

A Linear Bound on the Complexity of the Delaunay Triangulation of Points on Polyhedral Surfaces*

Dominique Attali
LIS, ENSIEG
Domaine Universitaire, BP 46,
38402 Saint Martin d'Hères Cedex, France
Dominique.Attali@inpg.fr

Jean-Daniel Boissonnat
INRIA,
2004 Route des Lucioles, BP 93,
06904 Sophia-Antipolis, France
Jean-Daniel.Boissonnat@sophia.inria.fr

ABSTRACT

Delaunay triangulations and Voronoi diagrams have found numerous applications in surface modeling, surface mesh generation, deformable surface modeling and surface reconstruction. Many algorithms in these applications begin by constructing the three-dimensional Delaunay triangulation of a finite set of points scattered over a surface. Their running-time therefore depends on the complexity of the Delaunay triangulation of such point sets.

Although the complexity of the Delaunay triangulation of points in \mathbb{R}^3 may be quadratic in the worst-case, we show in this paper that it is only linear when the points are distributed on a fixed number of well-sampled facets of \mathbb{R}^3 (e.g. the facets of a polyhedron). Our bound is deterministic and the constants are explicitly given.

Categories and Subject Descriptors

I.3.5 [Computing Methodologies]: Computational Geometry and Object Modeling—*Curve, surface, solid, and object representations*; F.2.2 [Theory of Computation]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Theory, Algorithms, Performance

Keywords

Delaunay triangulation, Voronoi diagram, complexity, polyhedral surfaces, sample, surface modeling, reconstruction.

1. INTRODUCTION

Delaunay triangulations and Voronoi diagrams are among the most thoroughly studied geometric data structures in computational geometry. Recently, they have found many applications in surface

*The work by the second author has been partially supported by the IST Programme of the EU under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SM'02 June 17-21, 2002, Saarbrücken, Germany.

Copyright 2002 ACM 1-58113-506-8/02/0006 ...\$5.00.

modeling, surface mesh generation [12], deformable surface modeling [22, 16], medial axis approximation [4, 9, 21], and surface reconstruction [3, 1, 8, 2, 7, 6]. Many algorithms in these applications begin by constructing the three-dimensional Delaunay triangulation of a finite set of points scattered over a surface. Their running-time therefore depends on the complexity of the Delaunay triangulation of such point sets.

It is well known that the complexity of the Delaunay triangulation of n points in \mathbb{R}^d , i.e. the number of its simplices, can be $\Omega(n^{\lceil \frac{d}{2} \rceil})$ [10]. In particular, in \mathbb{R}^3 , the number of tetrahedra can be quadratic. This is prohibitive for applications where the number of points is in the millions, which is routine nowadays. Although it has been observed experimentally that the complexity of the Delaunay triangulation of well-sampled surfaces is linear (see e.g. [8, 13]), no result close to this bound has been obtained yet. Our goal is to exhibit practical geometric constraints that imply subquadratic and ultimately linear Delaunay triangulations. Since output-sensitive algorithms are known for computing Delaunay triangulations [11], better bounds on the complexity of the Delaunay triangulation would immediately imply improved bounds on the time complexity of computing the Delaunay triangulation.

First results on Delaunay triangulations with low complexity have been obtained by Dwyer [14, 15] who proved that, if the points are uniformly distributed in a ball, the *expected* complexity of the Delaunay triangulation is only linear. Recently, Erickson [17, 18] investigated the complexity of three-dimensional Delaunay triangulations in terms of a geometric parameter called the spread, which is the ratio between the largest and the smallest interpoint distances. He proved that the complexity of the Delaunay triangulation of any set of n points in \mathbb{R}^3 with spread Δ is $O(\Delta^3)$.

Despite its practical importance, the case of points distributed on a surface has not received much attention. A first result has been obtained by Golin and Na [19]. They proved that the expected complexity of 3D Delaunay triangulations of random points on any fixed *convex* polytope is $\Theta(n)$. Very recently, they extended their proof to the case of general polyhedral surfaces of \mathbb{R}^3 and obtained a $O(n \log^4 n)$ bound on the expected complexity of the Delaunay triangulation [20]. Deterministic bounds have also been obtained. Attali and Boissonnat [5] proved that, for any fixed polyhedral surface S , any so-called “light-uniform ε -sample” of S of size n has only $O(n^{7/4})$ Delaunay tetrahedra. If the surface is convex, the bound reduces to $O(n^{3/2})$. Applied to a fixed C^2 uniformly-sampled surface, the result of Erickson mentioned above shows that the Delaunay triangulation has complexity $O(n^{3/2})$. This bound is tight in the worst-case. It should be noticed however that Erickson’s definition of a uniform sample is rather restrictive and does not allow two points to be arbitrarily close (in which case, the spread

would become infinite).

In this paper, we consider the case of points distributed on a fixed number of planar facets in \mathbb{R}^3 , e.g. the facets of a given polyhedron. Under a mild uniform sampling condition, we show that the complexity of the Delaunay triangulation of the points is linear. Our bound is deterministic and the constants are explicitly given.

2. DEFINITIONS AND NOTATIONS

2.1 Voronoi diagrams and Delaunay triangulations

Let $E = \{p_1, \dots, p_n\}$ be a set of points of \mathbb{R}^d . The Voronoi cell of p_i is

$$V(p_i) = \{x \in \mathbb{R}^d : \|x - p_i\| \leq \|x - p_j\| \quad \forall j = 1, \dots, n\}$$

where $\|x - y\|$ denotes the Euclidean distance between the two points x, y of \mathbb{R}^d . The collection of Voronoi cells is called the *Voronoi diagram* of E , denoted $\text{Vor}(E)$. The Delaunay triangulation of E , denoted $\text{Del}(E)$ is the dual complex of $\text{Vor}(E)$ (see Figure 1). If there is no sphere passing through $d + 2$ points of E , $\text{Del}(E)$ is a simplicial complex that can be obtained from $\text{Vor}(E)$ as follows. If E' is a subset of points of E whose Voronoi cells have a non empty intersection, the convex hull $\text{conv}(E')$ is a Delaunay face and all Delaunay faces are obtained this way. It is well known that the balls circumscribing the d -simplices in $\text{Del}(E)$ cannot contain a point of E in their interior. The *complexity* of $\text{Del}(E)$ is the number of its faces, which is also the number of faces of the dual Voronoi diagram.

A ball or a disk is said to be *empty* iff its interior contains no point of E . We also say that a sphere is *empty* if the associated ball is empty.




Figure 1: Voronoi diagram of a set of points on the left and its dual Delaunay triangulation on the right.

2.2 Notations

For a curve Γ , we denote by $l(\Gamma)$ its length. For a portion of a surface R , we denote by $a(R)$ its area, and by ∂R its boundary. We further denote by $B(x, r)$ ($\Sigma(x, r)$) the ball (sphere) of radius r centered at x , and by $D_p(x, r)$ the disk lying in plane P centered at $x \in P$ and of radius r .

Let $R \subset P$ be a region of P . The plane P containing R is called a *supporting plane* of R . We define:

$$R \oplus_p \varepsilon = \{x \in P : D_p(x, \varepsilon) \cap R \neq \emptyset\}$$

$$R \ominus_p \varepsilon = \{x \in P : D_p(x, \varepsilon) \subset R\}$$

$R \oplus_p \varepsilon$ is obtained by growing R by ε within its supporting plane P and $R \ominus_p \varepsilon$ is obtained by shrinking R by ε within its supporting plane P . When the supporting plane is unique or when it is clear from the context, we will simply note $R \oplus \varepsilon$ and $R \ominus \varepsilon$.

2.3 Polyhedral surfaces

We call *polyhedral surface* a finite collection of bounded polygons, any two of which are either disjoint or meet in a common edge or vertex. The polygons are called *facets*. Notice that we allow an arbitrary number of polygons to be glued along a common edge. In the mathematical literature, such an object is called a pure two-dimensional piece-wise linear complex. We prefer to use in this paper the term surface since surfaces are our primary concern.

In the rest of the paper, S denotes an arbitrary but fixed polyhedral surface. Three quantities C , A and L will express the complexity of the surface S : C denotes the number of facets of S , $A = a(S)$ its area, and L the sum of the lengths of the boundaries of the facets of S :

$$L = \sum_{F \subset S} l(\partial F).$$

Observe that, if an edge is incident to k facets, its length will be counted k times.

We consider two zones on the surface, the ε -singular zone that surrounds the edges of S and the ε -regular zone obtained by shrinking the facets.

Definition 1. Let $\varepsilon \geq 0$. The ε -regular zone of a facet $F \subset S$ consists of the points of F at distance greater than ε from the boundary of F . The ε -regular zone of S is the union of the ε -regular zones of its facets. We call ε -*singular zone* of F (resp. S) the set of points that do not belong to the ε -regular zone of F (resp. S).

Observe that the ε -regular zone of the facet F is $F \ominus \varepsilon$. The 0-singular zone of S consists exactly of the edges of S .

2.4 Sample

Any finite subset of points $E \subset S$ is called a *sample* of S . The points of E are called *sample points*. We impose two conditions on samples. First, the facets of the surface must be *uniformly* sampled. Second, the sample cannot be arbitrarily dense locally.

Definition 2. Let S be a polyhedral surface. $E \subset S$ is said to be a (ε, κ) -sample of S iff for every facet F of S and every point $x \in F$:

- the ball $B(x, \varepsilon)$ encloses at least one point of $E \cap F$,
- the ball $B(x, 2\varepsilon)$ encloses at most κ points of $E \cap F$.

The 2 factor in the second condition of the definition is not important and is just to make the constant in our bound simpler. Any other constant and, in particular 1, will lead to a linear bound.

In the rest of the paper, E denotes a (ε, κ) -sample of S and we provide asymptotic results when the sampling density increases, i.e. when ε tends to 0. As already mentioned, we consider κ and the surface S (and, in particular, the three quantities C , A and L) to be fixed and not to depend on ε .

Several related sampling conditions have been proposed. Amenta and Bern have introduced ε -samples [3] that fit locally the surface shape: the point density is high where the surface has high curvature or where the object or its complement is thin. However this definition is not appropriate for polyhedral surfaces since an ε -sample, as defined in [3], should have infinitely many points.

Erickson has introduced a notion of uniform sample that is related to ours but forbids two points to be too close [17]. Differently, our definition of a (ε, κ) -sample does not impose any lower bound on the minimal distance between two sample points.

In [5], Attali and Boissonnat use a slightly different definition of a (ε, κ) -sample. They assumed that for every point $x \in S$, the ball $B(x, \varepsilon)$ encloses at least one sample point and the ball $B(x, r)$ encloses $O(\frac{r^2}{\varepsilon^2})$ sample points. With this sampling condition, they proved that the complexity of the Delaunay triangulation is $O(n^{1.8})$ for general polyhedral surfaces and $O(n^{1.5})$ for convex polyhedral surfaces. Our definition of a (ε, κ) -sample is slightly more restrictive since the facets need to be sampled independently of one another, which leads to add a few more sample points near the edges. However, the two conditions are essentially the same and our linear bound holds also under the slightly more general sampling condition of [5].

Golin and Na [19, 20] assume that the sample points are chosen uniformly at random on the surface. The practical relevance of such a model is questionable since data are usually produced in a deterministic way.

3. PRELIMINARY RESULTS

S designates a polyhedral surface and $E \subset S$ a (ε, κ) -sample of S . Let $n(R) = |E \cap R|$ be the number of sample points in the region $R \subset S$. Let $n = |E|$ be the total number of sample points. We first establish two propositions relating $n(R)$ and n . We start with the following lemma:

LEMMA 1. *Let F be a facet of S . For any $R \subset F$, we have:*

$$\frac{a(R)}{4\pi\varepsilon^2} \leq n(R) \leq \frac{\kappa a(R \oplus \varepsilon)}{\pi\varepsilon^2}$$

PROOF. Let $\cup_{i=1}^{\lambda} D(m_i, \varepsilon)$ be a maximal set of λ non-intersecting disks lying inside $R \oplus \varepsilon$. Because the set of disks is maximal, no other disk can be added without intersecting one of the λ disks $D(m_i, \varepsilon)$. This implies that no point x of R is at distance greater than 2ε from a point m_i (see Figure 2). Therefore, $\cup_{i=1}^{\lambda} D(m_i, \varepsilon)$ is a packing of $R \oplus \varepsilon$ and $\cup_{i=1}^{\lambda} D(m_i, 2\varepsilon)$ is a covering of R . Consequently:

$$\frac{a(R)}{4} \leq \lambda \pi \varepsilon^2 \leq a(R \oplus \varepsilon)$$

The disks lie in $R \oplus \varepsilon$. Therefore, the centers of the disks lie in R . By assumption, the disk $D(m_i, \varepsilon)$ contains at least one sample point and the disk $D(m_i, 2\varepsilon)$ contains at most κ sample points. Hence :

$$\lambda \leq n(R) \leq \kappa \lambda$$

and

$$\frac{a(R)}{4\pi\varepsilon^2} \leq n(R) \leq \frac{\kappa a(R \oplus \varepsilon)}{\pi\varepsilon^2}$$

□

PROPOSITION 2. *Let F be a facet of S . For any $R \subset F$, we have:*

$$n(R) \leq 4\kappa \frac{a(R \oplus \varepsilon)}{A} n$$

PROOF. We first apply Lemma 1 to bound n from below. Summing over the facets of S , we get :

$$\frac{A}{4\pi\varepsilon^2} \leq n \tag{1}$$

We apply again Lemma 1 to bound $n(R)$ from above.

$$n(R) \leq \frac{\kappa a(R \oplus \varepsilon)}{\pi\varepsilon^2}$$

Eliminating ε from the two inequalities yields the result. □




Figure 2: A maximal set of non-intersecting disks contained in $R \oplus \varepsilon$ and the corresponding covering of R obtained by doubling the radii of the disks.

PROPOSITION 3. *Let F be a facet of S . Let $\Gamma \subset F$ be a curve contained in F . Let $k > 0$. We have:*

$$n(\Gamma \oplus k\varepsilon) \leq \frac{(2k+1)^2}{k} \kappa \frac{l(\Gamma)}{\varepsilon} \leq \frac{2(2k+1)^2}{k} \sqrt{\pi} \kappa \frac{l(\Gamma)}{\sqrt{A}} \sqrt{n}$$

PROOF. Arguing as in the proof of Lemma 1, we see that the region $\Gamma \oplus k\varepsilon$ can be covered by $\frac{l(\Gamma)}{k\varepsilon}$ disks of radius $2k\varepsilon$ centered on Γ and contained in the supporting plane of F .

Applying Lemma 1 to a disk R with radius $2k\varepsilon$, we get:

$$n(R) \leq \frac{\kappa \pi (2k\varepsilon + \varepsilon)^2}{\pi \varepsilon^2} = (2k+1)^2 \kappa$$

Therefore, we have :

$$n(\Gamma \oplus k\varepsilon) \leq (2k+1)^2 \kappa \frac{l(\Gamma)}{k\varepsilon}$$

From Equation 1, we get:

$$\frac{1}{\varepsilon} \leq \frac{2\sqrt{\pi}}{\sqrt{A}} \sqrt{n}$$

Combining the two inequalities leads to the result. □

LEMMA 4. Let x be a sample point in the ε -regular zone of S . Let P be the supporting plane of the facet through x . Any empty sphere passing through x intersects P in a circle whose radius is less than ε .




Figure 3: Assume Σ is an empty sphere passing through a point $x \in F \ominus \varepsilon$ and intersecting the supporting plane of F in a circle of radius greater than ε . Then, Σ contains an empty disk $D_p(y, \varepsilon)$ centered on F .

PROOF. The proof is by contradiction. Let P be the supporting plane of F . Consider an empty sphere Σ passing through x and intersecting P along a circle of radius greater than ε (see Figure 3). Let c be the center of this circle. Let y be the point on the segment $[xc]$ at distance ε from x . Because x belongs to the ε -regular zone of F , $y \in F$. The empty sphere Σ encloses the disk $D_p(y, \varepsilon)$. Therefore, $D_p(y, \varepsilon)$ is an empty disk of P , centered on F and of radius ε , which contradicts our assumption. \square

4. COUNTING DELAUNAY EDGES

Let S be a polyhedral surface and E be a (ε, κ) -sample of S . The Delaunay triangulation of E connects two points $p, q \in E$ iff there exists an empty sphere passing through p and q . The edge connecting p and q is called a *Delaunay edge*. We will also say that p and q are *Delaunay neighbours*.

The number of edges e_p and the number of tetrahedra t_p incident to a vertex p lying in the interior of the convex hull of E are related by Euler formula

$$t_p = 2e_p - 4$$

since the boundary of those tetrahedra is a simplicial polyhedron of genus 0. Using the same argument, if p lies on the boundary of the convex hull, we have:

$$t_p < 2e_p - 4$$

By summing over the n vertices, and observing that a tetrahedron has four vertices and an edge two, we get

$$t < e - n.$$

To bound the complexity of the Delaunay triangulation, it is therefore sufficient to count the Delaunay edges of E .

We distinguish three types of Delaunay edges : those with both endpoints in the ε -regular zone, those with both endpoints in the

ε -singular zone and those with an endpoint in the ε -regular zone and the other in the ε -singular zone. They are counted separately in the following subsections,

We denote by E_s the set of sample points in the ε -singular zone of S .

4.1 Delaunay edges with both endpoints in the ε -regular zone

In this section, we count the Delaunay edges joining two points in the ε -regular zone.

LEMMA 5. Let x be a sample point in the ε -regular zone and F the facet that contains x . x has at most κ Delaunay neighbours in F .

PROOF. By Lemma 4, any empty sphere passing through x intersects F in a circle whose radius is less than ε . Therefore, the Delaunay neighbours of x on F are at distance at most 2ε from x . By assumption, the disk centered at x with radius 2ε contains at most κ points of E . \square

LEMMA 6. Let x be a sample point in the ε -regular zone of a facet F . Let $F' \neq F$ be another facet of S . x has at most κ Delaunay neighbours in the ε -regular zone of facet F' .

PROOF. Refer to Figure 4. P and P' are the supporting planes of F and F' , y' is a Delaunay neighbour of x in the ε -regular zone of F' and Σ is an empty sphere passing through x and y' . Σ intersects the planes P and P' along two circles whose radii are respectively r and r' . By Lemma 4, $r \leq \varepsilon$ and $r' \leq \varepsilon$.




Figure 4: Any sphere passing through x and y' intersects one of the two planes P or P' in a circle whose diameter is at least $\|x' - y'\|$.

Let M be the bisector plane of P and P' . Let x' and y be the points symmetric to x and y' with respect to M . Consider the smallest sphere passing through x, x', y and y' . This sphere intersects P and P' in two disks of the same radius r_{\min} . We claim that $r_{\min} \leq \max(r, r')$. Indeed, let c be the circumcenter of Σ , and $H_{xy'}$ (resp. $H_{x'y}$) be the bisector plane of x and y' (resp. of x' and y). Observe that $v \in H_{xy'} \cap H_{x'y}$ and $c \in H_{xy'}$. If $c \in H_{xy'} \cap H_{x'y}$, $r_{\min} = r = r'$ and the claim is proved. Otherwise, c must belong to one of the two open halfspaces limited by $H_{x'y}$. If c belongs to the halfspace that contains x' , Σ encloses x' and therefore $r_{\min} \leq r'$ while in the second it encloses y and $r_{\min} \leq r$.

We therefore have :

$$\frac{\|x' - y'\|}{2} = r_{\min} \leq \max(r, r') \leq \varepsilon$$

and consequently:

$$\|x' - y'\| \leq 2\varepsilon.$$

The Delaunay neighbours of x in the ε -regular zone of F' lie in the disk $D_{p'}(x', 2\varepsilon)$. This disk contains at most κ points of E . \square

PROPOSITION 7. *There are at most $\frac{C}{2}\kappa n$ Delaunay edges with both endpoints in the ε -regular zone of \bar{S} .*

PROOF. The surface has C facets. Therefore, by Lemmas 5 and 6, a point x in the ε -regular zone of S has at most $C\kappa$ Delaunay neighbours. \square

4.2 Delaunay edges with both endpoints in the ε -singular zone

In this section, we count the Delaunay edges joining two points in the ε -singular zone.

PROPOSITION 8. *The number of Delaunay edges with both endpoints in the ε -singular zone is less than*

$$\frac{1}{2} 18^2 \pi \kappa^2 \frac{L^2}{A} n$$

PROOF. By Proposition 3, the number $|E_s|$ of sample points in the ε -singular zone is at most

$$18\sqrt{\pi} \kappa \frac{L}{\sqrt{A}} \sqrt{n}$$

Hence, the number of Delaunay edges in the ε -singular zone is at most $\frac{1}{2} |E_s| \times (|E_s| - 1) < \frac{1}{2} |E_s|^2$. \square




Figure 5: Example of a Delaunay triangulation of m points having a quadratic number of edges. Even if such a configuration can occur for a subset of the sample points, the number of Delaunay edges involved in this configuration is $O(n)$.

4.3 Delaunay edges joining the ε -regular and the ε -singular zones

In this section, we count the Delaunay edges with one endpoint in the ε -regular zone and the other in the ε -singular zone.

We first introduce a geometric construction of independent interest that will be useful.

Let P be a plane and E_s be a set of points. We assign to each point x of E_s the region $V(x) \subset P$ consisting of the points $p \in P$ for which the sphere tangent to P at p and passing through x encloses no point of E_s (see Figure 6). In other words, if $R(p, x)$ denotes the radius of the sphere tangent to P at p and passing through x , we have:

$$V(x) = \{p \in P : \forall y \in E_s, R(p, x) \leq R(p, y)\}.$$




Figure 6: The cell $V(x)$ is the set of contact points between a plane P and a sphere passing through x and tangent to P . The part of the paraboloid \mathcal{P}_x on the lower envelope of the paraboloids projects to the cell $V(x)$.

It is easy to see that the set of all $V(x)$, $x \in E_s$, is a subdivision of P we note \mathcal{V} (see Figure 9). Let \mathcal{P}_x be the paraboloid of revolution with focus x and director plane P . The paraboloid \mathcal{P}_x consists of the centers of the spheres passing through x and tangent to P . Assume that the points E_s are all located above plane P . If not, we replace x by the point symmetric to x with respect to P , which does not change \mathcal{V} . Let us consider the lower envelope of the collection of paraboloids $\{\mathcal{P}_x\}_{x \in E_s}$. Cell $V(x)$ is the projection of the portion of the lower envelope contributed by \mathcal{P}_x (see Figures 6 and 9).

Consider the bisector $V(x, y)$ of $x, y \in E_s$, i.e. the points $p \in P$ such that $R(p, x) = R(p, y)$. $V(x, y)$ is the projection on P of the intersection of the paraboloids \mathcal{P}_x and \mathcal{P}_y . As easy computations can show, the bisector $V(x, y)$ of x and y is a circle or a line (considered as a degenerated circle). Let $H(x, y)$

$$H(x, y) = \{p \in P : R(p, x) \leq R(p, y)\}.$$

Since $V(x, y)$ is a circle, $H(x, y)$ is either a disk, in which case we rename it $D(x, y)^+$, or the complementary set of a disk $D(x, y)^-$. We therefore have

$$V(x) = \bigcap_{y \in E_s, y \neq x} H(x, y) = (\cap D(x, y)^+) \setminus (\cup D(x, y)^-)$$

It follows that the edges $E(x, y)$ of $V(x)$ are circle arcs that we call *convex* or *concave* wrt x depending whether the disk $D(x, y)$ (whose boundary contains $E(x, y)$) is labelled $+$ or $-$ (see Figure 7). Observe that the convex edges of $V(x)$ are included in the boundary of the convex hull of $V(x)$.




Figure 7: The bold edges are the convex edges of the shaded cells. The edge $E(x, y)$ which is concave wrt x is convex wrt y . The convex edges of a cell lie on the boundary of its convex hull.

PROPOSITION 9. *The number of Delaunay edges with one endpoint in the ε -regular zone and the other in the ε -singular zone is at most :*

$$\left(1 + 450 \pi \kappa^2 \frac{L^2}{A}\right) n$$

PROOF. Let F be a facet of S and P the supporting plane of F . We bound the number of Delaunay edges with one endpoint in E_s and the other in $E \cap (F \ominus \varepsilon)$, i.e. the number of Delaunay edges joining the ε -singular zone and the ε -regular zone of F .

We denote by \mathcal{V}_F the restriction of the subdivision \mathcal{V} introduced above to F , and, for $x \in E_s$, we denote by $V(x)$ the cell of \mathcal{V}_F associated to x .

We first show that the Delaunay neighbours of x that belong to the ε -regular zone of F belong to $V(x) \oplus 2\varepsilon$. Consider a Delaunay edge (xf) with $x \in E_s$, $x \notin P$ and $f \in E_s \cap (F \ominus \varepsilon)$. Let Σ be an empty sphere passing through x and f , v its center (see Figure 8). By Lemma 4, Σ intersects P in a circle whose radius r is less than ε . For a point c on the segment $[vx]$, we note Σ_c the sphere centered at c and passing through x . Because Σ encloses Σ_c , Σ_c is an empty sphere. For $c = v$, Σ_c intersects P . For $c = x$, Σ_c does not intersect P . Consequently, there exists a position of c on $[vx]$ for which Σ_c is tangent to P . Let $p = \Sigma_c \cap P$ for such a point c . We have $p \in V(x)$ and $\|p - f\| \leq 2r \leq 2\varepsilon$. Hence, $f \in V(x) \oplus 2\varepsilon$. Now, let us consider a Delaunay edge (xf) with $x, f \in E_s \cap P$. Applying Lemma 4 leads to $f \in V(x) \oplus 2\varepsilon$.

Let N_F be the number of Delaunay edges between E_s and $F \ominus \varepsilon$. We have, using the fact that \mathcal{V}_F is a subdivision of F and Proposition 3 :

$$\begin{aligned} N_F &\leq \sum_{x \in E_s} n(V(x) \oplus 2\varepsilon) \\ &\leq n(F) + \sum_{x \in E_s} n(\partial V(x) \oplus 2\varepsilon) \\ &\leq n(F) + 25\sqrt{\pi} \kappa \frac{1}{\sqrt{A}} \sqrt{n} \sum_{x \in E_s} l(\partial V(x)) \end{aligned}$$

Let us bound $\sum_{x \in E_s} l(\partial V(x))$. Given a cell $V(x)$, we bound the length of its convex edges. By summing over all $x \in E_s$, all edges in \mathcal{V}_F will be taken into account.

The convex edges of x are contained in the boundary of the convex hull of $V(x)$. Since $V(x) \subset F$, the length of the boundary of




Figure 8: Every sphere Σ passing through x and $f \in P$ contains a sphere Σ_c passing through x and tangent to P .

the convex hull of $V(x)$ is at most the length of ∂F . Consequently:

$$\sum_{x \in E_s} l(\partial V(x)) \leq l(\partial F) \times |E_s|$$

Since, by Proposition 3, $|E_s| \leq 18\sqrt{\pi} \kappa \frac{L}{\sqrt{A}} \sqrt{n}$, we have:

$$N_F \leq n(F) + 450 \pi \kappa^2 \frac{l(\partial F) \times L}{A} n$$

By summing over all the facets, we conclude that the total number of Delaunay edges with one endpoint in the ε -regular zone and the other in ε -singular zone is at most :

$$\left(1 + 450 \pi \kappa^2 \frac{L^2}{A}\right) n$$

□

4.4 Main result

We sum up our results in the following theorem :

THEOREM 10. *Let S be a polyhedral surface and E a (ε, κ) -sample of S of size $|E| = n$. The number of edges in the Delaunay triangulation of E is at most :*

$$\left(1 + \frac{C \kappa}{2} + 612 \pi \kappa^2 \frac{L^2}{A}\right) n$$

It should be observed that the bound does not depend on the relative position of the facets (provided that their relative interiors do not intersect). Notice also that the bound is not meaningful when $A = 0$, which is the case of the quadratic example in Figure 5.

5. CONCLUSION

We have shown that, under a mild sampling condition, the Delaunay triangulation of points scattered over a fixed polyhedral surface or any fixed pure piece-wise linear complex has linear complexity. Our sampling condition does not involve any randomness (as in the work by Golin and Na [19]) and is less restrictive than Erickson's one [17].

Although the sampling condition has been expressed in a simple and intuitive way, the linear bound holds under a more general setting. Indeed, all we need for the proof is to subdivide the surface in two zones, an ε -regular zone where one can apply Lemma 4 and an ε -singular zone containing $O(\sqrt{n})$ points.

As mentioned in the introduction, Erickson has shown that the Delaunay triangulation of n points distributed on a cylinder may be quadratic. To understand where our analysis fails for such an example, one has to remember that our proof relies on Lemma 4 which states that empty balls intersect polyhedral surfaces in disks whose area is smaller than $\pi\epsilon^2$, which is not the case anymore in Erickson's example.

The main open question is of course to consider the case of smooth surfaces. The $O(n\sqrt{n})$ lower bound obtained by Erickson for cylinders show that a linear bound does not hold for arbitrary surfaces. We conjecture that, for generic surfaces, the complexity of the Delaunay triangulation is still linear. We say that a surface S is generic if 1. its maximal balls intersect S at a finite number of so-called contact points, and 2. the intersection of S with the union of the maximal balls with only one contact point form a set of curves of finite length on S . In particular, generic surfaces cannot contain spherical nor cylindrical pieces.

6. REFERENCES

- [1] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In *Proc. SIGGRAPH '98*, pages 415–412, July 1998.
- [2] N. Amenta, S. Choi, T.K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 213–222, 2000.
- [3] Nina Amenta and Marshall Bern. Surface reconstruction by voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.
- [4] Nina Amenta and Ravi Krishna Kolluri. Accurate and efficient unions of balls. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 119–128, 2000.
- [5] Dominique Attali and Jean-Daniel Boissonnat. Complexity of the delaunay triangulation of points on polyhedral surfaces. Technical Report RR-4232, INRIA, July 2001. <http://www-sop.inria.fr/rapports/sophia/RR-4232.html>.
- [6] F. Bernardini, C. L. Bajaj, J. Chen, and D. R. Schikore. Automatic reconstruction of 3D CAD models from digital scans. *Internat. J. Comput. Geom. Appl.*, 9:327–369, 1999.
- [7] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. on Visualization and Computer Graphics*, 5(4):349–359, 1999.
- [8] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 223–232, 2000.
- [9] Jean-Daniel Boissonnat and Frédéric Cazals. Natural neighbour coordinates of points on a surface. *Computational Geometry - Theory and Application*, 19(2-3):87–220, 2001.
- [10] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.
- [11] 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.
- [12] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.
- [13] S. Choi and N. Amenta. Delaunay triangulation programs on surface data. In *Proc. of 13th ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, CA, January 2002.
- [14] R. A. Dwyer. Higher-dimensional voronoi diagram in linear expected time. *Discrete Comput. Geom.*, 6:343–367, 1991.
- [15] R. A. Dwyer. The expected number of k -faces of a voronoi diagram. *Internat. J. Comput. Math.*, 26(5):13–21, 1993.
- [16] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.*, 21:87–115, 1999.
- [17] Jeff Erickson. Nice point sets can have nasty Delaunay triangulations. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 96–105, 2001.
- [18] Jeff Erickson. Dense point sets have sparse delaunay triangulations. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2002. To appear. <http://www.cs.ust.hk/tcsc/RR/>.
- [19] Mordecai J. Golin and Hyeon-Suk Na. On the average complexity of 3d-voronoi diagrams of random points on convex polytopes. In *Proc. 12th Canad. Conf. Comput. Geom.*, 2000. <http://www.cs.ust.hk/tcsc/RR/>.
- [20] Hyeon-Suk Na Mordecai J. Golin. The probabilistic complexity of the voronoi diagram of points on a polyhedron. In *Proc. ACM Sym. on Computational Geometry*, 2002. To appear.
- [21] D.J. Sheehy, C.G. Armstrong, and D.J. Robinson. Computing the medial surface of a solid from a domain delaunay triangulation. In *Proc. 3rd Symposium on Solid Modeling and Applications*, pages 201–212, May 1995.
- [22] Andrew Thall, Stephen Pizer, and Tom Fletcher. Deformable solid modeling using sampled medial surfaces: A multiscale approach. Technical Report TR00-005, January 2000. <http://citeseer.nj.nec.com/thall00deformable.html>.

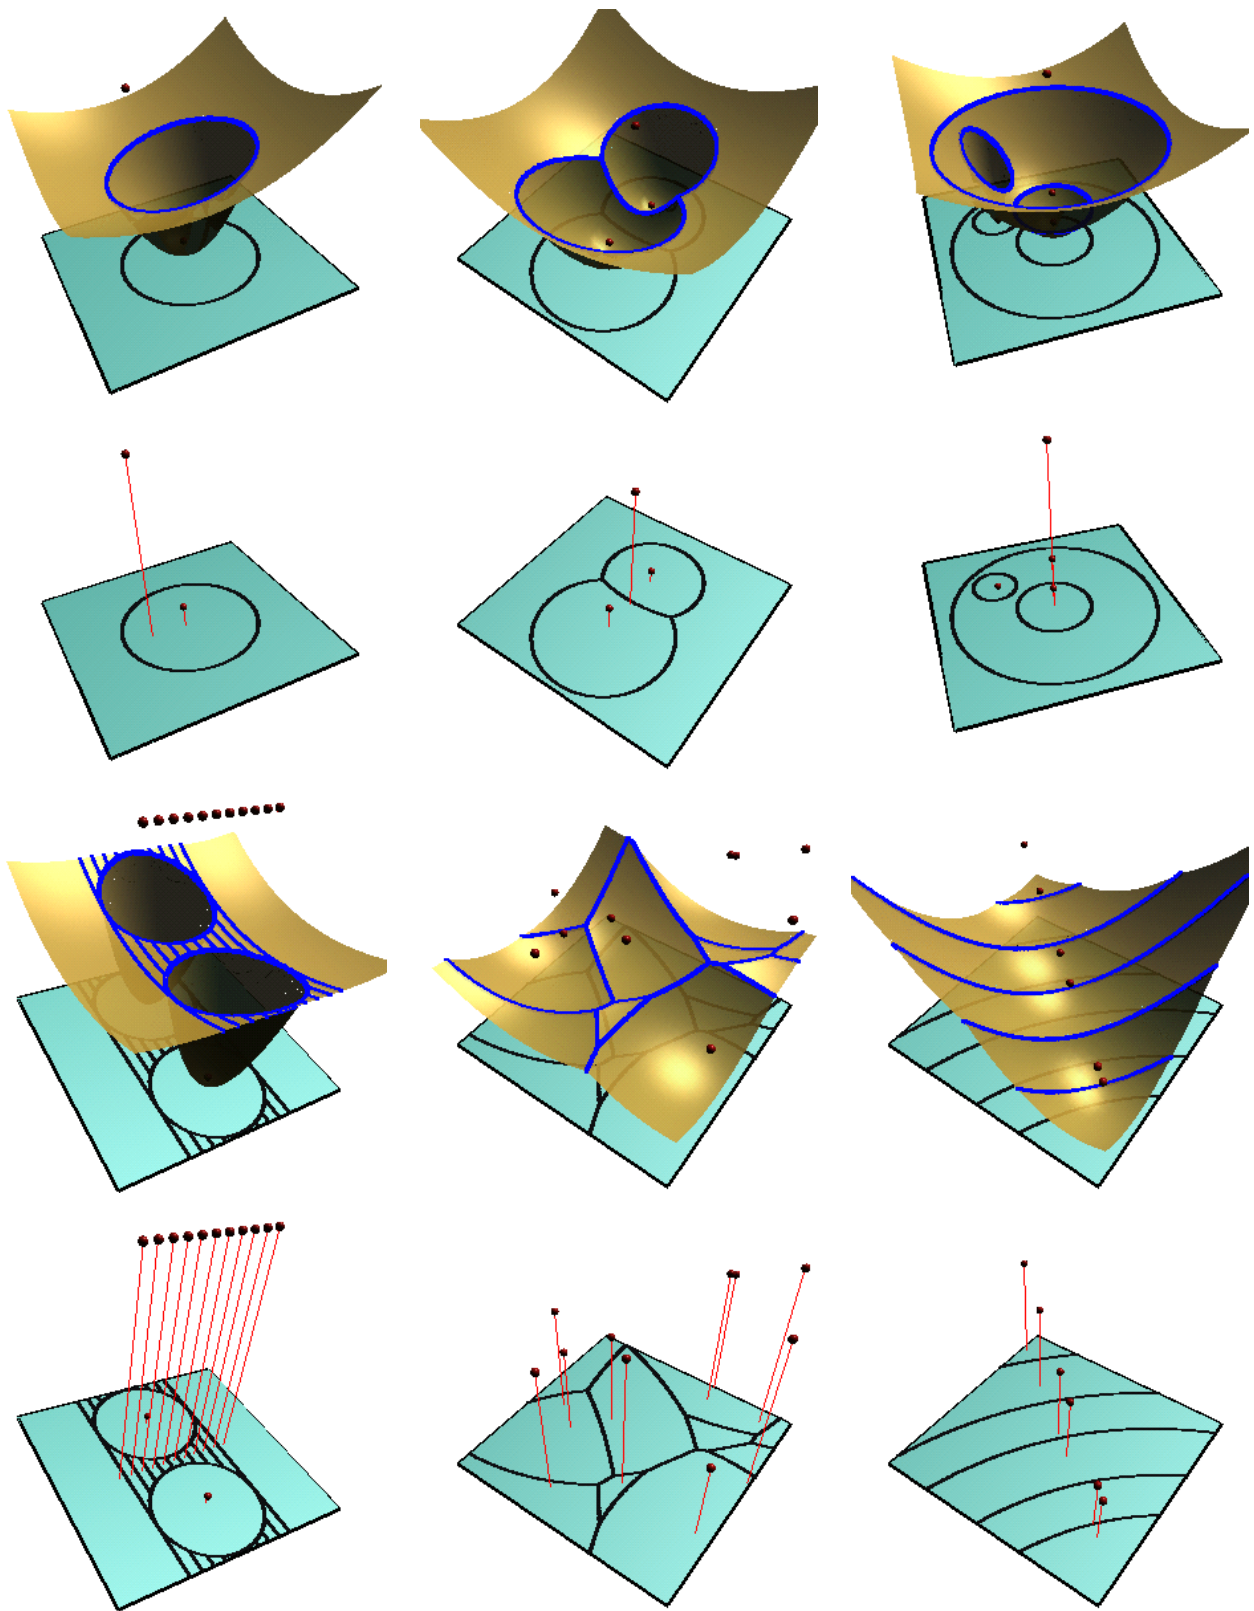


Figure 9: Decomposition of a facet F into cells for different set of points E_s . The lower envelope of the paraboloid $\{\mathcal{P}_x\}_{x \in E_s}$ has been represented. The red spheres represent the points of E_s and the red lines materialize the projection of the points of E_s on the plane P . The bisector of two points is a circle. The projection of x on P do not belong necessary to its cell. The decomposition of F can have a quadratic number of edges.