Eurographics Symposium on Geometry Processing (2003)
L. Kobbelt, P. Schröder, H. Hoppe (Editors)

Provably Good Surface Sampling and Approximation

J-D Boissonnat[1] and S. Oudot[1]

[1] INRIA, 2004 route des lucioles, F-06902 Sophia-Antipolis. {Jean-Daniel.Boissonnat, Steve.Oudot}@sophia.inria.fr

Abstract

We present an algorithm for meshing surfaces that is a simple adaptation of a greedy "farthest point" technique proposed by Chew. Given a surface S, it progressively adds points on S and updates the 3-dimensional Delaunay triangulation of the points. The method is very simple and works in $3d$-space without requiring to parameterize the surface. Taking advantage of recent results on the restricted Delaunay triangulation, we prove that the algorithm can generate good samples on S as well as triangulated surfaces that approximate S. More precisely, we show that the restricted Delaunay triangulation $Del_{|S}$ of the points has the same topology type as S, that the Hausdorff distance between $Del_{|S}$ and S can be made arbitrarily small, and that we can bound the aspect ratio of the facets of $Del_{|S}$. The algorithm has been implemented and we report on experimental results that provide evidence that it is very effective in practice. We present results on implicit surfaces, on CSG models and on polyhedra. Although most of our theoretical results are given for smooth closed surfaces, the method is quite robust in handling smooth surfaces with boundaries, and even non-smooth surfaces.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling — Boundary representations.

1. Introduction

A lot of applications dealing with surfaces require the use of discretized geometric models for fast and efficient computation. For instance, in computer graphics, most of the rendering techniques work on polyhedral approximations of the objects rather than on the objects themselves. In the same way, numerical simulations based on finite elements rely on discrete descriptions. Therefore it is an important issue to produce meshes to approximate geometric models. Here we deal exclusively with simplicial meshes and consider the following problem:

Given a 2-manifold S embedded in $\mathbb{R}^3$, a metric L and a constant $\varepsilon > 0$, build a triangulated 2-manifold P with a minimum number of vertices, such that S and P are homeomorphic and $L(S,P) \leq \varepsilon$.

A discrete version of this problem has been shown to be *NP*-hard[1]. To our knowledge, it is still an open question to approximate the solution to this problem for general surfaces in a both efficient and provably correct manner. The choice of the metric is an issue by itself. Hausdorff distance is a first candidate, but we would like to control also differential quantities like normals or curvatures, as well as the aspect ratio of the facets.

Previous work Previous work on surface approximation in the Computer Graphics community resorts to two types of techniques: grids and particle systems.

The former type relies on a tessellation of $3d$-space[20, 8]. A polygonal approximation of the surface is computed inside each cell. Its vertices are located where the function has opposite signs. A global approximation is obtained by gluing all the polygonal patches together. This approach yields volume-based approximations. However, the topology of the surface is not systematically preserved, although methods have been proposed to guarantee the topological consistency of the result[8]. In addition, the point distribution is hardly controllable.

The latter type makes a set of particles migrate along the surface[25, 27, 10], according to an equation of diffusion. The connectivity between the points is then built by various means which usually do not guarantee the topology of the output mesh. A step forward has been done by Hart and Stander[16], who proposed to use Morse theory to capture the correct topology. Unfortunately, the method is mostly

heuristic and the authors do not provide a proof of correctness of their algorithm.

A lot of work has been recently devoted to the related matter of remeshing polyhedral surfaces. Although this issue is quite different from the previous one, our method can be used for both problems. We briefly survey prominent remeshing techniques. Particles methods have been proposed[26], but the time required to compute the solution of the equation of diffusion makes them less efficient than methods based on heuristics. The latter combine mesh simplification[17] and vertex optimization, and allow to control the point density[15, 19]. Another class of algorithms is based on global parameterization[18, 2], which allows quick mesh generation. But the use of a global parameterization requires to cut the surface into patches, which is a non-trivial issue, and which produces artificial 1-manifolds – from the cut-graph – that are visible in the final mesh.

The surface approximation problem has been well-studied in the Computational Geometry community, in the recent years. Edelsbrunner and Shah[13] gave a sufficient condition, called Closed Ball Property, for the restricted Delaunay triangulation (see definition 1) to be homeomorphic to the surface. Amenta and Bern[3] introduced the concept of ε-sample (see definition 3) for smooth surfaces, and worked out a sufficient condition on the density of point sets for the Closed Ball Property to be verified. Cheng et al.[7] considered the special case of skin surfaces. They redefined the notion of ε-sample in this context and proposed an algorithm that can mesh such surfaces with certified topology and curvature-adapted vertex density.

Contributions In this paper we revisit a greedy "farthest point" technique based on the restricted Delaunay triangulation, which was originally proposed by Chew[9] for mesh refinement. Given a 2-manifold S embedded in $\mathbb{R}^3$ and a mesh M whose vertices are on S, the mesh refinement problem consists in inserting new points of S as vertices of M and rebuilding the connectivity accordingly so that all facets of M meet some criterion. This problem has been well-studied in the planar case and provably good methods based on the Delaunay triangulation have been proposed[23, 24]. These methods allow to control the size and the shape of the triangles. Our main contributions are:

- to give a variant of Chew's algorithm and show (theorem 3) that the algorithm terminates on a wide variety of input surfaces and allows to control both the size and the aspect ratio of the facets.
- taking advantage of recent theoretical results on the restricted Delaunay triangulation[13, 3, 6], to show how the algorithm can produce surface samplings and approximations. Specifically, we show (theorem 4) that it can generate good point samples on smooth closed surfaces, in the sense given in definition 3. From this result we deduce that the restricted Delaunay triangulation has the right topology type and approximates the original surface in terms of Hausdorff distance, normals and area. Under some additional assumptions, we show (theorem 5) that the point samples generated are sparse, that is, their size is optimal up to a multiplicative constant factor. We also give some precisions in the case of uniform samples (theorem 6). Finally, we show (theorem 7) that Chew's algorithm can also be used to generate meshes whose facets have a bounded aspect ratio.
- to provide experimental evidence that the algorithm works well in practice and is able to deal with boundaries. The algorithm also shows a good behaviour on piecewise smooth surfaces, except near singularities. Results on polyhedral surfaces indicate that it is also suitable for polyhedral surface remeshing, as far as there are no sharp edges.

2. Restricted Delaunay triangulation and point samples

In this section, S is a 2-manifold embedded in $\mathbb{R}^3$, and Y is a point sample of S, *ie* a finite set of points of S. By $\mathrm{Del}(Y)$ we denote the 3-dimensional Delaunay triangulation of Y. The algorithm and analysis both rely on a special data structure, called restricted Delaunay triangulation, which is a subcomplex of the 3-dimensional Delaunay triangulation, defined as follows:

Definition 1 The *Delaunay triangulation of Y restricted to S*, denoted by $\mathrm{Del}_{|S}(Y)$, is the sub-complex of $\mathrm{Del}(Y)$ that consists of the facets of $\mathrm{Del}(Y)$ whose dual Voronoi edges intersect S. For any facet f of $\mathrm{Del}_{|S}(Y)$, we call *surface Delaunay ball* of f any empty ball circumscribed to f whose center lies on S.

It can be useful to relate the local density of a given point set lying on a surface to the local curvature of that surface. More precisely, we shall relate it to the distance to the skeleton of the surface, which is smaller.

Definition 2
- we call *maximal ball* any ball that is maximal (with respect to the inclusion) among the set of open balls included in $\mathbb{R}^3 \setminus S$.
- the *skeleton of S*, denoted by σ, is the topological closure of the union of the centers of all maximal balls.
- for a point $x \in \mathbb{R}^3$, we call *distance to the skeleton* at x, and write $d_\sigma(x)$, the Euclidean distance from x to the skeleton of S.

According to lemma 1 of [Amenta, Bern][3], d_σ is 1-Lipschitz. Another useful property of the skeleton is the following:

Lemma 1 (from proposition 13 of [Boissonnat, Cazals][5])
Let B be a ball that intersects S (resp. the boundary ∂S of S). If the intersection is not a topological disc (resp. a topological arc), then B contains a point of the skeleton of S (resp. ∂S).

Now, we introduce the notion of "good sample", in relation to a given 1-Lipschitz function, which can be for instance the distance to the skeleton.

Definition 3 (from [Amenta, Bern][3] and [Attali et al.][4])
Given a 1-Lipschitz function $\phi : S \to \mathbb{R}$, Y is an ε-*sample of* S *with respect to* ϕ if $\forall x \in S,\ |Y \cap B(x, \varepsilon\,\phi(x))| \geq 1$. Most of the time $\phi = d_\sigma$, which will be assumed by default.
- If ϕ is constant (say $\phi = 1$), then Y is called a *uniform* ε-*sample*.
- If $\forall x \in S,\ 1 \leq |Y \cap B(x, \varepsilon\,\phi(x))| \leq \kappa$, then the sample is called an (ε, κ)-*sample*.

Erickson[14] has shown that $\Omega\left(\frac{\mu(S)}{\varepsilon^2}\right)$, with $\mu(S) = \int_S \frac{dx}{d_\sigma^2(x)}$, is a lower bound on the number of points of any ε-sample of S, with $\varepsilon < \frac{1}{5}$. If an ε-sample of S contains $O\left(\frac{\mu(S)}{\varepsilon^2}\right)$ points, then it is said to be *sparse*.

Now we introduce two results on the restricted Delaunay triangulation that are proved in [Boissonnat, Oudot][6]. They will be used in sections 5 and 6.

The first result gives a relationship between the size of the surface Delaunay balls and the density of the vertices of the restricted Delaunay triangulation. The idea is that, if one forces the surface Delaunay balls to be small enough, then one can show that their union covers the whole surface. It follows that any point of the surface is close to the center of a surface Delaunay ball, and hence close to the vertices of a facet of the restricted Delaunay triangulation.

Theorem 1 Assume that S is smooth, compact, without boundaries, and that the distance to its skeleton σ has a lower bound $d_\sigma^{min} > 0$. Assume also that $\mathrm{Del}_{|S}(Y)$ has at least one facet on each connected component of S.
- if $\forall f \in \mathrm{Del}_{|S}(Y)$ each surface Delaunay ball of f has a radius at most $\frac{\varepsilon}{4}$, where $\varepsilon \leq 0.36\, d_\sigma^{min}$, then the set of vertices of $\mathrm{Del}_{|S}(Y)$ is a uniform ε-sample of S
- if $\forall f \in \mathrm{Del}_{|S}(Y)$ each surface Delaunay ball $B(c_f, r_f)$ of f has a radius r_f not greater than $\frac{\varepsilon}{6+5\varepsilon}\, d_\sigma(c_f)$, where $\varepsilon \leq 0.327$, then the set of vertices of $\mathrm{Del}_{|S}(Y)$ is an ε-sample of S.

The second result gives an upper bound on the size of any point set lying on a surface, with respect to a given 1-Lipschitz function. In particular, this result can be applied with the distance to the skeleton, or any point density function that is 1-Lipschitz, as for instance a constant density.

Theorem 2 Assume that S is smooth and compact, and that Y is an ε-sample of S, with $\varepsilon \leq \frac{1}{2}$. If S has some boundaries, assume in addition that Y contains a μ-sample of ∂S, with $\mu < \frac{1}{1+2\sqrt{2}}$. Let K be a positive number and $\psi : S \to \mathbb{R}$ a 1-Lipschitz function such that

$$\forall v \in Y,\ \mathrm{dist}\,(v, Y \setminus \{v\}) \geq K\,\psi(v) \qquad (1)$$

Then

$$|Y| \leq C\,\frac{\left(1 + K(\sqrt{2}-1)\right)^2}{\pi(\sqrt{2}-1)^2 K^2} \int_S \frac{dx}{\psi^2(x)}$$

where $C = 10$ if S has some boundaries and $C = \frac{4}{3}$ otherwise.

In particular, if $\psi = d_\sigma$ and $K = \varepsilon$, then the above theorem states that Y is a sparse ε-sample of S.

3. Chew's algorithm

The algorithm takes as input a pair (S, X), where S is a compact surface and X is a set of points lying on S. If X has some points on the boundary ∂S of S, we call boundary edges the segments that join consecutive points of X on ∂S.

The algorithm iteratively constructs a set of points $\bar{X}$, and maintains its restricted Delaunay triangulation $\mathrm{Del}_{|S}(\bar{X})$ throughout the process. $\bar{X}$ is initialized to X. Procedure *insert(p)* adds point p to $\bar{X}$ and updates $\mathrm{Del}_{|S}(\bar{X})$. If e is a boundary edge that is "encroached" (*ie* there exists a point of $\bar{X}$ inside its diametral ball), $midpoint(e)$ returns an intersection point of the bisector of e and ∂S. When this point is inserted, the edges it forms with the vertices of e become boundary edges, whereas e is no longer a boundary edge. For a facet f of $\mathrm{Del}_{|S}(\bar{X})$, $circumcenter(f)$ returns the center of a surface Delaunay ball of f. The algorithm is templated by a criterion ρ on the facets of $\mathrm{Del}_{|S}(\bar{X})$.

ALGORITHM

```
INITIALISATION
  X̄ = X; compute Del_|S(X̄)
REPEAT
  WHILE there remains any encroached
   boundary edge e
    insert(midpoint(e))
  LET f be any facet of Del_|S(X̄) that does not
   meet ρ
  LET p = circumcenter(f)
  IF p encroaches boundary edges s_1, ..., s_k,
  THEN
    FOR i = 1 TO k
      insert(midpoint(s_i))
  ELSE
    insert(p)
UNTIL no boundary edge is encroached and
all facets of Del_|S(X̄) meet ρ
```

At the end of the process, the algorithm returns $\bar{X}$ and $\mathrm{Del}_{|S}(\bar{X})$. Originally, Chew's algorithm did not compute the 3-dimensional Delaunay triangulation of the points. However, there are several advantages in using a 3-dimensional triangulation: it allows to handle topology changes, which is important especially during the first steps of the algorithm, and it provides a location data structure that allows to insert new points fast[11].

As for ρ, we shall use two measures: the aspect ratio and the size of facets.

Definition 4 (some refinement criteria)
Let f be a facet of $\mathrm{Del}_{|S}(\bar{X})$. We call θ its smallest angle and $B_f = B(c_f, r_f)$ its biggest surface Delaunay ball:

($\rho_{aspect\ ratio}$) f meets the criterion if $\frac{1}{2\sin\theta} \leq \beta$, where β is a positive constant

(ρ_{size}) f meets the criterion if $r_f \leq g(c_f)$, where $g : S \to \mathbb{R}$ has a lower bound $h > 0$

An easy computation shows that for any triangle t with smallest angle θ, $\frac{1}{2\sin\theta}$ is equal to the ratio between the radius of the circumcircle of t and the length of its smallest edge.

The algorithm can use one of the above criteria or it can combine both. For instance, one can first remove facets that are too big and then remove facets with a bad aspect ratio.

4. Termination of the algorithm

We prove that Chew's algorithm terminates after a finite number of steps. This result holds for surfaces without boundaries or with smooth or piecewise linear boundaries.

Definition 5 Let v be a point inserted by the algorithm. We call *insertion radius* of v (denoted by r_v) the distance from v to the "current" set of points $\bar{X}$ at the time when v is inserted. By convention, the insertion radius of any vertex w of the input point sample X is $\text{dist}(w, X \setminus \{w\})$.

The proof of termination relies on the fact that the insertion radius remains greater than a constant during the whole process. It follows that the points inserted by the algorithm cannot get too close to one another, and thus cannot be infinitely many.

Theorem 3 Let S be a compact surface and X a point sample of S. The algorithm is applied to (S, X):
- If S has no boundary, then the algorithm will terminate provided that $\beta \geq 1$.
- If S has piecewise linear boundaries such that all angles between consecutive boundary edges are greater than $\pi/4$, then the algorithm will terminate provided that $\beta \geq \sqrt{2}$.
- If S has smooth boundaries and if the initial set of points X contains a μ-sample of ∂S, with $\mu < 1/3$, then the boundary edges make no angle less than $2\pi/3$ and the algorithm will terminate provided that $\beta \geq \frac{2\mu}{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}\sqrt{2}$.

Proof The proof of the first two statements is similar to Shewchuk's proof for the planar case[24] and is omitted here. As for the third statement, the proof differs slightly but keeps the same flavour. Consider a point v that is inserted by the algorithm. We distinguish three different cases:
• v is the "circumcenter" of a facet f. By definition, the Delaunay ball σ_f of f that is centered at v is empty. Hence, the insertion radius r_v is equal to the radius of σ_f. If v is inserted because f does not meet ρ_{size}, then the radius of σ_f is greater than h. If v is inserted because f does not meet $\rho_{aspect\ ratio}$, then the radius of σ_f is not less than the radius of the circumcircle of f. Let p be the vertex of the smallest edge of f that has been inserted last. By definition, its insertion radius r_p is not greater than the length of the smallest edge of f. Since f does not meet the criterion ($\rho_{aspect\ ratio}$), we have $\frac{r_v}{r_p} \geq \frac{\text{radius of circumcircle}}{\text{length of smallest edge}} > \beta$.

• v is the "midpoint" of a boundary edge $[a,b]$ that is encroached by a circumcenter p. We know that $[a,b]$ is encroached only by p and that p will not be inserted. The insertion radius r_p of p is less than $\|a-b\|\frac{\sqrt{2}}{2}$ since p lies in the diametral ball of $[a,b]$. So, according to lemma 5, the insertion radius r_v of v is such that

$$\begin{aligned} r_v &\geq \delta_v \geq \frac{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}{2\mu}\frac{\|a-b\|}{2} \\ &\geq \frac{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}{2\mu}\frac{1}{\sqrt{2}}\, r_p \end{aligned}$$

• v is the "midpoint" of a boundary edge that is encroached by the vertex of an adjacent boundary edge. In fact, this situation cannot occur: indeed, according to lemma 4, all angles between boundary edges are greater than $2\pi/3$, since $\mu < \frac{1}{3}$.

The different cases are summarized in the flow graph of figure 1:
- If S has no boundary, then only loop (1) may occur. Hence, with $\beta \geq 1$, the insertion radius does not decrease when a badly shaped facet is split, whereas it remains greater than h when a badly sized facet is split. Thus, during the whole process the insertion radius remains greater than $d = \min(e, h)$, where e is the minimal distance between any two vertices of the input set of points.
- If S has smooth boundaries, then taking $\beta \geq \frac{2\mu}{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}\sqrt{2}$ makes the coefficients of loops (1) and $(1)+(2)$ greater than 1, which implies that the insertion radius, throughout the process, remains greater than $d = \min\left(e, \frac{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}{2\mu}\frac{h}{\sqrt{2}}\right)$.

In both cases, every point v of $\bar{X}$ remains at distance at least d from any other point of $\bar{X}$. It follows that the open balls B_v of radius $\frac{d}{2}$ centered at the points of $\bar{X}$ are pairwise disjoint. Now consider the volume $V = \left\{x \in \mathbb{R}^3 | \text{dist}(x,S) \leq \frac{d}{2}\right\}$. This volume is clearly bounded since the surface is compact, and it contains $\cup_{v\in\bar{X}} B_v$, which implies that there can be only a finite number of pairwise disjoint balls. □

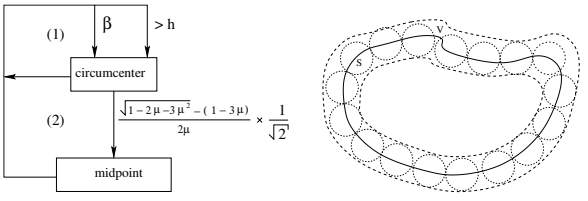

Figure 1: *Flow graph and volume V*

Theorem 3 asserts that, on surfaces without boundaries, the algorithm can generate meshes with no angle less than 30 degrees, provided that one sets $\beta = 1$. On surfaces with polygonal boundaries, no angle less than 20.7 degrees is created if one sets $\beta = \sqrt{2}$.

5. Approximating the surface

In this section, S denotes *a smooth compact surface without boundaries* such that, for any $x \in S$, $d_\sigma(x) \geq d_\sigma^{min} > 0$. Let X be a finite set of points on S. The algorithm is applied to (S, X) with the criterion (ρ_{size}), parameterized by a function g which is 1-Lipschitz. We call $\bar{X}$ the point sample generated by the algorithm.

We shall show that the algorithm can produce sparse ε-samples. We first consider the general case, and afterwards the case of unifom samples, for which better bounds can be provided. Approximation results come out as corollaries.

5.1. Generating ε-samples

To establish our results, we use theorem 1 which requires $\mathrm{Del}_{|S}(\bar{X})$ to be non-empty. The following lemma helps ensuring this condition.

Lemma 2 Assume that there exists a facet f_0 of $\mathrm{Del}_{|S}(X)$ that has a surface Delaunay ball $B_{f_0} = B(c_{f_0}, r_{f_0})$ such that $r_{f_0} \leq \frac{1}{3}\, g(c_{f_0})$. Then $f_0 \in \mathrm{Del}_{|S}(\bar{X})$.

Proof Assume that, at the end of the algorithm, $f_0 \notin \mathrm{Del}_{|S}(\bar{X})$. This implies that there exists a step at which the algorithm inserted a point v inside B_{f_0}. According to the size criterion, v is the center of a Delaunay ball B_f of a certain facet f, such that the radius r_f of B_f is greater than $g(v)$. Since v lies inside B_{f_0}, we have $\|v - c_{f_0}\| < \frac{1}{3}\, g(c_{f_0})$. And, since g is 1-Lipschitz,

$$g(v) \geq g(c_{f_0}) - \|v - c_{f_0}\| > \left(1 - \frac{1}{3}\right) g(c_{f_0})$$

Let a be one of the vertices of f_0. Since a is in B_{f_0}, we have

$$\|a - v\| \leq 2\, r_{f_0} \leq \frac{2}{3}\, g(c_{f_0}) < g(v) < r_f$$

which contradicts the fact that B_f is a Delaunay ball. □

Definition 6 f_0, defined as in the above lemma, is called a *seed-facet* – see figure 2. Lemma 2 claims that any seed-facet will remain a facet of $\mathrm{Del}_{|S}(\bar{X})$ throughout the course of the algorithm. Notice however that seed-facets may have also big surface Delaunay balls that will eventually be deleted.

From lemma 2 and theorem 1, we deduce that the algorithm can build ε-samples:

Theorem 4 Let ε be a positive constant such that $\varepsilon \leq 0.327$. We set $g \leq \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma$. Assume that $\mathrm{Del}_{|S}(X)$ has a seed-facet on each connected component of S. Then the vertices of $\mathrm{Del}_{|S}(\bar{X})$ form an ε-sample of S.

Proof Lemma 2 ensures that $\mathrm{Del}_{|S}(\bar{X})$ has at least one facet on each connected component of S. In addition, the size criterion ensures that every surface Delaunay ball $B(c_f, r_f)$ of any facet f of $\mathrm{Del}_{|S}(\bar{X})$ is such that $r_f \leq g(c_f) \leq \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma(c_f)$. So, the assumptions of theorem 1 are verified, which gives the result. □

Remark The vertices of $\mathrm{Del}_{|S}(\bar{X})$ belong to $\bar{X}$ (hence theorem 4 implies that $\bar{X}$ is an ε-sample of S), but the converse is not necessarily true.

5.2. A word about seed-facets

Figure 2 shows that seed-facets are preserved throughout the meshing process. The drawback of seed-facets is that, since they are smaller than the size criterion (at least three times as small), they force the algorithm to refine the mesh more than necessary in their vicinity. This problem can be avoided by using no seed-facet and choosing a few random points to start the process. In that case there is no guarantee that the final restricted Delaunay triangulation will be non-empty. However, this assumption is often verified in practice. For instance, experiments have shown that, with the torus of figure 2 and exactly three initial random points, one has a 20% chance that the meshing process succeeds. If one chooses an initial point set with a small enough diameter, then the chance of success grows up dramatically (92% with a diameter of 0.5 which is still far above the size criterion). In addition, if one chooses a bigger initial point set, then the chance of success also grows up dramatically (40% with 4 points, 78% with 5 points, and more than 94% with 6 points). In conclusion, it is usually unnecessary to use seed-facets in practice. For instance, the results shown in figures 3, 5, 6 and 7 were obtained without using any seed-facet. The initial point sets had various sizes, ranging from three points to a dozen of points.

Another possibility is to use seed-facets to mesh the surface, and then to decimate the mesh in their vicinity. The refinement process can then be restarted in order to guarantee that the set of vertices of the restricted Delaunay triangulation remains an ε-sample of the surface.

5.3. Approximation results

In this section, we give approximation results that are consequences of the fact that the algorithm can build ε-samples.

Topological guarantees Theorem 2 of [Amenta, Bern][3] states that the restricted Delaunay triangulation of an ε-sample of S, with $\varepsilon < 0.1$, is homeomorphic to S. From this result and theorem 4 we deduce the following corollary:

Corollary 1 Under the assumptions of theorem 4, with $\varepsilon < 0.1$, S and $\mathrm{Del}_{|S}(\bar{X})$ are homeomorphic.

Hausdorff distance The fact that $\bar{X}$ belongs to S is useful for bounding the Hausdorff distance between S and $\mathrm{Del}_{|S}(\bar{X})$. The following result says that this distance is $O(\varepsilon)$:

Corollary 2 Under the assumptions of theorem 4, with $\varepsilon \leq 0.327$, the Hausdorff distance between S and $\mathrm{Del}_{|S}(\bar{X})$ is bounded by $\varepsilon\, d_\sigma^{max}$, where $d_\sigma^{max} = \max\{d_\sigma(x),\ x \in S\}$.

Proof On the one hand, we use the criterion (ρ_{size}) with $g \leq \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma \leq \varepsilon\, d_\sigma$, which implies that every facet of $\mathrm{Del}_{|S}(\bar{X})$

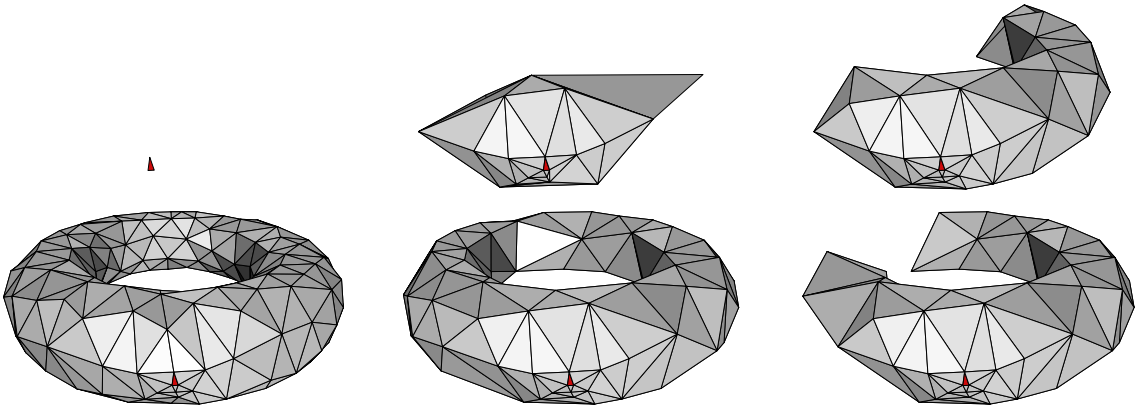

Figure 2: *Meshing process on the torus of equation* $(1.5-\sqrt{x^2+y^2}\,)^2+z^2-0.25=0$*, with a uniform size criterion of* 10^{-1} *and three starting points: the seed-facet remains in the restricted Delaunay triangulation throughout the process.*

has a radius less than $\varepsilon\, d_\sigma^{max}$. It follows that every point of $\mathrm{Del}_{|S}(\bar{X})$ is at distance less than $\varepsilon\, d_\sigma^{max}$ from $\bar{X} \subseteq S$.
On the other hand, theorem 4 says that the vertices of $\mathrm{Del}_{|S}(\bar{X})$ form an ε-sample of S. Thus every point of S is at distance less than $\varepsilon\, d_\sigma^{max}$ from those vertices, and hence from the restricted Delaunay triangulation. □

Normals approximation Lemma $7_{(b)}$ of [Amenta, Bern][3] bounds the angle between the normal to any facet f of $\mathrm{Del}_{|S}(\bar{X})$ and the normal to the surface at each of the vertices of f, when $\bar{X}$ is an ε-sample of S. From this result and theorem 4 we deduce the following:

Corollary 3 Under the assumptions of theorem 4, with $\varepsilon < \frac{1}{7}$, the angle between the normal to any facet f of $\mathrm{Del}_{|S}(\bar{X})$ and the normal to S at any of the vertices of f is less than $\frac{2\varepsilon}{1-7\varepsilon} + \arcsin\frac{\sqrt{3}\,\varepsilon}{1-\varepsilon}$.

Area approximation If $\bar{X}$ is an ε-sample of S, then the area of $\mathrm{Del}_{|S}(\bar{X})$ approximates the area of S – see [Morvan, Thibert][22]. From this result and theorem 4 we deduce the following:

Corollary 4 Under the assumptions of theorem 4, with $\varepsilon \leq 0.327$, there exist two constants C_1 and C_2 depending on ε, such that $C_1\ \mathrm{Area}(S) \leq \mathrm{Area}(\mathrm{Del}_{|S}(\bar{X})) \leq C_2\ \mathrm{Area}(S)$, and $\lim_{\varepsilon\to 0} C_1(\varepsilon) = \lim_{\varepsilon\to 0} C_2(\varepsilon) = 1$.

5.4. Sparse ε-samples

Provided that X contains few points, the algorithm actually produces sparse ε-samples. For simplicity, we assume that X contains exactly three points per connected component of S. This is no real loss of generality and the following result holds for any set X of constant size.

Theorem 5 Let ε be a positive constant such that $\varepsilon \leq \frac{\sqrt{57}-7}{4} \approx 0.14$. We set $g = \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma$. Assume that X has exactly 3 points per connected component of S, and that $\mathrm{Del}_{|S}(X)$ has a seed-facet on each connected component of S. Then $\bar{X}$ is a sparse ε-sample of S.

Proof Since $\varepsilon \leq \frac{\sqrt{57}-7}{4} \leq 0.327$ and $\mathrm{Del}_{|S}(X)$ has a seed-facet on each connected component of S, theorem 4 implies that $\bar{X}$ is an ε-sample of S.
To bound the cardinality of $\bar{X}$, we shall use theorem 2 with $Y = \bar{X} \setminus X$. We first show that $\bar{X} \setminus X$ is a $\frac{1}{2}$-sample of S. Let S_i be a connected component of S. We have $\varepsilon < 1$, thus at any point $x \in S_i$, $B(x, \varepsilon\, d_\sigma(x)) \cap S$ is a topological disc. Therefore, $B(x, \varepsilon\, d_\sigma(x))$ intersects S_i only. This means that, since $\bar{X}$ is an ε-sample of S, we have $\forall x \in S_i$, $|B(x, \varepsilon\, d_\sigma(x)) \cap S_i \cap \bar{X}| \geq 1$, *ie* $\bar{X} \cap S_i$ is an ε-sample of S_i with respect to d_σ. Let f_i be the seed-facet associated with S_i. Since $g = \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma < 0.17\, d_\sigma$, lemma 6 ensures that every edge of f_i is incident to another facet of $\mathrm{Del}_{|S}(\bar{X})$. Moreover, since $g < d_\sigma$, every facet of $\mathrm{Del}_{|S}(\bar{X})$ has its three vertices in the same connected component of S. In particular, all the facets of $\mathrm{Del}_{|S}(\bar{X})$ that are incident to f_i have their vertices in S_i. So, the assumptions of lemma 7 are verified, with $\Sigma = S_i$ and $\phi = d_\sigma$, hence $\bar{X} \cap S_i$ minus the set of vertices of f_i is a $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}$-sample of S_i with respect to d_σ. By assumption, the set of vertices of f_i is exactly $X \cap S_i$, thus $(\bar{X} \cap S_i) \setminus (X \cap S_i)$ is a $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}$-sample of S_i with respect to d_σ. Since this is true for every connected component of S, we conclude that $\bar{X} \setminus X$ is a $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}$-sample of S. Since $\varepsilon \leq \frac{\sqrt{57}-7}{4}$, we have $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon} \leq \frac{1}{2}$.
Now, let v be a point of $\bar{X} \setminus X$, *ie* a point that has been inserted by the algorithm. Let $w \neq v$ be the point of $\bar{X}$ that is closest to v. We distinguish two cases: 1. $w \in X$ or has been inserted before v, 2. w has been inserted after v. In the former case, we have, according to the size criterion,

$$\|v-w\| > \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma(v)$$

In the latter case, we have, for the same reason,

$$\|v-w\| > \frac{\varepsilon}{6+5\varepsilon}\, d_\sigma(w) \geq \frac{\varepsilon}{6+5\varepsilon}\, (d_\sigma(v) - \|v-w\|)$$

which gives

$$\|v-w\| > \frac{\varepsilon}{6(1+\varepsilon)}\, d_\sigma(v)$$

In both cases, we have $\text{dist}(v, (\bar{X} \setminus X) \setminus \{v\}) \geq \text{dist}(v, \bar{X} \setminus \{v\}) = \|v-w\| > \frac{\varepsilon}{6(1+\varepsilon)}\, d_\sigma(v)$. Since $\bar{X} \setminus X$ is a $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}$-sample of S with $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon} \leq \frac{1}{2}$, theorem 2 (with $Y = \bar{X} \setminus X$, $\psi = d_\sigma$ and $K = \frac{\varepsilon}{6(1+\varepsilon)}$) says that

$$|\bar{X} \setminus X| \leq \frac{4\left(6+(\sqrt{2}+5)\varepsilon\right)^2}{3\pi(\sqrt{2}-1)^2}\, \frac{\mu(S)}{\varepsilon^2}$$

where $\mu(S) = \int_S \frac{dx}{d_\sigma^2(x)}$. Now, by assumption, $|X| = 3k$, where k is the number of connected components of S. Thus $|\bar{X}| = 3k + |\bar{X} \setminus X|$. Note that for any connected component S_i of S, the distance to the skeleton σ_i of S_i (considered independently from the other connected components of S) is no less than $\frac{1}{2}\, d_\sigma$. Thus, $\int_{S_i} \frac{dx}{d_\sigma^2(x)} \geq \frac{1}{4} \int_{S_i} \frac{dx}{d_{\sigma_i}^2(x)}$, which is greater than π according to lemma 8. Hence $\mu(S) \geq \pi k$, which gives

$$|\bar{X}| \leq \left(\frac{3\varepsilon^2}{\pi} + \frac{4\left(6+(\sqrt{2}+5)\varepsilon\right)^2}{3\pi(\sqrt{2}-1)^2} \right) \frac{\mu(S)}{\varepsilon^2}$$

and, since $\varepsilon \leq \frac{\sqrt{57}-7}{4}$,

$$|\bar{X}| \;\leq\; C \;\times\; \frac{\mu(S)}{\varepsilon^2}$$

with

$$C \;=\; \frac{3\left(\sqrt{57}-7\right)^2}{16\pi} + \frac{\left(24+(\sqrt{2}+5)(\sqrt{57}-7)\right)^2}{12\pi(\sqrt{2}-1)^2} \;<\; 117.17$$

□

Remarks

Since the vertices of $\text{Del}_{|S}(\bar{X})$ belong to $\bar{X}$, the above result and theorem 4 imply that they also form a sparse ε-sample of S.

When ε tends to zero, our upper bound is equivalent to $\frac{48}{\pi(\sqrt{2}-1)^2}\, \frac{\mu(S)}{\varepsilon^2}$, which is about 1500 times the lower bound given by Erickson[14]. This suggests that our bound is still far from being tight.

5.5. Generating uniform samples

If one takes the function g of the size criterion to be constant (note that it is still 1-Lipschitz), then the algorithm will generate uniform samples. In fact, if $\bar{X}$ is an ε-sample, then it is a uniform $(\varepsilon\ d_\sigma^{max})$-sample, where $d_\sigma^{max} = \max\{d_\sigma(x),\ x \in S\}$). However, better bounds can be achieved when g is constant, as shown below.

Theorem 6 We set $g = h$, where h is a positive constant less than $0.09\ d_\sigma^{min}$. Assume that X has exactly three points per connected component of S, and that $\text{Del}_{|S}(\bar{X})$ has a seed-facet on each connected component of S. Then $\bar{X}$ is a uniform $(4h, 515)$-sample of S.

Proof Lemma 2 ensures that $\text{Del}_{|S}(\bar{X})$ has at least one facet on each connected component of S. Thus, since $4h \leq 0.36\ d_\sigma^{min}$, we can use theorem 1 which says that the vertices of $\text{Del}_{|S}(\bar{X})$ form a uniform $4h$-sample of S. It follows that $\bar{X}$ is also a uniform $4h$-sample of S.
Now, let v be a point of $\bar{X} \setminus X$, *i.e.* a point that has been inserted by the algorithm. Let w be any point of $(\bar{X} \setminus X) \setminus \{v\}$. If w has been inserted before v, then according to the size criterion we have $\|v-w\| > h$. If w has been inserted after v, then the size criterion also says that $\|v-w\| > h$. Thus, the points of $\bar{X} \setminus X$ are centers of pairwise disjoint balls of radius $\frac{h}{2}$. Hence, for every point $x \in S$, the number of points of $\bar{X} \setminus X$ that lie inside $B(x, 4h)$ is less than

$$\frac{\frac{4}{3}\,\pi\,(4h)^3}{\frac{4}{3}\,\pi\,\left(\frac{h}{2}\right)^3} \;=\; 512$$

Since $4h < d_\sigma^{min}$, at each point $x \in S$ the ball $B(x, 4h)$ intersects only the connected component of S where x lies. Thus, since X has exactly three points per connected component, we have $|B(x, 4h) \cap X| \leq 3$. Hence the number of points of $\bar{X}$ that lie inside $B(x, 4h)$ is less than $512 + 3 = 515$. □

Remark The fact that $\bar{X}$ is a uniform (ε, κ)-sample implies that, for a given surface, the size of $\bar{X}$ is $O\left(\frac{1}{\varepsilon^2}\right)$, which is optimal for uniform ε-samples.

6. Bounding the aspect ratio

Once an approximation of the surface has been obtained (or is given), the algorithm can be used to remove the skinny facets: this can be done with the help of the criterion ($\rho_{\textit{aspect ratio}}$). Theorem 3 gives bounds on the angles of the facets of the resulting triangulated surface. The following result provides a worst-case optimal upper bound on the size of the output with respect to the size of the input. It roughly shows that we can bound the aspect ratio of the triangles without significantly increasing the number of vertices. The upper bound we give in the next theorem is computed according to a special measure, called local feature size, defined as follows:

Definition 7 Let S be a surface and X a set of points sampled from S. Consider the graph G made of the vertices of X and of the boundary edges of S. The *local feature size* at x, denoted by $\text{lfs}_X(x)$, is defined as the radius of the smallest ball centered at x that intersects two nonincident vertices or edges of G. According to [Ruppert][23], lfs_X is 1-Lipschitz.

Theorem 7 Let S be a smooth compact surface and X an ε-sample of S, with $\varepsilon \leq \frac{1}{2}$. If S has some boundaries, assume that X contains a μ-sample of ∂S, with $\mu < \frac{1}{2(1+\sqrt{2})}$. If

one runs the algorithm with the criterion ($\rho_{aspect\ ratio}$) only, where $\beta > \frac{2\mu}{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}\sqrt{2}$, then the final number of points of $\bar{X}$ is at most

$$C \cdot \int_S \frac{dx}{\mathrm{lfs}^2_{\bar{X}}(x)}$$

where C is a constant that depends only on β and μ.

Proof This result is a simple application of theorem 2. However, we have to make sure that equation (1) is verified. As in the planar case[23, 24], one can start from the relations that we established in the proof of theorem 3 and show that there exists a constant $D_S \geq 1$ such that for every point $v \in \bar{X}$, $\frac{\mathrm{lfs}_X(v)}{\mathrm{dist}(v,\bar{X}\setminus\{v\})} \leq D_S + 1$.
Hence, (1) holds with $Y = \bar{X}$, $\psi = \mathrm{lfs}_X$ and $K = \frac{1}{D_S+1}$. We can then apply theorem 2 which gives the result. □

Remark The above upper bound is worst-case optimal. Indeed, if S is a planar surface then the problem becomes 2-dimensional, and according to [Mitchell][21], any solution requires $\Omega\left(\int_S \frac{dx}{\mathrm{lfs}^2_{\bar{X}}(x)}\right)$ points.

7. Experimental results – discussion

Implementation The algorithm has been implemented using the C^{++} library CGAL, which allowed us to derive the restricted Delaunay triangulation from the 3-dimensional Delaunay triangulation. Given a surface S and a point sample $\bar{X}$, $\mathrm{Del}_{|S}(\bar{X})$ is templated by the type of S – polyhedron, parametric or implicit surface etc. It is represented by marking the facets of $\mathrm{Del}(\bar{X})$ whose dual Voronoi edges intersect S.

An advantage of the algorithm is that, except for the *in_sphere()* predicate of the 3-dimensional Delaunay triangulation, the only predicate to implement is *does_segment_intersect_surface()*, which tells whether a segment intersects the surface or not. The related constructor, *intersect_segment_with_surface()*, is implemented in the same way. Both of them must be adapted to the type of the input surface. For instance, on an implicit surface given by the equation $P(x,y,z) = 0$, where P is a polynomial, one can use either an algebraic technique or a numerical technique. The former reduces to solving a univariate equation of degree $\deg(P)$, whereas the latter can be implemented as a simple binary search that uses the fact that if the signs of P at both endpoints are different then the segment intersects the surface. Although not exact and with a time complexity that depends highly on the precision of the search, the numerical approach is useful in practice since it can be extended to surfaces that are not algebraic nor semi-algebraic. However, the algebraic method is much more efficient in practice. Our implementation of the algebraic version uses the C^{++} library SYNAPS, which was designed to solve algebraic systems[12].

Our theoretical results require to know some estimate of d_σ. If the algorithm is used with a uniform size criterion, we only need a lower bound d_σ^{min}, which is easily available in many applications. If the algorithm is used with a local size criterion, we need an estimate of d_σ at all inserted points. One can approximate $d_\sigma(x)$ by the distance from x to the set of poles[3] of the Voronoi diagram of $\bar{X}$. Although we do not have theoretical guarantees then, good results have been observed in practice – see figure 3 (center).

Meshing surfaces Figure 5 shows some implicit surfaces that have been meshed by the algorithm, using as size criterion a linear function of d_σ. Up left is the "chair" of equation $(x^2+y^2+z^2-23.75)^2 - 0.8((z-5)^2-2x^2)((z+5)^2-2y^2) = 0$: degree 4 and genus 3. Down left is the "tanglecube" of equation $x^4-5x^2+y^4-5y^2+z^4-5z^2+10=0$: degree 4 and genus 5. Since both surfaces are closed and smooth, corollary 1 ensures that with a small enough size criterion the algorithm produces triangulated surfaces with the right topology type. This is observed here, even with a size criterion ($g = 10^{-1} d_\sigma$) that is far above the theoretical bound.

We also experimented with surfaces with boundaries or singularities. Figure 5 shows Barth's decic surface (up center) and Sheffer's surface (down center). The former has degree 10 and at least 40 singular points ! One can notice that these points have not been identified – see the zoom up right, so that the topology could not be captured. The result is satisfactory on the whole, except near the singularities, where d_σ tends to zero. Sheffer's surface is periodic, because its function is a trigonometric polynomial. Since it is not compact, we have only considered the portion of the surface that lies inside a ball. d_σ does not vanish on this surface and, as can be expected, the topology is preserved. However, the boundaries are jagged because they were not explicitely computed. Figure 6 shows some results on complex surfaces obtained from simple implicit primitives by applying boolean operations and offsets. Complex surfaces are meshed directly, without meshing first the primitives and then computing the boolean operations or the offsets on the triangulated primitives. The example in the middle is the solvent excluded surface of a molecule of alanin, with a probe of radius 20 picometers: this surface is a combination of spherical and toric patches – see the zoom on the left. The example on the right consists of two letters that were built as unions of cylinders.

Dealing with boundaries and singularities In the litterature, the usual method that is used to handle boundaries and singularities, samples them before the rest of the surface[2]. An advantage of Chew's algorithm is that it allows to handle boundaries and singularities at the same time as the rest of the surface. In section 3, we saw how to handle boundaries. The same treatment can be applied to singular curves. This results in meshes that include polygonal approximations of the singular curves. Figure 7 (up left), which presents the end of a tubular surface remeshed by the algorithm – the blue curve is the original boundary, shows that boundaries are correctly approximated. Figure 7 (down left) shows the

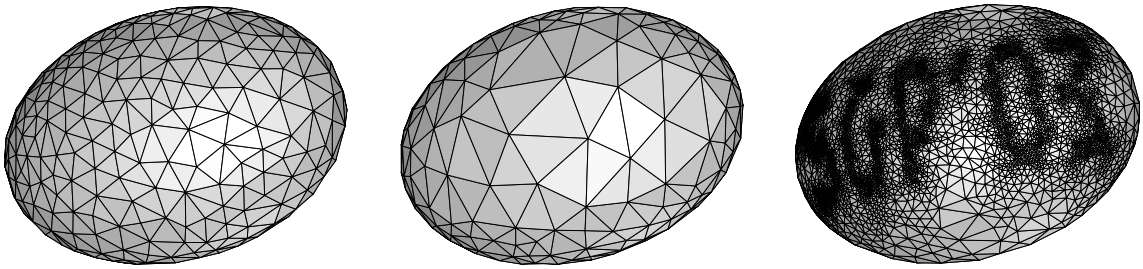

Figure 3: *Various size criteria: uniform (left), curvature-adapted (center), customized (right).*

result of the algorithm on Steiner's Roman surface, which is one of the three possible surfaces obtained by sewing a Moebius strip to the edge of a disk. This surface has one triple point, three singular lines and six pinch points. The singular lines were correctly approximated by the algorithm, except near the pinch points (circled in red).

Boundaries and singularities play different roles. Indeed, by definition, d_σ tends to zero near a singularity, whereas it does not near a boundary. Thus it should be possible to extend our theoretical results to smooth surfaces with boundaries, whereas the sampling theory we currently use clearly fails in the vicinity of singularities. Practicle experiments corroborate this observation, since the result of the algorithm is always a manifold near boundaries, whereas it is not always so near singularities.

The case of polyhedra. Theoretically, every edge of a polyhedron is problematic since it is a singularity. But in practice, it turns out that the results are manifold everywhere except near the sharp edges. A simple post-process that detects the edges with more than two incident facets and deletes the guilty facets often suffices to get a manifold surface. Note that taking the sharp edges as boundary edges greatly helps the algorithm to avoid non-manifold results. An example is shown in figure 7 (center and right): the octopus was designed at first by an artist, using a quad-dominant mesh. We show the output of the algorithm, with a very thin size criterion adapted to d_σ, and a dozen of random points to start with. One can notice that the final point density is bigger near the edges of the initial mesh – see figure 7 (up right), and that all the details have been remarkably captured: see for instance the eyes – figure 7 (down right) – which are separate discs floating in the air.

Timing and output complexity Figure 4 shows the evolution of the output size and computation time (with a processor at 900 Mhz) with respect to the bound h of the size criterion, on the sphere of equation $x^2+y^2+z^2-1=0$. The output size is given by the number of vertices of the restricted Delaunay triangulation, whereas the computation time is given in centiseconds. It turns out that the output size is $O\left(\frac{1}{h^2}\right)$, as predicted by theorems 5 and 6.

Concerning the time complexity: at each step, the algorithm inserts a point in the 3-dimensional Delaunay triangulation. To this end, it deletes the facets that are in conflict with the point that is being inserted, and then stars the hole. For every new facet, it checks whether the dual Voronoi edge intersects the surface. Thus the behaviour of the algorithm is that of the incremental Delaunay triangulation, except that no point location is needed since only "circumcenters" of known facets or "midpoints" of known edges are inserted. It has been shown recently[4] that the size of the 3-dimensional Delaunay triangulation of an (ε,κ)-sample of n points lying on a generic smooth surface is $O(n\log n)$. The worst-case time complexity of one step of the algorithm is thus $O(\phi\ n\log n)$, where n is the size of the current point sample, and ϕ is the time complexity of the predicate *does_segment_intersect_surface()*, which can be considered as a constant once the surface is given. This yields an overall worst-case time complexity of $O(N^2\log N)$, where N is the size of the output. Assuming that the points are inserted in random order, the expected running time reduces to $O(N\log N)$ since no point location is needed. This bound is

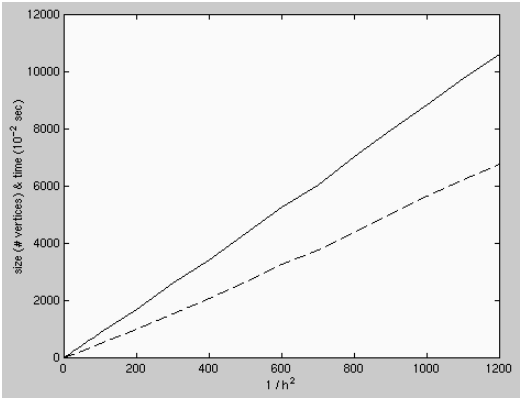

Figure 4: *Output size (solid line) and time complexity (dashed line) versus* $\frac{1}{h^2}$.

usually observed in practice, as illustrated in figure 4. This example also shows that the algorithm is able to generate

Surface (degree)	ellipsoid (2)	torus (4)	"chair" (4)	"tanglecube" (4)	Steiner (4)	Barth's Decic (10)
Criterion value	2.10^{-2}	10^{-2}	10^{-2}	10^{-2}	10^{-3}	10^{-5}
Output size (points)	166	1,297	6,323	4,162	8,461	142,605
Comput. time (sec.)	0.9	38	181	96	141	13,440
Rate (points per sec.)	184	34	35	43	60	11

Table 1: *Statistics on some implicit models presented in the paper.*

10,000 points on the sphere within less than 65 seconds. However, ϕ depends highly on the number of monomials of the polynomial that describes the surface, and the algorithm gets significantly slower when the degree of the polynomial increases – see table 1.

8. Conclusion – future work

We have revisited a greedy meshing algorithm. This algorithm can be applied to various types of surfaces and is quite simple. It is based on the 3-dimensional Delaunay triangulation and requires only few additional predicates. We have given some theoretical guarantees. The algorithm can produce various kinds of samples and construct good approximations with respect to several criteria: topology, Hausdorff distance, normals, area.

We have implemented the algorithm and presented experimental results that corroborate (and sometimes are better than) the theory.

Currently, the theoretical results are limited to closed smooth surfaces, but it should be possible to extend them to smooth surfaces with boundaries as the experimental results suggest.

Handling the singularities is a more difficult issue that should require the use of a more general theory of surface sampling. However, the algorithm appears to be extremely robust even in the presence of singularities.

Acknowledgments

The authors would like to thank David Cohen-Steiner and Marc Pouget for fruitful discussions, as well as Pierre Alliez for helpful remarks on the structure of the paper.

References

1. P. Agarwal and S. Suri. Surface approximation and geometric partitions. *Proc. 5th Annu. ACM Sympos. Discrete Algorithms*, 1994. pp 24-33.
2. P. Alliez, E. Colin de Verdière, O. Devillers and M. Isenburg. Isotropic Surface Remeshing. *Solid Modelling and Applications*, 2003.
3. N. Amenta and M. Bern. Surface Reconstruction by Voronoi Filtering. *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 1998, pp 39-48.
4. D. Attali, J-D Boissonnat and A. Lieutier. Complexity of the Delaunay Triangulation of Points on Surfaces: the Smooth Case. *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, 2003.
5. J-D Boissonnat and F. Cazals. Natural Neighbour Coordinates of Points on a Surface. *Comp. Geom. Theory and Appl.*, 155-174, 2001.
6. J-D Boissonnat and S. Oudot. Restricted Delaunay Triangulation and Point Samples. In preparation.
7. H-L Cheng, T. K. Dey, H. Edelsbrunner and J. M. Sullivan. Dynamic Skin triangulations. *Proc. 12th Annu. ACM Sympos. Discrete Algorithms*, 2001, pp 47-56.
8. E. V. Chernyaev. Marching Cubes 33: Construction of Topologically Correct Isosurfaces. Technical report CERN CN 95-17, 1995.
9. L. P. Chew. Guaranteed-Quality Mesh Generation for Curved Surfaces. *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, 1993, pp 274-280.
10. M. Desbrun, N. Tsingos and M-P Gascuel. Adaptive Sampling of Implicit Surfaces for Interactive Modeling and Animation. *Proc. Implicit Surfaces*, pp 171-185, 1995.
11. O. Devillers. The Delaunay hierarchy. *Internat. Journal Found. Comput. Sci.*, 2002, vol 283/1, pp 203-221.
12. G. Dos Reis, B. Mourrain, R. Rouillier and Ph. Trébuchet. An environment for Symbolic and Numeric Computation. *Proc. Internat. Conf. on Mathematical Software*, 2002, pp. 239–249.
13. H. Edelsbrunner, N.R. Shah. Triangulating topological spaces. *International Journal of Computational Geometry and Applications*, Vol. 7, No. 4 (1997), pp 365-378.
14. J. Erickson. Nice point sets can have nasty Delaunay Triangulations. *Proc. 17th Annu. Sympos. Comput. Geom.*, pp 96-105, 2001.
15. I. Guskov, W. Sweldens and P. Schröder. Multiresolution Signal Processing for Meshes. *Proc. SIGGRAPH*, pp 325-334, 1999.
16. J. Hart and B. T. Stander. Guaranteeing the Topology of an Implicit Surface Polygonization for Interactive Modeling. *Proc. SIGGRAPH*, 1997.
17. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stuetzle. Mesh Optimization. *Proc. SIGGRAPH*, pp 19-26, 1993.
18. K. Hormann, U. Labsik and G. Greiner. Remeshing Triangulated Surfaces with Optimal Parameterizations. *Computer-Aided Design 33*, pp 779-788, 2001.
19. K. Hormann, U. Labsik, M. Meister and G. Greiner. Hierarchical Extraction of Iso-Surfaces with Semi-Regular Meshes. *Solid Modelling and Applications*, 2002.
20. W. E. Lorensen and H. E. Cline. Marching Cubes: a High-Resolution 3D Surface Construction Algorithm. *Proc. SIGGRAPH*, 1987.
21. S. A. Mitchell. Cardinality bounds for triangulations with bounded minimum angle. *Proc. 6th Annu. Canadian Conference Comput. Geom.*, pp 326-331, 1994.
22. J-M Morvan and B. Thibert. On the Approximation of the Area of a Surface. INRIA RR-4375, 2002.
23. J. Ruppert. A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation. *Journal of Algorithms*, May 1995.
24. J. R. Shewchuk. Delaunay Refinement Algorithms for Triangular Mesh Generation. *Comput. Geom. Theory & Applications*, 22(1-3):21-74, may 2002.
25. G. Turk. Generating Textures for Arbitrary Surfaces using Reaction-diffusion. *Proc. SIGGRAPH*, pp 289-298, 1991.
26. G. Turk. Re-tiling Polygonal Surfaces. *Proc. SIGGRAPH*, pp 55-64, 1992.
27. A. Witkin and P. Heckbert. Using Particles to Sample and Control Implicit Surfaces. *Computer Graphics*, pp 269-278, 1994.

Appendix: Technical lemmata

Results used to prove the termination of the algorithm

Lemma 3 (proved in [Boissonnat, Oudot][6]) Let C be a curve in $\mathbb{R}^3$. We call d_ζ the distance to its skeleton ζ. Let Y be a μ-sample of C, with $\mu < 1/2$. Let a and b be two points of Y that are consecutive on C. Then

$$\|a-b\| \;\leq\; \frac{2\mu}{1-\mu}\, d_\zeta(a)$$

Lemma 4 (proved in [Boissonnat, Oudot][6]) Let C be a curve in $\mathbb{R}^3$, and let Y be a μ-sample of C, with $\mu < 1/2$. Let b, a and c be three points of Y that are consecutive on C. Then

$$\widehat{bac} \geq 2\arccos\frac{\mu}{1-\mu}$$

Lemma 5 Let S be a 2-manifold that has a smooth boundary. Let $\bar{X}$ be a finite set of points on S that contains a μ-sample $\widetilde{X}$ of ∂S, with $\mu < 1/3$. Let a and b be two points of $\widetilde{X}$ that are consecutive on ∂S. If the boundary edge $[a,b]$ is split by inserting point $v \in \partial S$, then

$$\delta_v \;\geq\; \frac{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}{2\mu}\;\frac{\|a-b\|}{2}$$

where δ_v is the distance from v to the diametral sphere of $[a,b]$.

Proof The point v which is inserted is at the intersection between ∂S and the bisector of $[a,b]$.

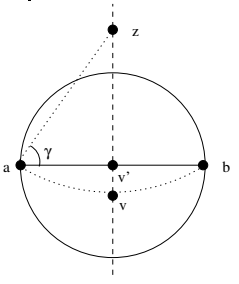

Let ζ be the skeleton of ∂S and let z be any point of the bisector line of $[a,b]$ in plane $\Pi = (a,b,v)$. If $\gamma = \widehat{baz}$, we have

$$\|z-a\| = \|z-b\| = \frac{\|a-b\|}{2\cos\gamma}$$

And, according to lemma 3,

$$\|a-b\| \leq \frac{2\mu}{1-\mu}\, d_\zeta(a)$$

since a and b are consecutive points of $\widetilde{X}$ on ∂S. Thus

$$\|z-a\| = \|z-b\| \leq \frac{\mu}{(1-\mu)\cos\gamma}\, d_\zeta(a)$$

which gives, since d_ζ is 1-Lipschitz,

$$d_\zeta(z) \geq d_\zeta(a) - \|z-a\| \geq \left(\tfrac{(1-\mu)\cos\gamma}{\mu} - 1\right)\|z-a\|$$

Now take point z so that $\cos\gamma = \frac{2\mu}{1-\mu}$ and z is on the side of (a,b) that does not contain v. γ is well-defined as far as $\mu \leq 1/3$. For this particular value of γ we have $d_\zeta(z) \geq \|z-a\|$. Thus point z is at the center of an open ball of radius $\|z-a\|$ which does not intersect ζ. According to lemma 1, the intersection between this ball and ∂S is a topological arc. And since points a and b are inside the ball, all points of ∂S between a and b, and in particular v, are also inside the ball. Hence the distance from v to the diametral sphere of $[a,b]$ is

$$\delta_v = \tfrac{\|a-b\|}{2} - \|v-v'\| \geq \tfrac{\|a-b\|}{2} - (\|z-a\| - \|z-v'\|)$$

Then, for $\mu < 1/3$ and $\gamma = \arccos\frac{2\mu}{1-\mu}$,

$$\begin{aligned}\|z-a\| - \|z-v'\| &= \tfrac{\|a-b\|}{2\cos\gamma} - \tfrac{\|a-b\|}{2}\tan\gamma \\ &= \tfrac{\|a-b\|}{2}\,\tfrac{(1-\mu)-\sqrt{1-2\mu-3\mu^2}}{2\mu}\end{aligned}$$

and

$$\begin{aligned}\delta_v &\geq \left(1 - \tfrac{(1-\mu)-\sqrt{1-2\mu-3\mu^2}}{2\mu}\right)\tfrac{\|a-b\|}{2} \\ &\geq \tfrac{\sqrt{1-2\mu-3\mu^2}-(1-3\mu)}{2\mu}\,\tfrac{\|a-b\|}{2}\end{aligned}$$

□

Results used to generate sparse ε-samples

Lemma 6 (proved in [Boissonnat, Oudot][6]) Let Σ be a smooth compact surface without boundaries. We call d_τ the distance to its skeleton τ. Let $\bar{X}$ be a set of points on Σ. Assume that $\forall f \in \mathrm{Del}_{|\Sigma}(\bar{X})$ every surface Delaunay ball $B(c_f, r_f)$ of f is such that $r_f \leq 0.17\; d_\tau(c_f)$. Then every edge that is incident to a facet of $\mathrm{Del}_{|\Sigma}(\bar{X})$ is actually incident to at least two facets of $\mathrm{Del}_{|\Sigma}(\bar{X})$.

Lemma 7 Let Σ be a smooth compact connected surface. Let $\bar{X}$ be an ε-sample of Σ, with respect to a given 1-Lipschitz function ϕ. Let f_0 be a facet of $\mathrm{Del}_{|\Sigma}(\bar{X})$ such that every edge of f_0 is incident to another facet of $\mathrm{Del}_{|\Sigma}(\bar{X})$. Call X the set of vertices of f_0. Then $\bar{X} \setminus X$ is a $\frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}$-sample of Σ, with respect to ϕ.

Proof Let $x \in \Sigma$. Let v be the point of $\bar{X}$ that is closest to x. Then $\|x-v\| \leq \varepsilon\,\phi(x)$, since $\bar{X}$ is an ε-sample of Σ with respect to ϕ. Now there are two cases: either $v \in \bar{X} \setminus X$, either $v \in X$.
- If $v \in \bar{X} \setminus X$, then

$$\mathrm{dist}\,(x, \bar{X} \setminus X) = \|x-v\| \leq \varepsilon\,\phi(x) \leq \frac{(3+\varepsilon)\varepsilon}{1-\varepsilon}\,\phi(x)$$

- If $v \in X$, then , since every edge of f_0 is incident to another facet of $\mathrm{Del}_{|\Sigma}(\bar{X})$, there exists a facet $f \neq f_0$ that is incident to v. At least one vertex of f is not a vertex of f_0, *ie* is in $\bar{X} \setminus X$. Let w be such a vertex. By definition, (v,w) is an edge of f, *ie* an edge of $\mathrm{Del}_{|\Sigma}(\bar{X})$, thus its dual Voronoi face intersects Σ. Let y be a point at the intersection. By definition of the Voronoi diagram, $\|y-v\| = \|y-w\| = \mathrm{dist}(y, \bar{X})$, which is less than $\varepsilon\,\phi(y)$ since $\bar{X}$ is an ε-sample of Σ with respect to ϕ. Thus we have $\|y-v\| \leq \varepsilon\,\phi(y) \leq \varepsilon\,(\phi(v) + \|y-v\|)$, that is, $\|y-v\| \leq \frac{\varepsilon}{1-\varepsilon}\,\phi(v)$. Finally, we get

$$\begin{aligned}\mathrm{dist}(x, \bar{X} \setminus X) &\leq \|x-w\| \leq \|x-v\| + \|v-w\| \\ &\leq \|x-v\| + 2\,\|v-y\| \\ &\leq \varepsilon\,\phi(x) + \tfrac{2\varepsilon}{1-\varepsilon}\,(\phi(x) + \|x-v\|) \\ &\leq \varepsilon\,\phi(x) + \tfrac{2\varepsilon}{1-\varepsilon}\,(1+\varepsilon)\,\phi(x) \\ &\leq \tfrac{(3+\varepsilon)\varepsilon}{1-\varepsilon}\,\phi(x)\end{aligned}$$

□

Lemma 8 Let Σ be a smooth compact connected surface without boundaries. We call d_τ the distance to its skeleton τ. We have

$$\int_\Sigma \frac{dx}{d_\tau^2(x)} \geq 4\pi$$

Proof Let $d_\tau^{max} = \max\{d_\tau(y),\; y \in \Sigma\}$. We have $\int_\Sigma \frac{dx}{d_\tau^2(x)} \geq \frac{\mathrm{Area}(\Sigma)}{\left(d_\tau^{max}\right)^2}$. Since Σ is connected and has no boundary, it bounds an open set Ω of $\mathbb{R}^3$. Let x be a point of Σ such that $d_\tau(x) = d_\tau^{max}$. Call B_x the maximal inner (*ie* that is included in Ω) ball that is tangent to Σ at x. Its center is on τ, thus its radius is greater than $d_\tau(x) = d_\tau^{max}$. Since B_x is a ball included in Ω, we have $\mathrm{Area}(\partial B_x) \leq \mathrm{Area}(\Sigma)$, which gives $4\pi\,(d_\tau^{max})^2 \leq \mathrm{Area}(\Sigma)$. Thus, $\int_\Sigma \frac{dx}{d_\tau^2(x)} \geq 4\pi$. □

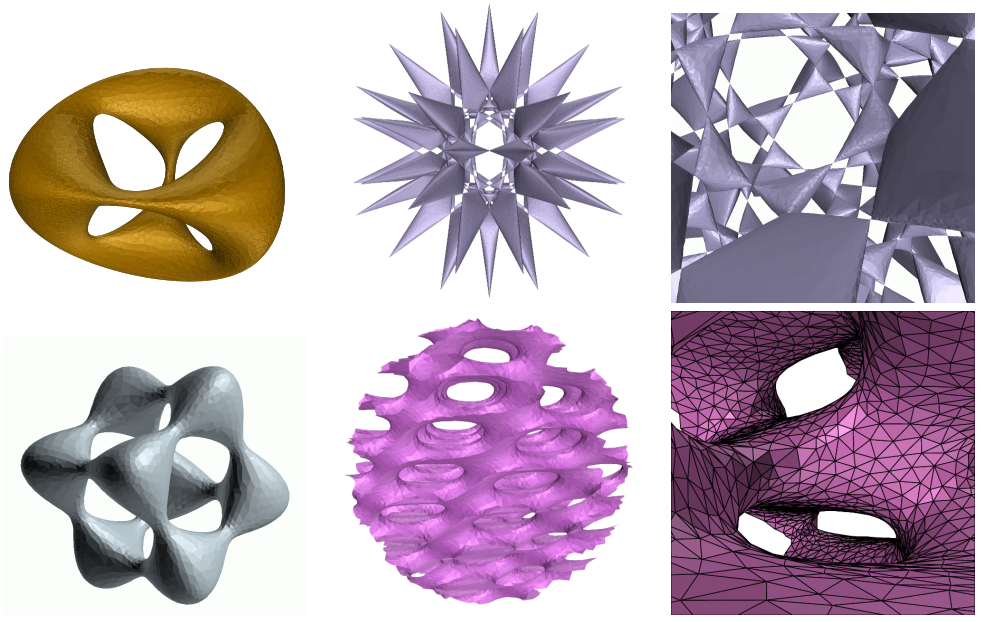

Figure 5: *Several implicit surfaces meshed by the algorithm.*

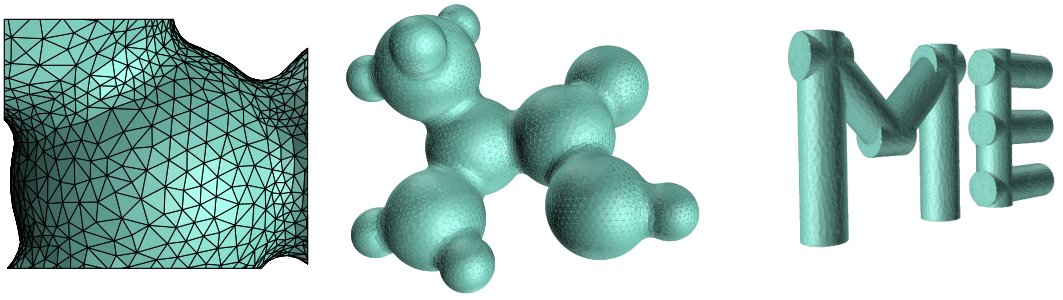

Figure 6: *Examples from biogeometry (left and center) and CAD (right).*

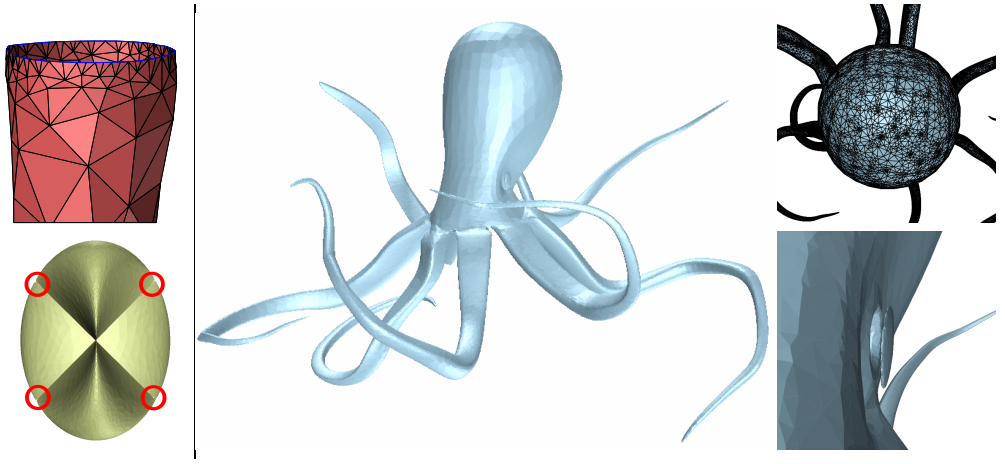

Figure 7: *Dealing with boundaries and singularities.*