Nice Point Sets Can Have Nasty Delaunay Triangulations*

Jeff Erickson

University of Illinois, Urbana-Champaign
jeffe@cs.uiuc.edu
http://www.cs.uiuc.edu/~jeffe

Submitted to *Discrete & Computational Geometry*: June 18, 2001

Abstract

We consider the complexity of Delaunay triangulations of sets of points in $\mathbb{R}^3$ under certain practical geometric constraints. The *spread* of a set of points is the ratio between the longest and shortest pairwise distances. We show that in the worst case, the Delaunay triangulation of n points in $\mathbb{R}^3$ with spread Δ has complexity $\Omega(\min\{\Delta^3, n\Delta, n^2\})$ and $O(\min\{\Delta^4, n^2\})$. For the case $\Delta = \Theta(\sqrt{n})$, our lower bound construction consists of a grid-like sample of a right circular cylinder with constant height and radius. We also construct a family of smooth connected surfaces such that the Delaunay triangulation of any good point sample has near-quadratic complexity.

1 Introduction

Delaunay triangulations and Voronoi diagrams are used as a fundamental tool in several geometric application areas, including finite-element mesh generation [17, 25, 37, 40], deformable surface modeling [16], and surface reconstruction [1, 3, 4, 5, 12, 35]. Many algorithms in these application domains begin by constructing the Delaunay triangulation of a set of n points in $\mathbb{R}^3$. Delaunay triangulations can have complexity $\Omega(n^2)$ in the worst case, and as a result, all these algorithms have worst-case running time $\Omega(n^2)$. However, this behavior is almost never observed in practice except for highly-contrived inputs. For all practical purposes, three-dimensional Delaunay triangulations appear to have linear complexity.

One way to explain this frustrating discrepancy between theoretical and practical behavior would be to identify geometric constraints that are satisfied by real-world input and to analyze Delaunay triangulations under those constraints. These constraints would be similar to the *realistic input models* such as fatness or simple cover complexity [8, 46], which many authors have used to develop geometric algorithms with good practical performance. Unlike these works, however, our (immediate) goal is not to develop new algorithms, but rather to formally explain the good practical performance of existing code.

*Portions of this work were done while the author was visiting INRIA, Sophia-Antipolis, with the support of a UIUC/CNRS/INRIA travel grant. This research was also partially supported by a Sloan Fellowship and by NSF CAREER grant CCR-0093348. An extended abstract of this paper was presented at the 17th Annual ACM Symposium on Computational Geometry [30]. See http://www.cs.uiuc.edu/~jeffe/pubs/spread.html for the most recent version of this paper.

Dwyer [23, 24] showed that if a set of points is generated uniformly at random from the unit ball, its Delaunay triangulation has linear expected complexity. Golin and Na [33] recently derived a similar result for random points on the surface of a three-dimensional convex polytope with constant complexity. Although these results are encouraging, they are unsatisfying as an explanation of practical behavior. Real-world point data generated by laser range finders, digital cameras, tomographic scanners, and similar input devices is often highly structured.

This paper considers the complexity of Delaunay triangulations under two types of practical geometric constraints. First, in Section 2, we consider the worst-case Delaunay complexity as a function of both the number of points and the *spread*—the ratio between its diameter and the distance between its closest pair. For any n and Δ, we construct a set of n points with spread Δ whose Delaunay triangulation has complexity $\Omega(\min\{\Delta^3, n\Delta, n^2\})$. When $\Delta = \Theta(\sqrt{n})$, our lower bound construction consists of a grid-like sample of a right circular cylinder with constant height and radius. We also show that the worst-case complexity of a Delaunay triangulation is $O(\min\{\Delta^4, n^2\})$. We conjecture that our lower bounds are tight, and sketch a possible technique to improve our upper bounds.

An important application of Delaunay triangulations that has received a lot of attention recently is surface reconstruction: Given a set of points from a smooth surface Σ, reconstruct an approximation of Σ. Several algorithms provably reconstruct surfaces if the input points satisfy certain sampling conditions [4, 5, 12, 35]. In Section 3, we consider the complexity of Delaunay triangulations of good samples of smooth surfaces. Not surprisingly, oversampling almost any surface can produce a point set whose Delaunay triangulation has quadratic complexity. We show that even surface data with *no* oversampling can have quadratic Delaunay triangulations and that there are smooth surfaces where *every* good sample has near-quadratic Delaunay complexity. We also derive similar results for randomly distributed points on non-convex smooth surfaces. An important tool in our proofs is the definition of *sample measure*, which measures the intrinsic difficulty of sampling a smooth surface for reconstruction.

Throughout the paper, we analyze the complexity of three-dimensional Delaunay triangulations by counting the number of edges. Two points are joined by an edge in the Delaunay triangulation of a set S if and only if they lie on a sphere with no points of S in its interior. Euler's formula implies that any three-dimensional triangulation with n vertices and e edges has at most $2e - 2n$ triangles and $e - n$ tetrahedra, since the link of every vertex is a planar graph.

2 Sublinear Spread

We define the *spread* Δ of a set of points (also called the *distance ratio* [18]) as the ratio between the longest and shortest pairwise distances. In this section, we derive upper and lower bounds on the worst-case complexity of the Delaunay triangulation of a point set in $\mathbb{R}^3$, as a function of both the number of points and the spread. The spread is minimized at $\Theta(n^{1/3})$ when the points are packed into a tight lattice, in which case the Delaunay triangulation has only linear complexity. On the other hand, all known examples of point sets with quadratic-complexity Delaunay triangulations, such as points on the moment curve or a pair of skew lines, have spread $\Omega(n)$. Thus, it is natural to ask how the complexity of the Delaunay triangulation changes as the spread varies between these two extremes.

The spread of a set of points is loosely related to its dimensionality. If a set uniformly covers a bounded region of space, a surface of bounded curvature, or a curve of bounded curvature, its spread is respectively $\Theta(n^{1/3})$, $\Theta(n^{1/2})$, or $\Theta(n)$. The case of surface data is particularly interesting in light of numerous algorithms that reconstruct surfaces using a subcomplex of the Delaunay triangulation. We will discuss surface reconstruction in more detail in the next section. Indyk *et al.* [36] observed that molecular data usually has sublinear spread, by examining a database of over 100,000 small drug molecules.

Several algorithmic and combinatorial bounds are known that depend favorably on spread, especially for *dense* point sets; a d-dimensional point set is dense if its spread is $O(n^{1/d})$. Edelsbrunner *et al.* [28] showed that a dense point set in the plane has at most $O(n^{7/6})$ halving lines[1], and a dense point set in $\mathbb{R}^3$ has at most $O(n^{7/3})$ halving planes. The best upper bounds known for arbitrary point sets are $O(n^{4/3})$ [19] and $O(n^{5/2})$ [41], respectively. Valtr [42, 43, 44] proved several other combinatorial bounds for dense planar sets that improve the corresponding worst-case bounds. Verbarg [45] describes an efficient algorithm to find approximate center points in dense point sets. Cardoze and Schulman [14], Indyk *et al.* [36], and Gavrilov *et al.* [31] describe algorithms for approximate geometric pattern matching whose running times depend favorably on the spread of the input set. Clarkson [18] describes data structures for nearest neighbor queries in arbitrary metric spaces which are efficient if the spread of the input is small.

Although our results are described in terms of spread, they also apply to other more robust quality measures. For example, we could define the *average spread* of a point set as the average (in some sense), over all points p, of the ratio between the distances from p to its farthest and nearest neighbors. In each of our constructions, the distance ratios of different points differ by at most a small constant factor, so our results apply to average spread as well.

2.1 Lower Bounds

The crucial special case of our lower bound construction is $\Delta = \Theta(\sqrt{n})$. For any positive integer x, let $[x]$ denote the set $\{1, 2, \ldots, x\}$. Our construction consists of n evenly spaced points on a helical space curve:

$$S_{\sqrt{n}} = \left\{ \left(\frac{2\pi k}{n}, \cos\frac{2\pi k}{\sqrt{n}}, \sin\frac{2\pi k}{\sqrt{n}} \right) \;\middle|\; k \in [n] \right\}.$$

See Figure 1. $S_{\sqrt{n}}$ is a grid-like uniform ε-sample (see Section 3) of a right circular cylinder, where $\varepsilon = \Theta(\sqrt{1/n})$. By adding additional points on two hemispherical caps at the ends of the cylinder, we can extend $S_{\sqrt{n}}$ into a uniform ε-sample of a smooth convex surface with bounded curvature and constant local feature size. The closest pair of points in $S_{\sqrt{n}}$ has distance $2\pi/\sqrt{n} + o(1)$, and the diameter of $S_{\sqrt{n}}$ is $2\pi - o(1)$, so the spread of $S_{\sqrt{n}}$ is $\sqrt{n} - o(1)$. We will show that the Delaunay triangulation of $S_{\sqrt{n}}$ has complexity $\Omega(n^{3/2})$.

Let $h_\alpha(t)$ denote the helix $(\alpha t, \cos t, \sin t)$, where the parameter $\alpha > 0$ is called the *pitch*. The combinatorial structure of the Delaunay triangulation depends entirely on the signs of certain insphere determinants. Using elementary trigonometric identities and matrix operations, we can

[1] Edelsbrunner *et al.* [28] only prove the upper bound $O(n^{5/4}/\log^* n)$; the improved bound follows immediately from their techniques and Dey's more recent $O(n^{4/3})$ worst-case bound [19].

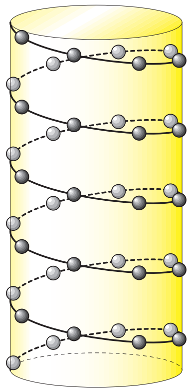

Figure 1. A set of n points whose Delaunay triangulation has complexity $\Omega(n^{3/2})$

simplify the insphere determinant for five points on this helix as follows.

$$\begin{vmatrix} 1 & \alpha t_1 & \cos t_1 & \sin t_1 & \alpha^2 t_1^2 + \cos^2 t_1 + \sin^2 t_1 \\ 1 & \alpha t_2 & \cos t_2 & \sin t_2 & \alpha^2 t_2^2 + \cos^2 t_2 + \sin^2 t_2 \\ 1 & \alpha t_3 & \cos t_3 & \sin t_3 & \alpha^2 t_3^2 + \cos^2 t_3 + \sin^2 t_3 \\ 1 & \alpha t_4 & \cos t_4 & \sin t_4 & \alpha^2 t_4^2 + \cos^2 t_4 + \sin^2 t_4 \\ 1 & \alpha t_5 & \cos t_5 & \sin t_5 & \alpha^2 t_5^2 + \cos^2 t_5 + \sin^2 t_5 \end{vmatrix} = \alpha^3 \begin{vmatrix} 1 & t_1 & \cos t_1 & \sin t_1 & t_1^2 \\ 1 & t_2 & \cos t_2 & \sin t_2 & t_2^2 \\ 1 & t_3 & \cos t_3 & \sin t_3 & t_3^2 \\ 1 & t_4 & \cos t_4 & \sin t_4 & t_4^2 \\ 1 & t_5 & \cos t_5 & \sin t_5 & t_5^2 \end{vmatrix}$$

We obtain the surprising observation that changing the pitch α does not change the combinatorial structure of the Delaunay triangulation of any set of points on the helix. (More generally, scaling any set of points on any circular cylinder along the cylinder's axis leaves the Delaunay triangulation invariant.) Thus, for purposes of analysis, it suffices to consider the case $\alpha = 1$. Let $h(t) = h_1(t) = (t, \cos t, \sin t)$.

Our first important observation is that any set of points on a single turn of any helix has a *neighborly* Delaunay triangulation, meaning that every pair of points is connected by a Delaunay edge. For any real value t, define the *bitangent sphere* $\beta(t)$ to be the unique sphere passing through $h(t)$ and $h(-t)$ and tangent to the helix at those two points.

Lemma 2.1. *For any $0 < t < \pi$, the sphere $\beta(t)$ intersects the helix h only at its two points of tangency.*

Proof: Symmetry considerations imply that the bitangent sphere must be centered on the y-axis, so it is described by the equation $x^2 + (y-a)^2 + z^2 = r^2$ for some constants a and r. Let γ denote the intersection curve of $\beta(t)$ and the cylinder $y^2 + z^2 = 1$. Every intersection point between $\beta(t)$ and the helix must lie on γ. If we project the helix and the intersection curve to the xy-plane, we obtain the sinusoid $y = \cos x$ and a portion of the parabola $y = \gamma(x) = (x^2 - r^2 + a^2 + 1)/2a$. These two curves meet tangentially at the points $(t, \cos t)$ and $(-t, \cos t)$.

The mean value theorem implies that $\gamma(x) = \cos x$ at most four times in the range $-\pi < x < \pi$. (Otherwise, the curves $y'' = -\cos x$ and $y'' = \gamma''(x) = 1/a$ would intersect more than twice in that range.) Since the curves meet with even multiplicity at two points, those are the only intersection points in the range $-\pi < x < \pi$. Since $\gamma(x)$ is concave, we have $\gamma(\pm\pi) < \cos \pm\pi = -1$, so there are no intersections with $|x| \geq \pi$. Thus, the curves meet only at their two points of tangency. $\square$

Corollary 2.2. *Any set S of n points on the helix $h(t)$ in the range $-\pi < t < \pi$ has a neighborly Delaunay triangulation.*

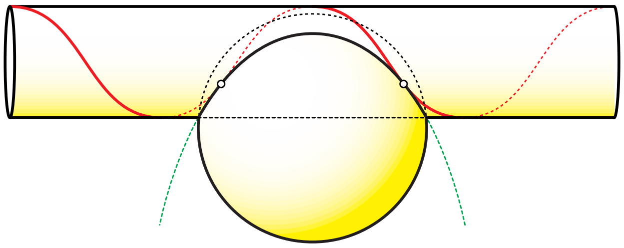

Figure 2. The intersection curve of the cylinder and a bitangent sphere projects to a parabola on the xy-plane.

Proof: Let p and q be arbitrary points in S, and let β be the unique ball tangent to the helix at p and q. By Lemma 2.1, β does not otherwise intersect the helix and therefore contains no point in S. Thus, p and q are neighbors in the Delaunay triangulation of S. $\square$

We can now easily complete the analysis of our helical point set $S_{\sqrt{n}}$. Lemma 2.1 implies that every point in $S_{\sqrt{n}}$ is connected by a Delaunay edge to every other point less than a full turn around the helix $h_{\sqrt{1/n}}(t)$. Each full turn of the helix contains at least $\lfloor\sqrt{n}\rfloor$ points. Thus, except for points on the first and last turn, every point has at least $2\lfloor\sqrt{n}\rfloor$ Delaunay neighbors, so the total number of Delaunay edges is more than $2\lfloor\sqrt{n}\rfloor(n-2\lfloor\sqrt{n}\rfloor) > 2n^{3/2}-4n$. This crude lower bound does not count Delaunay edges in the first and last turns of the helix, nor edges that join points more than one turn apart. It is not difficult to show that there are at most $O(n)$ uncounted Delaunay edges if $\sqrt{n}$ is an integer [29] and at most $O(n^{3/2})$ uncounted edges in general.

Theorem 2.3. *For any n, there is a set of n points in $\mathbb{R}^3$ with spread $\sqrt{n}$ whose Delaunay triangulation has complexity $\Omega(n^{3/2})$. Moreover, this point set is a uniform sample of a smooth convex surface with constant local feature size.*

We generalize our helix construction to other values of the spread Δ as follows.

Theorem 2.4. *For any n and $\Delta = \Omega(n^{1/3})$, there is a set of n points in $\mathbb{R}^3$ with spread Δ whose Delaunay triangulation has complexity $\Omega(\min\{\Delta^3, n\Delta, n^2\})$.*

Proof: There are three cases to consider, depending on whether the spread is at least n, between $\sqrt{n}$ and n, or at most $\sqrt{n}$. The first case is trivial.

For the case $\sqrt{n} \le \Delta \le n$, we take a set of evenly spaced points on a helix with pitch Δ/n:

$$S_\Delta = \left\{ \left(\frac{2\pi k}{n}, \cos\frac{2\pi k}{\Delta}, \sin\frac{2\pi k}{\Delta} \right) \,\middle|\, k \in [n] \right\}.$$

Every point in S_Δ is connected by a Delaunay edge to every other point less than a full turn away on the helix, and each turn of the helix contains $\Omega(\Delta)$ points, so the total complexity of the Delaunay triangulation is $\Omega(n\Delta)$.

The final case $n^{1/3} \le \Delta \le \sqrt{n}$ is somewhat more complicated. Our point set consists of several copies of our helix construction, with the helices positioned at the points of a square lattice, so the entire construction loosely resembles a mattress. Specifically,

$$S_\Delta = \left\{ \left(\frac{2\pi k}{r}, 6i + \cos\frac{2\pi k}{\sqrt{r}}, 6j + \sin\frac{2\pi k}{\sqrt{r}} \right) \,\middle|\, i, j \in [w];\, k \in [wr] \right\},$$

where r and w are parameters to be determined shortly. This set contains $n = w^3 r$ points. The diameter of S_Δ is $\Theta(w)$ and the closest pair distance is $\Theta(1/\sqrt{r})$, so its spread is $\Delta = \Theta(w\sqrt{r})$. Thus, given n and Δ, we have $w = \Theta(n/\Delta^2)$ and $r = \Theta(\Delta^6/n^2)$.

To complete our analysis, we need to show that Delaunay circumspheres from one helix do not interfere significantly with nearby helices. Let $\beta_\alpha(t)$ denote the sphere tangent to the helix h_α at $h_\alpha(t)$ and $h_\alpha(-t)$, for some $0 < \alpha \le 1$ and $0 < t \le \pi/2$. We claim that the radius of this sphere is less than 3. Since $\beta_\alpha(t)$ is centered on the y-axis (see Lemma 2.1), we can compute its radius by computing its intersection with the y-axis. The intersection points satisfy the determinant equation

$$\begin{vmatrix} 1 & \alpha t & \cos t & \sin t & 1+\alpha^2 t^2 \\ 0 & \alpha & -\sin t & \cos t & 2\alpha^2 t \\ 1 & -\alpha t & \cos t & -\sin t & 1+\alpha^2 t^2 \\ 0 & -\alpha & -\sin t & -\cos t & 2\alpha^2 t \\ 1 & 0 & y & 0 & y^2 \end{vmatrix} = 0.$$

For any $\alpha > 0$ and $t > 0$, this equation simplifies to

$$(\sin t)y^2 + 2\alpha^2 t y - (\sin t)(1+\alpha^2 t^2) = 0,$$

which implies that the radius of $\beta_\alpha(t)$ is

$$\sqrt{\frac{\alpha^4 t^2}{\sin^2 t} + \alpha^2 t^2 + 1}.$$

We easily verify that this is an increasing function of both α and t in the range of interest. Thus, to prove our claim, it suffices to observe that the radius of $\beta_1(\pi/2)$ is $\sqrt{\pi^2/2+1} \approx 2.4361 < 3$.

Since adjacent helices are separated by distance 6, every point in S_Δ is connected by a Delaunay edge to every point at most half a turn away in the same helix. Each turn of each helix contains $\Omega(\sqrt{r})$ points, so the Delaunay triangulation of S_Δ has complexity $\Omega(n\sqrt{r}) = \Omega(\Delta^3)$. $\square$

2.2 Upper Bounds

Let B be a ball of radius R in $\mathbb{R}^3$, and let $b_1, b_2, b_3, \ldots$ be balls of radius at least r, where $1 \le r \le R$. Our upper bound proof uses geometric properties of the 'Swiss cheese' $C = B \setminus \bigcup_i b_i$. See Figure 3(a). In our upper bound proofs, B will be a ball that contains a subset of the points, and each b_i will be an empty circumsphere of some Delaunay edge.

Lemma 2.5. *The surface area of* C *is* $O(R^3/r)$.

Proof: The outer surface $\partial C \cap \partial B$ clearly has area $O(R^2) = O(R^3/r)$, so it suffices to bound the surface area of the 'holes'. For each i, let $H_i = B \cap \partial b_i$ be the boundary of the ith hole, and let $H = \bigcup_i H_i = \partial C \setminus \partial B$. For any point $x \in H$, let s_x denote the open line segment of length r extending from x towards the center of the ball b_i with x on its boundary. (If x lies on the surface of more than one b_i, choose one arbitrarily.) Let $S = \bigcup_{x \in H} s_x$ be the union of all such segments, and for each i, let $S_i = \bigcup_{x \in H_i} s_x$. Each S_i is a fragment of a spherical shell of thickness r inside the ball b_i. See Figure 3(b).

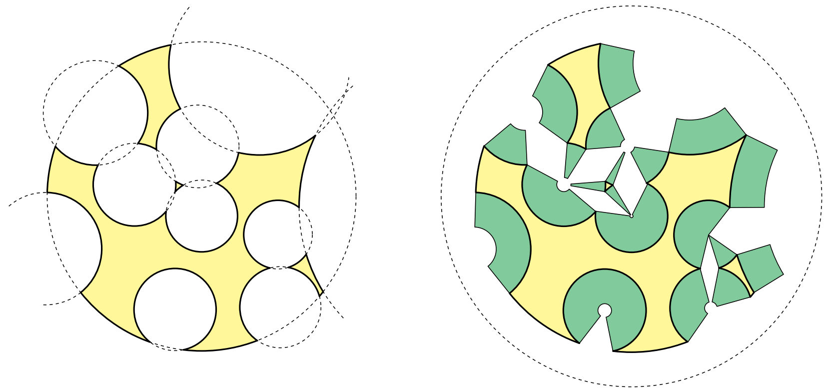

Figure 3. (a) Swiss cheese (in $\mathbb{R}^2$). (b) Shell fragments used to bound its surface area.

We can bound the volume of each shell fragment S_i as follows:

$$\mathrm{vol}(S_i) = \frac{r_i}{3}\,\mathrm{area}(H_i) - \frac{r_i - r}{3} \cdot \left(\frac{r_i - r}{r_i}\right)^2 \mathrm{area}(H_i) \geq \frac{r}{3}\,\mathrm{area}(H_i),$$

where $r_i \geq r$ is the radius of b_i. The triangle inequality implies that s_x and s_y are disjoint for any two points $x, y \in H$, so the shell fragments S_i are pairwise disjoint. Finally, since S fits inside a ball of radius $R + r \leq 2R$, its volume is $O(R^3)$. Thus, we have

$$\mathrm{area}(H) = \sum_i \mathrm{area}(H_i) \leq \sum_i \frac{3\,\mathrm{vol}(S_i)}{r} = \frac{3\,\mathrm{vol}(S)}{r} \leq \frac{4\pi R^3}{r}.$$

□

At this point, we would like to argue that any unit ball whose center is on the boundary of C contains a constant amount of surface area of C, so that we can apply a packing argument. Unfortunately, C might contain isolated components and thin handles with arbitrarily small surface area (like the small triangular component in Figure 3(a)). Thus, we consider balls centered slightly away from the boundary of C.

Lemma 2.6. *Let U be any unit ball whose center is in C and at distance $2/3$ from ∂C. Then U contains $\Omega(1)$ surface area of C.*

Proof: Without loss of generality, assume that U is centered at the origin and that $(0, 0, 2/3)$ is the closest point of ∂C to the origin. Let U' be the open ball of radius $2/3$ centered at the origin, let V be the open unit ball centered at $(0, 0, 5/3)$, and let W be the cone whose apex is the origin and whose base is the circle $\partial U \cap \partial V$. See Figure 4. U' lies entirely inside C, and since $r \geq 1$, we easily observe that V lies entirely outside C. Thus, the surface area of $\partial C \cap W \subseteq \partial C \cap U$ is at least the area of the spherical cap $\partial U' \cap W$, which is exactly $4\pi/27$. □

Theorem 2.7. *Let S be a set of points in $\mathbb{R}^3$ whose closest pair is at distance 2, and let r be any real number. Any point in S has $O(r^2)$ Delaunay neighbors at distance at most r.*

Proof: Let o be an arbitrary point in S, and let B be a ball of radius r centered at o. Call a Delaunay neighbor of o a *friend* if it lies inside B, and call a friend q *interesting* if there is another

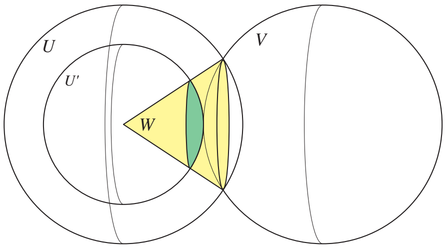

Figure 4. Proof of Lemma 2.6

point $p \in S$ (not necessarily a Delaunay neighbor of o) such that $|op| < |oq|$ and $\angle poq < 1/r$. A simple packing argument implies that o has at most $O(r^2)$ boring friends.

Let q be an interesting friend of o, and let p be a point that makes q interesting, as described above. Since q is a Delaunay neighbor of o, there is a ball d_q that has o and q on its boundary and no points of S in its interior. In particular, p is outside d_q, so the radius of d_q is greater than the radius of the circle passing through o, p, and q. Let c be the center of this circle. Since $\angle poq < 1/r$, we must have $\angle pcq < 2/r$, and since $|pq| \geq 2$, we must have $|cq| > r$. See Figure 5. We have just shown that the radius of d_q is at least r.

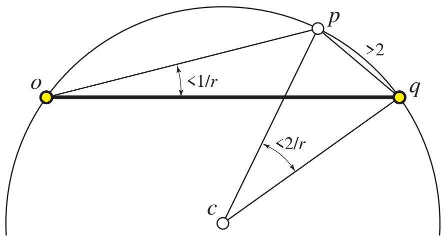

Figure 5. The radius of any interesting Delaunay ball is at least r.

For every interesting friend q, let b_q be the ball concentric with d_q with radius $2/3$ less than the radius of d_q, and let U_q be the unit-radius ball centered at q. We now have a set of unit balls, one for each interesting friend of o, whose centers lie at distance exactly $2/3$ from the boundary of the Swiss cheese $C = B \setminus \bigcup_q b_q$. By Lemma 2.5, C has surface area $O(r^2)$, and by Lemma 2.6, each unit ball U_q contains $\Omega(1)$ surface area of C. Since the unit balls are disjoint, it follows that o has at most $O(r^2)$ interesting friends. □

Theorem 2.8. *Let S be a set of points in $\mathbb{R}^3$ whose closest pair is at distance 2 and whose diameter is 2Δ, and let r be any real number. There are $O(\Delta^3/r)$ points in S with a Delaunay neighbor at distance at least r.*

Proof: Call a point *far-reaching* if it has a Delaunay neighbor at distance at least r, and let Q be the set of far-reaching points. Let B be a ball of radius 2Δ containing S. For each $q \in Q$, let f_q be a maximal empty ball containing q and its furthest Delaunay neighbor, and let b_q be the concentric ball with radius $2/3$ smaller than f_p. By construction, each ball b_q has radius at least $r/2 - 2/3$. Finally, for any far-reaching point q, let U_q be the unit-radius ball centered at q. By Lemma 2.5, the Swiss cheese $C = B \setminus \bigcup_{q \in Q} b_q$ has surface area $O(\Delta^3/r)$, and by Lemma 2.6, each unit ball U_q contains $\Omega(1)$ surface area of C. Since these unit balls are disjoint, there are at most $O(\Delta^3/r)$ of them. □

Corollary 2.9. *Let S be a set of points in $\mathbb{R}^3$ with spread Δ. The Delaunay triangulation of S has complexity $O(\Delta^4)$.*

Proof: For all r, let $F(r)$ be the number of far-reaching points in S, *i.e.*, those with Delaunay edges of length at least r. From Theorem 2.8, we have $F(r) = O(\Delta^3/r)$. By Theorem 2.7, if the farthest neighbor of a point p is at distance between r and $r+1$, then p has $O(r^2)$ Delaunay neighbors. Thus, the total number of Delaunay edges is at most

$$\sum_{r=0}^{\Delta} O(r^2) \cdot \big(F(r) - F(r+1)\big) = \sum_{r=0}^{\Delta} O(r) \cdot F(r) = \sum_{r=0}^{\Delta} O(\Delta^3) = O(\Delta^4).$$

□

2.3 Conjectured Upper Bounds

We conjecture that the lower bounds in Theorem 2.4 are tight, but Corollary 2.9 is the best upper bound known. Nearly matching upper bounds could be derived from the following conjecture, using a divide and conquer argument suggested by Edgar Ramos (personal communication).

Let S be a *well-separated* set of points with closest pair distance 1, lying in two balls of radius Δ that are separated by distance at least $c\Delta$ for some constant $c > 1$. Call an edge in the Delaunay triangulation of S a *crossing edge* if it has one endpoint in each ball.

Conjecture 2.10. *Some point in S is an endpoint of $O(\Delta)$ crossing edges.*

Lemma 2.11. *Conjecture 2.10 implies that S has $O(\min\{\Delta^3, \Delta n, n^2\})$ crossing edges.*

Proof: Theorem 2.8 implies that only $O(\Delta^2)$ points can be endpoints of crossing edges. Thus, we can assume without loss of generality that $n = O(\Delta^2)$. We compute the total number of crossing edges by iteratively removing the point with the fewest crossing edges and retriangulating the resulting hole, say by incremental flipping [27]. Conjecture 2.10 implies that we delete only $O(\Delta)$ crossing edges with each point, so altogether we delete $O(n\Delta) = O(\Delta^3)$ crossing edges. Not all of these edges are in the original Delaunay triangulation, but that only helps us. □

Theorem 2.12. *Conjecture 2.10 implies that the Delaunay triangulation of n points in $\mathbb{R}^3$ with spread Δ has complexity $O(\min\{\Delta^3 \log \Delta, n\Delta, n^2\})$.*

Proof: Assume Conjecture 2.10 is true, and let S be an arbitrary set of n points with diameter Δ, where the closest pair of points is at unit distance. S is contained in an axis-parallel cube C of width Δ. We construct a well-separated pair decomposition of S [13], based on a simple octtree decomposition of C. The octtree has $O(\log \Delta)$ levels. At each level i, there are 8^i cells, each a cube of width $\Delta/2^i$. Our well-separated pair decomposition includes, for each level i, the points in any pair of level-i cells separated by a distance between $c\Delta/2^i$ and $2c\Delta/2^i$. A simple packing argument implies that any cell in the octtree is paired with $O(1)$ other cells, all at the same level, and so any point appears in $O(\log \Delta)$ subset pairs. Every Delaunay edge of S is a crossing edge for some well-separated pair of cells.

Lemma 2.11 implies that the points in any well-separated pair of level-i cells have $O(\Delta^3/8^i)$ crossing Delaunay edges. Since there are $O(8^i)$ such pairs, the total number of crossing edges between level-i cells is $O(\Delta^3)$. Thus, there are $O(\Delta^3 \log \Delta)$ Delaunay edges altogether.

Lemma 2.11 also implies that for any well-separated pair of level-i cells, the average number of crossing edges per point is $O(\Delta/2^i)$. Since every point belongs to a constant number of subset pairs at each level, the total number of crossing edges at level i is $O(n\Delta/2^i)$. Thus, the total number of Delaunay edges is $O(n\Delta)$. □

This upper bound is still a logarithmic factor away from our lower bound construction when $\Delta = o(\sqrt{n})$. However, our argument is quite conservative; all crossing edges for a well-separated pair of subsets are counted, even though some or all of these edges may be blocked by other points in S. A more careful analysis would probably eliminate the final logarithmic factor.

3 Nice Surface Data

Let Σ be a smooth surface without boundary in $\mathbb{R}^3$. The *medial axis* of Σ is the closure of the set of points in $\mathbb{R}^3$ that have more than one nearest neighbor on Σ. The *local feature size* of a surface point x, denoted $\mathrm{lfs}(x)$, is the distance from x to the medial axis of Σ. Let S be a set of *sample points* on Σ. Following Amenta and Bern [1], we say that S is an ε*-sample* of Σ if the distance from any point $x \in \Sigma$ to the nearest sample point is at most $\varepsilon \cdot \mathrm{lfs}(x)$.

The first step in several surface reconstruction algorithms is to construct the Delaunay triangulation or Voronoi diagram of the sample points. Edelsbrunner and Mücke [26] and Bajaj *et al.* [7, 9] describe algorithms based on *alpha shapes*, which are subcomplexes of the Delaunay triangulation; see also [34]. Extending earlier work on planar curve reconstruction [2, 32], Amenta and Bern [1, 3] developed an algorithm to extract a certain manifold subcomplex of the Delaunay triangulation, called the *crust*. Amenta *et al.* [4] simplified the crust algorithm and proved that if S is an ε-sample of a smooth surface Σ, for some sufficiently small ε, then the crust is homeomorphic to Σ. Boissonnat and Cazals [12] and Hiyoshi and Sugihara [35] proposed algorithms to produce a smooth surface using natural coordinates, which are defined and computed using the Voronoi diagram of the sample points. Further examples can be found in [5, 6, 11, 16, 21].

We have already seen that even very regular ε-samples of smooth surfaces can have super-linear Delaunay complexity. In this section, we show that ε-samples of smooth surfaces can have Delaunay triangulations of quadratic complexity, implying that all these surface reconstruction algorithms take at least quadratic time in the worst case.

3.1 Sample Measure

We will analyze our lower bound constructions in terms of the *sample measure* of a smooth surface Σ, which we define as follows:

$$\mu(\Sigma) = \int_\Sigma \frac{dx}{\mathrm{lfs}^2(x)}.$$

Intuitively, the sample measure of a surface describes the intrinsic difficulty of sampling that surface for reconstruction.[2] The next lemma formalizes this intuition.

Lemma 3.1. *For any C^2 surface Σ and and $\varepsilon < 1/5$, every ε-sample of Σ contains $\Omega(\mu(\Sigma)/\varepsilon^2)$ points.*

[2] Ruppert and Seidel [39] use precisely this function—but with a different definition of local feature size—to measure the minimum number of triangles with bounded aspect ratio required to mesh a planar straight-line graph.

Proof: Let S be an arbitrary ε-sample of Σ for some $\varepsilon < 1/5$, and let $n = |S|$.

Amenta and Bern [1, Lemma 1] observe that the local feature size function is 1-Lipschitz, that is, $|\mathrm{lfs}(x) - \mathrm{lfs}(y)| < |xy|$ for any surface points $x, y \in \Sigma$. Thus, for any point $x \in \Sigma$, we have $|xp| \le \varepsilon\,\mathrm{lfs}(x) \le \varepsilon(\mathrm{lfs}(p) + |xp|)$, so $|xp| \le \frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)$, where $p \in S$ is the sample point closest to x.

It follows that we can cover Σ with spheres of radius $\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)$ around each sample point p. Call the intersection of Σ and the sphere around p the *neighborhood* of p, denoted $N(p)$. Similar Lipschitz arguments imply that

$$\mathrm{lfs}(x) \ge \frac{1-2\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p).$$

for any point $x \in N(p)$.

For any $x \in \Sigma$, let n_x denote the normal vector to Σ at x. Using the fact that local feature size is at most the minimum radius of principal curvature, Amenta and Bern [1, Lemma 2] prove that for any $x, y \in \Sigma$ where $\delta = |xy|/\mathrm{lfs}(x) < 1/3$, we have $\angle n_x n_y \le \frac{\delta}{1-3\delta}$. Thus,

$$\angle n_x n_p \le \frac{\varepsilon}{1-4\varepsilon} < 1 < \frac{\pi}{3}$$

for any $x \in N(p)$. It follow that $N(p)$ is monotone with respect to n_p, so we can compute its area by projecting it onto a plane normal to n_p. Since the projection fits inside a circle of radius $\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)$, we have

$$\mathrm{area}(N(p)) \le \frac{\pi\left(\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)\right)^2}{\min\limits_{x \in N(p)} \cos\angle n_x n_p} \le \frac{\pi\left(\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)\right)^2}{\cos(\pi/3)} = 2\pi\left(\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)\right)^2.$$

We can now bound the sample measure of each neighborhood as follows:[3]

$$\mu(N(p)) \le \frac{\mathrm{area}(N(p))}{\min\limits_{x \in N(p)} \mathrm{lfs}^2(x)} \le \frac{2\pi\left(\frac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)\right)^2}{\left(\frac{1-2\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(p)\right)^2} = \frac{2\pi\varepsilon^2}{(1-2\varepsilon)^2} < \frac{50\pi}{9}\varepsilon^2.$$

Finally, since Σ is covered by n such neighborhoods, $\mu(\Sigma) = O(n\varepsilon^2)$. □

We say that an ε-sample is *parsimonious* if it contains $O(\mu(\Sigma)/\varepsilon^2)$ points, that is, only a constant factor more than the minimum possible number required by Lemma 3.1.

3.2 Oversampling Is Bad

The easiest method to produce a surface sample with high Delaunay complexity is *oversampling*, where some region of the surface contains many more points than necessary. In fact, the only surface where oversampling cannot produce a quadratic-complexity Delaunay triangulation is the sphere, even if we only consider parsimonious samples.

The idea behind our construction is to find two skew (*i.e.*, not coplanar) lines tangent to the surface, place points on these lines in small neighborhoods of the tangent points to create a complete bipartite Delaunay graph, and then perturb the points onto the surface. The neighborhoods must be sufficiently small that the perturbation does not change the Delaunay structure. Also, the original tangent lines must be positioned so that the resulting Delaunay circumspheres are small,

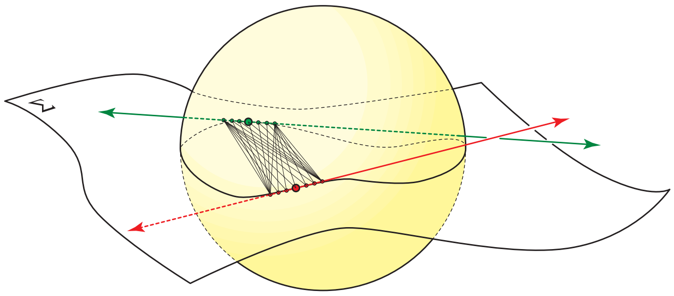

Figure 6. Parsimoniously oversampling a non-spherical surface.

so that we can uniformly sample the rest of the surface without destroying the local quadratic structure. See Figure 6.

To quantify our construction, we first establish a technical lemma about perturbations. Let $P = \{p_1, p_2, \ldots, p_m\}$ and $Q = \{q_1, q_2, \ldots, q_m\}$ be sets of $m \geq 4$ evenly spaced points on two skew lines ℓ_P and ℓ_Q. Every segment p_iq_j is an edge in the Delaunay triangulation of $P \cup Q$. An r-*perturbation* of $P \cup Q$ is a set $\tilde{P} \cup \tilde{Q} = \{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_m, \tilde{q}_1, \tilde{q}_2, \ldots, \tilde{q}_m\}$ such that $|p_i\tilde{p}_i| < r$ and $|q_j\tilde{q}_j| < r$ for all $1 \leq i, j \leq n$. Continuity arguments imply that if r is sufficiently small, the Delaunay triangulation of $\tilde{P} \cup \tilde{Q}$ also has quadratic complexity.

Let δ denote the distance between successive points in both P and Q. For each point $p_i \in P$ and $q_j \in Q$, let β_{ij} denote the empty ball whose boundary is tangent to ℓ_P at p_i and tangent to ℓ_Q at q_j, and let ρ denote the largest radius of any bitangent sphere β_{ij}.

Lemma 3.2. *Let $\tilde{P} \cup \tilde{Q}$ be an arbitrary r-perturbation of $P \cup Q$, where $r < \delta^2/9\rho$. Every pair of points $\tilde{p}_i$ and $\tilde{q}_j$ lies on an empty sphere with radius at most 2ρ and thus are neighbors in the Delaunay triangulation of $\tilde{P} \cup \tilde{Q}$.*

Proof: Since $m \geq 4$, we observe that $\rho > \delta$. The distance between any point p_i to any bitangent sphere β_{kj} with $k \neq i$ is at least $\sqrt{\delta^2 + \rho^2} - \rho \geq \delta^2/3\rho$. Let $\tilde{\beta}_{ij}$ be the ball concentric with β_{ij} with radius $\delta^2/9\rho$ larger than β_{ij}. This ball contains $\tilde{p}_i$ and $\tilde{q}_j$ but excludes every other point in $\tilde{P} \cup \tilde{Q}$, and thus contains an empty circumsphere of $\tilde{p}_i\tilde{q}_j$ with radius at most $\rho + \delta^2/9\rho < 2\rho$. □

Theorem 3.3. *For any non-spherical C^2 surface Σ and any $\varepsilon > 0$, there is a parsimonious ε-sample of Σ whose Delaunay triangulation has complexity $\Omega(n^2)$, where n is the number of sample points.*

Proof: Let S be any parsimonious ε-sample of Σ, and let $m = |S|$. Let σ be a sphere, centered at a point $x \in \Sigma$, with radius $\rho < \text{lfs}(x)/36m^2$, such that every point in S has distance at least 6ρ from σ, and the intersection curve $\gamma = \sigma \cap \Sigma$ is not planar (*i.e.*, not a circle). Such a sphere always exists unless Σ is itself a sphere; for example, we could take x to be any point whose principal curvatures are different.

Let ℓ_p and ℓ_q be skew lines tangent to γ at points p and q, respectively; these lines must exist since γ does not lie in a plane. Let $P = \{p_1, p_2, \ldots, p_m\}$ and $Q = q_1, q_2, \ldots, q_m$ be sets of evenly spaced points on ℓ_p and ℓ_q, respectively, in sufficiently small neighborhoods of p and q that every bitangent sphere β_{ij} (see above) has radius less than 2ρ. Such neighborhoods exist by continuity

[3] We can obtain slightly better constants using the fact that every surface point lies in the neighborhood of its closest sample point; see [1, Lemma 3].

arguments. Let δ denote the distance between successive points in P and Q and assume without loss of generality that $m\delta < \rho/2$. Finally, define $\tilde{P} \cup \tilde{Q} = \{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_m, \tilde{q}_1, \tilde{q}_2, \ldots, \tilde{q}_m\} \subset \Sigma$ where $\tilde{p}_i$ and $\tilde{q}_j$ are the surface points closest to p_i and q_j, respectively.

Without loss of generality, $P \cup Q \cup \tilde{P} \cup \tilde{Q}$ lies in a ball $\tilde{\sigma}$ of radius 2ρ centered at x. Lipschitz arguments imply that $\Sigma \cap \tilde{\sigma}$ lies between two balls of radius $R = \mathrm{lfs}(x) - 4\rho$ tangent to Σ at p (or at q). Thus, any point in P (or in Q) has distance at most

$$\sqrt{(m\delta)^2 + R^2} - R < \frac{(m\delta)^2}{R} < \frac{2(m\delta)^2}{\mathrm{lfs}(x)} < \frac{\delta^2}{18\rho}.$$

from Σ, since $\rho < \mathrm{lfs}(x)/36m^2$. Lemma 3.2 now implies that for all i, j, there is a ball of radius less than 4ρ that contains $\tilde{p}_i$ and $\tilde{q}_j$ and excludes every other point in $\tilde{P} \cup \tilde{Q}$. Since every other point in S has distance at least 6ρ away from σ, this ball also excludes every point in S.

We conclude that $S \cup \tilde{P} \cup \tilde{Q}$ is a parsimonious ε-sample of Σ consisting of $n = 3m$ points whose Delaunay triangulation has complexity $\Omega(n^2)$. □

The reconstruction algorithm of Amenta *et al.* [4] extracts a surface from a subset of the Delaunay triangles of the sample points. Their algorithm estimates the surface normal at each sample point p using the Voronoi diagram of the samples. The *cocone* at p is the complement of a very wide double cone whose apex is p and whose axis is the estimated normal vector at p. The algorithm extracts the Delaunay triangles whose dual Delaunay edges intersect the cocones of all three of its vertices, and then extracts a manifold surface from those cocone triangles. Usually only a small subset of the Delaunay triangles pass this filtering phase, but our construction shows that there are $\Omega(n^2)$ cocone triangles in the worst case.

3.3 Uniform Sampling Can Still Be Bad

Unfortunately, oversampling is not the only way to obtain quadratic Delaunay triangulations. Let S be a set of sample points on the surface Σ. We say that S is a *uniform ε-sample* of Σ if the distance form any point $x \in \Sigma$ to its second-closest sample point is between $(\varepsilon/c)\,\mathrm{lfs}(x)$ and $\varepsilon\,\mathrm{lfs}(x)$, for some constant $c > 2$.[4] We easily verify that a uniform ε-sample is in fact an ε-sample. A packing argument similar to the proof of Lemma 3.1 implies that uniform ε-samples are parsimonious.

Lemma 3.4. *For any n and $\varepsilon > \sqrt{1/n}$, there is a two-component surface Σ and an n-point uniform ε-sample S of Σ, such that the Delaunay triangulation of S has complexity $\Omega(n^2\varepsilon^2)$.*

Proof: The surface Σ is the boundary of two *sausages* Σ_x and Σ_y, each of which is the Minkowski sum of a unit sphere and a line segment. Specifically, let

$$\Sigma_x = U + \overline{(-w, 0, d+1), (w, 0, d+1)} \quad \text{and}$$
$$\Sigma_y = U + \overline{(0, -w, -d-1), (0, w, -d-1)},$$

where U is the unit ball centered at the origin, $w = n\varepsilon^2$, and $d = 4w/\varepsilon = 4n\varepsilon$. The local feature size of every point on Σ is 1, so any uniform ε-sample of Σ has $\Theta((w+1)/\varepsilon^2) = \Theta(n)$ points.

[4] Equivalently, following Dey *et al.* [22], we could define an ε-sample to be uniform if $|pq|/\mathrm{lfs}(p) \geq \varepsilon/c$ for any sample points p and q, for some constant $c > 1$.

Define the *seams* σ_x and σ_y as the maximal line segments in each sausage closest to the xy-plane:

$$\sigma_x = \overline{(-w, 0, d), (w, 0, d)} \qquad \text{and}$$
$$\sigma_y = \overline{(0, -w, -d), (0, w, -d)}.$$

Our uniform ε-sample S contains $2w/\varepsilon + 1$ points along each seam:

$$p_i = (i\varepsilon, 0, d) \quad \text{for all integers } -w/\varepsilon \le i \le w/\varepsilon, \text{ and}$$
$$q_j = (0, j\varepsilon, -d) \quad \text{for all integers } -w/\varepsilon \le j \le w/\varepsilon.$$

The Delaunay triangulation of these $\Theta(w/\varepsilon) = \Theta(n\varepsilon)$ points has complexity $\Theta(w^2/\varepsilon^2) = \Theta(n^2\varepsilon^2)$.

Let β_{ij} be the ball whose boundary passes through p_i and q_j and is tangent to both seams. The intersection $\Sigma_x \cap \beta_{ij}$ is a small oval, tangent to σ_x at p_i and symmetric about the plane $x = i\varepsilon$. See Figure 7(a).

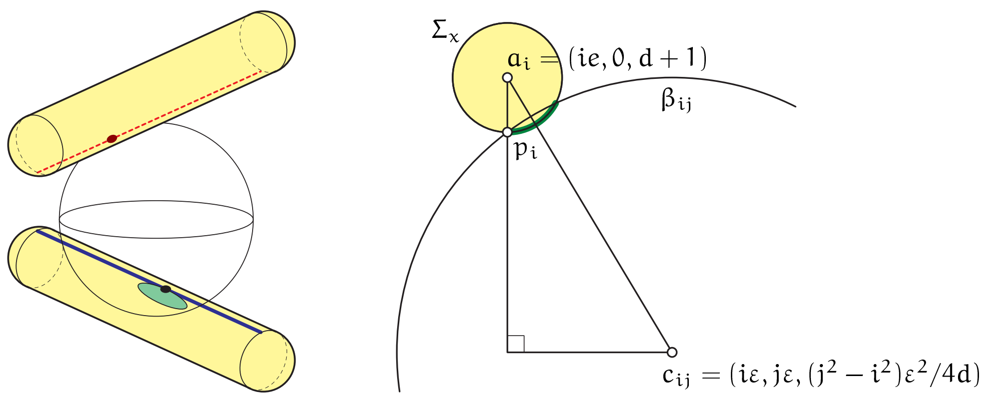

(a) (b)

Figure 7. (a) Two sausages and a sphere tangent to both seams. (b) Computing the width of the intersection oval.

We claim that this oval lies in a sufficiently thin strip around the seam of Σ_x that we can avoid it with the other sample points in S. We compute the width of the oval by considering the intersection of Σ_x and β_{ij} with the plane $x = i\varepsilon$. Simple calculations imply that β_{ij} is centered at the point $c_{ij} = (i\varepsilon, j\varepsilon, (j^2 - i^2)\varepsilon^2/4d)$. Let $a_i = (i\varepsilon, 0, d+1)$. The width of the intersection oval,measured along the surface of Σ_x, is exactly $2\angle p_i a_i c_{ij}$. From Figure 7(b), we see that

$$\tan \angle p_i a_i c_{ij} = \frac{j\varepsilon}{(d+1) - (j^2 - i^2)\varepsilon^2/4d}.$$

Thus, we can bound the width of the oval as follows:

$$\begin{aligned} 2\angle p_i a_i c_{ij} &= 2\tan^{-1}\left(\frac{4dj\varepsilon}{4d(d+1) + (i^2 - j^2)\varepsilon^2}\right) \\ &< \frac{8dj\varepsilon}{4d^2 + (i^2 - j^2)\varepsilon^2} \\ &< \frac{8dw}{4d^2 - 2w^2} \\ &< \frac{8w}{d} = \varepsilon. \end{aligned}$$

We conclude that $\Sigma_x \cap \beta_{ij}$ lies entirely within a strip of width less than 2ε centered along the seam σ_x. A symmetric argument gives the analogous result for $\Sigma_y \cap \beta_{ij}$. We can uniformly sample Σ so that no sample point lies within either strip except the points we have already placed along the seams. Each segment $p_i q_j$ is an edge in the Delaunay triangulation of the sample, and there are $\Omega(w^2/\varepsilon^2) = \Omega(n^2\varepsilon^2)$ such segments. □

Theorem 3.5. *For any n and any $\varepsilon > \sqrt{1/n}$, there is a connected surface Σ and an n-point uniform ε-sample S of Σ, such that the Delaunay triangulation of S has complexity $\Omega(n^2\varepsilon^2)$.*

Proof: Intuitively, we produce the surface Σ by pushing two sausages into a spherical balloon. These sausages create a pair of conical *wedges* inside the balloon whose *seams* lie along two skew lines. The local feature size is small near the seams and drops off quickly elsewhere, so a large fraction of the points in any uniform sample must lie near the seams. We construct a particular sample with points *exactly* along the seams that form a quadratic-complexity triangulation, similarly to our earlier sausage construction. Our construction relies on several parameters: the radius R of the spherical balloon, the width w and height h of the wedges, and the distance d between the seams.

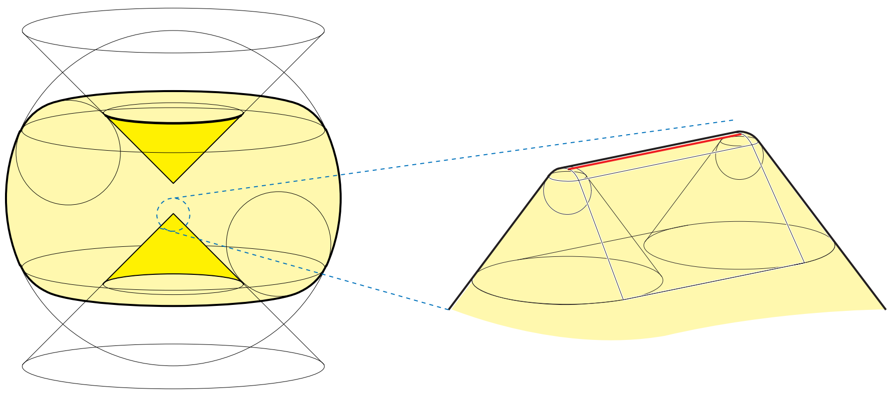

Figure 8. A smooth surface with a bad uniform ε-sample, and a closeup of one of its wedges.

Each wedge is the Minkowski sum of a unit sphere, a right circular cone with height h centered along the z-axis, and a line segment of length w parallel to one of the other coordinate axes. A wedge can be decomposed into cylindrical, spherical, conical, and planar facets. The cylindrical and spherical facets constitute the *blade* of the wedge, and the *seam* of the blade is the line segment of length w that bisects the cylindrical facet. The local feature size of any point on the blade is exactly 1, and the local feature size of any other wedge point is 1 plus its distance from the blade. Straightforward calculations imply that the sample measure of each wedge is $O(w + \log h + 1)$.

A first approximation $\widetilde{\Sigma}$ of the surface Σ is obtained by removing two wedges from a ball of radius R centered at the origin. One wedge points into the ball from below; its seam is parallel to the x-axis and is centered at the point $(0, 0, -R + h)$. The other wedge points into the ball from above; its seam is parallel to the y-axis and is centered at $(0, 0, R - h)$. Let $d = 2R - 2h - 2$ denote the distance between the wedges. Our construction has $1 \ll w \ll d \ll h$, so $R < 3h$.

To obtain the final smooth surface Σ, we round off the sharp edges by rolling a ball of radius $h/4$ inside $\widetilde{\Sigma}$ along the wedge/balloon intersection curves. We call the resulting warped toroidal patches the *sleeves*. The local feature size of any point on the sleeves or on the balloon is at least $h/4$. The surface area of the sleeves is $O(R^2) = O(h^2)$, so the sleeves have constant sample measure.

The local feature size of other surface points changes only far from the blades and by only a small constant factor. Thus, $\mu(\Sigma) = \Theta(w + \log h + 1)$. To complete the construction, we set $w = n\varepsilon^2$, $d = 4n\varepsilon$, and $h = 20n\varepsilon$. See Figure 8.

Finally, we construct a uniform ε-sample S with $\Theta(w/\varepsilon)$ sample points evenly spaced along each seam and every other point at least ε away from the seams. Setting $h > 5d$ (and thus $R > 10d$) ensures that the Delaunay spheres β_{ij} between seam points do not touch Σ except on the blades. By the argument in Lemma 3.4, there are $\Omega(w^2/\varepsilon^2) = \Omega(n^2\varepsilon^2)$ Delaunay edges between seam points. □

3.4 Some Surfaces Are Just Evil

In this section, we describe a family of surfaces for which *any* parsimonious ε-sample has a Delaunay triangulation of near-quadratic complexity. First we give a nearly trivial construction of a bad surface with several components, and then we join these components into a single connected surface using a method similar to Theorem 3.5.

Lemma 3.6. *For any n and any $\varepsilon > \sqrt{1/n}$, there is a smooth surface Σ such that the Delaunay triangulation of any parsimonious ε-sample of Σ has complexity $\Omega(n^2\varepsilon^4)$, where n is the size of the sample.*

Proof: Let P be a set containing the following k points:

$$\begin{aligned} p_i &= (ik, 0, k^2) && \text{for all integers } -k/4 \le i < k/4, \text{ and} \\ q_j &= (0, jk, -k^2) && \text{for all integers } -k/4 \le j < k/4. \end{aligned}$$

We easily verify that every pair of points p_i and q_j lie on the boundary of a ball β_{ij} with every other point in P at least unit distance outside. (See Lemma 3.2.)

Let $\Sigma = \bigcup_{p \in P} U_p$, where U_p is the unit-radius sphere centered at p. Clearly, $\text{lfs}(x) = 1$ for every point $x \in \Sigma$, so $\mu(\Sigma) = 4\pi k$. Let S be an arbitrary parsimonious ε-sample of Σ, let $n = |S| = \Theta(k/\varepsilon^2)$, and for any point $p \in P$, let $S_p = S \cap U_p$ be the sample points on its unit sphere.

Choose an arbitrary pair $p_i, q_j \in P$. By construction, β_{ij} contains only points in S_{p_i} and S_{q_j}. Shrink β_{ij} about its center until (without loss of generality) it has no points of S_{p_i} in its interior. Choose some point $p' \in S_{p_i}$ on the boundary of the shrunken ball, and then shrink the ball further about p' until it contains no point of S_{q_i}. The resulting ball has p' and some $q' \in S_{q_i}$ on its boundary, and no points of S in its interior. Thus, p' and q' are neighbors in the Delaunay triangulation of S. There are at least $\Omega(k^2) = \Omega(n^2\varepsilon^4)$ such pairs. □

To create a *connected* surface where good sample has a complicated Delaunay triangulation, we add 'teeth' to our earlier balloon and wedge construction. Unfortunately, in the process, we lose a polylogarithmic factor in the Delaunay complexity.

Theorem 3.7. *For any n and any $\varepsilon < \sqrt{1/n}$, there is a smooth connected surface Σ such that the Delaunay triangulation of any parsimonious ε-sample of Σ has complexity $\Omega(n^2\varepsilon^4/\log^2(n\varepsilon^2))$, where n is the size of the sample.*

Proof: As in Theorem 3.5, our surface Σ contains two wedges, but now each wedge has a row of small conical teeth. Our construction relies on the same parameters R, w, h of our earlier construction. We now have additional parameter t, which is simultaneously the height of the teeth, the distance between the teeth, and half the thickness of the 'blade' of the wedge.

Our construction starts with the (toothless) surface described in the proof of Theorem 3.5, but using a ball of radius t instead of a unit ball to define the wedges. We add w/t evenly-spaced teeth along the blade of each wedge, where each tooth is the Minkowski sum of a unit ball with a right circular cone of radius t. Each tooth is tangent to both planar facets of its wedge. To create the final smooth surface Σ, we roll a ball of radius $t/3$ over the blade/tooth intersection curves. The complete surface has sample measure $\Theta((w/t)(1+\log t)+\log h+1)$. Finally, we set the parameters $w = t^2$, $h = t^3$, and $R = 20t^3$, so that $\mu(\Sigma) = \Theta(t \log t)$.

Let S be a parsimonious ε-sample of Σ, and let $n = |S| = \Theta((t \log t)/\varepsilon^2)$. For any pair of teeth, one on each wedge, there is a sphere tangent to the ends of the teeth that has distance $\Omega(1)$ from the rest of the surface. We can expand this sphere so that it passes through one point on each tooth and excludes the rest of the points. Thus, the Delaunay triangulation of S has complexity $\Omega(t^2) = \Omega(n^2\varepsilon^4/\log^2(n\varepsilon^2))$. □

3.5 Randomness Doesn't Help Much

Golin and Na [33] proved that if S is a random set of n points on the surface of a convex polytope with a constant number of facets, then the expected complexity of the Delaunay triangulation of S is $O(n)$. Unfortunately, this result does not extend to nonconvex objects, even when the random distribution of the points is proportional to the sample measure.

Theorem 3.8. *For any n, there is a smooth connected surface Σ, such that the Delaunay triangulation of n independent uniformly-distributed random points in Σ has complexity $\Theta(n^2/\log^2 n)$ with high probability.*

Proof: Consider the surface Σ consisting of $\Theta(n/\log n)$ unit balls evenly spaced along two skew line segments, exactly as in the proof of Theorem 3.6, with extremely thin cylinders joining them into a single connected surface resembling a string of beads. With high probability, a random sample of n points contains at least one point on each ball, on the side facing the opposite segment. Thus, with high probability, there is at least one Delaunay edge between any ball on one segment and any ball on the other segment. □

Theorem 3.9. *For any n, there is a smooth connected surface Σ, such that the Delaunay triangulation of n independent random points in Σ, distributed proportionally to the sample measure, has complexity $\Theta(n^2/\log^4 n)$ with high probability.*

Proof: Let Σ be the surface used to prove Theorem 3.7, but with $\Theta(n/\log^2 n)$ teeth. With high probability, a random sample of Σ contains at least one point at the tip of each tooth. □

4 Conclusions

We have derived new upper and lower bounds on the complexity of Delaunay triangulations under two different geometric constraints: point sets with sublinear spread and good samples of smooth surfaces. Our results imply that with very strong restrictions on the inputs, most existing surface reconstruction algorithms are inefficient in the worst case.

Our results suggest several open problems, the most obvious of which is to tighten the spread-based bounds. We conjecture that our lower bounds are tight. Even the special case of dense point sets is open.

Another natural open problem is to generalize our analysis to higher dimensions. Dey *et al.* [22] describe a generalization of the cocone algorithm [4] that determines the dimension of a uniformly sampled manifold (in the sense of Section 3.3) in a space of any fixed dimension. Results developed in a companion paper [29] imply that for any n and $\Delta \geq \sqrt{n}$, there is a set of n points in $\mathbb{R}^d$ with spread Δ whose Delaunay triangulation has complexity $\Omega(n\Delta^{\lceil d/2 \rceil - 1})$. The techniques used in Section 2.2 generalize easily to prove that any d-dimensional Delaunay triangulation has $O(\Delta^{d+1})$ edges, but this implies a very weak bound on the overall complexity. We conjecture that the complexity is always $O(\Delta^d)$—in particular, $O(n)$ for all dense point sets—and can only reach the maximum $\Omega(n^{\lceil d/2 \rceil})$ when $\Delta = \Omega(n)$.

Our bad surface examples are admittedly quite contrived, since they have areas of very high curvature relative to their diameter. An interesting open problem is whether there are bad surfaces with smaller 'spread', *i.e.*, ratio between diameter and minimum local feature size. What is the worst-case complexity of the Delaunay triangulation of good surface samples as a function of the spread and sample measure of the surface? Is there a single surface such that for any ε, there is a uniform ε-sample with quadratic Delaunay complexity, or (as I conjecture) is the cylinder the worst case? Even worse, is there a "universally bad" surface such that *every* uniform sample of has super-linear Delaunay complexity?

Our surface results imply that most Delaunay-based surface reconstruction algorithms can be forced to take super-linear time, even for very natural surface data. It may be possible to improve these algorithms by adding a small number of Steiner points in a preprocessing phase to reduce the complexity of the Delaunay triangulation. In most of our bad surface examples, a single Steiner point reduces the Delaunay complexity to $O(n)$. Bern, Eppstein, and Gilbert [10] show that any Delaunay triangulation can be reduced to $O(n)$ complexity in $O(n \log n)$ time by adding $O(n)$ Steiner points; see also [15]. Unfortunately, the Steiner points they choose (the vertices of an octtree) may make reconstruction impossible. In order to be usable, any new Steiner points must either lie very close to or very far from the surface, and as our bad examples demonstrate, both types of Steiner points may be necessary. Boissonnat and Cazals (personal communication) report that adding a small subset of the original Voronoi vertices as Steiner points can significantly reduce the complexity of the resulting Voronoi diagram with only minimal changes to the smooth surface constructed by their algorithm [12].

After some of the results in this paper were announced, Dey *et al.* [20] developed a surface reconstruction algorithm that does not construct the entire Delaunay triangulation. Their algorithm runs in $O(n \log n)$ time if the sample is locally uniform, meaning (loosely) that the density of the sample points varies smoothly over the surface, but still requires quadratic time in the worst case. Even more recently, Ramos [38] discovered a fast algorithm to extract a locally uniform sample from any ε-sample, thereby producing a surface reconstruction algorithm that provably runs in $O(n \log n)$ time.

Finally, are there other natural geometric conditions under which the Delaunay triangulation provably has small complexity?

Acknowledgments. I thank Herbert Edelsbrunner for asking the (still open!) question that started this work, Kim Whittlesey for suggesting charging Delaunay features to area, Edgar Ramos for suggesting well-separated pair decompositions, and Tamal Dey and Edgar Ramos for sending me preliminary copies of their papers [20, 22, 38]. Thanks also to Sariel Har-Peled, Olivier Devillers, and Jean-Daniel Boissonnat for helpful discussions.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* 22(4):481–504, 1999.

[2] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graph. Models Image Process.* 60:125–135, 1998.

[3] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Proc. SIGGRAPH '98*, 415–412, 1998.

[4] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 213–222, 2000.

[5] N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. To appear in *Internat. J. Comput. Geom. Appl.* ⟨http://www.cs.utexas.edu/users/amenta/pubs/power.ps.gz⟩.

[6] D. Attali. r-regular shape reconstruction from unorganized points. *Comput. Geom. Theory Appl.* 10:239–247, 1998.

[7] C. Bajaj, F. Bernardini, and G. Xu. Automatic reconstruction of surface and scalar fields from 3d scans. *Comput. Graph.* 29:109–118, 1995. Proc. SIGGRAPH '95.

[8] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, 294–303, 1997.

[9] F. Bernardini and C. L. Bajaj. Sampling and reconstructing manifolds using alpha-shapes. *Proc. 9th Canad. Conf. Comput. Geom.*, 193–198. 1997.

[10] M. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. *J. Comput. Syst. Sci.* 48:384–409, 1994.

[11] J.-D. Boissonnat. Representing 2d and 3d shapes with the Delaunay triangulation. *Proc. 7th IEEE Internat. Conf. Pattern Recogn.*, 745–748, 1984.

[12] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 223–232, 2000.

[13] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM* 42:67–90, 1995.

[14] D. E. Cardoze and L. Schulman. Pattern matching for spatial point sets. *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.*, 156-165, 1998.

[15] B. Chazelle, H. Edelsbrunner, L. Guibas, J. Hershberger, R. Seidel, and M. Sharir. Selecting heavily covered points. *SIAM J. Comput.* 23:1138–1151, 1994.

[16] H.-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete Comput. Geom.* 25(4):525–568, 2001.

[17] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver exudation. *Proc. 15th Annu. Sympos. Comput. Geom.*, 1–13, 1999.

[18] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.* 22:63–93, 1999.

[19] T. K. Dey. Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.* 19(3):373–382, 1998.

[20] T. K. Dey, S. Funke, and E. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. *Proc. 17th European Workshop Comput. Geom.*, 2001.

[21] T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, 257–263, 2001.

[22] T. K. Dey, J. Giesen, S. Goswami, and W. Zhao. Shape dimension from samples. Unpublished manuscript, June 2001.

[23] R. Dwyer. The expected number of k-faces of a Voronoi diagram. *Internat. J. Comput. Math.* 26(5):13–21, 1993.

[24] R. A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.* 6:343–367, 1991.

[25] H. Edelsbrunner, X.-T. Li, G. Miller, A. Stathopolous, D. Talmor, S.-H. Teng, A. Üngör, and N. Wlkington. Smoothing and cleaning up slivers. *Proc. 32nd Annu. ACM Sympos. Theory Comput.*, 273–277, 2000.

[26] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.* 13(1):43–72, 1994.

[27] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica* 15:223–241, 1996.

[28] H. Edelsbrunner, P. Valtr, and E. Welzl. Cutting dense point sets in half. *Discrete Comput. Geom.* 17(3):243-257, 1997.

[29] J. Erickson. Arbitrarily large neighborly families of congruent symmetric convex 3-polytopes. Unpublished manuscript, June 2001. arXiv:math.CO/0106095. ⟨http://www.cs.uiuc.edu/˜jeffe/pubs/crum.html⟩.

[30] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, 96–105, 2001.

[31] M. Gavrilov, P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric pattern matching: A performance study. *Proc. 15th Annu. ACM Symops. Comput. Geom.*, 79–85, 1999.

[32] C. Gold and J. Snoeyink. A one-step crust and skeleton extraction algorithm. *Algorithmica* 30(2):144–163, 2001.

[33] M. Golin and H. Na. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. *Proc. 12th Canadian Conf. Comput. Geom.*, 127–135, 2000. ⟨http://www.cs.unb.ca/conf/cccg/eProceedings/44.ps.gz⟩.

[34] B. Guo, J. Menon, and B. Willette. Surface reconstruction using alpha shapes. *Comput. Graph. Forum* 16(4):177–190, 1997.

[35] H. Hiyoshi and K. Sugihara. Voronoi-based interpolation with higher continuity. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 242–250, 2000.

[36] P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric matching under noise: Combinatorial bounds and algorithms. Proc. 8th Annu. ACM-SIAM Sumpos. Discrete Algorithms, 457–465, 1999.

[37] X.-Y. Li and S.-H. Teng. Generating sliver-free three dimensional meshes. *Proc. 12th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 2001.

[38] E. Ramos. Smooth surface reconstruction in near linear time. Unpublished manuscript, May 2001.

[39] J. Ruppert and R. Seidel. A Delaunay refinement algorithm for quality 2-dimensional mesh generation *J. Algorithms* 18:548–585, 1995.

[40] J. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 86–95, 1998.

[41] M. Sharir, S. Smorodinsky, and G. Tardos. An improved bound for k-sets in three dimensions. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 43–49, 2000. *Discrete Comput. Geom.*, 2001, in press.

[42] P. Valtr. Convex independent sets and 7-holes in restricted planar point sets. *Discrete Comput. Geom.* 7:135–152, 1992.

[43] P. Valtr. Lines, line-point incidences and crossing families in dense sets. *Combinatorica* 16:269–294, 1996.

[44] P. Valtr. *Planar Point Sets with Bounded Ratios of Distances*. Doctoral thesis, Fachbereich Mathematik, Freie Universität Berlin, 1994.

[45] K. Verbarg. Approximate center points in dense point sets. *Inform. Proc. Letters* 61(5):271–278, 1997.

[46] J. Vleugels. *On Fatness and Fitness: Realistic Input Models for Geometric Algorithms*. Ph.D. Thesis, Dept. Comput. Sci., Univ. Utrecht, 1997.