.

Let p be the point in P furthest to the northwest, that is, maximizing the inner product $p \cdot (-1, 1)$. We claim that p has $O(\Delta)$ southeast triangle neighbors; this claim implies the lemma by induction.

Without loss of generality, assume that p is the origin, so that every other point in P lies on or below the line $y = x$. The neighbors of p lie in two right isosceles triangles: one with vertices $(0, 0)$, $(\Delta, 0)$, (Δ, Δ) , and the other with vertices $(0, 0)$, $(0, -\Delta)$, $(-\Delta, -\Delta)$. Let $(x_1, y_1), (x_2, y_2), \dots, (x_r, y_r)$ denote the neighbors of p in the first quadrant, sorted from left to right ($0 < x_1 \leq x_2 \leq \dots \leq x_r \leq \Delta$), or equivalently, by decreasing slope ($1 \geq y_1/x_1 \geq y_2/x_2 \geq \dots \geq y_r/x_r \geq 0$).

For all i , either $x_i > x_{i-1} + 1/3$ or $y_i/x_i < y_{i-1}/x_{i-1} + 1/3\Delta$. Otherwise, the point (x_i, y_i) would lie inside the convex quadrilateral defined by the inequalities

$$x_{i-1} \leq x_i \leq x_{i-1} + \frac{1}{3} \quad \text{and} \quad \frac{y_{i-1}}{x_{i-1}} \geq \frac{y_i}{x_i} \geq \frac{y_{i-1}}{x_{i-1}} + \frac{1}{3\Delta}, \quad (1)$$

which lies entirely within the unit circle centered at (x_i, y_i) , violating our assumption that no pair of points in P has distance less than 1. See Figure 3. Thus, p has at most 6Δ neighbors in the first quadrant, and thus by a symmetric argument at most 12Δ neighbors overall. \square

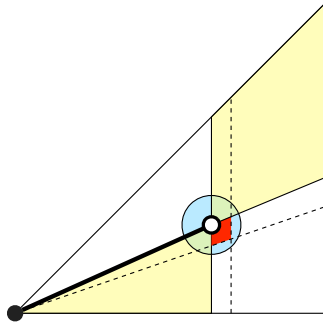


Figure 3. Proof of Lemma 2.1. No other neighbor lies in the shaded region.

Theorem 2.2. *The empty-ellipse graph of any set of n points in the plane with spread Δ has $O(\Delta n)$ edges.*

The following construction shows that this bound is tight in the worst case.

Theorem 2.3. *For any integer n and any real Δ such that $\sqrt{n} \ll \Delta \ll n$, there is a set of n points in general position in the plane whose spread is at most Δ and whose empty-ellipse graph has $\Omega(n\Delta)$ edges.*

Proof: First consider the case $\Delta = \Theta(\sqrt{n})$. Without loss of generality, assume that $n = m^2$ for some integer m . Let P be the tilted $m \times m$ square grid $\{q_{i,j} \mid 1 \leq i, j \leq m\}$, where $q_{i,j} = (i + mj, mi - j \bmod (n - 1))$. This set is clearly dense. Each point $q_{i,j}$ is an empty-ellipse neighbor of the following $2m + 2$ points: $q_{i+1,j'}$ for all $j' \geq j$, $q_{i-1,j'}$ for all $j' \leq j$, $q_{i',j-1}$ for all $i' \geq i$, and $q_{i',j+1}$ for all $i' \leq i$. Thus, the empty-ellipse graph of P has exactly $m^3 + m^2 = n^{3/2} + n$ edges. See Figure 4.

More generally, assume without loss of generality that both Δ and n/Δ are integers. Let P be the vertices of a $\Delta \times n/\Delta$ grid of squares, rotated very slightly away from the coordinate axes. P has spread $O(\Delta)$, because $\Delta \geq n/\Delta$. If the rotation is sufficiently small, each point in this grid

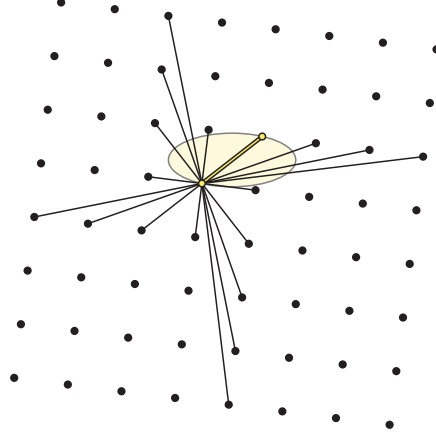


Figure 4. A dense set of points whose empty-ellipse graph has complexity $\Omega(n^{3/2})$. The $\Omega(\sqrt{n})$ empty-ellipse neighbors of one point are shown, along with the empty ellipse for one edge.

has exactly $\Delta + n/\Delta + 2 = \Omega(\Delta)$ empty-ellipse neighbors. Thus, the empty-ellipse graph of P has exactly $\Delta n/2 + n^2/2\Delta + n = \Omega(\Delta n)$ edges.

For each pair p, q of empty-ellipse neighbors in P , there is an empty ellipse that passes through p and q but no other points in P . Thus, we can perturb the set P into general position without removing any empty-ellipse edges. \square

2.1 Maximum Degree

Lemma 2.4. *The maximum degree in any southeast triangle graph is $O(\Delta \log \Delta)$, where Δ is the spread of the vertices.*

Proof: Let P be a set of n points in the plane with closest pair distance 1 and diameter Δ . Fix a point $p \in P$; without loss of generality, assume p is the origin. The neighbors of p in the southeast triangle graph of P lie entirely in the first and third quadrants. We explicitly bound the number of neighbors in the first quadrant; a symmetric argument bounds the third-quadrant neighbors.

The first-quadrant neighbors of p lie in an axis-aligned square of width Δ . Let $\delta = \lceil \lg \Delta \rceil$. We divide this square into $\delta + 1$ triangular zones $Z_0, Z_1, \dots, Z_\delta$ by the lines $y = 2^i x$ for all integers $0 \leq i \leq \delta - 1$. Z_0 is a right isosceles triangle pointing southeast; Z_δ is a thin right triangle pointing northwest; each of the remaining zones is a triangle with one horizontal edge. We claim that each zone contains $O(\Delta)$ neighbors.

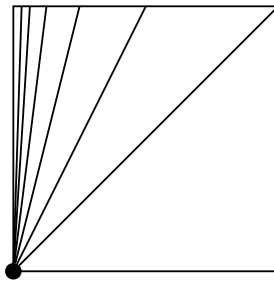


Figure 5. $\lg \Delta$ zones, each containing $O(\Delta)$ southeast triangle neighbors.

The proof of Lemma 2.1 already implies that there are $O(\Delta)$ neighbors in Z_0 . Zone Z_δ has area $O(\Delta)$ and thus contains at most $O(\Delta)$ points. For any intermediate zone Z_i , we argue as in the

proof of Lemma 2.1 except that the constants defining the inequalities (1) have to be scaled to keep the quadrilateral inside the unit disk. Let $(x_1, y_1), (x_2, y_2), \dots, (x_r, y_r)$ be the neighbors in Z_i , sorted from left to right. For each j , we have $0 \leq x_{j-1} \leq x_j \leq \Delta/2^{i-1}$ and $2^{i-1} \leq y_j/x_j \leq y_{j-1}/x_{j-1} \leq 2^i$. Because no pair of points is closer than unit distance, either

$$x_j > x_{j-1} + \frac{1}{3 \cdot 2^{i-1}} \quad \text{or} \quad \frac{y_j}{x_j} < \frac{y_{j-1}}{x_{j-1}} + \frac{2^{i-1}}{3\Delta}.$$

It follows that each zone Z_i contains at most $O(\Delta)$ neighbors, as claimed. \square

Theorem 2.5. *The maximum degree in any empty-ellipse graph is $O(\Delta \log \Delta)$, where Δ is the spread of the vertices.*

The following construction shows that this upper bound cannot be improved.

Theorem 2.6. *For any integer n and any real Δ such that $\sqrt{n} \ll \Delta \ll n/\log n$, there is a set of n points whose spread is at most Δ and whose empty-ellipse graph has maximum degree $\Omega(\Delta \log \Delta)$.*

Proof: We describe a set P of $\Theta(\Delta \log \Delta)$ points in the unit square $[0, 1]^2$, with spread $\Theta(\Delta)$, such that every point in P is a neighbor of the origin $o = (0, 0)$ in the empty-ellipse graph of $P \cup \{o\}$.

We define a template set P_0 , consisting of Δ points evenly distributed on the line segment σ_0 with endpoints $(1, 1)$ and $(1/2, 7/9)$. Each point $p \in \sigma_0$ lies on an axis-aligned ellipse that is both tangent to σ_0 at p and tangent to the y -axis at the origin; moreover, each of these ellipses lies entirely to the left of the line $y = 10$. In particular, the ellipse $(x - 5)^2/25 + 9y^2/25 = 1$ is tangent to σ_0 at $(1, 1)$ and tangent to the y -axis at $(0, 0)$. See Figure 6.

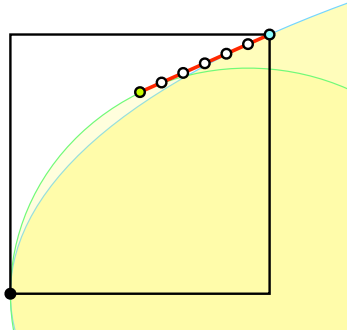


Figure 6. The template set P_0 .

For each integer i , let P_i denote the image of the template set P_0 under the map $(x, y) \mapsto (x/20^i, y)$. Each point in P_i lies on an empty ellipse that is tangent to the y -axis at the origin and is entirely disjoint from P_j for any $j \neq i$. Finally, let $P = \bigcup_{i=0}^{\ell} P_i$, where $\ell = \lfloor \log_{20} \Delta \rfloor$. This set contains $\Delta \log_{20} \Delta$ points, and the spread of P is at most $9\Delta/2$.

If n is larger than $\Delta \log_{20} \Delta$, we can add to P a dense set of $n - \Delta \log_{20} \Delta$ points in the square $[-1, 0] \times [0, 1]$. Because each point in P lies on an empty ellipse tangent to the y -axis, adding the extra points can only increase the number of empty-ellipse neighbors of o . Adding these points increases the spread of P to at most 9Δ .

Finally, perturbing P into general position can only add new empty-ellipse edges. \square

3 Random Points

We now prove that the empty-ellipse graph of a random point set is usually sparse. Specifically, if n points are drawn independently and uniformly at random from an axis-aligned rectangle, the expected number of empty-ellipse edges is $\Theta(n \log n)$. Unlike the bounds in the previous section, our bounds on the expected complexity of the graph follow directly from bounds on the expected degree of a single point.

We prove our results by sandwiching the empty-ellipse graph between two other empty-shape graphs. As in the previous section, we prove the upper bound in Section 3.1 by considering the *empty-right-triangle graph*. To prove the matching lower bound in Section 3.2, we analyze the *empty-doubled-box graph* of P , which contains an edge between two points $p, q \in P$ if scaling the axis-aligned bounding box of $\{p, q\}$ by a factor of $\sqrt{2}$ in each dimension yields a rectangle that contains no points in $P \setminus \{p, q\}$. The doubled bounding box of p and q contains a unique axis-aligned ellipse through p and q ; thus, the empty-doubled-box graph is a subgraph of the empty-ellipse graph.

A straightforward balls-and-bins argument [17] implies that a random set of n points in the square has expected spread $\Theta(n)$. Thus, the results in the previous section imply only trivial upper bounds in this randomized setting.

3.1 Upper Bound: Right Triangles

Lemma 3.1. *Let P be a set of n points chosen independently and uniformly at random in $[0, 1]^2$. In the southeast triangle graph of $P \cup \{o\}$, the expected number of neighbors of o is less than $2H_n$.*

Proof: With probability 1, no two points in P lie on a common horizontal line, vertical line, or line through the origin. Let p_i denote the point in P with the i th smallest x -coordinate. If j of the $i - 1$ points to the left of p_i have smaller y -coordinates than p_i , then p_i is a southeast neighbor of o with probability 2^{-j} . Because each value of j between 0 and $i - 1$ is equally likely, the expected number of southeast neighbors of o is precisely

$$\sum_{i=1}^n \sum_{j=0}^{i-1} \frac{2^{-j}}{i} = \sum_{i=1}^n \frac{2 - 2^{1-i}}{i} = 2H_n - 2 \sum_{i=1}^n \frac{1}{i2^i} < 2H_n. \quad \square$$

Corollary 3.2. *Let P be a set of n points chosen independently and uniformly at random in $[0, 1]^2$. The expected number of edges in the southeast triangle graph of P is less than $2nH_n$.*

Corollary 3.3. *Let P be a set of n points chosen independently and uniformly at random in $[0, 1]^2$. The expected number of edges in the empty-ellipse graph of P is less than $8nH_n$.*

3.2 Lower Bound: Doubled Boxes

To prove a lower bound on the expected complexity of the empty-ellipse graph, we introduce another related geometric graph. For any two points p and q in the plane, let $\square(p, q)$ be the axis-aligned box with p and q at opposite corners. Let $\boxplus(p, q)$ be the result of uniformly scaling $\square(p, q)$ by a factor of $\sqrt{2}$ about its center, so that the area of $\boxplus(p, q)$ is twice the area of $\square(p, q)$. Finally, two points $p, q \in P$ are neighbors in the *empty-doubled-box graph* of P if and only if $P \cap \boxplus(p, q) = \{p, q\}$. The doubled box graph of a point set may be empty in the worst (best?) case, but as we prove below, the graph is usually larger.

Let P be a set of points in the square $[-1, 1]^2$, and let o denote the origin $(0, 0)$.

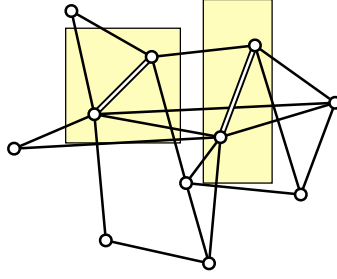


Figure 7. The empty-doubled-box graph for the same points as Figure 2.

Lemma 3.4. *Let P be a set of n points chosen independently and uniformly at random in $[-1, 1]$. In the doubled box graph of $P \cup \{o\}$, the expected number of neighbors of o is greater than $2H_n - 2 \ln 2$.*

Proof: For any point $p \in P$, let $\square(o, p)$ denote the axis-aligned rectangle with the same height and width as $\square(o, p)$ but centered at o . Call a point p *special* if $P \cap \square(o, p) = \emptyset$. We easily observe that the probability that p is special is at most the probability that p is a doubled-box neighbor of o . Thus, to prove the lemma, it suffices to bound the expected number of special points in P .

Because both the distribution of points and the boxes $\square(o, p)$ are symmetric about the origin, the expected number of special points is unchanged if we replace each point $p = (x, y) \in P$ with the point $(|x|, |y|)$, or equivalently, if we choose the points P uniformly at random from the upper right square $[0, 1]^2$. The area of $\square(o, p) \cap [0, 1]^2$ is exactly half the area of $\square(o, p)$. Thus, the proof of Lemma 3.1 implies that the expected number of special points is

$$2H_n - 2 \sum_{i=1}^n \frac{1}{i2^i} > 2H_n - 2 \sum_{i=1}^{\infty} \frac{1}{i2^i} = 2H_n - 2 \ln 2.$$

The last equality follows from the Taylor-Maclaurin series $\ln(1-x) = \sum_{i=1}^{\infty} \frac{x^i}{i}$ at $x = 1/2$. \square

Theorem 3.5. *Let P be a set of n points chosen independently and uniformly at random in $[-1, 1]$. The expected number of edges in the empty-doubled-box graph of P is $\Omega(n \log n)$.*

Proof: Let p be a point in the central unit square $[-1/2, 1/2]^2$. With very high probability², there are at least $n/5$ points inside the unit square centered at p . Thus, the previous lemma implies that the expected number of neighbors of p within this unit square is at least $2 \ln n - O(1)$. Similarly, with very high probability, there are at least $n/5$ points in $P \cap [-1/2, 1/2]^2$, so the expected number of edges in the doubled-box graph of P is at least $0.4 n \ln n - O(n)$. \square

Corollary 3.6. *Let P be a set of n points chosen independently and uniformly at random in $[-1, 1]$. The expected number of edges in the empty-ellipse graph of P is $\Omega(n \log n)$.*

4 Delaunay Complexity of Random Points on the Cylinder

In this section, we prove that the Delaunay triangulation of n random points on a right circular cylinder has expected complexity $\Theta(n \log n)$. Without loss of generality, we will assume that our cylinder C is aligned with the z -axis and has height and radius 1:

$$C = \{(\cos \theta, \sin \theta, z) \mid -\pi \leq \theta \leq \pi \text{ and } 0 \leq z \leq 1\}.$$

²that is, with probability $1 - c^{-n}$ for some constant $c > 1$

In two earlier papers [11, 12], the second author proved that any uniform sample of the cylinder has Delaunay complexity $\Theta(n^{3/2})$ in the worst case. The lower bound construction closely resembles the tilted grid used in the proof of Theorem 2.3. Specifically, let $n = m^2 + 1$, and let $P = \{p_1, p_2, \dots, p_n\}$, where $p_i = (\cos(2\pi i/m), \sin(2\pi i/m), i/n)$. This set contains the vertices of a slightly tilted grid wrapped around the cylinder. Every point in P is a Delaunay neighbor of every point in the adjacent rows and columns of the grid, so the complexity of the the Delaunay triangulation of P is $\Omega(n^{3/2})$.

As in the previous section, we bound the complexity of the Delaunay graph by relating it to two other empty-shape graphs. Both of the empty shapes become axis-aligned right triangles when the cylinder is unrolled into the plane. In a *standard* right triangle, the hypotenuse is a geodesic between p and q that wraps less than a full turn around the cylinder. In a *special* right triangle, one of the vertices, say p , is an endpoint of the hypotenuse, and the opposite edge lies on the horizontal line through q and has length exactly $\pi/2$. A generic pair of points on the cylinder determines four standard triangles and four special triangles, as shown in Figure 8.

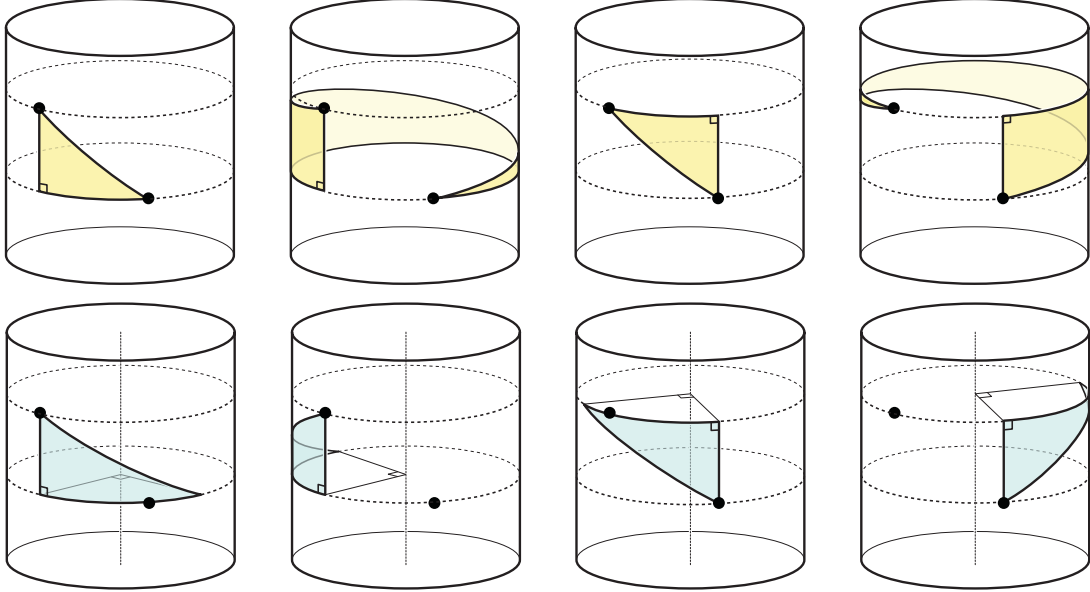


Figure 8. Standard (top row) and special (bottom row) right triangles determined by two points on a cylinder.

The *empty-standard-triangle graph* of a point set P on the cylinder has an edge between two points $p, q \in P$ if and only if p and q determine a standard right triangle with no point of P in its interior. The *empty-special-triangle graph* of P is defined analogously.

Lemma 4.1. *Let P be a set of n points chosen independently and uniformly at random on a right circular cylinder C . The expected number of edges in the empty-standard-triangle graph of P is $O(n \log n)$.*

Proof: This follows immediately from our earlier analysis of empty-right-triangle graphs of random points in the plane. \square

Lemma 4.2. *Let P be a set of n points chosen independently and uniformly at random on a right circular cylinder C . The expected number of edges in the empty-special-triangle graph of P is $O(n)$.*

Proof: Index the points in P as p_1, p_2, \dots, p_n in z -coordinate order. For any indices $i < j$, let $H(i, j)$ be the closed belt on C bounded by horizontal planes through p_i and p_j . The area of any special triangle determined by p_i and p_j is exactly $1/8$ the area of $H(i, j)$. Thus, the probability that p_i and p_j determine an *empty* special triangle is at most $4 \cdot (7/8)^{j-i-1}$. It follows that the expected number of edges in the empty-special-triangle graph of P is at most

$$\sum_{i=1}^n \sum_{j=i+1}^n 4 \cdot (7/8)^{j-i-1} < 4n \sum_{k \geq 0} (7/8)^k < 32n. \quad \square$$

Lemma 4.3. *Let C be a right circular cylinder with radius 1, let p and q be points on C , and let B be a generic closed ball with p and q on its boundary. The region $B \cap C$ contains either a standard right triangle determined by p and q , or a special right triangle determined by p and q .*

Proof: We parameterize C with the map $(\theta, z) \mapsto (\cos \theta, \sin \theta, z)$. For any point p or region R on the cylinder, let \bar{p} or \bar{R} denote its preimage in the parameter space $(-\pi, \pi] \times \mathbb{R}$.

To simplify our notation, let $R = B \cap C$. Without loss of generality, assume that B has radius r and center $(a, 0, 0)$ for some $a > 0$. The region R is the set of points satisfying the inequality

$$z^2 \leq 2a \cos \theta + r^2 - a^2 - 1.$$

The shape of this intersection region depends on the relative values of r and a . Only two cases are relevant for our proof. (In the remaining cases, either B and C meet tangentially or R is empty.)

1. If $|a - 1| < r < a + 1$, then R is homeomorphic to a disk. In this case the unfolded region \bar{R} is a convex, axis-symmetric oval centered at the origin.
2. If $r > a + 1$, then R is homeomorphic to an annulus. In this case, the unfolded region \bar{R} is an axis-symmetric band whose boundary curves are convex in the range $-\phi \leq \theta \leq \phi$ for some angle $\phi > \pi/2$, and concave outside that range.

We prove our convexity claims as follows. The top boundary of \bar{R} is described by the equation $z(\theta) = \sqrt{2a \cos \theta + r^2 - a^2 - 1}$. This curve has an inflection point wherever its second derivative

$$\frac{d^2 z}{d\theta^2}(\theta) = \frac{-a \cos \theta (2a \cos \theta + r^2 - a^2 - 1) - a^2 \sin^2 \theta}{(2a \cos \theta + r^2 - a^2 - 1)^{3/2}}$$

is equal to zero. The denominator of this function is positive whenever $z(\theta)$ is real, and $a > 0$, so we can simplify this equation to

$$a \cos^2 \theta + (r^2 - a^2 - 1) \cos \theta + a = 0.$$

The quadratic formula now implies that

$$\cos \theta = \frac{-r^2 + a^2 + 1 \pm \sqrt{(r^2 - a^2 - 1)^2 - 4a^2}}{2a}.$$

If $r < a + 1$, the expression under the radical is negative, so the equation has no real roots. We easily verify that

$$\frac{d^2 z}{d\theta^2}(0) = \frac{-a}{\sqrt{r^2 - (a - 1)^2}} < 0,$$

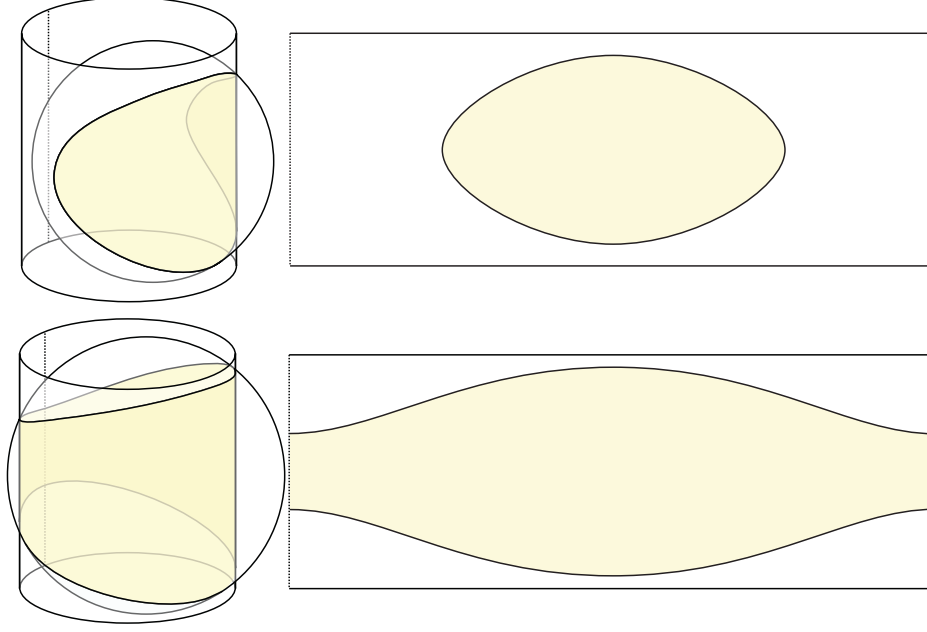


Figure 9. A cylinder and a generic ball intersect in either a disk or an annulus.

so the curve $z(\theta)$ is concave over its entire real domain. It follows that $\overline{B \cap C}$ is convex.

On the other hand, if $r > a + 1$, then both roots are negative reals, but one of them is less than -1 , so it can be discarded. In this case, the function z has exactly two inflection points, located at $\theta = \pm\phi$, where

$$\phi = \arccos \frac{-r^2 + a^2 + 1 + \sqrt{(r^2 - a^2 - 1)^2 - 4a^2}}{2a} > \pi/2.$$

Again, because the curve $z(\theta)$ is concave at 0, it is concave in the range $-\phi < \theta < \phi$ and convex outside that range.

If R is a topological disk, or if R is a topological annulus and both p and q lie in the strip $-\phi \leq \theta \leq \phi$, then \overline{R} contains the segment \overline{pq} . Because \overline{R} is axis-symmetric, it contains at least one axis-aligned right triangle with hypotenuse \overline{pq} .

On the other hand, if either p or q lies on a concave portion of the boundary of \overline{R} , the segment \overline{pq} may not lie inside \overline{R} . Let $\bar{p} = (\theta_1, z_1)$ and $\bar{q} = (\theta_2, z_2)$, and assume without loss of generality that $\theta_1 > \phi$ and $|\theta_2| < |\theta_1|$ (which implies that $|z_2| > |z_1|$). The line $z = z_1$ intersects \overline{R} in a line segment of length $2\theta_1 > 2\phi > \pi$. This segment contains the point (θ_2, z_1) , and therefore also contains either $(\theta_2 + \pi/2, z_1)$ or $(\theta_2 - \pi/2, z_1)$, depending on whether θ_2 is positive or negative. In either case, R contains a special right triangle determined by p and q . \square

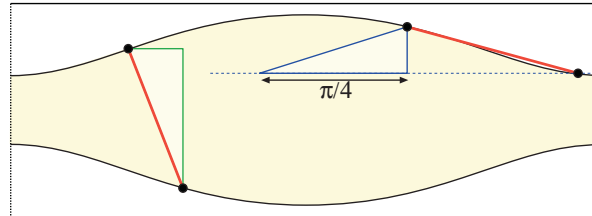


Figure 10. A standard right triangle and a special right triangle determined by pairs of points on an annulus.

Theorem 4.4. *The Delaunay triangulation of n points chosen independently and uniformly at random from a right circular cylinder has expected complexity $O(n \log n)$.*

Proof: Lemma 4.3 implies that every Delaunay edge is either an empty-standard-triangle edge or an empty-special triangle edge. \square

We can also prove a matching $\Omega(n \log n)$ lower bound on the expected Delaunay complexity by considering the *empty-quadrupled-box graph* of P . Two points $p, q \in P$ are joined by an edge in this graph if and only if (1) their angular coordinates differ by at most $\pi/2$ and (2) the box obtained by expanding the axis-aligned bounding box of \bar{p} and \bar{q} by a factor of 2 about its center contains no other points in \bar{P} . The analysis is almost identical to the proof of Theorem 3.5; we omit further details.

Experimental results suggest that the hidden constant in Theorem 4.4 is quite small. To obtain these results, we distributed 2^{20} points uniformly and independently at random on a right cylinder of radius 1 and height 1 and constructed their Delaunay triangulation incrementally using CGAL [7]. For each integer k between 1 and 20, we measure the number of tetrahedra in the Delaunay triangulation of the first 2^k inserted points. Figure 11 shows the average and standard deviation for 20 experiments. Asymptotic behavior is reached quickly, for about a hundred points. The relatively small standard deviation suggests that the $\Theta(n \log n)$ bound may even hold with high probability.

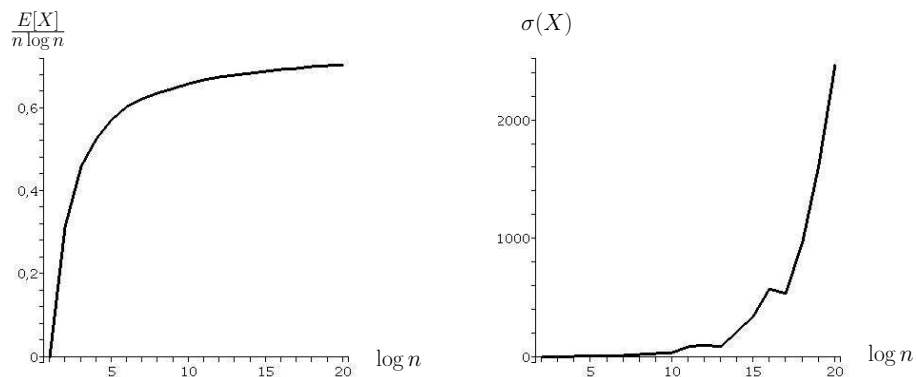


Figure 11. Experimentally measured complexity of the Delaunay triangulation of n random points on a cylinder. Left: ratio of the average complexity over $n \log n$. Right: standard deviation.

References

- [1] N. Amenta, D. Attali, and O. Devillers. Complexity of delaunay triangulation for points on lower-dimensional polyhedra. *Proc. 18th Annu. ACM-SIAM Sympos. Discrete Algorithms*. 2007.
- [2] D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces: the smooth case. *Proc. of the 19th ACM Symposium on Computational Geometry*, pp. 201–210. 2003.
- [3] J.-D. Boissonnat, D. Cohen-Steiner, B. Mourrain, G. Rote, and G. Vegter. Meshing of surfaces. *Effective Computational Geometry for Curves and Surfaces*, pp. 181–229. Springer-Verlag, Mathematics and Visualization, 2006.
- [4] J. Cardinal, S. Collette, and S. Langerman. Region counting graphs. *21st European Workshop on Computational Geometry (EuroCG 2005)*. 2005. (<http://www.win.tue.nl/EWCG2005/Proceedings/6.pdf>).

- [5] J. Cardinal, S. Collette, and S. Langerman. Empty region graphs. Preprint, Oct. 2006. (<http://www.ulb.ac.be/di/algo/secollet/ccl06b.pdf>).
- [6] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. *Effective Computational Geometry for Curves and Surfaces*, pp. 231–276. Springer-Verlag, Mathematics and Visualization, 2006.
- [7] The computational geometry algorithm library. (<http://www.cgal.org>).
- [8] T. Chan, J. Snoeyink, and C.-K. Yap. Primal dividing and dual pruning: Output-sensitive construction of four-dimensional polytopes and three-dimensional Voronoi diagrams. *Discrete and Computational Geometry* 18:433–454, 1997.
- [9] R. Dwyer. The expected number of k -faces of a Voronoi diagram. *Internat. J. Comput. Math.* 26(5):13–21, 1993.
- [10] H. Edelsbrunner. Geometry for modeling biomolecules. *Robotics: The Algorithmic Perspective (Proc. 3rd WAFR)*, pp. 265–277. A. K. Peters, Ltd., 1998.
- [11] J. Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete and Computational Geometry* 30(1):109–132, 2003.
- [12] J. Erickson. Dense point sets have sparse delaunay triangulations. *Discrete and Computational Geometry* 33:83–115, 2005.
- [13] M. Golin and N. H.-S. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. *Computational Geometry: Theory and Applications* 25:197–231, 2003.
- [14] M. Golin and H.-S. Na. The probabilistic complexity of the Voronoi diagram of points on a polyhedron. *Proc. of the 18th ACM Symposium on Computational Geometry*, pp. 209–216. 2002.
- [15] E. W. Jaromczyk and G. T. Toussaint. Relative neighborhood graphs and their relatives. *Proc. IEEE*, pp. 1502–1517. vol. 80(9), 1992.
- [16] D. J. Marchette. Neighborhood graphs. *Random Graphs for Statistical Pattern Recognition*, pp. 73–128. 2004.
- [17] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.
- [18] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd edition. John Wiley & Sons, Chichester, UK, 2000.
- [19] R. Seidel. A convex hull algorithm optimal for point sets in even dimensions. M.Sc. thesis, Dept. Comput. Sci., Univ. British Columbia, Vancouver, BC, 1981. Report 81/14.
- [20] R. Seidel. Exact upper bounds for the number of faces in d -dimensional Voronoi diagrams. *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, pp. 517–530. DIMACS Series in Discrete Mathematics and Theoretical Computer Science 4, AMS Press, 1991.
- [21] R. Seidel. Convex hull computations. *Handbook of Discrete and Computational Geometry*, chapter 19, pp. 361–376. CRC Press LLC, 1997.
- [22] D. Talmor. *Well-Spaced Points and Numerical Methods*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, August 1997.
- [23] S.-H. Teng. *Points, Spheres, and Separators: A Unified Geometric Approach to Graph Partitioning*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1992.

A Characterization and Construction

A.1 The Quadric Polytope

Every axis-aligned quadric has an equation of the form $A + Bx + Cy + Dx^2 + Ey^2 = 0$. The quadric is a (possibly imaginary) ellipse if and only if the coefficients D and E are non-zero and have the same sign. By rewriting this equation in inner-product form as $(A, B, C, D, E) \cdot (1, x, y, x^2, y^2) = 0$, we observe that the space of axis-aligned quadrics is a four-dimensional projective space. It follows that five points $(x_1, y_1), (x_2, y_2), \dots, (x_5, y_5)$ lie on an axis-aligned quadric if and only if

$$\begin{vmatrix} 1 & x_1 & y_1 & x_1^2 & y_1^2 \\ 1 & x_2 & y_2 & x_2^2 & y_2^2 \\ 1 & x_3 & y_3 & x_3^2 & y_3^2 \\ 1 & x_4 & y_4 & x_4^2 & y_4^2 \\ 1 & x_5 & y_5 & x_5^2 & y_5^2 \end{vmatrix} = 0.$$

If this determinant is not zero, its sign indicates whether one point is ‘inside’ or ‘outside’ the oriented axis-aligned quadric determined by the other four points.

Consider the lifting map $\hat{\cdot}: \mathbb{R}^2 \rightarrow \mathbb{R}^4$ that maps any point (a, b) to (a, b, a^2, b^2) . Under this transformation, our quadric determinant test becomes the standard simplex orientation test applied to five lifted points in \mathbb{R}^4 . We refer to the convex hull of the lifted points \hat{P} as the *quadric polytope* of P . For any subset $Q \subseteq P$, the convex hull of \hat{Q} is a face of the quadric polytope if and only if there is an axis-aligned quadric that passes through each point in Q and has every point in $P \setminus Q$ on one side. In particular, the empty-ellipse graph of P is a subgraph of the 1-skeleton of $\text{conv } \hat{P}$. An edge of $\text{conv } \hat{P}$ may also correspond to a pair of points on an axis-aligned ellipse with every other point in P in its interior, or to a pair of points that can be separated from the rest of P by an axis-aligned parabola, axis-aligned hyperbola, or an arbitrary straight line.

This characterization allows us to construct the empty-ellipse graph of P in $O(n^2)$ time, as follows. First, we construct the quadric polytope $\text{conv } \hat{P}$ in $O(n^2)$ time using an algorithm of Seidel [19, 21]. A facet of $\text{conv } \hat{P}$ corresponds to a set of points on an empty ellipse if and only if the last two coefficients of its outer normal vector are positive; we can check this condition in constant time per facet, or in $O(n^2)$ time overall. The empty-ellipse graph consists of the edges of these empty-ellipse facets. As we will prove in Appendix B, the empty-ellipse graph can have $\Omega(n^2)$ edges, so this algorithm is optimal in the worst case.

Theorem A.1. *The empty-ellipse graph of n points in the plane can be constructed in $O(n^2)$ time.*

A.2 Output-Sensitive Construction

Applying an output-sensitive convex hull algorithm does not necessarily give us an output-sensitive algorithm for building empty-ellipse graph, because the the quadric polytope could be a clique even when the empty-ellipse graph has only $O(n)$ edges. (See the proof of Theorem B.4.) However, by adding two points ‘at infinity’ to P , we can guarantee that a constant fraction of the edges of the quadric polytope are empty-ellipse edges.

Specifically, let $P_\infty = P \cup \{(0, \infty), (\infty, 0)\}$, where ∞ is a symbolic value larger than any positive real number.³ If q^+ is an open axis-aligned quadric halfspace that excludes P , then either q^+ is an

³More formally, we are considering the limit behavior of $P_z = P \cup \{(0, z), (z, 0)\}$ as z goes to infinity. The quadric polytope *polyhedron* of P_∞ is the convex hull of \hat{P} and two infinite rays with directions $(0, 0, 0, 1)$ and $(0, 0, 1, 0)$. Equiv-

open elliptical disk, or q^+ contains at least one of the infinite points $(0, \infty)$ and $(\infty, 0)$. Thus, every edge of the quadric polytope of P_∞ is either an empty-ellipse edge for P or has at least one infinite endpoint. In particular, if the empty-ellipse graph of P has e edges, then the quadric polytope of P_∞ has exactly $e + 2n + 1$ edges.

Euler's formula now implies that $\text{conv } \hat{P}_\infty$ has $\Theta(e + n) = \Theta(e)$ faces, so the output-sensitive algorithm of Chan, Snoeyink, and Yap [8] can be used to construct the polytope in time $O(e \log^2 e) = O(e \log^2 n)$. The empty-ellipse graph of P can be extracted easily from this polytope in $O(e)$ time.

Theorem A.2. *The empty-ellipse graph of n points in the plane can be constructed in $O(e \log^2 n)$ time, where e is the number of edges in the graph.*

An easy reduction from sorting, along with the trivial output size bound, establishes a lower bound of $\Omega(n \log n + e)$ on the time to construct empty-ellipse graphs in the worst case, in the algebraic decision tree or real RAM models of computation. Our algorithm matches this output-sensitive lower bound up to a polylogarithmic factor.

B Complexity Bounds

It is easy to prove that the empty-ellipse graph of a generic n -point set can have between $2n - 3$ and $\binom{n}{2}$ edges. In this section, we show by direct construction that neither of these bounds can be improved. To prove that the upper bound is tight, we describe a set of points whose quadric polytope is a *cyclic* polytope—in particular, its 1-skeleton is complete—such that every edge of the quadric polytope is an empty-ellipse edge. To prove that the lower bound is tight, we describe a point set whose empty-ellipse graph coincides with its Delaunay graph. (The quadric polytope of this point set is also a cyclic polytope.) Both point sets are obtained by choosing n arbitrary points on the positive branch of a rational curve: the unit cubic (t, t^3) for the upper bound, and the unit semi-cubic (t^2, t^3) for the lower bound.

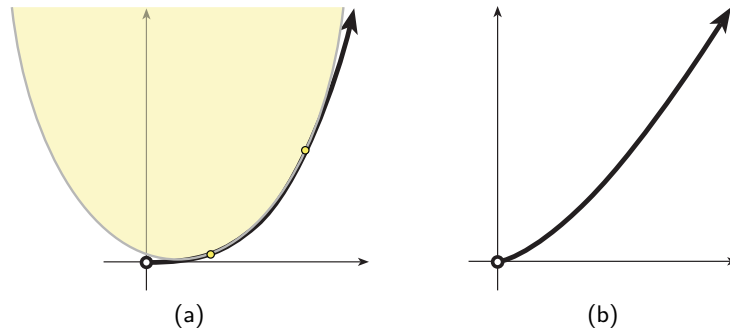


Figure 12. (a) Any set of points on the positive branch of the unit cubic (t, t^3) has a complete empty-ellipse graph. (b) Any set of points on the positive branch of the unit semi-cubic (t^2, t^3) has an outer-planar empty-ellipse graph.

B.1 Upper Bound

It is not hard to show that empty-ellipse graphs can have quadratic complexity in the worst case. Consider a set consisting of $n/2$ points on the line $x - y = 1$, all to the right and above the origin,

alently, we are taking the points from the oriented real projective plane. The two new points lie on the line at infinity, with homogeneous coordinates $[0, 0, 1]$ and $[0, 1, 0]$. Under the lifting map $[w, x, y] \mapsto [w^2, wx, wy, x^2, y^2]$, these two points lift to the points $[0, 0, 0, 0, 1]$ and $[0, 0, 0, 1, 0]$ on the hyperplane at infinity.

and $n/2$ points on the line $y - x = 1$, all to the left and below the origin. Each point in the first subset is an empty-ellipse neighbor of any point in the second subset, so the empty-ellipse graph has at least $n^2/4$ edges. The following more careful construction shows that the empty-ellipse graph can actually be complete in the worst case.

Theorem B.1. *For any positive integer n , there is a set of n points in the plane in general position whose empty-ellipse graph has $\binom{n}{2}$ edges.*

Proof: Let P be an arbitrary set of n points on the positive branch of the unit cubic $\gamma(t) = (t, t^3)$. I claim that the empty-ellipse graph of P is a clique. The argument is similar to Seidel's proof that any set of points on the positive branch of the three-dimensional moment curve (t, t^2, t^3) has a complete Delaunay graph [20].

Let $\gamma(a)$ and $\gamma(b)$ be two arbitrary points in P . The unique axis-aligned conic tangent to γ at these two points has the following equation:

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \begin{vmatrix} 1 & a & a^3 & a^2 & a^6 \\ 1 & a + \varepsilon & (a + \varepsilon)^3 & (a + \varepsilon)^2 & (a + \varepsilon)^6 \\ 1 & b & b^3 & b^2 & b^6 \\ 1 & b + \varepsilon & (b + \varepsilon)^3 & (b + \varepsilon)^2 & (b + \varepsilon)^6 \\ 1 & x & y & x^2 & y^2 \end{vmatrix} = \begin{vmatrix} 1 & a & a^3 & a^2 & a^6 \\ 0 & 1 & 3a^2 & 2a & 6a^5 \\ 1 & b & b^3 & b^2 & b^6 \\ 0 & 1 & 3b^2 & 2b & 6b^5 \\ 1 & x & y & x^2 & y^2 \end{vmatrix} = 0.$$

Expanding the determinant, removing the common factor $(a - b)^2$, and gathering and factoring terms transforms this equation into the following form:

$$3(a^2 + ab + b^2)(a^2 + 3ab + b^2)x^2 - 6ab(a + b)(a^2 + ab + b^2)x + y^2 - 2(a + b)(2a^2 + ab + 2b^2)y + a^2b^2(3a^2 + 4ab + 3b^2) = 0$$

The exact coefficients are not particularly important, but it should be clear that this equation describes an ellipse, because $a > 0$ and $b > 0$. Moreover, this ellipse passes through the points (a, a^3) and (b, b^3) , so it is not imaginary. Substituting $x = t$ and $y = t^3$ into this equation, we obtain a polynomial equation in t , which factors into

$$(a - t)^2(b - t)^2(3a^2 + 4ab + 3b^2 + 2at + 2bt + t^2) = 0.$$

The only positive solutions for this equation (in fact, the only real solutions) are $t = a$ and $t = b$. Thus, the ellipse tangent to $\gamma(a)$ and $\gamma(b)$ does not otherwise intersect the positive branch of the unit cubic. It follows that $\gamma(a)$ and $\gamma(b)$ are empty-ellipse neighbors, no matter what other points P contains. \square

B.2 Lower Bound

Theorem B.2. *The empty-ellipse graph of a generic set of n points in the plane has at least $2n - 3$ edges.*

Proof: Because circles are ellipses, the Delaunay graph is a subgraph of the empty-ellipse graph. The Delaunay triangulation of a generic planar point set P is a triangulation, which by Euler's formula has $3n - 3 - h$ edges, where h is the number of vertices on the convex hull of P . In particular, if P is in convex position, the Delaunay graph of P has exactly $2n - 3$ edges. \square

We will prove this lower bound is tight by relating the empty-ellipse graph to two other geometric graphs. Two points p and q in a planar set P are neighbors in the *empty-horizontal-parabola graph* of P , here denoted $C(P)$, if and only if p and q lie on a parabola whose axis of symmetry is horizontal, with every other point in P outside. The *empty-vertical-parabola graph* $U(P)$ is defined similarly. Both of these graphs contain the convex hull of P as a subgraph, because a line (in any orientation) is a degenerate axis-aligned parabola.

Lemma B.3. *Let P be a set of points in the plane whose empty-horizontal-parabola graph is equal to its empty-vertical-parabola graph. Then both of these graphs are equal to both the Delaunay graph of P and the empty-ellipse graph of P .*

Proof: We define a one-parameter family of subgraphs $D_t(P)$ of the empty ellipse graph, for all $0 \leq t \leq 1$, as follows. Two points $p, q \in P$ are neighbors in $D_t(P)$ if and only if they lie on an empty axis-aligned ellipse described by the equation $tAx^2 + Bx + (1-t)Ay^2 + Dy + E = 0$ for some real constants A, B, D, E . (If $A = 0$, the empty ellipse degenerates to an empty halfplane.) This definition encompasses all the graphs relevant to this lemma: The empty-ellipse graph is $\bigcup_{0 < t < 1} D_t(P)$; the Delaunay graph is $D_{1/2}(P)$; the empty-vertical-parabola graph is $D_0(P)$; and the empty-horizontal-parabola graph is $D_1(P)$. To prove the lemma, it suffices to show that $D_0(P) = D_1(P)$ implies that $D_t(P) = D_0(P)$ for all $0 \leq t \leq 1$.

$D_t(P)$ is the projection of the lower convex hull of the image of P under the lifting map $(x, y) \mapsto (x, y, tx^2 + (1-t)y^2)$. Thus, $D_t(P)$ is a planar graph that partitions the convex hull of P into convex faces. In particular, because $D_0(P) = D_1(P)$, these two graphs define the same convex decomposition.

Fix three points (a, b) , (c, d) , and (e, f) in P that appear in counterclockwise order. These three points are vertices of a common face of $D_t(P)$ if and only if

$$\Delta_t(x, y) := \begin{vmatrix} 1 & a & b & ta^2 + (1-t)b^2 \\ 1 & c & d & tc^2 + (1-t)d^2 \\ 1 & e & f & te^2 + (1-t)f^2 \\ 1 & x & y & tx^2 + (1-t)y^2 \end{vmatrix} \geq 0$$

for every point $(x, y) \in P$. Suppose these three points lie on a common face of $D_0(P) = D_1(P)$. We have both $\Delta_0(x, y) \geq 0$ and $\Delta_1(x, y) \geq 0$ for all $(x, y) \in P$. Thus, for any $0 < t < 1$, we also have $\Delta_t(x, y) = t\Delta_1(x, y) + (1-t)\Delta_0(x, y) \geq 0$, so our three fixed points must also lie on a common face of $D_t(P)$. We conclude that for all t , the graphs $D_t(P)$ and $D_0(P)$ define the same convex decomposition, and therefore are equal as graphs. \square

Theorem B.4. *For any positive integer n , there is a set of n points in the plane in general position whose empty-ellipse graph has $2n - 3$ edges.*

Proof: Let P be an arbitrary set of n points on the positive branch of the unit semi-cubic $\eta(t) = (t^2, t^3)$. A nearly identical argument similar to the previous proof implies that the quadric polytope of P is a cyclic polytope. Specifically, for any positive reals a and b , the unique axis-aligned conic tangent to η at two points $\eta(a)$ and $\eta(b)$ is a hyperbola that does not otherwise intersect the curve η . We omit the straightforward but tedious algebraic details.

We prove the theorem by appealing to Lemma B.3. To prove that the two empty-parabola graphs of P are equal, it suffices to show that the determinants

$$\Sigma_{2,3,4}(a, b, c, d) := \begin{vmatrix} 1 & a^2 & a^3 & a^4 \\ 1 & b^2 & b^3 & b^4 \\ 1 & c^2 & c^3 & c^4 \\ 1 & d^2 & d^3 & d^4 \end{vmatrix} \quad \text{and} \quad \Sigma_{2,3,6}(a, b, c, d) := \begin{vmatrix} 1 & a^2 & a^3 & a^6 \\ 1 & b^2 & b^3 & b^6 \\ 1 & c^2 & c^3 & c^6 \\ 1 & d^2 & d^3 & d^6 \end{vmatrix}$$

are positive for all $0 < a < b < c < d$. The functions $\Sigma_{2,3,4}$ and $\Sigma_{2,3,6}$ are non-zero anti-symmetric integer polynomials of total degree 9 and 11, respectively. Thus, they are both divisible by the discriminant $\Sigma_{1,2,3}(a, b, c, d) := (b-a)(c-a)(c-b)(d-a)(d-b)(d-c) > 0$, and the functions

$$\frac{\Sigma_{2,3,4}(a, b, c, d)}{\Sigma_{1,2,3}(a, b, c, d)} \quad \text{and} \quad \frac{\Sigma_{2,3,6}(a, b, c, d)}{\Sigma_{1,2,3}(a, b, c, d)}$$

are positive symmetric polynomials of total degree 3 and 5, respectively. Specifically,

$$\frac{\Sigma_{2,3,4}(a, b, c, d)}{\Sigma_{1,2,3}(a, b, c, d)} = abc + abd + acd + bcd > 0, \text{ and}$$

$$\begin{aligned} \frac{\Sigma_{2,3,6}(a, b, c, d)}{\Sigma_{1,2,3}(a, b, c, d)} = & a^3bc + ab^3c + abc^3 + a^2b^2c + a^2bc^2 + ab^2c^2 \\ & + a^3bd + ab^3d + abd^3 + a^2b^2d + a^2bd^2 + ab^2d^2 \\ & + a^3cd + ac^3d + acd^3 + a^2c^2d + a^2cd^2 + ac^2d^2 \\ & + b^3cd + bc^3d + bcd^3 + b^2c^2d + b^2cd^2 + bc^2d^2 \\ & + 3a^2bcd + 3ab^2cd + 3abc^2d + 3abcd^2 > 0. \end{aligned}$$

Lemma B.3 now implies that the empty-ellipse graph of P is equal to its Delaunay graph. P is in convex and general position, so its Delaunay graph is a triangulation with exactly $2n - 3$ edges. \square