

Complexity of Delaunay Triangulation for Points on Lower-dimensional Polyhedra.

Nina Amenta*
UC Davis

Dominique Attali†
CNRS-LIS Grenoble

Olivier Devillers‡
INRIA Sophia-Antipolis

Abstract

We show that the Delaunay triangulation of a set of points distributed nearly uniformly on a polyhedron (not necessarily convex) of dimension p in d -dimensional space is $O(n^{(d-1)/p})$. For all $2 \leq p \leq d-1$, this improves on the well-known worst-case bound of $O(n^{\lceil d/2 \rceil})$.

1 Introduction

The Delaunay triangulation of a set of points is a data structure, which in low dimensions has applications in mesh generation, surface reconstruction, molecular modeling, geographic information systems, and many other areas of science and engineering. Like many spatial partitioning techniques, however, it suffers from the “curse of dimensionality”: in higher dimensions, the complexity of the Delaunay triangulation increases exponentially. Its worst-case complexity is bounded precisely by the following theorem, known as the Upper Bound Theorem.

THEOREM 1.1. (MCMULLEN [13]) *The number of simplices in the Delaunay triangulation of n points in dimension d is at most*

$$\binom{n - \lfloor \frac{d+1}{2} \rfloor}{n-d} + \binom{n - \lfloor \frac{d+2}{2} \rfloor}{n-d} = O(n^{\lceil \frac{d}{2} \rceil})$$

This bound is achieved exactly by the vertices of a cyclic polytope, which all lie on a one-dimensional curve known as the *moment curve*. Indeed all of the examples that we have of point sets which have Delaunay triangulations of complexity $O(n^{\lceil d/2 \rceil})$ are distributed on one-dimensional curves. At the opposite extreme, points distributed uniformly at random inside

the sphere have Delaunay triangulations of complexity $O(n)$ in any fixed dimension, with a constant factor which is exponential in the dimension [9]. Our goal in this paper is to begin to fill in the picture for distributions between the two extremes, in which the points lie on manifolds of dimension $2 \leq p \leq d-1$.

As an easy first case, we consider a fixed polyhedral set (not necessarily convex) \mathbb{P} of dimension p in $d > p$ dimensional space. Our point set S is a *sparse ϵ -sample* from \mathbb{P} . Sparse ϵ -sampling is a model, sometimes used in computational geometry, in which the sampling can be neither too dense nor too sparse; we define it precisely below. Let n be the number of points in S . We consider how the complexity of the Delaunay triangulation of S grows, as $n \rightarrow \infty$, with \mathbb{P} remaining fixed. Our main result is that the number of simplices of all dimensions is $O(n^{(d-1)/p})$. The hidden constant factor depends, among other things, on the geometry of \mathbb{P} , which is constant since \mathbb{P} is fixed.

While our result is purely combinatorial, it has both potential and immediate algorithmic implications. The Delaunay triangulation can be computed in optimal worst-case time in dimension $d \geq 3$ by the standard randomized incremental algorithm [8, 16], or deterministically [7]. While our result does not immediately improve these running times for the special case of points distributed on lower-dimensional manifolds [3], it is of course a necessary step towards such an improvement. Our result shows that Seidel’s giftwrapping algorithm [15] runs in time $O(n^2 + n^{(d-1)/p} \lg n)$ in our special cases, which can be somewhat improved using more sophisticated data structures [4].

1.1 Prior work. The complexity of the Delaunay triangulation of a set of points on a two-manifold in \mathbb{R}^3 has received considerable recent attention, since such point sets arise in practice, and their Delaunay triangulations are found nearly always to have linear size. Golin and Na [12] proved that the Delaunay triangulation of a large enough set of points distributed uniformly at random on the surface of a fixed convex polytope in \mathbb{R}^3 has expected size $O(n)$. They later [11] established

*namenta@ucdavis.edu. Computer Science Department, University of California, One Shields Ave, Davis, CA 95616. Fax 1-530-752-5767. Supported by NSF CCF-0093378.

†Dominique.Attali@lis.inpg.fr. LIS, ENSIEG, Domaine Universitaire, BP 46, 38402 Saint Martin d’Hères, France. Supported by the EU under contract IST-2002-506766 (Aim@Shape) and CNRS under grant PICS 3416.

‡Olivier.Devillers@sophia.inria.fr. INRIA, BP 93, 06902 Sophia-Antipolis, France. Supported by the EU under contract IST-2002-506766 (Aim@Shape)

an $O(n \lg^6 n)$ upper bound with high probability for the case in which the points are distributed uniformly at random on the surface of a non-convex polyhedron.

Attali and Boissonnat considered the problem using a sparse ϵ -sampling model similar to the one we use here, rather than a random distribution. For such a set of points distributed on a polygonal surface \mathbb{P} , they showed that the size of the Delaunay triangulation is $O(n)$ [1]. Our proof gives the same bound, and is perhaps a little simpler; but, as we shall describe in a moment, our definition of sparse ϵ -sampling for polyhedra is a little more restrictive. In a subsequent paper with Lieutier [2] they considered “generic” surfaces, and got an upper bound of $O(n \lg n)$. Specifically, a “generic” surface is one for which each medial ball touches the surface in at most a constant number of points.

The genericity assumption is important. Erickson considered more general point distributions, which he characterized by the *spread*: the ratio of the largest inter-point distance to the smallest. The spread of a sparse ϵ -sample of n points from a two-dimensional manifold is $O(\sqrt{n})$. Erickson proved that the Delaunay triangulation of a set of points in \mathbb{R}^3 with spread Δ is $O(\Delta^3)$. Perhaps even more interestingly, he showed that this bound is tight for $\Delta = \sqrt{n}$, by giving an example of a sparse ϵ -sample of points from a cylinder that has a Delaunay triangulation of size $\Omega(n^{3/2})$ [10]. Note that this surface is not generic and has a degenerate medial axis.

To the best of our knowledge, there are no prior results for $d > 3$.

1.2 Overview of the proof. Our proof uses two samples, the original sparse ϵ -sample S from the polyhedron \mathbb{P} , and a sparse ϵ -sample M of a bounded subset \mathcal{M}^* of the medial axis of \mathbb{P} . We prove that any Delaunay ball circumscribing points of S intersects the polyhedron in a set of points that is contained in an enlarged medial ball centered at a medial sample point $z \in M$. We then prove that each sample z in M is assigned at most a constant number of Delaunay balls.

Since M is a sparse ϵ -sample from a fixed $(d - 1)$ -dimensional set of constant volume, its cardinality is $m = O(\epsilon^{-(d-1)})$. Similarly, S is a sparse ϵ -sample of \mathbb{P} and we get $n = \Omega(\epsilon^{-p})$. Eliminating ϵ gives $m = O(n^{(d-1)/p})$, and since each sample of M is charged for a constant number of Delaunay balls, this bound applies to the size of the Delaunay triangulation as well. This is the key insight: as a function of ϵ , the number of Delaunay balls depends only on the dimension of the medial axis, which is always $d - 1$. The number of samples, n , depends on the dimension p of \mathbb{P} . As p increases, n increases, but the complexity of the

Delaunay triangulation stays about the same. If written as a function of n , the complexity of the Delaunay triangulation goes down.

2 Statement of Theorem

In this section, we introduce the setting for our result. We first define simplicial complexes, Delaunay triangulations, polyhedra and medial axes.

2.1 Simplicial complexes. We refer to [14] for more details on simplicial complexes. A *geometric simplex*, σ , is the convex hull of a collection of affinely independent points in \mathbb{R}^d . If there are $k + 1$ points in the collection, we call σ a k -simplex and k its *dimension*. Any simplex spanned by a subset of the $k + 1$ points is called a *face* of σ . A *geometric simplicial complex* is a finite collection of simplices, K , satisfying the two following properties: (1) every face of a simplex in K is in K ; (2) the intersection of any two simplices of K is either empty or a face of each of them. Its *underlying space*, $|K|$, is the subset of \mathbb{R}^d , covered by the simplices with the subspace topology inherited from \mathbb{R}^d . We will also need abstract versions of those geometric notions. An *abstract simplicial complex*, K , is a collection of finite nonempty sets, such that if $\sigma \in K$, so is every nonempty subset of σ . The element σ of K is called an *abstract simplex*, its *dimension* is one less than its cardinality. A *face* of σ is any nonempty subset of σ . The vertex set of σ is the collection one-point elements of σ , which we denote as $\text{Vert } \sigma$. The *dimension* of a simplicial complex K , geometric or abstract, is the largest dimension of any simplex in K .

2.2 Delaunay triangulations. Let $S \subseteq \mathbb{R}^d$ be a finite set of points. The *Voronoi region* $V(s)$, of $s \in S$ is the set of points $x \in \mathbb{R}^d$ with $\|x - s\| \leq \|x - t\|$ for all $t \in S$. The Delaunay triangulation $\text{Del}(S)$ of S is the nerve of the Voronoi regions. Specifically, an abstract simplex $\sigma = \{s_0, \dots, s_k\} \subseteq S$ belongs to the Delaunay triangulation iff the Voronoi regions of its vertices have a nonempty common intersection, $\bigcap_{0 \leq i \leq k} V(s_i) \neq \emptyset$. Equivalently, a simplex σ is in the Delaunay triangulation iff there exists of a $(d - 1)$ -sphere, called *Delaunay sphere*, that passes through s_0, \dots, s_k and encloses no point of S . Afterwards, any d -ball bounded by a Delaunay sphere is referred to as a *Delaunay ball*. The Delaunay triangulation is an abstract simplicial complex. Notice that using this definition, on degenerate inputs in which $k + 1 > d + 2$ points are co-spherical, every $d + 1$ subset of these points defines a Delaunay simplex, showing that the Delaunay triangulation is not necessarily a geometric simplicial complex. In this paper, we don't assume

general position for points in S . The *complexity* (or *size*) of the Delaunay triangulation is the total number of its simplices of all dimensions. We express it as a function of n , the number of points in S .

2.3 Polyhedron. A *polyhedron* is the underlying space of a geometric simplicial complex. If the dimension of the simplicial complex K is p , we say that $\mathbb{P} = |K|$ is a p -polyhedron. The collection of simplices $\sigma \in K$, whose dimension is i or less is called the *i -skeleton* of K and denoted $K^{(i)}$. Given a point $x \in \mathbb{P}$, we let i be the largest dimension for which a small open d -ball centered at x intersects $|K^{(i)}|$ in an open i -ball $N(x)$ of $|K^{(i)}|$ (see Figure 1). The *tangent flat* to \mathbb{P} at x , denoted by TP_x , is the i -flat spanned by the open i -ball $N(x)$. A *face* F of \mathbb{P} is a maximal collection of points with identical tangent flat. If the dimension of the tangent flat is i , F is an i -face. The 0-faces are called *vertices*. Afterwards, f designates the number of faces of \mathbb{P} . We assume the dimension p of \mathbb{P} is less than d .




Figure 1: The tangent flat to \mathbb{P} at x is a line. Note that a polyhedron may have smaller-dimensional parts. Points that sample lower-dimensional parts can have a high-complexity Delaunay triangulation but still not damage the overall complexity, because they form a small subset of the set of sample points.

2.4 Medial axis. The medial axis is instrumental in expressing the constant in our main result and crucial for proving it. The *medial axis*, $\mathcal{M} = \mathcal{M}(\mathbb{P})$, is the set of points that have at least two closest points in \mathbb{P} . Formally, writing $\Pi(z)$ for the set of points in \mathbb{P} with minimum distance to $z \in \mathbb{R}^d$, we have

$$\mathcal{M} = \{z \in \mathbb{R}^d \mid \text{card} \Pi(z) \geq 2\}.$$

A point z on the medial axis is said to be *medial*. The smallest d -ball centered at z and containing $\Pi(z)$ is

called a *medial ball*. The medial axis of a polyhedron \mathbb{P} admits a stratification. In other words, it can be decomposed into a finite number of strata, each a connected i -manifold with boundary, for $i < d$. Furthermore, any point $x \in \mathcal{M}$ belongs to the closure of at least one stratum of dimension $d - 1$. Roughly speaking, the medial axis of a polyhedron is a $(d - 1)$ -dimensional surface that has a positive $(d - 1)$ -dimensional volume, possibly infinite.




Figure 2: The essential medial axis is solid and includes the black dots. The non-essential part is dotted. The point z is essential. Observe that the circle around z is tangent to one edge incident to a gray point.

We introduce a bounded subset of the medial axis, $\mathcal{M}^* = \mathcal{M}^*(\mathbb{P})$, called the *essential medial axis* and defined as follows. For $z \in \mathcal{M}$, we let $c(z)$ be the center of the smallest d -ball enclosing $\Pi(z)$. A point z for which $z = c(z)$ is called a *critical point* of the distance-to-polyhedron function [6, 5]. One can prove that $z = c(z)$ iff z lies in the convex hull of $\Pi(z)$. It follows that if z lies outside the convex hull of the polyhedron, then $z \neq c(z)$ and the line passing through $c(z)$ and z is well defined. We denote by $L(z)$ the half-line that starts at $c(z)$ and passes through z . We are now ready to define what it means for a point z to be *essential*. We consider two cases: if z belongs to the convex hull of \mathbb{P} , then z is essential; otherwise, z is essential iff the half-line $L(z)$ minus the segment connecting z and $c(z)$ is contained in the medial axis and z is the center of the smallest medial ball enclosing $\Pi(z)$ and centered on $L(z)$ (see Figure 2). An *essential medial ball* is a d -ball whose center is essential. The *essential medial axis* is defined as the set of essential medial points and, by construction, is bounded. Unlike the medial axis, it has a finite $(d - 1)$ -dimensional volume, which vanishes iff the polyhedron lies on an hyperplane of \mathbb{R}^d .

2.5 Sampling condition. We write $B(x, r)$ for the closed d -ball with center x and radius r . Given a polyhedron $\mathbb{P} \subseteq \mathbb{R}^d$, we say that a set of points $S \subseteq \mathbb{P}$ is a λ -sparse ϵ -sample of \mathbb{P} iff it satisfies the following two conditions:

Density: for every point x on a face F of \mathbb{P} , $B(x, \epsilon) \cap F$ contains at least one point of S ;

Sparsity: for every point $y \in \mathbb{R}^d$, $B(y, d\epsilon)$ contains at most λ points of S .

The number of samples on each i -face is $\Theta(\epsilon^{-i})$. Writing p for the dimension of \mathbb{P} , the number of points in the sample is $n = \Theta(\epsilon^{-p})$.

To state our main result, we let v_i be the i -dimensional volume of a unit i -ball. The i -dimensional volume of a i -ball of radius r is $v_i r^i$. Furthermore, writing $\text{vol}_i(X)$ for the i -dimensional volume of $X \subseteq \mathbb{R}^d$, we assume that there exist two positive constants $w_{\mathbb{P}} > 0$ and $r_0 > 0$ such that for every $0 < r \leq r_0$, for every point $x \in \mathcal{M}^*$

$$(2.1) \quad w_{\mathbb{P}} \leq \frac{\text{vol}_{d-1}(B(x, r) \cap \mathcal{M})}{r^{d-1}}.$$

To provide an intuition for the constant $w_{\mathbb{P}}$, suppose x is a point on a smooth hypersurface \mathbb{M} embedded in \mathbb{R}^d . In first approximation, the intersection of a d -ball $B(x, r)$ and the hypersurface has the volume of a $(d-1)$ -ball with radius r , showing that the ratio on the right side of Inequality (2.1) tends to v_{d-1} as r tends to 0. In our case, x does not live on an hypersurface but on the medial axis which is a set of patches of smooth hypersurfaces. If x lies inside a patch, the ratio tends to v_{d-1} . Inequality (2.1) takes into account the fact that if x lies on the boundary of some patches, the ratio tends to something which depends on the local geometry of the incident patches. Finally, we introduce the constant

$$C(\mathbb{P}) = 2^{2d-1} \frac{\text{vol}_{d-1}(\mathcal{M}^*(\mathbb{P}))}{\text{vol}_p(\mathbb{P})^{\frac{d-1}{p}}} \times \frac{(2v_p)^{\frac{d-1}{p}}}{w_{\mathbb{P}}}.$$

THEOREM 2.1. (MAIN THEOREM) *Let \mathbb{P} be a p -polyhedron in \mathbb{R}^d composed of f faces not all contained in an hyperplane. Let S be a set of n points that forms a λ -sparse ϵ -sample of \mathbb{P} . The Delaunay triangulation of S has complexity $O(n^{\frac{d-1}{p}})$. More precisely, for n large enough, the number of Delaunay k -simplices is bounded from above by*

$$C(\mathbb{P}) \binom{\lambda f}{k+1} n^{\frac{d-1}{p}} + \binom{f}{k+1} n.$$

Our proof assumes that neither the p -dimensional volume of the polyhedron nor the $(d-1)$ -dimensional volume of the essential medial axis vanishes, which ensures that $0 < C(\mathbb{P}) < \infty$.

3 Covering Delaunay balls

In this section, we establish preliminary results that bound the intersection of a Delaunay ball and the polyhedron. For this, we use a shape obtained by enlarging either a medial ball or a point on the polyhedron. Those results will be used in Section 4 to prove our bound on the complexity of Delaunay triangulations.

3.1 Almost empty balls. We first give a crucial property of Delaunay balls induced by our sampling condition. For this, we need notations and definitions. Recall that $B(z, r)$ is the closed d -ball with center z and radius r . We define the *penetration* $h(b)$ of $b = B(z, r)$ as the difference of square radii between b and the largest d -ball centered at z and whose interior does not intersect \mathbb{P}

$$h(b) = r^2 - d(z, \mathbb{P})^2.$$

Note that b intersects the polyhedron iff $h(b) \geq 0$. A ball b for which $h(b) \leq 0$ is said to be *empty*. Calling a ball b for which $h(b) \leq \epsilon^2$ an ϵ -almost empty ball, we have the fundamental property:

FUNDAMENTAL PROPERTY. Delaunay balls are ϵ -almost empty.

Proof. Let x be a point in \mathbb{P} with minimum distance to the center of b and $s \in S$ be the closest sample on the i -face that contains x . Because of our sampling condition, $\|x - s\| \leq \epsilon$. Because b contains no sample point in its interior, b intersects $\mathbb{T}\mathbb{P}_x$ in an i -ball of radius $\sqrt{h(b)} \leq \|x - s\|$. It follows that $h(b) \leq \epsilon^2$. \square

3.2 Contact. We now introduce the notion of ϵ -contact of a ball ν and use it to cover the intersection of ϵ -almost empty balls with the polyhedron. Given a d -ball ν with center z and radius r , we let $\nu^{\oplus a} = B(z, \sqrt{r^2 + a^2})$ be the d -ball with center z and radius $\sqrt{r^2 + a^2}$. Recalling that $\Pi(z)$ is the set of points in \mathbb{P} with minimum distance to z , we define the ϵ -contact of ν as

$$\text{Contact}(\nu, \epsilon) = (\nu^{\oplus \epsilon} \cap \mathbb{P}) \cup \bigcup_{x \in \Pi(z)} B(x, \epsilon).$$

Suppose b is an ϵ -almost empty ball. We show in the next lemma that there exists an empty ball ν whose ϵ -contact contains $b \cap \mathbb{P}$. In addition, this empty ball is either medial or a point of \mathbb{P} . To distinguish between the two cases, we classify balls as *ordinary* or *trivial*. Any medial ball is said to be ordinary. If a ball b is not medial, its center y has a unique closest point $x \in \mathbb{P}$.

If $x = y$, or equivalently if the ball b is centered on the polyhedron, b is said to be trivial. Suppose now that $x \neq y$. The half-line $L = L(y)$ beginning at x and passing through the center y of b either intersects the medial axis in a point z , or extends to infinity (when y and \mathbb{P} lie on both sides of the hyperplane that passes through x and is orthogonal to the straight-line xy). In the first case we say that b is *ordinary*, and in the second we say that b is *trivial*. Notice that when b is ordinary, y lies in the interior of segment xz (since x is the closest point on \mathbb{P} to y).

LEMMA 3.1. *Suppose b is an ϵ -almost empty ball.*

- (i) *If b is trivial, then there exists a point $x \in \mathbb{P}$ such that $b \cap \mathbb{P} \subseteq B(x, \epsilon)$.*
- (ii) *If b is ordinary, then there exists a medial ball ν such that $b \cap \mathbb{P} \subseteq \text{Contact}(\nu, \epsilon)$.*

Proof. Without loss of generality, we can assume that b intersects the polyhedron and $h(b) \geq 0$. If this is not the case, we replace b by the largest empty ball with same center. If b is medial, the result follows immediately by setting $\nu = b$. If the center y of b lies on \mathbb{P} , the result also follows immediately by observing that $b \subseteq B(y, \epsilon)$. Suppose now that b is not medial and its center y has a unique closest point $x \neq y$ on \mathbb{P} (see Figure 3). Consider the half-line, $L = L(y)$, with origin x and




Figure 3: Notations for the proof of Lemma 3.1. Dotted spheres are tangent to \mathbb{P} . Solid spheres penetrate \mathbb{P} from r and belong to the same pencil of spheres passing through the boundary of τ .

passing through y . The half-line L either intersects the medial axis or extends at infinity. In both cases, we let $z \in \mathcal{M}$ be the point on L (possibly at infinity) closest to y . Because b is ϵ -almost empty and contains x , the

hyperplane H that passes through x and is orthogonal to the straight-line xy intersects b in a $(d-1)$ -ball τ with center x and radius $0 \leq r \leq \epsilon$. Consider the pencil of $(d-1)$ -spheres that pass through the boundary of τ and let b_u be the d -ball bounded by a sphere in this pencil with center $u \in L$. Because the interval defined by x and z contains y , we have $b = b_y \subseteq b_x \cup b_z$. First, assume that z belongs to the medial axis (*i.e.* b is ordinary) and let ν be the medial ball centered at z . This medial ball passes through x , and therefore $b_z = \nu^{\oplus r} \subseteq \nu^{\oplus \epsilon}$. Observing that $x \in \Pi(z)$, it follows that

$$\begin{aligned} b \cap \mathbb{P} &\subseteq (b_x \cup b_z) \cap \mathbb{P} \\ &\subseteq (B(x, \epsilon) \cup \nu^{\oplus \epsilon}) \cap \mathbb{P} \\ &\subseteq \text{Contact}(\nu, \epsilon). \end{aligned}$$

Second, assume that z is at infinity (*i.e.* b is trivial). The ball b_z is the one half-space that H bounds and which contains y . Its interior does not intersect \mathbb{P} and it follows that

$$b \cap \mathbb{P} \subseteq b_x = B(x, r) \subseteq B(x, \epsilon). \quad \square$$

We strengthen the second part of Lemma 3.1 by showing that the medial ball ν whose ϵ -contact contains $b \cap \mathbb{P}$ can be chosen essential.

LEMMA 3.2. *For every non-essential medial ball μ , there exists an essential medial ball ν such that*

$$\text{Contact}(\mu, \epsilon) \subseteq \text{Contact}(\nu, \epsilon)$$

Proof. See Figure 4; let y be the center of μ . Let $c = c(y)$ be the center of the smallest ball enclosing $\Pi(y)$, the set of points in \mathbb{P} with minimum distance to y . Because y is not essential, it does not lie in the convex hull of \mathbb{P} and $y \neq c$. Consider H , the hyperplane passing through c and orthogonal to the vector $c - y$. \mathbb{P} and y lie on opposite sides of the hyperplane H . Let w and w' be the two closed half-spaces that H bounds and assume that $y \in w$ and $\mathbb{P} \in w'$. The $(d-1)$ -spheres bounding μ and w generate a pencil of spheres, whose centers lie on the straight-line passing through c and y . They intersect the hyperplane H in a common $(d-1)$ -ball τ , whose boundary contains $\Pi(y)$. Consider the smallest medial sphere in the pencil and let ν be the ball that this sphere bounds. The center z of ν lies on the segment cy . Note that $\Pi(y) \subseteq \Pi(z)$ and points in $\Pi(z) - \Pi(y)$ lie in the smallest ball centered at c and passing through $\Pi(y)$. It follows that the smallest ball enclosing $\Pi(z)$ is also centered at $c = c(z)$. Therefore, ν is an essential medial ball. Now, observe that $\mu^{\oplus \epsilon}$ belongs to the pencil




Figure 4: Notations for the proof of Lemma 3.2. The point z is essential and the point y is not. Inner spheres of gray annuli are medial. Outer spheres penetrate \mathbb{P} from ϵ .

of spheres generated by $\nu^{\oplus\epsilon}$ and w . Since the center y of the former lie on the half-line defined by the centers of the latter, we have $\mu^{\oplus\epsilon} \subseteq \nu^{\oplus\epsilon} \cup \omega$. It follows that

$$\mu^{\oplus\epsilon} \cap \mathbb{P} \subseteq \nu^{\oplus\epsilon} \cap \mathbb{P},$$

and because $\Pi(y) \subseteq \Pi(z)$, the ϵ -contacts of μ and ν are nested as claimed. \square

3.3 Extended contact. It will be convenient to bound the intersection of Delaunay balls and the polyhedron, using a slightly different notion of ϵ -contact. Given $z \in \mathbb{R}^d$, we let $\chi(z)$ be the set of orthogonal projections of z onto the tangent planes of \mathbb{P} . The set $\chi(z)$ contains at most f points, one for each face of the polyhedron. We define the *extended ϵ -contact* of the d -ball ν with center z as

$$\text{ExtendedContact}(\nu, \epsilon) = \bigcup_{x \in \chi(z)} B(x, \epsilon).$$

For any two medial balls whose centers are less than ϵ apart, we show that the ϵ -contact of one is contained in the extended $(d\epsilon)$ -contact of the other. This result will be used in the next section to reduce the amount of essential medial balls necessary to cover intersections of Delaunay balls with the polyhedron.

LEMMA 3.3. For any two medial balls μ and ν whose centers y and z satisfy $\|y - z\| \leq \epsilon$,

$$\text{Contact}(\mu, \epsilon) \subseteq \text{ExtendedContact}(\nu, d\epsilon).$$

Proof. Let $x \in \text{Contact}(\mu, \epsilon)$. First, suppose that $x \in B(q, \epsilon)$ for some $q \in \Pi(y)$. Writing $\pi_q(z)$ for the orthogonal projection of z onto the tangent plane passing through $q \in \mathbb{P}$ and using $\pi_q(y) = q$, we get

$$\begin{aligned} \|x - \pi_q(z)\| &\leq \|x - q\| + \|\pi_q(y) - \pi_q(z)\| \\ &\leq \epsilon + \|y - z\| \\ &\leq 2\epsilon, \end{aligned}$$

and $x \in \text{ExtendedContact}(\nu, 2\epsilon)$. Suppose now that $x \in \mu^{\oplus\epsilon}$ and consider the annulus $\alpha = \mu^{\oplus\epsilon} - \mu$. We construct a sequence of points in $\mathbb{P} \cap \alpha$ starting from x such that the last point q in the sequence satisfies $\pi_q(y) \in \alpha$ and any two consecutive points in the sequence are less than ϵ apart. Starting from x , we distinguish two cases:

1. if $\pi_x(y) \in \alpha$, we are done: the sequence ends at x .
2. if $\pi_x(y) \notin \alpha$, we consider the point $q \in \mathbb{P}$ on the segment $x\pi_x(y)$, which is closest to x and does not have the same tangent flat as x . The dimension of TP_q is strictly less than the dimension of TP_x and $\|x - q\| \leq \epsilon$, since the segment xq is contained in α . We now iterate, using q instead of x .

After a finite number of steps, we find a point $q \in \mathbb{P}$ such that $\pi_q(y) \in \alpha$. If q is a vertex of the polyhedron, $q = \pi_q(y)$ and at most $d - 1$ steps were necessary, showing that $\|x - q\| \leq (d - 1)\epsilon$. If q is not a vertex, $\|q - \pi_q(y)\| \leq \epsilon$ and at most $d - 2$ steps were necessary, showing that $\|x - q\| \leq (d - 2)\epsilon$. In all cases, $\|x - \pi_q(y)\| \leq (d - 1)\epsilon$ and applying triangular inequality, we get

$$\begin{aligned} \|x - \pi_q(z)\| &\leq \|x - \pi_q(y)\| + \|\pi_q(y) - \pi_q(z)\| \\ &\leq (d - 1)\epsilon + \|y - z\| \\ &\leq d\epsilon, \end{aligned}$$

and $x \in \text{ExtendedContact}(\nu, d\epsilon)$ as claimed. \square

4 Bounding the number of simplices

In this section, we bound the number of Delaunay k -simplices induced by λ -sparse ϵ -samples S of the polyhedron \mathbb{P} . A key step in our proof is to introduce a sample of the essential medial axis $M \subseteq \mathcal{M}^*$ consisting of $m = O(n^{\frac{d-1}{p}})$ points.

4.1 Sampling the medial axis. An ϵ -sample of the essential medial axis is a subset $M \subseteq \mathcal{M}^*$ such that every point $x \in \mathcal{M}^*$ has a point $z \in M$ at distance no more than ϵ , $\|x - z\| \leq \epsilon$. To construct such a sample, we consider a maximal collection of m non-overlapping d -balls $b_i = B(z_i, \frac{\epsilon}{2})$ whose centers z_i lie on the essential medial axis \mathcal{M}^* . Because the collection is maximal, no ball with center $z \in \mathcal{M}^*$ and radius $\frac{\epsilon}{2}$ can be added to the collection without overlapping $\bigcup b_i$. This implies that the set of centers z_i is an ϵ -sample of the essential medial axis, which we denote by M . On the other hand, the patches $\gamma_i = B(z_i, \frac{\epsilon}{2}) \cap \mathcal{M}$ pack the enlarged essential medial axis $\mathcal{M}_{\epsilon/2}^* = \{x \in \mathcal{M} \mid d(x, \mathcal{M}^*) \leq \epsilon/2\}$ and

$$m \min_i \text{vol}_{d-1}(\gamma_i) \leq \text{vol}_{d-1}(\mathcal{M}_{\epsilon/2}^*).$$

As ϵ tends to 0, the ratio between the two $(d-1)$ -dimensional volume of \mathcal{M}^* and $\mathcal{M}_{\epsilon/2}^*$ tends to 1 and for ϵ small enough, $\text{vol}_{d-1}(\mathcal{M}_{\epsilon/2}^*) \leq 2 \text{vol}_{d-1}(\mathcal{M}^*)$. Applying Inequality (2.1) yields to an upper bound on the number of points m in M :

$$m \leq 2^d \frac{\text{vol}_{d-1}(\mathcal{M}^*)}{w_{\mathbb{P}}} \epsilon^{-(d-1)}$$

In short, $m = O(\epsilon^{-(d-1)})$. We now establish that size of S is at least some constant times the p -th power of one over ϵ , $n = \Omega(\epsilon^{-p})$.

LEMMA 4.1. *Let S be a λ -sparse ϵ -sample of a p -polyhedron \mathbb{P} . For ϵ small enough, the number of points n in S satisfies*

$$2^{-(p+1)} \frac{\text{vol}_p(\mathbb{P})}{v_p} \epsilon^{-p} \leq n$$

Proof. Let F be a face of \mathbb{P} that spans a k -flat H . Let δ be the unit k -ball centered at the origin of \mathbb{R}^d and parallel to H . Consider a maximal collection of l non-overlapping k -balls $\delta_i = B(x_i, \epsilon) \cap H$ contained in F . Each ball δ_i contains at least a sample point of S . Thus, $l \leq \text{card}(S \cap F)$. On the other hand, the k -balls $B(x_i, 2\epsilon) \cap H$ cover

$$F - \epsilon\delta = \{x \in F \mid B(x, \epsilon) \cap H \subseteq F\}.$$

Translating this in terms of k -dimensional volume, we get $\text{vol}_k(F - \epsilon\delta) \leq l v_k (2\epsilon)^k$. Eliminating l and summing over all faces F of \mathbb{P} , we get

$$\sum_F \frac{\text{vol}_k(F - \epsilon\delta)}{v_k 2^k} \epsilon^{-k} \leq n.$$

The claim follows for ϵ small enough. \square

We combine our upper bound on the number of points m in M and our lower bound on the number of sample points n and get the following lemma:

LEMMA 4.2. *For ϵ small enough, there exists an ϵ -sample M of the essential medial axis with $m \leq C(\mathbb{P}) n^{\frac{d-1}{p}}$ points.*

4.2 Proof of the Main Theorem. We now give the proof of our main theorem. We consider the sparse ϵ -sample M of the medial axis with m points defined in the previous section. For $k \geq 1$, we map each Delaunay k -simplex σ to a point $z \in S \cup M$ and use z to locate the vertices of σ . To explain this, let b be one of the Delaunay ball that circumscribes σ . The vertices of σ are located on the boundary of b . By the fundamental property of Delaunay balls, b has the property to be ϵ -almost empty. We apply Lemma 3.1:

1. If b is trivial, by (i) of Lemma 3.1, there exists a point $x \in \mathbb{P}$ such that $b \cap \mathbb{P} \subseteq B(x, \epsilon)$. Because S is an ϵ -sample of \mathbb{P} , there exists a point $z \in S$ with $\|x - z\| \leq \epsilon$ and

$$\text{Vert } \sigma \subseteq B(z, 2\epsilon).$$

2. If b is ordinary, we apply (ii) of Lemma 3.1, combined with Lemma 3.2 and Lemma 3.3 to find that there exists a point $z \in M$ such that its associated medial ball ν satisfies

$$\text{Vert } \sigma \subseteq \text{ExtendedContact}(\nu, d\epsilon).$$

To summarize, any simplex in the Delaunay triangulation has its vertices contained either in a ball $B(z, 2\epsilon)$ with $z \in S$ or in the extended $(d\epsilon)$ -contact of a medial ball whose center belongs to M . Because S is λ -sparse, $B(z, 2\epsilon)$ contains at most λ sample points. The number of k -simplices σ that we can form by picking $k+1$ vertices among those λ points is $\binom{\lambda}{k+1}$. The extended $(d\epsilon)$ -contact is the union of at most f balls of radius $d\epsilon$ and therefore contains at most λf points. The number of k -simplices σ that we can form by picking $k+1$ vertices among those $f\lambda$ points is $\binom{f\lambda}{k+1}$. Therefore, the number of Delaunay k -simplices is bounded by

$$\binom{\lambda}{k+1} n + \binom{f\lambda}{k+1} m,$$

which concludes the proof of our main theorem.

5 Conclusion

In this paper, we proved that a size n sparse sample of a fixed p -polyhedron in \mathbb{R}^d for $p < d$ has a Delaunay

triangulation with size $O(n^{(d-1)/p})$ as n goes to infinity. In particular, this gives a linear bound for a $(d-1)$ -dimensional polyhedron in \mathbb{R}^d . This result is, to our knowledge, the first result of this kind for dimension greater than three. We see several directions in which this result could be improved.

The sparse ϵ -sample definition we use here is more restrictive than the definition in previous papers for the three dimensional case, in that we require all faces of all dimensions to be densely sampled and not only faces with highest dimension. If we sample only p -dimensional faces of the polyhedron, a first observation is that a Delaunay ball with radius r is no longer ϵ -almost empty but only $\sqrt{2r\epsilon}$ -almost empty. We conjecture, however, that the $O(n^{(d-1)/p})$ bound is still achievable in this setting.

Similar results in the random sampling model or when noise is added to the sample points should be attainable.

Another question of interest would be to bound the size of the Delaunay triangulation of sets of points sampled on or near smooth manifolds of dimension $2 \leq p \leq d-1$. Our results, and the prior work in dimension three, suggest that it may be possible to improve on the worst case bounds.

Acknowledgments

This research was initiated at the McGill-INRIA Workshop on Computational Geometry in Computer Graphics, February 2006, co-organized by H. Everett, S. Lazard, and S. Whitesides, and held at the Bellairs Research Institute of McGill University. The authors thank all the workshop participants for fruitful discussions. The authors would also like to thank Herbert Edelsbrunner for his useful comments.

References

- [1] D. Attali and J.-D. Boissonnat. A linear bound on the complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete and Computational Geometry*, 31(3):369–384, 2004.
- [2] D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces: the smooth case. In *Proc. of the 19th ACM Symposium on Computational Geometry*, pages 201–210, 2003.
- [3] D. Avis, D. Bremner, and R. Seidel. How good are convex hull algorithms? *Computational Geometry: Theory and Applications*, 7(5-6):265–306, 1997.
- [4] T. M. Chan. Output-sensitive results on convex hulls, extreme points, and related problems. In *Proc. 11th ACM Annual Symp. on Computational Geometry*, pages 10–19, 1995.
- [5] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. In *Proc. 11th Ann. Sympos. Comput. Geom.*, pages 319–326, 2006.
- [6] F. Chazal and A. Lieutier. Topology guaranteeing manifold reconstruction using distance function to noisy data. In *Proc. 11th Ann. Sympos. Comput. Geom.*, pages 112–118, 2006.
- [7] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete and Computational Geometry*, 10:377–409, 1993.
- [8] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4(1):387–421, 1989.
- [9] R.A. Dwyer. Higher-dimensional Voronoi diagrams in expected linear time. *Discrete and Computational Geometry*, 6:343–267, 1991.
- [10] J. Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete and Computational Geometry*, 30:109–132, 2003.
- [11] M. J. Golin and H.-S. Na. The probabilistic complexity of the Voronoi diagram of points on a polyhedron. In *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, pages 209–216, 2002.
- [12] M. J. Golin and H.-S. Na. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. *Computational Geometry: Theory and Applications*, 25:197–231, 2003.
- [13] P. McMullen. The maximum number of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.
- [14] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984.
- [15] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. In *ACM Symposium on the Theory of Computing*, pages 404–413, 1986.
- [16] ———. Small-dimensional linear programming and convex hulls made easy. *Discrete and Computational Geometry*, 6:423–434, 1991.