

# Dense Point Sets Have Sparse Delaunay Triangulations\*

## or "... But Not Too Nasty"

Jeff Erickson

University of Illinois, Urbana-Champaign

jeffe@cs.uiuc.edu

<http://www.cs.uiuc.edu/~jeffe>

### Abstract

The *spread* of a finite set of points is the ratio between the longest and shortest pairwise distances. We prove that the Delaunay triangulation of any set of  $n$  points in  $\mathbb{R}^3$  with spread  $\Delta$  has complexity  $O(\Delta^3)$ . This bound is tight in the worst case for all  $\Delta = O(\sqrt{n})$ . In particular, the Delaunay triangulation of any dense point set has linear complexity. On the other hand, for any  $n$  and  $\Delta = O(n)$ , we construct a regular triangulation of complexity  $\Omega(n\Delta)$  whose  $n$  vertices have spread  $\Delta$ .

### 1 Introduction

Delaunay triangulations and Voronoi diagrams are one of the most thoroughly studied objects in computational geometry, with applications to nearest-neighbor searching [3, 21, 25, 43], clustering [1, 46, 48, 54], finite-element mesh generation [20, 32, 49, 56], deformable surface modeling [19], and surface reconstruction [4, 5, 6, 7, 12, 45]. Many algorithms in these application domains begin by constructing the Delaunay triangulation or Voronoi diagram of a set of points in  $\mathbb{R}^3$ . Since three-dimensional Delaunay triangulations can have complexity  $\Omega(n^2)$  in the worst case, 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.

This frustrating discrepancy between theory and practice motivates our investigation of practical geometric constraints that imply low-complexity Delaunay triangulations. Previous works in this direction have studied *random* point sets under various distributions [28, 27, 35, 38]; *well-spaced* point sets, which are low-discrepancy samples of Lipschitz density functions [20, 49, 51, 52]; and *surface samples* with various density

constraints [8, 35]. (We will discuss the connections between these models and our results shortly.) Our efforts fall under the rubric of *realistic input models*, which have been primarily studied for inputs consisting of polygons or polyhedra [10, 63].

This paper investigates the complexity of three-dimensional Delaunay triangulations in terms of a geometric parameter called the *spread*, continuing our work in an earlier paper [35]. The spread of a set of points is the ratio between the largest and smallest interpoint distances. Of particular interest are *dense* point sets in  $\mathbb{R}^d$ , which have spread  $O(n^{1/d})$ . Valtr and others [33, 59, 60, 61, 62] have established several combinatorial results for dense point sets that improve corresponding bounds for arbitrary point sets. For other combinatorial and algorithmic results related to spread, see [15, 23, 37, 41, 42, 47].

In Section 2, we prove that the Delaunay triangulation of any set of  $n$  points in  $\mathbb{R}^3$  with spread  $\Delta$  has complexity  $O(\Delta^3)$ . In particular, the Delaunay triangulation of any dense point set in  $\mathbb{R}^3$  has only linear complexity. This bound is tight in the worst case for all  $\Delta = O(\sqrt{n})$  and improves an earlier upper bound of  $O(\Delta^4)$  [35]. We briefly mention some implications of our new upper bound in Section 3.

In Section 4, we show that our upper bound does not generalize to non-Delaunay triangulations. For any  $n$  and  $\Delta \leq n$ , we construct a set of  $n$  points with spread  $\Delta$  with a *regular* triangulation of complexity  $\Omega(n\Delta)$ . This worst-case lower bound was already known for Delaunay triangulations for all  $\sqrt{n} \leq \Delta \leq n$  [35]. In particular, there is a dense point set in  $\mathbb{R}^3$ , arbitrarily close to a cubical lattice, with a regular triangulation of complexity  $\Omega(n^{4/3})$ .

We conclude in Section 5 by suggesting several open problems.

Throughout the paper, we analyze the complexity of three-dimensional Delaunay triangulations by counting their edges. Since the link of every vertex in a three-dimensional triangulation is a planar graph, Euler's formula implies that any triangulation with  $n$  vertices

---

\*Portions of this work were done while the author was visiting The Ohio State University. This research was partially supported by a Sloan Fellowship and by NSF CAREER grant CCR-0093348. See <http://www.cs.uiuc.edu/~jeffe/pubs/screw.html> for the most recent version of this paper.

and  $e$  edges has at most  $2e - 2n$  triangles and  $e - n$  tetrahedra. 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.

## 1.1 Related Results

Our results compare favorably with several other types of point sets: points with small integer coordinates, random points, well-spaced points, and surface samples. First, we easily observe that any triangulation of an integer point set with coordinates between 1 and  $\Delta$  has complexity  $O(\Delta^3)$ , since each tetrahedron has volume at least  $1/6$ . (It is open whether this bound is tight for all  $n$  and  $\Delta$ .)

Dwyer [28, 27] showed that if a set of  $n$  points is generated uniformly at random from the unit ball, its Delaunay triangulation has complexity  $O(n)$  with high probability. Golin and Na [38, 39, 40] recently proved that if  $n$  points are chosen uniformly at random on the surface of any fixed<sup>1</sup> three-dimensional convex polytope, the expected complexity of their Delaunay triangulation is  $O(n)$ . The spread of a uniform set of points is  $\Omega(n)$  with high probability.

Miller *et al.* [51, 52] define a point set  $S$  in  $\mathbb{R}^d$  to be *well-spaced* with respect to a 1-Lipschitz spacing function  $f : \mathbb{R}^d \rightarrow \mathbb{R}^+$  if, for some fixed constants  $\beta > 1$  and  $\varepsilon > 0$ , the distance between any two points  $p, q \in S$  is at least  $(f(p) + f(q))/\beta$ , and the distance between any point  $x \in \mathbb{R}^d$  and its nearest neighbor  $p \in S$  is at most  $\varepsilon f(p)$ . (A nearly-equivalent notion of *locally uniform samples* of a smooth surface is used by Dey *et al.* [26] and Funke and Ramos [36]; see also [35].) Miller *et al.* show that Delaunay triangulations of well-spaced point sets have complexity  $O(n)$ ; in particular, any point in a well-spaced point set has  $O(1)$  Delaunay neighbors. Our results are formally incomparable with those of Miller *et al.* On the one hand, well-spaced point sets can have arbitrarily high spread. On the other hand, dense point sets are not necessarily well-spaced with respect to *any* spacing function—for example, consider a cubical lattice, minus the points in a large ball. Moreover, our upper bound degrades gracefully as the spread increases, while the upper bounds of Miller *et al.* depend exponentially on the spacing constants  $\beta$  and  $\varepsilon$ .

Very recently, Attali and Boissonnat [8] proved that for any fixed polyhedral surface  $P$ , any so-called “light uniform  $\varepsilon$ -sample” of  $P$  has only  $O(n^{7/4})$  Delaunay


edges, or  $O(n^{3/2})$  if the surface is convex, where  $n$  is the number of sample points. Our new upper bound has a similar corollary. Informally, a uniform sample of any fixed (*not* necessarily polyhedral or convex) surface has spread  $O(\sqrt{n})$ , so its Delaunay triangulation has complexity  $O(n^{3/2})$ . We describe this result more formally in Section 3.

## 2 Sparse Delaunay Triangulations

In this section, we prove the main result of the paper.

**Theorem 2.1.** *The Delaunay triangulation of any set of  $n$  points in  $\mathbb{R}^3$  with spread  $\Delta$  has complexity  $O(\Delta^3)$ .*

Our proof is structured as follows. We will implicitly assume that no two points are closer than unit distance apart, so that spread is synonymous with diameter. Two sets  $P$  and  $Q$  are *well-separated* if each set fits in a ball of radius  $r$ , and these two balls are separated by distance  $2r$ . Without loss of generality, we assume that the balls containing  $P$  and  $Q$  are centered at points  $(0, 0, 2r)$  and  $(0, 0, -2r)$ , respectively. Our argument ultimately reduces to counting the number of *crossing edges*—edges in the Delaunay triangulation of  $P \cup Q$  with one endpoint in each set. See Figure 1.



**Figure 1.** A well-separated pair of sets  $P \cup Q$  and a crossing edge intersecting a pixel.

Our proof has four major steps, each presented in its own subsection.

- We place a grid of  $O(r^2)$  circular *pixels* of constant radius  $\varepsilon$  on the plane  $z = 0$ , so that every crossing edge passes through a pixel. In Section 2.1, we prove that all the crossing edges stabbing any single pixel lie within a slab of constant width between two parallel planes. Our proof relies on the fact that the edges of a Delaunay triangulation have a consistent depth order from any viewpoint.
- We say that a crossing edge is *relaxed* if its endpoints lie on an empty sphere of radius  $O(r)$ . In Section 2.2, we show that at most  $O(r)$  relaxed edges pass through any pixel, using a generalization of the ‘Swiss cheese’ packing argument used to

<sup>1</sup>In upper bounds for points on *fixed* surfaces, hidden constants depend on geometric parameters of the fixed surface. For example, Golin and Na’s bound depends on the number of facets and the minimum face and dihedral angles of the polytope. Since the surface is fixed, all such parameters are considered constants.

prove our earlier  $O(\Delta^4)$  upper bound [35]. This implies that there are  $O(r^3)$  relaxed crossing edges overall.

- In Section 2.3, we show that there are a constant number of conformal (*i.e.*, sphere-preserving) transformations that change the spread of  $P \cup Q$  by at most a constant factor, such that every crossing edge of  $P \cup Q$  is a relaxed Delaunay edge in at least one image. It follows that  $P \cup Q$  has at most  $O(r^3)$  crossing edges.
- Finally, in Section 2.4, we count the Delaunay edges for an arbitrary point set  $S$  using an octree-based well-separated pair decomposition [14]. Every edge in the Delaunay triangulation of  $S$  is a crossing edge of some subset pair in the decomposition. However, not every crossing edge is a Delaunay edge; a subset pair contributes a Delaunay edge only if it is close to a large empty *witness* ball. We charge the pair's  $O(r^3)$  crossing edges to the  $\Omega(r^3)$  volume of this ball. We choose the witness balls so that any unit of volume is charged at most a constant number of times, implying the final  $O(\Delta^3)$  bound.

## 2.1 Nearly Concurrent Implies Nearly Coplanar

The first step in our proof is to show that the crossing edges intersecting any pixel are nearly coplanar. To do this, we use an important fact about *depth orders* of Delaunay triangulations, related to shellings of convex polytopes.

Let  $x$  be a point in  $\mathbb{R}^3$ , called the *viewpoint*, and let  $S$  be a set of line segments (or other convex objects). A segment  $s \in S$  is *behind* another segment  $t \in S$  with respect to  $x$  if  $t$  intersects  $\text{conv}\{x, s\}$ . If the transitive closure of this relation is a partial order, any linear extension is called a *consistent depth order* of  $S$  with respect to  $x$ . Otherwise,  $S$  contains a *depth cycle*—a sequence of segments  $s_1, s_2, \dots, s_k$  such that every segment  $s_i$  is directly behind its successor  $s_{i+1}$  and  $s_k$  is directly behind  $s_1$ . De Berg *et al.* [11] describe an algorithm to either compute a depth order for a set of segments or find a depth cycle, in  $O(n^{4/3+\epsilon})$  time. See [9, 18] for related results.

We say that three line segments form a *screw* if they form a depth cycle from some viewpoint. See Figure 2(a).

**Lemma 2.2.** *The edges of any Delaunay triangulation have a consistent depth order from any viewpoint. In particular, no three Delaunay edges form a screw.*




Figure 2. (a) A screw. (b) Front view of  $C$ , showing viewpoints where  $s_x$  appears behind  $s_z$ . (c) Every vertex of  $C$  sees a different depth order, two of which are inconsistent. See Lemma 2.3.

**Proof:** Edelsbrunner [29, 31] proved that a consistent depth order for the simplices in a Delaunay triangulation, with respect to any viewpoint  $x$ , can be obtained by sorting the distances of their circumcenters from  $x$ . This is precisely the order in which the Delaunay tetrahedra are computed by Seidel's shelling convex hull algorithm [55]. We can easily extract a consistent depth order for the Delaunay edges from this simplex order.  $\square$

The next lemma describes sufficient (but not necessary) conditions for three mutually orthogonal segments to form a screw.

**Lemma 2.3.** *Let  $c$  and  $C$  be concentric axis-aligned cubes of width 1 and  $w$ , respectively, for some  $w \geq 5$ . Three line segments, each parallel to a different coordinate axis, form a screw if they all intersect  $c$  but none of their endpoints lie inside  $C$ .*

**Proof:** Let  $s_x, s_y, s_z$  be the three segments, parallel to the  $x$ -,  $y$ -, and  $z$ -axes, respectively. Any ordered pair of these segments, say  $(s_z, s_x)$ , define an unbounded polyhedral region  $V_{xz}$  of viewpoints from which  $s_x$

appears behind  $s_z$ . The segment  $s_z$  is the only bounded edge of  $V_{xz}$ , and both of its endpoints are outside  $C$ . Thus, we can determine which vertices of  $C$  lie inside  $V_{xz}$  by considering the projection to the  $xy$ -plane. From Figure 2(b), we observe that if  $\omega \geq 2 + \sqrt{5} \approx 4.2361$ , then  $V_{xz}$  contains exactly half of the vertices of  $C$ , all with the same  $y$ -coordinate. A symmetric argument implies that the other four vertices lie in  $V_{zx}$ . Similarly,  $V_{xy}$  and  $V_{yx}$  partition the vertices of  $C$  by their  $z$ -coordinates, and  $V_{yz}$  and  $V_{zy}$  partition the vertices of  $C$  along their  $x$ -coordinates. Thus, each of the eight vertices of  $C$  sees one of the eight possible depth orders of the three segments. Since only six of these orders are consistent, two vertices of  $C$  see a depth cycle, implying that the segments form a screw. See Figure 2(c).  $\square$

Recall that a pixel is a circle of radius  $\varepsilon$  in the  $xy$ -plane.

**Lemma 2.4.** *The crossing edges passing through any pixel lie inside a slab of width  $\omega = 44\varepsilon$  between two parallel planes.*

**Proof:** Suppose to the contrary that three Delaunay edges  $e_1$ ,  $e_2$ , and  $e_3$  pass through a common pixel  $\pi$  but are not contained in any slab of width  $\omega$ . Without loss of generality, assume that the pixel  $\pi$  is centered at the origin.

Translate each edge  $e_i$  parallel to the  $xy$ -plane so that it passes through the origin, and call the resulting segment  $\tilde{e}_i$ . Since each segment moved a distance of at most  $\varepsilon/2$ , the segments  $\tilde{e}_i$  are not contained in any slab of width  $\omega - \varepsilon$ . For each  $i$ , let  $s_i$  be the intersection of  $\tilde{e}_i$  with the slab  $-r \leq z \leq r$ , and let  $p_i$  and  $q_i$  be its endpoints. The segments  $s_i$  are not contained in any slab of width  $(\omega - \varepsilon)/3$ .

Any slab  $\sigma$  containing the segments  $s_i$  intersects the plane  $z = r$  in a two-dimensional slab  $\sigma_p$  containing the triangle  $\Delta p_1 p_2 p_3$ . Symmetry arguments imply that if  $\sigma$  is the thinnest slab containing every segment  $s_i$ , then  $\sigma_p$  is the thinnest slab containing  $\Delta p_1 p_2 p_3$ . Thus,  $\Delta p_1 p_2 p_3$  has width at least  $(\omega - \varepsilon)/3$ , and therefore contains a circle of radius at least  $(\omega - \varepsilon)/9$ .

Let  $\mathcal{C}$  denote the Minkowski sum  $s_1 + s_2 + s_3$ .  $\mathcal{C}$  is a rhomboid—a hexahedron combinatorially equivalent to a cube, each of whose facets is a rhombus—centered at the origin, with edges parallel to the original crossing edges  $e_i$ . Since the projection of each segment  $s_i$  to the  $z$ -axis has length  $2r$ , the projection of  $\mathcal{C}$  to the  $z$ -axis has length  $6r$ . If we scale  $\mathcal{C}$  by a factor of  $1/3$ , the resulting rhomboid  $C$  fits exactly within the slab  $-r \leq z \leq r$ . The intersection of  $C$  with the  $xy$ -plane contains a circle of radius at least  $(\omega - \varepsilon)/9$ .

Let  $c$  be the smallest rhomboid homothetic and concentric with  $C$  that contains the pixel  $\pi$ . Since  $\pi$  is a circle of radius  $\varepsilon$  in the  $xy$ -plane,  $C$  is at least a factor of  $(\omega + \varepsilon)/9\varepsilon = 5$  larger than  $c$ . Thus, an appropriate linear transformation maps the edges  $e_i$  and the rhomboids  $c$  and  $C$  to segments and axis-aligned cubes satisfying the conditions of Lemma 2.3. Since screws are invariant under linear transformations, this implies that the Delaunay edges  $e_1, e_2, e_3$  form a screw, which by Lemma 2.2 is impossible.  $\square$

## 2.2 Slabs Contain Few Relaxed Edges

At this point, we would like to argue that any slab of constant width contains only  $O(r)$  crossing edges. Unfortunately, this is not true—a variant of our helix construction [35] implies that a slab can contain up to  $\Omega(r^3)$  edges,  $\Omega(r^2)$  of which can pass through a single, arbitrarily small pixel. However, most of these Delaunay edges have extremely large empty circumspheres.

We say that a crossing edge is *relaxed* if its endpoints lie on the boundary of an empty ball with radius less than  $4r$ , and *tense* otherwise. In this section, we show that few relaxed edges pass through any pixel. Once again, recall that that  $\varepsilon$  denotes the radius of a pixel.

**Lemma 2.5.** *If  $\varepsilon < 1/16$ , then for any pixel  $\pi$ , each point in  $P \cup Q$  is an endpoint of at most one relaxed edge passing through  $\pi$ .*

**Proof:** Suppose some point  $p \in P$  is an endpoint of two crossing edges  $pq$  and  $pq'$  passing through  $\pi$ , where  $|pq| \geq |pq'|$ . We immediately have  $\angle qpq' \leq 2 \tan^{-1}(\varepsilon/r)$  and  $|qq'| \geq 1$ . Thus, the circle through  $p$ ,  $q$ , and  $q'$  has radius at least  $1/(4 \tan^{-1}(\varepsilon/r)) \approx r/4\varepsilon > 4r$ . Any empty circumsphere of  $pq$  must have at least this radius, so it must be tense.  $\square$

**Lemma 2.6.** *The relaxed edges inside any slab of constant width are incident to at most  $O(r)$  endpoints.*

**Proof:** We will use a variant of our earlier ‘Swiss cheese’ packing argument [35]. Let  $\sigma$  be a slab of width  $\omega$  between two parallel planes, and let  $\sigma'$  be a parallel slab with the same central plane, with slightly larger width  $\omega + 1$ . For any point  $p \in P \cup Q$ , let  $U_p$  denote the unit-diameter ball centered at  $p$ ; these balls are pairwise disjoint.

Let  $C$  be an infinite circular cylinder whose axis is normal to the planes bounding  $\sigma$ , large enough to contain every ball  $U_p$  and  $U_q$ . The radius of  $C$  is at most  $(\sqrt{3} + 1)r + 1 < 4r$ . Finally, let  $D = C \cap \sigma'$ . The volume of this flat cylindrical disc is at most  $16\pi(\omega + 1)r^2 = O(r^2)$ .

Let  $E \subseteq \sigma \cap (P \cup Q)$  be the set of endpoints of relaxed edges that lie entirely within  $\sigma$ . For every endpoint  $p \in E$ , the ball  $U_p$  lies entirely within the disc  $D$ . For each endpoint  $p \in E$ , let  $B_p$  be the smallest Delaunay ball of some relaxed edge  $pq$ ; since  $pq$  is relaxed, the radius of  $B_p$  is at most  $4r$ . Let  $b_p$  be the open ball concentric with  $B_p$  but with radius smaller by  $1/3$ . Finally, let  $\Sigma$  denote the ‘Swiss cheese slice’  $D \setminus \bigcup_{p \in E} b_p$ .

For each point  $p \in E$ , let  $h_p = \partial \Sigma \cap \partial b_p$  be the surface of the corresponding ‘hole’ not eaten by any other ball, and let  $H = \bigcup_{p \in E} h_p$ . We claim that the surface area of  $H$  is only  $O(r)$ . The proof of this claim is elementary but tedious, and so we defer it to the appendix. By an argument identical to Lemma 2.6 of [35], each unit-diameter ball  $U_p$  contains at least  $\Omega(1)$  of this surface area—for completeness, we also include this argument in the appendix. Since the balls  $U_p$  are disjoint, and there is one for every endpoint  $p \in E$ , we conclude that  $E$  contains at most  $O(r)$  points.  $\square$

Together, Lemmas 2.4, 2.5, and 2.6 imply that  $O(r)$  relaxed edges intersect any pixel. Since there are  $O(r^2)$  pixels, we conclude that there are  $O(r^3)$  relaxed edges overall.

### 2.3 Tense Edges Are Easy to Relax

In order to count the tense crossing edge of  $P \cup Q$ , we will show that there are a constant number of transformations of space, such that every tense edge is mapped to a relaxed edge at least once.

A *Möbius transformation* is a continuous map from the extended Euclidean space  $\widehat{\mathbb{R}^d} = \mathbb{R}^d \cup \{\infty\} \simeq \mathbb{S}^d$  to itself that maps spheres to spheres. (A hyperplane in  $\mathbb{R}^d$  is a sphere through  $\infty$  in  $\widehat{\mathbb{R}^d}$ .) The space of Möbius transformations is generated by inversions. Examples include reflections, rotations, translations, dilations, and the well-known stereographic lifting map from  $\widehat{\mathbb{R}^d}$  to  $\mathbb{S}^d \subset \mathbb{R}^{d+1}$  relating  $d$ -dimensional Delaunay triangulations to  $(d+1)$ -dimensional convex hulls [13]. Möbius transformations are *conformal*, meaning they locally preserve angles. There are many other conformal maps in the plane—in fact, conformal maps are widely used in two-dimensional mesh-generation algorithms—but Möbius transformations are the only continuous conformal maps in dimensions three and higher. For further background on higher dimensional conformal transformations, see Hilbert and Cohn-Vossen [44], Thurston [58, 57], or Miller *et al.* [53].

Let  $S$  be a set of points in  $\mathbb{R}^3 \subset \widehat{\mathbb{R}^3}$ , let  $p, q, r, s \in S$  be the vertices of a Delaunay simplex, and let  $\Sigma$  be its empty circumsphere. For any conformal transformation  $\kappa$ , the points  $\kappa(p)$ ,  $\kappa(q)$ ,  $\kappa(r)$ , and  $\kappa(s)$  lie on

the sphere  $\kappa(\Sigma)$ , and this sphere either excludes every other point in  $\kappa(S)$  or contains every other point in  $\kappa(S)$ . In other words,  $\kappa(p)\kappa(q)\kappa(r)\kappa(s)$  is either a Delaunay simplex or an anti-Delaunay<sup>2</sup> simplex of  $\kappa(S)$ . Thus, the abstract simplicial complex consisting of Delaunay and anti-Delaunay simplices of any point set, which we call its *Delaunay polytope*, is invariant under conformal transformations.

In this section, we exploit this conformal invariance to count tense crossing edges. The main idea is to find a small collection of conformal maps, such that for any tense edge, at least one of the maps transforms it into a relaxed edge, by shrinking (but not inverting) its circumsphere. In order to apply our earlier arguments to count the transformed edges, we consider only conformal maps that map  $P \cup Q$  to another well-separated pair of sets with nearly the same spread.

Recall that  $P$  and  $Q$  lie inside balls of radius  $r$  centered at  $(0, 0, 2r)$  and  $(0, 0, -2r)$ , respectively. Call these balls  $\bigcirc P$  and  $\bigcirc Q$ . We say that an orientation-preserving conformal map  $\kappa$  is *rotary* if  $\kappa(\bigcirc P) = \bigcirc P$  and  $\kappa(\bigcirc Q) = \bigcirc Q$ . Rotary maps actually preserve a continuous one-parameter family of spheres centered on the  $z$ -axis, including the points  $p^* = (0, 0, \sqrt{3}r)$  and  $q^* = (0, 0, -\sqrt{3}r)$  and the plane  $z = 0$ . (In the space of spheres [24, 30], this family is just the line through  $\bigcirc P$  and  $\bigcirc Q$ .)

The image of  $P \cup Q$  under any rotary map is clearly well-separated. In order to apply our earlier arguments, we also require that these maps do not significantly change the spread.

**Lemma 2.7.** *For any rotary map  $\kappa$ , the closest pair of points in  $\kappa(P \cup Q)$  has distance between  $1/3$  and  $3$ .*

**Proof:** Consider the stereographic lifting map  $\lambda$  that takes  $p^*$  and  $q^*$  to opposite poles of  $\mathbb{S}^3$  and the plane  $z = 0$  to the equatorial sphere of  $\mathbb{S}^3$ . Any rotary map can be written as  $\lambda^{-1} \circ \rho \circ \lambda$ , where  $\rho$  is a simple rotation about the axis  $p^*q^*$ . (Thus, the space of rotary maps is isomorphic to  $SO(3)$ , the group of rigid motions of  $\mathbb{S}^2$ .) We can write  $\lambda$  as an inversion through a sphere of radius  $\sqrt{3}r$  as

$$\lambda(x, y, z) = \frac{3r^2(x, y, z, -\sqrt{3}r)}{x^2 + y^2 + z^2 + 3r^2}.$$

Consider a sphere  $\sigma$  of infinitesimal radius  $dr$ , centered in  $\bigcirc P$ . Tedious calculations (which we omit from this extended abstract) imply that  $\lambda(\sigma)$  is a sphere of infinitesimal radius between  $dr/4$  and  $3dr/4$ . It follows

<sup>2</sup>The anti-Delaunay triangulation is dual to the furthest point Voronoi diagram.

that any rotary map takes  $\sigma$  to a sphere of infinitesimal radius between  $dr/3$  and  $3dr$ . Thus, rotary maps locally change the metric at any point in  $\bigcirc P$  by at most a factor of 3.  $\square$

**Lemma 2.8.** *There is a set of  $O(1)$  rotary maps  $\{\pi_1, \pi_2, \dots, \pi_k\}$  such that any crossing edge of  $P \cup Q$  is mapped to a relaxed crossing edge of  $\pi_i(P \cup Q)$  by some  $\pi_i$ .*

**Proof:** Rotations about the  $z$ -axis are rotary maps, but since they do not actually change the radius of any sphere, we would like to ignore them. Two rotary maps  $\kappa_1$  and  $\kappa_2$  are *rotationally equivalent* if  $\kappa_1 = \rho \circ \kappa_2$  for some rotation  $\rho$  about the  $z$ -axis. The *rotation class* of a rotary map  $\kappa$ , denoted  $\langle \kappa \rangle$ , is the set of maps rotationally equivalent to  $\kappa$ . Since any rotation class  $\langle \kappa \rangle$  is uniquely identified by the point  $\kappa^{-1}(0, 0, 0)$ , the space of rotation classes is isomorphic to  $\widehat{\mathbb{R}^2} \simeq \mathbb{S}^2$ .

Let  $B_1, B_2, \dots, B_m$  be the smallest empty balls containing the crossing edges of  $P \cup Q$ . For each ball  $B_i$ , let  $\kappa_i$  denote any rotary map such that  $\kappa_i(B_i)$  is centered on the  $z$ -axis and is not everted, so  $\kappa_i(B_i)$  is an empty Delaunay ball of some crossing edge of  $\kappa_i(P \cup Q)$ . We easily observe that  $\kappa_i(B_i)$  has radius less than  $3r$ , so the corresponding crossing edge is relaxed. Thus, for each crossing edge, we have a point  $\langle \kappa_i \rangle$  on the sphere of rotation classes, corresponding to a rotation class of maps that relax that edge.

Our key observation is that we have a lot of 'wiggle room' in choosing our relaxing maps  $\kappa_i$ . To quantify this flexibility, consider the ball  $\bar{B}$  of radius  $3r$  centered at the origin; this is the smallest ball containing both  $\bigcirc P$  and  $\bigcirc Q$ . Let  $W$  be the set of rotation classes  $\langle w \rangle$  such that the radius of  $w(\bar{B})$  is at most  $4r$  and  $w(\bar{B})$  is not everted.  $W$  is a circular cap of some constant angular radius  $\omega$  on the sphere of rotation classes, centered at  $\langle 1 \rangle$ , the rotation class of the identity map.

For each  $i$ , the ball  $\kappa_i(B_i)$  lies entirely inside  $\bar{B}$ , so any rotation class in  $W$  transforms  $\kappa_i(B_i)$  into another ball of radius at most  $4r$ . Thus, *any* rotation class in the set  $W_i = \{\langle w \circ \kappa_i \rangle \mid \langle w \rangle \in W\}$  relaxes the  $i$ th crossing edge.  $W_i$  is a circular cap of angular radius  $\omega$  on the sphere of rotation classes, centered at the point  $\langle \kappa_i \rangle$ .

Since each of these  $m$  caps has constant angular radius, we can stab them all with a constant number of points. Specifically, let  $\Pi = \{\langle \pi_1 \rangle, \langle \pi_2 \rangle, \dots, \langle \pi_k \rangle\} \subset \mathbb{S}^2$  be a set of  $k = O(1/\omega^2)$  points on the sphere of rotation classes, such that any point in  $\mathbb{S}^2$  is within angular distance  $\omega$  of some point in  $\Pi$ . (In surface reconstruction terms,  $\Pi$  is a  $\omega$ -sample of the sphere.) Each disk  $W_i$  contains at least one point in  $\Pi$ , which

implies that each crossing edge is relaxed by some rotation class  $\langle \pi_j \rangle \in \Pi$ .

Finally, to satisfy the theorem, we can choose an arbitrary map  $\pi_j$  from each rotation class  $\langle \pi_j \rangle \in \Pi$ .  $\square$

It now follows immediately that  $P \cup Q$  has  $O(r^3)$  crossing edges.

## 2.4 Charging Delaunay Edges to Volume

In the last step of our proof, we count the Delaunay edges in an arbitrary point set  $S$  by decomposing it into a collection of subset pairs and counting the crossing edges for each pair.

Let  $S$  be an arbitrary set of points with diameter  $\Delta$ , where the closest pair of points is at unit distance.  $S$  is contained in an axis-parallel cube  $C$  of width  $\Delta$ . We construct a *well-separated pair decomposition* of  $S$  [14], based on a simple octree decomposition of  $C$ . (See [1] for a similar decomposition into subset pairs.) The octree has  $\log_2 \Delta$  levels; at each level  $\ell$ , there are  $8^\ell$  cubical cells, each a cube of width  $w_\ell = \Delta/2^\ell$ . Our well-separated pair decomposition  $\Xi = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_m, Q_m)\}$  contains, for each level  $\ell$ , the points in any pair of level- $\ell$  cells separated by a distance between  $3w_\ell$  and  $6w_\ell$ .

Each subset pair  $(P_i, Q_i) \in \Xi$  is well-separated: if the pair is at level  $\ell$  in our decomposition, then  $P_i$  and  $Q_i$  lie in a pair of balls of radius  $r_i$  separated by distance  $2r_i$ , for some  $r_i = \Theta(w_\ell)$ . Thus, by our earlier arguments,  $P_i \cup Q_i$  has at most  $O(w_\ell^3)$  crossing edges.

For any two points  $p, q \in S$ , we have  $p \in P_i$  and  $q \in Q_i$  for some index  $i$ . In particular, every Delaunay edge of  $S$  is a crossing edge between some subset pair in  $\Xi$ . A straightforward counting argument immediately implies that the total number of crossing edges, summed over all subset pairs in  $\Xi$ , is  $O(\Delta^3 \log \Delta)$  [35]. However, not every crossing edge appears in the Delaunay triangulation of  $S$ . We remove the final logarithmic factor by charging crossing edges to volume as follows.

We say that a subset pair  $(P_i, Q_i) \in \Xi$  is *relevant* if some pair of points  $p \in P_i$  and  $q \in Q_i$  are Delaunay neighbors in  $S$ . For each relevant pair  $(P_i, Q_i)$ , we define a large, close, and empty *witness ball*  $B_i$  as follows. Suppose  $P_i$  and  $Q_i$  are at level  $\ell$  in our decomposition, so the distance between  $p$  and  $q$  is at most  $(6 + \sqrt{3})w_\ell$ . Choose an arbitrary crossing edge  $pq$  of  $P_i \cup Q_i$ . Let  $\beta$  be the smallest ball with  $p$  and  $q$  on the boundary and no point of  $S$  in the interior; the radius of  $\beta$  is at least  $3w_\ell/2$ . Let  $\beta'$  be a ball concentric with  $\beta$  with radius smaller by  $w_\ell/2$ . Finally, let  $B_i$  be the (necessarily empty) ball of radius  $w_\ell$  inside  $\beta'$  whose center is closest to the midpoint  $m$  of segment  $pq$ . See Figure 3.




Figure 3. Defining the witness ball  $B_i$  for a relevant subset pair.

The distance from any point in  $B_i$  to any point in  $S$  is at least  $w_\ell/2$ , since  $B_i \subset \beta$  is empty. On the other hand, the triangle inequality implies that every point in  $B_i$  has distance less than  $(7 + \sqrt{3})w_\ell/2 < 5w_\ell$  either to  $p$  or to  $q$ . It follows that at most a constant number of witness balls overlap at any point.

For any relevant subset pair at level  $\ell$ , we can charge its  $O(w_\ell^3)$  crossing edges to its witness ball, which has volume  $\Omega(w_\ell^3)$ . Thus, the total number of relevant crossing edges is at most the sum of the volumes of all the witness balls. Since the witness balls have only constant overlap, the sum of their volumes is at most a constant factor larger than the volume of their union. Finally, every witness ball fits inside a cube of width  $8\Delta$  concentric with  $C$ . It follows that the number of Delaunay edges of  $S$  is at most  $O(\Delta^3)$ .

This completes the proof of Theorem 2.1.

### 3 Implications

Theorem 2.1 has several immediate algorithmic implications. For example, we can construct the Delaunay triangulation of a three-dimensional point set in  $O(\Delta^3 \log n)$  expected time using the standard randomized incremental algorithm [43], or in  $O(\Delta^3 \log^2 n)$  time using the deterministic algorithm of Chan, Snoeyink, and Yap [16]. Using the history graph of the randomized incremental algorithm, which has expected size  $O(\Delta^3 \log n)$ , we can answer nearest neighbor queries in  $O(\log^2 n)$  expected time. Since the Euclidean minimum spanning tree of a set of points is a subcomplex of the Delaunay triangulation, we can compute it in  $O(\Delta^3 \log n)$  time by first computing the Delaunay triangulation and then running any efficient minimum spanning tree algorithm on its  $O(\Delta^3)$  edges. We will discuss these results in greater detail in the full paper.

A somewhat less obvious implication concerns dense surface data. Let  $\Sigma$  be a  $C^2$  surface in  $\mathbb{R}^3$ . The *medial axis* of  $\Sigma$  is the set of points in  $\mathbb{R}^3$  with more than one nearest point in  $\Sigma$ . The *local feature size* of a point  $x \in \Sigma$ , denoted  $\text{lfs}(x)$ , is its distance to the medial axis.

A set of points  $S$  is a *uniform  $\varepsilon$ -sample* of  $\Sigma$  if the distance from any point  $x \in \Sigma$  to its *second* nearest neighbor in  $S$  is between  $\beta\varepsilon \text{lfs}(x)$  and  $\varepsilon \text{lfs}(x)$ , for some constant  $0 < \beta < 1$  [35, 36]. In earlier work [35], we showed that for any  $n$  and  $\varepsilon$ , there is a smooth surface  $\Sigma$  such that any uniform  $\varepsilon$ -sample of  $\Sigma$  has  $O(n)$  points and a Delaunay triangulation with complexity  $\Omega(n^2)$ . Our new results imply a more optimistic upper bound.

**Theorem 3.1.** *Let  $\Sigma$  be a fixed  $C^2$  surface in  $\mathbb{R}^3$ . The Delaunay triangulation of any uniform  $\varepsilon$ -sample of  $\Sigma$  has complexity  $O(n^{3/2})$ .*

**Proof:** Let  $S$  be a uniform  $\varepsilon$ -sample of  $\Sigma$ . This set contains  $n = \Theta(\mu/\varepsilon^2)$  points, where  $\mu$  is the *sample measure* of  $\Sigma$  [35, Lemma 3.1]. The spread of  $S$  is  $\Theta(\Delta/\varepsilon)$ , where  $\Delta$  is the *spread* of  $\Sigma$ , the ratio between the diameter of  $\Sigma$  and its minimum local feature size. Thus, by Theorem 2.1, the Delaunay triangulation of  $S$  has complexity  $O(n^{3/2}\Delta^3/\mu^{3/2}) = O(n^{3/2})$ .  $\square$

This bound is tight in the worst case, for example, when  $\Sigma$  is a circular cylinder with spherical caps [35]. Note that Theorem 3.1, which applies to any fixed surface, does *not* contradict our earlier  $\Omega(n^2)$  lower bound, which requires the surface to depend on  $n$  and  $\varepsilon$ .

Very recently, Attali and Boissonnat [8] showed that under certain sampling conditions, samples of *polyhedral* surfaces have subquadratic Delaunay triangulations. Unlike most surface-reconstruction results, their sampling conditions do not take local feature size into account (since otherwise samples would be infinite). Say that a point set  $S$  is a uniform *ambient  $\varepsilon$ -sample* of a surface  $\Sigma$  if the distance from any surface point to its second nearest sample is between  $\beta\varepsilon$  and  $\varepsilon$ , for some constant  $0 < \beta < 1$ . Similar arguments as Theorem 3.1 imply the following:

**Theorem 3.2.** *Let  $\Sigma$  be a fixed (not necessarily polyhedral or smooth) surface in  $\mathbb{R}^3$ . The Delaunay triangulation of any uniform ambient  $\varepsilon$ -sample of  $\Sigma$  has complexity  $O(n^{3/2})$ .*

### 4 Denser Regular Triangulations

In this section, we show that our  $O(\Delta^3)$  upper bound does not generalize to arbitrary triangulations, or even to *regular* triangulations. A regular triangulation (also called a weighted Delaunay triangulation) is the dual of a power diagram, or equivalently, the orthogonal projection of the lower convex hull of a set of points in on higher dimension [31].

**Theorem 4.1.** *For any  $n$  and  $\Delta < n$ , there is a set of  $n$  points with spread  $\Delta$  with a regular triangulation of complexity  $\Omega(n\Delta)$ .*

**Proof:** Any affine transformation of  $\mathbb{R}^d$  lifts to an essentially unique affine transformation of  $\mathbb{R}^{d+1}$  that preserves vertical lines and vertical distances. Since affine transformations preserve convexity, it follows that any affine transformation of a regular triangulation is another regular triangulation. Thus, to prove the theorem, it suffices to construct a set  $S$  of  $n$  points whose *Delaunay* triangulation has complexity  $\Omega(n\Delta)$ , such that some affine image of  $S$  has spread  $O(\Delta)$ .

Without loss of generality, assume that  $\sqrt{n/\Delta}$  is an integer. For each positive integer  $i, j \leq \sqrt{n/\Delta}$ , let  $s(i, j)$  be the line segment with endpoints  $(2i, 2j, 0) \pm ((-1)^{i+j}, (-1)^{i+j}, 1)$ . Let  $S$  be the set of  $n$  points containing  $\Delta$  evenly spaced points on each segment  $s(i, j)$ . Straightforward calculations imply that the Delaunay triangulation of  $S$  contains at least  $\Delta^2/4$  edges between any segment  $s(i, j)$  and any adjacent segment  $s(i \pm 1, j)$  or  $s(i, j \pm 1)$ . Thus, the overall complexity of the Delaunay triangulation of  $S$  is  $\Omega(n\Delta)$ . Applying the linear transformation  $f(x, y, z) = (x, y, \Delta z)$  results in a point set  $f(S)$  with spread  $O(\Delta)$ .  $\square$

## 5 Open problems

Our results suggest several open problems, the most obvious of which is to simplify our rather complicated proof of Theorem 2.1. The hidden constant in our upper bound is in the millions; the corresponding constant in the lower bound (which we think is closer to the true worst-case complexity) is close to 1.

We conjecture that Theorem 4.1 is tight for *arbitrary* triangulations. In fact, we believe that any complex of points, edges, and triangles, embedded in  $\mathbb{R}^3$  so that no triangle crosses an edge, has  $O(n\Delta)$  triangles. Even the following special case is still open: What is the minimum spread of a set of  $n$  points in  $\mathbb{R}^3$  in which *every* pair is joined by a Delaunay edge? We optimistically conjecture that the answer is exactly  $n/\pi$ .

What is the worst-case complexity of the convex hull of a set of  $n$  points in  $\mathbb{R}^4$  with spread  $\Delta$ ? Our earlier results [35] already imply a lower bound of  $\Omega(\min\{\Delta^3, n\Delta, n^2\})$ . This bound is *not* improved by Theorem 4.1, since our construction requires points large weights. The only known upper bound is  $O(n^2)$ .

Another interesting open problem is to generalize our results to higher dimensions. Our techniques almost certainly imply an upper bound of  $O(\Delta^d)$  on the number of Delaunay *edges*, improving our earlier upper

bound of  $O(\Delta^{d+1})$ . Unfortunately, this gives a very weak bound on the overall complexity, which we conjecture to be  $O(\Delta^d)$ . What is needed is a technique to directly count higher-dimensional Delaunay simplices: triangles in  $\mathbb{R}^4$ , tetrahedra in  $\mathbb{R}^6$ , and so on.

Standard range searching techniques can be used to answer nearest neighbor queries in  $\mathbb{R}^3$  in  $O(\log n)$  time using  $O(n^2/\text{polylog } n)$  space, or in  $O(\sqrt{n} \text{ polylog } n)$  time using  $O(n)$  space [2, 17, 22, 34, 50]. Using these data structures, we can compute the Euclidean spanning tree of a three-dimensional point set in  $O(n^{4/3+\epsilon})$  time [1]. All these results ultimately rely on the simple observation that the Delaunay triangulation of a random sample of a point set is significantly less complex (in expectation) than the Delaunay triangulation of the whole set. Unfortunately, if we try to reanalyze these algorithms in terms of the spread, this argument falls apart—in the worst case, a random sample of a point set with spread  $\Delta$  has expected spread close to  $\Delta$ , so we get no improvement in the Delaunay complexity. Can random sampling be integrated with our distance-sensitive bounds?

**Acknowledgments.** Once again, I thank Edgar Ramos for suggesting well-separated pair decompositions. Thanks also to Edgar Ramos, Herbert Edelsbrunner, Pat Morin, Sariel Har-Peled, Sheng-Hua Teng, Tamal Dey, and Timothy Chan for helpful discussions; to Pankaj Agarwal and Sariel Har-Peled for comments on an earlier draft of the paper; and to the anonymous reviewers for pointing to the work of Miller *et al.* [51, 53, 52].

## References

- [1] P. K. Agarwal, H. Edelsbrunner, O. Schwarzkopf, and E. Welzl. Euclidean minimum spanning trees and bichromatic closest pairs. *Discrete Comput. Geom.* 6(5):407–422, 1991.
- [2] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. *Advances in Discrete and Computational Geometry*, 1–56, 1999. Contemporary Mathematics 223, American Mathematical Society.
- [3] A. Aggarwal and P. Raghawan. Deferred data structures for the nearest-neighbor problem. *Inform. Process. Lett.* 40(3):119–122, 1991.
- [4] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* 22(4):481–504, 1999.
- [5] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Proc. SIGGRAPH '98*, 415–412, 1998. *Computer Graphics Proceedings, Annual Conference Series*.
- [6] 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.
- [7] N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Internat. J. Comput. Geom. Appl.* 19(2–3):127–153, 2001.
- [8] D. Attali and J.-D. Boissonnat. Complexity of Delaunay triangulations of points on polyhedral surfaces. Rapport de recherche 4015, INRIA Sophia-Antipolis, July 2001. (<http://www.inria.fr/rrrt/tr-4232.html>).
- [9] M. de Berg. *Ray Shooting, Depth Orders and Hidden Surface Removal*. Lecture Notes Comput. Sci. 703. Springer-Verlag, Berlin, Germany, 1993.
- [10] 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.



- [11] M. de Berg, M. Overmars, and O. Schwarzkopf. Computing and verifying depth orders. *SIAM J. Comput.* 23:437–446, 1994.
- [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] K. Q. Brown. Voronoi diagrams from convex hulls. *Inform. Process. Lett.* 9(5):223–228, 1979.
- [14] P. B. Callahan and S. R. Kosaraju. A decomposition of multi-dimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM* 42:67–90, 1995.
- [15] D. E. Cardoze and L. Schulman. Pattern matching for spatial point sets. *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.*, 156–165, 1998.
- [16] T. M. Chan, J. Snoeyink, and C. K. Yap. Primal dividing and dual pruning: Output-sensitive construction of 4-d polytopes and 3-d Voronoi diagrams. *Discrete Comput. Geom.* 18:433–454, 1997.
- [17] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.* 9(2):145–158, 1993.
- [18] B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Stolfi. Lines in space: Combinatorics and algorithms. *Algorithmica* 15:428–447, 1996.
- [19] H.-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete Comput. Geom.* 25(4):525–568, 2001.
- [20] 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.
- [21] K. L. Clarkson. A probabilistic algorithm for the post office problem. *Proc. 17th Annu. ACM Sympos. Theory Comput.*, 175–184, 1985.
- [22] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.* 17:830–847, 1988.
- [23] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.* 22:63–93, 1999.
- [24] O. Devillers, S. Meiser, and M. Teillaud. The space of spheres, a geometric tool to unify duality results on Voronoi diagrams. Report 1620, INRIA Sophia-Antipolis, 1992. (<http://www.inria.fr/rrrt/rr-1620.html>).
- [25] A. K. Dewdney and J. K. Vranč. A convex partition of  $\mathbb{R}^3$  with applications to Crum's problem and Knuth's post-office problem. *Utilitas Math.* 12:193–199, 1977.
- [26] T. K. Dey, S. Funke, and E. A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. *Abstracts 17th European Workshop Comput. Geom.*, 129–132, 2001. Freie Universität Berlin. (<http://www.cis.ohio-state.edu/~tamaldey/paper/recon-linear/surf.ps.gz>).
- [27] R. Dwyer. The expected number of k-faces of a Voronoi diagram. *Internat. J. Comput. Math.* 26(5):13–21, 1993.
- [28] R. A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.* 6:343–367, 1991.
- [29] H. Edelsbrunner. An acyclicity theorem for cell complexes in d dimensions. *Combinatorica* 10(3):251–260, 1990.
- [30] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.* 21:87–115, 1999.
- [31] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, Cambridge, England, 2001.
- [32] H. Edelsbrunner, X.-Y. Li, G. Miller, A. Stathopoulos, D. Talmor, S.-H. Teng, A. Üngör, and N. Walkington. Smoothing and cleaning up slivers. *Proc. 32nd Annu. ACM Sympos. Theory Comput.*, 273–277, 2000.
- [33] H. Edelsbrunner, P. Valtr, and E. Welzl. Cutting dense point sets in half. *Discrete Comput. Geom.* 17:243–255, 1997.
- [34] J. Erickson. Space-time tradeoffs for emptiness queries. *SIAM J. Comput.* 29(6):1968–1996, 2000.
- [35] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, 96–105, 2001. (<http://www.cs.uiuc.edu/~jeffe/pubs/spread.html>).
- [36] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near-linear time. *Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms*, p. to appear, 2002.
- [37] M. Gavrillov, P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric pattern matching: A performance study. *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, 79–85, 1999.
- [38] M. J. Golin and H. S. Na. On the average complexity of 3D-Voronoi diagrams of random points on convex polytopes. *Proc. 12th Canad. Conf. Comput. Geom.*, 127–135, 2000. (<http://www.cs.unb.ca/conf/cccg/eProceedings/>).
- [39] M. J. Golin and H. S. Na. On the average complexity of 3D-Voronoi diagrams of random points on convex polytopes. Tech. Rep. HKUST-TCSC-2001-08, Hong Kong Univ. Sci. Tech., June 2001. (<http://www.cs.ust.hk/tcsc/RR/2001-08.ps.gz>).
- [40] M. J. Golin and H. S. Na. On the proofs of two lemmas describing the intersections of spheres with the boundary of a convex polytope. Tech. Rep. HKUST-TCSC-2001-09, Hong Kong Univ. Sci. Tech., July 2001. (<http://www.cs.ust.hk/tcsc/RR/2001-09.ps.gz>).
- [41] J. E. Goodman, R. Pollack, and B. Sturmfels. Coordinate representation of order types requires exponential storage. *Proc. 21st Annu. ACM Sympos. Theory Comput.*, 405–410, 1989.
- [42] J. E. Goodman, R. Pollack, and B. Sturmfels. The intrinsic spread of a configuration in  $\mathbb{R}^d$ . *J. Amer. Math. Soc.* 3:639–651, 1990.
- [43] L. J. Guibas, D. E. Knuth, and M. Sharir. Randomized incremental construction of Delaunay and Voronoi diagrams. *Algorithmica* 7:381–413, 1992.
- [44] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. Chelsea Publishing Company, New York, NY, 1952.
- [45] H. Hiyoshi and K. Sugihara. Voronoi-based interpolation with higher continuity. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 242–250, 2000.
- [46] M. Inaba, N. Katoh, and H. Imai. Applications of weighted Voronoi diagrams and randomization to variance-based k-clustering. *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, 332–339, 1994.
- [47] P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric matching under noise: Combinatorial bounds and algorithms. *Proc. 8th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 457–465, 1999.
- [48] D. Krznaric and C. Levkopoulos. Computing hierarchies of clusters from the Euclidean minimum spanning tree in linear time. *Proc. 15th Conf. Found. Softw. Tech. Theoret. Comput. Sci.*, 443–455, 1995. Lecture Notes Comput. Sci. 1026, Springer-Verlag.
- [49] X.-Y. Li and S.-H. Teng. Generating well-shaped Delaunay meshes in 3D. *Proc. 12th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 28–37, 2001.
- [50] J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.* 10(2):157–182, 1993.
- [51] G. L. Miller, D. Talmor, and S.-H. Teng. Optimal coarsening of unstructured meshes. *J. Algorithms* 31(1):29–65, 1999.
- [52] G. L. Miller, D. Talmor, S.-H. Teng, and N. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation, and partition. *Proc. 27th Annu. ACM Sympos. Theory Comput.*, 683–692, 1995.
- [53] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM* 44:1–29, 1997.
- [54] T. Schreiber. A Voronoi diagram based adaptive k-means-type clustering algorithm for multidimensional weighted data. *Proc. Computational Geometry: Methods, Algorithms and Applications*, 265–275, 1991. Lecture Notes Comput. Sci. 553, Springer-Verlag.
- [55] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. *Proc. 18th Annu. ACM Sympos. Theory Comput.*, 404–413, 1986.
- [56] J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 86–95, 1998.
- [57] W. Thurston. *Three-Dimensional Geometry and Topology, Volume 1*. Princeton University Press, New Jersey, 1997.
- [58] W. P. Thurston. *The geometry and topology of 3-manifolds*. Mathematical Sciences Research Institute, Berkeley, CA, 1997. (<http://msri.org/publications/books/gt3m/>).
- [59] P. Valtr. Convex independent sets and 7-holes in restricted planar point sets. *Discrete Comput. Geom.* 7:135–152, 1992.
- [60] P. Valtr. *Planar point sets with bounded ratios of distances*. Ph.D. thesis, Fachbereich Mathematik, Freie Universität Berlin, Berlin, Germany, 1994.
- [61] P. Valtr. Lines, line-point incidences and crossing families in dense sets. *Combinatorica* 16:269–294, 1996.
- [62] K. Verbarg. Approximate center points in dense point sets. *Abstracts 13th European Workshop Comput. Geom.*, 55–56, 1997. Universität Würzburg.
- [63] J. Vleugels. *On Fairness and Fitness — Realistic Input Models for Geometric Algorithms*. Ph.D. thesis, Dept. Comput. Sci., Univ. Utrecht, Utrecht, The Netherlands, 1997.

## Appendix: Swiss Cheese Holes

In this appendix, we prove two key claims from the proof of Lemma 2.6. We refer the reader to Lemma 2.6 for definitions and notation.

**Claim 1.**  $\text{area}(H) = O(r)$ .

**Proof:** Let  $c_p$  be the common center of  $B_p$  and  $b_p$ , and let  $a_p$  be the axis line through  $c_p$  and normal to the planes bounding  $\sigma$ . For any point  $x \in h_p$ , let  $\bar{x}$  be its nearest neighbor on the axis  $a_p$ , and let  $H_p$  be the union of segments  $x\bar{x}$  over all  $x \in h_p$ . See Figure 4 for a two-dimensional example.

The triangle inequality implies that  $x\bar{x}$  and  $y\bar{y}$  have disjoint interiors for all  $x \neq y$ , so

$$\sum_{p \in E} \text{vol}(H_p) = \text{vol} \left( \bigcup_{p \in E} H_p \right) \leq \text{vol}(D).$$

We can compute the volume of each  $H_p$  as follows:

$$\begin{aligned} \text{vol}(H_p) &= \iint_{x \in h_p} \frac{|x\bar{x}| \cdot \cos \angle c_p x \bar{x}}{2} dx^2 \\ &= \iint_{x \in h_p} \frac{|x\bar{x}|^2}{2|xc_p|} dx^2 \\ &= \frac{1}{2|pc_p|} \iint_{x \in h_p} |x\bar{x}|^2 dx^2 \\ &\geq \frac{\text{area}(h_p)}{12r} \cdot \min_{x \in h_p} |x\bar{x}|^2 \end{aligned}$$

For notational convenience, let  $r_p = \min_{x \in h_p} |x\bar{x}|$ .




Figure 4. Proof of Claim 1 of Lemma 2.6. The slab  $\sigma'$  is shaded.

The intersection  $b_p \cap \partial\sigma'$  consists of two parallel disks, the smaller of which has radius  $r_p$ . Since both  $\sigma'$  and the boundary of  $B_p$  contain the endpoints of some crossing edge  $pq$  of length at least  $2r$ , the larger of these two disks has radius at least  $r - \omega - 4/3$ . Again referring

to Figure 4, we choose points  $y, z \in h_p \cap \partial\sigma'$  on the boundary of the larger and smaller disks, respectively.

$$\begin{aligned} |z\bar{z}|^2 &= |zc_p|^2 - |\bar{z}c_p|^2 \\ &= |zc_p|^2 - \left( |\bar{z}y| + \sqrt{|yc_p|^2 - |y\bar{y}|^2} \right)^2 \\ &\geq |y\bar{y}|^2 - 2|\bar{z}y| \sqrt{|yc_p|^2 - |y\bar{y}|^2} - |\bar{z}y|^2 \\ &\geq |y\bar{y}|^2 - 2|\bar{z}y| |yc_p| - |\bar{z}y|^2 \end{aligned}$$

Now substituting the known equations and inequalities

$$\begin{aligned} r_p &= |z\bar{z}|, & |y\bar{y}| &\geq r - \omega - 4/3, \\ |\bar{z}y| &= \omega + 1, & |yc_p| &< 4r - 1/3, \end{aligned}$$

we obtain the lower bound  $r_p^2 = \Omega(r^2)$ . We can now finally bound the area of  $H$ :

$$\begin{aligned} \text{area}(H) &= \sum_{p \in E} \text{area}(h_p) \\ &\leq \sum_{p \in E} \frac{12r}{r_p^2} \text{vol}(H_p) \\ &= O(1/r) \cdot \sum_{p \in E} \text{vol}(H_p) \\ &\leq O(1/r) \cdot \text{vol}(D) \\ &= O(r). \quad \square \end{aligned}$$

**Claim 2 (Lemma 2.6 of [35]).** *Let  $U$  be any unit-diameter ball contained in  $\Sigma$  whose center is distance  $1/3$  from  $H$ . Then  $U$  contains  $\Omega(1)$  surface area of  $H$ .*

**Proof:** Without loss of generality, assume that  $U$  is centered at the origin and that  $(0, 0, 1/3)$  is the closest point of  $H$  to the origin. Let  $U'$  be the open ball of radius  $1/3$  centered at the origin, let  $V$  be the open unit ball centered at  $(0, 0, 5/6)$ , and let  $W$  be the cone whose apex is the origin and whose base is the circle  $\partial U \cap \partial V$ . See Figure 5.




Figure 5. Proof of Claim 2 of Lemma 2.6

$U'$  lies entirely inside  $\Sigma$ , and since  $r \geq 1$ , we easily observe that  $V$  lies entirely outside  $\Sigma$ . Thus, the surface area of  $H \cap W \subseteq H \cap U$  is at least the area of the spherical cap  $\partial U' \cap W$ , which is exactly  $\pi/27$ .  $\square$