Discrete Comput Geom 31:369–384 (2004)
DOI: 10.1007/s00454-003-2870-4

Discrete & Computational
Geometry

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

Dominique Attali[1] and Jean-Daniel Boissonnat[2]

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

[2]INRIA, Unité de Recherche Sophia Antipolis,
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 set of well-sampled facets of $\mathbb{R}^3$ (e.g. the planar polygons in a polyhedron). Our bound is deterministic and the constants are explicitly given.

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 modeling, surface mesh generation [13], deformable surface modeling [22], [17], medial axis approximation [4], [9], [23] and surface reconstruction [1], [2], [10], [3], [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.

* 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).

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 $\Omega(n^{\lceil d/2 \rceil})$ [11]. 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. [10] and [14]), 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 [12], 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 [15], [16] who proved that if the points are uniformly distributed in a ball, then the *expected* complexity of the Delaunay triangulation is only linear. Recently, Erickson [18], [19] 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 [20]. They proved that the expected complexity of three-dimensional 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 [21]. 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 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 finite set of interior-disjoint planar regions whose total area is positive and whose total perimeter is finite. This includes the case of polyhedral surfaces. Under a mild uniform sampling condition (depending on a parameter κ), we show that the complexity of the Delaunay triangulation of the points is linear when κ is a constant. Our bound is deterministic. The constants are explicitly given and depend on κ and on the number of planar regions C_S, the total area A_S and the total perimeter L_S of the regions. More precisely, our main result states that the number of Delaunay edges is at most

$$\left(1 + \frac{C_S \kappa}{2} + 5300\pi\kappa^2 \frac{L_S^2}{A_S}\right) n.$$

Our bound holds for any $n > 0$.

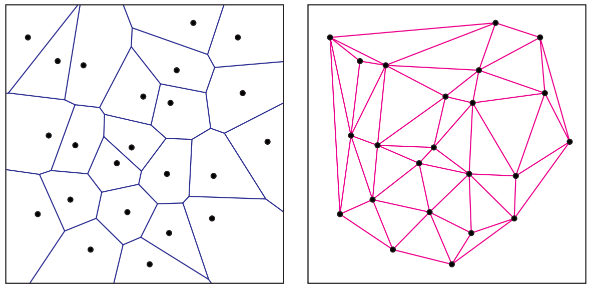

Fig. 1. Voronoi diagram of a set of points (left) and its dual Delaunay triangulation (right).

2. Definitions and Notations

2.1. *Voronoi Diagrams and Delaunay Triangulations*

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

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

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 P, denoted $\mathrm{Vor}(P)$. The Delaunay triangulation of P, denoted $\mathrm{Del}(P)$, is the dual complex of $\mathrm{Vor}(P)$ (see Fig. 1). If there is no sphere passing through $d+2$ points of P, $\mathrm{Del}(P)$ is a simplicial complex that can be obtained from $\mathrm{Vor}(P)$ as follows. If P' is a subset of points of P whose Voronoi cells have a nonempty intersection, the convex hull $\mathrm{conv}(P')$ is a Delaunay face and all Delaunay faces are obtained this way. It is well known that the balls circumscribing the d-simplices in $\mathrm{Del}(P)$ cannot contain a point of P in their interior. The *complexity* of $\mathrm{Del}(P)$ 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* if and only if its interior contains no point of P. We also say that a sphere is *empty* if the associated ball is empty.

2.2. *Notations*

For a curve Γ, we denote by $\mathrm{length}(\Gamma)$ its length. For a portion of a surface R, we denote by $\mathrm{area}(R)$ 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 $D_H(x, r)$ the disk lying in plane H centered at $x \in H$ and of radius r.

Let H be a plane and let $R \subset H$ be a region of H. The plane H containing R is called a *supporting plane* of R. We define

$$\begin{aligned} R \oplus_H \varepsilon &= \{x \in H : D_H(x, \varepsilon) \cap R \neq \emptyset\}, \\ R \ominus_H \varepsilon &= \{x \in H : D_H(x, \varepsilon) \subset R\}. \end{aligned}$$

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

2.3. *Polyhedral Surfaces*

In this paper we use the term *polyhedral surface* to denote a fixed finite set of interior-disjoint planar regions whose total area is positive and whose total perimeter is finite. Accordingly, the planar regions are called *facets* and the intersection between two facets is called an *edge*. This abuse of terminology is mainly for simplicity and to refer to what is probably the most important case in applications. It should be kept in mind however that our results hold for objects that are more general than usual polyhedral surfaces. In particular, we do not require our polyhedral surfaces to be connected or to be manifolds, we allow an arbitrary number of facets to be glued to a common edge, etc.

In the rest of the paper, S denotes an arbitrary but fixed polyhedral surface. Three quantities C_S, A_S and L_S express the complexity of the surface S: C_S denotes the number of facets of S, $A_S = \text{area}(S)$ denotes its area and L_S is the sum of the lengths of the boundaries of the facets of S:

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

Observe that if an edge is incident to k facets, then 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 $F \subset S$ is $F \ominus \varepsilon$. The ε-regular zone of S is the union of the ε-regular zones of its facets. The *ε-singular zone* of F (*resp*. S) is the set of points that do not belong to the ε-regular zone of F (*resp*. S).

Observe that the 0-singular zone of S consists exactly of the edges of S.

2.4. *Sample*

Any finite subset of points $P \subset S$ is called a *sample* of S. The points of P 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. $P \subset S$ is said to be an (ε, κ)-sample of S if and only if for every facet F of S and every point $x \in F$:

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

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, in particular 1, will lead to a linear bound.

In the rest of the paper, P denotes an (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_S, A_S and L_S) to be fixed and not to depend on ε.

Several related sampling conditions have been proposed.

Amenta and Bern have introduced ε-samples [1] that fit the surface shape locally: 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 [1], should have infinitely many points.

Erickson has introduced a notion of a uniform sample that is related to ours but forbids two points to be too close [18]. Differently, our definition of an (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 an (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 $O(r^2/\varepsilon^2)$ 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. In this paper our definition of an (ε, κ)-sample is slightly more restrictive since the facets need to be sampled independently of one another, which leads to adding 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 [20], [21] 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 $P \subset S$ designates an (ε, κ)-sample of S. We denote by $\sharp(A)$ the number of elements of A. Let $n(R) = \sharp(P \cap R)$ be the number of sample points in the region $R \subset S$. Let $n = \sharp(P)$ 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.

$$\frac{A_S}{4\pi\varepsilon^2} \leq n.$$

Proof. Let F be a facet of S. Let $\{D(x_i, \varepsilon)\}_{i\in\{1,\dots,\lambda\}}$ be a maximal set of λ nonintersecting disks lying inside $F \oplus \varepsilon$. Because the set of disks is maximal, no other disk can be added without intersecting $\bigcup_{i=1}^{\lambda} D(x_i, \varepsilon)$. This implies that no point m of F is at a distance greater than 2ε from a point x_i (see Fig. 2). Therefore, $\{D(x_i, 2\varepsilon)\}_{i\in\{1,\dots,\lambda\}}$ is a covering of F. We have $\text{area}(F)/4\pi\varepsilon^2 \leq \lambda$. Because of our sampling condition, every disk $D(x_i, \varepsilon)$ contains at least one sample point. Therefore, $\lambda \leq n(F)$ and

$$\frac{\text{area}(F)}{4\pi\varepsilon^2} \leq \lambda \leq n(F).$$

By summing over the facets of S, we get the result. $\square$

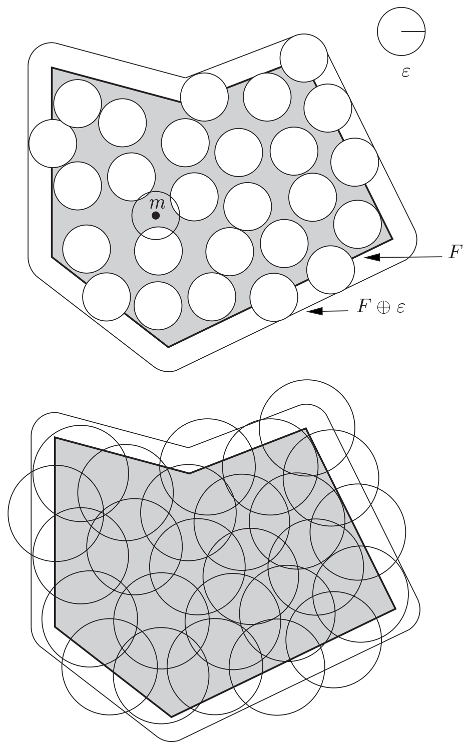

Fig. 2. A maximal set of nonintersecting disks contained in $F \oplus \varepsilon$ and the corresponding covering of F obtained by doubling the radii of the disks.

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

$$n(R) \leq \frac{4\kappa \operatorname{area}(R \oplus \varepsilon/2)}{\pi \varepsilon^2}.$$

Proof. Let $\{D(x_i, \varepsilon/2)\}_{i \in \{1,\ldots,\lambda\}}$ be a maximal set of λ nonintersecting disks lying inside $R \oplus \varepsilon/2$. Because the set of disks is maximal, no other disk can be added without intersecting $\bigcup_{i=1}^{\lambda} D(x_i, \varepsilon/2)$. This implies that no point m of R is at a distance greater than ε from a point x_i. Therefore, $\{D(x_i, \varepsilon)\}_{i \in \{1,\ldots,\lambda\}}$ is a covering of R. We have

$$n(R) \leq \kappa\lambda \leq \kappa \times \frac{\operatorname{area}(R \oplus \varepsilon/2)}{\pi \varepsilon^2/4}. \qquad \square$$

Proposition 3. *Let F be a facet of S. For any $R \subset F$, we have*

$$n(R) \leq 16\kappa \, \frac{\operatorname{area}(R \oplus \varepsilon/2)}{A_S} \, n.$$

Proof. By Lemma 1, we have

$$\frac{A_S}{4\pi\varepsilon^2} \le n.$$

We apply Lemma 2 to bound $n(R)$ from above:

$$n(R) \le \frac{4\kappa\,\mathrm{area}(R \oplus \varepsilon/2)}{\pi\varepsilon^2}.$$

Eliminating ε from the two inequalities yields the result. □

Proposition 4. *Let F be a facet of S. Let $\Gamma \subset F$ be a curve contained in F. Let $a > 0$. We have*

$$n(\Gamma \oplus a\varepsilon) \le \frac{(4a+1)^2}{a}\,\kappa\,\frac{\mathrm{length}(\Gamma)}{\varepsilon} \le \frac{2(4a+1)^2}{a}\sqrt{\pi}\,\kappa\,\frac{\mathrm{length}(\Gamma)}{\sqrt{A_S}}\sqrt{n}.$$

Proof. Arguing as in the proof of Lemma 2, we see that the region $\Gamma \oplus a\varepsilon$ can be covered by $\mathrm{length}(\Gamma)/a\varepsilon$ disks of radius $2a\varepsilon$ centered on Γ and contained in the supporting plane of F.

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

$$n(R) \le \frac{4\kappa\pi(2a\varepsilon + \varepsilon/2)^2}{\pi\varepsilon^2} = \kappa(4a+1)^2.$$

Therefore, we have

$$n(\Gamma \oplus a\varepsilon) \le \kappa\,\frac{(4a+1)^2}{a}\,\frac{\mathrm{length}(\Gamma)}{\varepsilon}.$$

From Lemma 1, we get

$$\frac{1}{\varepsilon} \le \frac{2\sqrt{\pi}}{\sqrt{A_S}}\sqrt{n}.$$

Combining the two inequalities leads to the result. □

Lemma 5. *Let x be a sample point in the ε-regular zone of S. Let H be the supporting plane of the facet through x. Any empty sphere passing through x intersects H in a circle whose radius is less than ε.*

Proof. The proof is by contradiction. Let H be the supporting plane of F. Consider an empty sphere Σ passing through x and intersecting H along a circle of radius greater than ε (see Fig. 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 $D_H(y, \varepsilon)$. Therefore, $D_H(y, \varepsilon)$ is an empty disk of H, centered on F and of radius ε, which contradicts our sampling condition. □

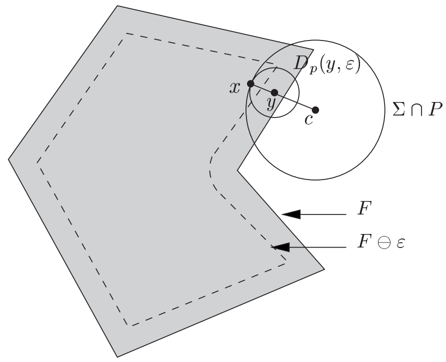

Fig. 3. Assume Σ is an empty sphere passing through a point $x \in F \ominus \varepsilon$ and intersecting the supporting plane of F in a circle of radius greater than ε. Then Σ contains an empty disk $D_H(y, \varepsilon)$ centered on F.

4. Counting Delaunay Edges

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

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 P are related by the Euler formula

$$t_p = 2e_p - 4,$$

since the boundary of the union 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 P.

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 P_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 6. *Let x be a sample point in the ε-regular zone and let F be the facet that contains x. x has at most κ Delaunay neighbors in F.*

Proof. By Lemma 5, any empty sphere passing through x intersects F in a circle whose radius is less than ε. Therefore, the Delaunay neighbors of x on F are at a distance at most 2ε from x. By assumption, the disk centered at x with radius 2ε contains at most κ points of P. □

Lemma 7. *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 neighbors in the ε-regular zone of facet F'.*

Proof. Refer to Fig. 4. H and H' are the supporting planes of F and F', y is a Delaunay neighbor of x in the ε-regular zone of F' and Σ is an empty sphere passing through x and y. Let B be the closed ball whose boundary is Σ. B intersects the planes H and H' along two disks whose radii are respectively r and r'. By Lemma 5, $r \leq \varepsilon$ and $r' \leq \varepsilon$.

Let M be the bisector plane of H and H'. Let x' and y' be the points symmetric to x and y with respect to M. Consider the sphere Σ_0 centered on M and passing through the four points x, x', y and y'. Let B_0 be the closed ball whose boundary is Σ_0. B_0 intersects H and H' along two disks D_0 and D_0' of the same radius r_0. We claim that $r_0 \leq \max(r, r')$. Indeed, let v_0 be the center of Σ_0 and let v be the center of Σ. Let M_{xy} (resp. $M_{x'y'}$) be the bisector plane of x and y (resp. of x' and y'). Observe that $v_0 \in M_{xy} \cap M_{x'y'}$ and $v \in M_{xy}$. If $v \in M_{xy} \cap M_{x'y'}$, $r_0 = r = r'$ and the claim is proved. Otherwise, v must belong to one of the two open halfspaces limited by $M_{x'y'}$. If v belongs to the halfspace that contains x', B encloses D_0' and therefore $r_0 \leq r'$ while in the second it encloses D_0 and $r_0 \leq r$.

We therefore have

$$\frac{\|x' - y\|}{2} = r_0 \leq \max(r, r') \leq \varepsilon$$

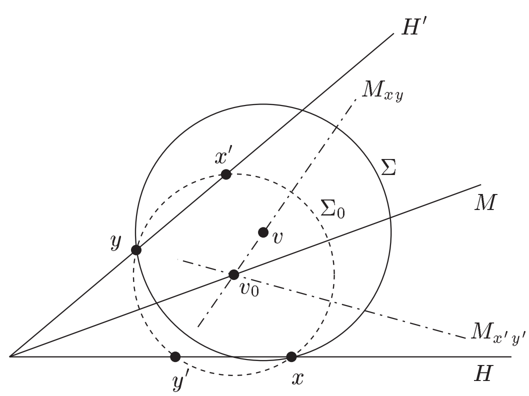

Fig. 4. Any sphere passing through x and y intersects one of the two planes H or H' in a circle whose diameter is at least $\|x' - y\|$.

and consequently

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

The Delaunay neighbors of x in the ε-regular zone of F' lie in the disk $D_{H'}(x', 2\varepsilon)$. This disk contains at most κ points of P. □

Proposition 8. *There are at most $\frac{1}{2}C_S\kappa n$ Delaunay edges with both endpoints in the ε-regular zone of S.*

Proof. The surface S has C_S facets. Therefore, by Lemmas 6 and 7, a point x in the ε-regular zone of S has at most $C_S\kappa$ Delaunay neighbors. □

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 (see Fig. 5).

Proposition 9. *The number of Delaunay edges with both endpoints in the ε-singular zone is less than*

$$\tfrac{1}{2}50^2\pi\kappa^2\frac{L_S^2}{A_S}n.$$

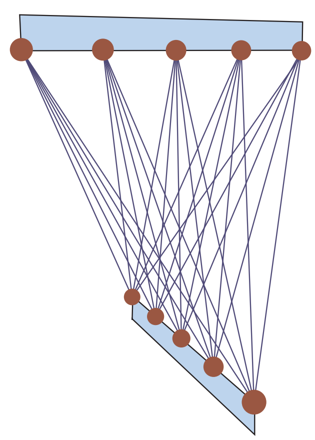

Fig. 5. Example of a Delaunay triangulation of m points having a quadratic number of edges. Even if such a configuration can occur for a subset of the sample points, the number m of sample points involved in this configuration is $O(\sqrt{n})$. Therefore, the number of Delaunay edges involved in this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

Proof. By Proposition 4, the number $\sharp(P_s)$ of sample points in the ε-singular zone is at most

$$50\sqrt{\pi}\,\kappa\,\frac{L_S}{\sqrt{A_S}}\,\sqrt{n}.$$

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

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 H be a plane in $\mathbb{R}^3$ and let $X \subseteq \mathbb{R}^3$ be a finite set of points. We assign to each point x of X the region $V(x) \subset H$ consisting of the points $h \in H$ for which the sphere tangent to H at h and passing through x encloses no point of X (see Fig. 6). In other words, if $R(h, x)$ denotes the radius of the sphere tangent to H at h and passing through x, we have

$$V(x) = \{h \in H : \forall y \in X,\ R(h, x) \le R(h, y)\}.$$

It is easy to see that the set of all $V(x)$, $x \in X$, is a subdivision of H which we denote $\mathcal{V}_H(X)$ (see Fig. 7). The diagram $\mathcal{V}_H(X)$ is a multiplicatively weighted power Voronoi diagram. Let $\mathcal{P}_x$ be the paraboloid of revolution with focus x and director plane H. The paraboloid $\mathcal{P}_x$ consists of the centers of the spheres passing through x and tangent to H. Assume that the points X are all located above plane H. If not, we replace x by the point symmetric to x with respect to H, which does not change $\mathcal{V}_H(X)$. We consider the lower envelope of the collection of paraboloids $\{\mathcal{P}_x\}_{x\in X}$. Cell $V(x)$ is the projection of the portion of the lower envelope contributed by $\mathcal{P}_x$ (see Figs. 6 and 7).

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

$$V(x, y) = \{h \in H : R(h, x) \le R(h, y)\}.$$

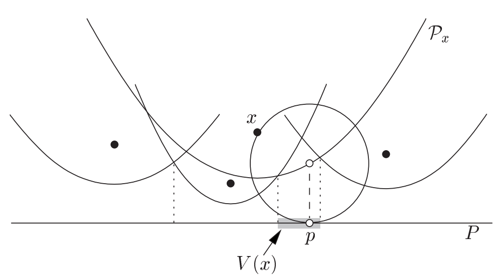

Fig. 6. The cell $V(x)$ is the set of contact points between a plane H and a sphere passing through x and tangent to H. The part of the paraboloid $\mathcal{P}_x$ on the lower envelope of the paraboloids projects to the cell $V(x)$.

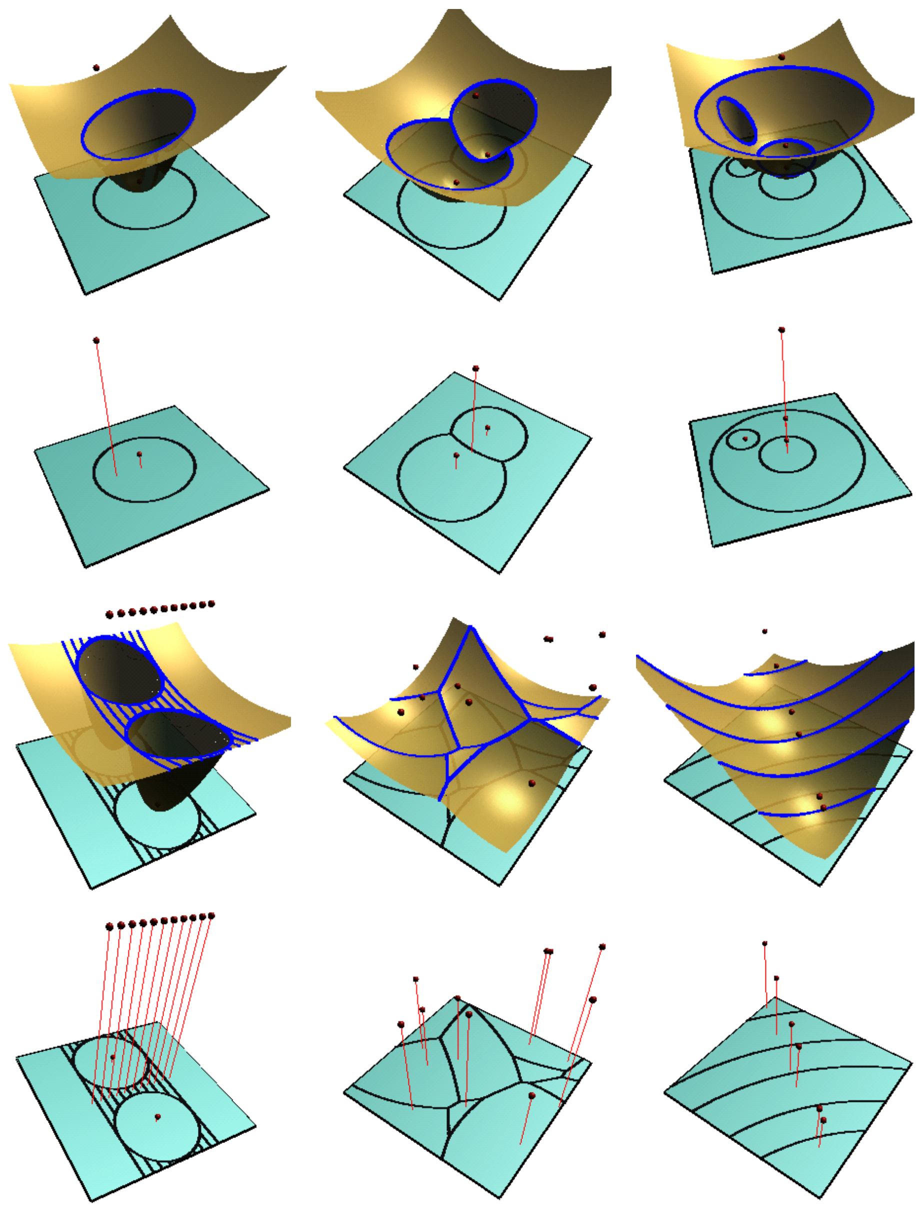

Fig. 7. Decomposition of a facet F into cells for different sets of points P_s. The lower envelope of the paraboloid $\{\mathcal{P}_x\}_{x\in P_s}$ has been represented. The red spheres represent the points of P_s and the red lines materialize the projection of the points of P_s on the plane H. The bisector of two points is a circle. The projection of x on H do not belong necessarily to its cell. The decomposition of F can have a quadratic number of edges.

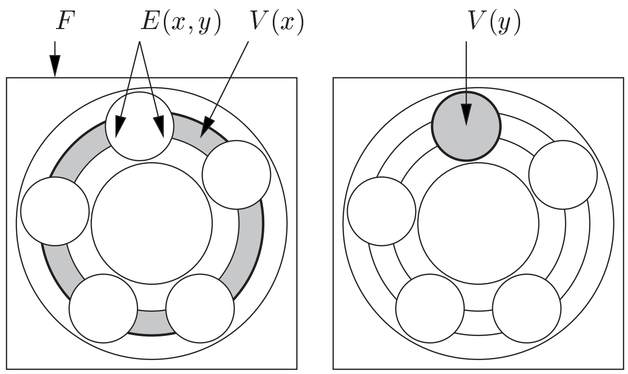

Fig. 8. The bold edges are the convex edges of the shaded cells. The edge $E(x, y)$, which is concave with respect to x, is convex with respect to y. The convex edges of a cell lie on the boundary of its convex hull.

Since $M(x, y)$ is a circle, $V(x, y)$ 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

$$V(x) = \bigcap_{y \in X, y \neq x} V(x, y) = \left(\bigcap D(x, y)^+\right) \setminus \left(\bigcup D(x, y)^-\right).$$

It follows that the edges $E(x, y)$ of $V(x)$ are circle arcs that we call *convex* or *concave* with respect to x depending on whether the disk $D(x, y)$ (whose boundary contains $E(x, y)$) is labeled $+$ or $-$ (see Fig. 8). Observe that the convex edges of $V(x)$ are included in the boundary of the convex hull of $V(x)$.

Proposition 10. *The number of Delaunay edges with one endpoint in the ε-regular zone and the other in the ε-singular zone is at most*

$$\left(1 + 4050\pi\kappa^2 \frac{L_S^2}{A_S}\right) n.$$

Proof. Let F be a facet of S and let H be the supporting plane of F. We bound the number of Delaunay edges with one endpoint in P_s and the other in $P \cap (F \ominus \varepsilon)$, 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 $\mathcal{V}_H(P_s)$ introduced above to F, and, for $x \in P_s$, we denote by $V(x)$ the cell of $\mathcal{V}_F$ associated to x.

We first show that the Delaunay neighbors of x that belong to the ε-regular zone of F belong to $V(x) \oplus 2\varepsilon$. Consider a Delaunay edge (xf) with $x \in P_s$, $x \notin H$ and $f \in P_s \cap (F \ominus \varepsilon)$. Let Σ be an empty sphere passing through x and f, and let v be its center (see Fig. 9). By Lemma 5, Σ intersects H in a circle whose radius r is less than ε. For a point c on the segment $[vx]$, we denote by 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 H. For $c = x$, Σ_c does not intersect H. Consequently, there exists a position of c on $[vx]$ for which Σ_c is tangent to H. Let $p = \Sigma_c \cap H$ 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, we consider a Delaunay edge (xf) with $x, f \in P_s \cap H$. Applying Lemma 5 leads to $f \in V(x) \oplus 2\varepsilon$.

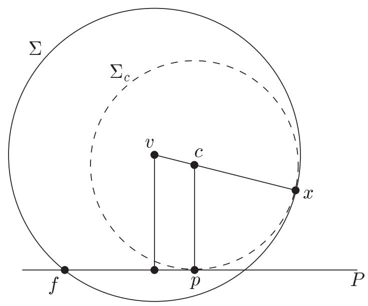

Fig. 9. Every sphere Σ passing through x and $f \in H$ contains a sphere Σ_c passing through x and tangent to H.

Let N_F be the number of Delaunay edges between P_s and $F \ominus \varepsilon$. We have, using the fact that $\mathcal{V}_F$ is a subdivision of F and Proposition 4,

$$\begin{aligned} N_F &\leq \sum_{x \in P_s} n(V(x) \oplus 2\varepsilon) \\ &\leq n(F) + \sum_{x \in P_s} n(\partial V(x) \oplus 2\varepsilon) \\ &\leq n(F) + 81\sqrt{\pi}\,\kappa\,\frac{1}{\sqrt{A_S}}\,\sqrt{n} \sum_{x \in P_s} \text{length}(\partial V(x)). \end{aligned}$$

We bound $\sum_{x \in P_s} \text{length}(\partial V(x))$. Given a cell $V(x)$, we bound the length of its convex edges. By summing over all $x \in P_s$, 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 the convex hull of $V(x)$ is at most the length of ∂F. Consequently,

$$\sum_{x \in P_s} \text{length}(\partial V(x)) \leq \text{length}(\partial F) \times \sharp(P_s).$$

Since, by Proposition 4, $\sharp(P_s) \leq 50\sqrt{\pi}\kappa(L_S/\sqrt{A_S})\sqrt{n}$, we have

$$N_F \leq n(F) + 4050\pi\kappa^2\,\frac{\text{length}(\partial F) \times L_S}{A_S}\,n.$$

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 the ε-singular zone is at most

$$\left(1 + 4050\pi\kappa^2\frac{L_S^2}{A_S}\right)n. \qquad \square$$

4.4. *Main Result*

We sum up our results in the following theorem:

Theorem 11. *Let S be a polyhedral surface and let P be a (ε, κ)-sample of S of size $\sharp(P) = n$. The number of edges in the Delaunay triangulation of P is at most*

$$\left(1 + \frac{C_S \kappa}{2} + 5300\pi\kappa^2 \frac{L_S^2}{A_S}\right) n.$$

Notice that our bound holds for any $n > 0$. It should be observed also that the bound does not depend on the relative position of the facets (provided that their relative interiors do not intersect). In particular, it does not depend on the dihedral angles between the facets. Notice also that the bound is not meaningful when $A_S = 0$, which is the case of the quadratic example in Fig. 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 piecewise linear complex has linear complexity. Therefore, we (partially) answered an old question of Boissonnat [8]. Our sampling condition does not involve any randomness (as in the work by Golin and Na [20]) and is less restrictive than Erickson's one [18].

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 into two zones, an ε-regular zone where one can apply Lemma 5 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 $\Omega(n\sqrt{n})$. To understand where our analysis fails for such an example, one has to remember that our proof relies on Lemma 5 which states that empty balls intersect polyhedral surfaces in disks whose area is smaller than $\pi\varepsilon^2$, which is not the case anymore in Erickson's example.

In our result the dependence on κ is quadratic. We left as an open question to see if it can be improved to linear. Another open question is of course to consider the case of smooth surfaces. The $\Omega(n\sqrt{n})$ 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.

References

1. N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.
2. N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In *Proc. SIGGRAPH '98*, pages 415–412, July 1998.

3. 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.
4. N. Amenta and R. K. Kolluri. Accurate and efficient unions of balls. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 119–128, 2000.
5. D. Attali and J.-D. Boissonnat. Complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete Comput. Geom.*, 30: 437–452, 2003.
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. Visual. Comput. Graphics*, 5(4):349–359, 1999.
8. J.-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graphics*, 3(4):266–286, October 1984.
9. J.-D. Boissonnat and F. Cazals. Natural neighbour coordinates of points on a surface. *Comput. Geom. Theory Appl.*, 19(2–3):155–173, July 2001.
10. J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comput. Geom. Theory Appl.*, 22(1):185–203, 2002.
11. J.-D. Boissonnat and M. Yvinec. *Algorithmic Geometry*. Cambridge University Press, Cambridge, 1998. Translated by Hervé Brönnimann.
12. 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.
13. L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.
14. S. Choi and N. Amenta. Delaunay triangulation programs on surface data. In *Proc. 13th ACM–SIAM Sympos. Discrete Algorithms*, pages 135–136, 2002.
15. R. A. Dwyer. Higher-dimensional Voronoi diagram in linear expected time. *Discrete Comput. Geom.*, 6:343–367, 1991.
16. R. A. Dwyer. The expected number of k-faces of a Voronoi diagram. *Internat. J. Comput. Math.*, 26(5):13–21, 1993.
17. H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.*, 21:87–115, 1999.
18. J. Erickson. Nice point sets can have nasty Delaunay triangulations. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 96–105, 2001.
19. J. Erickson. Dense point sets have sparse Delaunay triangulations. In *Proc. 13th ACM–SIAM Sympos. Discrete Algorithms*, pages 125–134, 2002.
20. M. J. Golin and H.-S. Na. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. In *Proc. 12th Canad. Conf. Comput. Geom.*, pages 127–135, 2000.
21. M. J. Golin and H.-S. Na. The probabilistic complexity of the Voronoi diagram of points on a polyhedron. In *Proc. ACM Sympos. Comput. Geom.*, pages 209–216, Barcelona, June 2002.
22. S. Joshi, S. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, and J. S. Marron. Multiscal deformable model segmentation and statistical shape using medial descriptions. *IEEE Trans. Medical Imaging*, 21(5):538–550, May 2002.
23. D. J. Sheehy, C. G. Armstrond, and D. J. Robinson. Computing the medial surface of a solid from a domain Delaunay triangulation. In *Proc. 3rd Sympos. Solid Modeling Appl.*, pages 201–212, May 1995.

Received April 30, 2002, *and in revised form May* 5, 2003. *Online publication February* 5, 2004.