A tight bound for the Delaunay triangulation of points on a polyhedron

Nina Amenta *　　Dominique Attali †

UC Davis　　Gipsa-lab, CNRS Grenoble

Olivier Devillers ‡

INRIA Sophia-Antipolis

May 5, 2008

Abstract

We show that the Delaunay triangulation of a set of n points distributed nearly uniformly on a p-dimensional polyhedron (not necessarily convex) in d-dimensional Euclidean space is $O(n^{\frac{d-k+1}{p}})$, where $k = \lceil \frac{d+1}{p+1} \rceil$. This bound is tight, and improves on

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

†Dominique.Attali@lis.inpg.fr. Gipsa-lab, 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)

the prior upper bound for most values of p.

1 Introduction

Overview. The Delaunay triangulation of a set of points is a fundamental geometric data structure, used, in low dimensions, in surface reconstruction, mesh generation, molecular modeling, geographic information systems, and many other areas of science and engineering. In higher dimensions, it is well-known [10] that the complexity of the Delaunay triangulation of n points is $O(n^{\lceil \frac{d}{2} \rceil})$ and that this bound is achieved by distributions of points along one-dimensional curves such as the moment curve. But points distributed uniformly in $\mathbb{R}^d$, for instance inside a d-dimensional ball, have Delaunay triangulations of complexity $O(n)$; the constant factor is exponential in the dimension, but the dependence on the number of points is linear. In an earlier paper [1], we began to fill in the picture in between these two extremes, that is, when the points are distributed on a manifold of dimension $2 \leq p \leq d-1$. We began with the easy case of a *p-dimensional polyhedron* P, and showed that for a particular (probably overly restrictive) sampling model the size of the Delaunay triangulation is $O(n^{(d-1)/p})$.

Main result. Here as in [1], we consider a fixed *p-dimensional polyhedron* P in d-dimensional Euclidean space $\mathbb{R}^d$. Our point set S is a *sparse ε-sample* from P. Sparse ε-sampling requires the sampling to be neither too sparse nor too dense. Let n be the number of points in S. We consider the complexity of the Delaunay triangulation of S, as $n \to +\infty$, while P remains fixed. The main result in this paper is that the number of simplices of all dimensions is $O(n^{\frac{d-k+1}{p}})$ where $k = \lceil \frac{d+1}{p+1} \rceil$. The hidden constant factor depends, among other things, on the geometry of P, which is constant since P is fixed. This bound is tight. It directly improves on the bound established in [1] for all $1 \leq p < \frac{d-1}{2}$.

At the coarsest level, the idea of this proof is the same as that of [1]: we map Delaunay simplices to the medial axis and then use a packing argument to count them. The key new idea is the observation that when $k = \lceil \frac{d+1}{p+1} \rceil > 2$, the vertices of any Delaunay simplex, which must span $\mathbb{R}^d$, have to be drawn from more than two faces of P. This allows us to map Delaunay simplices to only the lower-dimensional submanifolds of the medial axis, induced by k or more faces. This idea is embodied in Corollary 6. We structure the definition carefully so that we can avoid making any non-degeneracy assumptions on the input polyhedron. An important technical innovation is the introduction of a new geometric structure, the *quasi medial axis*, which replaces the centers of tangent balls defining the medial axis with the centers of tangent annuli.

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 [6] 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 [5] established 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 P, they showed that the size of the Delaunay triangulation is $O(n)$ [2]. In a subsequent paper with Lieutier [3] they considered "generic" smooth 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 $\Omega(n^{3/2})$ [4]. Note that this surface is not generic and has a degenerate medial axis.

To the best of our knowledge, ours [1] is the only prior result for $d > 3$.

Outline of the proof. We begin by introducing the *ε-quasi k-medial axis*, a variant of the medial axis based on tangent annuli rather than tangent balls. We then define the part of the ε-quasi k-medial axis to which Delaunay simplices will be mapped: the *essential ε-quasi k-medial axis* (considering only the parts induced by k or more faces, and lopping off the parts which extend to infinity). By definition, this object has dimension at most $d-k+1$ and we prove that its $(d-k+1)$-dimensional volume is bounded from above by a constant that does not depend on ε. It follows that we can construct an ε-sample M of the essential ε-quasi k-medial axis with $m = O(\varepsilon^{-(d-k+1)})$ points.

We then turn our attention to assigning Delaunay simplices to the samples in M. We define the cover of a point z as

$$\mathrm{Cover}(z) = \bigcup_{x \in \Pi(z)} B(x, 5d\varepsilon),$$

where $\Pi(z)$ is the set of all orthogonal projections of z onto the planes supporting faces of P, and $B(x,r)$ is the ball centered at x with radius r. For $k = \lceil \frac{d+1}{p+1} \rceil$, we map each Delaunay simplex σ to a point $z \in M$ in such a way that the vertices of σ are contained in the cover of z; this is done by associating an annulus with a Delaunay simplex, and then "growing" the annulus to increase the number of its tangent points. Since the cover of each point $z \in M$ contains a constant number of points in S, each point $z \in M$ can only be charged for a constant number of Delaunay simplices. It follows that the size of the Delaunay triangulation is bounded from above by the size of M which is $m = O(\varepsilon^{-(d-k+1)})$. Since our point set S is a sparse ε-sample from a p-dimensional polyhedron, its cardinality is $n = \Omega(\varepsilon^{-p})$. Eliminating ε gives the $O(n^{\frac{d-k+1}{p}})$ upper bound. The lower bound is a straightforward construction of a polyhedron P.

2 Statement of Theorem

In this section, we introduce the setting for our result. We assume that the reader is familiar with notions of abstract and geometric simplicial complexes [11] and we use those notions to define Delaunay triangulations and polyhedra.

2.1 Delaunay triangulation.

We call any finite non-empty collection of points $\sigma \subseteq \mathbb{R}^d$ an *abstract simplex*. The points in σ will be referred to as the *vertices* of σ. 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 $\mathrm{Del}(S)$ of S is the nerve of the Voronoi regions. Specifically, an abstract simplex $\sigma \subseteq S$ belongs to the Delaunay triangulation iff the Voronoi regions of its vertices have a non-empty common intersection, $\bigcap_{s \in \sigma} V(s) \neq \emptyset$. Equivalently, the simplex σ is in the Delaunay triangulation iff there exists of a $(d-1)$-sphere, called *Delaunay sphere*, that passes through all vertices of σ and encloses no point of S. In this paper, we allow $d+2$ or more points in S to be co-spherical. These points may create Delaunay simplices with dimension higher than d. The *complexity* (or *size*) of the Delaunay triangulation is the total number of its simplices of all dimensions. We express this as a function of n, the number of points in S.

2.2 Polyhedron.

We call the underlying space of any geometric simplicial complex of dimension p a *p-dimensional polyhedron*. To define the faces of a polyhedron, we need some definitions. Given $X \subseteq \mathbb{R}^d$, we define the affine space $\mathrm{Aff}(X)$ spanned by X to consist of all points x of $\mathbb{R}^d$ such that

$$x = \sum_{i \in I} \alpha_i x_i,$$

for some finite set of integers I, points $x_i \in X$ and scalars α_i with $\sum_{i \in I} \alpha_i = 1$. The dimension of $\mathrm{Aff}(X)$ is the smallest amount of points that span $\mathrm{Aff}(X)$ minus one.

An affine space of dimension i is called a *i-plane*. Given a polyhedron P and a point $x \in P$, the *tangent plane* to P at x is the largest i-plane H through x such that a small neighborhood of x in H is contained in P. A *face* F of P is a maximal collection of points with identical tangent plane. If the dimension of the tangent plane is i, F is an i-face. The 0-faces are the *vertices* of P. Note that with our definition, faces are relatively open and every point $x \in P$ belongs to a unique face that we denote by F_x.

2.3 Sampling and Theorem.

Given a polyhedron $P \subseteq \mathbb{R}^d$, we say that a set of points $S \subseteq P$ is a *λ-sparse ε-sample* of P iff it satisfies the following two conditions:

`Density`: Every point x in P is at distance ε or less to a point in S lying on the closure of F_x. In other words,

$$\forall x \in P,\ \exists s \in S \cap \mathrm{cl}(F_x),\ \|x - s\| \leq \varepsilon;$$

`Sparsity`: Every closed d-ball with radius $5d\varepsilon$ contains at most λ points of S.

Note that our density condition implies that all faces of all dimensions are uniformly sampled, not just faces with highest dimension. Afterwards, we consider λ to be a constant. The number n of points in a λ-sparse ε-sample of a p-dimensional polyhedron is related to ε by $n = \Theta(\varepsilon^{-p})$. Thus, as n tends to infinity, ε tends to zero. We are now ready to state our main result:

Theorem 1 *Let S be a λ-sparse ε-sample of a p-dimensional polyhedron P in $\mathbb{R}^d$, and let n be the number of points in S. The Delaunay triangulation of S has size $O(n^{\frac{d-k+1}{p}})$ where $k = \lceil \frac{d+1}{p+1} \rceil$ and λ is a constant.*

Note that our result requires no non-degeneracy assumption, neither on P nor on S.

3 Essential quasi medial axes

In this section, we first define the *ε-quasi k-medial axis* $\mathcal{M}^k(P, \varepsilon)$ which is the key geometric object in our proof. We shall see that because P might be degenerate, we must introduce tools to identify the parts of $\mathcal{M}^k(P, \varepsilon)$ which have dimension $d-k+1$ or less very carefully. We also rigorously characterize a "finite" part of $\mathcal{M}^k(P, \varepsilon)$ whose dimension is $d-k+1$. We call this finite part the *essential* ε-quasi k-medial axis $\bar{\mathcal{M}}^k(P, \varepsilon)$, and we prove that its volume is bounded by a constant that does not depend on ε. We will see in Section 4 that given these definitions and tools, it is not too difficult to map Delaunay spheres to points on $\bar{\mathcal{M}}^k(P, \varepsilon)$.

3.1 Quasi medial axes

We start by defining ε-quasi k-medial axes. If X is a subset of $\mathbb{R}^d$, we denote the closure of X by $\mathrm{cl}(X)$ and write $\mathrm{Aff}(X)$ for the affine space spanned by X. We say that a $(d-1)$-sphere Σ is *tangent* to a face F at point x if both $\mathrm{cl}(F)$ and $\mathrm{Aff}(F)$ intersect Σ in a unique point x. In other words, letting z and r designate respectively the center and radius of Σ, we have

$$d(z, \mathrm{cl}(F)) = d(z, \mathrm{Aff}(F)) = r.$$

Since faces are relatively open, a sphere Σ tangent to a face F at point x may have an empty intersection with F, *i.e.* $\Sigma \cap F = \emptyset$. Note also that a sphere can be tangent to several faces of the polyhedron P at x and the face F_x is the unique one amongst them which contains x.

An annulus with center z, inner radius r and outer radius R is the set of points x whose distance to the center satisfies $r \leq \|x - z\| \leq R$. The boundary of an annulus consists of two $(d-1)$-spheres and we call the smallest one the *inner sphere* and the largest one the *outer sphere*. Extending what we just defined for spheres, we say that an annulus A is *tangent* to F at x if one of the two spheres bounding A is tangent to F at x (see Figure 6). Point x is called a *tangency point* of A. An annulus is P*-empty* if its inner sphere bounds a d-ball whose interior does not intersect P. An annulus is called ε*-thin* if the difference between the outer and inner radii squared is $R^2 - r^2 = \varepsilon^2$. Note that ε is not the width of the annulus.

Definition 2 *The* ε-quasi k-medial axis $\mathcal{M}^k(P, \varepsilon)$ *of* P *is the set of points* $z \in \mathbb{R}^d$ *for which the largest* P*-empty* ε*-thin annulus centered at* z *is tangent to at least* k *faces of* P.

Afterwards, we write $A(z, \varepsilon)$ for the largest P-empty ε-thin annulus centered at z. It should be observed that the 0-quasi 2-medial axis is a superset of the medial axis. Indeed, the medial axis of the polyhedron is the set of points $z \in \mathbb{R}^d$ for which $A(z, 0)$ touches the polyhedron in two points or more, while the 0-quasi 2-medial axis is the set of points z for which $A(z, 0)$ is tangent to two faces of P or more. Figure 1 pictures an example of ε-quasi 2-medial axis in $\mathbb{R}^2$.

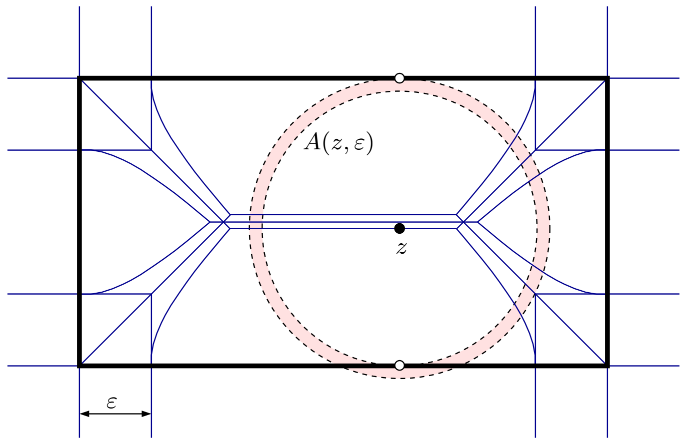

Figure 1: A rectangle and its ε-quasi 2-medial axis composed of 16 half-lines, 5 segments and 8 pieces of hyperbolas.

3.2 Stratification

Given a subset $X \subseteq \mathbb{R}^d$, a stratification of X is a filtration

$$\emptyset = X_{-1} \subseteq X_0 \subseteq \cdots \subseteq X_j = X$$

by subspaces such that the set difference $X_i \setminus X_{i-1}$ is -dimensional manifold, called the i-dimensional *stratum* of X. In particular, semi-algebraic sets admit a stratification [7] and since ε-quasi k-medial axes of polyhedra are piecewise semi-algebraic, they also admit a stratification. In this section, we give conditions under which a point $z \in \mathcal{M}^k(P, \varepsilon)$ belongs to a stratum of dimension $d - k + 1$ or less.

For this, let us break down $\mathcal{M}^k(P, \varepsilon)$ into pieces. Specifically, we define $\mathcal{S}^j(P, \varepsilon)$ as the set of points $z \in \mathbb{R}^d$ for which the annulus $A(z, \varepsilon)$ is tangent to exactly j faces of P. In particular,

$$\mathcal{S}^j(P, \varepsilon) = \mathcal{M}^j(P, \varepsilon) \setminus \mathcal{M}^{j+1}(P, \varepsilon).$$

Without loss of generality, we now focus on $\mathcal{S}^k(P,\varepsilon)$. As we will see shortly, $\mathcal{S}^k(P,\varepsilon)$ is not necessarily a $(d-k+1)$-dimensional stratum of $\mathcal{M}^k(P,\varepsilon)$ as one might have expected. We start by writing the equations that determine $\mathcal{S}^k(P,\varepsilon)$ locally around z. Since $A(z,\varepsilon)$ is tangent to exactly k faces $F_1,\ldots,F_k$, there exists $\delta > 0$ such that every face of the polyhedron not in $\{F_1,\ldots,F_k\}$ is at distance at least δ to the boundary of $A(z,\varepsilon)$. Using a compactness argument as in [9], it follows that for a point y close enough to z, the only faces possibly tangent to $A(y,\varepsilon)$ are $F_1,\ldots,F_k$. We set $e_i = -\varepsilon^2$ if F_i is tangent to the outer sphere of $A(z,\varepsilon)$ and $e_i = 0$ if F_i is tangent to the inner sphere of $A(z,\varepsilon)$. In a small neighborhood of z, $\mathcal{S}^k(P,\varepsilon)$ is thus determined by the following $k-1$ equations:

$$d(y,F_i)^2 - d(y,F_k)^2 + e_i - e_k = 0,$$

for $0 < i < k$. Each equation is the zero-set of a polynomial of second degree that identifies a quadric. It follows that $\mathcal{S}^k(P,\varepsilon)$ is piecewise a subset of the intersection of $k-1$ quadrics. In general, $k-1$ hypersurfaces meet at point z in a $(d-k+1)$-manifold. But in degenerate situations $\mathcal{S}^k(P,\varepsilon)$ can have dimension greater than $d-k+1$ as illustrated in Figure 2. We now give conditions under which such degeneracies cannot happen at z:

Lemma 3 *Suppose $z \in \mathcal{S}^k(P,\varepsilon)$ is the center of an annulus $A(z,\varepsilon)$ tangent to the polyhedron at k affinely independent points $x_1,\ldots,x_k$. Then, $\mathcal{S}^k(P,\varepsilon)$ is a $(d-k+1)$-manifold in a neighborhood of z. Furthermore, the tangent space to $\mathcal{S}^k(P,\varepsilon)$ at z is spanned by the set of vectors orthogonal to the $k-1$ vectors $x_k - x_1,\ldots,x_k - x_{k-1}$.*

PROOF. Let $F_i = F_{x_i}$ be the face to which x_i belongs. In a small neighborhood

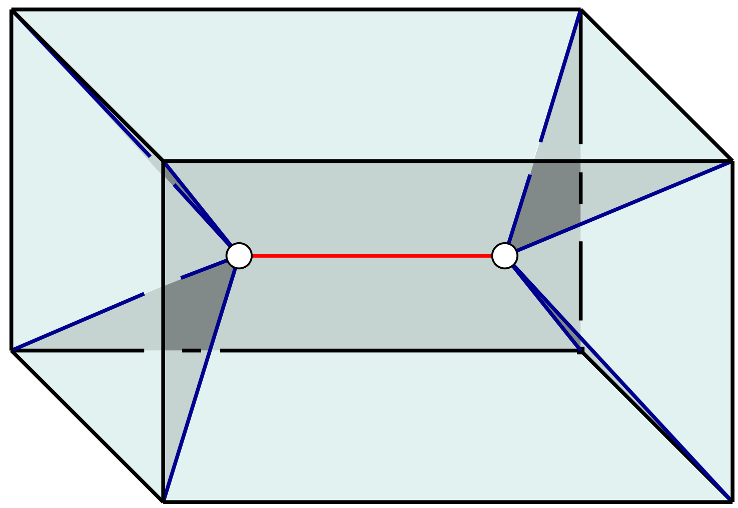

Figure 2: A box P. The set $\mathcal{S}^4(P, 0)$ is the segment connecting the hollow dots and has dimension 1.

of z, $\mathcal{S}^k(P, \varepsilon)$ coincides with the zero-set of the map $g : \mathbb{R}^d \to \mathbb{R}^{k-1}$ defined by $g(y) = (g^1(y), \ldots, g^{k-1}(y))$ with

$$g^i(y) = d(y, F_i)^2 - d(y, F_k)^2 + e_i - e_k.$$

The map g is differentiable and the ith component of the derivative of g at z is $Dg^i(z)(v) = 2(x_k - x_i) \cdot v$. We note that $\operatorname{rank}(Dg(z)) = k-1$ iff the k points $x_1, \ldots, x_k$ are affinely independent, which is true by assumption. Applying the implicit function theorem, we deduce that since the derivative $Dg(z) : \mathbb{R}^d \to \mathbb{R}^{k-1}$ has rank $k-1$, then $g^{-1}(0)$ is a $(d-k+1)$-dimensional manifold in a neighborhood of z. Furthermore, the tangent space of $g^{-1}(0)$ at z is precisely equal to the null space of the derivative $Dg(z)$, which is the set of vectors orthogonal to $x_k - x_1, \ldots, x_k - x_{k-1}$. $\square$

While the assumption that the tangency points of $A(z, \varepsilon)$ are independent is sufficient to show that z belongs to a stratum of dimension at most $d-k+1$, this condition is not necessary. We show that the stratum has dimension at most $d-k+1$ under the following weaker condition as well.

Definition 4 *We say that k faces $F_1, \ldots, F_k$ are* independent *if none of them is contained in the affine space spanned by the union of the others, that is for $1 \le i \le k$,*

$$F_i \quad \not\subseteq \quad \mathrm{Aff}(F_1 \cup \cdots \cup \widehat{F_i} \cup \cdots \cup F_k),$$

where the symbol $\widehat{}$ over F_i indicates that it is omitted in the union.

Lemma 5 *Suppose that $A(z, \varepsilon)$ is tangent to exactly k faces. If those k faces are independent, then $\mathcal{S}^k(P, \varepsilon)$ is a stratified space of dimension at most $d-k+1$ in a neighborhood of z.*

PROOF. We partition $\mathcal{S} = \mathcal{S}^k(P, \varepsilon)$ into k pieces possibly empty. More precisely, we write $\mathcal{S}_i = \mathcal{S}_i^k(P, \varepsilon)$ for the set of points $y \in \mathcal{S}$ which are the center of an annulus $A(y, \varepsilon)$ tangent to exactly k faces and whose tangency points span a space of dimension i. Thus we have $\mathcal{S} = \bigcup_i \mathcal{S}_i$. Each piece $\mathcal{S}_i$ is a semi-algebraic set and therefore admits a stratification. All we need to prove is that in a small neighborhood U of z, each stratified space $\mathcal{S}_i$ has dimension at most $d-k+1$ for all $0 \le i \le k$. By Lemma 3, we already know that $\mathcal{S}_k$ is a $(d-k+1)$-dimensional manifold. Let us assume $i < k$.

Let $F_1, \ldots, F_k$ be the k faces tangent to $A(z, \varepsilon)$. Given $y \in \mathbb{R}^d$, we denote the orthogonal projection of y onto $\mathrm{Aff}(F_i)$ by $x_i(y)$ (see Figure 3). Using the same compactness argument as before, there exists a small neighborhood U of z such that for every point $y \in U$, the only faces possibly tangent to the annulus $A(y, \varepsilon)$ are $F_1, \ldots, F_k$. Consider $y \in \mathcal{S}_i \cap U$. The tangency points of $A(y, \varepsilon)$ are $x_1(y), \ldots, x_k(y)$

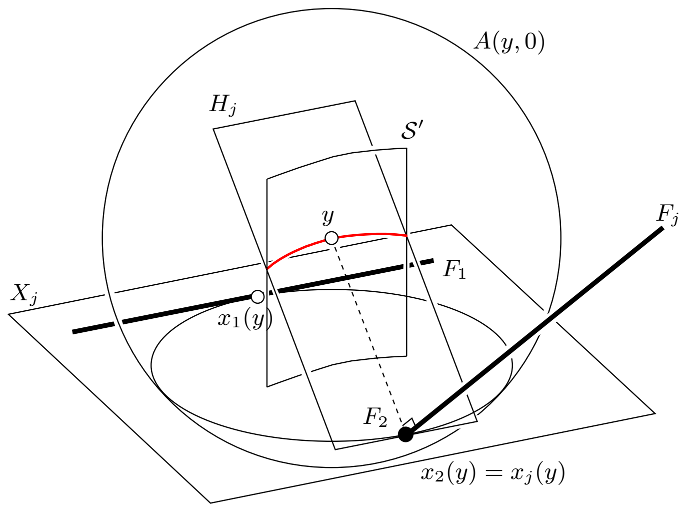

Figure 3: Notations for the proof of Lemma 5. $A(y,0)$ is tangent to F_1, F_2 and F_j.

and span an affine space of dimension i. Without loss of generality, we may assume that the first i tangency points $x_1(y), \ldots, x_i(y)$ are affinely independent. Let $P' = \mathrm{cl}(F_1) \cup \cdots \cup \mathrm{cl}(F_i)$ and write $\mathcal{S}'$ for the set of points which are the center of a P'-empty ε-thin annulus tangent to the i faces $F_1, \ldots, F_i$. By Lemma 3, $\mathcal{S}'$ is a $(d-i+1)$-manifold in a neighborhood of y. For $i < j \leq k$, $x_j(y)$ is an affine combination of $x_1(y), \ldots, x_i(y)$ and therefore belongs to

$$X_j \;=\; \mathrm{Aff}(F_1 \cup \cdots \cup \widehat{F_j} \cup \cdots F_{k-1}).$$

It follows that $F_j \cap X_j \neq \emptyset$ and we can define H_j as the set of points $w \in \mathbb{R}^d$ such that the nearest point to w on $\mathrm{Aff}(F_j)$ lies in $\mathrm{Aff}(F_j \cap X_j)$. Equivalently, H_j can be defined as the set of points equidistant to $\mathrm{Aff}(F_j)$ and $\mathrm{Aff}(F_j \cap X_j)$. It is an affine space whose dimension is $d - \dim F_j + \dim(F_j \cap X_j)$. We claim that in a neighborhood of y,

$$\mathcal{S}_i \;\subseteq\; \mathcal{S}' \cap H_{i+1} \cap H_{i+2} \cap \cdots \cap H_k.$$

By construction, $\mathcal{S}_i \subseteq \mathcal{S}'$ in a neighborhood of y. Since $x_1(y), \ldots, x_i(y)$ are affinely independent, there exists a neighborhood $U' \subseteq U$ of y such that for every point $y' \in U'$, the tangency points $x_1(y'), \ldots, x_i(y')$ are also affinely independent. Suppose $y' \in \mathcal{S}_i \cap U'$ and let us prove that $y' \in H_j$ for $i < j \leq k$. The dimension of the affine space spanned by $x_1(y'), \ldots, x_k(y')$ is i. It follows that $x_j(y')$ is an affine combination of $x_1(y'), \ldots, x_i(y')$ for $i < j \leq k$. Thus, $x_j(y') \in \mathrm{Aff}(F_j \cap X_j)$. Since by definition $x_j(y')$ is the orthogonal projection of y onto $\mathrm{Aff}(F_j)$, it follows that y' belongs to H_j. Therefore, $\mathcal{S}_i \cap U' \subseteq H_j$, for all $i < j \leq k$.

Let us prove that $\mathcal{S}' \cap H_{i+1} \cap H_{i+2} \cdots \cap H_k$ is a manifold of dimension at most $d - k + 1$ in a neighborhood of y. By Lemma 3, the normal space to $\mathcal{S}'$ is spanned by the $i - 1$ vectors $v_2 = x_1(y) - x_i(y), \ldots, v_i = x_{i-1}(y) - x_i(y)$. For $i + 1 \leq j \leq k$, we can always find a vector v_j in the normal space to H_j obtained by choosing v_j in the tangent plane to F_j and orthogonal to $F_j \cap X_j$. By construction, the $k - 1$ vectors $v_2, \ldots, v_k$ are linearly independent and all belong to the normal space of the intersection $\mathcal{S}' \cap H_{i+1} \cap H_{i+2} \cap \cdots \cap H_k$. It follows that the intersection is a manifold of dimension at most $d - k + 1$ and $\mathcal{S}_i$ is a stratified space of dimension at most $d - k + 1$.

□

We deduce immediately the following corollary:

Corollary 6 *Let $z \in \mathcal{M}^k(P, \varepsilon)$ and suppose that $A(z, \varepsilon)$ is tangent to j faces amongst which k faces are independent. Then, z lies on a i-dimensional stratum of $\mathcal{M}^k(P, \varepsilon)$ with $i \leq d - k + 1$.*

3.3 Essential part

In this section, we select a subset of the ε-quasi k-medial axis called the essential ε-quasi k-medial axis, $\bar{\mathcal{M}}^k(P,\varepsilon)$, and prove that all its strata have a finite volume bounded by a constant that does not depend on ε. We first define ε-essential points and show that the set of ε-essential points is contained in a d-ball $B(P)$ whose definition depends only on the geometry of P and does not depend on ε. For this, we need some definitions. We say that a hyperplane *supports* $X \subseteq \mathbb{R}^d$ if it has non-empty intersection with the boundary of X and empty intersection with the interior of X.

Definition 7 *A point z is* non ε-essential *if there exists a hyperplane supporting the convex hull of P and containing all faces tangent to $A(z,\varepsilon)$.*

It follows immediately that:

Lemma 8 *If the union of faces tangent to $A(z,\varepsilon)$ spans $\mathbb{R}^d$, then z is ε-essential.*

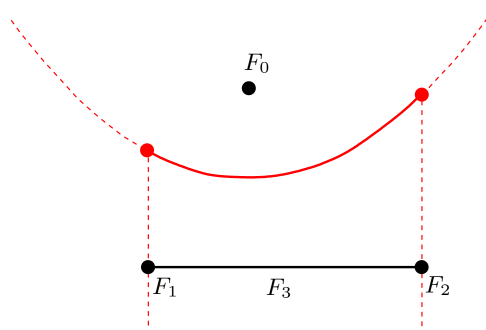

Figure 4: A polyhedron formed of four faces and its 0-quasi 2-medial axis. The set of 0-essential points is the closed piece of parabola consisting of points equidistant to F_0 and F_3.

Note that the set of ε-essential point is non-empty iff $\mathrm{Aff}(P) = \mathbb{R}^d$. Also, if two annuli $A(z,\varepsilon)$ and $A(z',\varepsilon)$ share the same set of faces, then z and z' are either both

ε-essential or both non ε-essential (see Figure 4). We start with a technical lemma that bounds the inner radius of an annulus based on the following observation: the only way a sphere Σ through a point q and tangent to a hyperplane H at x can have infinite radius is if either the distance of q to H is zero or the distance between q and x is infinite. Our technical lemma makes this idea precise and extends it to annuli:

Lemma 9 *Let A be an annulus tangent to a hyperplane H at point x and whose inner sphere does not enclose point q. Suppose q and the center of A lie on the same side of H. Let R and r be respectively the outer and inner radii of A. Suppose that there exist two scalars D and $h > 0$ such that $d(q, H) \geq h$, $\|x - q\| \leq D$ and $R - r \leq \frac{h}{2}$. Then, the inner radius of A satisfies $r \leq \frac{D^2}{h}$.*

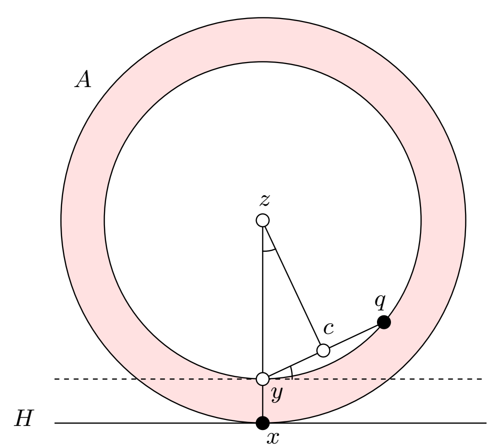

Figure 5: Notations for the proof of Lemma 9.

PROOF. We only consider what happens when the inner sphere of A passes through q and point x lies on the outer sphere of A (see Figure 5). Let y be the intersection the inner sphere of A with the segment connecting x to the center z of A. Let c be the midpoint of the segment yq. Since the angle between the two vectors $c - z$ and $q - z$ is equal to the angle between the vector $q - y$ and the hyperplane H, we have

$$\frac{\|q-y\|}{2r} = \frac{d(q,H)-(R-r)}{\|q-y\|}$$

The bound on r follows immediately. $\square$

Using this technical lemma, we are now able to establish that ε-essential points cannot be too far away from P, assuming ε is not too big.

Lemma 10 *Given a polyhedron P with diameter D, there exists a constant μ such that for $\varepsilon < D$, every ε-essential point is at distance μ or less to the polyhedron P.*

PROOF. We first give a characterization of ε-essential points. Let $\mathbb{S}^{d-1} = \{v \in \mathbb{R}^d \mid \|v\| = 1\}$. Given a face F of the polyhedron P, we associate to F the function $\delta_F : \mathbb{S}^{d-1} \to \mathbb{R}$ which maps every unit vector $v \in \mathbb{S}^{d-1}$ to

$$\delta_F(v) = \max\{\langle q - x, v\rangle \mid \forall x \in \mathrm{cl}(F), \forall q \in P\}.$$

Equivalently, $\delta_F(v)$ represents the distance between an extreme point in direction v on P and an extreme point in direction $-v$ on the closure of F. Note that δ_F is continous. Given a set of faces $\mathcal{F} = \{F_1, \ldots, F_k\}$, we introduce the map defined by

$$\delta_{\mathcal{F}}(v) = \frac{1}{k}\sum_{i=1}^{k} \delta_{F_i}(v).$$

It is continuous as a sum of continuous functions. Consider a point z which is the center of an annulus $A(z, \varepsilon)$ tangent to the set of faces $\mathcal{F} = \{F_1, \ldots, F_k\}$. We prove that z is non ε-essential iff there exists a unit vector v such that $\delta_{\mathcal{F}}(v) = 0$. Indeed, $\delta_{\mathcal{F}}(v) = 0$ iff $\delta_{F_i}(v) = 0$ for all i which happens iff the hyperplane supporting the convex hull of P and passing through an extreme point in direction v on P contains all faces F_i. We have just shown that a point z is ε-essential iff $\delta_{\mathcal{F}}(v) > 0$ for all unit vectors v. Since the map $\delta_{\mathcal{F}}$ is continuous and defined on a compact set, it attains a global minimum and this minimum is positive. We define

$$h = \frac{1}{2} \min_{\mathcal{F}} \min_{v} \delta_{\mathcal{F}}(v),$$

where v ranges over all unit vectors and $\mathcal{F}$ ranges over all subset of faces tangent to an annulus $A(z, \varepsilon)$ whose center z is ε-essential. We have $h > 0$.

Recall that z is the center of an annulus $A(z, \varepsilon)$ tangent to the set of faces $\mathcal{F} = \{F_1, \ldots, F_k\}$. Let x_i be the tangency point on the closure of face F_i and $v_i = \frac{z - x_i}{\|z - x_i\|}$. For every face F of the polyhedron, δ_F is uniformly continuous because defined on a compact set. Thus, there exists $\alpha_F > 0$ such that

$$\frac{\angle v_i v_j}{2} < \alpha_F \implies |\delta_F(v_i) - \delta_F(v_j)| < h.$$

We define $\alpha = \min_F \alpha_F$ over all faces F of the polyhedron. We now use the fact that for z sufficiently far away from P, the angle between v_i and v_j can be made arbitrarily small. Formally, let D be the diameter of P and let r be the distance of z to P. We have $\sin \frac{\angle v_i v_j}{2} \leq \frac{D}{r}$ and therefore

$$r > \frac{D}{\sin \alpha} \implies |\delta_F(v_i) - \delta_F(v_j)| < h, \tag{1}$$

for every face F and $1 \leq i, j \leq k$.

We are now ready to prove that for $\varepsilon < D$, every ε-essential point z is at distance $r \leq \max\{\frac{D^2}{h}, \frac{D}{\sin \alpha}\}$ to P. Suppose for a contradiction that this is not the case. Equivalently, suppose that the inner radius r of $A(z, \varepsilon)$ satisfies $r > \max\{\frac{D^2}{h}, \frac{D}{\sin \alpha}\}$. Writing

R for the outer radius of $A(z, \varepsilon)$, we have

$$R + r > 2r > \frac{2D^2}{h} > \frac{2\varepsilon^2}{h},$$

from which we deduce that

$$R - r = \frac{\varepsilon^2}{R + r} < \frac{h}{2}.$$

Let v_j be any of the k vectors $v_1, \ldots, v_k$. By definition of h, we have $\delta_{\mathcal{F}}(v_j) > 2h$ and therefore, at least one of the face F_i must satisfy $\delta_{F_i}(v_j) > 2h$. By Equation (1) and since $r > \frac{D}{\sin\alpha}$, we deduce that $\delta_{F_i}(v_i) > h$. Let H_i the hyperplane through F_i and normal to v_i. The inequality $\delta_{F_i}(v_i) > h$ implies that there exists a point $q \in P$ such that $d(q, H_i) \geq h$. Furthermore, $\|x_i - q\| \leq D$ and $R - r \leq \frac{h}{2}$. Therefore, we can apply Lemma 9 and get $r \leq \frac{D^2}{h}$, which leads to a contradiction. $\square$

Afterwards, $B(P)$ denotes the smallest ball containing the parallel body $P^\mu = \{x \in \mathbb{R}^d \mid d(x, P) < \mu\}$. We have just proved that $B(P)$ contains the set of ε-essential points for ε smaller than the diameter of P.

Definition 11 *The* essential ε-quasi k-medial axis*, $\bar{\mathcal{M}}^k(P, \varepsilon)$, is the set of ε-essential points lying on the i-dimensional strata of the ε-quasi k-medial axis for $i \leq d - k + 1$.*

We now prove that the i-dimensional stratum of the ε-quasi k-medial axis has an i-dimensional volume bounded by a constant that does not depend on ε. For this, we use a generalization of Crofton's formula that can be found in Santaló [12] on page 245. Writing M_ε for the i-dimensional stratum of the ε-quasi k-medial axis, we have

$$\mathrm{Vol}_i(M_\varepsilon) \quad = \quad \frac{O_{d-i} \cdots O_1}{O_d \cdots O_{i+1}} \int_{M_\varepsilon \cap H \neq \emptyset} N(M_\varepsilon \cap H) \, dH,$$

where the integration is over all $(d-i)$-planes H having a non-empty intersection with M_ε, $N(M_\varepsilon \cap H)$ denotes the number of points of the intersection $M_\varepsilon \cap H$ and O_j is the surface area of the j-dimensional unit sphere. The i-dimensional stratum M_ε of $\mathcal{M}^k(P,\varepsilon)$ can be decomposed in pieces, each piece being a subset of the intersection Q_ε of $d-i$ independent quadrics. Since a $(d-i)$-plane H is the intersection of i independent hyperplanes, it follows that $Q_\varepsilon \cap H$ is the solution of a system of d polynomial equations of degree two or one. By the higher-dimensional version of Bezout's theorem [8], the number of roots of a system of d polynomial equations in d variables is either infinite or the product of their degrees. It follows that the size of the intersection $Q_\varepsilon \cap H$ is either infinite or consists of at most 2^d points. Furthermore, the set of $(d-i)$-planes H for which $Q_\varepsilon \cap H$ is infinite has measure zero. Since the number of pieces forming M_ε can be bounded from above by a constant $c(P)$ that depends only on the number of faces of P, it follows that for ε smaller than the diameter of P,

$$\mathrm{Vol}_i(M_\varepsilon) \ \leq \ 2^d c(P)\, \frac{O_{d-i}\cdots O_1}{O_d\cdots O_{i+1}} \int_{B(P)\cap H\neq\emptyset} dH.$$

The integral on the right side is finite and represents the measure of all $(d-i)$-planes H that intersect the convex set $B(P)$ (see [12] page 233 for an expression of this integral). Hence, the right side of the above inequality does not depend on ε and we conclude that:

Lemma 12 *For ε smaller than the diameter of P, the i-dimensional stratum of the ε-quasi k-medial axis has an i-dimensional volume bounded by a constant, that does not depend on ε.*

4 Covering Delaunay spheres

The goal of this section is to prove that the intersection of a p-dimensional polyhedron P with any Delaunay sphere Σ is contained in the cover of some point z on the essential ε-quasi k-medial axis, for $k = \lceil \frac{d+1}{p+1} \rceil$. We first state crucial properties of Delaunay spheres and polyhedra before defining the cover of a point. The first property is induced by our sampling condition.

Definition 13 *We say that a sphere Σ with center z is* ε-almost P-empty *if* $\Sigma \subseteq A(z, \varepsilon)$.

Lemma 14 *Delaunay spheres are ε-almost P-empty.*

PROOF. For reader's convenience, we recall the proof given in [1]. Consider a Delaunay sphere Σ with center z. Let x be a point in P with minimum distance to z and let s be the sample point in $S \cap \mathrm{cl}(F_x)$ closest to x. Because of our sampling condition, $\|x - s\| \leq \varepsilon$ and therefore $s \in A(z, \varepsilon)$. Because Σ encloses no sample point, $\Sigma \subseteq A(z, \varepsilon)$. ◽

The second property concerns polyhedra.

Definition 15 *We say that a polyhedron P is* k-reductible *if for any collection of $k-1$ faces $\{F_1, \ldots, F_{k-1}\}$ of P, there exists a hyperplane that contains the union $\bigcup_{i=1}^{k-1} F_i$.*

Note that every polyhedron P that is k-reductible is also k'-reductible with $k' \leq k$.

Lemma 16 *Any p-dimensional polyhedron of $\mathbb{R}^d$ is $\lceil \frac{d+1}{p+1} \rceil$-reductible.*

PROOF. Let $k = \lceil \frac{d+1}{p+1} \rceil$. The dimension of the smallest affine space containing $k-1$ faces of P is bounded from above by the amount of affinely independent points that we can pick on each face minus 1. In other words, for any collection of $k-1$ faces $\{F_1, \ldots, F_{k-1}\}$ of P

$$\begin{aligned} \dim \mathrm{Aff}(\bigcup_{i=1}^{k-1} F_i) &\leq (k-1)(p+1) - 1 \\ &< (d+1) - 1. \end{aligned}$$

The claim follows. $\square$

We now define the cover of a point $z \in \mathbb{R}^d$. Writing $\pi_x(z)$ for the orthogonal projection of z onto the tangent plane of $x \in P$, we say that x is a critical point of the distance-to-z function if $\pi_x(z) = x$. We define $\chi(z, \varepsilon)$ as the set of critical points lying in $P \cap A(z, \varepsilon)$. Note that $\chi(z, \varepsilon)$ contains the tangency points of the annulus with the polyhedron but possibly other points of P located in the interior of $A(z, \varepsilon)$. Given a map $w : P \to \mathbb{R}^+$ that associates to each point $x \in P$ a positive real number $w(x)$, we define the cover of z as:

$$\mathrm{Cover}^w(z, \varepsilon) = \bigcup_{x \in \chi(z,\varepsilon)} B(x, w(x)\varepsilon),$$

Given a Delaunay sphere Σ, we show that it is possible to find a point z on the essential ε-quasi k-medial axis and a map w bounded from above by a constant in a such a way that $P \cap \Sigma \subseteq \mathrm{Cover}^w(z, \varepsilon)$. We prepare our result with a technical lemma, which says roughly that any point in $P \cap A(z, \varepsilon)$ must be close to a critical point in $\chi(z, \varepsilon)$. We then proceed in two steps, first finding a point in $\mathcal{M}^k(P, \varepsilon)$ and next in $\bar{\mathcal{M}}^k(P, \varepsilon)$.

Lemma 17 *For every point $x \in P \cap A(z, \varepsilon)$, there exists a point $y \in \chi(z, \varepsilon)$ in the closure of the face to which x belongs and such that*

$$\|x - y\| \leq (\dim F_x - \dim F_y + 1)\, \varepsilon.$$

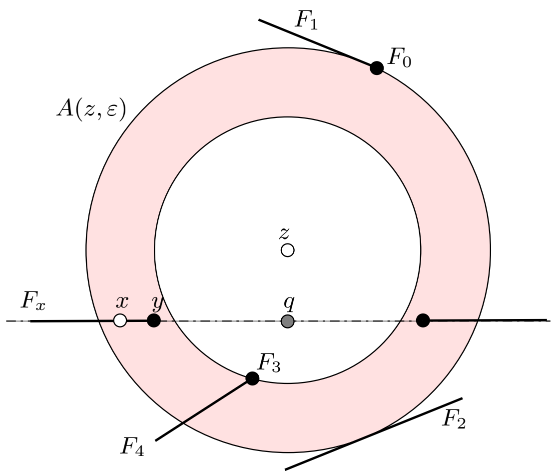

Figure 6: The annulus is tangent to the four faces F_0, F_1, F_2 and F_3. Notations for the proof of Lemma 17.

PROOF. The proof is by induction over the dimension $d_x = \dim F_x$ of the face F_x containing x. If $d_x = 0$, the result holds for $y = x$. Suppose $d_x > 0$ and let $q = \pi_x(z)$ be the orthogonal projection of z onto the tangent plane to F_x. We distinguish two cases: (1) if $q \in F_x$, the segment xq lies inside $A(z, \varepsilon)$ and therefore $\|x - q\| \leq \varepsilon$; (2) if $q \notin F_x$, we consider the point $y \in P$ on the segment xq, which is closest to x and does not have the same tangent plane as x (see Figure 6). Since the segment xy is contained in $A(z, \varepsilon)$, this implies $\|x - y\| \leq \varepsilon$. Furthermore, since y belongs to the boundary of the face to which x belongs, $d_y < d_x$. Therefore, we can apply our induction hypothesis to y and conclude. ◻

Lemma 18 *Consider a k-reductible polyhedron P that spans $\mathbb{R}^d$ and let $w_k(x) = \dim P - \dim F_x + 2k$. For every ε-almost P-empty sphere Σ, there exists a point $z \in \mathcal{M}^k(P, \varepsilon)$ such that*

$$\Sigma \cap P \subseteq \mathrm{Cover}^{w_k}(z, \varepsilon).$$

PROOF. For simplicity, we write $d_x = \dim F_x$. The proof is by induction over k. For $k = 1$, let z_1 be the center of Σ. The inner sphere of the annulus $A(z_1, \varepsilon)$ is tangent to the polyhedron in at least one point, showing that $z_1 \in \mathcal{M}^1(P, \varepsilon)$. Furthermore, $b \subseteq A(z_1, \varepsilon)$ and by Lemma 17, this implies that $b \cap P \subseteq \mathrm{Cover}^{w_1}(z_1, \varepsilon)$. Suppose the statement holds for $k = i$ and let us prove it for $k = i + 1$. By induction hypothesis, there exists a point $z \in \mathcal{M}^i(P, \varepsilon)$ such that $b \cap P \subseteq \mathrm{Cover}^{w_i}(z, \varepsilon)$. This means in particular that we can find i faces $F_1, \ldots, F_i$, each either tangent to the inner sphere of $A(z, \varepsilon)$ or to the outer sphere of $A(z, \varepsilon)$. Since P is i-reductible, we can find a hyperplane H that contains $\bigcup_{j=1}^{i} F_j$. Let L^+ be the half-line with origin z, orthogonal to H and avoiding H. Keeping the intersection $H \cap A(z, \varepsilon)$ fixed, we move the center z on L^+ until either the inner sphere of $A(z, \varepsilon)$ or the outer sphere $A(z, \varepsilon)$ becomes tangent to a new face. Let z' be the point of L^+ at which this happens. If z' does not exist, we repeat the search replacing L^+ by the half-line with origin z, orthogonal to H and intersecting H. In any case, the point z' must exist because we assumed that no hyperplane contains the polyhedron P. We have $z' \in \mathcal{M}^{i+1}(P, \varepsilon)$.

To establish the statement for $k = i + 1$, we only need to prove that for every $x \in \chi(z, \varepsilon)$, there exists $x' \in \chi(z', \varepsilon)$ such that

$$\|x - x'\| \leq (d_x - d_{x'} + 2)\varepsilon,$$

which will entail that

$$\mathrm{Cover}^{w_i}(z, \varepsilon) \subseteq \mathrm{Cover}^{w_{i+1}}(z', \varepsilon).$$

Let H^+ be the closed half-space that H bounds and that contains points at infinity on the half-line L^+. Let H^- be the complement of H^+. We consider two cases:

1. If $x \in H^+$, by construction, the annulus $A(y, \varepsilon)$ remains P-empty as the center y moves on the segment zz'. It follows that x cannot escape the annulus $A(y, \varepsilon)$ and therefore $x \in A(z', \varepsilon)$. By Lemma 17, there exists a point $x' \in \chi(z', \varepsilon)$ such that $\|x - x'\| \leq (d_x - d_{x'} + 1)\varepsilon$.

2. If $x \in H^-$, we consider the intersection D_y of $A(y, \varepsilon)$ with the tangent plane to the face containing x. Because $x \in \chi(z, \varepsilon)$, D_z is a d_x-ball of radius less than ε with center x. The restrictions of D_y to H^- form a nested family of sets. In particular, $D_y \cap H^- \subseteq D_z \cap H^-$, for all points $y \in zz'$. As the point y moves on the segment zz', the closure of the face F containing x cannot escape $D_y \cap H^-$. Indeed, if it were the case, it would mean that y passed a point at which the outer sphere of $A(y, \varepsilon)$ becomes tangent to F or to a face on the boundary of F, which is impossible unless $y = z'$. Therefore, we can always find a point $x'' \in D_z \cap D_{z'} \cap \mathrm{cl}(F)$. Because D_z is a ball of radius less than ε with center x, $\|x - x''\| \leq \varepsilon$. By Lemma 17, there exists a point $x' \in \chi(z', \varepsilon)$ such that $\|x'' - x'\| \leq (d_{x''} - d_{x'} + 1)\varepsilon$. Combining these two inequalities with $d_{x''} \leq d_x$, we get $\|x - x'\| \leq (d_x - d_{x'} + 2)\varepsilon$.

$\square$

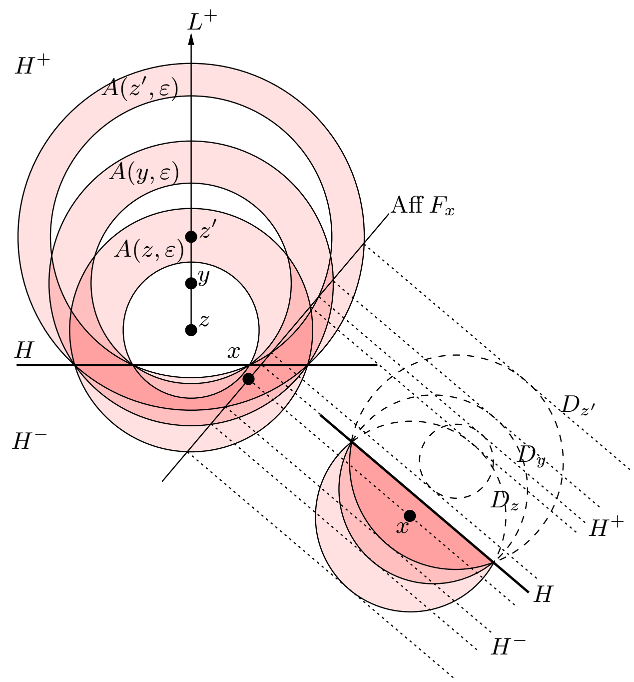

Figure 7: On the upper left, three annuli centered at z, y and z' that share the same intersection with a hyperplane H. On the lower right, intersection of the three annuli with the tangent space passing through x. The restriction of those intersections to H^- are nested.

Lemma 19 *Let P be a k-reductible polyhedron that spans $\mathbb{R}^d$. For every point $z \in \mathcal{M}^k(P, \varepsilon)$, there exists a point $\bar{z} \in \bar{\mathcal{M}}^k(P, \varepsilon)$ such that*

$$\mathrm{Cover}^{w_k}(z, \varepsilon) \subseteq \mathrm{Cover}^{w_{k+d}}(\bar{z}, \varepsilon).$$

PROOF. The proof is omitted. The intuition is that after at most d steps similar to those described in the previous Lemma, we are able to find a point $\bar{z}$ which is the center of an annulus tangent to a set of faces that span $\mathbb{R}^d$ and amongst which k faces are independent. Furthermore, the cover of z weighted by w_k is contained in the cover of $\bar{z}$ weighted by w_{k+d}. By Corollary 6 and Lemma 8, $\bar{z}$ belongs to the essential ε-quasi

k-medial axis. $\square$

We combine Lemma 18 and Lemma 19 and get the following lemma:

Lemma 20 *Consider a k-reductible polyhedron P that spans $\mathbb{R}^d$. For every ε-almost P-empty sphere Σ, there exists a point $z \in \bar{\mathcal{M}}^k(P,\varepsilon)$ such that*

$$\Sigma \cap P \subseteq \mathrm{Cover}^{4d+1}(z,\varepsilon).$$

In the next section, it will be convenient to use a slightly different notion of cover. Let $\Pi(z)$ be the set of orthogonal projections of z onto the planes supporting faces of P. We define the extended cover of point z as

$$\mathrm{ExtendedCover}^w(z,\varepsilon) = \bigcup_{x\in\Pi(z)} B(x, w(x)\varepsilon).$$

Lemma 21 *For every points z and z' with $\|z - z'\| \leq \varepsilon$:*

$$\mathrm{Cover}^w(z,\varepsilon) \subseteq \mathrm{ExtendedCover}^{w+1}(z',\varepsilon).$$

PROOF. Recalling that $\pi_x(z)$ is the orthogonal projection of z onto the tangent plane to P at x, we have $\|\pi_x(z) - \pi_x(z')\| \leq \|z - z'\| \leq \varepsilon$. The claim follows immediately. $\square$

5 Size of Delaunay triangulation

In this section, we collect results from previous sections and establish our upper bound on the number of Delaunay simplices. We then prove that our bound is tight. We recall that the number of points in a λ-sparse ε-sample S of a p-dimensional polyhedron P is $n = \Theta(\epsilon^{-p})$ and that the i-faces of P have $\Theta(\varepsilon^{-i})$ points of S [1].

5.1 Upper bound

Without loss of generality, we may assume $\mathrm{Aff}(P) = \mathbb{R}^d$. An ε-sample of the essential ε-quasi k-medial axis is a subset $M \subseteq \bar{\mathcal{M}}^k(P,\varepsilon)$ such that every point $x \in \bar{\mathcal{M}}^k(P,\varepsilon)$ is at distance no more than ε to a point $z \in M$, $\|x - z\| \leq \varepsilon$. We claim that we can construct an ε-sample M of $\bar{\mathcal{M}}^k(P,\varepsilon)$ in such a way that the i-dimensional stratum of the essential ε-quasi k-medial axis receives $O(\varepsilon^{-i})$ points of M and the number of points in M is $m = O(\varepsilon^{-(d-k+1)})$. This is a consequence of Lemma 12 which says that the i-dimensional volume of the i-dimensional stratum of $\bar{\mathcal{M}}^k(P,\varepsilon)$ is bounded by a constant that does not depend on ε. To establish our upper bound, we map each Delaunay simplex $\sigma \in \mathrm{Del}(S)$ to a point $z \in M$. Consider a Delaunay sphere Σ passing through the vertices of σ. By Lemma 14, Delaunay spheres are ε-almost P-empty. We can therefore combine Lemma 16, Lemma 20 and Lemma 21 and get that for $d \geq 2$ and $k = \lceil \frac{d+1}{p+1} \rceil$, there exists a point $z \in M$ such that

$$\Sigma \cap P \subseteq \mathrm{ExtendedCover}^{5d}(z,\varepsilon)$$

The extended cover of z is a union of d-balls of radius $5d\varepsilon$, one for each face of the polyhedron and therefore, it contains a constant number of points of S. It follows that the number of simplices that we can form by picking points in the extended cover of z is constant. Hence, each point $z \in M$ is charged with a constant number of Delaunay simplices and using $n = \Omega(\varepsilon^{-p})$, we get that the number of Delaunay simplices is

$$O(m) = O(\varepsilon^{-(d-k+1)}) = O(n^{\frac{d-k+1}{p}}),$$

where $k = \lceil \frac{d+1}{p+1} \rceil$.

5.2 The bound is tight

We now prove that our upper bound is tight. Consider a set of $d+1$ affinely independent points that we partition into $k = \lceil \frac{d+1}{p+1} \rceil$ groups $Q_1, \ldots, Q_k$ in such a way that (1) no group Q_i has more than $p+1$ points; (2) at least one of the group has $p+1$ points. Writing q_i for the dimension of the affine space spanned by Q_i, we have

$$\sum_{i=1}^{k} q_i = d - k + 1. \tag{2}$$

Letting C_i be the convex hull of Q_i, we consider the polyhedron $P = \bigcup_{i=1}^{k} C_i$ and S a λ-sparse ε-sample of P. The simplex $\sigma = \{s_1, \ldots, s_k\}$ obtained by picking a sample point $s_i \in S \cap C_i$ for $1 \leq i \leq k$ belongs to the Delaunay triangulation. Indeed, since the points $s_1, \ldots, s_k$ are affinely independent, there exists a $(d-1)$-sphere Σ tangent to P at s_i for $1 \leq i \leq k$, whose center lies on the 0-quasi k-medial axis of P. By construction, this sphere encloses no sample point of S in its interior, showing that σ is a Delaunay simplex. Since C_i contains $\Omega(\varepsilon^{-q_i})$ points of S, the amount of Delaunay simplices that we can construct this way is at least

$$\Omega(\varepsilon^{-q_1} \times \cdots \times \varepsilon^{-q_k}) = \Omega(\varepsilon^{-(d-k+1)}) = \Omega(n^{\frac{d-k+1}{p}}).$$

6 Conclusion

This paper answers only the first of many possible questions about the complexity of the Delaunay triangulations of points distributed nearly uniformly on manifolds. Similar bounds for smooth surfaces rather than polyhedra would be of more practical interest. The proof in this paper seems to relay on some properties specific to polyhedra, particularly that sample points on k faces are needed to form a simplex. On the other hand, the tight bound seems to be "right", at least in the sense that it agrees with the well-known bounds in the cases $p = 1$ and $p = d$.

Acknowledgments

We gratefully acknowledge André Lieutier for helpful suggestions. The work of the first author was supported by NSF awards CCF–0331736 and CCF–0635250.

References

[1] N. Amenta, D. Attali, and O. Devillers. Complexity of Delaunay triangulation for points on lower-dimensional polyhedra. In *Proc. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, New-Orleans, Lousiana, USA, January 7–9 2007.

[2] 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.

[3] 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.

[4] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Discrete and Computational Geometry*, 30:109–132, 2003.

[5] 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.

[6] 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.

[7] Gunnells. Stratified spaces twigs. TWIGS talks. `http://www.math.umass.edu/%7Egunnells/talks/singspc.pdf`.

[8] S. Katz. *Enumerative Geometry and String Theory*, volume 32 of *IAS/Park City Mathematical Subseries*. AMS, Student Mathematical Library edition, 2006.

[9] A. Lieutier. Any open bounded subset of has the same homotopy type as its medial axis. *Computer-Aided Design*, 36:1029–1046, 2004.

[10] P. McMullen. The maximum number of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[11] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984.

[12] L. Santaló. *Integral Geometry and Geometric Probablity*. Cambridge University Press, Cambridge Mathematical Library edition, 1976.