Sampling and Meshing Surfaces with Guarantees

Steve Y. Oudot

École Doctorale de l'École Polytechnique

THESE

présentée et soutenue par

Steve OUDOT

pour obtenir le titre de

DOCTEUR DE L'ÉCOLE POLYTECHNIQUE

Discipline : Informatique

ÉCHANTILLONNAGE ET MAILLAGE DE SURFACES AVEC GARANTIES

Thèse dirigée par Jean-Daniel Boissonnat

soutenue le 14 décembre 2005

Composition du jury :

M. Olivier Faugeras	PRÉSIDENT
M. Frédéric Chazal	RAPPORTEUR
M. Gert Vegter	RAPPORTEUR
M. Jean-Daniel Boissonnat	DIRECTEUR
M. Leonidas Guibas	EXAMINATEUR
M. Gilles Dowek	EXAMINATEUR

Remerciements

J'aimerais remercier en tout premier lieu Jean-Daniel Boissonnat, qui a dirigé ma thèse de manière à la fois fructueuse et conviviale tout au long de ces trois années. Le temps passé sous sa direction a été pour moi un véritable enrichissement intellectuel et personnel.

Je tiens également à remercier particulièrement Frédéric Chazal, dont les nombreuses remarques et suggestions ont contribué à améliorer significativement la qualité scientifique du mémoire, en particulier les chapitres 2 et 3. Les échanges que nous avons eus m'ont permis de mieux mettre en perspective mon travail, et d'établir des connexions avec d'autres branches des mathématiques.

Les chapitres 8 et 9 ne seraient sûrement pas ce qu'ils sont sans le concours de mes autres co-auteurs : Leo Guibas pour le probing, Mariette Yvinec et son étudiant Laurent Rineau pour le maillage de volumes à bords lisses. Qu'ils en soient remerciés.

Plus généralement, j'aimerais remercier les membres de l'équipe Géométrica pour les nombreuses discussions scientifiques et techniques que nous avons eues et pour la qualité constante de leurs remarques. Merci donc à Pierre Alliez, Frédéric Cazals, Raphaëlle Chaine, David Cohen-Steiner, Olivier Devillers, Sylvain Pion et Monique Teillaud.

J'aimerais aussi remercier personnellement les doctorants de l'équipe pour leur sérieux et leur entrain, ainsi que pour tous les bons moments passés ensemble. Camille, Christophe, Laurent, Marc, Marie, Thomas, un grand merci à vous tous. Enfin, *last but not least*, merci à Agnès pour son amitié et son aide précieuse.

Sur un plan plus personnel, je tiens à remercier Béatrice, ainsi que ma famille et mes proches, pour leur soutien et leur dévouement constants. En particulier, je leur sais infiniment gré d'avoir cru en mon travail aux moments où moi-même n'y croyais plus...

Contents

Introduction

In this thesis, we focus on one aspect of the problem of representing a continuous shape by a finite number of parameters. This issue finds applications in many areas of Science and engineering, where the goal is to perform computations and simulations on objects from the real world. Due to the finite amount of memory available on a computer, only discretized versions of these objects can be represented informatically.

It turns out that the choice of a specific representation depends highly on the application considered. Let us illustrate this claim with an example from the medical world. Recent technology allows to represent the surface of any organ in a human body through a series of parallel slices, each slice containing the image of the contours of the organ in a certain plane. This kind of model is provided for instance by an MRI scanner. The problem with such a representation is that it is difficult to exploit, in contexts such as tumor or lesion tracking. Consequently, there is a real need for new kinds of representations. A whole area of research in image processing is devoted to this issue. Some solutions consist in converting the slice-to-slice model into a whole 3-dimensional image. Although this new representation is well-adapted to the tracking problem, it remains unsuitable for further processing. Moreover, the size of the data is a problem, since images are usually large, with a high definition.

Some representations are generically more effective than others. It is the case for simplicial meshes. Roughly speaking, a simplicial mesh is a collection of simplices of pairwise disjoint relative interiors, such that two simplices of the mesh either do not intersect, or intersect along a simplex of lower dimension. We recall that a simplex of dimension d (or d-simplex, for short) is the affine hull of $(d+1)$ points. For instance, a 0-simplex is a point, a 1-simplex is a segment, a 2-simplex a triangle, a 3-simplex a tetrahedron, and so on. In this thesis, we focus mainly on simplicial meshes that approximate surfaces, a case in which all simplices have dimension at most 2. The mesh is then called a triangular mesh. Similarly, if all simplices have dimension at most 3, then the mesh is said to be a tetrahedral mesh.

Simplicial meshes are one of the most popular representations for surfaces, volumes, scalar fields and vector fields. Their success finds its origin in the fact that simplicial meshes are very well suited for many applications, such as visualization or numerical simulation. However, constructing a simplicial mesh that approximates a continuous object can be time-consuming, especially when the geometry of the object is complex. In this case, mesh generation becomes the pacing phase in the computational simulation cycle. Roughly speaking, the more the user is involved in the mesh generation process, the longer the latter is. An appealing example is given in [88], where the mesh generation time is shown to be 45 times that required to compute the solution. This motivates the search for fully-automated mesh-generation techniques.

This thesis addresses mainly the problem of constructing a triangular mesh to approximate a given surface. This problem is stated below in an informal manner. Some precisions will be given thereafter:

Surface mesh generation *Given a surface S without boundary, construct a triangular mesh $\hat{S}$ of optimal size, that is both topologically equivalent to S and geometrically close to S.*

The main concern here is that $\hat{S}$ be of same topological type as the input surface S. Since the surface S is assumed to have no boundary, this means that $\hat{S}$ and S must have the same number of connected components and the same number of handles.

The surface mesh generation paradigm states also that $\hat{S}$ must be geometrically close to the input surface S. This requires to define an accuracy measure. The Hausdorff distance is a good candidate and is therefore usually chosen. Recall that the Hausdorff distance between $\hat{S}$ and S is ε if every point of $\hat{S}$ is at distance at most ε from S, and every point of S is closer than ε to $\hat{S}$.

Concerning the size of the mesh, Agarwal and Suri [1] proved that, given a surface S, a threshold ε and a positive number k, it is NP-hard to decide whether there exists a mesh of size at most k that lies at Hausdorff distance at most ε from S. This means that it is hopeless to find effective methods to construct size-optimal meshes to approximate surfaces. Therefore, instead of constructing meshes of optimal size, our goal is rather to construct meshes that are *as small as possible*, typically of size within a constant factor of the optimal.

It is a well-known fact that, once the topology of $\hat{S}$ is fixed, the numbers of faces of each dimension in the triangular mesh are fully determined by the number of its vertices. Therefore, bounding the size of $\hat{S}$ comes down to bounding the number of its vertices. The mesh generation process reduces then to two steps:

S1 Construct a finite set of points E sampled from the surface S.

S2 Connect the points of E with triangles, so that the underlying triangulation $\hat{S}$ has the same topological type as S.

The quality of a surface meshing algorithm is clearly measured by the quality of its output. The better the approximation, the better the algorithm. But another criterion has to be taken into account: the amount of information on the input surface S needed by the algorithm. Meshing strategies that use few geometric predicates will be easier to implement, and their algorithmic complexity will be less dependent on the topological and geometric complexity of the surface S. Moreover, the nature of the information available on S varies with the context. Geometric queries are more or less difficult, depending on the way the surface S is defined. For instance, computing the principal curvatures of S at a given point is easy when S is defined as a level-set of some function f whose closed formula is known. However, this task becomes very difficult when no closed formula of f is known, as it is the case for instance when S is defined as a level of grey in a 3-dimensional image or as the result of a PDE. In this respect, algorithms that use fewer predicates are not only easier to implement, but they turn out to be also more generic since their predicates can be implemented in a wider range of applications.

We distinguish between three variants of the surface meshing problem, depending on the way the surface S is defined and on the amount (and nature) of information available as input. In fact, there are

many more variants, but to our knowledge, each of them can be reduced to one of the following: surface reconstruction, surface sampling, and surface probing.

I.1 Surface reconstruction: working out sampling conditions

In surface reconstruction, the surface S is known only through a finite set of points E sampled from S. This point set can come from various sources, such as for instance from range scanning data. The point sample E is given as input, and the goal is simply to build a connectivity between the points of E, so that the resulting mesh $\hat{S}$ approximates S topologically and geometrically. This reduces the problem to the step S2 described above.

Several provably good methods have been proposed to solve the smooth surface reconstruction problem. We refer the reader to [30, 31] for a comprehensive survey. The proofs of correctness of these methods rely on sampling conditions that control the local density of the input point set E.

Closed Ball Property The first sampling condition was introduced by Edelsbrunner and Shah [63] and is called the *Closed Ball Property*. It states that every d-face of the Voronoi diagram of E has an intersection with S that is either empty or topologically equivalent to a d-ball. Edelsbrunner and Shah proved that, under this condition, some subcomplex of the Delaunay triangulation of E, called $\mathrm{Del}_{|S}(E)$, has the same topological type as S. $\mathrm{Del}_{|S}(E)$ will be described more in detail in the next chapter. Unfortunately, the Closed Ball Property guarantees only the topology of the reconstruction, not its geometry.

μ-samples Amenta and Bern [6] proposed then the *μ-sampling* condition. Roughly speaking, a finite subset E of a surface S is a μ-sample of S if every point p of S is closer to E than $\mu\ \mathrm{d}_{\mathrm{M}}(p)$, where $\mathrm{d}_{\mathrm{M}} : \mathbb{R}^3 \to \mathbb{R}$ denotes the distance to the so-called medial axis M of S, defined in Chapter 1. This condition requires that the surface S be smooth.

Amenta and Bern proved that, if E is a μ-sample of S, for a sufficiently small μ, then the Closed Ball Property of Edelbrunner and Shah is satisfied. Hence, $\mathrm{Del}_{|S}(E)$ has the same topological type as S. They also proved that $\mathrm{Del}_{|S}(E)$ lies at Hausdorff distance $c(S)\ \mu^2$ from S, where the constant $c(S)$ depends only on S. In addition, $\mathrm{Del}_{|S}(E)$ provides good approximations of the normals of S [6], of its area [92], and of its curvatures [49]. Therefore, if the input point set E is a μ-sample of S, then the problem of reconstructing S from E comes down to finding which faces of the Delaunay triangulation of E belong to $\mathrm{Del}_{|S}(E)$. A number of provably good reconstruction algorithms are based on this approach [6, 7, 8, 19, 53, 54, 55].

Several variants of the concept of μ-sample appeared since [6], in particular the so-called uniform ε-samples [10], whose density is specified by a uniform sizing field equal to a constant ε. In order to use a single concept, we introduce the following

Definition I.1 *Given a surface S and a positive function $\sigma : S \to \mathbb{R}$, a finite point set $E \subset S$ is a* σ-sample *of S if $\forall p \in S,\ E \cap B(p, \sigma(p)) \neq \emptyset$.*

From now on, a μ-sample in the sense of [6] will be called a $\mu\mathrm{d}_\mathrm{M}$-sample, whereas a uniform ε-sample in the sense of [10] will be referred to as an ε-sample, according to Definition I.1. Note that both concepts are closely related. Indeed, any ε-sample of S is a $\mu\mathrm{d}_\mathrm{M}$-sample for some μ depending on ε, and reciprocally, any $\mu\mathrm{d}_\mathrm{M}$-sample of S is an ε-sample for some ε depending on μ. Therefore, the theoretical guarantees offered by $\mu\mathrm{d}_\mathrm{M}$-samples hold for ε-samples as well, provided that ε is small enough. Specifically, ε must be less than a fraction of the *reach* of S, which is the infimum of d_M over S. This infimum is positive when d_M is positive. The surface S is then said to have a positive reach.

The noisy case The ε-sampling condition, as defined above, assumes that the points of E belong to S. In practice, due to the fact that scanners can only measure points within a given precision, the points of E usually do not lie exactly on the surface S.

Dey and Goswami [56] proposed an extension of the cocone algorithm [55] that solves this variant of the reconstruction problem, provided that the input point set E is a so-called *noisy* $\mu\mathrm{d}_\mathrm{M}$*-sample* of S, which is equivalent to an ε-sample, with the difference that the points may not lie on S but must be closer to S than a fraction of their distance to the medial axis of S.

The nonsmooth case The μ-sampling condition assumes also that the surface is smooth. The distance to the medial axis of a nonsmooth surface S vanishes at the points where S is not differentiable. As a consequence, the reach of S is zero, and no finite subset of S is a $\mu\mathrm{d}_\mathrm{M}$-sample of S. This makes μ-samples useless in this context. Nevertheless, ε-samples are still well-defined if ε is positive, but the theoretical guarantees of [6] no longer apply since $\varepsilon > 0$ is not smaller than a fraction of the reach of S.

Chazal and Lieutier [35] introduced the so-called *weak feature size*, or wfs for short. Let $\mathrm{d}_S : \mathbb{R}^3 \to \mathbb{R}^+$ map each point of $\mathbb{R}^3$ to its distance to S. The weak feature size is simply the smallest critical value of d_S, in the sense of Riemannian geometry [69, 74].

The advantage of the weak feature size, over the distance to the medial axis, is that the class of objects with positive wfs is much larger than the class of objects with positive reach. This allows to construct ε-samples, with $\varepsilon \leq$ wfs, on a wide variety of smooth and nonsmooth shapes.

Chazal and Lieutier did not exhibit any subcomplex of $\mathrm{Del}_{|S}(E)$ with the same topological type as S. However, they showed that the homology groups of the 3-dimensional object $\mathcal{O}$ bounded by S can be retrieved from any noisy wfs-sample[1] E of S. The approach consists in applying persistence techniques [62] on a specific filtration of $\mathrm{Del}(E)$, to compute the homology groups of $\mathcal{O}$. Furthermore, an approximation of the medial axis of the object $\mathcal{O}$ can be computed from the Voronoi diagram of E.

The approach of Chazal and Lieutier is very generic and works in the nonsmooth and noisy setting. However, it provides information exclusively on the topology of the object, not on its geometry.

Other approaches To complete our overview of previous work on surface reconstruction, let us emphasize that many other methods have been proposed to solve the reconstruction problem. Some of them interpolate the point set E, others define surfaces that approximate E. However, almost all these methods come with no theoretical guarantees regarding the topology of the approximation. Two noticeable exceptions are the approach of [79], based on Moving Least Squares techniques [83], and the method

[1] In the same sense as Dey and Goswami, except that d_M is replaced by wfs.

of [19], based on natural neighbor coordinates [18]. Both methods can guarantee the topology of their approximating surface, under ε-sampling conditions.

I.2 Surface sampling: satisfying the sampling conditions

Here, the surface S is given in a form that allows to sample new points from S. For instance, S can be defined implicitely, as a level set of some real-valued function defined over $\mathbb{R}^3$.

This issue is somehow dual to the surface reconstruction problem. Indeed, in surface reconstruction, the point set E is given as input, and the goal is to find sampling conditions that guarantee a correct reconstruction. In surface sampling, the goal is to build a point set E that satisfies some sampling condition, so that correct reconstruction is then ensured. In other words, surface sampling focusses on step S1 of the mesh generation process, whereas surface reconstruction focusses on step S2.

The main drawback of the ε-sampling condition is that, given a surface S and a positive value ε, it is difficult to check whether a sample is an ε-sample of S, and even more difficult to construct an ε-sample of S. This is due to the fact that a direct application of the definition of an ε-sample leads to complicated operations like cutting the surface S with balls. This is why previous work on surface sampling does not rely on this condition.

It is only recently that provably good sampling and meshing techniques appeared in the literature. Some of these techniques handle only restricted types of shapes, such as piecewise parametric CAD models [104, 112], Van der Waals and solvent-excluded molecular surfaces [81], or skin surfaces [38, 39, 80]. We do not consider them in the sequel.

The other provably good algorithms can only handle smooth surfaces. Moreover, to our knowledge, all these algorithms require that the surface S be defined implicitely, as a level-set of some function f of known closed formula. Although it is well-known that every surface S without boundary is the zero-set of some real-valued function defined over $\mathbb{R}^3$ (take for instance the signed distance to S), the assumption that a closed formula of f is known is very restrictive in practice, since in many applications (for instance sampling an isosurface in a 3-dimensional image) no closed formula of f is known *a priori*. We classify the existing algorithms with respect to the amount of information on S they require:

1. The implicit surface mesher of Plantinga and Vegter [96] generates an adaptive grid and then applies the Marching Cubes algorithm [87]. Using interval arithmetics, Plantinga and Vegter can certify the topology of $\hat{S}$. Moreover, by refining the grid sufficiently, they can achieve any given bound on the Hausdorff distance between $\hat{S}$ and S. This is a significant step since the Marching Cubes algorithm and its variants [44] usually come without any topological or geometric guarantees. However, the use of interval arithmetics requires to now the gradient of f.
2. Algorithms based on the Closed Ball Property of Edelsbrunner and Shah [63], like the implicit surface mesher of Cheng *et al.* [42], require to be able to compute the critical points of height functions on the restrictions of S to some hyperplanes. The topology of the output mesh is ensured thanks to the Closed Ball Property.
3. Methods based on critical points theory [20, 76] require to compute the critical points of f, as well as their indices, which is an even more involved computation, although it can be performed using interval arithmetics.

The above techniques work only in the smooth setting. Recently, Dey *et al.* adapted the method of [42] to the case where S is a polyhedron that approximates a smooth surface [57]. This assumption on S allows the authors to use the same mathematical tools as in the smooth setting, but it is quite restrictive for practical applications. To our knowledge, no provably good sampling algorithm has ever been proposed for the nonsmooth non-polyhedral case.

I.3 Surface probing: reducing the required knowledge of the surface

Surface probing, also known as blind surface approximation or interactive surface reconstruction, consists of discovering the shape of an unknown object $\mathcal{O}$ through an adaptive process of probing its surface S from the exterior. A probe is issued along a ray whose origin lies outside $\mathcal{O}$ and returns the first point of $\mathcal{O}$ hit by the ray. Successive probes may require the probing device to be moved through the free space outside $\mathcal{O}$. The goal is to find a strategy for the sequence of probes that guarantees a precise approximation of S after a minimal number of probes.

This problem belongs to the class of geometric probing problems, pioneered by Cole and Yap [50]. Geometric probing is motivated mainly by applications in robotics. In this context, our probe model described above is called a tactile or finger probe. Geometric probing also finds applications in other areas and gave rise to several variants. In particular, other probe models have been studied in the literature, *e.g.* line probes (a line moving perpendicular to a direction), X-ray probes (measuring the length of intersection between a line and the object), as well as their counterparts in higher dimensions.

The existing probing algorithms can be classified into two main categories, exact or approximate, depending on whether they return the exact shape of the probed object or an approximation. An exact probing algorithm can only be applied to shapes that can be described by a finite number of parameters like polygons and polyhedra. Therefore, exact probing is too restrictive for most practical applications. Approximate probing algorithms overcome this deficiency by considering the accuracy of the desired reconstruction as a parameter. The goal is to find a strategy that can discover a guaranteed approximation of the object using a minimal number of probes. This issue is closely related to surface sampling. However, it differs in an essential way: here, the surface S is known through an oracle, the probing device, which can answer only very specific geometric questions.

Most of the work on exact geometric probing is for convex polygons and polyhedra – see [106] for a survey of the computational literature on the subject. Nevertheless, it has been shown that, using enhanced finger probes, a large class of non convex polyhedra can be exactly determined [2, 24]. Concerning approximate probing, an important class is the class of convex shapes. Probing strategies have been proposed for planar convex objects using line probes [86, 97] and some other probe models are analyzed in [98]. However, as far as we know, designing provably good techniques to probe non convex non polyhedral objects has not been considered prior to this thesis.

I.4 Overview of the thesis and contributions

In this thesis, the surface S can be smooth or nonsmooth, but the devices used for taking measures on S provide exact information. As a consequence, the points of E belong to S.

Sampling conditions In Part A, we introduce the concept of loose ε-sample, which can be viewed as a weak version of the notion of ε-sample. The main advantage of loose ε-samples over ε-samples is that they are easier to check and to construct. Indeed, checking that a sample is a loose ε-sample reduces to checking whether a finite number of spheres have radii at most ε.

When the surface S has a positive reach, we prove that, for sufficiently small ε, ε-samples are loose ε-samples, and reciprocally. As a consequence, loose ε-samples offer the same topological and geometric guarantees as ε-samples.

We also focus on the case where S is nonsmooth. In order to make geometric claims, we restrict our study to a subclass of the objects with positive weak feature size: the so-called k-Lipschitz surfaces, defined in Chapter 2. This subclass contains all objects with positive reach, but also all sufficiently *smooth* polyhedra (see Definition 2.7), as well as a number of other nonsmooth shapes with non-trivial topology. We prove that, if S is a k-Lipschitz surface and E is a (loose) ε-sample of S, for sufficiently small k and ε, then $\mathrm{Del}_{|S}(E)$ has the same topological type as S and is close to S for the Hausdorff distance. Our theoretical results hold provided that the inner angles of the facets of the Delaunay triangulation of E are not too small, which is ensured by assuming that the points of E are farther than a fraction of ε from one another. This sparseness condition is known to be a bit restrictive in the context of surface reconstruction. However, it can be easily satisfied in the contexts of surface sampling and surface probing.

A simple surface sampling algorithm In Part B, we show how our sampling condition can be turned into a simple and efficient surface mesh generator. The latter is based on a *Delaunay refinement* technique, which consists in constructing an initial mesh and then refining iteratively the elements of the mesh that do not meet some user-defined size or shape criteria. This greedy tehnique was pioneered by Ruppert [99] in the plane, and then extended by Chew to surfaces in 3-space [45]. Our mesher derives from Chew's algorithm. It takes as input a user-defined parameter ε and an initial point set $E_I \subset S$, and it outputs a loose ε-sample E_F of S, together with $\mathrm{Del}_{|S}(E_F)$.

Chew did not provide his algorithm with any topological guarantees. Here, taking advantage of the theoretical results of Part A, we can prove that the output mesh is a manifold without boundary, with the same topological type as S and close to S for the Hausdorff distance, provided that ε is chosen sufficiently small. Specifically, ε must be less than a fraction of the reach of S when the latter is smooth, and less than a fraction of the so-called *k-Lipschitz radius of S* (defined in Chapter 2) when S is k-Lipschitz, for sufficiently small k. It follows that the algorithm generates provably good meshes on a wide class of smooth and nonsmooth shapes. Moreover, we show that the number of points sampled from S by the algorithm lies within a constant factor of the optimal.

Let us emphasize that our mesh generator maintains the 3-dimensional Delaunay triangulation of the point sample throughout the process. This is mandatory for guaranteeing the topology of the output mesh. Maintaining a whole 3-dimensional triangulation can be time-consuming in general. However, we prove in our case that the size of the data structure remains bounded, which implies that the space and time complexities of the algorithm are quite reasonable. Further detail is provided in Chapter 4.

A unified solution to the meshing problem A noticeable feature of the algorithm is that it needs only to know the surface S through an oracle that can compute the intersection of any given segment with S. Therefore, our mesh generator is generic enough to be applied in a wide variety of contexts. This genericity is illustrated in Part C, where we show that the algorithm can be used to mesh implicit surfaces, remesh polyhedra, or reconstruct surfaces from scattered data points. Our approach to surface reconstruction is similar to that of [19]. It consists of defining an implicit function f' from the input point set, and then to mesh the zero-set S' of f' using our mesh generator. Combining our theoretical results with those of [19], we can certify the topology and geometry of the output mesh. Note that no closed formula of f' is available, hence none of the implicit surface sampling techniques presented in Section I.2 can be applied in this context.

We also show that the algorithm can be easily adapted to probe unknown objects. The approach consists of using the probing device as an oracle for our mesh generator. This oracle is weaker than the one used above, since it can detect only the first intersection point of a given segment with the surface S. Moreover, before checking the intersection of a given segment s with S, the probing device must first be moved to an endpoint of s. Therefore, we cannot check the intersections of all the segments with S. We prove however that this version of the algorithm comes with the same theoretical guarantees as the original version, regarding the quality and the size of the output.

Finally, we show that our meshing technique can be extended to construct tetrahedral meshes approximating 3-dimensional objects with curved boundaries, such that the mesh elements (tetrahedra and triangles) conform to some user-defined size and shape criteria. The idea is to exploit the fact that the algorithm maintains a whole 3-dimensional Delaunay triangulation. Whenever a tetrahedron does not meet the size or shape requirements, it is refined by inserting its circumcenter. The output point set is no longer a subset of S. Moreover, the output mesh contains all Delaunay tetrahedra whose circumcenters lie in the object $\mathcal{O}$ to mesh. Using our theoretical results on the approximation of the boundary of $\mathcal{O}$, we can certify the output of the algorithm.

Our contributions at a glance We list our main contributions below:

- A new sampling condition, called *loose ε-sampling*, which offers the same topological and geometric guarantees as the classical ε-sampling condition in the smooth setting, but which is much easier to check and to construct.

- A theoretical analysis in the nonsmooth Lipschitz setting, which proves that (loose) ε-samples offer the same guarantees in this context as in the smooth setting, provided that an additional sparseness condition is satisfied. As a consequence, ε-sampling and loose ε-sampling are the first sampling conditions to provide both topological and geometric guarantees in the nonsmooth setting.

- A simple meshing algorithm that can provably well approximate smooth or Lipschitz surfaces, with a number of points that lies within a constant factor of the optimal. This algorithm is one of the few certified methods for the smooth and the polyhedral cases, and the first certified method for the nonsmooth non-polyhedral case. We have implemented it in a number of practical sit-

uations: isosurface extraction from 3D medical images, polygonal surface remeshing, etc. Our experimental results provide evidence that the approach is very effective in practice.

- An easy (yet certified) adaptation of our algorithm that solves the surface probing problem in the case of a convex or non-convex object with curved boundaries. No certified solution was known for this problem prior to this work.

- A natural extension of our algorithm that constructs provably good tetrahedral meshes to approximate 3-dimensional objects with curved boundaries. It is the first certified solution ever proposed for this problem. Moreover, our preliminary experimental results are quite promising, regarding the practicality of the approach.

Preliminaries

This chapter introduces most of the mathematical concepts that will be used in the thesis.

Basic notations

Throughout the thesis, $\mathbb{R}^3$ is the ambient affine space. Points are written in italic font, *e.g.* p, q, and vectors in bold font, *e.g.* $\mathbf{n}, \mathbf{n}'$.

$\mathbb{R}^3$ is endowed with a canonical frame $(o, \mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z)$. For any point p, we call $\mathbf{p}$ the vector that goes from the origin o of the frame to p. The usual inner product of two vectors $\mathbf{n}, \mathbf{n}'$ is denoted $\mathbf{n} \cdot \mathbf{n}'$, and defined by

$$\mathbf{n} \cdot \mathbf{n}' = \mathbf{n}_x \mathbf{n}'_x + \mathbf{n}_y \mathbf{n}'_y + \mathbf{n}_z \mathbf{n}'_z,$$

where $\mathbf{n}_x, \mathbf{n}_y, \mathbf{n}_z$ and $\mathbf{n}'_x, \mathbf{n}'_y, \mathbf{n}'_z$ are the coordinates of $\mathbf{n}$ and $\mathbf{n}'$ in the canonical basis $(\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z)$. We abuse notations and write $\mathbf{n}^2$ the inner product of $\mathbf{n}$ with itself. In addition, we denote $\mathbf{n} \times \mathbf{n}'$ the outer product of $\mathbf{n}, \mathbf{n}'$, defined by

$$\mathbf{n} \times \mathbf{n}' = \begin{vmatrix} \mathbf{n}_y & \mathbf{n}'_y \\ \mathbf{n}_z & \mathbf{n}'_z \end{vmatrix} \mathbf{e}_x + \begin{vmatrix} \mathbf{n}_z & \mathbf{n}'_z \\ \mathbf{n}_x & \mathbf{n}'_x \end{vmatrix} \mathbf{e}_y + \begin{vmatrix} \mathbf{n}_x & \mathbf{n}'_x \\ \mathbf{n}_y & \mathbf{n}'_y \end{vmatrix} \mathbf{e}_z$$

Finally, we denote by $(\mathbf{n}, \mathbf{n}')$ the modulus of the angle (measured in $[-\pi, \pi]$) between vectors $\mathbf{u}$ and $\mathbf{v}$.

Distances and topology

We call $\|.\|$ the Euclidean norm:

$$\forall \mathbf{n} \in \mathbb{R}^3, \qquad \|\mathbf{n}\| = \sqrt{\mathbf{n}^2} = \sqrt{\mathbf{n} \cdot \mathbf{n}}$$

For any points p, q, $\mathrm{d}(p, q)$ denotes the Euclidean distance from p to q, defined by

$$\mathrm{d}(p, q) = \|\mathbf{p} - \mathbf{q}\|$$

Given a point c and a positive constant r, $B(c, r)$ denotes the ball of center c and radius r. Unless otherwise specified, the ball $B(c, r)$ is closed.

The affine space $\mathbb{R}^3$ is canonically endowed with the topology generated by the set of open balls. Given a subset X of $\mathbb{R}^3$, $\mathrm{int}(X)$, $\bar{X}$, and ∂X denote respectively the interior, the closure, and the boundary of X. Moreover, we call $\mathrm{diam}(X)$ the Euclidean diameter of X:

$$\mathrm{diam}(X) = \sup\{\mathrm{d}(p, q),\ p \in X,\ q \in X\}$$

We now introduce two topological concepts that will be widely used in the thesis. We refer the reader to [14] for further detail.

Definition P.2 *Two topological spaces X, Y are* homeomorphic *if there exists a continuous and bijective map $h : X \to Y$ such that h^{-1} is continuous. The map h is called a* homeomorphism *from X to Y.*

If X, Y are subsets of $\mathbb{R}^3$, and if h is a homeomorphism from $\mathbb{R}^3$ to $\mathbb{R}^3$, such that $h(X) = Y$, then h is said to be an *ambient homeomorphism* from X to Y.

Definition P.3 *Two topological spaces X, Y embedded in $\mathbb{R}^3$ are* isotopic *if there exists a continuous map $i : [0,1] \times X \to \mathbb{R}^3$ such that $i(0,.)$ is the identity over X, $i(1, X) = Y$, and for any $t \in [0,1]$, $i(t,.)$ is a homeomorphism from X onto its image. The map i is called an* isotopy *from X to Y.*

If i is an isotopy from $\mathbb{R}^3$ to $\mathbb{R}^3$ such that $i(1, X) = Y$, then i is said to be an *ambient isotopy* from X to Y.

We need also to define distances between subsets of $\mathbb{R}^3$. Given a point p and a subset X of $\mathbb{R}^3$, the distance from p to X is denoted $\mathrm{d}(p, X)$ and defined as follows:

$$\mathrm{d}(p, X) = \inf\{\mathrm{d}(p, q),\ q \in X\}$$

If $X = \emptyset$, then $\mathrm{d}(p, X)$ is infinite. Otherwise, $\mathrm{d}(p, X)$ is finite and non-negative. Moreover, $\mathrm{d}(p, X) = 0$ iff (*if and only if*) p belongs to the closure $\bar{X}$ of X.

Definition P.4 *Given two subsets X and Y of $\mathbb{R}^3$, the Hausdorff distance between X and Y is*

$$\mathrm{d}_{\mathcal{H}}(X, Y) \ = \ \max\{\sup_{p \in X} \mathrm{d}(p, Y),\ \sup_{q \in Y} \mathrm{d}(q, X)\}$$

Note that $\mathrm{d}_{\mathcal{H}}(X, Y) = 0$ iff $\bar{X} = \bar{Y}$.

Definition P.5 *Given two homeomorphic subsets X and Y of $\mathbb{R}^3$, the Fréchet distance between X and Y is*

$$\mathrm{d}_{\mathcal{F}}(X, Y) \ = \ \inf_{h} \sup_{p \in X} \mathrm{d}(p, h(p))$$

where h ranges over all homeomorphisms from X to Y.

Observe that, for any homeomorphism $h : \ X \to Y$ and any point $p \in X$, we have

$$\mathrm{d}(p, h(p)) \geq \max\{\mathrm{d}(p, Y),\ \mathrm{d}(h(p), X)\},$$

thus $\mathrm{d}_{\mathcal{F}}(X, Y) \geq \mathrm{d}_{\mathcal{H}}(X, Y)$. It follows that any upper bound on the Fréchet distance of X, Y is also an upper bound on their Hausdorff distance.

Geometric measures

We will use two geometric measures: the Lebesgue measure $\mathcal{L}^2$ on $\mathbb{R}^2$, and the 2-dimensional Hausdorff measure $\mathcal{H}^2$ on $\mathbb{R}^2$ and $\mathbb{R}^3$. We borrow their definitions from [91].

Definition P.6 *There is a unique Borel-regular, translation-invariant measure on $\mathbb{R}^2$, such that the measure of the unit cube $[0,1]^2$ is 1. This measure is called the* Lebesgue measure *and noted $\mathcal{L}^2$.*

Note that $\mathcal{L}^2$ is defined only on $\mathbb{R}^2$. Hence, it tells nothing about the area of a surface of $\mathbb{R}^3$, a notion that will be extensively used in this thesis, especially in Chapter 3.

Definition P.7 *Given $X \subset \mathbb{R}^n$ ($n \geq 2$), the 2-dimensional Hausdorff measure of X, or $\mathcal{H}^2(X)$ for short, is defined as follows:*

$$\mathcal{H}^2(X) = \lim_{\delta \to 0^+} \inf_{\substack{X \subset \bigcup_i X_i \\ \operatorname{diam}(X_i) \leq \delta}} \sum_{i=0}^{+\infty} \pi \left(\frac{\operatorname{diam}(X_i)}{2} \right)^2 \tag{1}$$

The infimum in (1) is taken over all countable coverings $\{X_i\}$ of X whose members have Euclidean diameter at most δ. It is proved in [91, §2.3] that the limit always exists, and that $\mathcal{H}^2$ is a Borel-regular measure over $\mathbb{R}^n$, for any $n \geq 2$. A more comprehensive proof can be found in [67, §2.10].

The 2-dimensional Hausdorff measure coincides with $\mathcal{L}^2$ on $\mathbb{R}^2$. Moreover, it generalizes the notion of area of a surface. Specifically, $\mathcal{H}^2(S)$ coincides with the usual area of S, for any C^1-continuous 2-dimensional submanifold of $\mathbb{R}^n$, $n \geq 2$.

Surfaces

Throughout the thesis, S is a C^0-continuous 2-dimensional submanifold of $\mathbb{R}^3$. Informally, this means that, for any point $p \in S$, there exists an open neighborhood $\mathcal{N}$ of p in $\mathbb{R}^3$ that can be mapped to the unit open ball B by some homeomorphism h, such that $h(p)$ is the origin o and $h(\mathcal{N} \cap S) = B \cap \mathbb{R}^2$. See Figure 1 for an illustration. We refer the reader to [14, §2.1.1] for a formal definition.

For simplicity, we say that S is a surface without boundary. We call M its medial axis, defined in Section 1.1.1. If S is differentiable at p, then $T(p)$ and $\mathbf{n}(p)$ denote respectively the tangent plane and the unit normal vector (pointing outwards) of S at p.

Complexity

Given two positive functions f, g defined over $\mathbb{R}^+$, we use the notation $f = O(g)$ to indicate that there exists a non-negative constant ν such that $\forall x \geq 0,\ f(x) \leq \nu\ g(x)$. Moreover, $f = \Omega(g)$ stands for $g = O(f)$, and $f = \Theta(g)$ means that we have both $f = O(g)$ and $f = \Omega(g)$.

Unless explicitely mentioned, the constant ν in the above statements is absolute. Some of our results use these notations in a context where the constant ν depends on the surface S. In such a case, we write $f = O_S(g)$, $f = \Omega_S(g)$, $f = \Theta_S(g)$, instead of $f = O(g)$, $f = \Omega(g)$, $f = \Theta(g)$.

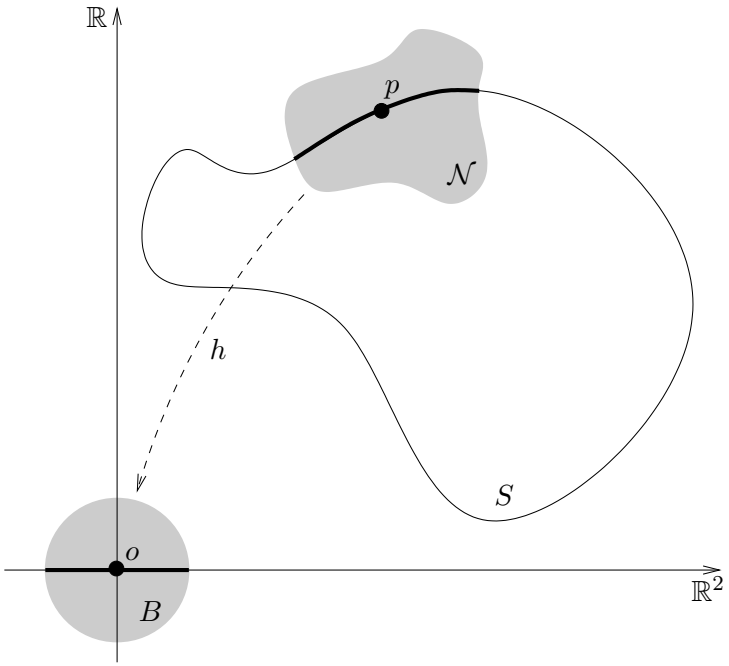

Figure 1: A surface without boundary.

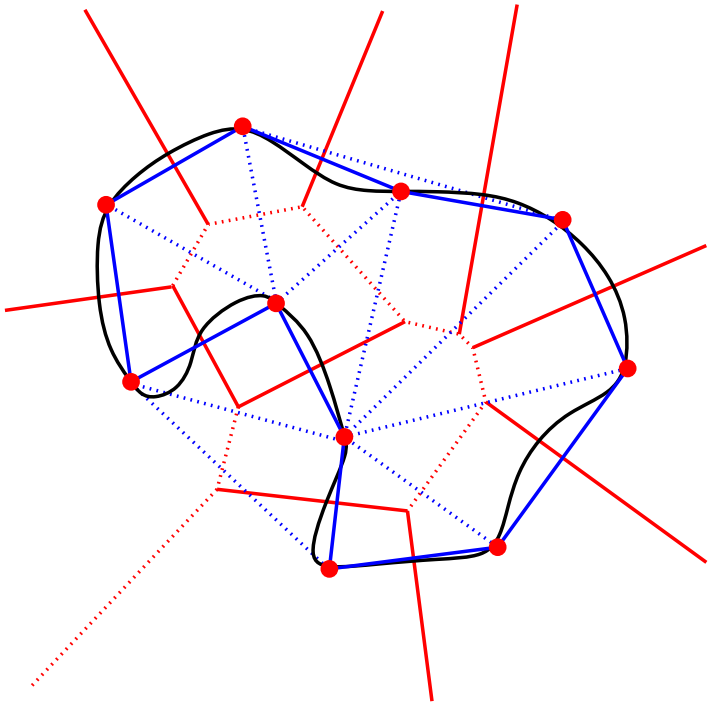

Figure 2: Example of a Delaunay triangulation (in blue) restricted to a curve (in black).

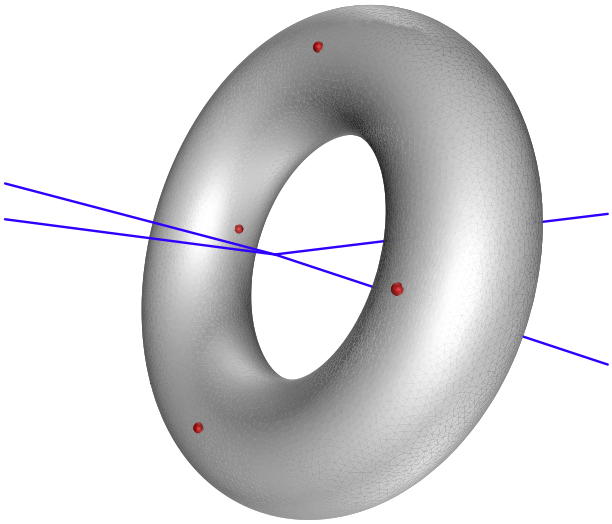

Figure 3: A case where $\mathrm{Del}_{|S}(E)$ is empty: the four points of E are placed on a torus, such that the Voronoi edges pass through the hole.

Voronoi diagrams and Delaunay triangulations

We now present several objects of the Computational Geometry that will be used throughout the thesis. Let E be a finite set of points of $\mathbb{R}^3$.

- The *Voronoi cell* of $p \in E$ is the set of all points of $\mathbb{R}^3$ that are closer to p than to any other point of E. The *Voronoi diagram* of E, $\mathrm{Vor}(E)$, is the cellular complex formed by the Voronoi cells of the points of E.

- The dual complex of $\mathrm{Vor}(E)$ is a tetrahedrization of the convex hull of E, called the *Delaunay triangulation* of E and noted $\mathrm{Del}(E)$. For any face f (vertex, edge, facet or cell) of $\mathrm{Del}(E)$, $\mathrm{V}(f)$ denotes the face of $\mathrm{Vor}(E)$ dual to f.

- By $\mathrm{VG}(E)$ we denote the 1-skeleton graph of $\mathrm{Vor}(E)$, also called *Voronoi graph of* E.

The mesh generator introduced in Part B of this thesis relies on a subcomplex of $\mathrm{Del}(E)$, defined below:

- We call Delaunay triangulation of E *restricted to* S, and we note $\mathrm{Del}_{|S}(E)$, the subcomplex of $\mathrm{Del}(E)$ that consists of the facets of $\mathrm{Del}(E)$ whose dual Voronoi edges intersect S. An example is given in Figure 2. An edge or vertex of $\mathrm{Del}(E)$ belongs to $\mathrm{Del}_{|S}(E)$ if it is incident to at least one facet of $\mathrm{Del}_{|S}(E)$. Notice that we depart from the usual definition [45, 63] and do not consider vertices and edges with no incident facet of $\mathrm{Del}_{|S}(E)$. See Figure 3 for an illustration.

- A facet (resp. edge, vertex) of $\mathrm{Del}_{|S}(E)$ is called a *restricted Delaunay facet* (resp. *restricted Delaunay edge*, *restricted Delaunay vertex*). For a restricted Delaunay facet f, we call *surface Delaunay ball* of f any ball circumscribing f centered at some point of $S \cap \mathrm{V}(f)$. Notice that the centers of the surface Delaunay balls are precisely the intersection points of S with $\mathrm{VG}(E)$.

- Given a vertex v of $\mathrm{Del}_{|S}(E)$, we call *star of* v the union of all the facets of $\mathrm{Del}_{|S}(E)$ incident to v.
- Given a facet f of $\mathrm{Del}_{|S}(E)$, we call *neighborhood of* f, of $\mathrm{N}(f)$ for short, the union of all the facets of $\mathrm{Del}_{|S}(E)$ that are non-disjoint with f (including f itself).

In the sequel, u, v, w denote vertices of $\mathrm{Del}_{|S}(E)$, whereas p, q denote points of $\mathbb{R}^3$.

Quality measures for mesh elements

In this thesis, our aim is to generate triangular and tetrahedral meshes with well-sized and well-shaped elements:

- Given a positive sizing field σ, a simplex of circumcenter c is said to be well-sized if its circumradius is at most $\sigma(c)$.
- Following [103], we say that a simplex is well-shaped if its aspect ratio is less than a given threshold. The aspect ratio of a simplex is the ratio between its circumradius and the radius of its inscribed sphere.

Another shape quality measure will be used in the thesis: the so-called radius-edge ratio. The radius-edge ratio of a simplex is the ratio between its circumradius and the length of its shortest edge. Any triangle with a small radius-edge ratio has a small aspect ratio. This property does not hold however for tetrahedra: the so-called *slivers* have small radius-edge ratios but a large aspect ratios. Roughly speaking, a sliver is a tetrahedron whose vertices are close to a great circle of its circumsphere and equally spaced along this circle. See [41] for additional information on this topic.

Part A

Sampling Conditions

Introduction

In this part of the thesis, we introduce the concept of loose σ-sample and we study its various properties on smooth and on Lipschitz surfaces. Our theoretical results will be instrumental in proving the correctness of our sampling algorithm, in Part B.

The concept of loose σ-sample can be viewed as a weak version of the notion of σ-sample. The main advantage of loose σ-samples over σ-samples is that they are easier to check and to construct. Indeed, checking that a sample is a loose σ-sample reduces to checking whether a finite number of spheres are small enough with respect to the value of the sizing field σ at their centers.

Definition A.1 *Given a surface S and a positive function $\sigma : S \to \mathbb{R}$, a finite point set $E \subset S$ is a* loose σ-sample *of S if the following conditions are satisfied:*

1. $\forall p \in S \cap \mathrm{VG}(E)$, $E \cap B(p,\ \sigma(p)) \neq \emptyset$,
2. $\mathrm{Del}_{|S}(E)$ *has vertices on all the connected components of S.*

Since the centers of the surface Delaunay balls are precisely the intersection points of S with the Voronoi edges, Condition 1 of Definition A.1 is verified iff every surface Delaunay ball $B(c, r)$ has a radius $r \leq \sigma(c)$.

Observe that Condition 1 alone is not sufficient to control the density of E. Indeed, according to our definition of the restricted Delaunay triangulation, a point of E is a vertex of $\mathrm{Del}_{|S}(E)$ only if at least one edge of the boundary of its Voronoi cell intersects S. It follows that some of the points of E may not be vertices of $\mathrm{Del}_{|S}(E)$. In some situations (see Figure 3 for an example), $\mathrm{Del}_{|S}(E)$ may even be empty, in which case Condition 1 is trivially verified for any positive function σ, but not Condition 2.

Note to the reader: *In the sequel, σ is assumed to be bounded from above by some constant ε. As a consequence, any (loose) σ-sample is a (loose) ε-sample. This assumption allows us to use a single concept throughout the thesis, namely the concept of (loose) ε-sample. Moreover, it simplifies the proofs in the smooth setting.*

In Chapter 1, we focus on the case where the surface S is smooth. In this context, we show that loose ε-samples of S enjoy the same topological and geometric properties as ε-samples of S, provided that ε is small enough. These properties trivially hold for (loose) σ-samples as well, provided that $\sigma \leq \varepsilon$. In addition, we prove that loose ε-samples and ε-samples are equivalent asymptotically, when ε goes to zero.

Chapter 2 deals with the more general case where S is a k-Lipschitz surface, smooth or nonsmooth. We show that, for sufficiently small k and ε, (loose) ε-samples offer the same guarantees in this context as in the smooth setting, provided that an additional sparseness condition is satisfied – see Section 2.2. These properties trivially hold for (loose) σ-samples as well, provided that σ is at most ε and satisfies the sparseness condition of Section 2.2. In addition, like in the smooth setting, we prove that loose ε-samples and ε-samples are equivalent asymptotically, when ε goes to zero.

In Chapter 3, we work out a lower bound on the size of (loose) ε-samples. This bound holds both in the smooth and in the Lipschitz settings. Moreover, we show that, under a mild sparseness condition[2] introduced in Definition 3.2, the size of any (loose) ε-sample is also bounded from above.

[2] This sparseness condition is different from that of Section 2.2, but it is satisfied when the latter is.

Chapter 1

The Smooth Case

In this chaper, we focus on the case where the surface S is smooth. By *smooth*, we mean that S is $C^{1,1}$-continuous, *i.e.* it is continuously differentiable and its normal satisfies a Lipschitz condition. We assume without loss of generality that the normal of S is oriented consistently, say it always points outwards.

In Section 1.1, we recall a few known facts about smooth surfaces and their medial axis, defined below. We also prove several local properties of loose ε-samples. In Section 1.2, we review the main global properties of loose ε-samples. Specifically, we prove that their restricted Delaunay triangulation is a manifold without boundary (1.2.1), isotopic to S (1.2.2) and at Fréchet distance $O_S(\varepsilon^2)$ from S (1.2.3). In Section 1.3, we show that loose ε-samples and ε-samples are closely related. To complete our study, we give additional properties in Section 1.4, that will not be used in this thesis but which can be useful in some applications. Finally, in Section 1.5, we recall analogous results that were proved in [23] for (loose) $\mu\mathrm{d}_{\mathrm{M}}$-samples.

1.1 Preliminaries

1.1.1 Positive reach

The *medial axis* of S, noted M, is the topological closure of the set of points of $\mathbb{R}^3$ that have more than one nearest neighbor on S. For a point $p \in \mathbb{R}^3$, we call *distance to the medial axis* at p, and write $\mathrm{d}_{\mathrm{M}}(p)$, the Euclidean distance from p to the medial axis of S. As noticed by Amenta and Bern [6], d_{M} is 1-Lipschitz, i.e.

$$\forall p, q \in \mathbb{R}^3,\ |\mathrm{d}_{\mathrm{M}}(p) - \mathrm{d}_{\mathrm{M}}(q)| \leq \mathrm{d}(p, q)$$

The *reach* of S is the infimum over S of the distance to M:

$$\mathrm{rch}(S) = \inf \{\mathrm{d}_{\mathrm{M}}(p),\ p \in S\}$$

As proved in [66], $\mathrm{rch}(S)$ is positive when S is $C^{1,1}$-continuous.

On surfaces with positive reach, ε-samples enjoy many beautiful properties, for sufficiently small values of ε. The following result, due to Amenta and Bern [6, Th. 2], is of particular interest in our context:

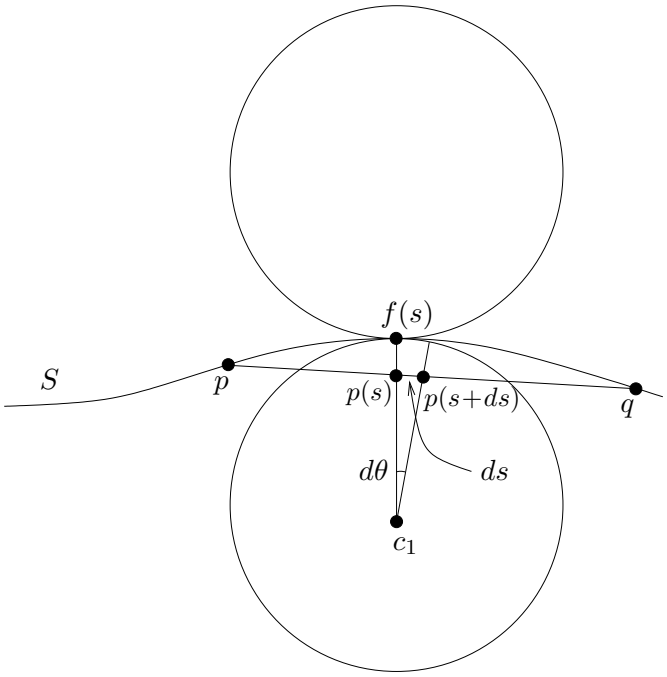

Figure 1.1: For the proof of Lemma 1.2.

Theorem 1.1 *If E is an ε-sample of S, with $\varepsilon < 0.1\ \mathrm{rch}(S)$, then $\mathrm{Del}_{|S}(E)$ is homeomorphic to S.*

Note that $\mathrm{Del}_{|S}(E)$ in the above statement refers to the classical notion of restricted Delaunay triangulation. However, it is proved in [6] that both notions coincide under the hypothesis of the theorem. We list below the main other properties of ε-samples, for sufficiently small ε:

- *Normals:* the angle between the normal to a facet f of $\mathrm{Del}_{|S}(E)$ and the normal to S at the vertices of f is $O(\varepsilon)$ [6].
- *Area:* the area of $\mathrm{Del}_{|S}(E)$ approximates the area of S [92].
- *Curvatures:* the curvature tensor of S can be estimated from $\mathrm{Del}_{|S}(E)$ [49].
- *Hausdorff distance:* the Hausdorff distance between S and $\mathrm{Del}_{|S}(E)$ is $O(\varepsilon^2)$ [23].
- *Reconstruction:* several algorithms can reconstruct from E a surface that is homeomorphic [6, 7, 19, 55] or even ambient isotopic [9] to S.

We will prove in this Chapter that the above properties hold for loose ε-samples as well.

1.1.2 Local properties

We will now review several local properties that will be instrumental in the next sections. The first result is an adaptation of Lemma 3 of [6]. We recall its proof, for completeness.

Lemma 1.2 (Normal Variation)
Let p and q be two points of S with $\mathrm{d}(p,q) \leq \mu\ \mathrm{rch}(S)$, $\mu < 1$. The angle $(\boldsymbol{n}(p), \boldsymbol{n}(q))$ between the normals $\boldsymbol{n}(p)$ and $\boldsymbol{n}(q)$ is at most $\frac{\mu}{1-\mu}$.

Proof. We parameterize the line segment $[p,q]$ by arc length. Let $p(s)$ denote the point on $[p,q]$ with parameter value s. We have $p(0) = p$ and $p(\mathrm{d}(p,q)) = q$. Let $f(s)$ denote the point of S closest to

$p(s)$. Point $f(s)$ is unique, because otherwise $p(s)$ would be a point of the medial axis, contradicting $\mathrm{d}(p,q) \leq \mu \ \mathrm{rch}(S)$ with $\mu < 1$.

Let $\mathbf{n}(s)$ denote the unit normal to S at $f(s)$, and let $\|\mathbf{n}'(s)\|$ denote the magnitude of the derivative of $\mathbf{n}(s)$ with respect to s. The change in normal between p and q is at most $\int_{[p,q]} \|\mathbf{n}'(s)\| \, ds$, which is at most $\mathrm{d}(p,q) \ \sup_s \|\mathbf{n}'(s)\|$.

The surface S passes between two empty balls B_1 and B_2 centered on the medial axis and tangent to S at $f(s)$. Let c_1 and c_2 be the centers of these tangent balls. We have $\mathrm{d}(f(s), c_i) \geq \mathrm{d_M}(f(s)) \geq \mathrm{rch}(S)$, for $i = 1, 2$. Moreover, since $f(s)$ is the only nearest neighbor of $p(s)$ on S, $p(s)$ lies either on $[c_1, f(s)]$ or on $[c_2, f(s)]$, as depicted in Figure 1.1. We assume without loss of generality that $p(s) \in [c_1, f(s)]$.

Since B_1 is an empty ball tangent to S at $f(s)$, the greater of the two principal curvatures of S at $f(s)$ is bounded by the curvature of B_1. Moreover, the rate at which the normal of S changes with $f(s)$ is at most the greater principal curvature, hence $\|\mathbf{n}'(s)\|$ is at most the rate at which the normal turns (as a function of s) on B_1. Referring to Fig. 1.1, we have:

$$ds \geq (\mathrm{rch}(S) - \mathrm{d}(p(s), f(s))) \ \sin d\theta$$

Now, $\sin d\theta$ approaches $d\theta$ as the latter goes to zero. Since $\mathbf{n}(s)$ is a unit vector, we get:

$$\|\mathbf{n}'(s)\| = \frac{d\theta}{ds} \leq 1/(\mathrm{rch}(S) - \mathrm{d}(p(s), f(s)))$$

Moreover, we have:

$$\mathrm{d}(p(s), f(s)) \leq \mathrm{d}(p(s), p) \leq \mu \ \mathrm{rch}(S)$$

Altogether, we obtain: $\sup_s \|\mathbf{n}'(s)\| \leq \frac{1}{(1-\mu) \ \mathrm{rch}(S)}$, which yields the lemma. □

Lemma 1.3 (Line)
Let $\boldsymbol{n}$ be a vector and Ω be a convex such that $\forall p \in S \cap \Omega$, the angle $(\boldsymbol{n}(p), \boldsymbol{n})$ is less than $\frac{\pi}{2}$. Then any line l aligned with $\boldsymbol{n}$ intersects $S \cap \Omega$ at most once.

Proof. Let us assume for a contradiction that there exists a line l aligned with $\mathbf{n}$ and such that $|l \cap S \cap \Omega| \geq 2$. Let p and q be two points of intersection that are consecutive along l. Since Ω is convex, $\Omega \cap l$ is a segment of l, hence p and q are consecutive among the points of $S \cap l$. It follows that the open segment $]p, q[$ is included in one component of $\mathbb{R}^3 \setminus S$, and that $\mathbf{n}(p)$ or $\mathbf{n}(q)$ has a negative or zero inner product with $\mathbf{n}$, which contradicts the hypothesis of the lemma. □

Let p be a point of S and μ a positive constant. We call B_p the closed ball centered at p of radius $\mu \ \mathrm{rch}(S)$, and we set $D_p = S \cap B_p$.

Lemma 1.4 (Cocone)
If $\mu < \frac{\pi}{2+\pi}$, then, for any $q \in D_p$, D_p lies outside the double cone $K(q)$ of apex q, of axis aligned with $\boldsymbol{n}(p)$ and of half-angle $\frac{\pi}{2} - \frac{\mu}{1-\mu}$.

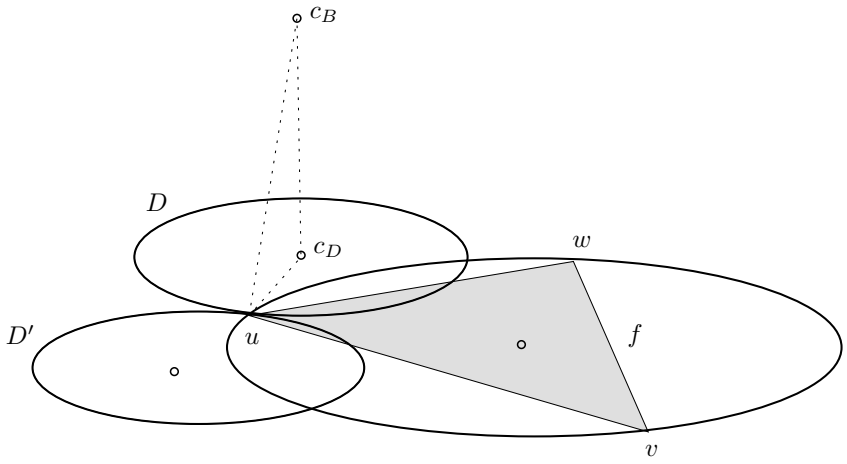

Figure 1.2: For the proof of Lemma 1.5.

Proof. Let l be a line included in $K(q)$, and $\mathbf{l}$ be a vector aligned with l such that $\mathbf{n}(p) \cdot \mathbf{l} \geq 0$. By definition of $K(q)$, we have $(\mathbf{n}(p), \mathbf{l}) < \frac{\pi}{2} - \frac{\mu}{1-\mu}$. Moreover, by the normal variation lemma 1.2, we have $(\mathbf{n}(p), \mathbf{n}(q')) \leq \frac{\mu}{1-\mu}$, for any $q' \in D_p$. Thus, $(\mathbf{n}(q'), \mathbf{l}) < \frac{\pi}{2} - \frac{\mu}{1-\mu} + \frac{\mu}{1-\mu} = \frac{\pi}{2}$. Since this is true for any point $q' \in D_p$, l intersects D_p at most once, by the line lemma 1.3. Hence, $l \cap D_p = \{q\}$. □

Let now $f = (u, v, w)$ be a triangle whose vertices u, v and w belong to S. Let $\mathbf{n}(f)$ denote the unit vector orthogonal to f such that $\mathbf{n}(f) \cdot \mathbf{n}(u) \geq 0$. The next result is an adaptation of Lemma 7(a) of [6]. We recall its proof, for completeness.

Lemma 1.5 (Triangle Normal)
If the circumradius of f is at most $\mu \operatorname{rch}(S)$, with $\mu < 1$, and if the inner angle $\hat{u} = \angle vuw$ is at least $\frac{\pi}{3}$, then $(\boldsymbol{n}(f), \boldsymbol{n}(u)) \leq \mu \sqrt{3}$.

Proof. Let B and B' denote the two balls of radius $\operatorname{rch}(S)$ tangent to S at u, $D = B \cap \operatorname{aff}(f)$ and $D' = B' \cap \operatorname{aff}(f)$. We call c_B and c_D the centers of B and D respectively. Moreover, $r_f \leq \mu \operatorname{rch}(S)$ denotes the circumradius of f, and $r_D \leq \operatorname{rch}(S)$ denotes the radius of D and D' (which have same radius, by a symmetry argument).

Let r_f be fixed. Since the interiors of D and D' do not intersect S, they contain no vertex of f. Therefore, r_D is maximized with respect to r_f when v and w are farthest from u and lie on the boundaries of D and D'. Since we assumed that $\hat{u} \geq \frac{\pi}{3}$, v and w are farthest from u when f is equilateral. It follows that $r_D \leq r_f \sqrt{3}$.

In addition, according to the hypothesis, $r_f \leq \mu \operatorname{rch}(S)$. We then have:

$$\sin(\mathbf{n}(f), \mathbf{n}(u)) = \sin(\angle u c_B c_D) \leq \frac{\mathrm{d}(u, c_D)}{\mathrm{d}(u, c_B)} = \frac{r_D}{\operatorname{rch}(S)} \leq \mu \sqrt{3} \tag{1.1}$$

□

1.2 Global properties

In this section, we review the main global properties of $\mathrm{Del}_{|S}(E)$ in the case where E is a loose ε-sample of S with sufficiently small ε. Let $\mu_0 = 0.16$ and $\nu = 0.52$. These two constants are set up so that the bounds in the following statements are as tight as possible. Throughout the section, E denotes a loose ε-sample of S. Let $\mu = \varepsilon/\mathrm{rch}(S)$.

For any facet f of $\mathrm{Del}_{|S}(E)$, we call B_f the surface Delaunay ball of smallest radius that circumscribes f. Let c_f and r_f denote respectively the center and radius of B_f. We set $D_f = S \cap B_f$.

Orientation convention 1.6 *For any facet $f \in \mathrm{Del}(E)$ of circumradius less than $(1+\nu)\varepsilon$, we orient f such that $\boldsymbol{n}(f) \cdot \boldsymbol{n}(u_f) > 0$, where u_f is the vertex of f with largest inner angle (if it is not unique, we choose any such vertex).*

Notice that it is not necessary to orient all the facets of $\mathrm{Del}(E)$, because only those of circumradius less than $(1+\nu)\varepsilon$ will be considered in the sequel. Among these are the facets of $\mathrm{Del}_{|S}(E)$, which are included in surface Delaunay balls of radius at most ε.

In Section 1.2.1 (Th. 1.7) we will show that, if $\mu \leq \mu_0$, then $\mathrm{Del}_{|S}(E)$ is an oriented manifold without boundary. The proof does not rely on the fact that the surface S is smooth, therefore Theorem 1.7 will be used in Chapter 2 as well, for the nonsmooth case. However, we have to work under slightly more general hypotheses. Specifically, we prove that, if E is a loose ε-sample of S (for some ε which can be greater than $\mu_0\ \mathrm{rch}(S)$) satisfying the following assumption:
For any vertex v of $\mathrm{Del}_{|S}(E)$, there is a point $p_v \in S$ where S is differentiable, such that:

- **M1** *For any $q \in S \cap B(v,\varepsilon)$, $S \cap B(v,\varepsilon)$ lies outside the double cone of apex q, of axis aligned with $\boldsymbol{n}(p_v)$ and of half-angle $\frac{\pi}{2} - \theta_1$.*
- **M2** *For any restricted Delaunay facet f incident to v, $(\boldsymbol{n}(f), \boldsymbol{n}(p_v)) \leq \theta_2$.*
- **M3** *For any Delaunay facet f incident to v, of circumradius less than $(1+\nu)\varepsilon$, $(\boldsymbol{n}(f), \boldsymbol{n}(p_v)) \leq \theta_3$.*

where θ_1, θ_2 and θ_3 depend only on ε and verify:

- **M4** $2\theta_1 + \theta_2 < \frac{\pi}{2}$.
- **M5** $2\ \sin\theta_1 < \nu\ \cos\theta_2$.
- **M6** $\sin\theta_1 < \cos\theta_3$.

then $\mathrm{Del}_{|S}(E)$ is an oriented manifold without boundary.

When S is smooth and $\mu \leq \mu_0$, for any vertex v of $\mathrm{Del}_{|S}(E)$ we take $p_v = v$, and we set $\theta_1 = \frac{\mu}{1-\mu}$, $\theta_2 = \mu\ \sqrt{3} + \frac{2\mu}{1-2\mu}$ and $\theta_3 = (1+\nu)\ \mu\ \sqrt{3} + \frac{2(1+\nu)\,\mu}{1-2(1+\nu)\,\mu}$. By the cocone lemma 1.4, M1 is satisfied. Moreover, by the triangle normal lemma 1.5 and the normal variation lemma 1.2, M2 and M3 are satisfied as well. Finally, since $\mu \leq \mu_0$, M4, M5 and M6 are also satisfied. Therefore, $\mathrm{Del}_{|S}(E)$ is an oriented manifold without boundary.

In Sections 1.2.2 and 1.2.3, we will prove that $\mathrm{Del}_{|S}(E)$ is isotopic to S and at Fréchet distance $O(\varepsilon^2)$ from S, provided that $\mu \leq \mu_0$. The proofs use the fact that $\mathrm{Del}_{|S}(E)$ is a manifold, as guaranteed above. Moreover, they hold in a slightly more general setting. Specifically, we will show that, for any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ verifying:

I1 *$\hat{S}$ is a compact surface without boundary, consistently oriented by the orientation convention 1.6,*

I2 *$\hat{S}$ has vertices on all the connected components of S,*

I3 *For any facet f of $\hat{S}$, B_f has a radius at most $\varepsilon \leq \mu_0 \operatorname{rch}(S)$,*

$\hat{S}$ is isotopic to S and at Fréchet distance $O(\varepsilon^2)$ from S. Our arguments rely heavily on the fact that the surface S is smooth, hence the results cannot be used in Chapter 2, for the nonsmooth case. Nevertheless, the proofs of Chapter 2 will keep the same spirit.

The fact that $\hat{S}$ is a subcomplex of $\mathrm{Del}_{|S}(E)$ (and not $\mathrm{Del}_{|S}(E)$ itself) will be instrumental in proving the correctness of certain Delaunay refinement algorithms in several meshing applications. See Section 5.1 and Chapter 8.

1.2.1 Manifold

This section is dedicated to the proof of the following result:

Theorem 1.7 *If E is a loose ε-sample of S, such that M1–M6 are satisfied, then $\mathrm{Del}_{|S}(E)$ is a compact oriented surface without boundary.*

In Section 1.2.1.1, we show that every edge of $\mathrm{Del}_{|S}(E)$ is incident to exactly two facets of $\mathrm{Del}_{|S}(E)$. In Section 1.2.1.2, we show that every vertex of $\mathrm{Del}_{|S}(E)$ is incident to exactly one cycle of facets of $\mathrm{Del}_{|S}(E)$. Such a cycle will be called an *umbrella*. These two properties imply that $\mathrm{Del}_{|S}(E)$ is a 2-manifold without boundary, because the relative interiors of the faces of $\mathrm{Del}_{|S}(E)$ are pairwise disjoint due to the fact that $\mathrm{Del}_{|S}(E)$ is a simplicial complex. Finally, in Section 1.2.1.3, we prove that the orientation convention 1.6 induces a valid orientation of $\mathrm{Del}_{|S}(E)$.

1.2.1.1 Edges

We need an intermediate result:

Lemma 1.8 (Projection)
Under M1, M2 and M4, for any facets f, f' of $\mathrm{Del}_{|S}(E)$ with a common edge e, for any vertex v of e, the orthogonal projections of f and f' onto $T(p_v)$ do not overlap, i.e. their interiors are disjoint.

Proof. Let u be the second vertex of e. For convenience, we note $B = B_f$, $r = r_f$ the radius of B_f, and $c = c_f$ the center of B_f. Similarly, we call $B' = B_{f'}$, $r' = r_{f'}$, $c' = c_{f'}$.

Since u and v are vertices of f and f', ∂B and $\partial B'$ have a non-empty intersection. Let P be the radical plane of B, B'. P contains $\partial B \cap \partial B'$, and therefore also u and v. Let w be the third vertex of f, and w' be the third vertex of f'. P is perpendicular to the line (c, c'). On one side of P, B is included in the interior of B', whereas on the other side of P, B' is included in the interior of B. Since B and B' are Delaunay balls, B cannot contain w' and B' cannot contain w. Hence, w and w' cannot lie on the same side of P.

Let Π be the plane that contains u, v and $\mathbf{n}(p_v)$. Since Π is orthogonal to $T(p_v)$, its orthogonal projection onto $T(p_v)$ is a line containing the projection of e. Hence, to prove that the projections of f and f' onto $T(p_v)$ do not overlap, it suffices to show that w and w' lie on different sides of Π.

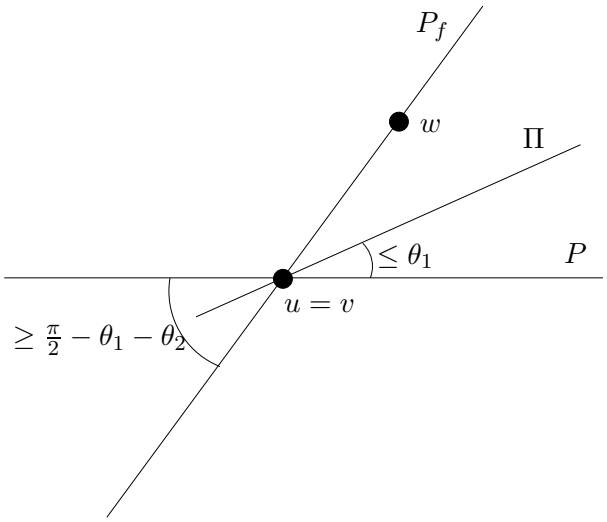

Figure 1.3: For the proof of Lemma 1.8.

P and Π intersect each other along the line (u, v). We define an oriented frame in $\mathbb{R}^3$ such that P is the horizontal plane and that w lies above P. The line $(u, v) = P \cap \Pi$ is then horizontal. We will show that w lies above Π while w' lies below, which will conclude the proof of the lemma.

Since $r \leq \varepsilon$ and $r' \leq \varepsilon$, c and c' belong to $S \cap B(v, \varepsilon)$. Thus, by M1 (applied with $q = c$), the angle between $\mathbf{n}(p_v)$ and the line (c, c') is at least $\frac{\pi}{2} - \theta_1$. This implies that the angle between $\mathbf{n}(p_v)$ and P is at most θ_1, since P is orthogonal to the line (c, c'). Hence, the angle $\angle \Pi P$ between planes Π and P is bounded by θ_1. In addition, by M2 we have $(\mathbf{n}(f), \mathbf{n}(p_v)) \leq \theta_2$. It follows that the angle between $\mathbf{n}(f)$ and P is at most $\theta_1 + \theta_2$, or equivalently, that the angle $\angle P_f P$ between the supporting plane of f and P is at least $\frac{\pi}{2} - \theta_1 - \theta_2$. By M4, this quantity is greater than θ_1, which means that $\angle P_f P$ is greater than $\angle \Pi P$. Hence, w lies above Π – see Fig. 1.3. By the same arguments, w' lies below Π, which concludes the proof of the lemma. □

Remark 1.9 *Observe that the proof of the projection lemma still holds if B and B' are two surface Delaunay balls circumscribing the same facet $f = f'$. Hence, B_f is the only surface Delaunay ball that circumscribes f. Equivalently, the Voronoi edge $\mathrm{V}(f)$ dual to f intersects S only once. Moreover, the vertices of $\mathrm{V}(f)$ belong to distinct connected components of $\mathbb{R}^3 \setminus S$ since otherwise a small perturbation of $\mathrm{V}(f)$ would intersect $S \cap B_f$ twice, hereby contradicting the projection lemma 1.8 (with $f = f'$).*

From the projection lemma 1.8, we deduce the following result:

Lemma 1.10 *Under M1, M2 and M4, every edge of* $\mathrm{Del}_{|S}(E)$ *is incident to exactly two facets of* $\mathrm{Del}_{|S}(E)$.

Proof. Let e be an edge of $\mathrm{Del}_{|S}(E)$. By definition, e is incident to a Delaunay facet f whose dual Voronoi edge $\mathrm{V}(f)$ intersects S. $\mathrm{V}(f)$ is an edge of $\partial \mathrm{V}(e)$, the boundary of the Voronoi facet dual to e. By Remark 1.9, the vertices of $\mathrm{V}(f)$ lie in distinct connected components of $\mathbb{R}^3 \setminus S$. Hence, $\partial \mathrm{V}(e)$

intersects S at least twice, since it is a topological circle that intersects two distinct connected coponents of $\mathbb{R}^3 \setminus S$. As a consequence, e is incident to at least two facets of $\mathrm{Del}_{|S}(E)$.

In addition, e cannot be incident to more than two facets of $\mathrm{Del}_{|S}(E)$. Indeed, by the projection lemma 1.8, the projections onto $T(p_v)$ (where v is any vertex of e and p_v is defined as in M1, M2, M4) of the restricted Delaunay facets incident to e pairwise do not overlap, thus they must lie on different sides of the line supporting the projection of e.

In conclusion, the number of facets of $\mathrm{Del}_{|S}(E)$ incident to e is two. □

1.2.1.2 Umbrellas

Consider a vertex v of $\mathrm{Del}_{|S}(E)$. Since every edge of $\mathrm{Del}_{|S}(E)$ incident to v has two incident facets of $\mathrm{Del}_{|S}(E)$, the star of v in $\mathrm{Del}_{|S}(E)$ consists of one or more cycles of facets, called *umbrellas*. Each umbrella is a triangulated topological disk. All umbrellas of v have v in common, but two distinct umbrellas have distinct edges and facets. We call $\bar{v}$ the orthogonal projection of v onto $T(p_v)$.

Lemma 1.11 *Under M1, M2, M3, M5 and M6, v has exactly one umbrella.*

Proof. Let U be an umbrella incident to v. We call $\bar{U}$ its orthogonal projection onto $T(p_v)$.

Claim 1.11.1 *$\bar{v}$ belongs to the interior of $\bar{U}$.*

Proof. Let us assume the contrary. Then U has a silhouette edge $[v, u]$ whose projection onto $T(p_v)$ belongs to the boundary of $\bar{U}$. Since, by lemma 1.10, $[v, u]$ is incident to two facets of U, these two facets project onto the same side of the line supporting the projection of $[u, v]$, therefore they operlap, which contradicts the projection lemma 1.8. This proves the claim. □

Assume now, for a contradiction, that v has several umbrellas. Let U and U' be two consecutive umbrellas (recall that two umbrellas of v intersect only in v). Let $\mathcal{T}$ be the set of tetrahedra of $\mathrm{Del}(E)$ incident to v that lie between U and U'.

Let l be a line parallel to $\mathbf{n}(p_v)$ and passing through a point of $\mathrm{Del}_{|S}(E)$ very close to v. We orient l as $\mathbf{n}(p_v)$. Let $T = (t_1, \ldots, t_s)$ be the sequence of tetrahedra of $\mathcal{T}$ that are intersected by the oriented line l, and let $F = (f_0, \ldots, f_s)$ be the facets of T pierced by l. We assume without loss of generality that l does not pass through an edge of one of the t_i. By definition, $f_0 \in U$ and $f_s \subset U'$ are facets of $\mathrm{Del}_{|S}(E)$. We rename them f and f' respectively, and call c and c' the centers of their surface Delaunay balls. Let $c_1, \ldots, c_s$ be the centers of the Delaunay balls circumscribing $t_1, \ldots, t_s$, and γ be the polygonal chain $\gamma = c, c_1, \ldots, c_s, c'$. For convenience, we write $c = c_0$ and $c' = c_{s+1}$. Clearly, γ is a path in the 1-skeleton graph of $\mathrm{Vor}(E)$.

Claim 1.11.2 *γ is monotone with respect to the oriented line l.*

Proof. l intersects t_i before t_{i+1}, for $i = 1, \ldots, s-1$. Thus, $\mathbf{n}(p_v) \cdot \mathbf{n}_i > 0$, where $\mathbf{n}_i$ is the unit normal vector of f_i, oriented from t_i to t_{i+1}. Since t_i and t_{i+1} are Delaunay tetrahedra, we have $\mathbf{n}_i = \frac{\mathbf{c}_{i+1} - \mathbf{c}_i}{\|\mathbf{c}_{i+1} - \mathbf{c}_i\|}$.

Hence, $(\mathbf{c}_{i+1} - \mathbf{c}_i) \cdot \mathbf{n}(p_v) > 0$, for $i = 1, \ldots, s-1$. By the same argument, $(\mathbf{c}_1 - \mathbf{c}) \cdot \mathbf{n}(p_v) > 0$ and $(\mathbf{c}' - \mathbf{c}_s) \cdot \mathbf{n}(p_v) > 0$, which ends the proof of the claim. □

By the monotonicity property above, γ lies in the slab W defined by the two planes Π and Π' orthogonal to $\mathbf{n}(p_v)$ and passing respectively through c and c'. Since c is the center of the surface Delaunay ball of a facet of $\mathrm{Del}_{|S}(E)$ incident to v, we have $\mathrm{d}(c, v) \leq \varepsilon$ and c belongs to $S \cap B(v, \varepsilon)$. It follows, by M1 (applied with $q = v$), that the angle between $\mathbf{n}(p_v)$ and the line (c, v) is at least $\frac{\pi}{2} - \theta_1$. Hence, the distance from v to Π is at most $h = \varepsilon \ \sin\theta_1$. Similarly, the distance from v to Π' is at most h.

We now bound the circumradii of the tetrahedra t_i. Let $\mathrm{V}(U)$ be the set of Voronoi facets dual to the edges of U that are incident to v, and let $\mathrm{V}(U)_v^+$ be the intersection of the halfspaces that are bounded by the affine hulls of those facets and contain v. The circumcenters c_i of the tetrahedra t_i are vertices of the Voronoi cell of v, which is convex. Hence, they belong to $\mathrm{V}(U)_v^+$. Moreover, since γ is monotone w.r.t. the oriented line l (Claim 1.11.2), the c_i also belong to W. The intersections of $\mathrm{V}(U)_v^+$ with Π and Π' respectively are convex polygons whose vertices are the intersection points between Π or Π' and the affine hulls of the dual Voronoi edges of the facets of U. Any such Voronoi edge e contains the center c_e of a surface Delaunay ball, such that $\mathrm{d}(c_e, v) \leq \varepsilon$. By M2, the angle between e and $\mathbf{n}(p_v)$ is at most θ_2. It follows that $e \cap \Pi$ and $e \cap \Pi'$ are not farther than $2h/\cos\theta_2$ from c_e. Hence, the distance from v to any vertex of $\mathrm{V}(U)_v^+ \cap \Pi$ or $\mathrm{V}(U)_v^+ \cap \Pi'$ is bounded by:

$$\eta \ = \ \varepsilon + \frac{2h}{\cos\theta_2} \ = \ \varepsilon \left(1 + \frac{2\sin\theta_1}{\cos\theta_2}\right)$$

Since the circumcenters c_i lie inside $\mathrm{V}(U)_v^+ \cap W$, their distance to v is at most η, which less than $(1+\nu)\varepsilon$, by M5. Hence, for all i, the circumradius of f_i is less than $(1+\nu)\varepsilon$, which implies that f_i is oriented by the orientation convention 1.6. Moreover, by M3, $(\mathbf{n}(f_i), \mathbf{n}(p_v)) \leq \theta_3$, which is less than $\frac{\pi}{2}$, by M6. According to Claim 1.11.2, the unit vector $n_i = \frac{\mathbf{c}_{i+1} - \mathbf{c}_i}{\|\mathbf{c}_{i+1} - \mathbf{c}_i\|}$ orthogonal to f_i and oriented from t_i to t_{i+1} has a positive inner product with $\mathbf{n}(p_v)$. Therefore, $\mathbf{n}_i = \mathbf{n}(f_i)$, which yields:

$$(\mathbf{c}_{i+1} - \mathbf{c}_i) \cdot \mathbf{n}(p_v) \geq \cos\theta_3 \ \|\mathbf{c}_{i+1} - \mathbf{c}_i\|, \tag{1.2}$$

and, by summing over all edges of γ,

$$\begin{aligned} \|\mathbf{c}' - \mathbf{c}\| \ \cos\theta_3 \quad &\leq \quad \sum_{i=0}^{s} \|\mathbf{c}_{i+1} - \mathbf{c}_i\| \ \cos\theta_3 \\ &\overset{(1.2)}{\leq} \quad \sum_{i=0}^{s} (\mathbf{c}_{i+1} - \mathbf{c}_i) \cdot \mathbf{n}(p_v) \\ &= \quad (\mathbf{c}' - \mathbf{c}) \cdot \mathbf{n}(p_v) \end{aligned}$$

In addition, since $\mathrm{d}(c, v) \leq \varepsilon$ and $\mathrm{d}(c', v) \leq \varepsilon$, c and c' belong to $S \cap B(v, \varepsilon)$. Hence, by M1 (applied with $q = c$), the angle $(\mathbf{n}(p_v), (\mathbf{c}' - \mathbf{c}))$ is at least $\frac{\pi}{2} - \theta_1$, which implies that $(\mathbf{c}' - \mathbf{c}) \cdot \mathbf{n}(p_v) \leq \|\mathbf{c}' - \mathbf{c}\| \ \sin\theta_1$. This raises a contradiction with $\|\mathbf{c} - \mathbf{c}'\| \ \cos\theta_3 \leq (\mathbf{c}' - \mathbf{c}) \cdot \mathbf{n}(p_v)$, by M6. Lemma 1.11 is proved. □

We deduce from Lemmas 1.10 and 1.11 that $\mathrm{Del}_{|S}(E)$ is a 2-manifold without boundary. Moreover, it is compact since it is included in the convex hull of E, which is bounded. In addition, by Lemmas 1.8 and 1.11, the facets of $\mathrm{Del}_{|S}(E)$ that share a vertex v do not overlap in projection onto $T(p_v)$.

1.2.1.3 Orientation

The next result concludes the proof of Theorem 1.7:

Lemma 1.12 *The orientation convention induces a valid orientation of the closed 2-manifold* $\mathrm{Del}_{|S}(E)$.

Proof. Let f and f' be two facets of $\mathrm{Del}_{|S}(E)$ with a common edge. These facets are oriented according to the orientation convention. Let v be a common vertex of f and f'. We call f_p and ${f'}_p$ the respective orthogonal projections of f and f' onto the tangent plane $T(p_v)$. Inside $T(p_v)$, we endow triangles f_p and ${f'}_p$ with the same orientations as f and f' respectively. By M2, $\mathbf{n}(f) \cdot \mathbf{n}(p_v)$ and $\mathbf{n}(f') \cdot \mathbf{n}(p_v)$ are both positive. So are $\mathbf{n}(f_p) \cdot \mathbf{n}(p_v)$ and $\mathbf{n}({f'}_p) \cdot \mathbf{n}(p_v)$. Moreover, triangles f_p and ${f'}_p$ share an edge e, and their interiors are disjoint, by the projection lemma 1.8. Thus, their orientations are consistent with e, *i.e.* they induce opposite orientations of e. Hence, the orientations of f and f' are consistent with the common edge of f and f'. □

1.2.2 Isotopy

Let $E \subset S$ be a finite set of points. In this section we show that every subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that verifies I1–I3 is isotopic to S. The proof relies on a classical result of differential geometry, stated as Theorem 1.13, which exploits the topological properties of the so-called *tubular neighborhood* of S. We introduce this notion in Section 1.2.2.1, then we present our proof in Section 1.2.2.2.

1.2.2.1 Tubular neighborhood

The *tubular neighborhood* of width ε around S, or $\mathcal{T}_\varepsilon$ for short, is defined as follows:

$$\mathcal{T}_\varepsilon = \bigcup_{p \in S} B(p, \varepsilon) = \left\{ q \in \mathbb{R}^3 \mid \mathrm{d}(q, S) < \varepsilon \right\} \tag{1.3}$$

Let $\pi_S : \mathbb{R}^3 \to S$ map each point to its nearest neighbor on S. Since $\varepsilon < \mathrm{rch}(S)$, $\mathcal{T}_\varepsilon$ does not intersect the medial axis of S, hence the restriction of π_S to $\mathcal{T}_\varepsilon$ is a well-defined continuous function.

Let p be a point of $\mathcal{T}_\varepsilon$ and $\tilde{p}$ be its image by π_S. Since S is smooth, the line $(p, \tilde{p})$ is aligned with the normal of S at $\tilde{p}$. Moreover, the segment $[p, \tilde{p}]$ is included in $\mathcal{T}_\varepsilon$ since, for any $p' \in [p, \tilde{p}]$, $\mathrm{d}(p', S) = \mathrm{d}(p', \tilde{p}) \leq \mathrm{d}(p, \tilde{p}) < \varepsilon$. It follows that $\mathcal{T}_\varepsilon \cap \pi_S^{-1}(\tilde{p})$ is an open line segment. We call this segment the *fiber* of $\tilde{p}$, or simply $\mathrm{Fib}(\tilde{p})$. Since $\mathcal{T}_\varepsilon \cap \mathrm{M} = \emptyset$, the fibers are pairwise disjoint [66, Th. 4.8(13)].

Theorem 1.13 *If $\hat{S}$ satisfies the following two conditions, where $\varepsilon < \mathrm{rch}(S)$:*

1. $\hat{S} \subset \mathcal{T}_\varepsilon$,

2. $\forall p \in S$, $|\hat{S} \cap \mathrm{Fib}(p)| = 1$,

then the restriction of π_S to $\hat{S}$ induces an isotopy that maps $\hat{S}$ to S. The isotopy does not move the points by more than ε.

This is a well-known result, based on classical arguments of differential geometry [77, Chapter 5]. We recall its proof briefly, for completeness:

Sketch of proof. The idea is to *push* $\hat{S}$ onto S, along the fibers, in a continuous and one-to-one fashion. To this end, we introduce the following map:

$$\begin{array}{rl} \phi: & [0,1] \times \mathcal{T}_\varepsilon \to \mathcal{T}_\varepsilon \\ & (t,p) \mapsto t\,\tilde{p} + (1-t)\,p \end{array} \tag{1.4}$$

where $\tilde{p}$ is the image of p by π_S. Since π_S is well-defined and continuous over $\mathcal{T}_\varepsilon$, ϕ defines a continuous flow towards $\phi(1, \mathcal{T}_\varepsilon) = S$.

Since $\hat{S} \subset \mathcal{T}_\varepsilon$ (Hyp. 1), we can consider the restriction $\phi_{\hat{S}}$ of ϕ to $\hat{S}$. $\phi_{\hat{S}}(0,.)$ is the identity on $\hat{S}$. Moreover, $\phi_{\hat{S}}(1,\hat{S}) = S$, by Hyp. 2. Let $t \in [0,1]$. According to Hyp. 2, the points of $\hat{S}$ belong to different fibers. It follows that $\phi_{\hat{S}}(t,.)$ is injective, since the fibers are pairwise disjoint. Hence, $\phi_{\hat{S}}(t,.)$ is a homeomorphism from $\hat{S}$ onto its image $\phi_{\hat{S}}(t,\hat{S})$, because ϕ is continuous and $\hat{S}$ is compact. As a conclusion, $\phi_{\hat{S}}$ is an isotopy from $\hat{S}$ to S.

The fact that $\phi_{\hat{S}}$ does not move the points of $\hat{S}$ by more than ε comes from the fact that every point $p \in \hat{S}$ is mapped to its nearest neighbor on S, which is closer than ε since $p \in \mathcal{T}_\varepsilon$. □

1.2.2.2 Our result

Theorem 1.14 *For any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that satisfies I1–I3, the restriction of π_S to $\hat{S}$ induces an isotopy that maps $\hat{S}$ to S. The isotopy does not move the points by more than ε.*

Specifically, we will prove successively that, under I1–I3, $\hat{S}$ verifies Hyp. 1 and Hyp. 2 of Theorem 1.13, which will yield the result.

Lemma 1.15 *$\hat{S}$ is included in $\mathcal{T}_\varepsilon$. As a consequence, Hyp. 1 of Theorem 1.13 is satisfied.*

Proof. Let $p \in \hat{S}$. Let f be a facet of $\hat{S}$ that contains p, and $B_f = B(c_f, r_f)$ be the surface Delaunay ball of f. If p is a vertex of $\hat{S}$, then p belongs to S, which is included in $\mathcal{T}_\varepsilon$ because $\varepsilon > 0$. Otherwise, p belongs to the interior of B_f, which means that $\mathrm{d}(p, c_f) < r_f$. Since $c_f \in S$ and $r_f \leq \varepsilon$ by I3, p belongs to $\mathcal{T}_\varepsilon$. □

Let S_1 be the subset of S made of the points whose fibers intersect $\hat{S}$ exactly once. These points are images by π_S of exactly one point of $\hat{S}$. Note that S is endowed with the topology induced by $\mathbb{R}^3$.

Lemma 1.16 *The vertices of $\hat{S}$ belong to* $\mathrm{int}(S_1)$.

Proof. Let v be a vertex of $\hat{S}$. We first show that $v \in S_1$. Since $v \in \hat{S} \cap S$, $|\mathrm{Fib}(v) \cap \hat{S}| \geq 1$. Let $m_v \in \mathrm{M}$ be the center of any medial ball B_v tangent to S at v. Assume that $]v, m_v[$ intersects a facet f of $\hat{S}$. ∂B_v and ∂B_f necessarily intersect since the vertices of f do not belong to the interior of B_v. Let Π be the plane containing $\partial B_v \cap \partial B_f$, Π^+ be the halfspace limited by Π and containing f. Since v does

not belong to the interior of B_f and is on the boundary of B_v, v must belong to Π^- (possibly in Π). On the other hand, by I3, the radius r_f of B_f is less than $\mathrm{rch}(S)$, hence $m_v \in \mathrm{M}$ does not belong to B_f. It follows that m_v is located in the open half space Π^-. Hence, the open segment $]v, m_v[$ is contained in the open half space Π^-, and therefore cannot intersect f which is contained in Π^+. As a consequence, $]v, m_v[\cap \hat{S} = \emptyset$, which means that $|\mathrm{Fib}(v) \cap \hat{S}| = 1$.

We now show that $v \in \mathrm{int}(S_1)$. By Lemmas 1.8 and 1.11 (with $p_v = v$), the facets of the umbrella U of v do not overlap in projection onto $T(v)$. Therefore, the projection from U to $T(v)$ is one-to-one, and there exists a small neighborhood $\mathcal{N}$ of v on S such that every line crossing $\mathcal{N}$ and parallel to $\mathbf{n}(v)$ intersects U once. Moreover, the planes supporting the facets of U cannot be perpendicular to $T(v)$, since otherwise the projection onto $T(v)$ could not be one-to-one. Hence, there exists a $\beta > 0$ such that any line crossing $\mathcal{N}$ and making an angle less than β with $\mathbf{n}(v)$ intersects U once. Since S is C^1-continuous, its Gauss map $p \mapsto \mathbf{n}(p)$ is continuous. Thus, $\mathcal{N}$ can be chosen sufficiently small, so that the fibers of the points of $\mathcal{N}$ make angles less than β with $\mathbf{n}(v)$. As a consequence, for any $p \in \mathcal{N}$, $|\mathrm{Fib}(p) \cap U| = 1$. In addition, since $|\mathrm{Fib}(v) \cap \hat{S}| = 1$, the fiber of v intersects $\hat{S}$ only at v. Hence, its distance to $\hat{S} \setminus U$ is positive. It follows that $\mathcal{N}$ can be chosen sufficently small, so that the fibers of the points of $\mathcal{N}$ do not intersect $\hat{S} \setminus U$. As a consequence, $\mathcal{N} \subseteq S_1$ and $v \in \mathrm{int}(S_1)$. □

It follows from Lemma 1.16 that $S_1 \neq \emptyset$. Moreover, by I2, S_1 intersects all the connected components of S. If we can show that S_1 has an empty boundary, then we will have $S_1 = S$, which implies that Hyp. 2 of Theorem 1.13 is satisfied, hereby concluding the proof of Theorem 1.14.

Lemma 1.17 $\partial S_1 = \emptyset$.

Proof. We assume that $\partial S_1 \neq \emptyset$ and look for a contradiction. Let $p \in \partial S_1$. Let $f_1, \cdots, f_k$ be the facets of $\hat{S}$ intersected by $\mathrm{Fib}(p)$. For any $q \in f_i \cap \mathrm{Fib}(p)$, the nearest neighbor of q on S is p, thus $\mathrm{d}(p, q) \leq \varepsilon$. Hence, the distance from p to the center c_i of the surface Delaunay ball of f_i is at most 2ε, by I3. From I1, from the triangle normal lemma 1.5 and the normal variation lemma 1.2, and from the fact that $\mu \leq \mu_0$, we deduce:

$$(\mathbf{n}(f_i), \mathbf{n}(p)) \leq 2\mu\sqrt{3} + \frac{4\mu}{1 - 4\mu} < \frac{\pi}{2} \tag{1.5}$$

Therefore, every f_i is intersected transversally by $\mathrm{Fib}(p)$. Let $p_i = f_i \cap \mathrm{Fib}(p)$.
- If p_i belongs to the relative interior of f_i, then in a small neighborhood of p on S the fibers intersect f_i transversally, since the Gauss map of S is continuous. Therefore, the number of intersections of a fiber with f_i does not change at p.
- If p_i is a vertex of $\hat{S}$, then p_i belongs to S, which implies that $p_i = p$. Hence, by Lemma 1.16, $p \in \mathrm{int}(S_1)$, which contradicts the assumption that $p \in \partial S_1$.

It follows from the two above observations that at least one of the p_i (say p_1) belongs to the relative interior of an edge e of $\hat{S}$, such that the number of intersections of a fiber with f_1 and its neighbor (through e) f_1' changes at p. This means that $\mathbf{n}(f_1) \cdot \mathbf{n}(p)$ and $\mathbf{n}(f_1') \cdot \mathbf{n}(p)$ have different signs, since the Gauss map of S is continuous. This raises a contradiction with Eq. (1.5) and hereby concludes the proof of Lemma 1.17. □

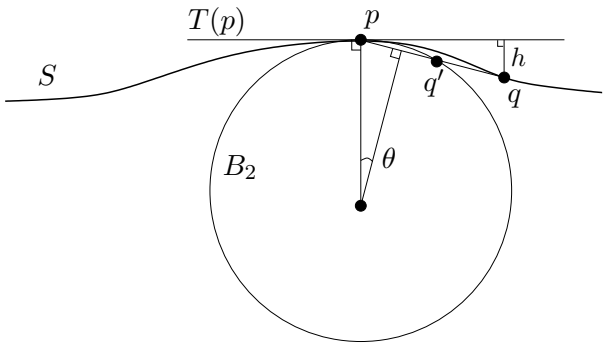

Figure 1.4: For the proof of Lemma 1.19.

1.2.3 Fréchet distance

The next result bounds the Fréchet distance (and hence also the Hausdorff distance) between S and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that verifies I1–I3:

Theorem 1.18 *For any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that verifies I1–I3, the Fréchet distance between S and $\hat{S}$ is at most $\frac{\varepsilon^2}{\mathrm{rch}(S)}$.*

According to Theorem 1.14, the projection onto S, π_S, induces an isotopy between $\hat{S}$ and S. In particular, the restriction of π_S to $\hat{S}$ is a homeomorphism between $\hat{S}$ and its image S. Hence, in view of Definition P.5, to control the Fréchet distance between $\hat{S}$ and S it suffices to bound $\mathrm{d}(p, \pi_S(p)) = \mathrm{d}(p, S)$ for any $p \in \hat{S}$.

Lemma 1.19 *Let $p \in S$. For any $q \in S$, the distance from q to $T(p)$ is at most $\frac{\mathrm{d}^2(p,q)}{2\,\mathrm{rch}(S)}$.*

Proof. Let B_1 and B_2 be the two balls of radius $\mathrm{rch}(S)$ tangent to S at p. Their interiors cannot intersect S and therefore do not contain q. Let q' be the intersection point other than p of the boundary of $B_1 \cup B_2$ with segment $[p, q]$. Let h be the distance of q to $T(p)$ and θ the angle between the line (p, q) and the plane $T(p)$. Referring to Figure 1.4, we have: $2 \sin\theta\ \mathrm{rch}(S) = \mathrm{d}(p, q') \leq \mathrm{d}(p, q)$. Therefore, $\sin\theta \leq \frac{\mathrm{d}(p,q)}{2\,\mathrm{rch}(S)}$ and $h = \mathrm{d}(p, q) \sin\theta \leq \frac{\mathrm{d}^2(p,q)}{2\,\mathrm{rch}(S)}$. □

Lemma 1.20 *Let $p \in S$. For any $q \in T(p)$, the distance of q to S is at most $\frac{\mathrm{d}^2(p,q)}{2\,\mathrm{rch}(S)}$.*

Proof. Let B_1 and B_2 be the two balls of radius $\mathrm{rch}(S)$ tangent to S at p. Their interiors cannot intersect S and therefore do not contain q. Moreover, since S has no boundary, it separates the tangent balls B_1 and B_2, in the sense that they belong to different connected components of $\mathbb{R}^3 \setminus S$.

In the plane conatining p, q and the centers of B_1 and B_2, let q_1 and q_2 be the points of intersection of the sphere $\partial B(p, \mathrm{d}(p, q))$ with ∂B_1 and ∂B_2 respectively. Let q'_1 and q'_2 be the orthogonal projections of q onto the lines (p, q_1) and (p, q_2) respectively. Note that q'_1 belongs to the segment $[p, q_1]$ and q'_2 belongs to $[p, q_2]$. Thus, the segments $[q, q'_1]$ and $[q, q'_2]$ intersect ∂B_1 and ∂B_2 respectively. Since S separates B_1 and B_2, one of these segments (say $[q, q'_2]$) intersects S. Hence, $\mathrm{d}(q, S) \leq \mathrm{d}(q, q'_2)$.

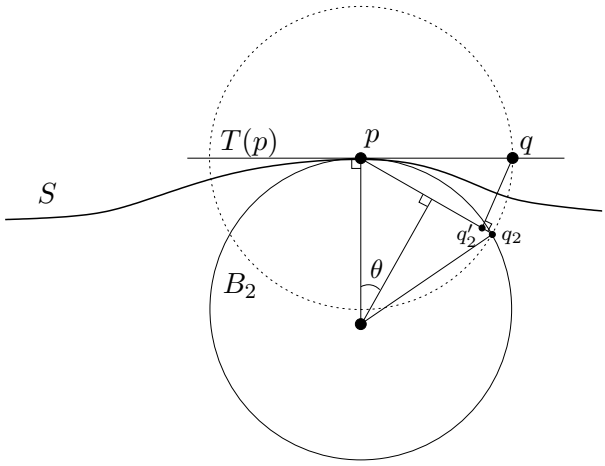

Figure 1.5: For the proof of Lemma 1.20.

Referring to Figure 1.5, we have: $\mathrm{d}(q, q_2') = \mathrm{d}(p, q)\ \sin\theta$. Moreover, $2\sin\theta\ \mathrm{rch}(S) = \mathrm{d}(p, q_2) = \mathrm{d}(p, q)$. Therefore, $\mathrm{d}(q, q_2') = \frac{\mathrm{d}^2(p,q)}{2\ \mathrm{rch}(S)}$. The result follows. □

Using Lemmas 1.19 and 1.20, we can bound $\mathrm{d}(p, S) = \mathrm{d}(p, \pi_S(p))$ for any $p \in \hat{S}$, which concludes the proof of Theorem 1.18:

Lemma 1.21 *$\forall p \in \hat{S}$, $\mathrm{d}(p, S)$ is at most $\frac{\varepsilon^2}{\mathrm{rch}(S)}$.*

Proof. Let $p \in \hat{S}$. Let f be a facet of $\hat{S}$ that contains p. We call $B(c, r)$ the surface Delaunay ball of f. Let p' be the orthogonal projection of p onto $T(c)$. By Lemma 1.19, for any vertex v of f, we have $\mathrm{d}(v, T(c)) \leq \frac{\varepsilon^2}{2\mathrm{rch}(S)}$. Since p is a convex combination of the vertices of f, $\mathrm{d}(p, p')$ is also bounded by $\frac{\varepsilon^2}{2\mathrm{rch}(S)}$. In addition, we have: $\mathrm{d}(p', c) \leq \mathrm{d}(p, c) \leq \varepsilon$. Thus, by Lemma 1.20, the distance from p' to S is at most $\frac{\varepsilon^2}{2\mathrm{rch}(S)}$. It follows that the distance from p to S is bounded by $\frac{\varepsilon^2}{\mathrm{rch}(S)}$. □

1.3 Loose ε-samples and ε-samples

As emphasized in the introduction of Part A, loose ε-samples and ε-samples are closely related. We will first prove that ε-samples are loose ε-samples, for sufficiently small ε. Then, we will prove that loose ε-samples are ε-samples asymptotically.

Trivially, an ε-sample satisfies Condition 1 of Definition A.1. Therefore, all we have to do is to prove that, for sufficiently small ε, Condition 2 is satisfied as well:

Lemma 1.22 *If E is an ε-sample of S, with $\varepsilon < \frac{\pi}{6+\pi}\ \mathrm{rch}(S) \approx 0.34\ \mathrm{rch}(S)$, then every point of E is a vertex of $\mathrm{Del}_{|S}(E)$. As a consequence, E is a loose ε-sample of S.*

Proof. Let $v \in E$. We consider the cell $\mathrm{V}(v)$ in the Voronoi diagram of E. We will prove that at least one edge of the boundary of $\mathrm{V}(v)$ intersects S, which means that v is incident to at least one facet of $\mathrm{Del}_{|S}(E)$.

Since E is an ε-sample of S, no point of $S \cap \mathrm{V}(v)$ is farther than ε from v. Thus, $S \cap \mathrm{V}(v)$ is included in $B(v, \varepsilon)$. Since $\varepsilon < \mathrm{rch}(S)$, the connected component of S on which v lies is not contained in $B(v, \varepsilon)$. Hence, this component intersects $\partial B(p, \varepsilon)$ and therefore $\partial \mathrm{V}(v)$. Let f be a facet of $\partial \mathrm{V}(v)$ that intersects S. We will show that $f \cap S$ contains no cycle, which implies that $S \cap \partial f \neq \emptyset$, hereby proving the lemma.

Let us assume for a contradiction that $f \cap S$ contains a cycle C. We call w the point of E such that $f = \mathrm{V}(v) \cap \mathrm{V}(w)$, and we consider a point p of $f \cap S$. Since E is an ε-sample, we have $\mathrm{d}(p, v) = \mathrm{d}(p, w) \leq \varepsilon$. Hence, v and w belong to $S \cap B(p, \varepsilon)$. According to the cocone lemma 1.4, the angle between $\mathbf{n}(p)$ and the line (v, w) is at least $\frac{\pi}{2} - \frac{\mu}{1-\mu}$, where $\mu = \frac{\varepsilon}{\mathrm{rch}(S)}$. Equivalently, the angle between $\mathbf{n}(p)$ and the plane $\mathrm{aff}(f)$ is at most $\frac{\mu}{1-\mu}$. Moreover, by the normal variation lemma 1.2, $(\mathbf{n}(p), \mathbf{n}(v))$ is at most $\frac{\mu}{1-\mu}$. Thus, the angle between $\mathbf{n}(v)$ and $\mathrm{aff}(f)$ is at most $\frac{2\mu}{1-\mu}$.

Let $\mathbf{l}$ be the unit vector of $\mathrm{aff}(f)$ whose angle with $\mathbf{n}(v)$ is minimal. We have $(\mathbf{n}(v), \mathbf{l}) \leq \frac{2\mu}{1-\mu}$. Let l be a line parallel to $\mathbf{l}$ that intersects C at least twice. Such a line exists because C is a cycle. Let q and q' be two points of intersection. We assume without loss of generality that $\mathbf{l} \cdot (\mathbf{q}' - \mathbf{q}) > 0$. Then, we have:

$$(\mathbf{n}(v),\ \mathbf{q}' - \mathbf{q}) \leq \frac{2\mu}{1-\mu} \tag{1.6}$$

Now, q and q' both belong to $\mathrm{V}(v)$, and hence to $B(v, \varepsilon)$. Thus, by the cocone lemma 1.4, the angle between $\mathbf{n}(v)$ and the line (q, q') is at least $\frac{\pi}{2} - \frac{\mu}{1-\mu}$. This contradicts (1.6) because $\frac{3\mu}{1-\mu} < \frac{\pi}{2}$, by hypothesis. Therefore, the lemma is proved. □

Conversely, Theorem 1.18 allows us to prove that loose ε-samples are asymptotically ε-samples.

Lemma 1.23 *If E is a loose ε-sample of S, with $\varepsilon \leq \mu_0\ \mathrm{rch}(S)$, then E is an $\varepsilon\sqrt{1 + \frac{\varepsilon^2}{\mathrm{rch}(S)^2}}$-sample of S.*

Proof. Let p be a point of S and p' a nearest neighbor of p on $\mathrm{Del}_{|S}(E)$. By Theorem 1.18, $\mathrm{d}(p, p')$ is at most $\frac{\varepsilon^2}{\mathrm{rch}(S)}$. If p' is a vertex of $\mathrm{Del}_{|S}(E)$, then $p' \in E$ and the result is proved. Else, let f be a facet of $\mathrm{Del}_{|S}(E)$ that contains p', and $v \neq p'$ a vertex of f closest to p'. By hypothesis, the circumradius of f is at most ε. Thus, $\mathrm{d}(p', v) \leq \varepsilon$, since $p' \in f$. If p' belongs to an edge e of f, then the lines $\mathrm{aff}(e)$ and (p, p') are perpendicular. Thus,

$$\mathrm{d}(p, E) \leq \mathrm{d}(p, v) = \sqrt{\mathrm{d}^2(p, p') + \mathrm{d}^2(p', v)} \leq \sqrt{\frac{\varepsilon^4}{\mathrm{rch}(S)^2} + \varepsilon^2}$$

and the lemma is proved. Else, p' belongs to the relative interior of f, thus the plane $\mathrm{aff}(f)$ is perpendicular to (p, p'). By the same computation as above, we get $\mathrm{d}(p, E) \leq \sqrt{\frac{\varepsilon^4}{\mathrm{rch}(S)^2} + \varepsilon^2}$, which proves the lemma. □

1.4 Covering

Note to the reader: *This section is not crucial for the coherence of the thesis, since its results will not be used in the sequel. Therefore, it can be skipped in a first reading.*

We assume in this section that the surface S is C^2-continuous, *i.e.* it is twice differentiable and its second derivative is continuous. Let $\mu_1 = \frac{\pi}{16+2\pi} \approx 0.14$. We will prove the following structural result:

Theorem 1.24 (Covering) *If E is a loose ε-sample of S, with $\varepsilon < \mu_1 \operatorname{rch}(S)$, then the surface Delaunay balls cover S.*

Let E be a loose ε-sample of S, with $\varepsilon < \mu_1 \operatorname{rch}(S)$. Let $\mu = \varepsilon/\operatorname{rch}(S) < \mu_1$. Before proving the theorem, we will study the geometric properties of the surface patches defined as the intersections of the surface Delaunay balls with S.

1.4.1 Pseudo-disks

Let f be a facet of $\operatorname{Del}_{|S}(E)$, $B_f = B(c_f, r_f)$ its surface Delaunay ball, $D_f = S \cap B_f$ its surface patch. Since E is a loose ε-sample of S, we have $r_f \leq \varepsilon = \mu \operatorname{rch}(S)$.

Lemma 1.25 (Terrain)
Since $\mu < \frac{\pi}{2+\pi}$, $S \cap B(c_f, \mu \operatorname{rch}(S))$ is a terrain over $T(c_f)$. Therefore, D_f is also a terrain over $T(c_f)$.

Proof. By the normal variation lemma 1.2, for any point $p \in B(c_f, \mu \operatorname{rch}(S))$, we have $(\mathbf{n}(c_f), \mathbf{n}(p)) \leq \frac{\mu}{1-\mu} < \frac{\pi}{2}$. Thus, by the line lemma 1.3, any line aligned with $\mathbf{n}(c_f)$ intersects $S \cap B(c_f, \mu \operatorname{rch}(S))$ at most once. Hence, $S \cap B(c_f, \mu \operatorname{rch}(S))$ is a terrain over the plane $T(c_f)$. □

Combining the cocone lemma 1.4 and the terrain lemma 1.25, we can show that D_f is a topological disk as soon as $\mu < \frac{\pi}{2+\pi}$. However, in the smooth case, a better bound on μ has been achieved by Boissonnat and Cazals [18], thus we simply recall their result here:

Lemma 1.26 (Topological Disk)
Since $\mu < 1$, D_f is a topological disk.

Definition 1.27 *Topological disks are pseudo-disks if they pairwise intersect along topological disks (that may be empty or reduced to a point) and if their boundaries pairwise intersect in at most two points.*

Observe that the boundaries of two pseudo-disks either do not intersect, or intersect in one point tangentially, or intersect in two points transversally. The main result of this section is the following one:

Proposition 1.28 *The surface patches of the facets of $\operatorname{Del}_{|S}(E)$ are pseudo-disks.*

Note to the reader: *The proof is organized in a hierarchical way, with several lemmas and claims with independent proofs. These proofs can be skipped in a first reading.*

Proof of Proposition 1.28. Let $B = B(c, r)$ and $B' = B(c', r')$ be two surface Delaunay balls. By the topological disk lemma 1.26, $D = B \cap S$ and $D' = B' \cap S$ are topological disks. Their boundaries C and C' are topological circles. Let us assume that balls B and B' intersect, the other case being trivial. Notice that none of them can be contained in the other one, since they are Delaunay balls. Thus, their

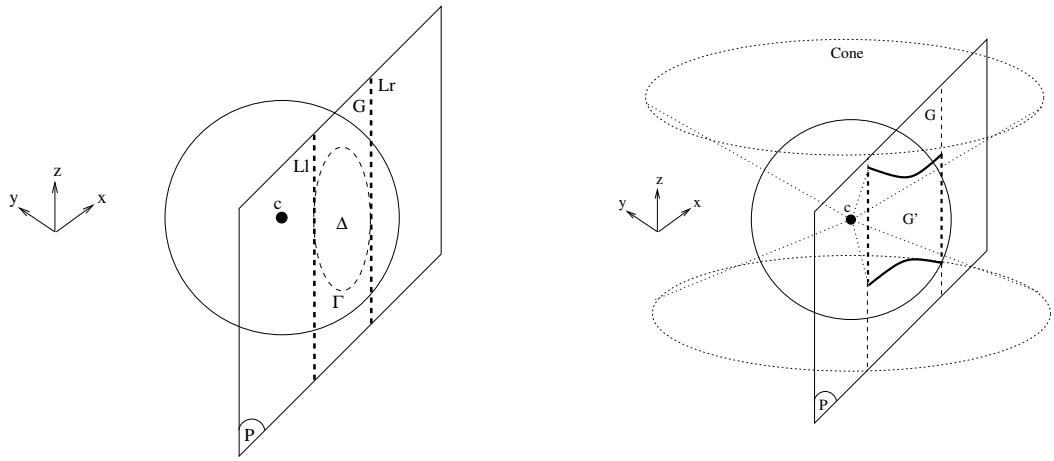

Figure 1.6: Definitions of G and G'

bounding spheres ∂B and $\partial B'$ also intersect. Let Γ be the circle $\partial B \cap \partial B'$, λ its radius ($\lambda < \min\{r, r'\}$) and P its supporting plane. We define $\Delta = B \cap P$ and notice that $\Gamma = \partial\Delta$. Since S is a closed surface, we have $C \subset \partial B$ and $C' \subset \partial B'$, which implies that:

$$C \cap C' \subseteq S \cap \Gamma \tag{1.7}$$

Inside the plane containing $\mathbf{n}(c)$ and line (c, c'), let $\mathbf{n}$ be the unit vector orthogonal to (c, c') that has a positive inner product with $\mathbf{n}(c)$. Since $\mathrm{d}(c, c') \leq r + r' \leq 2\mu \ \mathrm{rch}(S)$, the angle between $\mathbf{n}(c)$ and (c, c') is at least $\frac{\pi}{2} - \frac{2\mu}{1-2\mu}$, by the cocone lemma 1.4. Hence $(\mathbf{n}(c), \mathbf{n}) \leq \frac{2\mu}{1-2\mu}$.

Let $B^+ = B(c, 2r)$. Applying the normal variation lemma 1.2, we get:

$$\forall p \in S \cap B^+, \ (\mathbf{n}(p), \mathbf{n}) \leq (\mathbf{n}(p), \mathbf{n}(c)) + (\mathbf{n}(c), \mathbf{n}) \leq \frac{4\mu}{1-2\mu} < \frac{\pi}{4} \tag{1.8}$$

We endow $\mathbb{R}^3$ with an oriented orthonormal frame of origin c, of y-axis directed along **c'** $-$ **c**, and of z-axis directed along $\mathbf{n}$. We call L_l and L_r the two lines of P, parallel to the z-axis, that are tangent to Γ. The region of P bounded by L_l and L_r is called G – see Figure 1.6. In the sequel, ξ denotes $S \cap B^+ \cap G$.

Lemma 1.29 *ξ is a connected x-monotone arc.*

Proof. According to (1.8), we have $\forall p \in S \cap B^+$, $(\mathbf{n}(p), \mathbf{n}) < \frac{\pi}{4}$. Thus, by the line lemma 1.3, $S \cap B^+$ is xy-monotone, *i.e.* it is a terrain over the xy-plane in our frame. Hence, ξ is x-monotone. Moreover, again by the line lemma 1.3, $S \cap B^+$ lies outside the double cone of apex c, of vertical axis and of half-angle $\frac{\pi}{4}$. The equation of the double cone in our frame is $z^2 = x^2 + y^2$. It intersects P along two hyperbolic arcs of equations $z = \pm\sqrt{x^2 + d^2}$, where $d \leq r$ is the distance from c to P. Consider the subregion G' of G that is bounded vertically by the two hyperbolic arcs – see Figure 1.6. Since $S \cap B^+$ lies outside the double cone, ξ is included in G'.

The points of G' that are farthest from c are the points $(\pm\lambda, -d, \pm\sqrt{\lambda^2 + d^2})$. Their distance to c is

$$\sqrt{2(\lambda^2 + d^2)} < 2r$$

As a consequence, G' is included in the interior of B^+. It follows that ξ is also included in the interior of B^+ and therefore it cannot intersect ∂B^+. Its endpoints must then lie on the vertical lines L_r and L_l. But there can be only one endpoint per vertical line, since ξ is x-monotone. Hence, ξ has at most two endpoints and is thus connected. □

Lemma 1.30 $|S \cap \Gamma| \leq 2$.

Proof. Let us assume for a contradiction that $|S \cap \Gamma| > 2$. First, we show that there exists a point where the curvature of ξ is high and hence the distance to the medial axis of S is small. Then we work out a contradiction with the fact that E is a loose ε-sample, with $\varepsilon < \mu_1 \operatorname{rch}(S)$.

Claim 1.30.1 *There exists a point q at which the curvature of ξ is at least $\frac{1}{\lambda}$.*

Proof. We made the assumption that $|S \cap \Gamma| > 2$. Since $\Gamma \subset G$ and $\Gamma \subset B^+$, ξ also intersects Γ more than twice. And since ξ is connected by Lemma 1.29, there is a subarc $\overline{ab}$ of ξ that lies outside Δ and whose endpoints a and b lie on Γ. This subarc may be reduced to a point ($a = b$), since ξ may be tangent to Γ. But in this case, in the vicinity of a, ξ is locally included in Δ and tangent to Γ at a. Thus, its curvature at a is at least $\frac{1}{\lambda}$, which proves the claim with $q = a$. So now we assume that arc $\overline{ab}$ of ξ is not reduced to a point. Since ξ is x-monotone by Lemma 1.29, a and b lie on the same half of Γ, upper half or lower half (say upper half). Thus, the smaller arc of Γ that joins a and b is also x-monotone, and it lies below the arc $\overline{ab}$ of ξ. Therefore, there is a point q of the arc $\overline{ab}$ of ξ at which the curvature of ξ is at least the curvature of Γ, *i.e.* at least $\frac{1}{\lambda}$, which proves the claim. □

Claim 1.30.2 $\mathrm{d_M}(q) \leq \lambda\sqrt{2}$.

Proof. Let $\mathbf{n}_\xi(q)$ denote the normal to the planar curve ξ at point q. By (1.8), $\mathbf{n}(q)$ is not orthogonal to P, thus $\mathbf{n}_\xi(q)$ is oriented along the projection of $\mathbf{n}(q)$ onto P. Hence, we have $(\mathbf{n}(q), \mathbf{n}_\xi(q)) \leq (\mathbf{n}(q), \mathbf{n})$, which is less than $\frac{\pi}{4}$, by (1.8). According to Meusnier's Theorem [29, §3.2, Prop. 2], we then have at q:

$$I\!I(\xi', \xi') \geq \cos\frac{\pi}{4} \, \|\xi''\|$$

where $I\!I$ is the second fundamental form of S at q, ξ' is the unit tangent vector of ξ (parameterized by arclength) at q, and $\|\xi''\|$ is the curvature of ξ at q, which is more than $\frac{1}{\lambda}$ according to Claim 1.30.1. Thus, at q we have:

$$I\!I(\xi', \xi') \geq \frac{1}{\lambda\sqrt{2}} \tag{1.9}$$

Recall that $I\!I$ is a symmetric bilinear form, thus it can be diagonalized in an orthonormal frame, and its eigenvalues are the minimum and maximum curvatures of S at q. Let us call these values $\kappa_{\min}(q)$ and $\kappa_{\max}(q)$ respectively. Since ξ' is a unit vector, we have $I\!I(\xi', \xi') \leq \max\{|\kappa_{\min}(q)|, |\kappa_{\max}(q)|\}$. It follows, according to (1.9), that $\max\{|\kappa_{\min}(q)|, |\kappa_{\max}(q)|\} \geq \frac{1}{\lambda\sqrt{2}}$, or, equivalently, that the minimal radius of curvature of S at q is at most $\lambda\sqrt{2}$. The claim follows. □

The end of the proof of Lemma 1.30 is now immediate. The radius r of ball B is greater than λ, which is at least $\frac{1}{\sqrt{2}}\,\mathrm{d_M}(q) \geq \frac{1}{\sqrt{2}}\,\mathrm{rch}(S)$. This contradicts the assumption that B is a surface Delaunay ball of radius $r < \mu_1\,\mathrm{rch}(S)$, and hereby concludes the proof of Lemma 1.30. □

From Eq. (1.7) and Lemma 1.30, it follows immediately that $|C \cap C'| \leq 2$.

Lemma 1.31 *$S \cap \Delta$ is not reduced to two points.*

Proof. Let us assume that S intersects Δ in two points exactly, say a and b. Then, the subarc of ξ that joins points a and b lies outside Δ. It follows, by the same argument as in the proof of Claim 1.30.1, that there exists some point q of ξ at which the curvature of ξ is at least $\frac{1}{\lambda}$. Then, by Claim 1.30.2, $\mathrm{d_M}(q) \leq \lambda\sqrt{2}$, which leads to a contradiction, as in the end of the proof of Lemma 1.30. □

We can now complete the proof of Proposition 1.28, by showing that D and D' intersect each other either along a single point or along a topological disk. The result is clear if $D \subseteq D'$ or if $D' \subseteq D$. Otherwise, by Lemma 1.30, we have $|C \cap C'| \leq 2$. If $|C \cap C'| = 0$, then $D \cap D'$ is empty. If $|C \cap C'| = 1$, then $D \cap D'$ is a single point. If $|C \cap C'| = 2$, then $D \cap D'$ is either a topological disk or equal to $C \cap C'$. But if $D \cap D' = C \cap C'$, then $S \cap \Delta = C \cap C'$ since $C \cap C' \subseteq S \cap \Delta \subseteq D \cap D'$. This implies that $|S \cap \Delta| = 2$, which contradicts Lemma 1.31. Hence, $D \cap D'$ is not equal to $C \cap C'$ and is therefore a topological disk. This ends the proof of Proposition 1.28. □

1.4.2 Proof of the theorem

We can now prove the covering theorem 1.24. Let $\bigcup_{f \in \mathrm{Del}_{|S}(E)} B_f$ (or $\bigcup_f B_f$, for short) denote the union of the surface Delaunay balls.

Let f_0 be a facet of $\mathrm{Del}_{|S}(E)$. Our goal is to prove that the boundary C_{f_0} of the surface patch D_{f_0} of f_0 is included in the interior of $\bigcup_f B_f$. In fact, we will prove a slightly more precise result, stated as Lemma 1.32.

We consider the neighborhood $\mathrm{N}(f_0)$ of f_0 in $\mathrm{Del}_{|S}(E)$. By Theorem 1.7, $\mathrm{Del}_{|S}(E)$ is a manifold without boundary. Thus, $\mathrm{N}(f_0) \setminus \{f_0\}$ contains one facet of $\mathrm{Del}_{|S}(E)$ adjacent to f_0 through each edge of f_0. We call $R(f_0)$ the union of the surface patches associated with the facets of $\mathrm{N}(f_0)$. Let $\mathrm{int}(R(f_0))$ denote the relative interior of $R(f_0)$.

Lemma 1.32 *$C_{f_0} \subseteq \mathrm{int}(R(f_0))$.*

Proof. Let u, v and w be the vertices of f_0. We call f_{uv}, f_{vw} and f_{wu} the three facets of $\mathrm{N}(f_0)$ that are incident to f_0 through edges uv, vw and wu respectively. By Proposition 1.28, arcs $\overline{uv}$, $\overline{vw}$ and $\overline{wu}$ of C_{f_0} are included in $D_{f_{uv}}$, $D_{f_{vw}}$ and $D_{f_{wu}}$ respectively, and only their endpoints may possibly lie on $C_{f_{uv}}$, $C_{f_{vw}}$ or $C_{f_{wu}}$. Thus, the three arcs are included in the interior of $R(f_0)$, except for their endpoints which may possibly lie on the boundary of $R(f_0)$.

We claim that u, v and w also belong to $\mathrm{int}(R(f_0))$. Let U be the umbrella of facets of $\mathrm{Del}_{|S}(E)$ incident to u, and $R(u)$ be the union of surface patches of the facets of U. Notice that $R(u) \subseteq R(f_0)$, since $U \subseteq \mathrm{N}(f_0)$. Let $B(c,r)$ be the ball of largest radius, among the surface Delaunay balls of the

facets of U. For any point $p \in R(u)$, we have $\mathrm{d}(p,u) \leq 2r$, hence $R(u)$ is included in $S \cap B(u,2r)$. Since $2r \leq 2\varepsilon < 2\mu_1 \ \mathrm{rch}(S)$, $S \cap B(u,2r)$ is a terrain over $T(u)$, by the terrain lemma 1.25. It follows that the surface patches of the facets of U, which are pseudo-disks in $\mathbb{R}^3$, are also pseudo-disks in projection onto $T(u)$. Then, u lies inside the projection of $\mathrm{int}(R(u))$, since by Claim 1.11.1 u lies inside the projection of $\mathrm{int}(U)$. It follows that $u \in \mathrm{int}(R(u))$, since $R(u)$ is a terrain over $T(u)$. Thus, $u \in \mathrm{int}(R(f_0))$. The argument holds for v and w as well, which completes the proof of the lemma. □

Let $\bigcup_f D_f$ denote the union of the surface patches of the facets of $\mathrm{Del}_{|S}(E)$. A consequence of Lemma 1.32 is that the boundary of $\bigcup_f D_f$ is empty. Thus, $\bigcup_f D_f$ is topologically open and closed at the same time (for the topology of S induced by $\mathbb{R}^3$), which implies that $\bigcup_f D_f$ is the union of one or more connected components of S. Moreover, since $\mathrm{Del}_{|S}(E)$ has vertices on all the connected components of S, $\bigcup_f D_f$ intersects all the connected components of S. It follows that $\bigcup_f D_f = S$, which concludes the proof of the covering theorem 1.24.

1.5 Non-uniform samples

The previous sections deal exclusively with point samples whose density is driven by a uniform sizing field ε. Specifically, ε is an upper bound on the Hausdorff distance between the surface S and the point sample E. Therefore, the local density of E may be non-uniform, but it is everywhere lower-bounded by the same constant, which does not take into account the fact that some regions of S can be flat while others can be highly curved.

In [23], we focus on loose $\mu\mathrm{d}_{\mathrm{M}}$-samples and prove several properties similar to the ones stated in this chapter for loose ε-samples. We summarize our results in the next theorem:

Theorem 1.33 *If E is a loose $\mu\mathrm{d}_{\mathrm{M}}$-sample of S, with $\mu \leq 0.09$, then:*

- $\mathrm{Del}_{|S}(E)$ *is isotopic to S;*
- *the Hausdorff distance between* $\mathrm{Del}_{|S}(E)$ *and S is at most* $4.5\ \mu^2\ \sup\{\mathrm{d}_{\mathrm{M}}(p),\ p \in S\}$*;*
- $\mathrm{Del}_{|S}(E)$ *approximates S, in terms of normals and area, within an error of $O(\mu)$;*
- *S is covered by the surface Delaunay balls;*
- *E is a $\mu(1+8.5\mu)\mathrm{d}_{\mathrm{M}}$-sample of S.*

We refer the reader to [23] for comprehensive proofs of the above statements. In fact, the proofs of [23] are quite similar to the ones presented in this chapter, with the additional subtelty that the variation of d_{M} must be taken into account. Therefore, the bounds on μ in the above theorem are slightly deteriorated, compared to the bounds achieved in the previous sections of this chapter. Moreover, as we will see in Chapter 3, $\mu\mathrm{d}_{\mathrm{M}}$-samples and ε-samples with $\varepsilon = \mu\ \mathrm{rch}(S)$ have the same size, up to a constant factor that depends only on S.

Chapter 2

The Nonsmooth Case

In this chapter, we focus on the case where the surface S is nonsmooth. In order to make topological and geometric claims, we restrict our study to the class of Lipchitz surfaces, defined in Section 2.1. In Section 2.2, we review the main global properties of loose ε-samples. Specifically, we prove that their restricted Delaunay triangulation is a manifold without boundary (2.2.1), at Hausdorff distance $O_S(\varepsilon)$ from S (2.2.2) and isotopic to S (2.2.3). In Section 2.3, we show that loose ε-samples and ε-samples are closely related on Lipschitz surfaces.

2.1 Preliminaries

2.1.1 Lipschitz surfaces

Following [94, §1.3], we define Lipschitz surfaces as follows:

Definition 2.1 *Given a positive constant k, $S \subset \mathbb{R}^3$ is a k-Lipschitz surface if there exists some bounded open set $\Omega \subset \mathbb{R}^3$ such that $S = \partial\Omega$ and that Ω is locally the epigraph of a k-Lipschitz function, that is: there exist two positive constants η and ε, a finite number m of orthonormal systems of coordinates (x^i, y^i, z^i) of $\mathbb{R}^3$, and m k-Lipschitz maps $g^i : (x^i, y^i) \to \mathbb{R}$ defined over the open square $\Delta =]-\eta, \eta[^2$, such that*

$$\forall (x^i, y^i) \in \Delta, \begin{cases} \left]\left(x^i,\ y^i,\ g^i(x^i, y^i)\right),\ \left(x^i,\ y^i,\ g^i(x^i, y^i) + \varepsilon\right)\right[\subset \Omega \\ \left]\left(x^i,\ y^i,\ g^i(x^i, y^i) - \varepsilon\right),\ \left(x^i,\ y^i,\ g^i(x^i, y^i)\right)\right[\subset \mathbb{R}^3 \setminus \bar{\Omega} \end{cases}$$

and that any point $p \in S$ can be written in at least one of the coordinate systems under the form $p = (x^i_p,\ y^i_p,\ g^i(x^i_p, y^i_p))$.

Since the g^i are defined over open subsets of $\mathbb{R}^2$, around each point $p \in S$ there is a small neighborhood $\mathcal{N}(p) \subset \mathbb{R}^3$ such that $\Omega \cap \mathcal{N}(p)$ is the intersection of $\mathcal{N}(p)$ with the epigraph of some g^i . This property, illustrated in Figure 2.1, implies that our definition is equivalent to the one given by Clarke [47, §7.3], which states that $S = \partial\Omega$ is a k-Lipschitz surface if around each point $p \in S$ there is a neighborhood $\mathcal{N}(p) \subset \mathbb{R}^3$ such that $\Omega \cap \mathcal{N}(p)$ is the intersection of $\mathcal{N}(p)$ with the epigraph of a k-Lipschitz bivariate function.

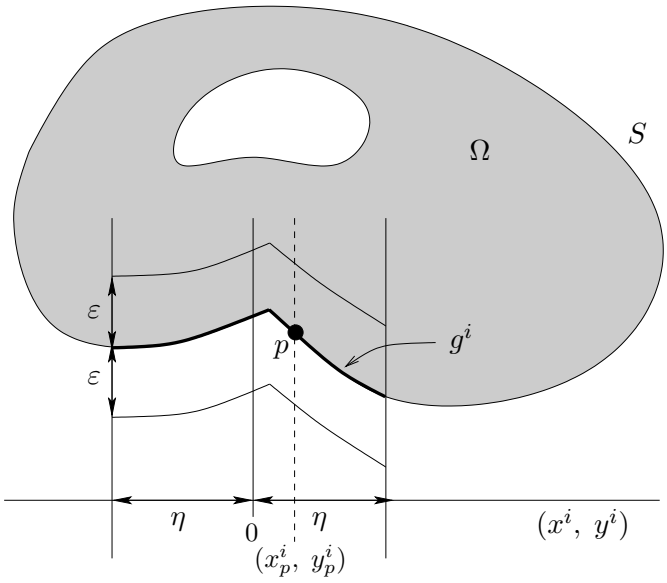

Figure 2.1: Being locally the epigraph of a Lipschitz function g^i (from [94, Fig. 2]).

Lipschitz surfaces are sometimes called *strongly Lipschitz surfaces* [13] in the literature, as opposed to *weakly Lipschitz surfaces*, which are not locally graphs of Lipschitz functions but whose local coordinate charts are Lipschitz homeomorphisms. As reported in [13, 93], both notions are identical when $k \geq 2$ but distinct when $k \leq 1$.

Let S be a k-Lipschitz surface, for some positive constant k to be specified later on. It follows from Definition 2.1 that S is an embedded compact C^0-continuous surface without boundary. In the sequel, we call $\mathcal{O}^-$ the open set Ω introduced in Definition 2.1, and $\mathcal{O}^+$ the open set $\mathbb{R}^3 \setminus \bar{\Omega}$. Moreover, we set $\mathcal{O} = \mathbb{R}^3 \setminus S$, which is equal to $\mathcal{O}^- \cup \mathcal{O}^+$.

By Rademacher's theorem [67, §3.1.6], each k-Lipschitz function $g^i : (x^i, y^i) \to \mathbb{R}$ of Definition 2.1 is differentiable $\mathcal{L}^2$-almost everywhere in Δ, *i.e.* the set Σ^i of singular points of g^i has zero Lebesgue measure in Δ. Since $\mathcal{L}^2$ coincides with the 2-dimensional Hausdorff measure $\mathcal{H}^2$ in $\mathbb{R}^2$ [67, §2.10.35], we have $\mathcal{H}^2(\Sigma^i) = 0$. And since g^i is Lipschitz, the map $(x^i,\ y^i) \mapsto (x^i,\ y^i,\ g^i(x^i, y^i))$ is also Lipschitz, hence it maps Σ^i onto a subset of $\mathbb{R}^3$ of zero $\mathcal{H}^2$ measure [67, §2.10.11]. Therefore, S is differentiable $\mathcal{H}^2$-almost everywhere. In particular, we have the following corollary, where $\tilde{S}$ denotes the set of points where S is differentiable:

Corollary 2.2 *$\tilde{S}$ is* dense *in S, that is: $\forall p \in S,\ \forall \eta > 0,\ \tilde{S} \cap B(p, \eta) \neq \emptyset$.*

Note that $S \setminus \tilde{S}$ may as well be dense in S, despite the fact that it has zero $\mathcal{H}^2$ measure.

Given $p \in \tilde{S}$, we call $T(p)$ the tangent plane of S at p, and $\mathbf{n}(p)$ the unit vector orthogonal to $T(p)$ that points towards $\mathcal{O}^+$. This vector is called the normal of S at p.

2.1.2 Lipschitz radius

Let S be a k-Lipschitz surface.

Definition 2.3 *Given a point $p \in S$, the k-Lipschitz radius of S at p, or $\mathrm{lr}_k(p)$ for short, is the maximum radius r such that $\mathcal{O}^- \cap B(p,r)$ is the intersection of $B(p,r)$ with the epigraph of a k-Lipschitz bivariate function.*

It is clear that $\mathrm{lr}_k(p)$ is positive for any $p \in S$, by Definition 2.1. Moreover, since S is compact without boundary, S is not the graph of a bivariate function. Therefore, $\mathrm{lr}_k(p)$ is finite, for any $p \in S$.

Lemma 2.4 lr_k *is 1-Lipschitz.*

Proof. Let p, q be two points of S. By definition of $\mathrm{lr}_k(p)$, for any $\eta > 0$, $\mathcal{O}^- \cap B(p, \mathrm{lr}_k(p) + \eta)$ is not the intersection of $B(p, \mathrm{lr}_k(p) + \eta)$ with the epigraph of a k-Lipschitz bivariate function. Now, $B(p, \mathrm{lr}_k(p) + \eta)$ is contained in the ball $B(q, \mathrm{d}(p,q) + \mathrm{lr}_k(p) + \eta)$. Thus, $\mathrm{lr}_k(q) \leq \mathrm{d}(p,q) + \mathrm{lr}_k(p) + \eta$. Since this is true for any $\eta > 0$, $\mathrm{lr}_k(q)$ is at most $\mathrm{d}(p,q) + \mathrm{lr}_k(p)$. □

It follows from Lemma 2.4 that lr_k is continuous over S. Since S is compact, lr_k reaches its minimum at some point $p_0 \in S$, with $\mathrm{lr}_k(p_0) > 0$. Hence,

Theorem 2.5 *The infimum of* lr_k *over S is positive.*

We call *k-Lipschitz radius of S* this infimum, noted $\mathrm{lr}_k(S)$. According to the above theorem, k-Lipschitz surfaces belong to the class of compact objects with positive k-Lipschitz radius. Conversely, every such object is a k-Lipschitz surface, by Clarke's definition. Like the distance to the medial axis in the smooth case, lr_k allows to predict the local behaviour of any k-Lipschitz surface, as we will see in Section 2.1.5.

2.1.3 The smooth case

Let S be a surface with positive reach. In this particular case, $\mathrm{lr}_k(S)$ is closely related to $\mathrm{rch}(S)$, for any positive value of k.

Theorem 2.6 *If S is a surface with positive reach, then S is k-Lipschitz for any $k > 0$, and we have* $\mathrm{lr}_k(S) \geq \frac{\arctan k}{1+\arctan k}\ \mathrm{rch}(S)$.

Proof. Let $\mu = \frac{\arctan k}{1+\arctan k}$. Given a point p of S, we call B_p the ball of center p and radius $\mu\ \mathrm{rch}(S)$. Let $D_p = S \cap B_p$. By the smooth cocone lemma 1.4, for any two points $q, q' \in D_p$, with $q \neq q'$, the line (q, q') lies outside the double cone of apex q, of axis aligned with $\mathbf{n}(p)$ and of half-angle $\frac{\pi}{2} - \frac{\mu}{1-\mu}$. Hence, the angle between (q, q') and the plane $T(p)$ is at most $\frac{\mu}{1-\mu} = \arctan k$. Since this is true for any pair of points of D_p, the surface patch D_p is the graph of a k-Lipschitz map from $T(p)$ to $\mathbb{R}$. Hence, $\mathrm{lr}_k(p) \geq \frac{\arctan k}{1+\arctan k}\ \mathrm{rch}(S)$. The theorem follows. □

Theorem 2.6 is illustrated in Figure 1.1: since S passes between two tangent balls of radius $\mathrm{rch}(S)$ at point $f(s)$, any Lipschitz coefficient can be achieved in a sufficiently small neighborhood of $f(s)$.

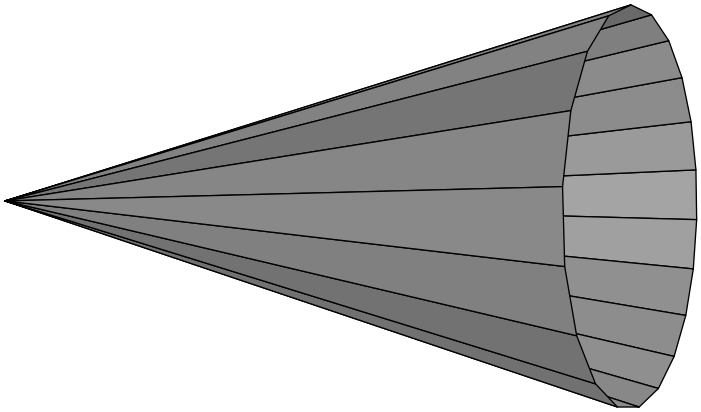

Figure 2.2: Bounded dihedral angles do not ensure to be Lipschitz.

2.1.4 The polyhedral case

As we have just seen, the class of k-Lipschitz surfaces includes all surfaces with positive reach. But it is in fact much broader than that. Another interesting subclass is the set of sufficiently *smooth* polyhedra, as defined below:

Definition 2.7 *Let $\theta < \frac{\pi}{2}$. An oriented polyhedron S is* θ-smooth *if, for any two non-disjoint facets f and f' of S, we have $(\boldsymbol{n}(f), \boldsymbol{n}(f')) \leq \theta$.*

Note that bounding the dihedral angles between adjacent facets of S is not sufficient to guarantee that S is Lipschitz. In Figure 2.2 for instance, by moving the apex of the cone away, on can make the cone as narrow as needed, while the dihedral angles between adjacent facets remain lower-bounded. This is why the bound in Definition 2.7 applies to all pairs of facets sharing at least one vertex.

We will prove that θ-smooth polyhedra are $O(\theta)$-Lipschitz surfaces. Let us start with a technical result:

Lemma 2.8 *Let $\theta < \frac{\pi}{2}$. Let S be an oriented polyhedron without boundary, p a point of S, B a ball centered at p, and $\boldsymbol{n}$ a unit vector such that, for any facet f of S intersecting B, $(\boldsymbol{n}(f), \boldsymbol{n}) \leq \theta$. Then, $S \cap B$ is the graph of a $\tan\theta$-Lipschitz function defined over the orthogonal projection of $S \cap B$ onto the plane orthogonal to $\boldsymbol{n}$.*

Proof. Let l be a line intersecting $S \cap B$ at least twice, and such that the angle between $\mathbf{n}$ and l is less than $\frac{\pi}{2} - \theta$. We call $\mathbf{l}$ the unit vector aligned with l such that $\mathbf{l} \cdot \mathbf{n} \geq 0$. We have

$$(\mathbf{l}, \mathbf{n}) < \frac{\pi}{2} - \theta \tag{2.1}$$

For any facet f of S intersecting B, $(\mathbf{n}(f), \mathbf{n})$ is at most θ, hence the angle between $\mathbf{n}$ and the supporting plane of f is at least $\frac{\pi}{2} - \theta$. As a consequence, l intersects $\mathrm{aff}(f)$ transversally. Up to a small perturbation of l, we can assume without loss of generality that two consecutive points of intersection, say p_1 and p_2, of l with $S \cap B$ belong to the relative interiors of two facets, say f_1 and f_2, of S. Since B is convex, p_1 and p_2 are also consecutive points of intersection of l with S. Therefore, $\mathbf{n}(f_1) \cdot \mathbf{l}$ and $\mathbf{n}(f_2) \cdot \mathbf{l}$ have

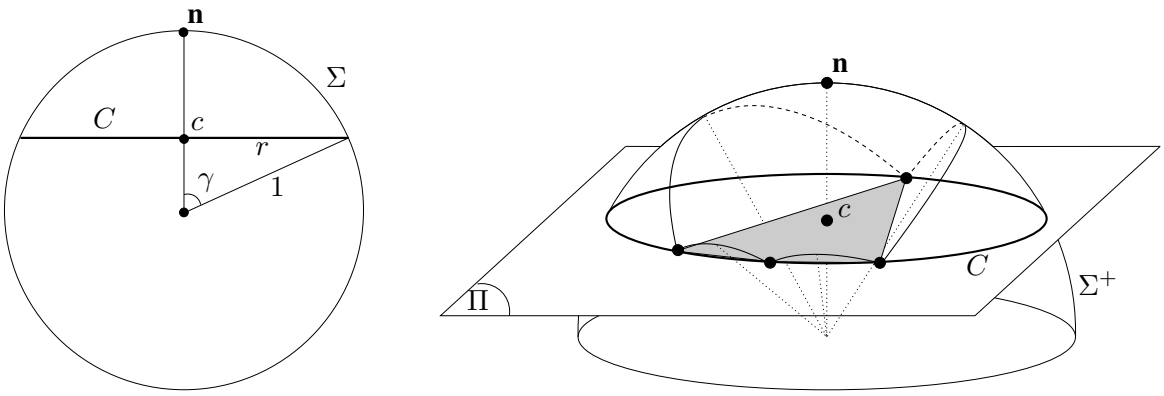

Figure 2.3: For the proof of Theorem 2.9.

different signs, say $\mathbf{n}(f_1) \cdot \mathbf{l} \leq 0$. It follows by Eq. (2.1) that $(\mathbf{n}(f_1), \mathbf{n}) > \theta$, which contradicts the assumption of the lemma. Hence, every line that intersects $S \cap B$ more than once makes an angle of at least $\frac{\pi}{2} - \theta$ with $\mathbf{n}$. As a consequence, for any point $q \in S \cap B$, $S \cap B$ lies outside the double cone of apex q, of axis aligned with $\mathbf{n}$, and of half-angle $\frac{\pi}{2} - \theta$. The result follows. □

From Lemma 2.8 we deduce the following

Theorem 2.9 *Any θ-smooth polyhedron is a k-Lipschitz surface, with $k = \frac{2\sin\theta/2}{\sqrt{3-4\sin^2\theta/2}}$.*

Notice that, as θ goes to zero, $\frac{2\sin\theta/2}{\sqrt{3-4\sin^2\theta/2}}$ is equivalent to $\frac{\theta}{\sqrt{3}}$, which means that θ-smooth polyhedra are asymptotically $\frac{\theta}{\sqrt{3}}$-Lipschitz surfaces.

Proof of the theorem. Let S be a θ-smooth polyhedron. Let $p \in S$. We call $f_1, \cdots, f_n$ the facets of S that contain p. Let $r(p) = \mathrm{d}(p, S \setminus \bigcup_i f_i) > 0$. We call B_p the open ball $B(p, r(p))$. By Definition 2.7, we have $\forall i$, $(\mathbf{n}(f_1), \mathbf{n}(f_i)) \leq \theta$. Hence, by Lemma 2.8, we know that S is a $\tan\theta$-Lipschitz surface. However, by a more thorough analysis we can improve on the constant.

We represent the unit normal vectors $\mathbf{n}(f_i)$ as points on the unit sphere Σ. Let D be the smallest disk of Σ that encloses the set $\{\mathbf{n}(f_i)\}$. We call $\mathbf{n}$ its center on Σ, and γ its spherical radius, *i.e.* the angle such that $\forall \mathbf{n}' \in D,\ (\mathbf{n}, \mathbf{n}') \leq \gamma$ – see Figure 2.3 (left). We will bound γ with respect to θ. By hypothesis, the set $\{\mathbf{n}(f_i)\}$ is included in the disk of center $\mathbf{n}(f_1)$ and spherical radius θ. As a consequence, we have $\gamma \leq \theta < \frac{\pi}{2}$, which implies that D is included in the disk of center $\mathbf{n}$ and spherical radius $\frac{\pi}{2}$, which is a hemisphere of Σ, called Σ^+ – see Figure 2.3 (right).

The bounding circle C of D is the intersection of Σ^+ with some Euclidean plane Π. Let c and r denote respectively the center and radius of C in Π. Referring to Figure 2.3 (left), we have:

$$r = \sin\gamma \tag{2.2}$$

Since D is the smallest disk of Σ^+ enclosing $\{\mathbf{n}(f_i)\}$, C contains at least two elements of $\{\mathbf{n}(f_i)\}$, and the radial projection of $\mathbf{n}$ onto Π (namely, c) belongs to the convex hull of $C \cap \{\mathbf{n}(f_i)\}$ – see Figure 2.3 (right).

– If C contains exactly two elements of $\{\mathbf{n}(f_i)\}$, then these elements are diametrally placed on C. It follows that the diameter of $C \cap \{\mathbf{n}(f_i)\}$ is $2r$.

– Else, C contains three or more elements of $\{\mathbf{n}(f_i)\}$. Let us triangulate $C \cap \{\mathbf{n}(f_i)\}$ in Π. At least one triangle t contains c. Since $c \in t$ is the circumcenter of t, the longest edge of t is longer than $r\sqrt{3}$. Hence, the diameter of $C \cap \{\mathbf{n}(f_i)\}$ is at least $r\sqrt{3}$.

In both cases, the diameter d of $C \cap \{\mathbf{n}(f_i)\}$ verifies:

$$d \geq r\sqrt{3} \tag{2.3}$$

Let $\mathbf{n}(f_j)$ and $\mathbf{n}(f_k)$ be two elements of $C \cap \{\mathbf{n}(f_i)\}$ whose distance in Π is d. From (2.2) and (2.3) we deduce:

$$\sin\gamma\ \sqrt{3} = r\ \sqrt{3} \leq d = 2\ \sin\frac{(\mathbf{n}(f_j), \mathbf{n}(f_k))}{2} \leq 2\ \sin\frac{\theta}{2}$$

Hence, for any facet f_i, $(\mathbf{n}(f_i), \mathbf{n})$ is bounded by $\gamma \leq \arcsin\left(\frac{2}{\sqrt{3}}\ \sin\frac{\theta}{2}\right)$. It follows by Lemma 2.8 that $S \cap B_p$ is the graph of a $\tan\gamma$-Lipschitz function, with $\tan\gamma \leq \frac{2\sin\theta/2}{\sqrt{3-4\sin^2\theta/2}}$. This concludes the proof of Theorem 2.9. □

2.1.5 The general case

From now on, and for the rest of the chapter, S is a general k-Lipschitz surface, for some fixed k. S may not necessarily be smooth or polyhedral. For convenience, we define $\theta = 2\arctan k \in [0, \pi[$. In addition to being positive and 1-Lipschitz, lr_k allows to predict the local behaviour of S.

Let $p \in S$, $D_p = S \cap B(p, \mathrm{lr}_k(p))$ and $\tilde{D}_p = \tilde{S} \cap B(p, \mathrm{lr}_k(p))$, where $\tilde{S}$ is the set of regular points of S. Note that $\tilde{D}_p$ is not empty because $\tilde{S}$ is dense in S, by Corollary 2.2. By definition of $\mathrm{lr}_k(p)$, D_p is the intersection of $B(p, \mathrm{lr}_k(p))$ with the graph of a k-Lipschitz bivariate function f_p. We choose a frame $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that f_p is defined over the xy-plane, and we orient $\mathbf{z}$ such that at every point of $\tilde{D}_p$ the normal of S has a non-negative inner product with $\mathbf{z}$. This is possible because D_p is the graph of a bivariate function.

Lemma 2.10 *For any $q \in D_p$, D_p lies outside the double cone of apex q, of axis aligned with the z-axis and of half-angle $\frac{\pi}{2} - \frac{\theta}{2}$. Moreover, if $q \in \tilde{D}_p$, then the angle $(\boldsymbol{n}(q), \boldsymbol{z})$ is at most $\frac{\theta}{2}$.*

Proof. Let $q \in D_p$ and $q' \in D_p \setminus \{q\}$. We call $\bar{q}$ and $\bar{q}'$ their orthogonal projections onto the xy-plane. The angle α between the line (q, q') and the xy-plane is given by:

$$\tan\alpha = \frac{|f_p(\bar{q}) - f_p(\bar{q}')|}{\mathrm{d}(\bar{q}, \bar{q}')} \leq k = \tan\frac{\theta}{2} \tag{2.4}$$

Hence, we have $\alpha \leq \frac{\theta}{2}$, which means that q' lies outside the double cone of apex q, of axis aligned with $\mathbf{z}$ and of half-angle $\frac{\pi}{2} - \frac{\theta}{2}$.

Let us now assume that $q \in \tilde{D}_p$. Eq. (2.4) holds for any $q' \in D_p \setminus \{q\}$. In particular, as q' approaches q, the angle between line (q, q') and the xy-plane remains bounded by $\frac{\theta}{2}$. As a consequence, the angle between the tangent plane $T(q)$ and the xy-plane is at most $\frac{\theta}{2}$. Since $\mathbf{z}$ is oriented such that $\mathbf{z} \cdot \mathbf{n}(q) \geq 0$, we get: $(\mathbf{n}(q), \mathbf{z}) \leq \frac{\theta}{2}$, which concludes the proof of the lemma. □

The above result has two straightforward corollaries:

Lemma 2.11 (Normal Variation)
For any two points q, q' of $\tilde{D}_p$, the angle between $\boldsymbol{n}(q)$ and $\boldsymbol{n}(q')$ is at most θ.

Lemma 2.12 (Cocone)
For any $q \in \tilde{D}_p$ and any $q' \in D_p$, D_p lies outside the double cone of apex q', of axis aligned with $\boldsymbol{n}(q)$ and of half-angle $\frac{\pi}{2} - \theta$.

Note that Lemma 2.12 makes sense only for $\theta \leq \frac{\pi}{2}$, *i.e.* for $k \leq 1$. Let now $f = (u, v, w)$ be a triangle whose vertices u, v and w belong to D_p. Assume that the radius-edge ratio of f is at most ϱ, for some constant $\varrho \geq 1$. By an easy computation, every inner angle of f is at least $\beta = \arcsin \frac{1}{2\varrho}$.

Lemma 2.13 (Triangle Normal)
With the above notations, for any point $q \in \tilde{D}_p$, the angle α between $\boldsymbol{n}(q)$ and the line orthogonal to the plane aff(u, v, w) verifies $\sin\alpha \leq 2\varrho \sin\theta$. Therefore, if $\varrho \leq \frac{1}{2\sin\theta}$, then $\alpha \leq \arcsin(2\varrho \sin\theta)$.

Proof. We assume without loss of generality that u is a vertex of f with largest inner angle. By the cocone lemma 2.12, v and w lie outside the double cone $K(u)$ of apex u, of axis aligned with $\mathbf{n}(q)$ and of half-angle $\frac{\pi}{2} - \theta$. Since aff(u, v, w) passes through the apex of $K(u)$, it intersects $K(u)$ either along a single point, or along a single line, or along a double wedge. If $K(u) \cap$ aff(u, v, w) is a single point or a single line, then $\alpha \leq \theta$, which implies that $\sin\alpha < 2\varrho \sin\theta$, since $\varrho \geq 1$.

If $K(u) \cap$ aff(u, v, w) is a double wedge $K'(u)$, then the half-angle θ' of this double wedge depends on α and θ. We endow $\mathbb{R}^3$ with an oriented orthonormal frame $(u, \mathbf{x}, \mathbf{y}, \mathbf{z})$, such that the z-axis is aligned with $\mathbf{n}(q)$ and that the line of intersection between aff(u, v, w) and the xy-plane is aligned with the y-axis. In this frame, the equation of the boundary of $K(u)$ is $z^2 = \tan^2\theta\ (x^2 + y^2)$, and the equation of aff(u, v, w) is $z = x\ \tan\alpha$. Thus, inside aff(u, v, w) (which we endow with an oriented orthonormal frame $(u, \mathbf{X}, \mathbf{y})$), the equation of the boundary of $K'(u)$ is $y = \pm\frac{1}{\sin\theta}\sqrt{\sin^2\alpha - \sin^2\theta}\ X$. Hence, the half-angle of $K'(u)$ is

$$\theta' = \arctan\left(\frac{1}{\sin\theta}\sqrt{\sin^2\alpha - \sin^2\theta}\right). \tag{2.5}$$

Since v and w lie outside $K(u)$, inside aff(u, v, w) they do not belong to $K'(u)$. Now, aff$(u, v, w) \setminus K'(u)$ is a double wedge delimited by the lines of the boundary of $K'(u)$.
- If v and w lie inside the same wedge, then $\hat{u}$ is at most $\pi - 2\theta'$. Since u is the vertex of f with largest inner angle, we have $\hat{u} \geq \frac{\pi}{3}$, which implies that $\theta' \leq \frac{\pi}{3}$. Since $\varrho \geq 1$, we have $\beta \leq \frac{\pi}{6}$, hence $\frac{\pi}{2} - \beta$ is at least $\frac{\pi}{3}$.
- If v and w do not lie in the same wedge, then $\hat{u}$ is at least $2\theta'$. Since all the angles of the triangle are at least β, we have $\hat{u} \leq \pi - 2\beta$, which implies that $\theta' \leq \frac{\pi}{2} - \beta$.
- In both cases, θ' is at most $\frac{\pi}{2} - \beta$, which implies that $\sin\alpha \leq \frac{\sin\theta}{\sin\beta}$, by (2.5). The lemma follows, since $\beta = \arcsin\frac{1}{2\varrho}$. □

2.1.6 Weak feature size

Chazal and Lieutier [35] have introduced the notion of *weak feature size*, which specifies how well the topology of the surface can be captured from a Hausdorff approximation. We will show that, when

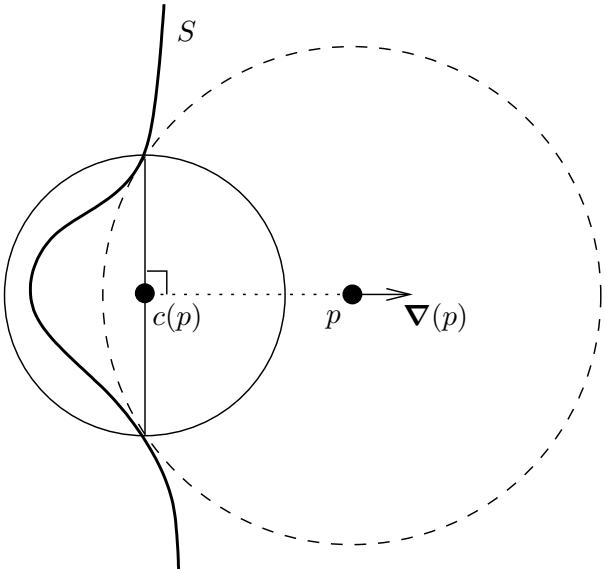

Figure 2.4: Definition of $\boldsymbol{\nabla}(p)$.

the surface S is k-Lipschitz, lr_k is closely related to the weak feature size. Let d_S denote the function *distance to* S:

$$\forall p \in \mathcal{O},\ \mathrm{d}_S(p) \ =\ \min\{\mathrm{d}(p,q) \mid q \in S\}$$

In [85, §3.1], the author derives from d_S a vector field $\boldsymbol{\nabla} :\ \mathcal{O} \to \mathbb{R}^3$ defined as follows:

$$\forall p \in \mathcal{O},\ \boldsymbol{\nabla}(p) \ =\ \frac{1}{\mathrm{d}_S(p)}\ (\mathbf{p} - \mathbf{c}(p)) \tag{2.6}$$

where $c(p)$ is the center of the smallest ball $B(c(p), r(p))$ that contains all the points of S that are closest to p. The definition of $\boldsymbol{\nabla}(p)$ is illustrated in Figure 2.4. The following result, introduced in [85, §5.1], will be useful:

$$\forall p \in \mathcal{O},\ \|\boldsymbol{\nabla}(p)\|^2 \ =\ 1 - \frac{r(p)^2}{\mathrm{d}_S(p)^2} \tag{2.7}$$

In particular, if p does not belong to the medial axis of S, then $c(p)$ is the only nearest neighbor of p on S. In this case, we have $r(p) = 0$ and the norm of $\boldsymbol{\nabla}(p)$ is 1.

A point $p \in \mathcal{O}$ is a critical point of $\boldsymbol{\nabla}$ if $\|\boldsymbol{\nabla}(p)\| = 0$. As emphasized in [35], p is critical iff it belongs to the convex hull of its nearest neighbors on S. We call *weak feature size of* $\mathcal{O}$, or simply $\mathrm{wfs}(\mathcal{O})$, the minimal distance from S to the set Φ of critical points of $\boldsymbol{\nabla}$:

$$\mathrm{wfs}(\mathcal{O}) \ =\ \min\{\mathrm{d}(p,q) \mid p \in S,\ q \in \Phi\}$$

The main topological properties of the weak feature size are reviewed in [36]. In the sequel, we assume that $k < 1$, *i.e.* $\theta < \frac{\pi}{2}$. We will show that, if S is k-Lipschitz, then the weak feature size of $\mathcal{O}$ is positive. We start with two technical results, stated as Lemmas 2.14 and 2.15, which control the orientation and norm of the vector field $\boldsymbol{\nabla}$ in the vicinity of S.

Let $p \in S$, $B_p^1 = B(p, \mathrm{lr}_k(p))$ and $B_p^{1/2} = B(p, \frac{1}{2}\ \mathrm{lr}_k(p))$. Let also $q \in \mathcal{O} \cap B_p^{1/2}$ and $p' \in \tilde{S} \cap B_p^{1/2}$. Such a point exists since $\tilde{S}$ is dense in S, by Corollary 2.2. We call $q_1, \cdots, q_k$ the nearest neighbors of q on S.

Lemma 2.14 *With the above notations, we have:*
If $q \in \mathcal{O}^+$, *then* $\forall i, \quad (\boldsymbol{q} - \boldsymbol{q}_i,\ \boldsymbol{n}(p')) \leq \theta$
If $q \in \mathcal{O}^-$, *then* $\forall i, \quad (\boldsymbol{q} - \boldsymbol{q}_i,\ -\boldsymbol{n}(p')) \leq \theta$

Proof. We deal with the case $q \in \mathcal{O}^+$, the case $q \in \mathcal{O}^-$ being symmetric. Since $p \in S$, we have $\forall i$, $\mathrm{d}(p, q) \geq \mathrm{d}(q_i, q)$. It follows that the distance from p to the q_i is at most $2\ \mathrm{d}(p, q) < \mathrm{lr}_k(p)$. Hence, the q_i lie in B_p^1.

If q_i belongs to $\tilde{S}$, then we have $(\mathbf{q} - \mathbf{q}_i,\ \mathbf{n}(p')) = (\mathbf{n}(q_i), \mathbf{n}(p'))$, since the normals of S point towards $\mathcal{O}^+$. Moreover, the fact that q_i belongs to B_p^1 implies that $(\mathbf{n}(q_i), \mathbf{n}(p')) \leq \theta$, by the normal variation lemma 2.11. Hence, $(\mathbf{q} - \mathbf{q}_i,\ \mathbf{n}(p')) \leq \theta$.

Consider now the case where q_i does not belong to $\tilde{S}$. By the cocone lemma 2.12, $S \cap B_p^1$ lies outside the double cone $K(q_i)$ of apex q_i, of axis aligned with $\mathbf{n}(p')$ and of half-angle $\frac{\pi}{2} - \theta$. Let $K^+(q_i)$ be the cone of $K(q_i)$ such that $K^+(q_i) \cap B_p^1$ lies on the same side of S as q. Let $K^-(q_i)$ be the other cone of $K(q_i)$, and let q' be a point of $K^-(q_i)$ closest to q. We claim that $q' = q_i$. Indeed, by assumption, q lies in $\mathcal{O}^+$ whereas $K^-(q_i) \setminus \{q_i\}$ is included in $\mathcal{O}^-$. Thus, if $q' \neq q_i$, the open line segment $]q, q'[$ intersects S. For any $q'' \in S \cap]q, q'[$, we then have $\mathrm{d}(q, q'') < \mathrm{d}(q, q') \leq \mathrm{d}(q, q_i)$, which contradicts the fact that q_i is a nearest neighbor of q on S. Hence, $q' = q_i$, which means that q_i is the point of $K^-(q_i)$ closest to q. It follows that q lies in the cone of apex q_i, of axis aligned with $\mathbf{n}(p')$, and of half-angle θ. Hence, $(\mathbf{q} - \mathbf{q}_i,\ \mathbf{n}(p')) \leq \theta$, which proves the lemma. □

The next result is an immediate consequence of Lemma 2.14 and of the fact that $\boldsymbol{\nabla}(q)$ is a convex combination of the $(\mathbf{q} - \mathbf{q}_i)$.

Lemma 2.15 *With the same notations as above, we have* $\|\boldsymbol{\nabla}(q)\| \geq \cos\theta > 0$. *Moreover,*
If $q \in \mathcal{O}^+$, *then* $(\boldsymbol{\nabla}(q), \boldsymbol{n}(p')) \leq \theta$
If $q \in \mathcal{O}^-$, *then* $(\boldsymbol{\nabla}(q), -\boldsymbol{n}(p')) \leq \theta$

Proof. We assume without loss of generality that $q \in \mathcal{O}^+$, the case $q \in \mathcal{O}^-$ being symmetric. By Lemma 2.14, we have

$$\forall i,\ (\mathbf{q} - \mathbf{q}_i) \cdot \mathbf{n}(p') \geq \|\mathbf{q} - \mathbf{q}_i\| \cos\theta = \mathrm{d}_S(q) \cos\theta$$

Since $c(q)$ is the center of the smallest ball containing the q_i, $c(q)$ lies in the convex hull of the q_i. Hence, $(\mathbf{q} - \mathbf{c}(q)) \cdot \mathbf{n}(p')$ is a convex combination of the $(\mathbf{q} - \mathbf{q}_i) \cdot \mathbf{n}(p')$, which are all at least $\mathrm{d}_S(q)\ \cos\theta$. Using Eq. (2.6), we get

$$\|\boldsymbol{\nabla}(q)\| = \frac{\|\mathbf{q} - \mathbf{c}(q)\|}{\mathrm{d}_S(q)} \geq \frac{(\mathbf{q} - \mathbf{c}(q)) \cdot \mathbf{n}(p')}{\mathrm{d}_S(q)} \geq \cos\theta$$

In addition, by Lemma 2.14 the q_i lie in the cone of apex q, of axis aligned with $-\mathbf{n}(p)$ and of half-angle θ. Since this cone is convex, it contains $c(q)$, which is a convex combination of the q_i. Hence, $(\boldsymbol{\nabla}(q), \mathbf{n}(p')) \leq \theta$, which ends the proof of the lemma. □

Our main result is an easy corollary of Lemma 2.15.

Theorem 2.16 *If S is a k-Lipschitz surface, for some $k < 1$, then $\mathcal{O} = \mathbb{R}^3 \setminus S$ has a positive weak feature size. More precisely,*

$$\forall p \in S,\ \mathrm{lr}_k(p) \leq 2\ \mathrm{d}(p, \Phi) \tag{2.8}$$

which implies that wfs$(\mathcal{O}) \geq \frac{1}{2}\ \mathrm{lr}_k(S)$.

Proof. Let $p \in S$ and $p \in \tilde{S} \cap B(p, \frac{1}{2}\ \mathrm{lr}_k(p))$. By Lemma 2.15, for any $q \in \mathcal{O} \cap B(p, \frac{1}{2}\ \mathrm{lr}_k(p))$, we have $\|\boldsymbol{\nabla}(q)\| \neq 0$. Hence, $B(p, \frac{1}{2}\ \mathrm{lr}_k(p)) \cap \Phi = \emptyset$, which implies that $\mathrm{d}(p, \Phi) \geq \frac{1}{2}\ \mathrm{lr}_k(p)$. □

Before proceeding to the next section, we would like to emphasize on an elementary but useful result:

Lemma 2.17 *The weak feature size of $\mathcal{O}$ is at most half the minimal distance between any two connected components of S.*

Proof. Consider the smallest segment $[p, p']$ whose vertices lie on different connected components of S. The interior of the diametral ball of $[p, p']$ cannot intersect S, since otherwise there would be a smaller segment satisfying the hypothesis. As a consequence, the midpoint q of $[p, p']$ is the center of the smallest ball containing its nearest neighbors on S (including p and p'). Therefore, $c(q) = q$, which implies by Eq. (2.6) that $\|\boldsymbol{\nabla}(q)\| = 0$, or equivalently, that $q \in \Phi$. As a consequence, $\mathrm{wfs}(\mathcal{O}) \leq \mathrm{d}(p, q) = \frac{\mathrm{d}(p,p')}{2}$. □

This lemma, combined with Theorem 2.16, yields the following corollary, which will be useful in the next section:

Corollary 2.18 *A triangle of circumradius less than $\frac{1}{2}\ \mathrm{lr}_k(S)$ has its three vertices on the same connected component of S.*

2.2 Global properties

Throughout the section, S is a k-Lipschitz surface. For convenience, we define $\theta = 2 \arctan k$. Since S is fixed, k and θ are fixed constants. Moreover, we set $\nu = 0.52$, as in Section 1.2. In the sequel, we consider point samples E of S that verify the following conditions:

H1 E is a loose ε-sample of S, with $\varepsilon < \frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S)$.

H2 The points of E are farther than $\frac{\varepsilon}{\varrho}$ from one another, where $\varrho \geq 1$ verifies:

H2.0 $\varrho < \dfrac{\sin\left(\frac{\pi}{3} - \theta\right)}{2\sin\theta}$

H2.1 $\varrho < \dfrac{\cos 3\theta}{2\sin\theta}$

H2.2 $\varrho < \dfrac{\sin\left(\arccos\left(\frac{2\sin\theta}{\nu}\right) - \theta\right)}{2\sin\theta}$

H2.3 $\varrho < \dfrac{\cos 2\theta}{2(1+\nu)\sin\theta}$

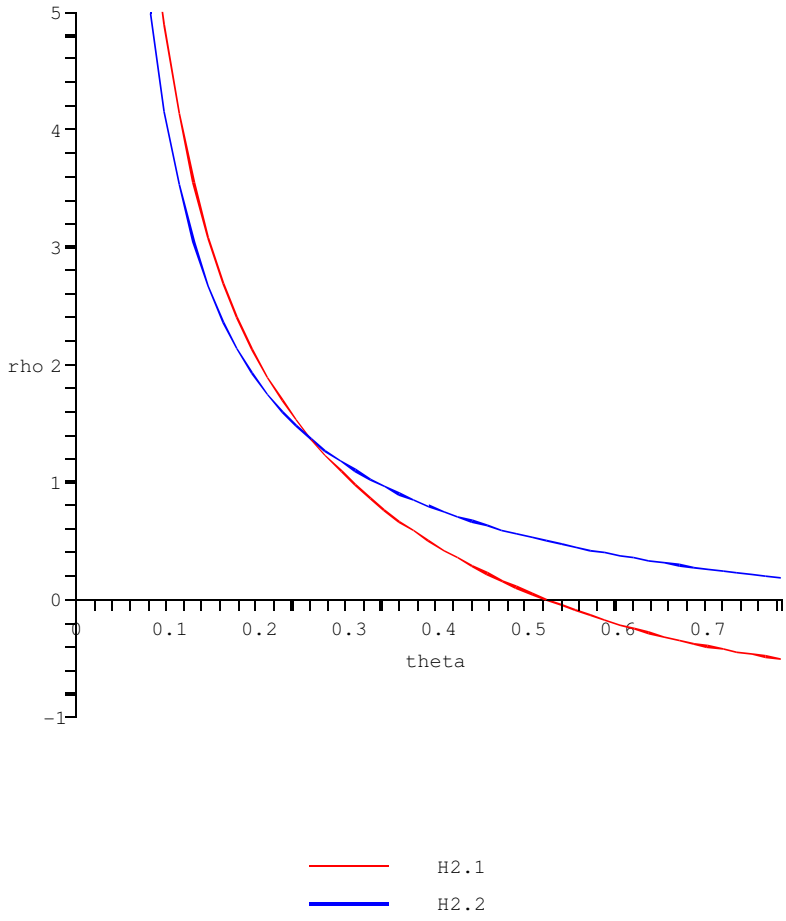

Figure 2.5: Conditions on ϱ: H2.0 (red), H2.1 (magenta), H2.2 (blue), H2.3 (cyan).

Let us give a few explanations about the meaning of the above conditions. H1 imposes that E be dense with respect to $\mathrm{lr}_k(S)$. H2 imposes that E be sparse with respect to ε. Once the surface S (and hence the angle θ) is fixed, H2.0–H2.3 give upper bounds on ϱ. As will be shown in Section 2.2.1, H2.1–H2.3 imply that Conditions M4–M6 of Section 1.2 are satisfied. Specifically, H2.1 yields M4, H2.2 yields M5 and H2.3 yields M6. As for H2.0, it will be used in Section 2.2.2 to bound the Hausdorff distance between S and $\mathrm{Del}_{|S}(E)$.

If the surface S has a positive reach, then it is k-Lipschitz for any $k > 0$ (Theorem 2.6). Therefore, H2.0–H2.3 become simply $\varrho < +\infty$, which means that E is not required to be sparse. In this respect, H1 and H2 extend the sampling conditions introduced in Chapter 1 for the smooth case.

Notice that H2.1 is fulfilled as soon as H2.2 and H2.3 are satisfied – see Figure 2.5. For this reason, H2.1 could as well be removed from the list of conditions. However, in the sequel we will consider two types of point sets: those that verify H1 and H2 (which means that H2.0–H2.3 are all verified), and those that verify only H1 and H2.0–H2.1. This is why we keep H2.1 here. In the sequel, we will prove that, if E satisfies H1 and H2, then $\mathrm{Del}_{|S}(E)$ is a manifold without boundary, isotopic to S and at Hausdorff distance $O(\varepsilon)$ from S. If E satisfies H1 and H2.0–H2.1, then there is a subset of $\mathrm{Del}_{|S}(E)$ that has the above properties.

Our theoretical results hold only for small enough θ (and hence small enough k). This can be inferred from the hypotheses. Indeed, we must assume in H2 that $\varrho \geq 1$ because our proofs use the triangle normal lemma 2.13. Therefore, θ must be less than $12°$ for H2.0–H2.4 to be satisfiable, and less than $17.6°$ for H2.0–H2.1 to be satisfiable. Note that these bounds on θ are pessimistic: our experimental results show that loose ε-samples of S yield good topological and geometric approximations of S for values of θ ranging from 0 to $\frac{\pi}{2}$ – see Chapters 6 and 7.

Let us now review the guarantees we obtain under our sampling conditions. Let E be a point sample of S. For any facet f of $\mathrm{Del}_{|S}(E)$, we call B_f the surface Delaunay ball of smallest radius that circumscribes f. Let c_f and r_f denote respectively the center and radius of B_f. We set $D_f = S \cap B_f$, and $\tilde{D}_f = \tilde{S} \cap B_f$.

Given a point v of E, we consider the set F of facets of $\mathrm{Del}(E)$ incident to v whose circumradius is less than $(1+\nu)\varepsilon$. We call $r < (1+\nu)\varepsilon$ the maximum circumradius of a facet of F. Let p_v be a point of $\tilde{S} \cap B(v, (1+\nu)\varepsilon - r)$. Such a point exists since $\tilde{S}$ is dense in S, by Corollary 2.2. We have $\mathrm{d}(v, p_v) < (1+\nu)\varepsilon$. Moreover, for any facet $f \in F$, p_v belongs to the ball of radius $(1+\nu)\varepsilon$ centered at the circumcenter of f.

Orientation convention 2.19 *For any facet $f \in \mathrm{Del}(E)$ of radius less than $(1+\nu)\varepsilon$, we choose arbitrarily a vertex v of f and we orient f such that $\boldsymbol{n}(f) \cdot \boldsymbol{n}(p_v) > 0$.*

Like in the smooth case – see Section 1.2, it is not necessary to orient all the facets of $\mathrm{Del}(E)$, since only those of circumradius less than $(1+\nu)\varepsilon$ will be considered in the sequel. Among these are the facets of $\mathrm{Del}_{|S}(E)$, which are included in surface Delaunay balls of radius at most ε, by H1. In Section 2.2.1 (Th. 2.21), we will prove that, under H1 and H2.1, every edge of $\mathrm{Del}_{|S}(E)$ is incident to two facets of $\mathrm{Del}_{|S}(E)$. Moreover, under H1 and H2, $\mathrm{Del}_{|S}(E)$ is a compact surface without boundary, consistently oriented by the orientation convention 2.19. The proof relies on Theorem 1.7.

In Sections 2.2.2 and 2.2.3, we will show that, under H1 and H2, $\mathrm{Del}_{|S}(E)$ is isotopic to S and at Hausdorff distance $O(\varepsilon)$ from S. Like in the smooth case, our proofs use the fact that $\mathrm{Del}_{|S}(E)$ is a manifold, and they hold in a slightly more general setting. Specifically, we will show that, for any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ verifying:

- **Î1** *$\hat{S}$ is a compact surface without boundary, consistently oriented by the orientation convention 2.19,*
- **Î2** *$\hat{S}$ has vertices on all the connected components of S,*
- **Î3** *For any facet of $\hat{S}$, B_f has a radius at most $\varepsilon < \frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S)$,*
- **Î4** *The radius-edge ratios of the facets of $\hat{S}$ are at most ϱ, where $\varrho \geq 1$ verifies H2.0 and H2.1,*

$\hat{S}$ is isotopic to S and at Hausdorff distance $O_S(\varepsilon)$ from S. Referring to the smooth case, we observe that Î1–Î3 are quite similar to Hyp. I1–I3 of Section 1.2. In the nonsmooth case however, we need to assume further that the radius-edge ratios of the facets of $\hat{S}$ are bounded by some constant $\varrho \geq 1$ (Î4), in order to control the normals of $\hat{S}$ by the triangle normal lemma 2.13.

Observe that, if in some application $\mathrm{Del}_{|S}(E)$ is known to be a manifold without boundary, then, to guarantee that $\mathrm{Del}_{|S}(E)$ approximates S topologically and geometrically, it suffices to check that Î2–Î4 are satisfied with $\hat{S} = \mathrm{Del}_{|S}(E)$. Î2 and Î3 are guaranteed if E verifies H1. As for Î4, it is less restrictive than H2 since it does not assume that E is sparse. It follows that, when $\mathrm{Del}_{|S}(E)$ is a manifold without boundary, only H1 and Î4 are needed, and the point set E does not have to be sparse. As we will see in Section 2.2.1, the sparseness condition on E is only mandatory to prove that Hyp. M3 of Theorem 1.7 is satisfied. Therefore,

Open question 2.20 *Can the proof of Theorem 1.7 be rewritten, so that Hyp. M3 is no longer needed? In the affirmative, a loose ε-sample E of a k-Lipschitz surface S does not need to be sparse w.r.t. ε, for $\mathrm{Del}_{|S}(E)$ to be isotopic to S and at Hausdorff distance $O_S(\varepsilon)$ from S.*

To conclude our remarks, we recall that considering a subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ (and not $\mathrm{Del}_{|S}(E)$ itself) in our theoretical results will be instrumental in proving the correctness of certain Delaunay refinement algorithms in several meshing applications. See Section 5.1 and Chapter 8.

2.2.1 Manifold and layers

Theorem 2.21 *Let S be a k-Lipschitz surface and $E \subset S$ be a finite point set. If E verifies H1 and H2.1, then every edge of $\mathrm{Del}_{|S}(E)$ is incident to two facets of $\mathrm{Del}_{|S}(E)$. If E verifies H1 and H2, then $\mathrm{Del}_{|S}(E)$ is a compact surface without boundary, consistently oriented by the orientation convention 2.19.*

As explained in Section 1.2.1.2, if every edge of $\mathrm{Del}_{|S}(E)$ is incident to two facets of $\mathrm{Del}_{|S}(E)$, then the star of each vertex v consists of one or several umbrellas touching only at v. Consider the dual graph G of $\mathrm{Del}_{|S}(E)$, where each node represents a facet and two nodes are linked if their facets share an edge. Every connected component of G represents a subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ whose edges have two incident facets and whose vertices have one umbrella. Hence, $\hat{S}$ is a manifold without boundary, consistently oriented by the orientation convention 2.19. Moreover, two such complexes are disjoint

except possibly at some common vertices. Therefore, we call them *layers*. The next result follows immediately:

Corollary 2.22 *Let S be a k-Lipschitz surface and $E \subset S$ be a finite point set. If E verifies H1 and H2.0–H2.1, then every layer of $\mathrm{Del}_{|S}(E)$ satisfies Î1, Î3 and Î4.*

Notice that H2.0 is here only to ensure that Î4 is satisfied. Indeed, Î3 is an direct consequence of H1, and only H1 and H2.1 are necessary to ensure that every layer of $\mathrm{Del}_{|S}(E)$ verifies Î1, as stated in Theorem 2.21.

Proof of the theorem. We will show that Hyp. M1–M6 of Section 1.2 are satisfied. Then, by Theorem 1.7 (whose proof does not rely on the fact that the surface S is smooth), we will get the result.

Let v be a vertex of $\mathrm{Del}_{|S}(E)$. By H1, we have $\varepsilon < \frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S) < \frac{1}{4}\ \mathrm{lr}_k(S)$, thus $(1+\nu)\varepsilon < \mathrm{lr}_k(S)$. Hence, the point p_v defined in the orientation convention 2.19 lies in $B(v, \mathrm{lr}_k(S))$, since it belongs to $\tilde{S} \cap B(v,(1+\nu)\varepsilon)$. By the cocone lemma 2.12, for any $q \in S \cap B(v,(1+\nu)\varepsilon)$, $S \cap B(v,(1+\nu)\varepsilon)$ lies outside the double cone of apex q, of axis aligned with $\mathbf{n}(p_v)$ and of half-angle $\frac{\pi}{2} - \theta$. Therefore, M1 is satisfied with $\theta_1 = \theta$.

Let f be a facet of $\mathrm{Del}_{|S}(E)$ incident to v. By H1, the circumradius of f is at most ε, hence its radius-edge ratio is bounded by ϱ, by H2. Let u be the vertex of f that was used to orient f. If $u = v$, then, since f and p_v are both included in $B(v, 2\varepsilon) \subset B(v, \mathrm{lr}_k(S))$, the triangle normal lemma 2.13 states that the angle between $\mathbf{n}(f)$ and $\mathbf{n}(p_v)$ verifies:

$$\sin(\mathbf{n}(f), \mathbf{n}(p_v)) \leq 2\varrho \sin\theta \tag{2.9}$$

Moreover, according to H2.1, we have $2\varrho\sin\theta < 1$, hence $(\mathbf{n}(f), \mathbf{n}(p_v)) \leq \arcsin(2\varrho\sin\theta)$. If $u \neq v$, then Eq. (2.9) holds with p_v replaced by p_u. Moreover, since $\mathrm{d}(v,u) \leq 2\varepsilon$ and $\mathrm{d}(p_v, v) \leq (1+\nu)\varepsilon$, p_v belongs to $B(u,(3+\nu)\varepsilon)$ which is included in $B(u, \mathrm{lr}_k(S))$ since $\varepsilon < \frac{1}{4}\ \mathrm{lr}_k(S)$, by H1. Therefore, by the normal variation lemma 2.13, we have $(\mathbf{n}(p_u), \mathbf{n}(p_v)) \leq \theta$. It follows that $(\mathbf{n}(f), \mathbf{n}(p_v))$ is bounded by $\theta + \arcsin(2\varrho\sin\theta)$. Hence, M2 is satisfied with $\theta_2 = \theta + \arcsin(2\varrho\sin\theta)$.

It follows from the above observations that $2\theta_1 + \theta_2 \leq 3\theta + \arcsin(2\varrho\sin\theta)$, which is less than $\frac{\pi}{2}$, by H2.1. Therefore, under H1 and H2.1, M1, M2 and M4 are fulfilled, which implies that every edge of $\mathrm{Del}_{|S}(E)$ is incident to exactly two facets of $\mathrm{Del}_{|S}(E)$, by Lemma 1.10.

Note to the reader: *Observe that the above arguments still hold if H2 is replaced by Î4, with $\hat{S} = \mathrm{Del}_{|S}(E)$. As a consequence, the sparseness condition on E is only mandatory to prove M3.*

Let now f be a facet of $\mathrm{Del}(E)$ incident to v, of circumradius less than $(1+\nu)\varepsilon$. The radius-edge ratio of f is bounded by $(1+\nu)\varrho$, by H2. Let u be the vertex of f that was used to orient f. If $u = v$, then, since f and p_v are both included in $B(v, 2(1+\nu)\varepsilon) \subset B(v, \mathrm{lr}_k(S))$, the triangle normal lemma 2.13 states that:

$$\sin(\mathbf{n}(f), \mathbf{n}(p_v)) \leq 2(1+\nu)\varrho \sin\theta \tag{2.10}$$

Moreover, by H2.3, we have $2(1+\nu)\varrho\sin\theta < 1$, hence $(\mathbf{n}(f), \mathbf{n}(p_v)) \leq \arcsin(2(1+\nu)\varrho\sin\theta)$. If $u \neq v$, then Eq. (2.10) holds with p_v replaced by p_u. Moreover, u and p_v both belong to the ball of radius $(1+\nu)\varepsilon$ centered at the circumcenter of f. Hence, p_v belongs to $B(u, 2(1+\nu)\varepsilon) \subseteq$

$B(u, \mathrm{lr}_k(S))$. Therefore, by the normal variation lemma 2.13, we have $(\mathbf{n}(p_u), \mathbf{n}(p_v)) \leq \theta$. It follows that $(\mathbf{n}(f), \mathbf{n}(p_v))$ is bounded by $\theta + \arcsin(2(1+\nu)\varrho \sin\theta)$. Hence, M3 is satisfied with $\theta_3 = \theta + \arcsin(2(1+\nu)\varrho \sin\theta)$.

From the above observations we deduce that $2\sin\theta_1 - \nu\cos\theta_2 \leq 2\sin\theta - \nu\cos(\theta + \arcsin(2\varrho\sin\theta))$, which by H2.2 is negative. Moreover, $\sin\theta_1 - \cos\theta_3 \leq \sin\theta - \cos(\theta + \arcsin(2(1+\nu)\varrho\sin\theta))$, which is also negative, by H2.3. Therefore, M5 and M6 are fulfilled. It follows that $\mathrm{Del}_{|S}(E)$ is a compact oriented manifold without boundary, by Theorem 1.7. This concludes the proof of Theorem 2.21. □

If H1 and H2.1 are satisfied, then, by Remark 1.9, for any facet $f \in \mathrm{Del}_{|S}(E)$, B_f is the only surface Delaunay ball circumscribing f. Equivalently, the Voronoi edge dual to f intersects S only once.

2.2.2 Hausdorff distance

The goal of this section is to prove the following

Theorem 2.23 *Let S be a k-Lipschitz surface, and $\theta = 2\arctan k$. For any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that verifies Î1–Î4, the Hausdorff distance between $\hat{S}$ and S is bounded by $\frac{\varepsilon}{\cos^2\theta}$.*

Theorem 2.23, combined with Corollaries 2.18 and 2.22, yields the following

Corollary 2.24 *Let S be a k-Lipschitz surface and $E \subset S$ be a finite point set. If E verifies H1 and H2.0–H2.1, then every layer of $\mathrm{Del}_{|S}(E)$ is at Hausdorff distance at most $\frac{\varepsilon}{\cos^2\theta}$ from the connected component of S on which its vertices lie. As a consequence, $\mathrm{Del}_{|S}(E)$ is at Hausdorff distance at most $\frac{\varepsilon}{\cos^2\theta}$ from S.*

Proof of Corollary 2.24.. Under H1, the vertices of a facet of $\mathrm{Del}_{|S}(E)$ cannot lie on distinct connected components of S, by Corollary 2.18. Therefore, the vertices of a layer of $\mathrm{Del}_{|S}(E)$ lie on the same connected component of S. The result follows, by Corollary 2.22 and Theorem 2.23. □

To prove Theorem 2.23, we will first bound the distance from $\hat{S}$ to S (Proposition 2.25), and then the distance from S to $\hat{S}$ (Proposition 2.26). Let $\mathcal{T}_\varepsilon$ be the tubular neighborhood of S of width ε, defined as in Eq. (1.3).

Proposition 2.25 *$\hat{S}$ is included in $\mathcal{T}_\varepsilon$. Therefore, the semi-Hausdorff distance from $\hat{S}$ to S is less than ε.*

Proof. This result is already known as Lemma 1.15. The proof is identical. □

Proposition 2.26 *The semi-Hausdorff distance from S to $\hat{S}$ is at most $\frac{\varepsilon}{\cos^2\theta}$.*

Propositions 2.25 and 2.26 prove Theorem 2.23. Our strategy for proving Proposition 2.26 is similar in spirit to the one adopted for the smooth case – see Section 1.2.2, Theorem 1.13. The approach consists in *pushing* the points of $\hat{S}$ along some continuous flow towards S, and showing that every point of S is eventually reached by some point of $\hat{S}$. The drawback of the flow of Section 1.2.2.1 is that it is not

defined on the medial axis of S, which, in the present case, may intersect $\mathcal{T}_\varepsilon$ for any positive value of ε, since S is not smooth. Therefore, this flow is not well defined over $\mathcal{T}_\varepsilon$ and cannot be used in our context.

In Section 2.2.2.1, we define a new flow ϕ that was first introduced by Lieutier [85]. This flow has the advantage of being well-defined and continuous over $\mathcal{T}_\varepsilon \setminus S$. However, ϕ is not defined over S. Consequently, we cannot adopt the strategy used in the smooth setting, which consisted in showing that the set of points of S whose flow lines intersect $\hat{S}$ is S itself. Here, we will proceed in three steps:

- First, in Section 2.2.2.2, we consider the set $\mathcal{I}$ of points of $\mathcal{O}$ whose flow lines intersect $\hat{S}$. We prove that $\mathcal{I}$ is a union of connected components of $\mathbb{R}^3 \setminus (S \cup \hat{S})$.
- Then, in Section 2.2.2.3, we consider the sets $S \cap \partial\mathcal{I}$ and $S \cap \hat{S}$, and we prove that their union covers S. As a consequence, every point $p \in S$ belongs either to $\hat{S}$ or to $\partial\mathcal{I}$. In the latter case, arbitrarily close to p, one can find points whose flow lines intersect $\hat{S}$.
- Using this last observation, we can conclude the proof of Proposition 2.26 in Section 2.2.2.4, by bounding the distance travelled along a flow line before reaching $\hat{S}$.

2.2.2.1 The flow

In the sequel, $\mathcal{O}$ denotes the open set $\mathbb{R}^3 \setminus S$. For any $d > 0$, we define $\mathcal{O}_d$ as the set of points of $\mathcal{O}$ that are farther than d from S. We have: $\mathcal{O}_d = \mathcal{O} \setminus \bar{\mathcal{T}}_d$, where $\mathcal{T}_d$ is the tubular neighborhood of S of width d, as defined in Section 1.2.2.1.

It is proved in [85] that Euler schemes, using the vector field $\boldsymbol{\nabla}$ defined in Section 2.1.6, converges uniformly towards a continuous flow $\phi : \mathbb{R}^+ \times \mathcal{O} \to \mathcal{O}$ that verifies:

$$\forall t \in \mathbb{R}^+,\ \forall p \in \mathcal{O},\ \phi(t,p) = p + \int_{t'=0}^{t} \boldsymbol{\nabla}(\phi(t',p))\, dt' \tag{2.11}$$

Intuitively, the real variable t stands for the time, while the other variable is the position in space. It follows from Eq. (2.11) that the *stationary points* of ϕ (*i.e.* the points $p \in \mathcal{O}$ such that $\phi(t,p) = p$ $\forall t \in \mathbb{R}^+$) are the critical points of $\boldsymbol{\nabla}$, *i.e.* the points of Φ.

For any $p \in \mathcal{O}$, we call *flow line of* p and note $\Lambda(p)$, the trajectory of p along ϕ:

$$\Lambda(p) = \phi(\mathbb{R}^+, p) = \{\phi(t,p) \mid t \geq 0\}$$

The flow ϕ enjoys several properties, including:

F1 [85, Lemma 4.12]
For any $p \in \mathcal{O} \setminus \Phi$, the distance to S increases strictly along $\Lambda(p)$, that is, the map $t \mapsto \mathrm{d}_S(\phi(t,p))$ is strictly increasing. Moreover,

$$\forall p \in \mathcal{O},\ \forall t \in \mathbb{R}^+,\ \mathrm{d}_S(\phi(t,p)) = \mathrm{d}_S(p) + \int_{t'=0}^{t} \|\boldsymbol{\nabla}(\phi(t',p))\|^2\, dt' \tag{2.12}$$

F2 [85, Lemma 4.13]
For any $p \in \mathcal{O}$, the map $t \mapsto \phi(t,p)$ is 1-Lipschitz. Moreover,

$$\forall d > 0,\ \forall t \geq 0,\ \forall p,q \in \mathcal{O}_d,\ \mathrm{d}(\phi(t,p),\phi(t,q)) \leq e^{t/d}\, \mathrm{d}(p,q) \tag{2.13}$$

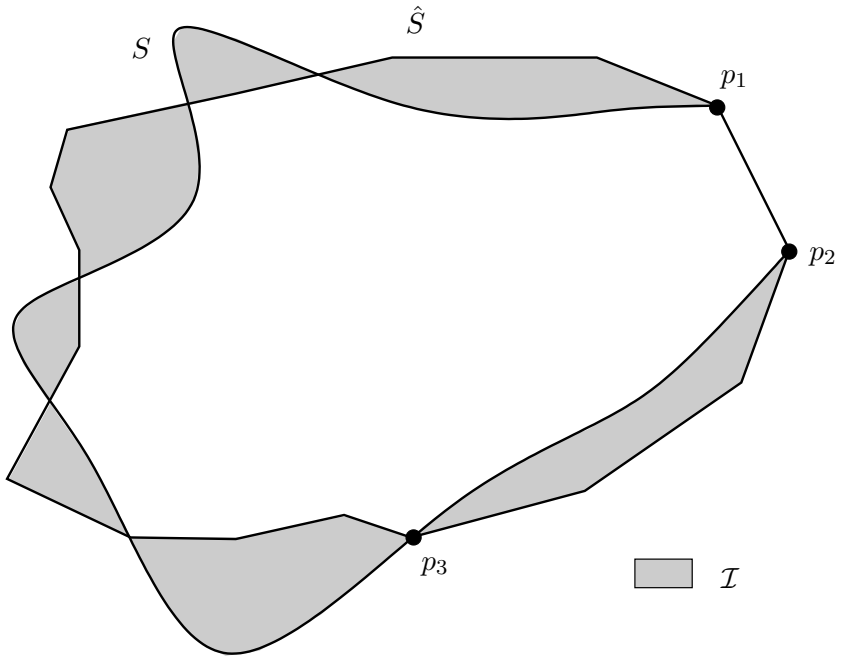

Figure 2.6: $\mathcal{I}$ is a union of connected components of $\mathcal{O} \setminus \hat{S}$.

The fact that ϕ is continuous implies that $\Lambda(p)$ is a connected arc, for any $p \in \mathcal{O}$. If $p \in \Phi$, then $\Lambda(p)$ is reduced to a point. Otherwise, by F1, the distance to S increases strictly along $\Lambda(p)$, thus $\Lambda(p)$ does not self-intersect. It follows also from F1 that $\Lambda(p)$ cannot leave and then re-enter $\mathcal{T}_\varepsilon$. Therefore, for $p \in \mathcal{T}_\varepsilon$, $\Lambda(p) \cap \mathcal{T}_\varepsilon$ is a simple arc, and for $p \notin \mathcal{T}_\varepsilon$, $\Lambda(p) \cap \mathcal{T}_\varepsilon$ is empty. The next result bounds the time spent before a point moving along a flow line leaves $\mathcal{T}_\varepsilon$:

Lemma 2.27

$$\forall p \in \mathcal{T}_\varepsilon \setminus S, \quad \forall t \geq \tfrac{\varepsilon - \mathrm{d}_S(p)}{\cos^2\theta}, \quad \phi(t,p) \notin \mathcal{T}_\varepsilon$$
$$\forall p \in \mathcal{O} \setminus \mathcal{T}_\varepsilon, \quad \forall t \geq 0, \quad \phi(t,p) \notin \mathcal{T}_\varepsilon$$

Proof. Given $p \in \mathcal{T}_\varepsilon \setminus S$ and $t \in \mathbb{R}^+$ such that $\phi(t,p) \in \mathcal{T}_\varepsilon$, we know by F1 that $\phi(t',p)$ belongs to $\mathcal{T}_\varepsilon$ for any $t' \in [0,t]$. By Lemma 2.15 and Eq. (2.12), we then have: $\mathrm{d}_S(\phi(t,p)) \geq \mathrm{d}_S(p) + t\ \cos^2\theta$. Hence, the time t_ε at which $\mathrm{d}_S(\phi(t_\varepsilon,p)) = \varepsilon$ is at most $\frac{\varepsilon - \mathrm{d}_S(p)}{\cos^2\theta}$. This means that $\phi(t,p) \notin \mathcal{T}_\varepsilon$ for all $t \geq \frac{\varepsilon - \mathrm{d}_S(p)}{\cos^2\theta}$, hereby proving the lemma for $p \in \mathcal{T}_\varepsilon$.

Given $p \in \mathcal{O} \setminus \mathcal{T}_\varepsilon$, F1 states that $\forall t \in \mathbb{R}^+$, $\mathrm{d}_S(\phi(t,p)) \geq \mathrm{d}_S(p) \geq \varepsilon$. Hence, $\phi(t,p) \notin \mathcal{T}_\varepsilon$, which proves the lemma for $p \in \mathcal{O} \setminus \mathcal{T}_\varepsilon$. □

2.2.2.2 Flow lines intersecting $\hat{S}$

We define $\mathcal{I}$ as the set of points of $\mathcal{O}$ whose flow lines intersect $\hat{S}$. For convenience, we exclude the points of $\hat{S}$ from $\mathcal{I}$:

$$\mathcal{I} = \{p \in \mathcal{O} \setminus \hat{S} \mid \Lambda(p) \cap \hat{S} \neq \emptyset\}$$

Our aim is to prove that $\mathcal{I}$ is a union of connected components of $\mathcal{O} \setminus \hat{S}$, as illustrated in Figure 2.6. Since $\mathcal{I} \subseteq \mathcal{O} \setminus \hat{S}$, this comes down to proving that the boundary of $\mathcal{I}$ is included in $S \cup \hat{S}$. We first show that the boundary of $\mathcal{I}$ lies in $S \cup \hat{S} \cup \partial\mathcal{T}_\varepsilon$:

Lemma 2.28 *For any $p \in \mathcal{I}$, there exists a positive value $r(p)$, vanishing only as p approaches S or $\hat{S}$ or $\partial\mathcal{T}_\varepsilon$, such that $B(p, r(p)) \subseteq \mathcal{I}$. As a consequence, $\partial\mathcal{I} \subseteq S \cup \hat{S} \cup \partial\mathcal{T}_\varepsilon$.*

Proof. Three major steps of the proof are stated as Claims 2.28.1, 2.28.2 and 2.28.3. To ease the reading, their proofs are deferred to the end of Section 2.2.2.

Let $p \in \mathcal{I}$. Since $p \notin S$, $\mathrm{d}_S(p)$ is positive. By F2, the restriction of ϕ to $\left[0, \frac{\varepsilon - \mathrm{d}_S(p)/2}{\cos^2\theta}\right] \times \mathcal{O}_{\mathrm{d}_S(p)/2}$ is 1-Lipschitz as a function of time, and $\exp\left(\frac{2\varepsilon - \mathrm{d}_S(p)}{\mathrm{d}_S(p)\ \cos^2\theta}\right)$-Lipschitz as a function of space. Let $\kappa = \exp\left(\frac{2\varepsilon - \mathrm{d}_S(p)}{\mathrm{d}_S(p)\ \cos^2\theta}\right)$. Since $\hat{S} \subset \mathcal{T}_\varepsilon$, p belongs to $\mathcal{T}_\varepsilon$, by Lemma 2.27. Hence, $\mathrm{d}_S(p) < \varepsilon$, which implies that $\kappa > 1$.

The function $q \mapsto \mathrm{d}(q, \partial\mathcal{T}_\varepsilon)$ is continuous over $\hat{S}$, thus it reaches its minimum δ since $\hat{S}$ is compact. This minimum is positive because $\hat{S} \subset \mathcal{T}_\varepsilon$. In addition, for any facet f of $\hat{S}$, the function $u \mapsto \mathrm{d}(u, \hat{S} \setminus \mathrm{N}(f))$ (where $\mathrm{N}(f)$ is the neighborhood of f on $\hat{S}$) is positive and continuous over f, hence its minimum $m(f)$ over f is positive. Let $m = \min\{m(f),\ f \in \hat{S}\}$. For any point $w \in \hat{S}$ and any facet f containing w, the distance of w to $\hat{S} \setminus \mathrm{N}(f)$ is at least m.

We set $r(p)$ as follows:

$$r(p) \ = \ \min\left\{\frac{\mathrm{d}_S(p)}{2\kappa},\ \varepsilon - \mathrm{d}_S(p),\ \frac{1}{3\kappa}\,\mathrm{d}(p, \hat{S}),\ \frac{\delta}{2\kappa},\ \frac{m}{2\kappa}\right\} \tag{2.14}$$

Notice that $r(p)$ vanishes only if $\mathrm{d}_S(p) \to 0$ (p approaches S), or if $(\varepsilon - \mathrm{d}_S(p)) \to 0$ (p approaches $\partial\mathcal{T}_\varepsilon$), or if $\mathrm{d}(p, \hat{S}) \to 0$ (p approaches $\hat{S}$). We will prove that $B(p, r(p)) \subseteq \mathcal{I}$.

Let q lie in the open ball $B(p, r(p))$. Since $r(p) \leq \varepsilon - \mathrm{d}_S(p)$, q belongs to $\mathcal{T}_\varepsilon$. Moreover, $r(p)$ is less than $\min\{\mathrm{d}_S(p),\ \mathrm{d}(p, \hat{S})\}$, since $\kappa > 1$. Thus, $q \notin S \cup \hat{S}$. Let us prove that $\Lambda(q)$ intersects $\hat{S}$.

Since $p \in \mathcal{I}$, $\Lambda(p)$ intersects $\hat{S}$. Let $p' \in \Lambda(p) \cap \hat{S}$. We have $p' \neq p$ since $p \notin \hat{S}$. Let d' be defined by

$$d' = \min\left\{\frac{\mathrm{d}_S(p)}{2},\ \frac{\mathrm{d}(p, \hat{S})}{3},\ \frac{\delta}{2},\ \frac{m}{2}\right\}$$

We call $B^1_{p'}$ and $B^2_{p'}$ the open balls centered at p', of radii d' and $2d'$ respectively. Observe that $2d' \leq \frac{2}{3}\,\mathrm{d}(p, \hat{S}) \leq \frac{2}{3}\,\mathrm{d}(p, p')$. Moreover, since $\kappa > 1$, $r(p)$ is less than $\frac{1}{3}\,\mathrm{d}(p, \hat{S}) \leq \frac{1}{3}\,\mathrm{d}(p, p')$. Hence,

$$B(p, r(p)) \cap B^2_{p'} = \emptyset \tag{2.15}$$

Claim 2.28.1 *$\Lambda(q)$ pierces $B^1_{p'}$,* i.e. *it enters and then leaves $B^1_{p'}$. Similarly, $\Lambda(q)$ pierces $B^2_{p'}$.*

Let f be a facet of $\hat{S}$ that contains p'. Let $B(c_f, r_f)$ be the surface Delaunay ball of f. By $\hat{\mathrm{I}}3$, we have $r_f \leq \varepsilon$. Moreover, $2d'$ is at most $\mathrm{d}_S(p)$, which by F1 is less than $\mathrm{d}_S(p') \leq \mathrm{d}(p', c_f) \leq r_f \leq \varepsilon$. Therefore, $B^2_{p'}$ is included in $B(c_f, 2\varepsilon)$. Let p_f be a point of $\tilde{S} \cap B(c_f, 2\varepsilon)$. Such a point exists because $\tilde{S}$ is dense in S, by Corollary 2.2.

By Claim 2.28.1, $\Lambda(q) \cap B^1_{p'}$ is not empty. Let $q' \in \Lambda(q) \cap B^1_{p'}$. We call $K(q')$ the double cone of apex q', of axis aligned with $\mathbf{n}(p_f)$ and of half-angle θ. Since $q' \in B^1_{p'} \subset B^2_{p'}$, $K(q')$ intersects $\partial B^2_{p'}$

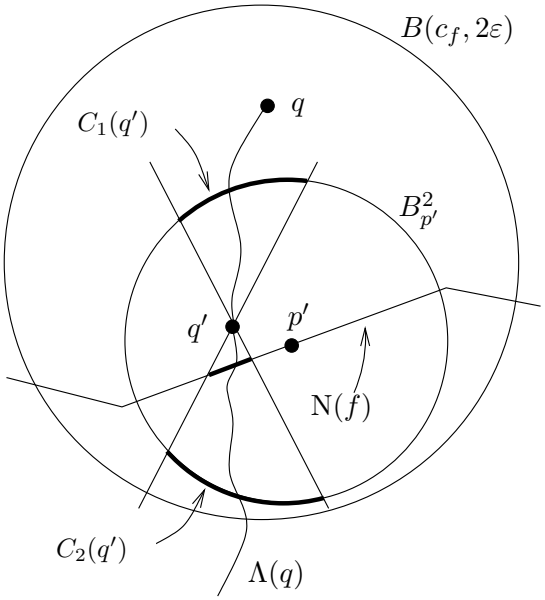

Figure 2.7: For the proof of Lemma 2.28.

along two spherical patches $C_1(q')$ and $C_2(q')$, such that every connected curve included in $K(q')$ and joining $C_1(q')$ to $C_2(q')$ passes through q'. One arc of $\Lambda(q) \cap B^2_{p'}$ has this property, as stated in the next claim and illustrated in Figure 2.7:

Claim 2.28.2 *Let $\Lambda'(q)$ be the arc of $\Lambda(q) \cap B^2_{p'}$ that contains q'. $\Lambda'(q)$ lies in $K(q')$ and joins $C_1(q')$ to $C_2(q')$, with one endpoint in $C_1(q')$ and the other endpoint in $C_2(q')$.*

The next step is to show that such an arc intersects $\mathrm{N}(f)$:

Claim 2.28.3 *Inside $B^2_{p'}$, $C_1(q')$ and $C_2(q')$ are separated by* $\mathrm{N}(f)$, i.e. *every connected curve included in $B^2_{p'}$ and joining $C_1(q')$ to $C_2(q')$ intersects* $\mathrm{N}(f)$.

It follows from Claims 2.28.2 and 2.28.3 that $\Lambda(q)$ intersects $\mathrm{N}(f)$. Hence, $\Lambda(q) \cap \hat{S} \neq \emptyset$, which means that $q \in \mathcal{I}$. This ends the proof of Lemma 2.28. □

By Lemma 2.28, the boundary of $\mathcal{I}$ is included in $S \cup \hat{S} \cup \partial\mathcal{T}_\varepsilon$. We now prove that, in fact, $\partial\mathcal{I}$ does not touch $\partial\mathcal{T}_\varepsilon$:

Lemma 2.29 $\partial\mathcal{I} \cap \partial\mathcal{T}_\varepsilon = \emptyset$.

Proof. Since $\hat{S}$ is compact and d_S is continuous, the restriction of d_S to $\hat{S}$ reaches its maximum. Let $p \in \hat{S}$ be such that $\forall p' \in \hat{S}$, $\mathrm{d}_S(p') \leq \mathrm{d}_S(p)$. By Proposition 2.25, $\hat{S}$ is included in $\mathcal{T}_\varepsilon$, thus $\mathrm{d}_S(p)$ is less than ε. It follows that $\hat{S}$ is in fact included in $\mathcal{T}_{\delta'}$, for any δ' such that $\mathrm{d}_S(p) < \delta' < \varepsilon$. By Lemma 2.27, for any $q \notin \mathcal{T}_{\delta'}$, $\Lambda(q) \cap \mathcal{T}_{\delta'} = \emptyset$, hence $\mathcal{I}$ is included in $\mathcal{T}_{\delta'}$. As a consequence, $\partial\mathcal{I}$ is included in the topological closure of $\mathcal{T}_{\delta'}$, which does not intersect $\partial\mathcal{T}_\varepsilon$ since $\delta' < \varepsilon$. □

Lemmas 2.28 and 2.29 imply that the boundary of $\mathcal{I}$ is included in $S \cup \hat{S}$, which concludes the proof of the main result of Section 2.2.2.2:

Lemma 2.30 *$\mathcal{I}$ is a union of connected components of $\mathcal{O} \setminus \hat{S}$.*

We end Section 2.2.2.2 by proving the three claims stated in the proof of Lemma 2.28. The notations are the same.

Proof of Claim 2.28.1. Since $B^1_{p'} \subset B^2_{p'}$, it suffices to show that $\Lambda(q)$ intersects $B^1_{p'}$ and pierces $B^2_{p'}$.

Let t' be the time at which $\Lambda(p)$ reaches p'. We have $t' > 0$, since $p' \neq p$. Moreover, $p' \in \hat{S} \subset \mathcal{T}_\varepsilon$, thus $t' < \frac{\varepsilon - \mathrm{d}_S(p)}{\cos^2\theta}$, by Lemma 2.27. Since $\kappa > 1$, we have $\mathrm{d}(p,q) < r(p) < \mathrm{d}_S(p)/2$. Hence, p and q belong to $\mathcal{O}_{\mathrm{d}_S(p)/2}$. Therefore,

$$\mathrm{d}(\phi(t',p),\ \phi(t',q)) \leq \kappa\ \mathrm{d}(p,q) < \kappa\ r(p) \leq d',$$

since ϕ is κ-Lipschitz as a function of space over $\left[0, \frac{\varepsilon - \mathrm{d}_S(p)/2}{\cos^2\theta}\right] \times \mathcal{O}_{\mathrm{d}_S(p)/2}$. It follows that $\phi(t',q) \in B^1_{p'}$, which means that $\Lambda(q)$ intersects $B^1_{p'}$ and hence $B^2_{p'}$.

Since $q \in B(p, r(p))$, q does not belong to $B^2_{p'}$, by Eq. (2.15). Moreover, since $p' \in \hat{S}$ and $2d' \leq \delta$, $B^2_{p'}$ is included in $\mathcal{T}_\varepsilon$. Hence, for sufficiently large t, $\phi(t,q) \notin B^2_{p'}$, by Lemma 2.27. It follows that $\Lambda(q)$ enters and then leaves $B^2_{p'}$, which concludes the proof of the claim. □

Proof of Claim 2.28.2. We assume without loss of generality that $p \in \mathcal{O}^+$, the case $p \in \mathcal{O}^-$ being symmetric. This assumption implies that $B^2_{p'} \subseteq \mathcal{O}^+$. Since, by Î3, we have $\varepsilon \leq \frac{1}{4}\ \mathrm{lr}_k(S)$, $B^2_{p'}$ is included in $B(c_f, \frac{1}{2}\ \mathrm{lr}_k(S))$. Thus, for any $t > 0$ such that $\phi(t,q) \in B^2_{p'}$, $\boldsymbol{\nabla}(\phi(t,q))$ makes an angle of at most θ with $\mathbf{n}(p_f)$, by Lemma 2.15. Hence, by Eq. (2.11), $\Lambda'(q)$ lies in $K(q')$. The fact that $\Lambda'(q)$ joins the two connected components of $K(q') \cap \partial B^2_{p'}$ is an immediate consequence of the fact that $\Lambda(q)$ pierces $B^2_{p'}$, by Claim 2.28.1. □

Proof of Claim 2.28.3. Let α be the maximum angle between $\mathbf{n}(p_f)$ and the normals of the facets of $\mathrm{N}(f)$. According to Î4, the radius-edge ratios of the facets of $\mathrm{N}(f)$ are at most ϱ. Therefore, by the triangle normal lemma 2.13 and the orientation convention 2.19, we have $\sin\alpha \leq 2\varrho\sin\theta$. By Î4, ϱ verifies H2.1, hence $2\varrho\sin\theta < 1$. It follows that $\alpha \leq \arcsin(2\varrho\sin\theta)$.

Moreover, the radius $2d'$ of $B^2_{p'}$ is at most m, thus $\hat{S} \cap B^2_{p'} \subset \mathrm{N}(f)$. Since, by Î1, $\hat{S}$ is a polyhedron without boundary, we have $\partial(\hat{S} \cap B^2_{p'}) = \hat{S} \cap \partial B^2_{p'}$. Moreover, by Lemma 2.8, $\hat{S} \cap B^2_{p'}$ is the graph of a $(\tan\alpha)$-Lipschitz bivariate function defined over the plane orthogonal to $\mathbf{n}(p_f)$. Hence, by Lemma[1] 2.10, $\hat{S} \cap B^2_{p'}$ lies outside the double cone $K(p')$ of apex p', of axis aligned with $\mathbf{n}(p_f)$ and of half-angle $\frac{\pi}{2} - \alpha$. Therefore, inside $B^2_{p'}$, $\mathrm{N}(f)$ separates the two cones of $K(p')$, that is, every connected curve included in $B^2_{p'}$ and joining the two cones of $K(p')$ intersects $\mathrm{N}(f)$ – see Figure 2.8. To prove the claim, it suffices then to show that $C_1(q')$ is included in one cone of $K(p')$ whereas $C_2(q')$ is included in the other cone.

Since $K(p')$ and $K(q')$ have parallel axes, $K(q')$ intersects the two cones of $K(p')$. Let K_1 and K_2 be the intersections of $K(q')$ with the two cones of $K(p')$. According to Î4, ϱ verifies H2.1, hence we have

$$\theta + \alpha \leq \theta + \arcsin(2\varrho\sin\theta) \leq \frac{\pi}{2} - 2\theta < \frac{\pi}{2}$$

[1] Notice that we could use the cocone Lemma 2.12 as well, but it would give an angle of $\frac{\pi}{2} - 2\alpha$.

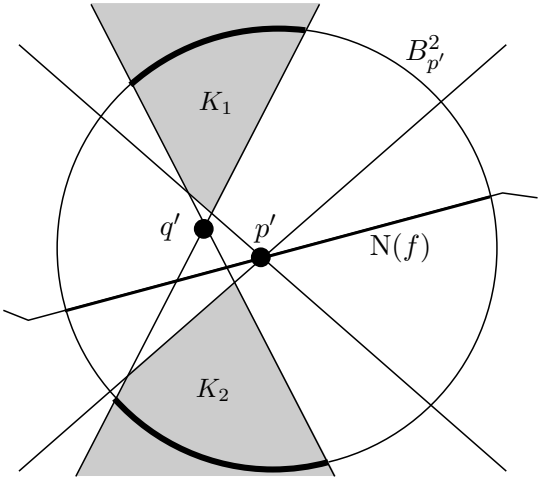

Figure 2.8: For the proof of Claim 2.28.3.

Thus, the half-angle $\frac{\pi}{2} - \alpha$ of $K(p')$ is larger than the half-angle θ of $K(q')$. Therefore, $K(q') \setminus K(p')$ is bounded. Moreover, the boundaries of K_1 and K_2 touch $K(q') \setminus K(p')$.

Let q'' be the point of $K(q') \setminus K(p')$ that is farthest from p'. By a symmetry argument, q'' lies in the plane defined by p', q' and $\mathbf{n}(p_f)$. Moreover, q'' lies on $\partial K(q') \cap \partial K(p')$. The intersection of $K(q')$ with the plane aff$(p', q', \mathbf{n}(p_f))$ is a double wedge of apex q', of axis $\mathbf{n}(p_f)$ and of half-angle θ. By definition, q' lies in the ball $B^1_{p'}$. For any position of q' in this ball, we consider the two lines of aff$(p', q', \mathbf{n}(p_f))$ that pass through q' and that make angles of θ with $\mathbf{n}(p_f)$. Among the set of such lines, the ones that lie farthest from p' are tangent to $\partial B^1_{p'}$. Hence, the distance between q'' and p' is maximal when q' lies on $\partial B^1_{p'}$, with (q', q'') tangent $\partial B^1_{p'}$. The angle between $\mathbf{n}(p_f)$ and (p', q') is then $\frac{\pi}{2} - \theta$, which implies that the distance between q'' and p' is $\frac{\mathrm{d}(p',q')}{\cos(\theta+\alpha)} = \frac{d'}{\cos(\theta+\alpha)}$.

By Î4, ϱ verifies H2.0, hence $\frac{d'}{\cos(\theta+\alpha)}$ is less than $2d'$. This means that $K(q') \setminus K(p')$ is included in $B^2_{p'}$. As a consequence, $K(q') \cap \partial B^2_{p'}$ belongs to $K(p') \cap \partial B^2_{p'}$, as illustrated in Figure 2.8. Moreover, K_1 and K_2 intersect $B^2_{p'}$ because their boundaries touch $K(q') \setminus K(p')$, which is included in $B^2_{p'}$. Since K_1 and K_2 are unbounded, they intersect $\partial B^2_{p'}$ as well. Therefore, one component of $K(q') \cap \partial B^2_{p'}$ (say $C_1(q')$) belongs to K_1, whereas the other component (namely, $C_2(q')$) belongs to K_2. This concludes the proof of the claim. □

2.2.2.3 $\hat{S} \cup \partial\mathcal{I}$ covers S

Let $S_\mathcal{D} = S \cap \hat{S}$ and $S_\mathcal{I} = S \cap \partial\mathcal{I}$. Our aim is to prove the following lemma:

Lemma 2.31 $S = S_\mathcal{D} \cup S_\mathcal{I}$.

The sets $S_\mathcal{D}$ and $S_\mathcal{I}$ can be visualized in Figure 2.6: the arc $\overline{p_1p_2}$ of S belongs to $S_\mathcal{D}$, while the arc $\overline{p_2p_3}$ of S belongs to $S_\mathcal{I}$. We will show that $S_\mathcal{D} \cup S_\mathcal{I}$ has no boundary, for the topology of S induced by $\mathbb{R}^3$. This will imply that $S_\mathcal{D} \cup S_\mathcal{I}$ is a union of connected components of S. Since, by Î2, $S_\mathcal{D}$ intersects every connected component of S, we will then have $S_\mathcal{D} \cup S_\mathcal{I} = S$, which proves Lemma 2.31.

The boundary of $S_{\mathcal{D}} \cup S_{\mathcal{I}}$ is included in $\partial S_{\mathcal{D}} \cup \partial S_{\mathcal{I}}$. According to Lemma 2.30, $\mathcal{I}$ is a union of connected components of $\mathbb{R}^3 \setminus (S \cup \hat{S})$. Thus, $\partial S_{\mathcal{I}}$ is included in $\partial S_{\mathcal{D}}$. Therefore, to prove that the boundary of $S_{\mathcal{D}} \cup S_{\mathcal{I}}$ is empty, we need only to show that $\partial S_{\mathcal{D}}$ is included in the interior of $S_{\mathcal{D}} \cup S_{\mathcal{I}}$:

Lemma 2.32 *For any point $p \in \partial S_{\mathcal{D}}$, there exists a small neighborhood $\mathcal{N}$ of p on S such that $\mathcal{N} \subseteq S_{\mathcal{D}} \cup S_{\mathcal{I}}$.*

Proof. Let $\delta > 0$ and $m > 0$ be defined as in the proof of Lemma 2.28. Recall that δ is the minimum distance from a point of $\hat{S}$ to $\partial\mathcal{T}_\varepsilon$, and that m is a lower bound on $\mathrm{d}(v, \hat{S} \setminus \mathrm{N}(f))$ for any point $v \in \hat{S}$ and any facet of $\hat{S}$ containing v.

Let $p \in \partial S_{\mathcal{D}}$, and let f be a facet of $\hat{S}$ that contains p. We set $d' = \frac{1}{2}\ \min\{\delta,\ m,\ r_f\}$, and we call B_p^1 and B_p^2 the balls of center p and radii d' and $2d'$ respectively. Let p_f be a point of $\tilde{S}$ lying in the surface Delaunay ball of f. Such a point exists since $\tilde{S}$ is dense in S, by Corollary 2.2.

Let q be a point of $S \cap B_p^1$ that does not belong to $S_{\mathcal{D}}$. We will show that q belongs to $S_{\mathcal{I}}$, which will prove the lemma. Let $K(q)$ be the double cone of apex q, of axis aligned with $\mathbf{n}(p_f)$ and of half-angle θ. Since $q \in B_p^2$, $K(q)$ intersects ∂B_p^2 along two spherical patches $C_1(q)$ and $C_2(q)$, such that every connected curve included in $K(q)$ and joining $C_1(q)$ to $C_2(q)$ passes through q.

Like in the proof of Lemma 2.28, we would like to show that q lies on a flow line included in $K(q)$ and joining $C_1(q)$ to $C_2(q)$. However, no flow line passes through q, because the latter lies on S. Therefore, we will build a curve $L(q)$ that will serve as flow line.

Since $q \notin S_{\mathcal{D}}$, $\mathrm{d}(q, \hat{S})$ is positive. Let η be any positive value less than $\min\{d' - \mathrm{d}(p,q),\ \mathrm{d}(q, \hat{S})\}$. We pick up two points q_+ and q_- from $(K(q) \cap B(q,\eta)) \setminus S$, one on either side of S (say $q_+ \in \mathcal{O}^+$ and $q_- \in \mathcal{O}^-$). Note that two such points exist. Indeed, we have

$$B(q,\eta) \subseteq B_p^1 \subseteq B(c_f, \frac{3}{2}\, r_f)$$

which by $\hat{\mathrm{I}}3$ is included in $S \cap B(c_f, \mathrm{lr}_k(c_f))$. As a consequence, $S \cap B(q,\eta)$ lies outside the double cone $K'(q)$ of apex q, of axis $\mathbf{n}(p_f)$ and of half-angle $\frac{\pi}{2} - \theta$, by the cocone lemma 2.12. Therefore, $S \cap B(q,\eta)$ lies outside $K(q) \cap K'(q)$, and q_+ and q_- are well defined.

Let $L(q)$ be the union of $\Lambda(q_+)$, $\Lambda(q_-)$, and of the line segments $[q, q_+]$ and $[q, q_-]$:

$$L(q) = \Lambda(q_+) \cup [q_+, q] \cup [q, q_-] \cup \Lambda(q_-)$$

We call $L'(q)$ the connected component of $L(q) \cap B_p^2$ that contains q. A straightforward adaptation of the proof of Claim 2.28.2, with $q' = q$ and $\Lambda'(q)$ replaced by $L'(q)$, shows that $L'(q)$ lies in $K(q)$ and joins $C_1(q)$ to $C_2(q)$. Moreover, Claim 2.28.3, applied here with $q' = q$, shows that every connected curve included in B_p^2 and joining $C_1(q)$ with $C_2(q)$ intersects the neighborhood $\mathrm{N}(f)$ of f in $\hat{S}$. This is the case in particular for $L'(q)$, thus $L(q) \cap \hat{S} \neq \emptyset$. It follows that $\Lambda(q_+)$ or $\Lambda(q_-)$ intersects $\hat{S}$, since $[q_+, q]$ and $[q_-, q]$ do not because they are included in $B(q, \mathrm{d}(q, \hat{S}))$. Therefore, $q_+ \in \mathcal{I}$ or $q_- \in \mathcal{I}$, which means that $B(q,\eta) \cap \mathcal{I} \neq \emptyset$. Since this is true for any sufficiently small positive value η, q belongs to the topological closure of $\mathcal{I}$. Notice that, since $S \cap \mathcal{I} = \emptyset$, $S_{\mathcal{I}}$ is also the intersection of S with the topological closure of $\mathcal{I}$. Therefore, $q \in S_{\mathcal{I}}$, which concludes the proof of Lemma 2.32. □

2.2.2.4 End of the proof of Theorem 2.23

According to Lemma 2.31, for any point $p \in S$:

- either p belongs to $\hat{S}$, which means that $\mathrm{d}(p, \hat{S}) = 0$;
- or p belongs to $\partial\mathcal{I}$, which means that for any $\eta > 0$ there exists some point $p_\eta \in B(p, \eta)$ such that $\Lambda(p_\eta) \cap \hat{S} \neq \emptyset$. Let $p'_\eta \in \Lambda(p_\eta) \cap \hat{S}$ and let $t_\eta \geq 0$ be the time at which $\Lambda(p_\eta)$ reaches p'_η. By Propositon 2.25 and Lemma 2.27, we have $t_\eta \leq \frac{\varepsilon - \mathrm{d}_S(p_\eta)}{\cos^2\theta} < \frac{\varepsilon}{\cos^2\theta}$. Since by F2 ϕ is 1-Lipschitz as a function of time, we deduce that $\mathrm{d}(p_\eta, p'_\eta) < \frac{\varepsilon}{\cos^2\theta}$, which implies that $\mathrm{d}(p, \hat{S}) < \eta + \frac{\varepsilon}{\cos^2\theta}$. Since this is true for any $\eta > 0$, $\mathrm{d}(p, \hat{S})$ is at most $\frac{\varepsilon}{\cos^2\theta}$.

Therefore, no point of S is farther than $\frac{\varepsilon}{\cos^2\theta}$ from $\hat{S}$, which concludes the proofs of Proposition 2.26 and Theorem 2.23.

2.2.3 Isotopy

Theorem 2.33 *Let S be a k-Lipschitz surface. For any finite point set $E \subset S$ and any subcomplex $\hat{S}$ of $\mathrm{Del}_{|S}(E)$ that verifies Î1–Î4, $\hat{S}$ is isotopic to S.*

This theorem, combined with Corollary 2.22, yields the following

Corollary 2.34 *Let S be a k-Lipschitz surface and $E \subset S$ be a finite point set. If E verifies H1 and H2.0–H2.1, then every layer of $\mathrm{Del}_{|S}(E)$ is isotopic to the connected component of S on which its vertices lie. As a consequence, any subset of $\mathrm{Del}_{|S}(E)$ made of one layer per connected component of S is isotopic to S.*

This result is of high practical interest. Indeed, if an algorithm can generate a point sample satisfying H1 and H2.1, then, after a simple manifold extraction procedure, the result is isotopic to the surface S (and still at Hausdorff distance at most $\frac{\varepsilon}{\cos^2\theta}$ from S, by Corollary 2.24). See Part B for more detail.

Let us now prove Theorem 2.33. By Theorem 2.23, the Hausdorff distance between $\hat{S}$ and S is bounded by $\frac{\varepsilon}{\cos^2\theta}$. To show that $\hat{S}$ is isotopic to S, we will use the following result, proved by Chazal and Lieutier[2]:

Theorem 2.35 [34, Th. 6.2]
Let $\mathcal{O}$ and $\hat{\mathcal{O}}$ be two open subsets of $\mathbb{R}^3$ of positive weak feature size, whose boundaries $\partial\mathcal{O}$ and $\partial\hat{\mathcal{O}}$ are compact embedded surfaces. If $\mathrm{d}_{\mathcal{H}}(\partial\mathcal{O}, \partial\hat{\mathcal{O}}) < \frac{1}{2} \min\{wfs(\mathcal{O}),\ wfs(\hat{\mathcal{O}})\}$, then $\partial\mathcal{O}$ and $\partial\hat{\mathcal{O}}$ are isotopic.

In our context, all we have to do is to show that the Hausdorff distance between $\hat{S}$ and S is less than half the weak feature size of $\mathcal{O} = \mathbb{R}^3 \setminus S$ and less than half the weak feature size of $\hat{\mathcal{O}} = \mathbb{R}^3 \setminus \hat{S}$, and then to apply Theorem 2.35. Using Theorem 2.23, it suffices to prove the two following lemmas:

[2]Note that the original theorem [34, Th. 6.2] requires that the open sets $\mathcal{O}$ and $\hat{\mathcal{O}}$ be bounded. However, since $\partial\mathcal{O}$ and $\partial\hat{\mathcal{O}}$ are compact, it is possible to bound $\mathcal{O}$ and $\hat{\mathcal{O}}$ with a sufficiently large sphere Σ while keeping wfs($\mathcal{O}$) and wfs($\hat{\mathcal{O}}$) unchanged, according to Lemma 2.17. Then, by [34, Th. 6.2], $\Sigma \cup \partial\mathcal{O}$ and $\Sigma \cup \partial\hat{\mathcal{O}}$ are isotopic, which means that $\partial\mathcal{O}$ and $\partial\hat{\mathcal{O}}$ are also isotopic.

Lemma 2.36 *$\frac{\varepsilon}{\cos^2\theta}$ is less than half the weak feature size of $\mathcal{O}$.*

Proof. Since by $\hat{\text{I}}3$ we have $\varepsilon < \frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S)$, $\frac{\varepsilon}{\cos^2\theta}$ is less than $\frac{1}{4}\ \mathrm{lr}_k(S)$. Moreover, by Theorem 2.16, $\mathrm{lr}_k(S)$ is at most 2 wfs($\mathcal{O}$). Thus, $\frac{\varepsilon}{\cos^2\theta} < \frac{1}{2}$ wfs($\mathcal{O}$). □

Lemma 2.37 *$\frac{\varepsilon}{\cos^2\theta}$ is less than half the weak feature size of $\hat{\mathcal{O}}$.*

Proof. Let p be a point of $\hat{S}$ and let f be a facet of $\hat{S}$ that contains p. By $\hat{\text{I}}3$, the surface Delaunay ball $B(c_f, r_f)$ of f has a radius $r_f \leq \varepsilon$. Let p_f be a point of $\tilde{S} \cap B(c_f, r_f)$. Such a point exists since $c_f \in S$ and $\tilde{S}$ is dense in S, by Corollary 2.2.

For any point $q \in \hat{S} \cap B(p, \mathrm{lr}_k(S) - 3\varepsilon)$ and any facet $f' \in \hat{S}$ containing q, f' is included in a surface Delaunay ball of radius at most ε (by $\hat{\text{I}}3$), hence $f' \subset B(p, \mathrm{lr}_k(S) - \varepsilon) \subset B(c_f, \mathrm{lr}_k(S))$. Moreover, the radius-edge ratio of f is bounded by ϱ, by $\hat{\text{I}}4$. Therefore, by the triangle normal lemma 2.13 and the orientation convention 2.19, we have $(\mathbf{n}(f'), \mathbf{n}(p_f)) < \arcsin(2\varrho \sin\theta)$. Since this is true for any $q \in \hat{S} \cap B(p, \mathrm{lr}_k(S) - 3\varepsilon)$, and since by $\hat{\text{I}}1$ $\hat{S}$ is a polyhedron without boundary, Lemma 2.8 states that $\hat{S} \cap B(p, \mathrm{lr}_k(S) - 3\varepsilon)$ is the graph of a k'-Lipschitz function, for some constant k'. It follows that $\hat{S}$ is a k'-Lipschitz surface, with $\mathrm{lr}_{k'}(\hat{S}) \geq \mathrm{lr}_k(S) - 3\varepsilon$. By Theorem 2.16, wfs($\hat{\mathcal{O}}$) is then at least $\frac{1}{2}\ \mathrm{lr}_{k'}(\hat{S}) \geq \frac{1}{2}\ (\mathrm{lr}_k(S) - 3\varepsilon)$, which is greater than $\frac{2\,\varepsilon}{\cos^2\theta}$ since $\varepsilon < \frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S)$. □

Observe that, to bound the weak feature size of $\hat{\mathcal{O}}$ from below, it is necessary to control the normals of the facets of $\hat{S}$ (Lemma 2.37). This requires that the radius-edge ratios of the facets of $\hat{S}$ be bounded (Lemma 2.13). As a consequence, Hyp. $\hat{\text{I}}4$ is mandatory for our approach to work.

2.3 Loose ε-samples and ε-samples

We showed in the smooth case (Lemmas 1.22 and 1.23) that, for sufficiently small ε, any ε-sample is a loose ε-sample and any loose ε-sample is an $\varepsilon(1 + O(\varepsilon))$-sample, where the constant hidden behind the O depends on the surface S. We will prove analogous results for the nonsmooth case. The proofs are quite similar to those of Section 1.3. Nevertheless, we recall them for completeness.

First, we show that any ε-sample is a loose ε-sample. Trivially, an ε-sample satisfies Condition 1 of Definition A.1. Therefore, we need only to show that Condition 2 is also satisfied:

Lemma 2.38 *Let E be an ε-sample of a k-Lipschitz surface S. If $k < \tan\frac{\pi}{8}$ and $\varepsilon \leq \frac{1}{2}\ \mathrm{lr}_k(S)$, then every point of E is a vertex of $\mathrm{Del}_{|S}(E)$. As a consequence, E is a loose ε-sample of S.*

Proof. Let $v \in E$. We consider the cell $\mathrm{V}(v)$ in the Voronoi diagram of E. We will prove that at least one edge of the boundary of $\mathrm{V}(v)$ intersects S, which means that v is incident to at least one facet of $\mathrm{Del}_{|S}(E)$.

Since E is an ε-sample of S, no point of $S \cap \mathrm{V}(v)$ is farther than ε from v. Thus, $S \cap \mathrm{V}(v)$ is included in $B(v, \varepsilon)$, which is a subset of $B(v, \mathrm{lr}_k(S))$ since $\varepsilon < \mathrm{lr}_k(S)$, by hypothesis. This implies that $S \cap \mathrm{V}(v)$ is the graph of a bivariate function. Since it is compact, it has a non-empty boundary. Moreover, its boundary is included in $\partial\mathrm{V}(v)$ because S has no boundary. Therefore, $S \cap \partial\mathrm{V}(v) \neq \emptyset$.

Let f be a facet of $\partial \mathrm{V}(v)$ that intersects S. We will show that $f \cap S$ contains no cycle, which implies that $S \cap \partial f \neq \emptyset$, hereby proving the lemma.

Let us assume for a contradiction that $f \cap S$ contains a cycle C. We call w the point of E such that $f = \mathrm{V}(v) \cap \mathrm{V}(w)$, and we consider a point p of $f \cap S$. Since E is an ε-sample, we have $\mathrm{d}(p, v) = \mathrm{d}(p, w) \leq \varepsilon$. Hence, $\mathrm{d}(v, w) \leq 2\varepsilon \leq \mathrm{lr}_k(S)$. Let $p_v \in \tilde{S} \cap B(v, \mathrm{lr}_k(S))$. Such a point exists since $\tilde{S}$ is dense in S, by Corollary 2.2. According to the cocone lemma 2.12, the angle between $\mathbf{n}(p_v)$ and the line (v, w) is at least $\frac{\pi}{2} - \theta$. Equivalently, the angle between $\mathbf{n}(p_v)$ and the plane $\mathrm{aff}(f)$ is at most θ.

Let $\mathbf{l}$ be the unit vector of $\mathrm{aff}(f)$ whose angle with $\mathbf{n}(p_v)$ is minimal. We have $(\mathbf{n}(p_v), \mathbf{l}) \leq \theta$. Let l be a line parallel to $\mathbf{l}$ that intersects C at least twice. Such a line exists because C is a cycle. Let q and q' be two points of intersection. We assume without loss of generality that $\mathbf{l} \cdot (\mathbf{q}' - \mathbf{q}) > 0$. Then, we have:

$$(\mathbf{n}(p_v),\ \mathbf{q}' - \mathbf{q}) \leq \theta \tag{2.16}$$

Now, by the cocone lemma 2.12, the angle between $\mathbf{n}(p_v)$ and the line (q, q') is at least $\frac{\pi}{2} - \theta$. This contradicts (2.16) because $2\theta = 4 \arctan k < \frac{\pi}{2}$. Therefore, the lemma is proved. □

Conversely, Corollary 2.24 allows us to prove that loose ε-samples are $O(\varepsilon)$-samples.

Lemma 2.39 *Let E be a finite subset of a k-Lipschitz surface S. If E verifies H1 and H2.0–H2.1, then E is an $\varepsilon\sqrt{1 + \frac{1}{\cos^4 \theta}}$-sample of S.*

Proof. Let p be a point of S and p' a nearest neighbor of p on $\mathrm{Del}_{|S}(E)$. By Corollary 2.24, $\mathrm{d}(p, p')$ is at most $\frac{\varepsilon}{\cos^2 \theta}$. If p' is a vertex of $\mathrm{Del}_{|S}(E)$, then $p' \in E$ and the result is proved. Else, let f be a facet of $\mathrm{Del}_{|S}(E)$ that contains p', and $v \neq p'$ a vertex of f closest to p'. By hypothesis, the circumradius of f is at most ε. Thus, $\mathrm{d}(p', v) \leq \varepsilon$, since $p' \in f$. If p' belongs to an edge e of f, then the lines $\mathrm{aff}(e)$ and (p, p') are perpendicular. Thus,

$$\mathrm{d}(p, E) \leq \mathrm{d}(p, v) = \sqrt{\mathrm{d}^2(p, p') + \mathrm{d}^2(p', v)} \leq \sqrt{\frac{\varepsilon^2}{\cos^4 \theta} + \varepsilon^2}$$

and the lemma is proved. Else, p' belongs to the relative interior of f, thus the plane $\mathrm{aff}(f)$ is perpendicular to (p, p'). By the same computation as above, we get $\mathrm{d}(p, E) \leq \sqrt{\frac{\varepsilon^2}{\cos^4 \theta} + \varepsilon^2}$, which proves the lemma. □

Open question 2.40 *In the above lemmas, and more generally in the results of Chapter 2, the point set E is assumed to be a (loose) ε-sample of S. Can similar results be proved for (loose) $\mu\mathrm{lr}_k$-samples of S, whose densities depend on the 1-Lipschitz function lr_k?*

Chapter 3

Size of (Loose) ε-Samples

Erickson [65] has shown that $\Omega\left(\frac{1}{\mu^2}\iint_S \frac{dS}{\mathrm{d}_{\mathrm{M}}^2}\right)$ is a lower bound on the number of points of any $\mu\mathrm{d}_{\mathrm{M}}$-sample of a smooth surface S, with $\mu < \frac{1}{5}$. In this chapter, we extend Erickson's result to the more general case where S is a Lipschitz surface. Instead of restricting ourselves to a specific sizing field, we consider a generic κ-Lipschitz sizing field $\sigma : S \to \mathbb{R}^+$, with $\kappa < 1$. Our bound holds for sizing fields that are sufficiently small with respect to lr_k.

We also introduce the notion of *sparse sample* (Definition 3.2), and we work out an upper bound on the size of such samples. In the specific case where σ is uniform and equal to a constant ε, the upper and lower bounds coincide, up to a constant factor. Thus, sparse (loose) ε-samples have an optimal size.

3.1 Surface integral

In the sequel, S is a k-Lipschitz surface. Let Δ, m and $g^i : (x^i, y^i) \to \mathbb{R}$, $1 \leq i \leq m$, be as in Definition 2.1. We say that a function $f : S \to \mathbb{R}$ is integrable if $\forall i$, $\iint_\Delta |f(x^i,\ y^i,\ g^i(x^i, y^i))|\ dx^i dy^i < +\infty$. Let us explain what we mean by the *surface integral* of f.

Since g^i is differentiable almost everywhere in Δ, it has partial derivatives $\frac{\partial g^i}{\partial x^i}$ and $\frac{\partial g^i}{\partial y^i}$ almost everywhere in Δ. So does the map $\psi^i : (x^i, y^i) \mapsto (x^i,\ y^i,\ g^i(x^i, y^i))$. Therefore, at almost every point of Δ we can define a *surface element* dS:

$$dS = \left\| \frac{\partial \psi^i}{\partial x^i} \times \frac{\partial \psi^i}{\partial y^i} \right\| dx^i dy^i = \sqrt{1 + \left(\frac{\partial g^i}{\partial x^i}\right)^2 + \left(\frac{\partial g^i}{\partial y^i}\right)^2}\, dx^i dy^i \tag{3.1}$$

Note that dS depends on the choice of the coordinate systems $\{(x^i, y^i, z^i)\}$ and their associated Lipschitz functions g^i in Definition 2.1. However, it is proved in [94, §3.1] that, given any partition of unity $\{\phi^i\}$ on S subordinate to the covering[1] $\{g^i(\Delta)\}$, the following quantity:

$$\sum_{i=1}^{m} \iint_\Delta f(x^i,\ y^i,\ g^i(x^i, y^i))\ \phi^i(x^i,\ y^i,\ g^i(x^i, y^i))\ dS$$

does not depend on the partition of unity $\{\phi^i\}$ nor on the coordinate systems $\{(x^i, y^i, z^i)\}$. We call this quantity the surface integral of f over S, noted $\iint_S f\ dS$, and we notice that it is well-defined since f is integrable.

[1] See [14, §3.2] for a definition of a partition of unity subordinate to a given covering.

This definition coincides with the usual definition of surface integral in the smooth case [14, §6.6]. Moreover, by the area formula [67, §3.2.3], $\iint_S 1 \, dS$ equals $\mathcal{H}^2(S)$, the 2-dimensional Hausdorff measure of S. As mentioned in the introduction of the thesis, $\mathcal{H}^2$ generalizes the notion of area of a surface, hence in the sequel $\mathcal{H}^2(S)$ will be denoted $\mathrm{Area}(S)$, by analogy with the smooth case.

3.2 Lower bound

Theorem 3.1 *If E is a σ-sample of S, with $\sigma \leq (1-\kappa)\,\mathrm{lr}_k$, then*

$$|E| \geq \frac{1}{\pi} \left(\frac{1-\kappa}{1+\kappa}\right)^2 \frac{1}{1+k^2} \iint_S \frac{dS}{\sigma^2}$$

Proof. For any $v \in E$, we call $\mathrm{V}_S(v)$ the intersection of S with the cell of v in $\mathrm{Vor}(E)$. $\mathrm{V}_S(v)$ is the set of all points of S that have v as nearest neighbor among the points of E. Since $S = \bigcup_{v\in E} \mathrm{V}_S(v)$, we have:

$$\iint_S \frac{dS}{\sigma^2} = \sum_{v\in E} \iint_{\mathrm{V}_S(v)} \frac{dS}{\sigma^2} \tag{3.2}$$

Since E is a σ-sample, for any $v \in E$ and any $p \in \mathrm{V}_S(v)$, $\mathrm{d}(p,v)$ is bounded by $\sigma(p)$. Hence, $\sigma(p)$ is at least $\frac{\sigma(v)}{1+\kappa}$ and at most $\frac{\sigma(v)}{1-\kappa}$, because σ is κ-Lipschitz. Therefore,

$$\forall v \in E,\ \iint_{\mathrm{V}_S(v)} \frac{dS}{\sigma^2} \leq (1+\kappa)^2 \frac{\mathrm{Area}(\mathrm{V}_S(v))}{\sigma^2(v)} \tag{3.3}$$

Moreover, $\mathrm{V}_S(v)$ is included in $B\left(v,\ \frac{\sigma(v)}{1-\kappa}\right) \subseteq B\left(v,\ \mathrm{lr}_k(v)\right)$. By definition of $\mathrm{lr}_k(v)$, $\mathrm{V}_S(v)$ is the graph of a k-Lipschitz function f_v defined over the orthogonal projection $\mathrm{V}'_S(v)$ of $\mathrm{V}_S(v)$ onto some plane Π_v. Since f_v is k-Lipschitz, the function that maps every point $(x,y) \in \mathrm{V}'_S(v)$ onto $(x,y,f_v(x,y)) \in \mathrm{V}_S(v)$ is $\sqrt{1+k^2}$-Lipschitz. Thus, by [67, §2.10.11] we have:

$$\mathrm{Area}(\mathrm{V}_S(v)) \leq (1+k^2)\ \mathrm{Area}(\mathrm{V}'_S(v))$$

In addition, $\mathrm{V}'_S(v)$ is included in a disk of radius $\frac{\sigma(v)}{1-\kappa}$, hence

$$\mathrm{Area}(\mathrm{V}_S(v)) \leq \frac{\pi\ \sigma^2(v)\ (1+k^2)}{(1-\kappa)^2} \tag{3.4}$$

It follows from (3.2), (3.3) and (3.4) that

$$\iint_S \frac{dS}{\sigma^2(x)} \leq \sum_{v\in E} \pi \frac{(1+\kappa)^2}{(1-\kappa)^2}\ (1+k^2) = \pi \frac{(1+\kappa)^2}{(1-\kappa)^2}\ (1+k^2)\ |E|$$

□

In the specific case where S is $C^{1,1}$, S is k-Lipschitz for any $k > 0$ (Theorem 2.6). Therefore, without loss of generality we take $k = 1$. If $\sigma = \mu\ \mathrm{d_M}$, then $\kappa = \mu$ and σ verifies the hypothesis of Theorem 3.1 for all $\mu \leq \frac{\pi}{4+2\pi}$, by Theorem 2.6. Moreover, we have $\left(\frac{1-\mu}{1+\mu}\right)^2 \geq \left(\frac{4+\pi}{4+3\pi}\right)^2$ and $\frac{1}{1+k^2} = \frac{1}{2}$. Hence, by Theorem 3.1, any $\mu\mathrm{d_M}$-sample of S contains $\Omega\left(\frac{1}{\mu^2} \iint_S \frac{dS}{\mathrm{d}_\mathrm{M}^2}\right)$ points, where the constant hidden in the Ω does not depend on S nor on μ. This bound coincides with Erickson's bound.

In the other specific case where σ is uniform and equal to a constant ε, we have $\kappa = 0$, hence Theorem 3.1 yields $|E| \geq \frac{1}{\pi} \frac{1}{1+k^2} \frac{\text{Area}(S)}{\varepsilon^2}$. If k is bounded by some constant k_0 (which is the case when one wants to have guarantees on the topology and geometry of $\text{Del}_{|S}(E)$, as observed in Chapter 2), then $|E| = \Omega\left(\frac{\text{Area}(S)}{\varepsilon^2}\right)$, where the constant hidden in the Ω does not depend on S nor on ε. This bound holds for loose ε-samples as well, since loose ε-samples are $O(\varepsilon)$-samples, by Lemma 2.39.

3.3 Upper bound

Since adding points to an ε-sample results in another ε-sample, we cannot hope for an upper bound on ε-samples without making some additional assumptions.

Definition 3.2 *A point set $E \subset S$ is said to be* σ-sparse *if there exists some positive constant h, independent from S and from σ, such that:*

$$\forall p \in S,\ |E \cap B(p, \sigma(p))| \leq h.$$

If in addition E is a σ-sample of S, then we say that E is a sparse σ-sample *of S.*

It is clear that every σ-sparse sample is *a fortiori* a σ'-sparse sample, for any $\sigma' \leq \sigma$. Moreover,

Lemma 3.3 *If $E \subset S$ is σ-sparse, then E is also $\nu\sigma$-sparse, for any constant $\nu < \frac{1}{2\kappa}$ independent from S and from σ.*

Proof. Let $p \in S$. We will bound the number of points of E that lie in $B(p, \nu\sigma(p))$. Let $\sigma_m > 0$ be the infimum of σ over $S \cap B(p, \nu\sigma(p))$. Since σ is κ-Lipschitz, for any $q \in S \cap B(p, \nu\sigma(p))$ we have $\sigma(q) \geq \sigma(p) - \kappa\, \mathrm{d}(p, q) \geq (1 - \kappa\nu)\, \sigma(p) > \frac{\sigma(p)}{2}$. Therefore, $\sigma_m \geq \frac{\sigma(p)}{2}$.

Let $q_1 \in S \cap B(p, \sigma(p))$, and for any $i \geq 1$ let $q_{i+1} \in S \setminus \bigcup_{i'=0}^{i} B(q_{i'}, \frac{\sigma_m}{2})$ (if such a point exists). By construction, the balls $B(q_i, \frac{\sigma_m}{2})$ cover $S \cap B(p, \nu\sigma(p))$. Moreover, the q_i are farther than $\frac{\sigma_m}{2}$ from one another, hence they are centers of pairwise disjoint balls of radius $\frac{\sigma_m}{4}$. Therefore, their number is at most:

$$\frac{\text{Vol}\left(B\left(p, \nu\sigma(p)\right)\right)}{\text{Vol}\left(B\left(p, \frac{\sigma_m}{4}\right)\right)} = 64\, \nu^3\, \frac{\sigma^3(p)}{\sigma_m^3} \leq 512\, \nu^3$$

In addition, for any q_i, if $B(q_i, \frac{\sigma_m}{2})$ does not intersect S, then $|E \cap B(q_i, \frac{\sigma_m}{2})| = 0$. Otherwise, let $q_i' \in S \cap B(q_i, \frac{\sigma_m}{2})$. $B(q_i, \frac{\sigma_m}{2})$ is included in $B(q_i', \sigma_m) \subseteq B(q_i', \sigma(q_i'))$, thus $|E \cap B(q_i, \frac{\sigma_m}{2})| \leq n$, because E is σ-sparse. Since this is true for any q_i, and since the balls $B(q_i, \frac{\sigma_m}{2})$ cover $S \cap B(p, \nu\sigma(p))$, we get $|S \cap B(p, \nu\sigma(p))| \leq 512\, \nu^3\, n$. □

The above result will be useful for proving that our sampling algorithm generates sparse samples – see Section 4.2.

Theorem 3.4 *If $E \subset S$ is σ-sparse, with $\sigma \leq 2\, \text{lr}_k$ and $k < 1$, then there exists some positive constant*[2] *h, independent from S and from σ, such that*

$$|E| \leq \frac{h}{\pi}\, \left(\frac{1+k^2}{1-k^2}\right)^2 (2+\kappa)^2\, \iint_S \frac{dS}{\sigma^2}$$

[2] This constant is the same as in Definition 3.2.

Proof. Let E be a σ-sparse point set included in S. Since S is compact, Definition 3.2 implies that E is finite. For any point $p \in E$, we define $B'(p)$ as the ball centered at p of radius $\sigma(p)$. We construct a subset $E' \subset E$ iteratively as follows: let $p_1 \in E$, and for any $i > 1$, let $p_i \in E \setminus \bigcup_{1 \leq j < i} B'(p_j)$. This process terminates since E is finite. Upon termination, every point of E belongs to $B'(p_i)$, for some $p_i \in E'$. Since E is σ-sparse, there exists a constant h such that every $B'(p_i)$ contains at most h points of E (including p_i itself). Therefore, the size of E' is at least $|E|/h$.

Hence, bounding the size of E comes down to bounding the size of E'. To do so, we proceed as in the case of planar meshes [60, 99] and lower-bound the integral over S of $1/\sigma^2$. For any $p_i \in E'$, we consider the open ball $B^{1/2}(p_i)$ of center p_i and radius $\sigma(p_i)/2$. Let $D(p_i) = S \cap B^{1/2}(p_i)$. Since $D(p_i) \subseteq S$ for any $p_i \in E'$, we have

$$\iint_S \frac{dS}{\sigma^2} \geq \iint_{\bigcup\limits_{p_i \in E'} D(p_i)} \frac{dS}{\sigma^2}$$

Given $p_i \neq p_j$, p_i does not belong to $B'(p_j)$ and p_j does not belong to $B'(p_i)$. Therefore, $\mathrm{d}(p_i, p_j)$ is greater than $\max\{\sigma(p_i),\ \sigma(p_j)\}$, which is at least $\frac{1}{2}\left(\sigma(p_i) + \sigma(p_j)\right)$. It follows that $D(p_i) \cap D(p_j) = \emptyset$. Since this is true for any $p_i \neq p_j$, we have

$$\iint_{\bigcup\limits_{p_i \in E'} D(p_i)} \frac{dS}{\sigma^2} = \sum_{p_i \in E'} \iint_{D(p_i)} \frac{dS}{\sigma^2}$$

In addition, since σ is κ-Lipschitz, $\forall q \in D(p_i),\ \sigma(q) \leq \sigma(p_i) + \kappa\ \mathrm{d}(q, p_i) \leq (1 + \kappa/2)\ \sigma(p_i)$. It follows that

$$\sum_{p_i \in E'} \iint_{D(p_i)} \frac{dS}{\sigma^2} \geq \sum_{p_i \in E'} \frac{\mathrm{Area}(D(p_i))}{(1 + \kappa/2)^2\, \sigma^2(p_i)}$$

Let $p_i \in E'$. Let q_i be a point of $\tilde{S} \cap B^{1/2}(p_i)$. Such a point exists because $\tilde{S}$ is dense in S, by Corollary 2.2. Since the radius of $B^{1/2}(p_i)$ is $\frac{\sigma(p_i)}{2} \leq \mathrm{lr}_k(p_i)$, the cocone lemma 2.12 states that $D(p_i)$ lies outside the double cone $K(p_i)$ of apex p_i, of axis aligned with $\mathbf{n}(q_i)$ and of half-angle $\frac{\pi}{2} - 2\arctan k$. Let $K^+(p_i)$ and $K^-(p_i)$ be the two cones of $K(p_i)$. Since S passes through the apex p_i and has no boundary, it encloses $K^+(p_i) \cap B^{1/2}(p_i)$ or $K^-(p_i) \cap B^{1/2}(p_i)$. Therefore, the orthogonal projection of $D(p_i)$ onto $T(q_i)$ contains the orthogonal projection of $K(p_i) \cap B^{1/2}(p_i)$, which is a disk of radius $\frac{\sigma(p_i)}{2}\ \cos(2\arctan k) = \frac{\sigma(p_i)}{2}\ \frac{1-k^2}{1+k^2}$. Hence,

$$\sum_{p_i \in E'} \frac{\mathrm{Area}(D(p_i))}{(1 + \kappa/2)^2\, \sigma^2(p_i)} \ \geq\ \pi\ \left(\frac{1 - k^2}{1 + k^2}\right)^2 \frac{1}{(2 + \kappa)^2}\ |E'|$$

which yields the theorem, since $|E'| \geq |E|/h$. □

In the specific case where S is $C^{1,1}$, S is k-Lipschitz for any $k > 0$ (Theorem 2.6). Therefore, without loss of generality we take $k = \frac{1}{\sqrt{3}}$. If $\sigma = \mu\ \mathrm{d_M}$, then $\kappa = \mu$ and σ verifies the hypothesis of Theorem 3.1 for all $\mu \leq \frac{\pi}{18+3\pi}$, by Theorem 2.6. Moreover, we have $(2 + \mu)^2 \leq \left(\frac{36+7\pi}{18+3\pi}\right)^2$ and $\left(\frac{1+k^2}{1-k^2}\right)^2 = 4$. Hence, by Theorem 3.4, any $\mu\mathrm{d_M}$-sparse sample of S contains $O\left(\frac{1}{\mu^2}\ \iint_S \frac{dS}{\mathrm{d}_\mathrm{M}^2}\right)$ points,

where the constant hidden in the O does not depend on S nor on μ. This bound is optimal in view of Theorem 3.1.

In the other specific case where σ is uniform and equal to a constant ε, and k is bounded by some constant k_0, Theorem 3.4 yields $|E| = O\left(\frac{\mathrm{Area}(S)}{\varepsilon^2}\right)$, where the constant hidden in the O does not depend on S nor on ε. This bound is also optimal in view of Theorem 3.1. Therefore, sparse (loose) ε-samples have an optimal size.

Part B

Sampling and Meshing Algorithm

Introduction

In this part of the thesis, we show how our sampling condition can be turned into a simple and efficient surface mesh generator. The latter is based on a *Delaunay refinement* technique, which consists in constructing an initial mesh and then refining iteratively the elements of the mesh that do not meet some user-defined size or shape criteria. This greedy tehnique was pioneered by Ruppert [99] in the plane, and then extended (without proof of correctness) by Chew to the case of surfaces in 3-space [45]. Our mesher derives from Chew's algorithm[3]. It takes as input a user-defined parameter ε and an initial point set $E_I \subset S$, and it outputs a loose ε-sample E_F of S, together with $\mathrm{Del}_{|S}(E_F)$.

Chew did not provide any theoretical guarantees with his algorithm. Here, taking advantage of the theoretical results of Part A, we can prove that the output mesh is a manifold without boundary, with the same topological type as S and close to S for the Hausdorff distance, provided that ε is chosen sufficiently small. Specifically, ε must be less than a fraction of $\mathrm{rch}(S)$ (introduced in Section 1.1.1) when S is smooth, and less than a fraction of $\mathrm{lr}_k(S)$ (see Section 2.1.1) when S is a more general k-Lipschitz surface, for sufficiently small k. It follows that the algorithm generates provably good meshes on a wide class of shapes. Moreover, we show that the number of points sampled from S by the algorithm lies within a constant factor of the optimal.

Let us emphasize that our mesh generator maintains the 3-dimensional Delaunay triangulation of the point sample throughout the process. This is mandatory for guaranteeing the topology of the output mesh. Maintaining a whole 3-dimensional triangulation can be time-consuming in general. However, we prove in our case that the size of the data structure remains bounded throughout the process, which implies that the space and time complexities of the algorithm are quite reasonable.

A noticeable feature of the algorithm is that it needs only to know the surface S through an oracle ω that can compute the intersection of any given segment with S. Therefore, our mesh generator is generic enough to be applied in a wide variety of contexts. This claim will be illustrated in Part C.

In Chapter 4, we present our mesher, provide theoretical guarantees on its output, and analyze its space and time complexities. We also describe how to construct an initial point set $E_I \subset S$ to be used as input by the algorithm.

In Chapter 5, we introduce several modifications that can be brought to the algorithm, to enhance the quality of its output or to improve its time complexity in important special cases. In particular, we present an oracle that does not compute all the intersection points of a given segment with the surface S. This oracle can be efficiently implemented for implicit surfaces, and we prove that the algorithm has the same theoretical guarantees with this oracle as it does with the exhaustive oracle.

[3] In the sequel, we identify our mesher with Chew's algorithm, for simplicity.

Chapter 4

Chew's Algorithm

Chew's surface mesh generator [45] is a greedy incremental algorithm that inserts sample points on S and maintains the Delaunay triangulation of the sample E restricted to S, $\mathrm{Del}_{|S}(E)$. In this chapter, we describe a modified version of Chew's algorithm and use our previous results to analyze its complexity and to give guarantees on its output.

4.1 The algorithm

4.1.1 Input

Chew's algorithm takes as input:

- a surface S,
- a point set $E_I \subset S$,
- a positive parameter ε.

The surface is known only through an oracle ω that, given a line segment s, can compute a (possibly empty) subset of the intersection points of s with S: $\forall s \in \mathrm{Seg}(\mathbb{R}^3),\ \omega(s) \subseteq s \cap S$.

4.1.2 Data structure

Recall that, given a point set $E \subset S$, the Delaunay triangulation of E restricted to S, $\mathrm{Del}_{|S}(E)$, is the subcomplex of the 3-dimensional Delaunay triangulation $\mathrm{Del}(E)$ made of the facets whose dual Voronoi edges intersect S. Every point of intersection of a Voronoi edge with S is the center of a surface Delaunay ball, also called *ball of* $\mathrm{Del}_{|S}(E)$ in the sequel.

By querying the oracle ω on every Voronoi edge, the algorithm can compute a subset of $\mathrm{Del}_{|S}(E)$, called $\mathrm{Del}^{\omega}_{|S}(E)$. Notice that $\mathrm{Del}^{\omega}_{|S}(E)$ may be different from $\mathrm{Del}_{|S}(E)$, since ω is not assumed to be able to detect all the intersection points of S with the edges of the Voronoi diagram. The balls of $\mathrm{Del}_{|S}(E)$ whose centers are detected by ω are called *balls of* $\mathrm{Del}^{\omega}_{|S}(E)$.

$\mathrm{Del}^{\omega}_{|S}(E)$ is stored as a subcomplex of $\mathrm{Del}(E)$. Each time a point is added to E, the part of the Voronoi diagram that has changed after the insertion of the point has to be investigated, by querying the oracle ω on every Voronoi edge that has appeared or has been shortened.

4.1.3 Course of the algorithm

The algorithm constructs iteratively a point set E and maintains its Delaunay triangulation restricted to S.

Initially, the algorithm takes $E = E_I$ and computes $\mathrm{Del}^{\omega}_{|S}(E)$ by querying the oracle on every edge of the Voronoi diagram of E. A ball of $\mathrm{Del}^{\omega}_{|S}(E)$ is said to be *bad* if its radius is greater than ε. A facet of $\mathrm{Del}^{\omega}_{|S}(E)$ is called *bad* if it is circumscribed by a bad ball of $\mathrm{Del}^{\omega}_{|S}(E)$. The bad balls of $\mathrm{Del}^{\omega}_{|S}(E)$ are stored in a priority queue $\mathcal{Q}$ where they are sorted by decreasing radii.

Then, at each iteration, the algorithm inserts a new point in E and updates $\mathrm{Del}^{\omega}_{|S}(E)$. Each point inserted in E is the center of a bad ball of $\mathrm{Del}^{\omega}_{|S}(E)$ of largest radius, stored at the top of the priority queue $\mathcal{Q}$. The process is described in the following loop:

while $\mathcal{Q}$ is not empty {
  take the largest element of $\mathcal{Q}$;
  insert its center into E and update $\mathrm{Del}(E)$;
  update $\mathrm{Del}^{\omega}_{|S}(E)$ by querying ω on all the Voronoi edges that have changed or appeared:
    delete from $\mathrm{Del}^{\omega}_{|S}(E)$ the former Delaunay facets that no longer belong to $\mathrm{Del}^{\omega}_{|S}(E)$;
    add to $\mathrm{Del}^{\omega}_{|S}(E)$ the new Delaunay facets that belong to $\mathrm{Del}^{\omega}_{|S}(E)$;
  update $\mathcal{Q}$ by:
    deleting all the elements of $\mathcal{Q}$ which are no longer balls of $\mathrm{Del}^{\omega}_{|S}(E)$;
    adding all the new balls of $\mathrm{Del}^{\omega}_{|S}(E)$ whose radius is greater than ε;
}

The algorithm stops at the end of the loop, when there are no more bad balls of $\mathrm{Del}^{\omega}_{|S}(E)$. Upon termination, the point set E is renamed E_F, and it is returned as well as $\hat{S} = \mathrm{Del}^{\omega}_{|S}(E_F)$.

Notice that the number of intersection points of a Voronoi edge with the surface S may be arbitrarily large. To prevent $\mathcal{Q}$ from growing too much, we store at most one ball of $\mathrm{Del}^{\omega}_{|S}(E)$ (namely, the one of largest radius) per Voronoi edge. This way, the size of $\mathcal{Q}$ is bounded by the number of Voronoi edges.

4.2 Guarantees on the output

In this section, we assume that the algorithm terminates, and we show how the structural results of Part A give theoretical guarantees on its output. One crucial condition is that $\mathrm{Del}^{\omega}_{|S}(E_F)$ have vertices on all the connected components of S. This condition is called I2 in Chapter 1 and $\hat{\text{I}}2$ in Chapter 2. In Section 4.2.1 we will introduce the so-called *persistent facets*, which will help to fulfill I2 and $\hat{\text{I}}2$.

Another condition is that the oracle be able to compute all the intersection points of S with any segment: $\forall s \in \mathrm{Seg}(\mathbb{R}^3)$, $\omega(s) = s \cap S$. Such an oracle is said to be *exhaustive*. Non-exhaustive oracles will be introduced later, in Section 5.1 and Chapter 8. From now on, and until the end of Chapter 4, we assume that the oracle is exhaustive, which means that $\mathrm{Del}^{\omega}_{|S}(E) = \mathrm{Del}_{|S}(E)$ for any point set E.

To review the properties of the output of the algorithm, we will distinguish between the case where the surface S has a positive reach (Section 4.2.2) and the more general case where S is k-Lipschitz (Section 4.2.3). Our theoretical results hold for ε less than a fraction of $\mathrm{rch}(S)$ when S has a positive

reach, and ε less than a fraction of $\mathrm{lr}_k(S)$ when S is k-Lipschitz. Therefore, it is necessary in practice to have at our disposal a positive lower bound ε_S on $\mathrm{rch}(S)$ or $\mathrm{lr}_k(S)$. We will see in Part C how such a bound can be computed in several applications.

Before going more into detail, let us first notice that the set of points inserted by the algorithm satisfies a sparseness condition, stated below:

Lemma 4.1 *The points of $E_F \setminus E_I$ are farther than ε from one another and from the points of E_I.*

Proof. Let p belong to $E_F \setminus E_I$. When p is inserted in E by the algorithm, it is the center of a Delaunay ball of radius greater than ε, hence its distance to E_I and to the points of $E_F \setminus E_I$ inserted before is greater than ε. The result follows. □

Lemma 4.1 will be instrumental in proving that the algorithm generates sparse samples.

4.2.1 Persistent facets

Since the algorithm is greedy, its behaviour depends highly on the initial point sample E_I. Figure 3 shows an example where $\mathrm{Del}_{|S}(E_I)$ is empty. In this case, the priority queue $\mathcal{Q}$ is empty at the beginning of the process, and the algorithm stops immediately. As a consequence, the output point sample E_F is equal to E_I. To prevent $\mathrm{Del}_{|S}(E)$ from vanishing during the process, we introduce the notion of persistent facet.

Definition 4.2 *Given a positive constant λ, a* λ-persistent facet *is a facet of $\mathrm{Del}_{|S}(E_I)$ circumscribed by a ball of $\mathrm{Del}_{|S}(E_I)$ of radius at most $\frac{\lambda}{2}$.*

Persistent facets are interesting in our context because they share a nice persistence property, illustrated in Figure 4.1 and stated in the following lemma:

Lemma 4.3 *During the course of the algorithm, all λ-persistent facets remain in $\mathrm{Del}_{|S}(E)$ as long as the radius of the largest ball of $\mathrm{Del}_{|S}(E)$ remains greater than λ.*

Proof. Let f be a λ-persistent facet and $B = B(c, r)$ be a ball of $\mathrm{Del}_{|S}(E_I)$ circumscribing f such that $r \leq \frac{\lambda}{2}$. Assume that, at some stage, the algorithm inserts in E a point c' that belongs to the interior of B. We know that c' is the center of a surface Delaunay ball B' circumscribing some facet f' of $\mathrm{Del}_{|S}(E)$. Notice that we may possibly have $B' = B$ and $f' = f$, but it is not necessarily the case. Since B' is a Delaunay ball, its interior does not contain the vertices of f. Hence, the radius r' of B' must be at most $2r \leq \lambda$. Moreover, B' is the ball of $\mathrm{Del}_{|S}(E)$ of largest radius, since its center is being inserted in E. Therefore, all the balls of $\mathrm{Del}_{|S}(E)$ have radii at most λ. □

We will see that assuming that E_I contains λ-persistent facets, for sufficiently small λ, ensures that $\mathrm{Del}_{|S}(E)$ remains non-empty throughout the process.

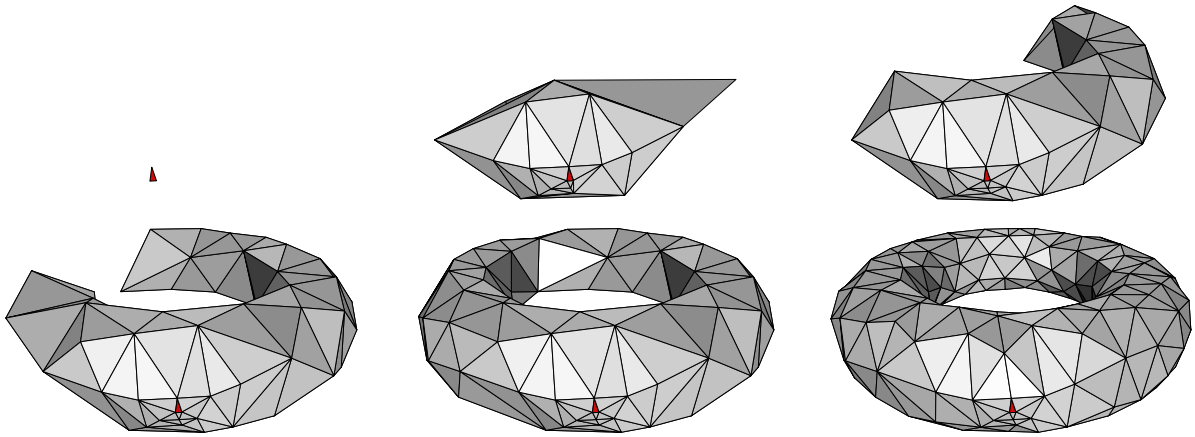

Figure 4.1: Meshing of a torus: the persistent facet (upper-left corner) remains in the restricted Delaunay triangulation throughout the process.

4.2.2 Surfaces with positive reach

Let S have a positive reach and let $\mu_0 = 0.16$ and $\mu_1 = \frac{\pi}{16+2\pi} \approx 0.14$ be the constants defined in Chapter 1, Sections 1.2 and 1.4. Recall that the oracle ω is assumed to be exhaustive, hence $\mathrm{Del}^{\omega}_{|S}(E) = \mathrm{Del}_{|S}(E)$ for any point set E.

Theorem 4.4 *If E_I and ε verify:*

(α) $\mathrm{Del}_{|S}(E_I)$ *has λ-persistent facets on all the connected components of S, with $\lambda \leq \mu_0 \operatorname{rch}(S)$,*

(β) $\varepsilon \leq \mu_0 \operatorname{rch}(S)$,

then E_F is a loose ε-sample of S. As a consequence,

(a) $\mathrm{Del}_{|S}(E_F)$ *is isotopic to S,*

(b) the Fréchet distance between $\mathrm{Del}_{|S}(E_F)$ *and S is at most* $\frac{\varepsilon^2}{\operatorname{rch}(S)}$,

(c) E_F is an $\varepsilon\sqrt{1+\frac{\varepsilon^2}{\operatorname{rch}(S)^2}}$*-sample of S, hence* $|E_F| = \Omega\left(\frac{\mathrm{Area}^2(S)}{\varepsilon^2}\right)$.

(d) $\mathrm{Del}_{|S}(E_F)$ *approximates S in terms of normals and area within an error $O(\varepsilon)$.*

(e) if $\varepsilon < \mu_1 \operatorname{rch}(S)$, then S is covered by the balls of $\mathrm{Del}_{|S}(E_F)$,

Moreover, if E_I is ε-sparse, then E_F is $\varepsilon\sqrt{1+\frac{\varepsilon^2}{\operatorname{rch}(S)^2}}$*-sparse, hence* $|E_F| = O\left(\frac{\mathrm{Area}(S)}{\varepsilon^2}\right)$.

Proof. Since the oracle ω is exhaustive, every ball of $\mathrm{Del}_{|S}(E_F)$ has a radius bounded by ε, which means that E verifies Condition 1 of Definition A.1.

- If all the λ-persistent facets remain in $\mathrm{Del}_{|S}(E)$ until the end of the process, then $\mathrm{Del}_{|S}(E_F)$ has vertices on all the connected components of S, hence Condition 2 of Definition A.1 is satisfied and E_F is a loose ε-sample of S.

- Otherwise, let p be the first center of a surface Delaunay ball of radius at most λ that is inserted in E. By Lemma 4.3, before p is inserted in E, all λ-persistent facets remain in $\mathrm{Del}_{|S}(E)$. We consider the point set E right before the insertion of p. At that stage, p is the center of a ball of $\mathrm{Del}_{|S}(E)$ of largest radius, hence all the balls of $\mathrm{Del}_{|S}(E)$ have radii at most λ. Moreover, all the λ-persistent facets are still in $\mathrm{Del}_{|S}(E)$. Therefore, E is a loose λ-sample of S. Since $\lambda \leq \mu_0 \operatorname{rch}(S)$, E is a

$\lambda\sqrt{1+\frac{\lambda^2}{\mathrm{rch}(S)^2}}$-sample, by Lemma 1.23. It follows that $E_F \supseteq E$ is also a $\lambda\sqrt{1+\frac{\lambda^2}{\mathrm{rch}(S)^2}}$-sample. Since $\lambda\sqrt{1+\frac{\lambda^2}{\mathrm{rch}(S)^2}} \leq \mu_0\sqrt{1+\mu_0^2}\,\mathrm{rch}(S) < \frac{\pi}{6+\pi}\,\mathrm{rch}(S)$, E_F is a loose ε-sample, by Lemma 1.22.

This property, combined with Theorems 1.14, 1.18 and Lemma 1.23, proves (a), (b) and (c). The lower bound on $|E_F|$ in (c) comes from Theorem 3.1. Moreover, (d) follows from the triangle normal lemma 1.5 and from a theorem by Morvan and Thibert [92]. Finally, Theorem 1.24 proves (e).

Let us now assume that E_I is ε-sparse and prove that E_F is ε'-sparse, with $\varepsilon' = \varepsilon\sqrt{1+\frac{\varepsilon^2}{\mathrm{rch}(S)^2}}$. To show that E_F is ε'-sparse, we will bound the number of points of E_F that lie in $B(p,\varepsilon')$, for any $p \in S$. By (β), we have $\varepsilon' \leq \varepsilon\sqrt{1+\mu_0^2}$. Thus, by Lemma 3.3, E_I is ε'-sparse[1], which implies that there exists a positive constant h, independent from S and ε, such that $|E_I \cap B(p,\varepsilon')| \leq h$. Moreover, by Lemma 4.1, the points of $E_F \setminus E_I$ are farther than ε from one another, hence they are centers of pairwise disjoint open balls of radius $\frac{\varepsilon}{2}$. For any $q \in B(p,\varepsilon') \cap E_F \setminus E_I$, the open ball $B(q,\frac{\varepsilon}{2})$ is included in $B(p,\varepsilon'+\frac{\varepsilon}{2})$. Therefore, the number of points of $E_F \setminus E_I$ that lie in $B(p,\varepsilon')$ is bounded by:

$$\frac{\mathrm{Vol}\left(B\left(p,\varepsilon'+\frac{\varepsilon}{2}\right)\right)}{\mathrm{Vol}\left(B\left(q,\frac{\varepsilon}{2}\right)\right)} \;\leq\; \left(1+2\sqrt{1+\mu_0^2}\right)^3$$

Hence, $|E_F \cap B(p,\varepsilon')| \leq h + \left(1+2\sqrt{1+\mu_0^2}\right)^3$, which means that E_F is ε'-sparse. The upper bound on its size follows then from Theorem 3.4. □

According to Theorem 4.4, to construct a PL approximation of S within a Fréchet distance $\delta > 0$, it suffices to compute a positive constant $\varepsilon_S \leq \mathrm{rch}(S)$, and then to apply the algorithm with $\varepsilon = \sqrt{\delta\,\varepsilon_S}$. The size of the output point set is then $\Theta_S\left(\varepsilon^{-2}\right) = \Theta_S\left(\delta^{-1}\right)$. This bound is optimal (up to a constant factor depending on S) when the object bounded by S is convex, since in this case any PL approximation of S within Hausdorff distance δ has $\Omega_S\left(\delta^{-1}\right)$ vertices [75].

Remark that λ is not lower-bounded in Condition (α) of Theorem 4.4. As a consequence, it may be (far) smaller than ε. In such a case, the facets of the output mesh are smaller in the vicinity of persistent facets. An illustration is given in Figure 4.1, where we took $\lambda = \frac{\varepsilon}{7}$ to emphasize the phenomenon. However, once the input point sample E_I is fixed, ε can be chosen arbitrarily small. In particular, if the user chooses an ε less than $\frac{\lambda}{2}$, then the λ-persistent facets are not visible in the output mesh. This is the case for the models presented in Part C, *e.g.* in Figure 6.1. In Section 4.4, we will explain how to construct persistent facets that are not too small compared to $\mu_0\,\mathrm{rch}(S)$.

4.2.3 Lipschitz surfaces

Let S be a k-Lipschitz surface. We define $\theta = 2\arctan k$, for convenience. Using the theoretical results of Chapter 2, we can guarantee the quality of the output of Chew's algorithm for $\theta \leq \arctan\frac{1}{3\sqrt{3}} \approx 10.9°$, *i.e.* for $k \leq \tan\left(\frac{1}{2}\arctan\frac{1}{3\sqrt{3}}\right) \approx 0.1$. Recall that the oracle ω is assumed to be exhaustive, hence $\mathrm{Del}^{\omega}_{|S}(E) = \mathrm{Del}_{|S}(E)$ for any point set E.

Theorem 4.5 *If E_I and ε verify:*

($\hat{\alpha}$) $\mathrm{Del}_{|S}(E_I)$ has λ-persistent facets on all the connected components of S, with $\lambda < \frac{\cos^2\theta}{4+3\cos^2\theta}\,\mathrm{lr}_k(S)$,

[1] The sizing field $p \mapsto \varepsilon$ is κ-Lipschitz for any $\kappa > 0$, hence the hypothesis of Lemma 3.3 is satisfied in our context.

($\hat{\beta}$) $\varepsilon \leq \frac{\lambda}{2}$,

($\hat{\gamma}$) the points of E_I are farther than $\frac{\lambda}{2\varrho}$ from one another, where $\varrho \geq 1$ verifies H2.2 and H2.3, and $2\varrho \geq 2$ verifies H2.0 and H2.1.

then E_F verifies H1 and H2. As a consequence,

(â) $\mathrm{Del}_{|S}(E_F)$ *is isotopic to S,*

($\hat{b}$) the Hausdorff distance between $\mathrm{Del}_{|S}(E_F)$ *and S is at most* $\frac{\varepsilon}{\cos^2\theta}$,

($\hat{c}$) E_F is an $\varepsilon\sqrt{1+\frac{1}{\cos^4\theta}}$*-sample of S, hence* $|E_F| = \Omega\left(\frac{\mathrm{Area}(S)}{\varepsilon^2}\right)$,

($\hat{d}$) E_F is $\frac{\varepsilon}{\varrho}$*-sparse, hence* $|E_F| = O\left(\varrho^2\,\frac{\mathrm{Area}(S)}{\varepsilon^2}\right)$.

($\hat{e}$) the facets of $\mathrm{Del}_{|S}(E)$ *have inner angles greater than* $\arcsin\frac{1}{2\varrho}$.

Before proving the theorem, let us explain where the bound on θ comes from. In Section 4.4, we will explain how to construct λ-persistent facets whose vertices are farther than $\frac{\lambda}{2}$ from one another. However, we are currently unable to guarantee that the vertices are farther than $\frac{\lambda}{2}(1+\eta)$ from one another, for any $\eta > 0$. Therefore, we have to assume that $\varrho \geq 1$ in ($\hat{\gamma}$). Consequently, ($\hat{\gamma}$) implies that θ verifies H2.2–H2.3 with ϱ replaced by 1, and H2.0–H2.1 with ϱ replaced by 2. Among these conditions, H2.0 is the most restrictive one and it yields $\theta \leq \arctan\frac{1}{3\sqrt{3}} \approx 10.9°$.

So, when $\theta \leq \arctan\frac{1}{3\sqrt{3}}$, we can set up E_I and ε so that they satisfy ($\hat{\alpha}$), ($\hat{\beta}$), ($\hat{\gamma}$). Then, Theorem 4.5 guarantees the quality of the output mesh. Note that the bound on θ is overly pessimistic since in practice the algorithm produces meshes of good quality for values of θ ranging from 0 to $\frac{\pi}{2}$ – see Chapters 6 and 7. From now on, and until the end of the section, θ is assumed to be at most $\arctan\frac{1}{3\sqrt{3}}$.

Let us now prove Theorem 4.5. The proof is very similar to the one of Theorem 4.4, although with a few additional subtelties.

Proof of the theorem. Since the oracle ω is exhaustive, all the balls of $\mathrm{Del}_{|S}(E_F)$ have radii bounded by ε. Therefore, if the λ-persistent facets remain in $\mathrm{Del}_{|S}(E_F)$, then E_F is a loose ε-sample of S. Otherwise, if p denotes the first center of surface Delaunay ball of radius at most λ that is inserted in E, then, as in the proof of Theorem 4.4, right before the insertion of p, E is a loose λ-sample. By ($\hat{\alpha}$), λ is less than $\frac{\cos^2\theta}{4+3\cos^2\theta}\ \mathrm{lr}_k(S)$. Moreover, all the points inserted before p were centers of Delaunay balls of radius greater than λ. These points are farther than λ from one another and from E_I, by Lemma 4.1 (applied with ε replaced by λ). It follows from ($\hat{\gamma}$) that the points of E are farther than $\frac{\lambda}{2\varrho}$ from one another. Hence, E verifies H1 and H2.0–H2.1, with ε replaced by λ, and ϱ replaced by 2ϱ. Therefore, according to Lemma 2.39, E is a $\lambda\sqrt{1+\frac{1}{\cos^4\theta}}$-sample, where $\lambda\sqrt{1+\frac{1}{\cos^4\theta}}$ is less than $\frac{1}{2}\ \mathrm{lr}_k(S)$, by ($\hat{\alpha}$). As a consequence, $E_F \supseteq E$ is a $\frac{1}{2}\mathrm{lr}_k(S)$-sample of S. Moreover, as mentioned above, ($\hat{\gamma}$) implies that θ is at most $\arctan\frac{1}{3\sqrt{3}} < \frac{\pi}{4}$, which means that $k < \tan\frac{\pi}{8}$. Thus, every point of E_F is a vertex of $\mathrm{Del}_{|S}(E_F)$, by Lemma 2.38. It follows that E_F is a loose ε-sample, since the radii of the balls of $\mathrm{Del}_{|S}(E_F)$ are all bounded by ε. By ($\hat{\alpha}$) and ($\hat{\beta}$), E_F verifies H1. Moreover, by ($\hat{\beta}$) and ($\hat{\gamma}$), the points of E_I are farther than $\frac{\varepsilon}{\varrho}$ from one another, where ϱ verifies H2.0–H2.3. Since by Lemma 4.1 the points of $E_F \setminus E_I$ are farther than $\varepsilon \geq \frac{\varepsilon}{\varrho}$ from one another and from E_I, we deduce that E_F verifies H2. Then, assertion (â) follows from Theorem 2.33, ($\hat{\mathrm{b}}$) from Theorem 2.23 and (ĉ) from Lemma 2.39 and Theorem 3.1. In addition, since E_F verifies H2, it is $\frac{\varepsilon}{\varrho}$-sparse, by the same argument as in the proof of Theorem 4.4. Therefore, by Theorem 3.4, we have $|E_F| = O\left(\varrho^2\,\frac{\mathrm{Area}(S)}{\varepsilon^2}\right)$. Finally, since E_F verifies

H1 and H2, the radius-edge ratios of the facets of $\mathrm{Del}_{|S}(E_F)$ are less than ϱ, which implies that their inner angles are greater than $\arcsin \frac{1}{2\varrho}$. □

Remark that there is an additional subtelty in Condition ($\hat{\beta}$), compared to (β): here, ε is supposed to be at most $\frac{\lambda}{2}$, and not simply at most a fraction of $\mathrm{lr}_k(S)$. The reason is that, to ensure that $\mathrm{Del}_{|S}(E_F)$ has vertices on all the connected components of S, E_F has to verify H2, which is impossible if the points of E_I are too close to one another, compared to ε. In compensation, the output point set E_F is ε-sparse, and the inner angles of the facets of $\mathrm{Del}_{|S}(E_F)$ are bounded from below.

Notice also that, since our theoretical results hold only for $\theta \leq \arctan \frac{1}{3\sqrt{3}}$ and $\varrho \geq 1$, we can replace θ by $\arctan \frac{1}{3\sqrt{3}}$ and ϱ by 1 in ($\hat{\alpha}$), ($\hat{\beta}$) and ($\hat{\gamma}$), which gives the following (more restrictive) conditions:

- ($\hat{\alpha}'$) $\mathrm{Del}_{|S}(E_I)$ has λ-persistent facets on all the connected components of S, with $\lambda < \frac{27}{193}\ \mathrm{lr}_k(S) \approx 0.14\ \mathrm{lr}_k(S)$,
- ($\hat{\beta}'$) $\varepsilon \leq \frac{\lambda}{2}$,
- ($\hat{\gamma}'$) the points of E_I are farther than $\frac{\lambda}{2}$ from one another.

If $\theta \leq \arctan \frac{1}{3\sqrt{3}}$, then ($\hat{\alpha}'$), ($\hat{\beta}'$), ($\hat{\gamma}'$) yield ($\hat{\alpha}$), ($\hat{\beta}$), ($\hat{\gamma}$), hence Theorem 4.5 applies and the quality of the output of the algorithm is guaranteed. If $\theta > \arctan \frac{1}{3\sqrt{3}}$, then ($\hat{\gamma}$) cannot be fulfilled anyway. Therefore, ($\hat{\alpha}$), ($\hat{\beta}$), ($\hat{\gamma}$) and ($\hat{\alpha}'$), ($\hat{\beta}'$), ($\hat{\gamma}'$) offer the same theoretical guarantees, but ($\hat{\alpha}'$), ($\hat{\beta}'$), ($\hat{\gamma}'$) have simpler expressions. In particular, they do not involve θ, which means that it is not necessary to know k in practice to check whether they are satisfied. This is why in the sequel we use them instead of ($\hat{\alpha}$), ($\hat{\beta}$), ($\hat{\gamma}$).

Open question 4.6 *The theoretical limit imposed on θ for our guarantees to hold is far below the practical limit, which seems to lie around $\frac{\pi}{2}$. This gap between theory and practice comes from the proofs of Part A, in which the bounds are not always tight. An appealing example is the proof of Theorem 1.7, which has been written in such a way that it holds both in the smooth and in the nonsmooth settings. Moreover, at several stages of the proofs, we have given priority to simplicity versus optimality, for pedagogical purpose. Therefore, it would be profitable to rewrite the proofs of Part A in a more optimized manner, to improve on the theoretical bound on θ.*

4.3 Termination and complexity

Lemma 4.7 *The algorithm terminates with any input compact surface S, any input point sample E_I and any input parameter $\varepsilon > 0$.*

Proof. At each iteration, the center of some bad surface Delaunay ball $B(c, r)$ is inserted. At this time, the distance from c to E is r, which is greater than ε since $B(c, r)$ is bad. Therefore, the distance between any two points inserted during the main loop is at least ε, and, as a consequence, the open balls of radius $\frac{\varepsilon}{2}$, centered at the points inserted during the main loop, are pairwise disjoint. Since S is compact, there can be only a finite number of such balls. Thus, a finite number of points are inserted during the main loop of the algorithm, which terminates since it inserts one point at each iteration. □

We will now analyze the space and time complexities of the algorithm. We assume that the surface S and the input point set E_I are fixed. Then, the size of E_I and the minimum distance d_I between any two of its points are fixed constants. The analysis assumes further that the input parameter ε is sufficiently small and that the oracle is exhaustive.

4.3.1 Space complexity

Recall that we store the restricted Delaunay triangulation as a subcomplex of the 3-dimensional Delaunay triangulation. Moreover, by storing in the priority queue $\mathcal{Q}$ only one ball of $\mathrm{Del}_{|S}(E)$ per Voronoi edge dual to a bad facet, we enforce the size of $\mathcal{Q}$ to be at most the number of Voronoi edges. Thus, the space complexity of the algorithm is bounded by the size of the 3-dimensional Delaunay triangulation, which is known to be at most quadratic w.r.t. the size of the output point sample. However, in our case the vertices of $\mathrm{Del}(E)$ lie on the surface S, thus we can work out better bounds.

To achieve these bounds, we will use the three lemmas stated below. The first one deals with the smooth case and assumes that S verifies a genericity condition. Let Z be the subset of the ridge of S made of all the points of S that admit an osculating ball whose interior does not intersect S. We assume that Z is a set of curves of finite total length. As mentioned in [12], this property of Z is verified generically. In particular, S cannot contain patches of spheres or cylinders with empty osculating spheres.

Lemma 4.8 *[12, Th. 22]*
If S has a positive reach and verifies the above genericity condition, then there exists a constant ε_0, depending only on S, such that, for any sparse ε-sample E of S, with $\varepsilon \leq \varepsilon_0$, the number of edges of $\mathrm{Del}(E)$ *is* $O_S(|E| \log |E|)$.

The second lemma deals with the polyhedral case.

Lemma 4.9 *[10, Th. 11]*
If S is a polyhedron and E is a sparse ε-sample of S, then the number of edges of $\mathrm{Del}(E)$ *is* $O_S(|E|)$.

The third lemma deals with the general case, smooth or nonsmooth. It makes a stronger assumption on the sparseness of E. Specifically, it assumes that E is a so-called *ambient ε-sample*, as defined by Erickson [64]. This means that there exists a constant $\nu \in]0, 1[$ such that the distance from any point of S to its second nearest neighbor in E lies between $\nu\varepsilon$ and ε.

Lemma 4.10 *[64, Th. 3.2]*
For any surface S and any ambient ε-sample of S, the number of edges of $\mathrm{Del}(E)$ *is* $O_S(|E|\sqrt{|E|})$.

Let $N = O_S(\varepsilon^{-2})$ be the size of the output point set E_F. We will prove the following

Theorem 4.11 *If the surface S has a positive reach and verifies the genericity condition of Lemma 4.8, then the space complexity of the algorithm is $O_S(N \log N) = O_S(\varepsilon^{-2} \log \frac{1}{\varepsilon})$. If S is a polyhedron, then the space complexity is $O_S(N) = O_S(\varepsilon^{-2})$. In the more general case where S is Lipschitz, the space complexity is $O_S(N\sqrt{N}) = O_S(\varepsilon^{-3})$.*

We give a proof only in the case of a generic smooth surface S, the polyhedral and general Lipschitz cases being similar. For any iteration $i > 0$ of the algorithm, we call $E(i)$ the point set E at the end of iteration i. Similarly, we set $E(0) = E_I$. $E(i) \setminus E(i-1)$ contains precisely the point $p(i)$ inserted in E at iteration i, and $E(i-1) \setminus E_I$ is the set of all points inserted before iteration i. We call $r(i)$ the radius of the largest ball of $\mathrm{Del}_{|S}(E(i))$. Since the algorithm always inserts the center of the ball of $\mathrm{Del}_{|S}(E)$ of largest radius, $p(i)$ is at a distance $r(i-1)$ from $E(i-1)$.

We define ε_0 as the minimum of $\mu_0 \ \mathrm{rch}(S)$, of $\frac{d_I}{2}$, and of the (unknown) constant of Lemma 4.8. The assumption that ε_0 is at most $\frac{d_I}{2}$ is not necessary here, but it will be used in Section 4.3.2 to bound the time complexity of the algorithm.

Let i_0 be the first iteration of the algorithm at the end of which all the balls of $\mathrm{Del}_{|S}(E)$ have radii less than $\frac{\varepsilon_0}{4}$. In other words, i_0 is the first iteration such that $r(i_0) < \frac{\varepsilon_0}{4}$. Since $\varepsilon_0 \leq \mu_0 \ \mathrm{rch}(S) \leq \mu_0 \ \mathrm{rch}(S)$, $E(i_0)$ is an $\frac{\varepsilon_0}{4}\sqrt{1+\frac{\mu_0^2}{16}}$-sample of S, by Theorem 4.4 (applied with ε replaced by $\frac{\varepsilon_0}{4}$). Since it is not our aim here to work out tight constants, for simplicity we replace $\sqrt{1+\frac{\mu_0^2}{16}}$ by 2 in the sequel. Hence, $E(i_0)$ is an $\frac{\varepsilon_0}{2}$-sample of S.

Lemma 4.12 *For any two iterations i and j such that $j \geq i \geq i_0$, we have $r(j) \leq 2r(i)$.*

Proof. Let $B(j)$ be a ball of $\mathrm{Del}_{|S}(E(j))$ of largest radius. Its center $c(j)$ lies on S. Since $i \geq i_0$, we have $E(i_0) \subseteq E(i)$. Thus, $E(i)$ is an $\frac{\varepsilon_0}{2}$-sample of S, which implies that the balls of $\mathrm{Del}_{|S}(E(i))$ have radii at most $\frac{\varepsilon_0}{2}$. Therefore, $r(i) \leq \frac{\varepsilon_0}{2}$. It follows that $E(i)$ is a $2r(i)$-sample, by Theorem 4.4 (applied with ε replaced by $r(i)$). Hence, $\mathrm{d}(c(j), E(i)) \leq 2r(i)$. Moreover, since $i \leq j$, $E(i)$ is included in $E(j)$. It follows that $r(j) = \mathrm{d}(c(j), E(j)) \leq \mathrm{d}(c(j), E(i)) \leq 2r(i)$, which concludes the proof of the lemma. □

Lemma 4.13 *For any iteration $i > i_0$, $E(i)$ is a $2r(i)$-sample of S, with $2r(i) \leq \varepsilon_0$, and the points of $E(i) \setminus E_I$ are farther than $\frac{r(i-1)}{2}$ from one another and from E_I.*

Proof. Let i be any iteration of the algorithm such that $i > i_0$. As mentioned in the proof of Lemma 4.12, $E(i)$ is a $2r(i)$-sample, with $2r(i) \leq \varepsilon_0$. In addition, by definition of i_0, every point of $E(i_0) \setminus E_I$, when inserted in E, is the center of a Delaunay ball of radius greater than $\frac{\varepsilon_0}{4} \geq r(i_0)$, which is at least $\frac{1}{2}\ r(i-1)$, by Lemma 4.12. Moreover, at any iteration k such that $i_0 < k \leq i$, the point inserted in E is the center of a Delaunay ball of radius $r(k-1)$, which is at least $\frac{1}{2}\ r(i-1)$, by Lemma 4.12. Therefore, the points of $E(i) \setminus E_I$ are at least $\frac{1}{2}\ r(i-1)$ away from one another and from E_I. □

It follows from Lemmas 4.12 and 4.13 that the points of $E(i) \setminus E_I$ are farther than $\frac{r(i)}{4}$ from one another. Since E_I is fixed, by the same argument as in the proof of Theorem 4.4 we get that the number of points of $E(i)$ that lie in $B(p, 2r(i))$ is bounded by a constant, for any $p \in S$. Hence,

Corollary 4.14 *For any $i > i_0$, $E(i)$ is a sparse $2r(i)$-sample of S.*

As a consequence, the size of $\mathrm{Del}(E(i))$ is $O_S(|E(i)| \log |E(i)|) = O_S(N \log N)$, by Lemma 4.8. Therefore, the space complexity of the algorithm is $O_S(N \log N)$. Moreover, by Theorem 4.4, we have $O_S(N \log N) = O_S(\varepsilon^{-2} \log \frac{1}{\varepsilon})$, which concludes the proof of Theorem 4.11 in the smooth case.

In the polyhedral case, the proof is identical. In the Lipschitz case however, the proof differs slightly since Lemma 4.10 requires that $E(i)$ be an ambient $2r(i)$-sample, which is stronger than being a sparse $2r(i)$-sample.

4.3.2 Time complexity

Let T be the overall number of Delaunay tetrahedra created during the course of the algorithm.

Lemma 4.15 *The time complexity of the algorithm is $O(T \log T)$.*

Proof. - At each iteration, the algorithm inserts in E a point p that lies on a Voronoi edge. The algorithm knows on which edge e the point p lies, since it queried the oracle on e to find p. Therefore, no point location is needed to insert p in $\mathrm{Del}(E)$. It follows that updating $\mathrm{Del}(E)$ can be done in linear time with respect to the number of Delaunay tetrahedra created and deleted during the insertion of p. Since a Delaunay tetrahedron is created and deleted at most once, the overall cost of maintaining $\mathrm{Del}(E)$ is $O(T)$.

- The cost of updating $\mathrm{Del}_{|S}(E)$ is also $O(T)$ since the oracle is queried at most once on each Voronoi edge.

- Since we store only one ball of $\mathrm{Del}_{|S}(E)$ per Voronoi edge in the priority queue $\mathcal{Q}$, the total number of balls of $\mathrm{Del}_{|S}(E)$ inserted in $\mathcal{Q}$ (and then retrieved from it) is at most the total number of Voronoi edges created during the process. Hence, the cost of maintaining $\mathcal{Q}$ is $O(T \log T)$. □

To bound the time complexity of the algorithm, it is then sufficient to control the total number T of tetrahedra created during the process. By Theorem 4.11, the size of $\mathrm{Del}(E)$ is controlled throughout the process. However, this fact alone does not prevent T from being quadratic w.r.t. the size of the output point set. Consider for instance the case of a set of N points sampled from a parabola $\mathcal{P}$ in $\mathbb{R}^2$. The size of the Delaunay triangulation is $O(N)$, but if the vertices are inserted in their order along $\mathcal{P}$ (see Figure 4.2, left), then at each iteration a linear number of Delaunay edges is created, which makes T quadratic w.r.t. N. However, this kind of situation cannot occur with our algorithm because it inserts the points in an order defined by the *largest empty ball* criterion (see Figure 4.2, right).

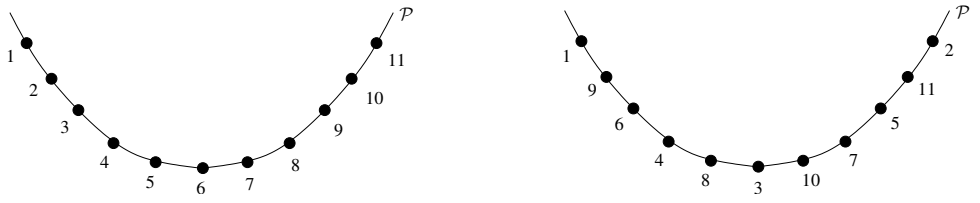

Figure 4.2: Two orders of insertion on a parabola.

We assume here that the surface S has a positive reach. We will say a few words about the polyhedral case and the more general Lipschitz case at the end of the section. Let Z, ε_0 and i_0 be defined as in Section 4.3.1. We assume that S verifies the genericity condition of Lemma 4.8. We will show that $T = O_S(N \log N)$, where N is the size of the output point set. To this end, we introduce the following result, stated as Lemma 17 in [12]:

Lemma 4.16 *There exist three constants c_0, k_1 and k_2, depending only on S, such that, for any sparse ε-sample E of S, with $\varepsilon \leq \varepsilon_0$, the number of Delaunay edges incident to a vertex p of $\mathrm{Del}(E)$ is at most $k_1\ \varepsilon^{-1/2}$ if $\mathrm{d}(p,Z) \leq c_0\ \sqrt{\varepsilon}$ and at most $k_2\ /\ \mathrm{d}(p,Z)$ otherwise.*

Assuming as in Section 4.3.1 that the initial point set E_I is fixed, we deduce from Lemma 4.12, Corollary 4.14 and Lemma 4.16, that the algorithm creates $O_S\left(r(i)^{-1/2}\right) = O_S\left(\varepsilon^{-1/2}\right)$ new Delaunay edges at each iteration $i > i_0$. As a consequence, the overall number of Delaunay edges created after iteration i_0 (which depends only on S) is $O_S\left(N\ \varepsilon^{-1/2}\right) = O_S\left(N^{5/4}\right)$. However, by summing more carefully the contributions of the points inserted after iteration i_0, we can work out a $O_S\left(N \log N\right)$ bound.

Let i be an iteration of the algorithm, such that $i \geq i_0$. Let $j > i$ be the first iteration such that $r(j) < \frac{r(i)}{8}$. Our goal is to bound the number of Delaunay edges created between iterations i and j. By Corollary 4.14, $E(j)$ is a $2r(j)$-sample of S, with $2r(j) < \frac{r(i)}{4}$. Moreover, by Lemma 4.12, we have $2r(j) \leq 4r(i_0) < \varepsilon_0$. We call $E(i,j)$ the set of the points inserted by the algorithm between iterations i (excluded) and j (included). We have $E(i,j) = E(j) \setminus E(i)$.

Lemma 4.17 *For any k such that $i < k \leq j$, $E(k)$ is a sparse $4r(i)$-sample of S.*

Proof. By Corollary 4.14, $E(k)$ is a $2r(k)$-sample of S. Since $2r(k) \leq 4r(i)$ (Lemma 4.12), $E(k)$ is a $4r(i)$-sample. To prove that $E(k)$ is sparse, we count the points of $E(k)$ that lie in $B(p, 4r(i))$, for any $p \in S$.
- Since E_I is fixed, $|E_I \cap B(p, 4r(i))|$ is a constant independent from S and ε.
- By Lemma 4.13, the points of $E(k) \setminus E_I$ are farther than $\frac{r(k-1)}{2}$ from one another. Now, $\frac{r(k-1)}{2}$ is at least $\frac{r(i)}{16}$, since $i < k \leq j$. It follows that the points of $E(k) \setminus E_I$ are centers of pairwise-disjoint balls of radius $\frac{1}{32}\ r(i)$. For every such ball B whose center lies in $B(p, 4r(i))$, B is included in $B(p, (4 + \frac{1}{32})\ r(i))$. It follows that the number of points of $E(k) \setminus E_I$ that lie in $B(p, 4r(i))$ is bounded by a constant, which shows that $E(k)$ is sparse and hereby concludes the proof of Lemma 4.17. □

Lemma 4.18 *$E(i,j)$ is a sparse $4r(i)$-sample of S.*

Proof. We first show that $E(i,j)$ is a $\frac{3}{4}r(i)$-sample, which implies *a fortiori* that it is a $4r(i)$-sample. Let $p \in S$ and let u be the nearest neighbor of p among the points of $E(j)$. We know that $E(j)$ is a $2r(j)$-sample, thus $\mathrm{d}(p,u) \leq 2r(j) < \frac{r(i)}{4}$. If u belongs to $E(i,j)$, then $\mathrm{d}(p, E(i,j)) = \mathrm{d}(p,u) < \frac{r(i)}{4}$. If u does not belong to $E(i,j)$, then we must show that there is some point of $E(i,j)$ nearby.

Since $2r(j) < \varepsilon_0 \leq \mu_0\ \mathrm{rch}(S)$, by Lemma 1.22 u is a vertex of $\mathrm{Del}_{|S}(E(j))$. Let v be one of its neighbors in $\mathrm{Del}_{|S}(E(j))$. For any point $q \in S$ lying in the Voronoi face $\mathrm{V}(u) \cap \mathrm{V}(v)$, we have $\mathrm{d}(u,q) \geq \frac{\mathrm{d}(u,v)}{2}$. Since u is a nearest neighbor of q in $E(j)$, $\mathrm{d}(u,q) \leq 2r(j) < \min\{\varepsilon_0, \frac{r(i)}{4}\}$. Thus, $\mathrm{d}(u,v) < 2\varepsilon_0 \leq d_I$, and $\mathrm{d}(u,v) < \frac{r(i)}{2}$, which is less than the minimal distance between any two points of $E(i+1) \setminus E_I$, by Lemma 4.13. Therefore, v cannot belong to $E(i+1)$, which contains both E_I and $E(i)$. This implies that

$$\mathrm{d}(p, E(i,j)) \leq \mathrm{d}(p,v) \leq \mathrm{d}(p,u) + \mathrm{d}(u,v) \leq 6r(j) \leq \frac{3}{4}r(i)$$

It follows that $E(i,j)$ is a $\frac{3}{4}r(i)$-sample, and hence a $4r(i)$-sample.

We now show that $E(i,j)$ is $4r(i)$-sparse. By definition of j, every point of $E(i,j)$, right before its insertion in E, is the center of a Delaunay ball of radius at least $\frac{r(i)}{8}$. It follows that the points of $E(i,j)$ are farther than $\frac{r(i)}{8}$ from one another. Hence, by the same argument as in the proof of Lemma 4.17, $E(i,j)$ is sparse. □

We can now combine Lemma 4.16 with Lemmas 4.17 and 4.18, to bound the number of Delaunay edges created between iterations i and j:

Lemma 4.19 *The overall number of Delaunay edges created during the insertion of the points of $E(i,j)$ is $O_S\left(|E(i,j)|\log|E(i,j)|\right)$.*

Proof. Let $\varepsilon_i = 4r(i)$. The reasoning is similar in spirit to that of Lemma 18 of [12], although with an additional subtelty. We decompose S into strips parallel to Z, of width $c_0\sqrt{\varepsilon_i}$, where c_0 is defined as in Lemma 4.16. Recall that c_0 depends on S but not on ε_i. Let Z_k denote the k^{th} strip ($k \geq 0$). The points of Z_k lie at a distance of Z ranging from $k\ c_0\sqrt{\varepsilon_i}$ to $(k+1)\ c_0\sqrt{\varepsilon_i}$.

As stated in Lemma 18 of [12], since $E(i,j)$ is a sparse ε_i-sample of S (Lemma 4.18), there exists some constant $c(S)$ depending only on S, such that the number of points of $E(i,j)$ that lie in a given strip Z_k is at most $c(S)\ \varepsilon_i^{-3/2}$. Moreover, for any i' such that $i < i' \leq j$, $E(i')$ is a sparse ε_i-sample of S (Lemma 4.17), thus Lemma 4.16 applies to the point inserted at iteration i'. Summing the contributions of all the points of $E(i,j)$, we find that the number n of Delaunay edges created by the insertion of the points of $E(i,j)$ is at most:

$$c(S)\ \varepsilon_i^{-3/2}\frac{k_1}{\sqrt{\varepsilon_i}} + \sum_{k\geq 1} c(S)\ \varepsilon_i^{-3/2}\frac{k_2}{k\ c_0\sqrt{\varepsilon_i}} = k_1\ c(S)\ \varepsilon_i^{-2} + \frac{k_2\ c(S)\ \varepsilon_i^{-2}}{c_0}\sum_{k\geq 1}\frac{1}{k}$$

The number of strips Z_k is $c'(S)\ /\ c_0\sqrt{\varepsilon_i}$, where $c'(S)$ depends only on S. Moreover, by Theorem 3.1, the size of $E(i,j)$ is at least $c''(S)\ \varepsilon_i^{-2}$, for some constant $c''(S)$ depending only on S. It follows that n is bounded by:

$$k_1\ \frac{c(S)}{c''(S)}\ |E(i,j)| + \frac{k_2}{c_0}\ \frac{c(S)}{c''(S)}\ |E(i,j)| \sum_{1\leq k\leq \frac{c'(S)}{c_0\ c''(S)^{1/4}}|E(i,j)|^{1/4}} \frac{1}{k}$$

i.e. $O_S\left(|E(i,j)|\right) + O_S\left(|E(i,j)|\right)\ O_S\left(\log|E(i,j)|\right)$. □

Finally, by subdividing the output point set into subsets of type $E(i,j)$, we can bound the overall number of Delaunay edges created after iteration i_0:

Theorem 4.20 *If S has a positive reach and verifies the genericity condition of Lemma 4.8, then the overall number T of Delaunay tetrahedra created during the course of the algorithm is $O_S\left(N\log N\right)$. As a consequence, by Lemma 4.15, the time complexity of the algorithm is $O_S(N\log^2 N) = O_S\left(\varepsilon^{-2}\log^2\frac{1}{\varepsilon}\right)$.*

Proof. Let i_0 be defined as above, and for any $k \geq 1$, let $i_k > i_{k-1}$ be the first iteration such that $r(i_k) < \frac{r(i_{k-1})}{8}$. We assume without loss of generality that the last iteration of the algorithm is i_K, for

some K. We have $E_F = E(i_0) \cup \bigcup_{0 \leq k < K} E(i_k, i_{k+1})$. By Lemma 4.19, every $E(i_k, i_{k+1})$ generates $c^*(S)\ |E(i_k, i_{k+1})| \log |E(i_k, i_{k+1})|$ Delaunay edges, where $c^*(S)$ is a constant that depends only on S. It follows that the overall number of Delaunay edges created is at most:

$$\begin{array}{rl} & |E(i_0)|^2 + \sum_{0 \leq k < K} c^*(S)\ |E(i_k, i_{k+1})| \log |E(i_k, i_{k+1})| \\ \leq & |E(i_0)|^2 + \sum_{0 \leq k < K} c^*(S)\ |E(i_k, i_{k+1})| \log |E_F| \\ \leq & |E(i_0)|^2 + c^*(S)\ |E_F| \log |E_F| \end{array}$$

Since E_I is fixed and since the points of $E(i_0) \setminus E_I$ are at least $\frac{\varepsilon_0}{4}$ away from one another, $E(i_0)$ is ε_0-sparse, by the same argument as in the proof of Theorem 4.4. Thus, by Theorem 3.4, we have $|E(i_0)| = O\left(\mathrm{Area}(S) \,/\, \varepsilon_0^2\right)$, which depends only on S. This concludes the proof of Theorem 4.20, since the number of Delaunay tetrahedra is linear w.r.t. the number of Delaunay edges. □

To our knowledge, there is no result analogous to Lemma 4.16 in the polyhedral or Lipschitz case. Indeed, in [10] and [64], the size of the Delaunay triangulation is bounded without bounding explicitely the number of neighbors of each vertex. Therefore, in the polyhedral or Lipschitz case, our bound on the time complexity of the algorithm is still $O(N^2 \log N)$ at the moment.

Open question 4.21 *Given a polyhedral or Lischitz surface S and a point sample E of S, is it possible to work out non-trivial bounds on the number of neighbors of each vertex in* $\mathrm{Del}(E)$*, under some mild assumption on E?*

4.4 Pre-conditioning the input point set

As emphasized in Section 4.2, our theoretical guarantees require that $\mathrm{Del}_{|S}(E_I)$ has λ-persistent facets on every connected component of S, for some $\lambda \leq \mu_0\ \mathrm{rch}(S) \approx 0.16\ \mathrm{rch}(S)$ in the smooth case and $\lambda < \frac{27}{193}\ \mathrm{lr}_k(S) \approx 0.14\ \mathrm{lr}_k(S)$ in the k-Lipschitz case (Conditions (α) and ($\hat{\alpha}'$)). Another requirement in the Lipschitz case is that the points of E_I be farther than $\frac{\lambda}{2}$ from one another (Condition ($\hat{\gamma}'$)). Once these requirements are fulfilled, plus the additional condition that $2 \arctan k \leq \arctan \frac{1}{3\sqrt{3}}$ in the k-Lipschitz case, it suffices to choose $\varepsilon \leq \frac{\lambda}{2}$ to have guarantees on the output of the algorithm.

Assume that we are given a positive constant ε_S, less than $\mathrm{rch}(S)$ if S has positive reach, and less than $\mathrm{lr}_k(S)$ if S is k-Lipschitz. In Part C, we will explain how to compute such a constant ε_S in several applications. Assume further that we are given a set E'_I containing at least one point per connected component of S. As explained in Part C, to construct E'_I in practice, we compute the critical points of the height function along some direction $\mathbf{n}$. Hence, we can assume without loss of generality that $\mathbf{n}$ is given here. A point $p \in S$ is called ***n*-extremal** if S lies locally on one side of the plane P perpendicular to $\mathbf{n}$ that passes through p. Note that not all the points of E'_I are $\mathbf{n}$-extremal. However, E'_I contains at least two $\mathbf{n}$-extremal points per connected component of S, since S is compact.

We will now describe briefly how we can construct a superset of E_I that satisfies (α), ($\hat{\alpha}'$) and ($\hat{\gamma}'$), by using ε_S, E'_I, $\mathbf{n}$, and the exhaustive oracle. We are mainly interested in the construction itself, not on its proof of correctness, which is technically easy but quite tedious. Therefore, we will only give the intuition, without going too much into detail. Moreover, we assume for simplicity that the surface S is

k-Lipschitz, for some k such that $\theta = 2\arctan k$ is at most $\arctan\frac{1}{3\sqrt{3}}$. Since by Theorem 2.6 every smooth surface is k-Lipschitz, our construction will work on smooth surfaces as well.

Our construction proceeds as follows. We first set $\lambda = \min\{\frac{d'_I}{2},\ \frac{27}{193}\,\varepsilon_S\}$, where d'_I is the minimum distance between any two points of $E_I \cup E'_I$. Then, for each point $p \in E'_I$, we construct a λ-persistent facet incident to p. Specifically, inside the plane P perpendicular to $\mathbf{n}$ that passes through p, we consider the circle of center p and radius $\frac{\lambda}{2}$, and we choose two points p_1 and p_2 on this circle, such that the triangle (p, p_1, p_2) is equilateral. Then, we use the exhaustive oracle to detect the intersection points of S with the lines l_1, l_2 orthogonal to $T(p)$ that pass through p_1 and p_2 respectively. Let p'_1 be the point of $S \cap l_1$ that is closest to p, and p'_2 its analogous on l_2. If p'_1 and p'_2 exist and are not farther than λ from p, then we insert them in E_I, together with p.

It can be proved that this procedure builds at least one λ-persistent facet per connected component of S. Here, we give the flavor of the proof. Since λ is smaller than $\mathrm{lr}_k(S)$, by the cocone lemma 2.12 $S \cap B(p, \lambda)$ lies outside the double cone $K^1(p)$ of apex p, of axis aligned with $\mathbf{n}(p')$ and of half-angle $\frac{\pi}{2} - \theta$, where p' is any point of $\tilde{S} \cap B(p, \lambda)$. Let K^1_1 and K^1_2 be the two cones of $K^1(p)$. Since S has no boundary and passes through p, $K^1_1 \cap B(p, \lambda)$ and $K^1_2 \cap B(p, \lambda)$ belong to distinct connected components (say Ω_1 and Ω_2) of $\mathbb{R}^3 \setminus S$.

Assume now that p is $\mathbf{n}$-extremal, which implies that S lies locally on one side of P. Then, so does Ω_1 or Ω_2 (say Ω_1). Notice that Ω_1 may still intersect the two half-spaces bounded by P, even inside $B(p, \lambda)$. However, sufficiently close to p, Ω_1 is included in one half-space only. This is sufficient to ensure that P does not intersect $K^1_1 \cap B(p, \lambda)$, which is included in Ω_1. It follows, by a symmetry argument, that P does not intersect the double cone $K^1(p)$, which means that the angle between P and $\mathbf{n}(p')$ is at least $\frac{\pi}{2} - \theta$. Equivalently, we have $(\mathbf{n}(p'), \mathbf{n}) \leq \theta$, which implies that $S \cap B(p, \lambda)$ lies outside the double cone $K^2(p)$ of apex p, of axis aligned with $\mathbf{n}$ and of half-angle $\frac{\pi}{2} - 2\theta$. Using this cone, with the same machinery as in Chapter 2, we can show that l_1 and l_2 intersect S in the vicinity of p, which implies that p'_1 and p'_2 exist and are sufficiently close to p to form a λ-persistent facet with p. The proof is long and does not present any particular technical difficulty; we skip it here, for simplicity.

It follows that our construction creates at least one λ-persistent facet per $\mathbf{n}$-extremal point of E'_I. Note that E'_I may not necessarily contain only $\mathbf{n}$-extremal points. However, E'_I contains at least two such points per connected component of S. Therefore, the output of our procedure contains at least one λ-persistent facet per connected component of S. Moreover, the fact that the points are farther than $\frac{\lambda}{2}$ from one another comes from two things: first, we chose λ such that the points of $E_I \cup E'_I$ are farther than 2λ from one another; second, for any $p \in E'_I$, we have $\mathrm{d}(p, p'_1) \geq \mathrm{d}(p, p_1) = \frac{\lambda}{2}$ and $\mathrm{d}(p, p'_2) \geq \mathrm{d}(p, p_2) = \frac{\lambda}{2}$.

Chapter 5

Improvements

In this chapter, we introduce several modifications that can be brought to the algorithm to enhance the quality of its output or to avoid some geometric computations.

In Section 5.1, we introduce a non-exhaustive oracle, called *bipolar oracle*, that performs simpler operations than the exhaustive oracle, while still guaranteeing the topological and geometric quality of the output mesh. The guarantees hold for smooth surfaces as well as Lipschitz surfaces.

In Section 5.2 we present a way to avoid the pre-conditioning of the input point set. This modification comes with no guarantee regarding the time complexity of the algorithm. However, termination is still ensured, together with the size of the output mesh and its topological and geometric properties.

In Section 5.3, we describe how the algorithm can be modified to increase the minimum inner angle of the facets of the output mesh.

Finally, in Section 5.4, we explain how to generate non-uniform samples.

5.1 Non-exhaustive oracle

The results of Part A, as used in Chapter 4, rely on the assumption that the algorithm uses an exhaustive oracle able to compute all the intersection points of any line segment with the surface. In practice, *e.g.* for implicit surfaces, computing all the intersection points can be time-consuming. In this section, we introduce a new oracle that offers the same guarantees as the exhaustive oracle, while performing simpler operations. Specifically, the new oracle is able:

- to compute the parity of the number of transversal intersections between s and S;
- to find one point of $s \cap S$ when s intersects S transversally an odd number of times.

We will use this oracle only in Chew's algorithm, not in the pre-conditioning phase of Section 4.4. Indeed, for technical reasons explained below, this phase requires the use of an exhaustive oracle. However, since it is only a precomputation, we can afford to keep the exhaustive oracle.

Let Ω be the compact object bounded by S. A Voronoi edge that intersects S transversally an odd number of times is called *bipolar*, because one of its endpoints lies in Ω while the other endpoint lies in $\mathbb{R}^3 \setminus \bar{\Omega}$. On the contrary, a Voronoi edge intersecting S transversally an even number of times has both vertices in Ω or in $\mathbb{R}^3 \setminus \bar{\Omega}$.

The Delaunay facet dual to a bipolar Voronoi edge is called a *bipolar restricted Delaunay facet*. The subcomplex of $\mathrm{Del}_{|S}(E)$ made of the bipolar Delaunay facets is called the *bipolar restricted Delaunay triangulation* and denoted by $\mathrm{Del}^b_{|S}(E)$. The new oracle, which detects only the bipolar Voronoi edges and, for each bipolar edge e, computes only one point of $e \cap S$, is called *bipolar oracle*, as opposed to the exhaustive oracle which can detect all the intersection points between the Voronoi edges and the surface. A Delaunay ball centered at some point computed by the bipolar oracle is called a *bipolar surface Delaunay ball*.

Persistent facets To pre-condition E_I, we use the exhaustive oracle and proceed in two steps:

1. We perform the construction described in Section 4.4, which provides us with one λ-persistent facet per connected component of S. However, some of the λ-persistent facets may not be bipolar at this stage.
2. For any λ-persistent facet f, we insert in E_I the centers of all the surface Delaunay balls of f whose radii are greater than $\frac{\lambda}{2}$. This requires to know all the points of intersection of S with $\mathrm{V}(f)$. Therefore, the exhaustive oracle is needed here. We do this for all the λ-persistent facets and take the result as the new input point sample E_I.

Termination Since $\mathrm{Del}^b_{|S}(E)$ is a subcomplex of $\mathrm{Del}_{|S}(E)$, the proof of Lemma 4.7 holds when the algorithm is run with the bipolar oracle instead of the exhaustive oracle. Hence, the algorithm terminates.

Output guarantees At any stage of the course of the algorithm, $\mathrm{Del}^b_{|S}(E)$ is a subcomplex of $\mathrm{Del}_{|S}(E)$. Hence, a natural way to have guarantees on the output of the algorithm, when used with the bipolar oracle, would be to show that the output mesh coincides with the restricted Delaunay triangulation of the output point set, that is, $\mathrm{Del}^b_{|S}(E_F) = \mathrm{Del}_{|S}(E_F)$. The theoretical results of Part A could then be applied, as in Section 4.2. Unfortunately, we could not prove this. Instead, we proved the following

Lemma 5.1 *Theorems 4.4 and 4.5 hold with "*$\mathrm{Del}_{|S}$*" replaced by "*$\mathrm{Del}^b_{|S}$*".*

For simplicity, we focus on the smooth case, the Lipschitz case being quite similar. From now on, we assume that the surface S has a positive reach. Let $\mu_0 = 0.16$ be defined as in Section 4.2.2.

Proof of the lemma. We will show that, under Hyp. (α) and (β) (with $\mathrm{Del}_{|S}(E_I)$ replaced by $\mathrm{Del}^b_{|S}(E_I)$), $\mathrm{Del}^b_{|S}(E_F)$ verifies assertions I1–I3 of Section 1.2. The lemma will then follow.

Proof of I1. Assertions M1 and M3 of Section 1.2 still hold in our context, since they do not involve restricted Delaunay facets. Moreover, since bipolar surface Delaunay balls are surface Delaunay balls and since they have small radii upon termination of the algorithm, assertion M2 holds with $\mathrm{Del}_{|S}(E_F)$ replaced by $\mathrm{Del}^b_{|S}(E_F)$. As a consequence, the projection lemma 1.8 holds for the facets of $\mathrm{Del}^b_{|S}(E_F)$. It follows that every edge of $\mathrm{Del}^b_{|S}(E_F)$ is incident to at most two facets of $\mathrm{Del}^b_{|S}(E_F)$.

Note however that Remark 1.9 does not hold since the bipolar oracle does not allow to control the radii of all the surface Delaunay balls. Nevertheless, it remains true that the dual Voronoi edge of any facet of $\mathrm{Del}^b_{|S}(E_F)$ intersects at least two distinct connected components of $\mathbb{R}^3 \setminus S$, since otherwise a small perturbation of this Voronoi edge would intersect S twice near E_F, which would contradict

the projection Lemma 1.8. Hence, the proof of Lemma 1.10 holds for $\mathrm{Del}^b_{|S}(E_F)$, and every edge of $\mathrm{Del}^b_{|S}(E_F)$ is incident to exactly two facets of $\mathrm{Del}^b_{|S}(E_F)$.

In addition, the proof of Lemma 1.11 relies only on M1, M3, M5, M6, and on the properties of the umbrellas made of small restricted Delaunay facets, without any additional argument, thus it holds when we replace $\mathrm{Del}_{|S}(E_F)$ by $\mathrm{Del}^b_{|S}(E_F)$. It follows that every vertex of $\mathrm{Del}^b_{|S}(E_F)$ has one umbrella.

As a conclusion, $\mathrm{Del}^b_{|S}(E_F)$ is a compact surface without boundary. The fact that it is consistently oriented comes from Lemma 1.12, whose proof uses only the projection lemma 1.8. □

Proof of I2. Notice first that, independently from any oracle, Remark 1.9 states that a Delaunay facet cannot be circumscribed by two surface Delaunay balls of radii less than $\mu_0\ \mathrm{rch}(S)$. Since by (α) we have $\lambda \leq \mu_0\ \mathrm{rch}(S)$, any λ-persistent facet f is circumscribed by exactly one ball $B_f = B(c_f, r_f)$ of radius $r_f \leq \frac{\lambda}{2}$, while the other balls circumscribing f have radii greater than λ. Hence, after the pre-processing phase, f is circumscribed by at most one surface Delaunay ball. Moreover, every point inserted in E_I during the second step of the pre-processing phase is farther than λ from E_I. Thus, by Lemma 4.3, B_f remains a surface Delaunay ball throughout the pre-processing phase.

During the course of Chew's algorithm, as long as every point inserted in E is farther than λ from E, B_f remains a surface Delaunay ball, by Lemma 4.3. Hence, f remains bipolar since B_f is the only surface Delaunay ball of f. Therefore, by the same argument as in the proof of Theorem 4.4, I2 is verified by $\mathrm{Del}^b_{|S}(E_F)$. □

Proof of I3. Upon termination of the algorithm, the radius of every bipolar surface Delaunay ball is at most ε, which is bounded by $\mu_0\ \mathrm{rch}(S)$, by (β). Hence, I3 is verified by $\mathrm{Del}^b_{|S}(E_F)$. □

So, under (α) and (β), $\mathrm{Del}^b_{|S}(E_F)$ verifies I1–I3. Thus, we can apply the various results of Chapter 1 to $\mathrm{Del}^b_{|S}(E)$, which proves that Theorem 4.4 holds with "$\mathrm{Del}_{|S}$" replaced by "$\mathrm{Del}^b_{|S}$". This concludes the proof of Lemma 5.1. □

In the smooth case, if $\varepsilon \leq 0.98\ \mu_0\ \mathrm{rch}(S)$, then, by Theorem 4.4, the output point set E_F is a $\varepsilon\sqrt{1+\frac{\varepsilon^2}{\mathrm{rch}(S)^2}}$-sample of S, with $\varepsilon\sqrt{1+\frac{\varepsilon^2}{\mathrm{rch}(S)^2}} \leq \mu_0\ \mathrm{rch}(S)$. By Lemma 1.22, E_F is a loose $\varepsilon\sqrt{1+\frac{\varepsilon^2}{\mathrm{rch}(S)^2}}$-sample. Thus, $\mathrm{Del}_{|S}(E_F)$ is isotopic to S, by Theorem 1.14. This implies that $\mathrm{Del}^b_{|S}(E_F) = \mathrm{Del}_{|S}(E_F)$, since $\mathrm{Del}^b_{|S}(E_F)$ and $\mathrm{Del}_{|S}(E_F)$ are isotopic and since $\mathrm{Del}^b_{|S}(E_F)$ is a subcomplex of $\mathrm{Del}_{|S}(E_F)$. The same argument holds in the Lipschitz case, with Theorem 4.4 replaced by Theorem 4.5, Lemma 1.22 by Lemma 2.38, and Theorem 1.14 by Theorem 2.33.

In conclusion, replacing the exhaustive oracle by the bipolar oracle implies that the algorithm works with $\mathrm{Del}^b_{|S}(E)$ instead of $\mathrm{Del}_{|S}(E)$. Note that the exhaustive oracle is still used for the preprocessing phase. For small enough ε, $\mathrm{Del}^b_{|S}(E_F)$ verifies all the properties stated in Theorems 4.4 and 4.5 for $\mathrm{Del}_{|S}(E_F)$. Moreover, for slightly smaller ε, $\mathrm{Del}_{|S}(E_F)$ and $\mathrm{Del}^b_{|S}(E_F)$ are equal. As explained in Chapter 6 and illustrated in Table 6.1, the bipolar oracle is much more efficient than the exhaustive oracle when run on implicit surfaces.

5.2 Getting rid of persistent facets

In this section, S is assumed to have no nested connected components, which means that the compact object bounded by S has no void. We are given a positive constant ε_S less than $\mathrm{rch}(S)$ if S has a positive reach, and less than $\mathrm{lr}_k(S)$ if S is k-Lipschitz. Moreover, the oracle is assumed to be exhaustive. With these restrictions, we can avoid the construction of persistent facets during the pre-processing phase.

Recall that persistent facets are used only to ensure that $\mathrm{Del}_{|S}(E_F)$ verifies Hyp. I2 of Section 1.2, that is, that $\mathrm{Del}_{|S}(E_F)$ has vertices on every connected component of S. The idea here is to use the set E'_I of Section 4.4 as a set of watch points that will allow us to check whether I2 is verified.

Specifically, the pre-conditioning phase consists now simply in inserting the points of E'_I in E_I without constructing any persistent facet. Then, throughout the course of Chew's algorithm, we check whether the points of E'_I are vertices of $\mathrm{Del}_{|S}(E)$. If so, then we know that I2 is satisfied since E'_I has vertices on each connected component of S. Otherwise, we cannot decide whether I2 is satisfied or not because we do not know on which connected components of S the vertices of $\mathrm{Del}_{|S}(E)$ lie. Hence, we insert additional points of S in E and go on running the main loop of the algorithm until no more surface Delaunay balls are bad and all points of E'_I have become vertices of $\mathrm{Del}_{|S}(E)$.

The additional points of S are computed by casting rays in some specific directions from the points of E'_I that are not yet vertices of $\mathrm{Del}_{|S}(E)$. New directions are determined by sampling iteratively the sphere of directions, using for instance Chew's algorithm. More precisely, let $(\mathbf{v}_i)_i$ be the sequence of points of the sphere of directions sampled by Chew's algorithm, with a sizing field equal to zero. For any point $p \in E'_I$, we maintain a counter i_p. Each time the algorithm has to probe new points of S from p, it increments the counter i_p and then chooses the unit vector $\mathbf{v}_{i_p}$ as the new direction in which to probe. The point that is inserted in E is the first intersection point of the ray $(p, \mathbf{v}_{i_p})$ with S. Detecting this point can be done with the exhaustive oracle, but not with the bipolar oracle. Remark that, in order to prevent the points of E from getting too close to one another, the new point is inserted in E only if it is farther than $\frac{\lambda}{2}$ from E, where λ is defined as in Section 4.4.

The following loop summarizes this version of the algorithm:

while $\mathcal{Q}$ is not empty **or some point** p **of** E'_I **is not a vertex of** $\mathrm{Del}_{|S}(E)$ {
  if $\mathcal{Q} \neq \emptyset$ { *// proceed as before*
    take the largest element of $\mathcal{Q}$;
    insert its center into E and update $\mathrm{Del}(E)$;
    update $\mathrm{Del}_{|S}(E)$;
    update $\mathcal{Q}$;
  }
  else { *// $p \in E'_I$ is not a vertex of* $\mathrm{Del}_{|S}(E)$
    choose vector $\mathbf{v}_{i_p}$ as the new direction in which to probe;
    insert in E the first point of intersection of S with the ray $(p,\ \mathbf{v}_{i_p})$ (if it exists);
    update $\mathrm{Del}_{|S}(E)$;
    update $\mathcal{Q}$;
  }
}

Lemma 5.2 shows that this version of the algorithm terminates and has the same output guarantees as the one that uses persistent facets. The proof relies on the fact that, in the "worst-case" scenario, the algorithm creates persistent facets on all the connected components of S, which implies that I2 is eventually verified.

Lemma 5.2 *The modified algorithm terminates and the quality of its output is guaranteed, provided that the input parameter ε is at most $\frac{\lambda}{2}$.*

We give the proof in the k-Lipschitz case. It holds in the smooth case as well, since smooth surfaces are k-Lipschitz, by Theorem 2.6.

Proof. Let $E_c \subset E_F \setminus E_I$ be the set of centers of bad surface Delaunay balls inserted by the algorithm. Since ε is positive, E_c is finite, by the same packing argument as in the proof of Lemma 4.7.

Let $p \in E'_I$. We call S_p the connected component of S on which p lies. Let $p' \in \tilde{S} \cap B(p, \mathrm{lr}_k(p))$. By the cocone lemma 2.12, $S \cap B(p, \mathrm{lr}_k(p))$ lies outside the double cone $K(p)$ of apex p, of axis aligned with $\mathbf{n}(p')$ and of half-angle $\theta = 2 \arctan k$. It follows that, inside one of the cones of $K(p)$ (say $K_1(p)$), all vectors point towards the interior of the compact object bounded by S. Therefore, any ray r cast in $K_1(p)$ from p is such that its first intersection point with S lies on S_p, since Ω has no void.

As a consequence, if at some stage the algorithm detects that p is not a vertex of $\mathrm{Del}_{|S}(E)$, then it inserts a new point of S_p in E, provided that the direction $\mathbf{v}_{i_p}$ lies in $K_1(p)$. Since the density of the set $\{\mathbf{v}_i\}$ of directions sampled by Chew's algorithm increases uniformly with i, the algorithm will eventually pick up a direction $\mathbf{v}_{i_p}$ that belongs to $K_1(p)$. Therefore, after a finite number of iterations of the main loop, either p has become and will remain a vertex of $\mathrm{Del}_{|S}(E)$, or there are enough points of S_p in E to ensure that $E \cap S_p$ is locally a $\frac{\lambda}{2}$-sample on some part S'_p of S_p that is visible from p.

At this stage, $\mathrm{Del}_{|S'_p}(E)$ has at least one facet, by the same argument as in the proof of Lemma 2.38. This facet is circumscribed by a ball of radius at most $\frac{\lambda}{2}$, since $E \cap S'_p$ is a $\frac{\lambda}{2}$-sample of S'_p. Therefore, $\mathrm{Del}_{|S}(E)$ contains at least one λ-persistent facet on S_p. Since this is true for every $p \in E'_I$, and since E'_I intersects all the connected components of S, $\mathrm{Del}_{|S}(E)$ has λ-persistent facets on every connected component of S. Moreover, every surface Delaunay ball has a radius of at most ε, since the algorithm refines surface Delaunay balls before casting rays from the points of E'_I. Therefore, by the same arguments as in the proof of Theorem 4.5, we know that E is a $\varepsilon\sqrt{1 + \frac{1}{\cos^4\theta}}$-sample of S. Since $\varepsilon \leq \frac{\lambda}{2} < \frac{1}{2}\,\mathrm{lr}_k(S)$, E is *a fortiori* a $\frac{1}{2}\mathrm{lr}_k(S)$-sample of S. Hence, by Lemma 2.38, all the points of E (and in particular those of E'_I) are vertices of $\mathrm{Del}_{|S}(E)$, which makes the algorithm terminate. The guarantees on its output come from the fact that E is a $\varepsilon\sqrt{1 + \frac{1}{\cos^4\theta}}$-sample of S upon termination. □

Although the size of E_F is bounded from above, we have no upper bound on the number of iterations of the algorithm, because not all the points of S probed from the points of E'_I are inserted in E. But still, we know that the algorithm terminates. We could reduce the number of probes by using the vector $\mathbf{n}$ introduced in section 4.4 as an approximate of the normal of S in the vicinity of the points of E'_I.

In the case where no input point set E_I is specified, we can replace the pre-processing step by just computing a few random points on S by shooting along random lines of $\mathbb{R}^3$. As reported in Table 5.1, with this simple procedure, $\mathrm{Del}_{|S}(E_F)$ has vertices on all the connected components of S with a

[1] These surfaces are shown in Chapter 6, Figures 6.1 and 6.2.

Surface \ $\lvert E_I \rvert$	5	7	10	12	15	20	50	100
Torus	16	32	75	80	100	100	100	100
Chair	6	12	63	92	100	100	100	100
Tanglecube	2	6	16	52	92	100	100	100
Trefoil knot						2	36	89

Table 5.1: Success rates (in %) of the algorithm with random initial point samples, on various surfaces[1].

probability that increases dramatically with the number of initial random points. The probability depends also on how well S fills space, as shown by the example of the trefoil knot, which fills space very little and therefore requires a large number of initial random points to be intersected by $\mathrm{VG}(E_F)$.

5.3 Removing the skinny facets

Once an ε-sample E_I of S has been built or is given, one can remove the skinny facets from the mesh, by simply running the algorithm with E_I as the initial point sample and by using a new definition of a bad surface Delaunay facet. In this section, a surface Delaunay facet is called *bad* if it is skinny, *i.e.* if at least one of its inner angles is less than a user-defined parameter $\tau \geq 0$. The output point sample E_F contains E_I and hence is an ε-sample of S. Moreover, all the facets of $\mathrm{Del}_{|S}(E_F)$ are well-shaped, *i.e.* no angle is less than τ.

This variant of the algorithm is due to Ruppert [99]. Although it can be applied indifferently on smooth and nonsmooth surfaces, with the same guarantees, it is only interesting in the smooth setting. Indeed, as observed in Section 4.2.3 (Theorem 4.5), our theoretical guarantees in the nonsmooth setting require that the density of the output point sample be uniform, which ensures automatically that the inner angles of the facets of the output mesh are lower-bounded.

Recall that any point p inserted by the algorithm is the center of some surface Delaunay ball $B(p)$. We call $r(p)$ the radius of $B(p)$, which equals the distance from p to the point sample E maintained by the algorithm, right before the insertion of p in E. Moreover, we call $r_2(p)$ the distance from p to its second nearest neighbor in E_I. We extend the definition to the points p of E_I by setting $r(p) = r_2(p) = \mathrm{d}(p, E_I \setminus \{p\})$. The (positive) map r_2 was introduced by Ruppert [99] and proved to be 1-Lipschitz.

From now on, we assume that $\tau < \frac{\pi}{6}$ and we define $\varrho_\tau = \frac{1}{2\sin\tau} > 1$. Note that ϱ_τ is closely related to the ϱ of Theorem 4.5. However, both quantities differ in the sense that ϱ yields a lower bound on the distance between any two points of E_F, whereas ϱ_τ yields a bound on the distance between any two neighbors in $\mathrm{Del}_{|S}(E_F)$.

Lemma 5.3 $\forall p \in E_F$, $r(p) \geq \frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(p)$.

Proof. We reason by induction. We first prove that the result holds for the points of E_I, then we show that it holds also for every point inserted by the algorithm.
- Let $p \in E_I$. By definition, we have $r(p) = r_2(p) > \frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(p)$.

- Let $p \in E_F \setminus E_I$. Right before the insertion of p in E, p is the center of a ball of $\mathrm{Del}_{|S}(E)$ that circumscribes a facet f whose smallest inner angle τ_f is less than τ. Let r be the circumradius of f. We have $r \leq r(p)$. Let e be the smallest edge of f, and v be the vertex of e that was inserted last (if both points belong to E_I, then we choose v to be any of them). Let w be the other vertex of e. We have $\mathrm{d}(v,w) \geq r(v)$. Moreover, a quick computation shows that the sine of the angle τ_f (which is opposite to e) is equal to half the ratio between $\mathrm{d}(v,w)$ and the circumradius of f. Hence,

$$\sin\tau = \sin\tau_f \geq \frac{\mathrm{d}(v,w)}{2r} \geq \frac{r(v)}{2r(p)}$$

It follows that $r(p) \geq \varrho_\tau\, r(v)$. Now, according to the induction hypothesis, we have $r(v) \geq \frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(v)$, which is at least $\frac{\varrho_\tau - 1}{\varrho_\tau}\, (r_2(p) - \mathrm{d}(p,v))$ since r_2 is 1-Lipschitz. Thus, $r(p) \geq \varrho_\tau\, \frac{\varrho_\tau - 1}{\varrho_\tau}\, (r_2(p) - r(p))$, *i.e.* $r(p) \geq \frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(p)$. □

Let $r_2^{\inf} > 0$ be the infimum of r_2 over S. It follows from Lemma 5.3 that the distance between any two points of E_F is at least $r_2^{\inf}$. Then, by the same packing argument as in the proof of Lemma 4.7, the algorithm terminates.

Corollary 5.4 *E_F is $\frac{\varrho_\tau - 1}{2\varrho_\tau - 1} r_2$-sparse. Hence, by Theorem 3.4,*

$$|E_F| = O\left(\left(2 + \frac{1}{\varrho_\tau - 1}\right)^2 \iint_S \frac{dS}{r_2^2}\right)$$

Proof. Let $p \in S$. We will bound the number of points of $E_F \cap B_p$, where B_p is the ball of center p and radius $\frac{\varrho_\tau - 1}{2\varrho_\tau - 1}\, r_2(p)$. Let q, q' be any two points of $E_F \cap B_p$. We assume without loss of generality that q and q' both belong to E_I, or that q has been inserted after q' by the algorithm. Then, we have $\mathrm{d}(q,q') \geq r(q)$, which is at least $\frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(q)$, by Lemma 5.3. Since r_2 is 1-Lipschitz and q lies in B_p, we have

$$\frac{\varrho_\tau - 1}{\varrho_\tau}\, r_2(q) \geq \frac{\varrho_\tau - 1}{\varrho_\tau}\, (r_2(p) - \mathrm{d}(p,q)) \geq \frac{\varrho_\tau - 1}{2\varrho_\tau - 1}\, r_2(p)$$

Hence, the points of $E_F \cap B_p$ are centers of pairwise disjoint balls of radius $\frac{\varrho_\tau - 1}{2(2\varrho_\tau - 1)}\, r_2(p)$. Moreover, for any $q \in E_F \cap B_p$, the ball $B\left(q, \frac{\varrho_\tau - 1}{2(2\varrho_\tau - 1)}\, r_2(p)\right)$ is included in $B\left(p, \frac{3(\varrho_\tau - 1)}{2(2\varrho_\tau - 1)}\, r_2(p)\right)$. It follows that the number of points $q \in E_F$ that lie in B_p is at most

$$\frac{\mathrm{Vol}\left(B\left(p, \frac{3(\varrho_\tau - 1)}{2(2\varrho_\tau - 1)}\, r_2(p)\right)\right)}{\mathrm{Vol}\left(B\left(q, \frac{\varrho_\tau - 1}{2(2\varrho_\tau - 1)}\, r_2(p)\right)\right)} \;=\; 27$$

□

The constant hidden in the O in the bound of Corollary 5.4 is absolute, *i.e.* it depends neither on S, nor on E_I, nor on ϱ_τ. As ϱ_τ tends to 1 (*i.e.* τ approaches $\frac{\pi}{6}$ from below), the bound goes to infinity. This means that, the better the angle criterion, the more points the algorithm has to insert. On the contrary, as ϱ_τ tends to infinity (*i.e.* τ goes down to zero), the bound becomes $O\left(\iint_S \frac{dS}{r_2^2}\right) = O(|E_I|)$, which is coherent with the fact that the algorithm does not have to insert any point since E_I already meets the angle criteron.

5.4 Non-uniform sizing fields

In some situations, it may be interesting to generate point sets with a local density prescribed by a user-defined sizing field σ. Recall that a point set $E \subset S$ is a σ-sample of S if no point p of S is farther than $\sigma(p)$ from E.

Assuming the infimum $\sigma^{\inf}$ of σ over S is positive and smaller than $\mu_0 \operatorname{rch}(S) \approx 0.16 \operatorname{rch}(S)$ in the smooth case and $\frac{27}{193} \operatorname{lr}_k(S) \approx 0.14 \operatorname{lr}_k(S)$ in the k-Lipschitz case, we can run Chew's algorithm with $\varepsilon = \frac{1}{2}\sigma^{\inf}$. By Theorems 4.4 and 4.5, the output point set E_F is a $\sigma^{\inf}$-sample of S, which *a fortiori* is a σ-sample of S. However, E_F is uniform, *i.e.* its local density does not vary with σ. As a consequence, its size is $\Theta\left(\iint_S \frac{dS}{\left(\sigma^{\inf}\right)^2}\right)$, whereas in the case where σ is Lipschitz the optimal size is $\Theta\left(\iint_S \frac{dS}{\sigma^2}\right)$, by Theorems 3.1 and 3.4.

In this section we show how to generate sparse σ-samples, for a given positive Lipschitz sizing field σ. Let $\kappa < 1$ denote the Lipschitz constant of σ. Since σ is not uniform, the theoretical results of Chapter 2 no longer apply. Therefore, in the sequel, we assume that S has a positive reach.

To generate sparse σ-samples of S, we run Chew's algorithm with ε replaced by $\frac{\sigma}{2+\kappa}$. A surface Delaunay ball $B(c,r)$ is then *bad* if $r > \frac{\sigma(c)}{2+\kappa}$. During the pre-processing phase we generate $0.09\mathrm{d}_{\mathrm{M}}$-persistent facets, that is, facets circumscribed by surface Delaunay balls $B(c,r)$ such that $r \leq \frac{0.09\ \mathrm{d}_{\mathrm{M}}(c)}{3} = 0.03\ \mathrm{d}_{\mathrm{M}}(c)$. The construction is identical to the one described in Section 4.4. Notice that we now have a factor of $\frac{1}{3}$ in the definition of a persistent facet, instead of a factor of $\frac{1}{2}$ in Section 4.2.1. This is because d_{M} is 1-Lipschitz, and not constant as ε.

If $\sigma \leq 0.09\ (2+\kappa)\ \mathrm{d}_{\mathrm{M}}$, then the output point set E_F is a σ-sample of S. Indeed, upon termination of the algorithm, for every surface Delaunay ball $B(c,r)$ we have $r \leq \frac{\sigma(c)}{2+\kappa} \leq 0.09\ \mathrm{d}_{\mathrm{M}}$. Moreover, a straightforward adaptation of Lemma 4.3 shows that our $0.09\mathrm{d}_{\mathrm{M}}$-persistent facets remain in $\mathrm{Del}_{|S}(E)$ until E becomes a loose $0.09\mathrm{d}_{\mathrm{M}}$-sample of S. At that stage, by Theorem 1.33, E is a $0.1\mathrm{d}_{\mathrm{M}}$-sample, and it remains such until the end of the process. As a consequence, $\mathrm{Del}_{|S}(E_F)$ has vertices on every connected component of S, by Lemma 1.22. It follows that E_F is a loose $\frac{\sigma}{2+\kappa}$-sample of S, which, by Theorem 1.33, implies that the balls of $\mathrm{Del}_{|S}(E_F)$ cover S. Then, for any point $p \in S$, there exists some surface Delaunay ball $B(c,r)$ containing p. Since σ is κ-Lipschitz, we have

$$\sigma(p) \geq \sigma(c) - \kappa\ \mathrm{d}(c,p) \geq \left(1 - \frac{\kappa}{2+\kappa}\right)\ \sigma(c) = \frac{2}{2+\kappa}\ \sigma(c)$$

Thus, the distance from p to E_F is bounded by

$$2\mathrm{d}(p,c) \leq \frac{2}{2+\kappa}\ \sigma(c) \leq \sigma(p)$$

which means that E_F is a σ-sample of S.

In addition, once the input point set E_I is fixed, E_F is σ-sparse. Indeed, every point $p \in E_F \setminus E_I$ is farther than $\frac{\sigma(p)}{2+\kappa}$ from E at the time of its insertion. Therefore, for any two points $p, q \in E_F \setminus E_I$ (with q inserted after p, without loss of generality) we have

$$\mathrm{d}(p,q) \geq \frac{\sigma(q)}{2+\kappa} \geq \frac{\sigma(p)}{2+\kappa} - \frac{\kappa}{2+\kappa}\ \mathrm{d}(p,q)$$

which yields

$$\mathrm{d}(p,q) \geq \frac{\sigma(p)}{2(1+\kappa)}$$

Hence, the points $p \in E_F \setminus E_I$ are centers of pairwise disjoint balls of radii $\frac{\sigma(p)}{4(1+\kappa)}$. It follows, by the same argument as the in the proof of Theorem 4.4, that $E_F \setminus E_I$ is σ-sparse, which implies that E_F is also σ-sparse since E_I is fixed.

As a conclusion, for sufficiently small σ, the output point set is a sparse σ-sample of S, hence its size is $\Theta\left(\iint_S \frac{dS}{\sigma^2}\right)$, by Theorems 3.1 and 3.4.

Since our results hold only for σ less than a fraction of $\mathrm{d_M}$, it is mandatory to have an estimate of $\mathrm{d_M}$, to know whether the theoretical guarantees apply or not. One can also take $\min\{\sigma, \mathrm{d_M}\}$ as input sizing field of the algorithm, in order to both certify the quality of the output mesh and make the density of its vertices depend on the user's sizing field σ.

In some particular situations, *e.g.* when the surface S is an isosurface in a 3D greyscaled image, $\mathrm{d_M}$ can be computed exactly or at least approximated well from the input – see Section 6 for more detail. However, in most cases computing $\mathrm{d_M}$ is a formidable task. Even computing the radius of the maximal empty ball B tangent to S at a given point p can be difficult, since it involves non-local operations such as finding another tangency point of B and S, which may lie far away from p.

Instead, we remark that computing a positive constant ε_S less than $\mathrm{rch}(S)$ is easier than estimating $\mathrm{d_M}$ at a given set of points. Moreover, the computation can be done once for all, in a pre-processing phase. Hence, one can use Chew's algorithm with $\varepsilon = \mu_0 \mathrm{rch}(S)$ to generate a uniform point sample, from which an approximation of $\mathrm{d_M}$ can be computed. The problem of this approach is that, although it produces good results in practice, it is not theoretically certified since none of the current results on medial axis approximation from finite sets of points guarantees a full Hausdorff approximation [8, 11, 35, 59].

Open question 5.5 *Given a smooth surface S and an ε-sample E of S, is it possible to bound the Hausdorff distance between S and some subcomplex of* $\mathrm{Vor}(E)$ *by a function of ε?*

Part C

Applications

Introduction

In this part of the thesis, we illustrate the universality of our mesh generator through several practicle applications. Our implementations are done in C++, using the geometric library CGAL [33], which provided us with a robust and efficient implementation of the 3-dimensional Delaunay triangulation.

In Chapter 6, we apply our algorithm to mesh implicit surfaces. Our experimental results provide evidence that our mesher has several advantages over the highly celebrated Marching Cubes algorithm [44, 87]: topological guarantees, approximation of normals, facets with bounded aspect ratio, optimal size. We also argue that, as an implicit surface mesher, our algorithm can be used for surface reconstruction. The approach is similar to that of [19]. It consists of defining an implicit function f' from the input point set, and then to mesh the zero-set S' of f' using our algorithm. Combining our theoretical results with those of [19], we can certify the topology and geometry of the output mesh.

In Chapter 7, we use our algorithm to remesh polygonal surfaces. Compared to many existing methods (see [4, 5]), our mesher does not require to perform complicated operations on the input polyhedron S, such as cutting S into patches homeomorphic to a disk. Moreover, it comes with topological and geometric guarantees, together with a bound on the aspect ratios of the triangles. Finally, differently from the polyhedron remesher of Dey *et al.* [57], we do not assume that S approximates a smooth surface.

In Chapter 8, we adapt our algorithm to the probing problem. The approach consists of using the probing device as an oracle to be used by our mesher. This oracle is weaker than the exhaustive oracle, since it can detect only the first intersection point of a given segment with the boundary S of the object to discover. Moreover, before checking the intersection of a given segment s with S, the probing device must first be moved to an endpoint of s. Therefore, we cannot check the intersections of all the segments with S. We prove however that this version of the algorithm comes with the same theoretical guarantees as the original version, regarding the quality and size of the output. We also provide a thorough analysis of its complexity.

In Chapter 9, we present an extension of the algorithm that can construct tetrahedral meshes approximating 3-dimensional objects with curved boundaries, such that the mesh elements (tetrahedra and triangles) conform to some user-defined size and shape criteria. The idea is to exploit the fact that the algorithm maintains a whole 3-dimensional Delaunay triangulation. Whenever a tetrahedron does not meet the size or shape requirements, it is refined by inserting its circumcenter. The output point set is no longer a subset of S. Moreover, the output mesh contains all Delaunay tetrahedra whose circumcenters lie in the object $\mathcal{O}$ to mesh. Using our theoretical results on the approximation of the boundary of $\mathcal{O}$, we can certify the output of the algorithm.

Chapter 6

Implicit Surface Meshing

6.1 Introduction

Implicit surfaces are widely used to model geometric objects like molecules in biomolecular modeling, smooth surfaces reconstructed from MRI data in image-guided surgery, or car parts in CAD/CAM-systems. The representation of a geometric object as an implicit surface may be very convenient for such physical or technical applications – see for instance [17, 114]. However, for further geometry processing, like rendering the object on a computer screen or performing numerical simulations, it may be more convenient to have a PL approximation of the object. Therefore, fast and reliable meshing algorithms have immediate applications in these practical situations.

The goal of an implicit surface mesher is to generate a PL approximation that is close to the original surface, both in a geometric and in a topological sense. Although many meshing techniques produce good geometric approximations [16, 113], very few algorithms are guaranteed to construct topologically correct meshes. Some polygonization schemes achieve topological consistency, that is, they ensure that the result is a manifold, by taking more or less arbitrary decisions when a topologically ambiguous configuration is encountered. Nevertheless, topology is still not guaranteed, since the algorithm may miss some sonnected components or handles of the original object. Several adaptations of the Marching Cubes algorithm [87], such as the one proposed in [44], belong to this category.

The first topology-preserving implicit surface mesher was designed by Hart and Stander [76]. Unfortunately, they gave no proof of correctness and, although their method has an intuitive justification and seems to work well on simple cases, it is not clear that it is robust on complex shapes. A more recent algorithm, proposed by Boissonnat *et al.* [20], is based on the same mathematical tools and comes with a proof of correctness. Plantinga and Vegter [96] introduced a variant of the Marching Cubes algorithm with topological guarantees, based on interval arithmetics. Finally, Cheng *et al.* [42] proposed an algorithm that guarantees the homeomorphism between the implicit surface and its PL approximation. All these algorithms work only in the smooth setting.

6.1.1 Statement of the problem

The input surface S is described as a level-set (say the zero-set) of some function $f : \mathbb{R}^3 \to \mathbb{R}$, that is: $S = f^{-1}(\{0\})$. We assume in first place that f is C^2-continuous and that its gradient does not vanish on

S, which implies, by the implicit function theorem, that S is a C^2-continuous surface without boundary – see for instance [14, §2.1.2]. The nonsmooth case is deferred to Section 6.4.

Given a positive constant δ, the purpose is to construct a PL approximation $\hat{S}$ of S with the same topology type as S and at Hausdorff distance δ from S. Moreover, the number of vertices of $\hat{S}$ is required to be as small as possible.

6.1.2 Our approach

We use our implementation of Chew's algorithm. Following Section 4.2.2, we set the input parameter ε of the algorithm to $\sqrt{\delta\ \varepsilon_S}$ where ε_S is less than $\mathrm{rch}(S)$. Moreover, using the construction of Section 4.4, we build an initial point set E_I such that $\mathrm{Del}_{|S}(E_I)$ has at least one λ-persistent facet per connected component of S, where[1] $\lambda = \mu_0\ \varepsilon_S$.

If $\delta \leq \mu_0^2\ \mathrm{rch}(S)$, then we have $\varepsilon \leq \mu_0\ \mathrm{rch}(S)$, and by Theorem 4.4 the output of the algorithm is a PL approximation of S, isotopic to S and at Fréchet distance at most δ from S. Moreover, the size of the output point set is $\Theta_S\left(\varepsilon^{-2}\right) = \Theta_S\left(\delta^{-1}\right)$, which is optimal when the object bounded by S is convex, as detailed in Section 4.2.2. Notice that the same guarantees apply if the exhaustive oracle is replaced by the bipolar oracle, according to Section 5.1.

If the density of the vertices of $\hat{S}$ is prescribed by a user-defined positive and κ-Lipschitz sizing field σ_u, with $\kappa < 1$, then we adopt the strategy of Section 5.4 and run Chew's algorithm with $\sigma = \frac{\sigma_u}{2+\kappa}$ as input sizing field and with an initial point set composed of the vertices of at least one $0.03\mathrm{d_M}$-persistent facet per connected component of S.

As reported in Section 5.4, if σ_u is sufficiently small (specifically, if $\sigma_u \leq 0.09\ (2+\kappa)\ \mathrm{d_M}$), then the output point set E_F is guaranteed to be a sparse σ_u-sample of S. As a consequence, the Hausdorff distance between $\mathrm{Del}_{|S}(E_F)$ and S is at most $\mathrm{d}_{\mathcal{H}}(E_F, S) \leq \sup\{\sigma_u(p),\ p \in S\}$. Moreover, by Theorem[2] 1.33, $\mathrm{Del}_{|S}(E_F)$ has the same topology type as S. Here again, the same guarantees apply if we replace the exhaustive oracle by the bipolar oracle ($\mathrm{Del}_{|S}(E_F)$ becomes then $\mathrm{Del}^b_{|S}(E_F)$).

For both strategies, the prerequisites on the knowledge of the surface are the following ones:

- **P1** we can pick up at least one point from each connected component of S (to construct E_I),
- **P2** we know or we can compute a positive constant ε_S less than or equal to $\mathrm{rch}(S)$,
- **P3** we can implement the exhaustive oracle, and optionally the bipolar oracle.

6.1.3 Overview

In Section 6.2, we describe some implementations that fulfill P1, P2 and P3. Notice that P3 is used throughout the algorithm, whereas P1 is involved only in the initialization phase and P2 is a simple precalculation. Therefore, algorithmic issues arise mainly from P3.

Experimental results on smooth surfaces are presented in Section 6.3. The nonsmooth case is addressed in Section 6.4. Finally, Section 6.5 presents an application of implicit surface meshing in the context of surface reconstruction.

[1] $\mu_0 = 0.16$, as in Section 4.2.2.

[2] The bound on the Hausdorff distance in this theorem holds only for $\sigma_u = \mathrm{d_M}$, hence we cannot use it here.

6.2 Satisfying the prerequisites

The bipolar oracle is quite easy to implement for implicit surfaces. Indeed, the parity of the number of transversal intersections between S and a line segment $[p, q]$ is given by the signs of f at p and q: if the signs are equal, then the number of transversal intersections is even, otherwise it is odd. In the latter case, we can find a transversal intersection point by binary search. The computation time is $O(1)$ for the bipolarity test, and $O\left(\log \frac{\mathrm{d}(p,q)}{\eta}\right)$ for finding an intersection point within a precision of $\eta > 0$. The constant in the O depends on the time spent to compute the sign of f at a given point. When f is a polynomial, we use Descartes' rule instead of evaluating the signs of f at p and q, which reduces the computation time.

The exhaustive oracle can be implemented using a divide-and-conquer strategy. However, computing all the intersection points of S with a segment $[p, q]$ within a precision of $\eta > 0$ takes $\Theta\left(\frac{\mathrm{d}(p,q)}{\eta}\right)$ time. Hence, for implicit surfaces, the exhaustive oracle is far less efficient than the bipolar oracle. Timings are reported in Table 6.1.

We will now focus on P1 and P2. In Section 6.2.2, we discuss the case where a closed formula of f is known. In Section 6.2.2, we analyze the other case through a specific application.

6.2.1 When a closed formula of f is known

P1 We compute the points of S that have a horizontal tangent plane. Each connected component of S has at least two such points, since S has no boundary. These points are the critical points of the height function, which means that they are solutions of the following system:

$$\begin{cases} f(x,y,z) = 0 \\ \frac{\partial f}{\partial x}(x,y,z) = 0 \\ \frac{\partial f}{\partial y}(x,y,z) = 0 \end{cases}$$

which is generically zero-dimensional. If f is a polynomial, then the system is algebraic and can be solved by various means. Our approach consists in computing the generalized normal form modulo the ideal generated by the three polynomials of the system, and then finding the roots from eigencomputation – this method was developed in [111] and implemented in C^{++} as part of the SYNAPS library [107], which we use in our implementation. If f is not a polynomial but still continuous, then we compute the solutions of the system using interval arithmetics. Notice that this computation may be quite complicated but is invoked only once, during the initialization phase of the algorithm.

P2 Since S is compact, there exists a point $p \in S$ such that $\mathrm{d}_{\mathrm{M}}(p) = \mathrm{rch}(S)$. Let c be a nearest neighbor of p on M. We have $\mathrm{d}(c, S) = \mathrm{d}(c, p)$, since $\mathrm{d}(c, p) = \mathrm{rch}(S)$. Hence, p is a nearest neighbor of c on S. If c belongs to the boundary of M, then $\mathrm{d}(c, p)$ equals the minimum radius of curvature of S at p. Otherwise, c has another nearest neighbor on S, say q. Since $\mathrm{d}(c, q) = \mathrm{d}(c, p) = \mathrm{rch}(S)$, c is a nearest neighbor of q on M. Therefore, the balls $B(p, \mathrm{d}(c, p))$ and $B(q, \mathrm{d}(c, q))$ are both tangent to M in c. It follows that p, c and q are collinear, since c belongs to the relative interior of M. This implies that c is the midpoint of the line segment $[p, q]$, since $p \neq q$. So, the ball $B(c, \mathrm{rch}(S))$ is tangent to S in two diametral points. Therefore, $\mathrm{rch}(S)$ can be computed by finding:

1. the point of S at which the smallest radius of curvature is minimal, which reduces to solving some low-dimensional optimization problem over S;
2. the smallest sphere bitangent to S with diametral contact points, which reduces to finding the smallest real positive root of the following system:

$$\begin{cases} f(p) = 0 \\ f(q) = 0 \\ (\mathbf{p} - \mathbf{q}) \times \boldsymbol{\nabla} f(p) = \mathbf{0} \\ (\mathbf{p} - \mathbf{q}) \times \boldsymbol{\nabla} f(q) = \mathbf{0} \\ \lambda \, \mathrm{d}^2(p, q) = 1 \end{cases}$$

This system has seven unknowns (the coordinates of p, those of q, and λ) and seven independent equations generically, therefore it is zero-dimensional. If f is a polynomial of degree d ($d \geq 2$ since S is compact), then the system is algebraic, of degree d. Note that λ and the last equation of the system have been introduced only to ensure that $p \neq q$.

The two above issues can be solved using the same tools as for P1.

6.2.2 When no closed formula of f is known

The computations involved in the previous section require to know a closed formula of f. However, in many applications no such formula is available or even exists. In this case, P1 and P2 may still possibly be fulfilled but the strategy depends highly on the context.

We will illustrate this through a specific application related to medical imaging, in which the surface S is defined as a level-set in a $3D$ greyscaled image. Let f be the function that maps every point of space to its grey value. Due to the nature of the image, the values of f are known only at the centers of the pixels, *i.e.* at a discrete but regular set of points. We can then retrieve the value of $f(p)$ for any $p \in \mathbb{R}^3$ by interpolation.

P1 We compute the points of S that have a horizontal tangent plane, by sweeping a horizontal plane vertically across the image. Since S has no boundary, this gives us at least two points per connected component of S.

P2 We use a thinning algorithm [37] to compute a discrete approximation $\tilde{\mathrm{M}}$ of the medial axis of S. An estimate of $\mathrm{rch}(S)$ can be easily computed from $\tilde{\mathrm{M}}$. Moreover, using the strategy of Section 5.4, we can generate point sets whose densities depend on d_{M}, since a good estimate of d_{M} can be computed from $\tilde{\mathrm{M}}$.

6.3 Experimental results

Results on smooth closed algebraic surfaces are reported in Figure 6.1. The top row shows the inputs, the bottom row shows the outputs. From left to right, we have a torus, a genus-three surface of degree 4 called "chair", and a genus-five surface of degree 4 called "tanglecube". These surfaces have been

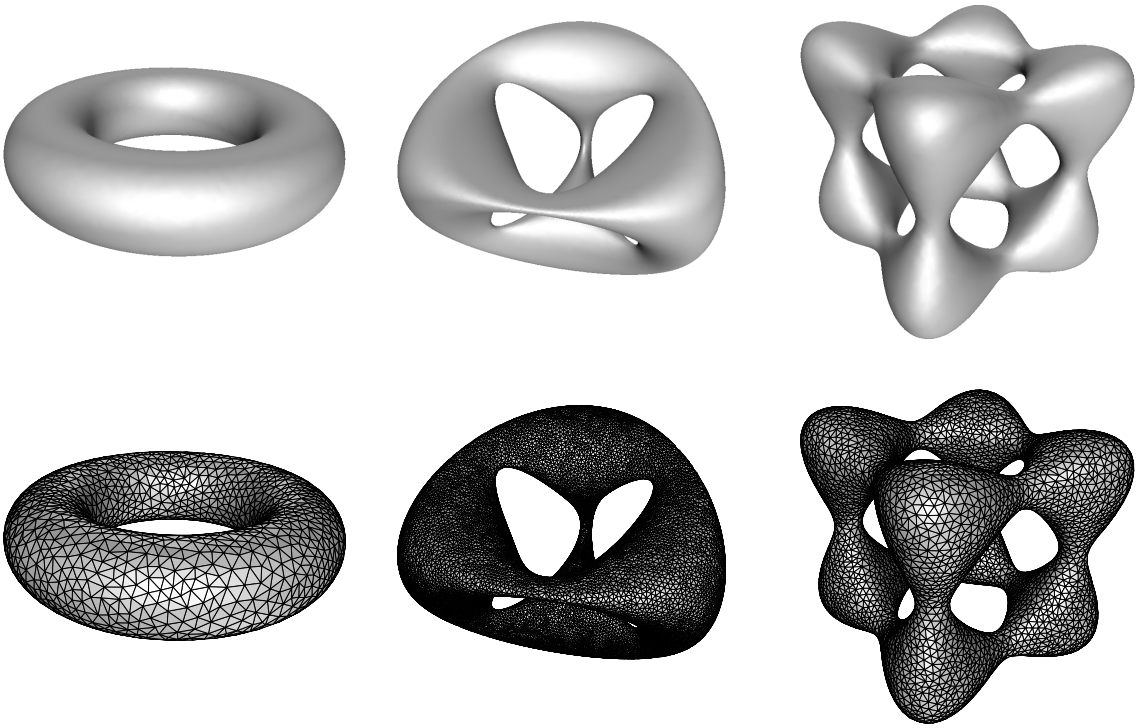

Figure 6.1: Results on smooth closed algebraic surfaces, with $\varepsilon = \mu_0 \, \varepsilon_S$.

meshed with $\varepsilon = \mu_0 \, \varepsilon_S$ as input parameter. As predicted by Theorem 4.4, we obtained good topological and geometric approximations of the surfaces.

Figures 6.2 and 6.3 show the results of the algorithm respectively on the standard left trefoil knot and on a more intricate knot with high self-entanglement. In both cases, the result of the algorithm is isotopic to the original algebraic model, as predicted by the theory. Notice that every knot is represented in "sausage" format, *i.e.* as the boundary of a thickening of some knotted curve. Figure 6.2 (left) shows that such a curve can be defined as the intersection of two algebraic surfaces, $f_1 = 0$ and $f_2 = 0$, which are the images through a stereographic projection of two 2-manifolds embedded in the unit sphere of $\mathbb{R}^4$, as explained in [28]. One possible thickening of the curve $f_1 = f_2 = 0$ is $f_1^2 + f_2^2 < \nu$, whose boundary $f_1^2 + f_2^2 = \nu$ is a smooth closed surface, for a sufficiently small constant ν. Notice that this thickening does not have a constant radius, as can be observed in Figure 6.2 (right).

Timings for the above algebraic models are reported in Table 6.1, normalized with respect to the 3D incremental Delaunay triangulation algorithm implemented in CGAL [33]. The precalculations associated with P1 and P2 are not taken into account here. We tried two versions of the oracle: exhaustive and bipolar. For each version, we separated the time spent in the oracle from the time spent in the rest of the program (the *engine*).

The first observation is that the ratio between the timings of the two versions of the oracle is quite large. In the table, it ranges from several units for small models (*e.g.* 5.8 for the sphere), to hundreds for bigger models (*e.g.* 187 for the chair). As explained previously, this phenomenon is due to our implementation, which makes the running time of the exhaustive oracle linear w.r.t. $1/\eta$, where η is a user-defined precision threshold (10^{-6} here), while the time complexity of the bipolar oracle is logarithmic w.r.t. $1/\eta$.

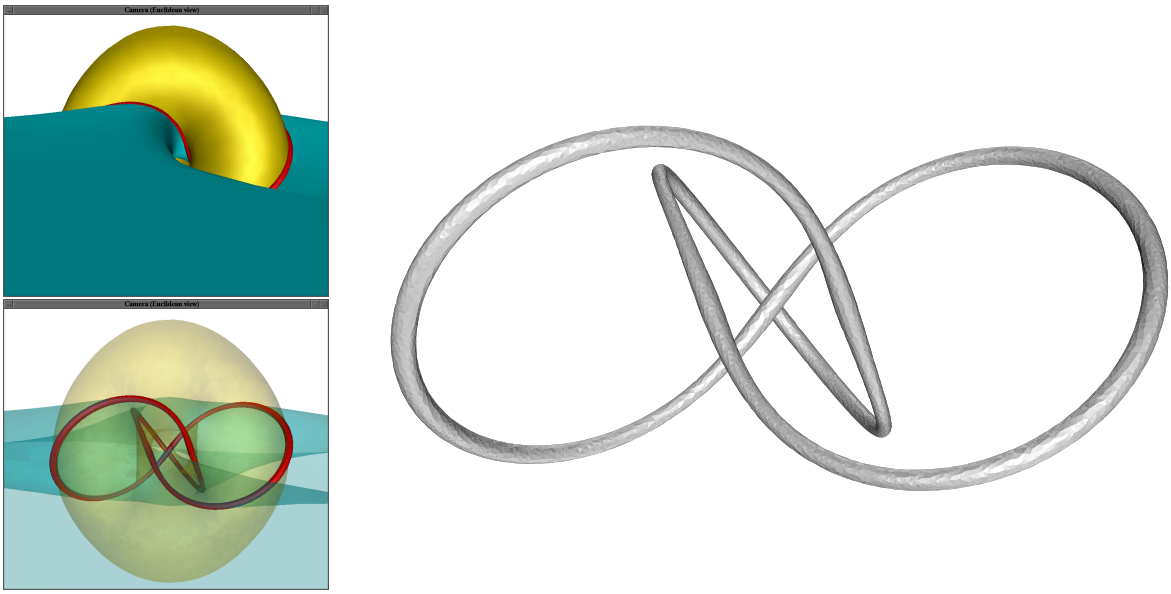

Figure 6.2: The standard left trefoil knot, in "sausage" format, meshed with $\varepsilon = \mu_0\ \varepsilon_S$.

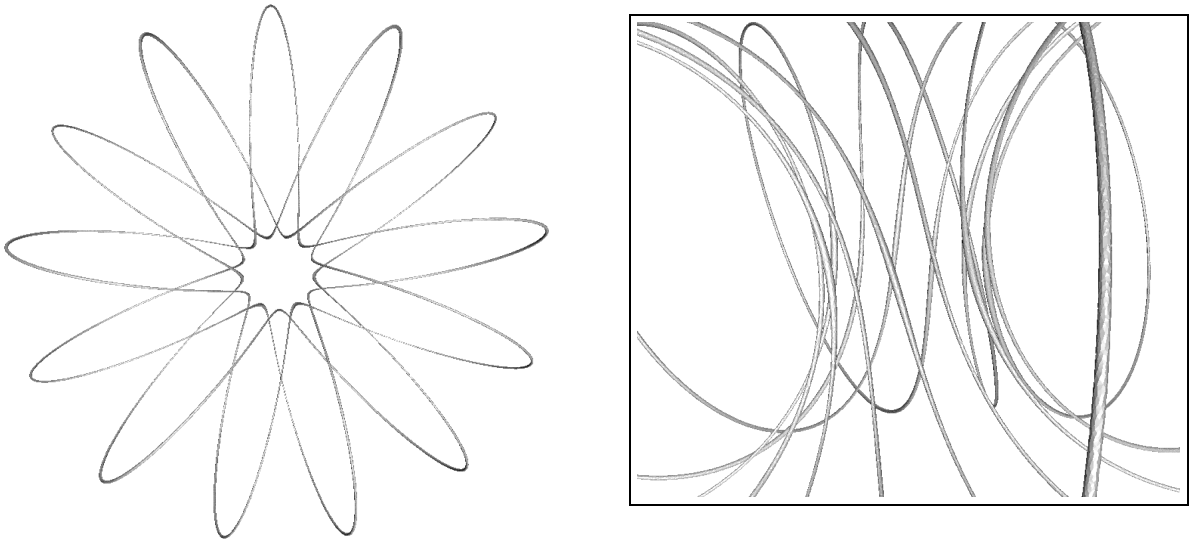

Figure 6.3: An intricate knot, in "sausage" format, meshed with $\varepsilon = \mu_0\ \varepsilon_S$.

Surface	Exhaustive oracle (uses $\mathrm{Del}_{\mid S}(E)$)							Bipolar oracle (uses $\mathrm{Del}^b_{\mid S}(E)$)						
	Engine		Oracle			Total time / Del.	Output size (# pts)	Engine		Oracle			Total time / Del.	Output size (# pts)
	Time/Del.	%	# calls	Time/Del.	%			Time/Del.	%	# calls	Time/Del.	%		
Sphere	1.92	14.6	$8,929$	11.4	85.4	13.3	445	1.98	86.8	$3,226$	0.3	13.2	2.28	445
Ellipsoid	7.5	6.3	$7,410$	111	93.7	118	380	8.25	91.7	$8,170$	0.75	8.3	9	380
Torus	0.87	0.2	$58K$	474	99.8	475	$1,268$	9.18	88.7	$63K$	1.17	11.3	10.35	$1,307$
Chair	8.01	0.5	$260K$	$1,686$	99.5	$1,694$	$6,619$	8.1	89.4	$255K$	0.96	10.6	9.06	$6,461$
Tanglecube	8.31	2.9	$168K$	354	97.1	362	$4,225$	7.71	90.5	$165K$	0.81	9.5	8.52	$4,242$
Trefoil knot	12.54	1.6	$545K$	783	98.4	796	$8,329$	12.54	93	$592K$	0.93	7	13.47	$8,317$
Intricate knot	7.68	2.3	$5.2M$	330	97.7	338	$133K$	6.93	89.9	$7M$	0.78	10.1	7.71	$148K$
Barth's octic	6.93	1.7	$547K$	396	98.3	403	$13,928$	7.05	83.9	$603K$	1.35	16.1	8.4	$14,168$
Heart	6.72	7	$252K$	89	93	96	$8,445$	6.45	91.8	$263K$	0.57	8.2	7.02	$8,539$
Klein's bottle	0.45	3.5	$101K$	123	96.5	123	$3,424$	4.74	91.9	$102K$	0.42	8.1	5.16	$3,445$

Table 6.1: Timings (normalized w.r.t. 3D incremental Delaunay triangulation algorithm) and output size for various algebraic surfaces.

The second observation is that the timings of the engine are comparable to the timings of the incremental Delaunay triangulation algorithm (the ratio is almost always less than 10). This observation is not surprising since both algorithms are similar. The difference between the performances comes mainly from the oracle. With the bipolar oracle, the difference is small and the performances of both algorithms are comparable. To give an idea, on a Pentium IV at 3.6 GHz, the mesher with the bipolar oracle takes 5.14 seconds to sample $8,317$ points from the trefoil knot, while the incremental Delaunay triangulation algorithm of CGAL takes 0.41 seconds to triangulate the $8,317$ points.

Figure 6.4 shows some results of the algorithm on $3D$ greyscaled images. The input images are courtesy of the Epidaure and Odyssee teams, at INRIA. The first one has size $256^3 \approx 16.7M$ voxels, the second one $50 \times 60 \times 60 = 180K$ voxels. We used the strategy of Section 5.4, with an input sizing field σ equal to $0.09\ \mathrm{d}_{\tilde{\mathrm{M}}}$, where $\mathrm{d}_{\tilde{\mathrm{M}}}$ is the distance to the approximation $\tilde{\mathrm{M}}$ of M computed by the thinning algorithm, as described in Section 6.2.2. According to the theory, the output point set is a $0.27\mathrm{d}_{\tilde{\mathrm{M}}}$-sample, since $\mathrm{d}_{\tilde{\mathrm{M}}}$ is 1-Lipschitz.

The first mesh has about $25K$ vertices, which is far less than the size of the output of the Marching Cubes algorithm (several millions of vertices). The advantage of our mesher over the Marching Cubes algorithm is that the user can specify, through the input sizing field σ, at which level of detail he wants to work. The second mesh has approximately $100K$ vertices, which is comparable to the number of voxels of the input image. The reason for this is that $\mathrm{rch}(S)$ is small and S is very space-filling, which means that generating a $0.27\mathrm{d}_{\mathrm{M}}$-sample of S requires a lot of points.

6.4 The nonsmooth case

Let us now assume that $f : \mathbb{R}^3 \to \mathbb{R}$ is k-Lipschitz. If the so-called generalized gradient of f (as defined in [46]) does not vanish on S, then S is a k-Lipschitz surface, by Clarke's implicit function theorem [47, §7.1]. This case includes in particular the case where f is smooth and its gradient vanishes on S but not its generalized gradient.

The strategy used for smooth f can still be applied when f is Lipschitz, with the same theoretical guarantees. The only difference is that the input of Chew's algorithm must satisfy Conditions ($\hat{\alpha}'$), ($\hat{\beta}'$), ($\hat{\gamma}'$) of Theorem 4.5 instead of Conditions (α), (β) of Theorem 4.4, an issue that has already been addressed in Section 4.4. The prerequisites on the knowledge of the surface are P1, P3, and the following variant of P2:

$\hat{\mathbf{P}}$2 we know or we can compute a positive constant ε_S less than or equal to $\mathrm{lr}_k(S)$.

Compared to the smooth case, the implementations of the exhaustive and bipolar oracles remain the same. This fulfills P3. Moreover, when a closed formula of f is given, we can fulfill P1 by performing the same computations as in Section 6.2.1, using the generalized gradient of f instead of the usual gradient. When no closed formula is given, the solution depends on the context. In the case where S is a level-set in a 3D greyscaled image, the plane-sweep procedure described in Section 6.2.2 still works.

Concerning $\hat{\mathrm{P}}2$, we note that estimating $\mathrm{lr}_k(S)$ is a formidable task. Therefore, in practice we replace this calculation by a simple iterative process: we make a first guess of $\mathrm{lr}_k(S)$ and set up ε_S accordingly. If no initial guess can be made, we set ε_S to 1 by default. Then, we run Chew's algorithm

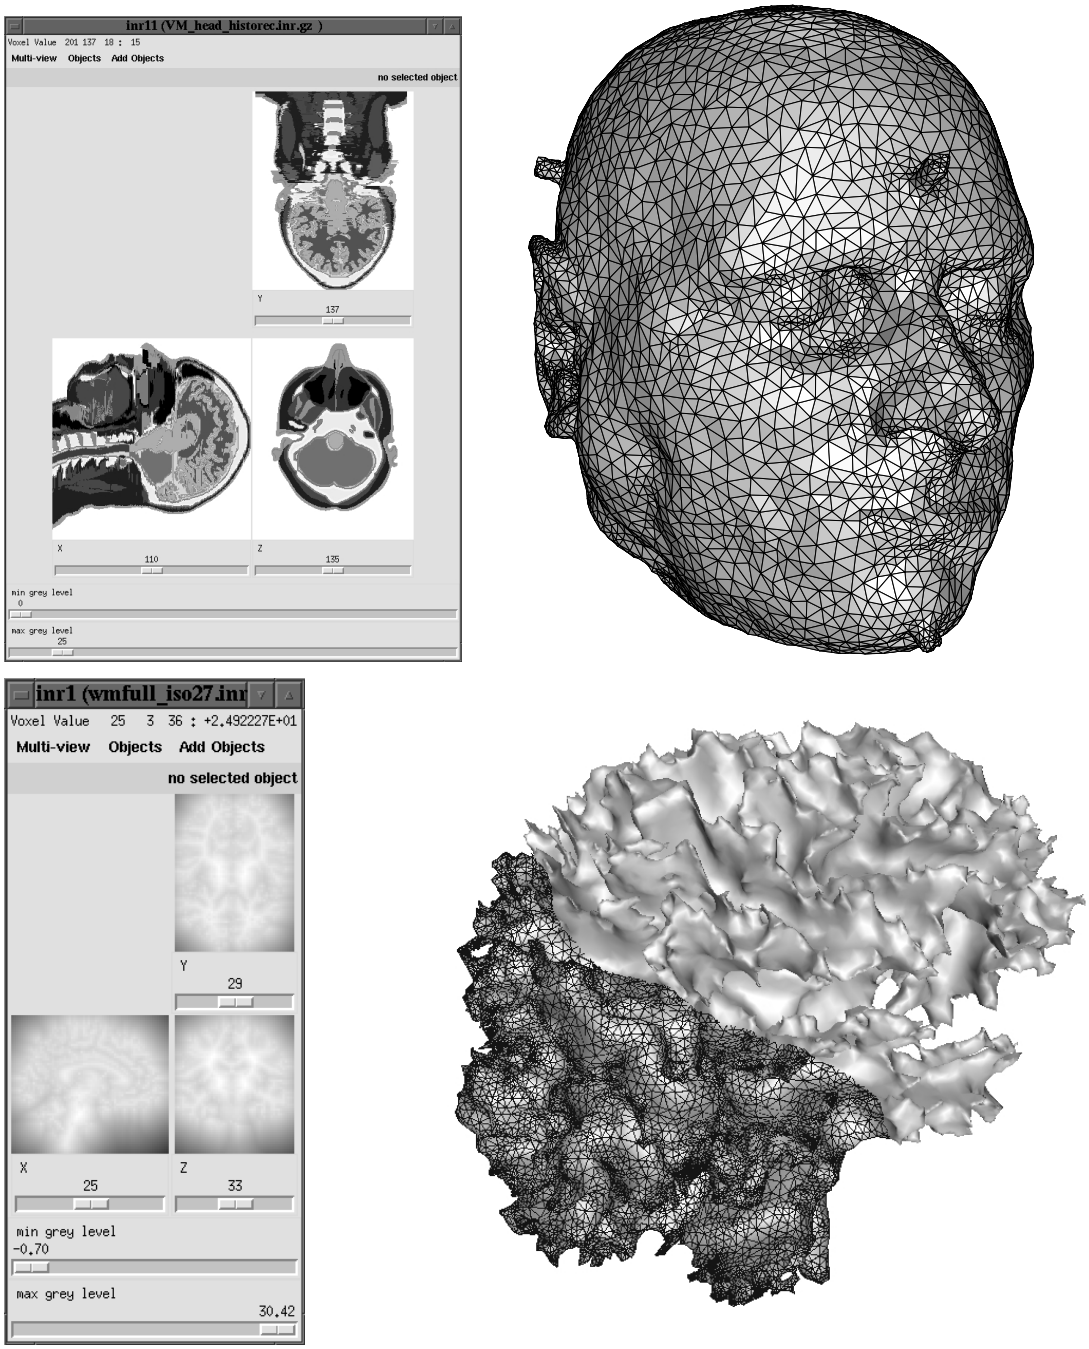

Figure 6.4: Results on $3D$ greyscaled images, with $\sigma = 0.09\ \mathrm{d}_{\tilde{\mathrm{M}}}$.

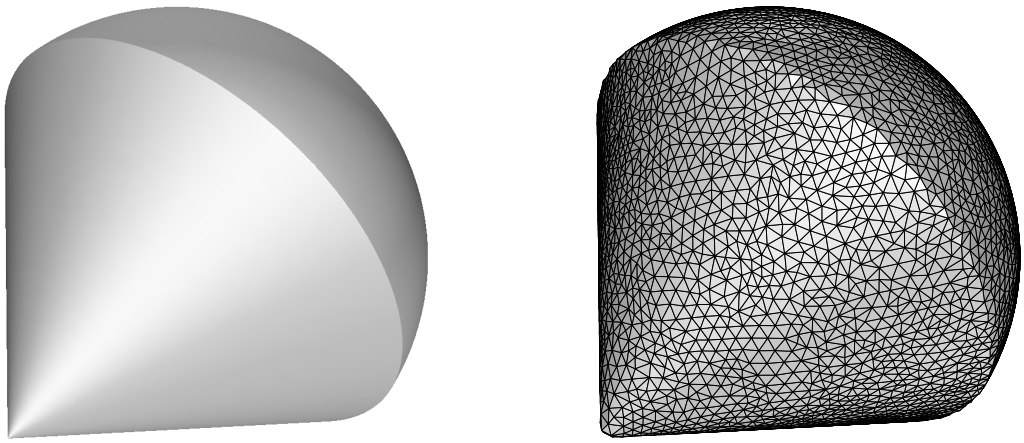

Figure 6.5: Result on a 1-Lipschitz surface, with $\varepsilon = 0.07\ \varepsilon_S$.

with $\varepsilon = 0.07\ \varepsilon_S$ (as prescribed by $(\hat{\alpha}')$ and $(\hat{\beta}')$) and look at the output. If the mesh is a manifold without boundary, isotopic to the original surface[3], then we stop. Otherwise, we divide ε_S by 2 and restart the process. Since S is Lipschitz, there exists a k such that $\mathrm{lr}_k(S)$ is positive. Therefore, after a finite number of iterations, ε_S is small enough to guarantee that the topology of the output mesh if correct, thanks to Theorem 4.5. In practice, only a few iterations are necessary after a good first guess.

Figure 6.5 presents the result of the algorithm on a semi-algebraic surface called "spinning top". This surface is 1-Lipschitz and its 1-Lipschitz radius is $\sqrt{2}/2$ (while its diameter is 2). The left column contains the input model and the right column the output mesh. On this example, we took $\varepsilon_S = \frac{1}{2}$ and $\varepsilon = 0.07\ \varepsilon_S = 0.035$. Although the Lipschitz constant ($k = 1$) is far beyond our theoretical bound, we obtained a good topological and geometric approximation of the original surface.

Figure 6.6 presents the result of the algorithm on a mechanical part. The CAD data is shown in the upper row, the output of the algorithm is shown in the bottom row. This surface is $\sqrt{2}$-Lipschitz, and its $\sqrt{2}$-Lipschitz radius is very small compared to its diameter, due to the presence of thin parts. Here again, we obtained a good topological and geometric result, although the Lipschitz constant ($k = \sqrt{2}$) is far beyond our theoretical bound.

Other results are presented in Figure 6.7. The left column contains the input models, the center column the output meshes, and the right column some close-ups of both objects. From top to bottom, we have:

- a degree-six algebraic surface called "heart", with two pinch points, one at the top and the other at the bottom. Because of these pinch points, the surface is not Lipschitz. However, in practice the result of the algorithm is topologically correct, provided that ε is small enough: here, we took $\varepsilon = 0.05$, while the diameter of the surface is 4.
- Barth's octic surface, of degree 8 (according to its name), made of eight pillows placed at the vertices of a cube and connected along the edges of the cube by means of two singular points. Because of these singular points, the surface is not an embedded 2-manifold. We took $\varepsilon = 0.03$, while the diameter of the surface is 4.

[3]Checking this last condition requires to know the topology of the original surface, which is the case in most applications.

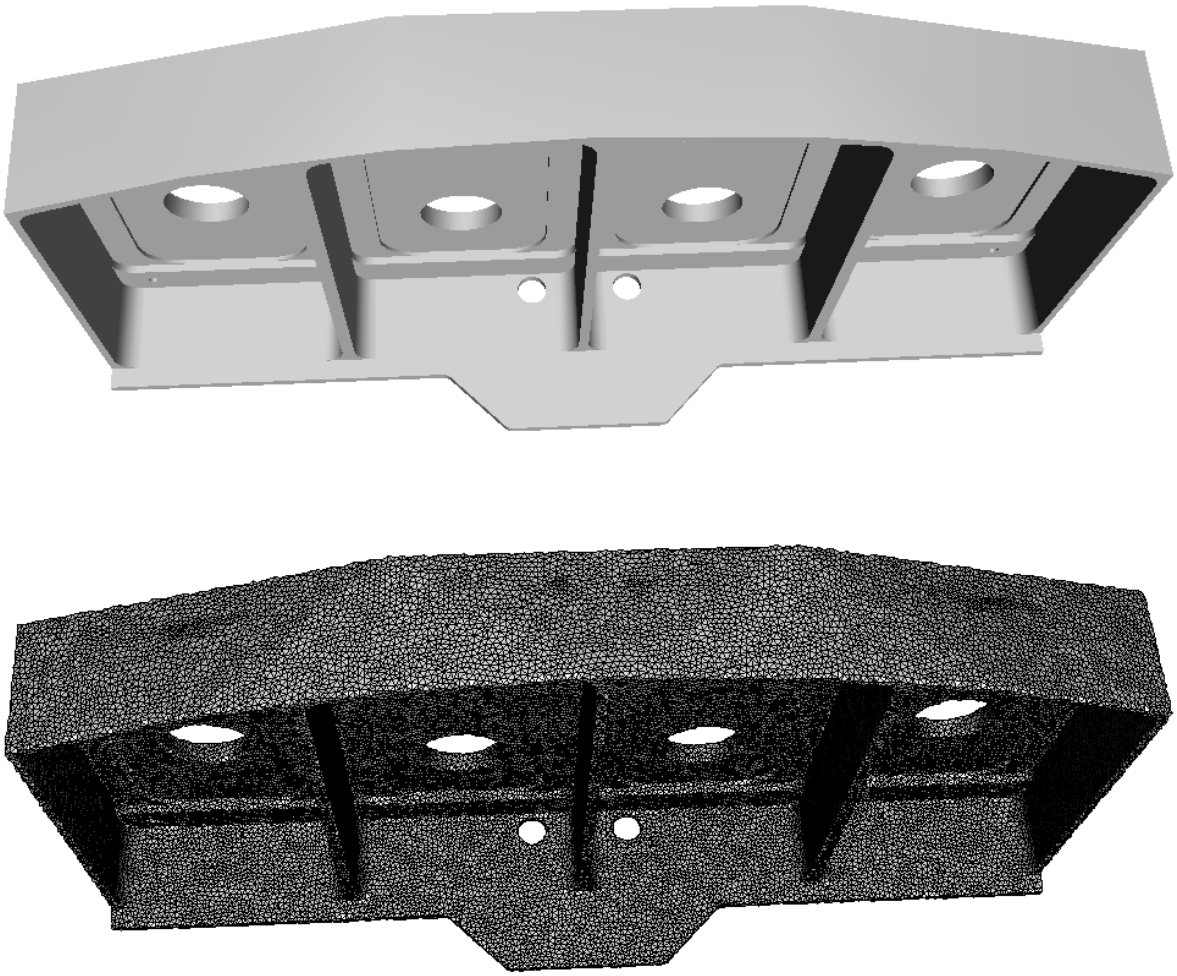

Figure 6.6: Result on a $\sqrt{2}$-Lipschitz CAD model.

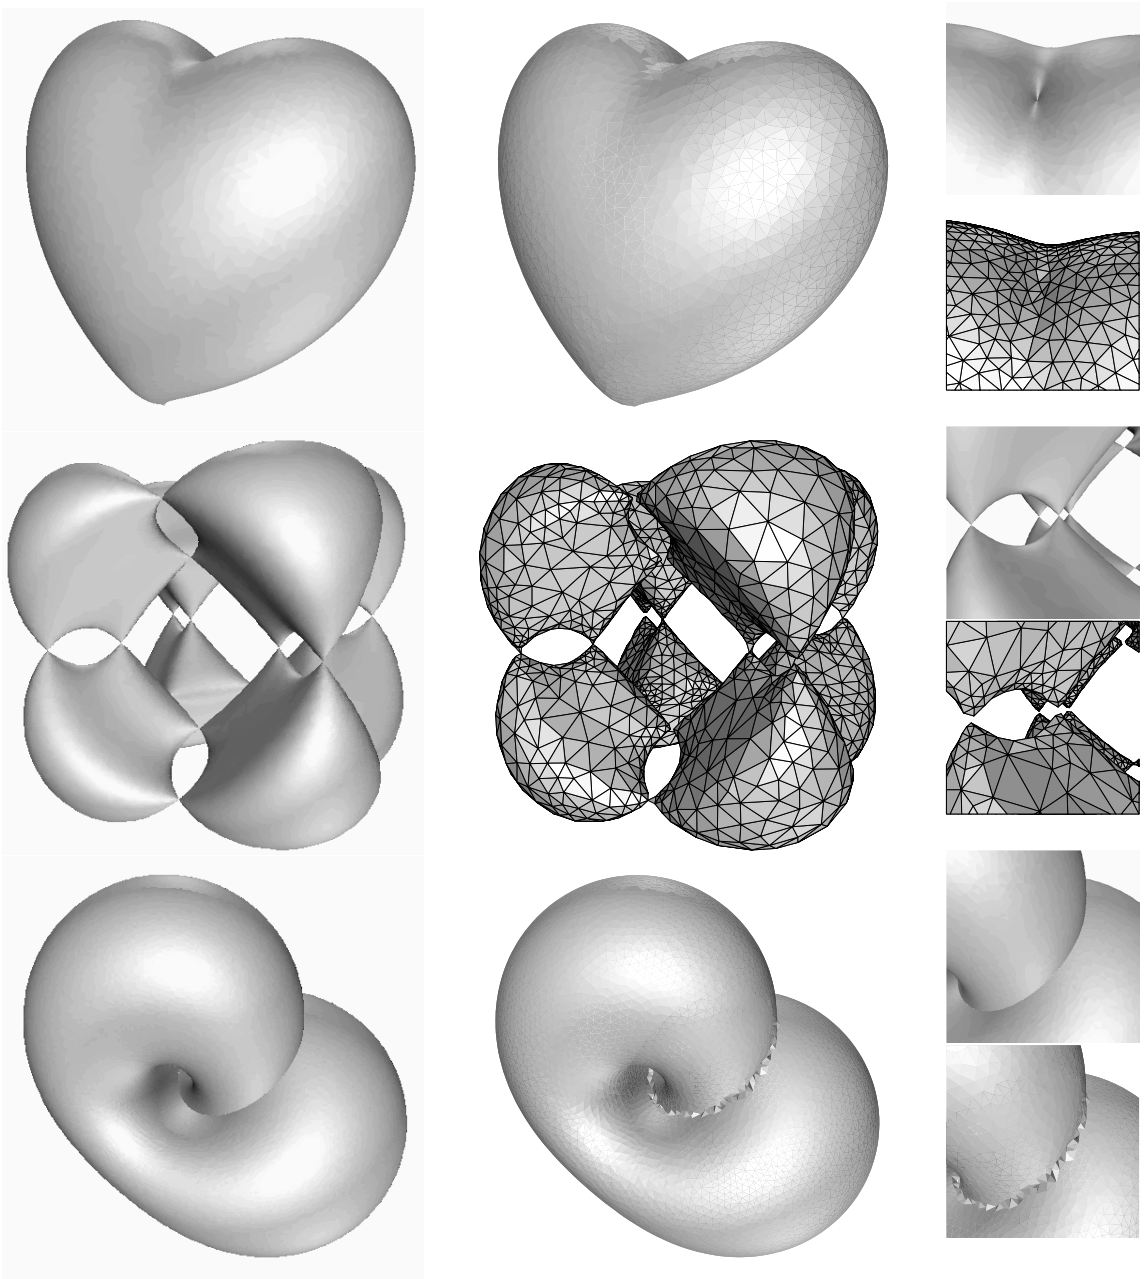

Figure 6.7: Results on non-Lipschitz or non-manifold algebraic surfaces.

– Klein's bottle, which cannot be embedded in $\mathbb{R}^3$ either. Here, we took $\varepsilon = 0.1$, while the diameter of the surface is 8.

Figure 6.7 illustrates two things. First, the algorithm terminates on all kinds of compact objects, as reported in Lemma 4.7. Second, the algorithm usually does a good job far from the non-Lipschitz or non-manifold parts of the surface, and sometimes also in their vicinity. We have added to our implementation a patch that checks whether $\mathrm{Del}_{|S}(E)$ is a manifold and that goes on refining the mesh in the negative. This patch can be toggled by the user, which allows him to force the algorithm to generate a manifold. This option has been activated for Barth's octic, and the output mesh is indeed manifold. However, the singular points of the original surface are missing in our triangulated version.

6.5 Application to surface reconstruction

Sampling and meshing implicit surfaces with guaranteed topology and geometry can be useful for reconstructing surfaces from point clouds. The issue is the following one: given an input point cloud E_I sampled from an unknown surface S, construct a PL surface that approximates S topologically and geometrically.

The literature on this topic is vast, and we refer the reader to [30, 31] for a comprehensive survey. Roughly speaking, previous work on surface reconstruction reduces to two main approaches:

- The first approach is combinatorial: it consists in constructing some specific data structure from E_I, such as its Delaunay triangulation, and then in extracting from this data structure a PL approximation of S. A number of provably good algorithms based on this approach have been proposed [6, 7, 8, 19, 53, 54, 55, 56].

- The second approach consists in using the input point cloud E_I to define an implicit function f', whose zero-set S' is then approximated using some implicit surface mesher. The first algorithm of this kind was designed by Hoppe *et al.* [78], who define f' as the signed distance function to a local estimate of the tangent plane of S, and mesh the zero-set of f' using the Marching Cubes algorithm. Other algorithms have been proposed ever since, based on natural neighbors [18, 19, 105], Moving Least Squares [58, 79, 83], Radial Basis Functions [100], or similar techniques.

The success of the second approach relies heavily on the choice of the implicit mesher. To this extend, our mesher is interesting since it comes with topological and geometric guarantees (even on nonsmooth Lipschitz surfaces), together with a tight bound on the number of points that will be inserted to achieve a given Hausdorff approximation. The remaining issues are the following ones:

1. Define a function f' whose zero-set S' is a smooth or Lipschitz surface without boundary. This issue is carefully addressed in the literature, especially for the Moving Least Squares surface, which is proved to be homeomorphic to S and close to S for the Hausdorff distance, under mild sampling assumptions.
2. Estimate $\mathrm{rch}(S')$ if S' is smooth, or $\mathrm{lr}_k(S')$ if S' is k-Lipschitz. To our knowledge, this remains an open theoretical question. However, as shown by our experiments, this is not a real issue in practice.

Figures 6.8 through 6.10 present some results of the reconstruction based on natural neighbor interpolation. Images are courtesy of F. Cazals and J.-D. Boissonnat [19]. In Figure 6.8, no input parameter

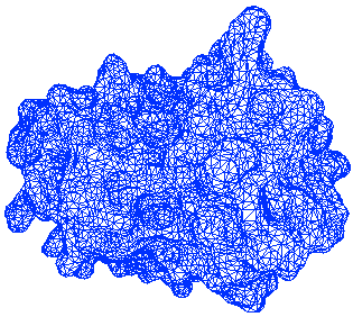

Figure 6.8: Result of natural neighbor interpolation on the papaine model.

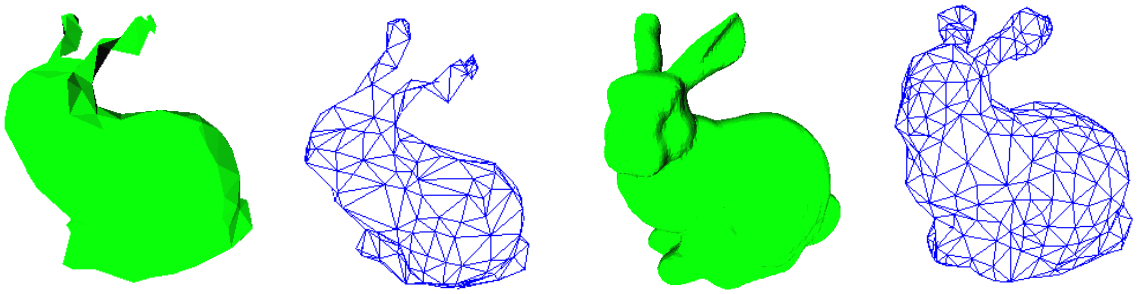

Figure 6.9: Results of natural neighbor interpolation on the Stanford bunny (250 pts).

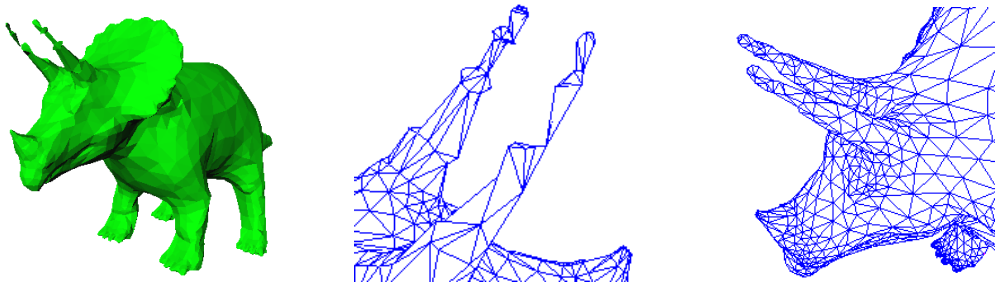

Figure 6.10: Results of natural neighbor interpolation on the triceratops model.

ε has been specified to our mesher, hence the output mesh is $\mathrm{Del}_{|S'}(E_I)$. Since S is smooth and E_I is a $\mu \mathrm{d}_{\mathrm{M}}$-sample of S, for some sufficiently small μ, the theoretical results of [19] ensure that $\mathrm{Del}_{|S'}(E_I)$ is a 2-manifold homeomorphic to S' and close to S' for the Hausdorff distance. Moreover, it is observed that S' is a good topological and geometric approximation of S.

In Figures 6.9 (left) and 6.10 (left), the input point sample E_I is not dense enough to ensure that $\mathrm{Del}_{|S'}(E_I)$ is a good approximation of S'. Nevertheless, S' is guaranteed to be a smooth surface without boundary. Therefore, when a sufficiently small parameter ε is specified, the outcome of our mesher is a 2-manifold that approximates S' correctly – see Figures 6.9 (right) and 6.10 (right). In the present case, the topology became correct after adding a few points to E_I. Moreover, it is observed that S' is not too far from S and has the same topology type.

Figure 6.11 presents the result of the reconstruction based on MLS approximation. The input point cloud E_I has $42K$ points and is shown in the upper-left corner. The output mesh is in the bottom-right corner, and the (smooth) MLS surface is in the bottom-left corner. For comparison, we have added in the upper-right corner the output of the reconstruction based on nearest neighbor interpolation. Both outputs have the right topology type. However, the interpolation strategy leads to a wrong approximation of the normals of the original surface[4], due to the presence of noise in the input point cloud E_I. This is not the case with the approximation strategy.

[4]Ondulations are clearly visible in the upper-right corner of Figure 6.11.

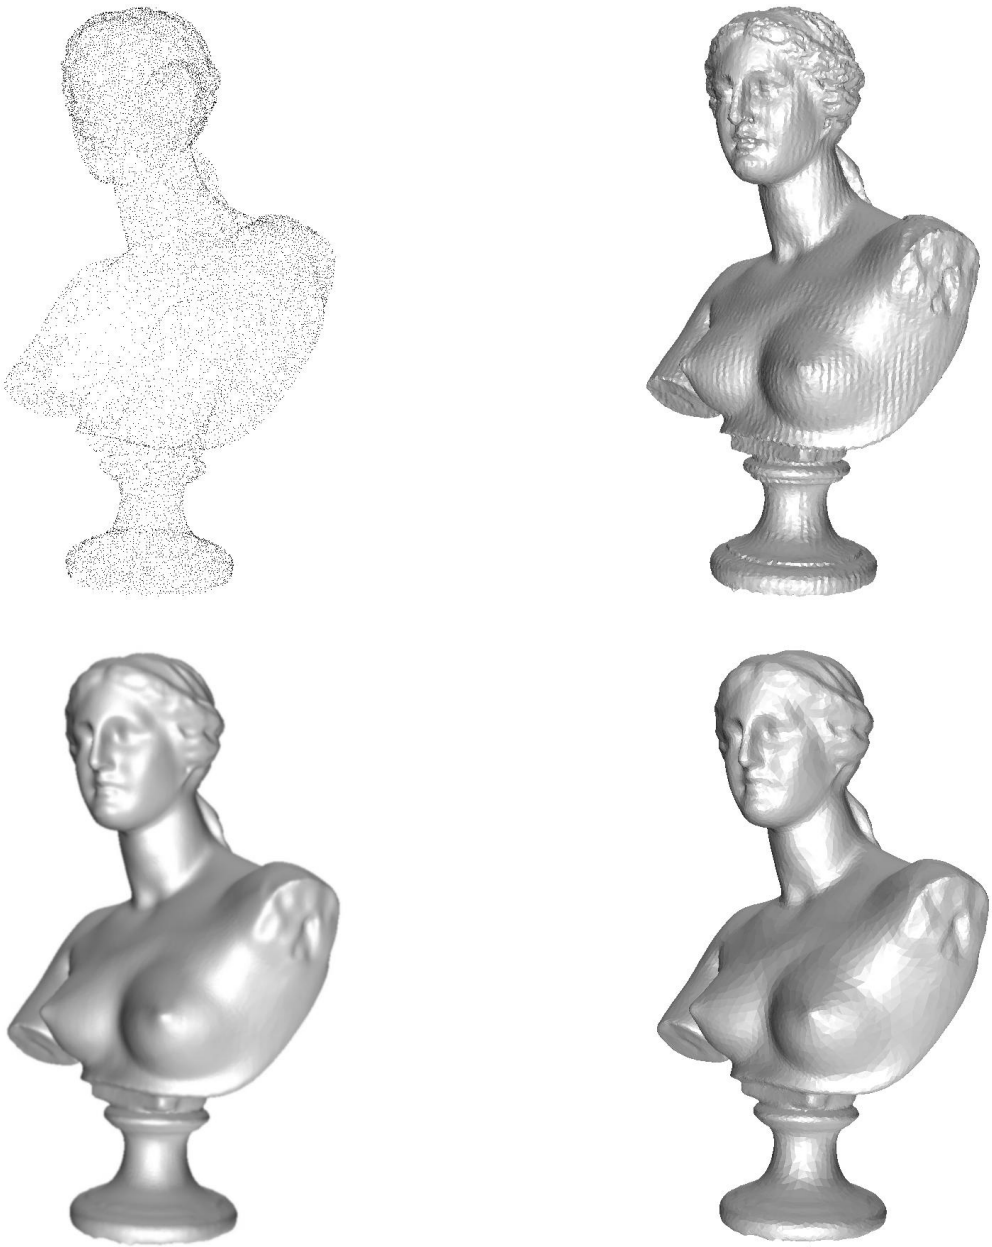

Figure 6.11: Comparison between natural neighbor interpolation and MLS approximation.

Chapter 7

Polygonal Surface Remeshing

7.1 Introduction

Polygonal meshes are used in a wide variety of applications of science and engineering. They can be the result of careful design using CAD software, or they may come as an output of a scanning device, associated with reconstruction or computer vision algorithms. Such meshes usually lack properties that are useful for subsequent numerical processing [95]. In particular, the application may need a grading of the mesh elements, with respect to some specific sizing field, together with a certified quality on their shape. In the case of a simplicial surface mesh, the quality is measured by the radius-edge ratios of the triangles, or equivalently by their smallest inner angle.

The problem of remeshing polygonal surfaces has received a lot of attention in the recent years, especially among the numerical analysis community, where the application needs are highest. Previous work on the topic can be divided into two categories: techniques that use a parameter space [51, 112] and techniques that perform iterative refinements on an explicit mesh [26, 27, 68].

The problem has also been extensively studied by the Computer Graphics community. We refer the reader to [4, 5] and the references therein for a complete survey. Although the methods proposed by the Graphics Community come with no theoretical guarantee concerning the topology of the output mesh or the quality of its elements, they are usually very effective in practice.

Recently, a provably good polyhedron remeshing algorithm has been proposed by Dey *et al.* [57]. Their work assumes that the vertices of the polyhedron have been sampled from a smooth surface and form a $\mu \mathrm{d}_{\mathrm{M}}$-sample of this surface, with μ sufficiently small. Their algorithm is an extension of previous work on smooth surface meshing [42], and the proof of its correctness is based on the same ideas.

7.1.1 Statement of the problem

The input surface S is an oriented polyhedron without boundary. We assume that the normals of any two non-disjoint facets of S form an angle less than $\frac{\pi}{2}$. In the context of volume meshing, a polyhedron verifying this condition is said to have no small angle [102].

Given two constants $\delta > 0$ and $\beta < \frac{\pi}{6}$, the goal is to construct a PL approximation $\hat{S}$ that is isotopic to S and at Hausdorff distance at most δ of S. Moreover, the facets of $\hat{S}$ are required to have inner angles greater than β. Unless otherwise specified, the density of the vertices of $\hat{S}$ is required to be uniform.

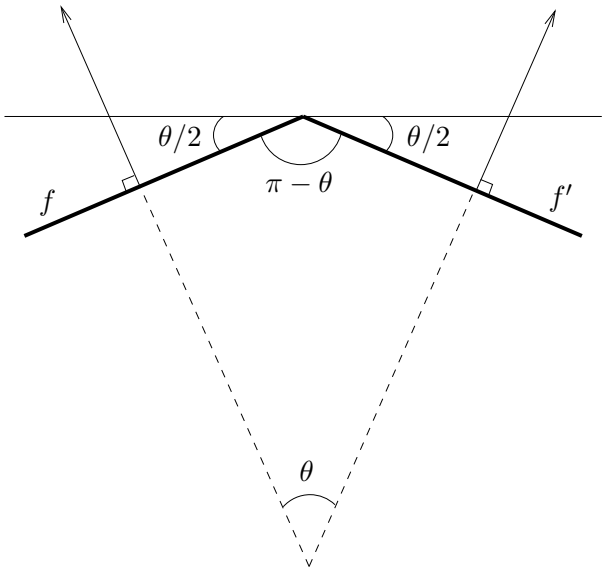

Figure 7.1: Lower bound on k'.

7.1.2 Our approach

Contrary to the implicit case, it is quite feasible to estimate k such that S is k-Lipschitz and to compute a positive lower bound ε_S on $\mathrm{lr}_k(S)$ when S is a polyhedron. This step is described in Section 7.2.2.

We can then follow the same strategy as for implicit surfaces: we use our implementation of Chew's algorithm, setting the input parameter ε to $\left(\frac{1-k^2}{1+k^2}\right)^2 \delta$. If $k \leq \tan\left(\frac{1}{2}\arctan\frac{1}{3\sqrt{3}}\right) \approx 0.1$ and $\delta \leq 0.07 \left(\frac{1+k^2}{1-k^2}\right)^2 \varepsilon_S$, then $\varepsilon \leq 0.07\, \varepsilon_S$ and, by Theorem 4.5, the output of the algorithm is a PL approximation $\hat{S}$ of S, isotopic to S and at Hausdorff distance at most $\frac{\varepsilon}{\cos^2(2\arctan k)} = \varepsilon\left(\frac{1+k^2}{1-k^2}\right)^2 = \delta$ of S. Moreover, since we use the construction of Section 4.4 to build the initial point set E_I, the facets of $\mathrm{Del}_{|S}(E_F)$ have inner angles of at least $\frac{\pi}{6} \geq \beta$.

The prerequisites on the knowledge of the surface are P1, $\hat{\text{P}}$2 and P3.

7.2 Satisfying the prerequisites

7.2.1 Computing one point per connected component of S

Consider the 1-skeleton graph G of S, where each node represents a vertex of S and two nodes are linked iff they are vertices of the same edge of S. Fulfilling P1 comes down to computing one point per connected component of G, which can be easily done by traversing G.

7.2.2 Estimating k and $\mathrm{lr}_k(S)$

For every couple (f, f') of non-disjoint facets of S, we compute the angle $(\mathbf{n}(f), \mathbf{n}(f'))$. Let θ be the largest angle computed. Since S is assumed to have no small angle, θ is less than $\frac{\pi}{2}$. By Theorem 2.9, we know that S is $\frac{2\sin\frac{\theta}{2}}{\sqrt{3-4\sin^2\frac{\theta}{2}}}$-Lipschitz, with $\frac{2\sin\frac{\theta}{2}}{\sqrt{3-4\sin^2\frac{\theta}{2}}} \leq \tan\frac{3\theta}{5}$ since $\theta \leq \frac{\pi}{2}$. Moreover, since at

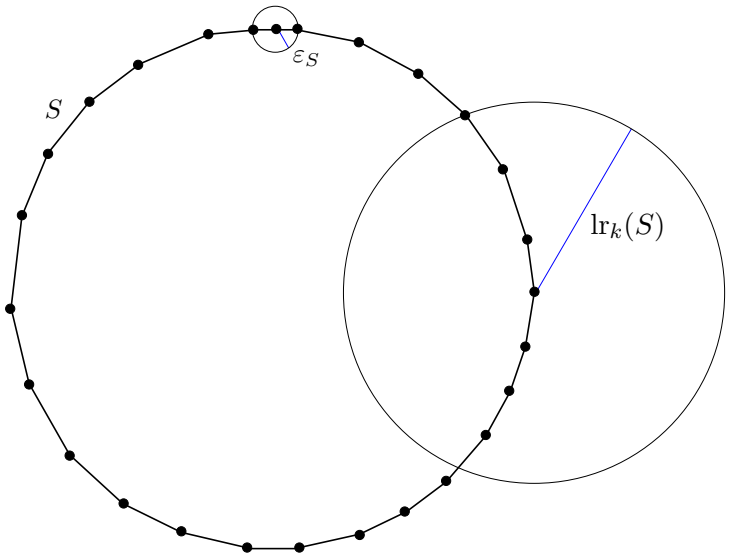

Figure 7.2: A case where ε_S is a bad estimate of $\mathrm{lr}_k(S)$.

least two non-disjoint facets f, f' of S have normals that make an angle of θ, k cannot be less than $\tan\frac{\theta}{2}$ – see Figure 7.1 for an illustration in the plane. Hence, we have $\tan\frac{\theta}{2} \leq k \leq \tan\frac{3\theta}{5}$.

We must also compute a lower bound ε_S of $\mathrm{lr}_k(S)$. Let f be a facet of S. The normals of the facets of $\mathrm{N}(f)$ make angles at most θ with $\mathbf{n}(f)$. Thus, by Lemma 2.8,

$$\forall p \in f,\ \mathrm{lr}_k(p) \geq \mathrm{d}(p, S \setminus \mathrm{N}(f)) \tag{7.1}$$

Computing the smallest distance from a point of f to $S \setminus \mathrm{N}(f)$ is easy: it suffices to compute $\mathrm{d}(f, f')$ for every facet $f' \in S \setminus \mathrm{N}(f)$ and to take the minimum. We do this for every facet f of S and we set ε_S to the minimum. By construction, ε_S is bounded by $\mathrm{d}(p, S \setminus \mathrm{N}(f))$, for any facet f of S and any point $p \in f$. Hence, $\varepsilon_S \leq \mathrm{lr}_k(S)$, by (7.1). Moreover, $\varepsilon_S > 0$ since the smallest distance from a point of a facet f to $S \setminus \mathrm{N}(f)$ is positive.

Notice that in some cases ε_S may be a bad estimate of $\mathrm{lr}_k(S)$, since it depends highly on the size of the facets. An illustration is given in Figure 7.2. Other strategies for computing ε_S could be considered as well, to make ε_S depend more on the discrete curvature of S and less on the size of the facets. However, such approaches usually involve more complex calculations. As an extreme, in the case where S is a PL approximation of a C^2-continuous surface W, one could use a curvature tensor estimator [32, 49] to make ε_S depend directly on W and as little as possible on S.

7.3 Implementing the oracle

The exhaustive oracle can be easily implemented by means of a naive procedure which, given a line segment s, checks the intersection of s with every facet of S. The segment-surface-intersection test is then performed in linear time with respect to the number of facets of S. Our implementation of the oracle uses some code developped by M. Samozino, from an original idea by P. Bhattacharya [15]. This code uses an octree to speed-up the segment-surface-intersection test.

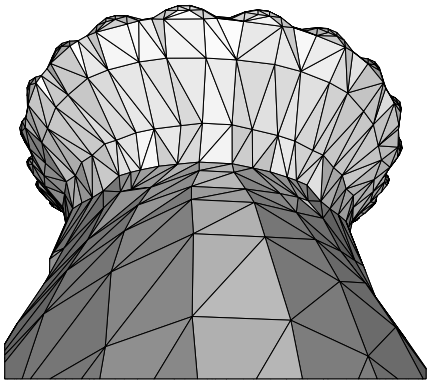

Input

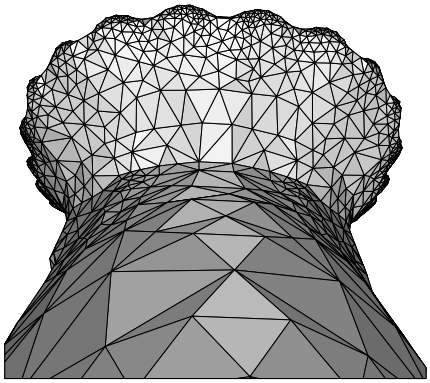

Output

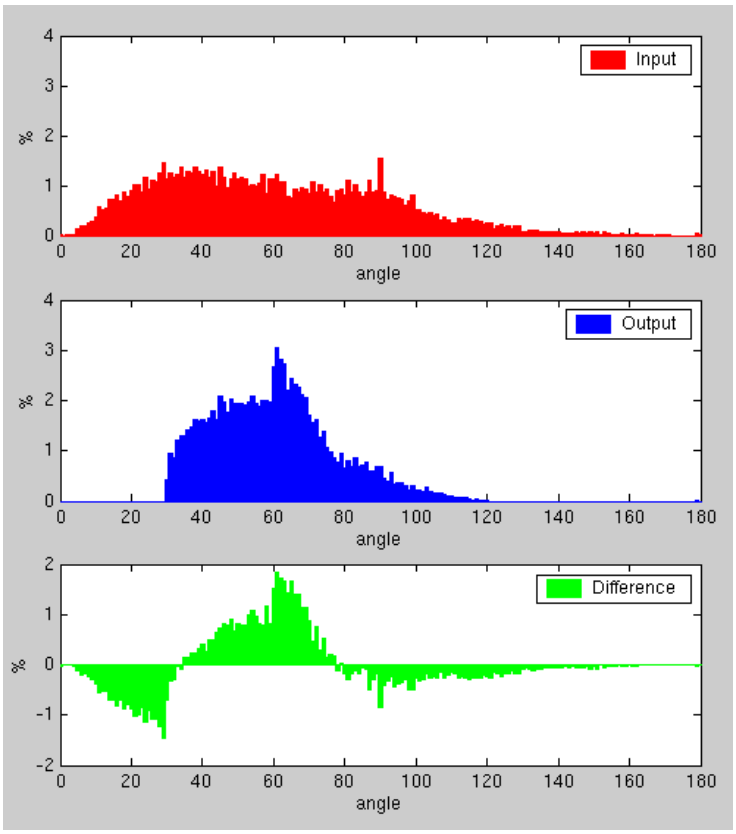

Figure 7.3: Removing the skinny facets.

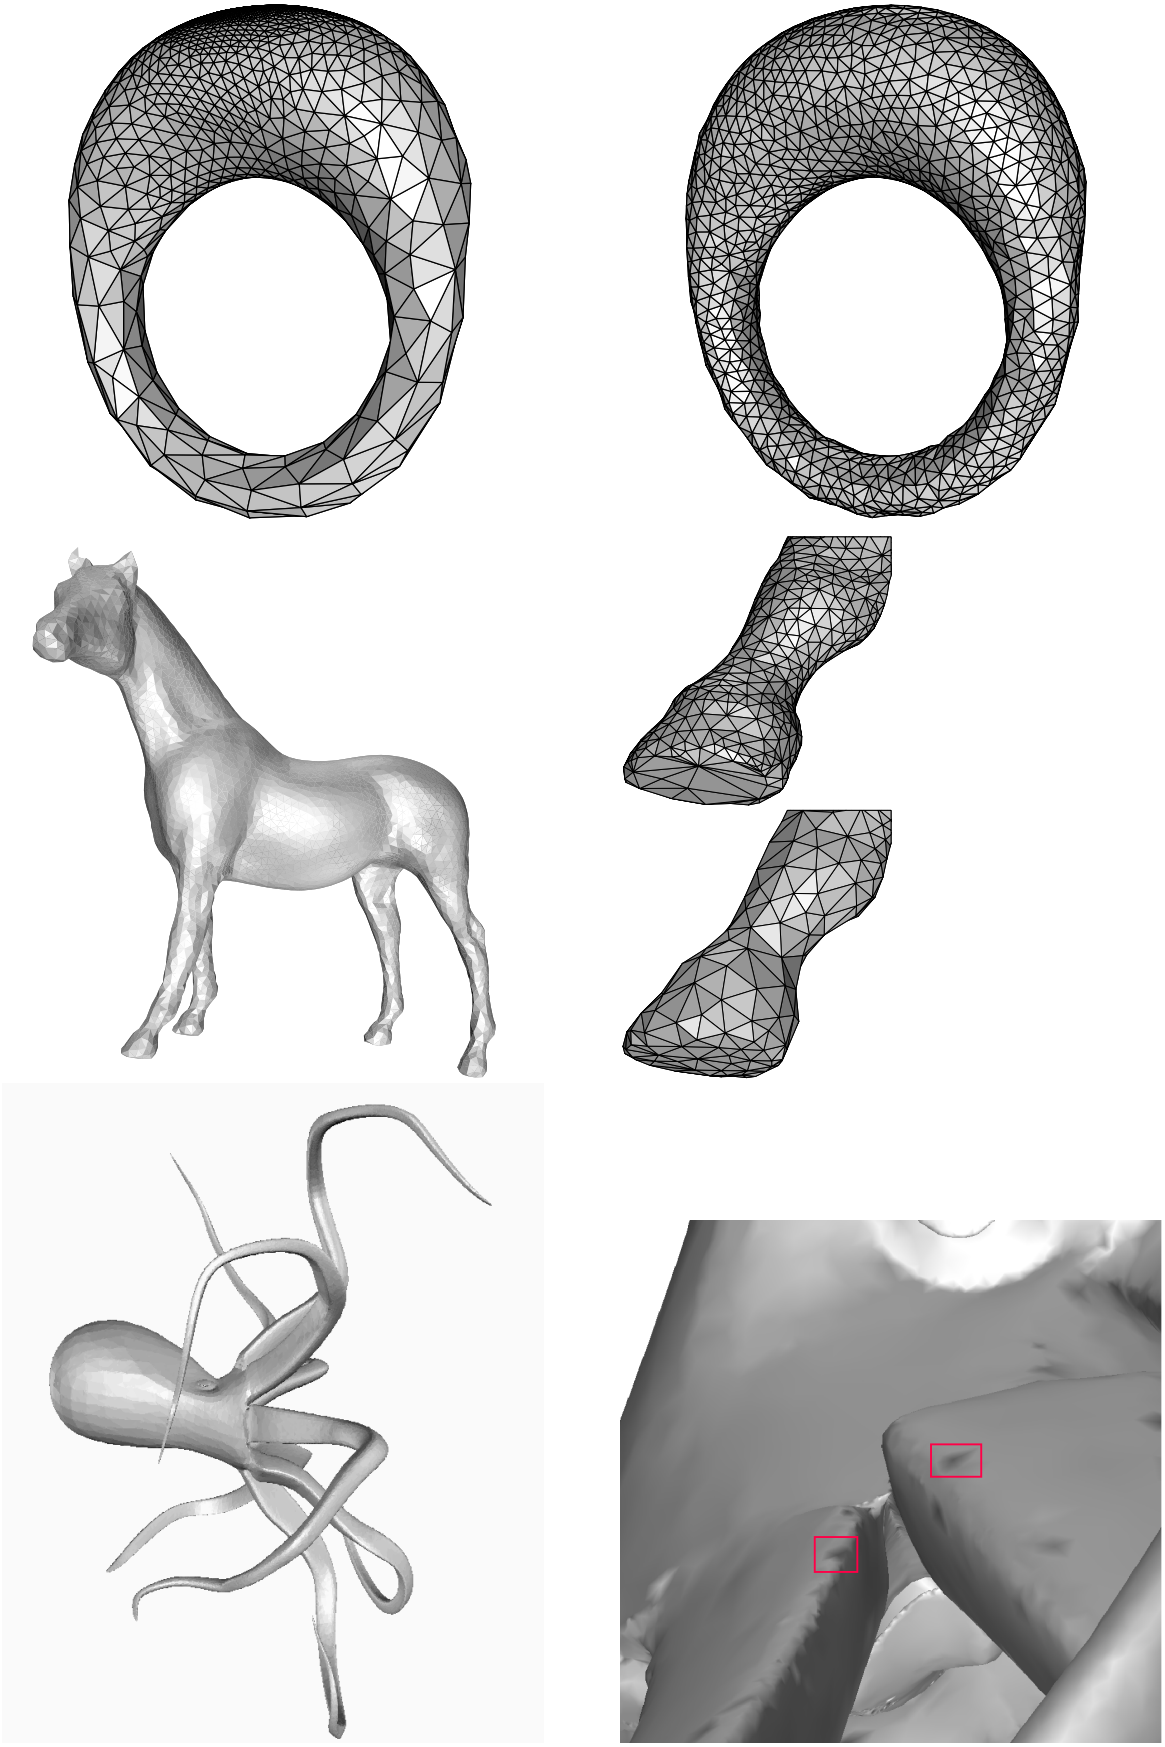

Figure 7.4: Results on various polyhedra.

7.4 Experimental results

Figure 7.3 illustrates the fact that the inner angles of the output mesh are guaranteed to be larger than $\frac{\pi}{3}$.

Figure 7.4 shows some results on various polyhedra. As predicted by the theory, when the polyhedron S is sufficiently *smooth*, the quality of the output of the algorithm is guaranteed. It is the case for the ring model (shown in the upper-left corner), from which the algorithm generated a manifold (in the upper-right corner) with the right topology type and close to the original surface for the Hausdorff distance.

The algorithm still works well in practice when the Lipschitz constant of the surface lies above the bound of Theorem 4.5 and below 1 (which corresponds to a dihedral angle of $\frac{\pi}{2}$). On the horse model for instance, the topology of the surface was well captured by the algorithm, whereas the minimum dihedral angle lies around $\frac{\pi}{2}$, as shown by the closeups of the hoofs[1]. This gives evidence that our theoretical bound on k is pessimistic.

When the input polyhedron has *sharp edges*, *i.e.* edges where the dihedral angle is smaller than $\frac{\pi}{2}$, the algorithm usually fails to capture its topology. It may not even produce a manifold. An illustration is given in the close-up of the octopus, in which the mesh is shown with Gouraud shading. The consequence of the non-orientability of the mesh is that dark stains appear in places where the normal is oriented in the wrong direction. In such places, the surface is not locally a manifold.

[1] The original model is above, the output of the algorithm is below.

Chapter 8

Surface Probing

8.1 Introduction

In this chapter we consider the problem of discovering the shape of an unknown object $\mathcal{O}$ of $\mathbb{R}^3$ through an adaptive process of probing its surface from the exterior. A probe is issued along a ray whose origin lies outside $\mathcal{O}$ and returns the first point of $\mathcal{O}$ hit by the ray. Successive probes may require the probing device to be moved through the free space outside $\mathcal{O}$. The goal is to find a strategy for the sequence of probes that guarantees a precise approximation of $\mathcal{O}$ after a minimal number of probes. Note that this problem involves an interesting bootstrapping issue, as the underlying surface is only known to the probing algorithm through the samples already taken. Thus, differently from most existing work in surface reconstruction, the data are not given all at once prior to the reconstruction phase but must instead be computed iteratively, each new probe depending on the outcomes of the previous probes. Furthermore, collision avoidance between the probing device and $\mathcal{O}$ must be observed at all times.

8.1.1 Previous work

The above problem belongs to the class of geometric probing problems, pioneered by Cole and Yap [50]. Geometric probing, also known as blind approximation or interactive reconstruction, is motivated by applications in robotics. In this context, our probe model described above is called a tactile or finger probe. Geometric probing finds applications in other areas and gave rise to several variants. In particular, other probe models have been studied in the literature, *e.g.* line probes (a line moving perpendicular to a direction), X-ray probes (measuring the length of intersection between a line and the object), as well as their counterparts in higher dimensions.

We classify the probing algorithms into two main categories, exact or approximate, depending on whether they return the exact shape of the probed object or an approximation. An exact probing algorithm can only be applied to shapes that can be described by a finite number of parameters like polygons and polyhedra. In fact, most of the work on exact geometric probing is for convex polygons and polyhedra. See [106] for a survey of the computational literature on the subject. Although it has been shown that, using enhanced finger probes, a large class of non convex polyhedra can be exactly determined [2, 24], exact probing is too restrictive for most practical applications.

Approximate probing algorithms overcome this deficiency by considering the accuracy of the desired

reconstruction as a parameter. The goal is to find a strategy that can discover a guaranteed approximation of the object using a minimal number of probes. The general problem is ill-posed, since we cannot conclude anything about the shape of the object if we have only local information about the shape. Some global information or prior knowledge is required to restrict the class of shapes being approximated. An important class is the class of convex shapes. Probing strategies have been proposed for planar convex objects using line probes [86, 97] and some other probe models are analyzed by Rote [98]. Observe that approximating a convex object using hyperplane probes is nothing else than approximating its supporting function.

As far as we know, probing non convex (non polyhedral) objects has not been studied. The problem has some similarity with surface approximation. In particular, the Marching Cubes algorithm [87] and Chew's algorithm provide blind approximations of a surface since the surface needs to be known only through an oracle that typically decides whether a line segment intersects the surface or not. However, the probing problem differs from surface approximation in an essential way: we cannot place the probing device at will anywhere but need to plan the motion of the probing device to its next probing location. Differently from the convex case, we cannot simply probe from infinity and need to determine finite positions outside the object where to place the probing device. Moreover, in order to reach such positions, we need to determine paths along which the probing device can be safely moved without colliding with the object.

8.1.2 Statement of the problem

Let $\mathcal{O}$ be a bounded open set of $\mathbb{R}^3$ and S its boundary. The goal is to approximate S by a probing tool that can locate points on S. The following assumption allows us to localize $\mathcal{O}$ within $\mathbb{R}^3$, preventing indefinite searches.

A1 *For every connected component $\mathcal{O}_i$ of $\mathcal{O}$, we know a point o_i that belongs to $\mathcal{O}_i$.*

For simplicity, we assume in this chapter that the surface S is connected. A1 tells then that we know a point $o \in \mathcal{O}$. The case of a surface with more than one connected component is analyzed in [21]. Assumption A2 bounds the area of interest and allows us to obtain initial locations and paths for the probing device without bumping into $\mathcal{O}$.

A2 *We know a convex and compact subset Ω of $\mathbb{R}^3$ that contains S (and hence also $\mathcal{O}$). We denote by $\partial\Omega$ the boundary of Ω.*

We have at our disposal a *probing device*, which is an oracle that, once placed at some point p of $\mathbb{R}^3 \setminus \mathcal{O}$, can be oriented towards any direction $\mathbf{v}$ and then tasked to return the first point of transverse intersection between S and the ray defined by $(p, \mathbf{v})$. The probing device can move freely in $\mathbb{R}^3 \setminus \mathcal{O}$ but cannot penetrate $\mathcal{O}$. Such a device can be constructed in practice, using for instance a laser with three DOFs of displacement and two DOFs of rotation, that can cast a ray in any direction and measure its distance to the point where the ray hits the object.

We assume that the probing device provides *exact information*. The outcome of a probe is a point on the boundary of the object.

We need also to define the *accuracy measure* for our reconstruction. The accuracy will be measured by the Hausdorff distance. Since the measured points are on the boundary S of the object, the accuracy of the reconstruction will be ε iff any point of S is at distance at most ε from a measured point. In such a case, the set of measured points is an ε-sample of S.

As mentioned above, to be able to make any reconstruction claims, we need to restrict the class of shapes we probe. We consider here those with positive reach or positive k-Lipschitz radius.

A3 *We know a positive constant ε_S less than $\mathrm{rch}(S)$ if S has a positive reach, and less than $\mathrm{lr}_k(S)$ if S is k-Lipschitz.*

Finally, we need a *model of computation* to analyze the complexity of our algorithm. Following the perception-action-cognition paradigm, we distinguish between the information or probing cost, the displacement cost, and the combinatorial cost. This distinction is also reminiscent of the difference made between combinatorial and informational complexity in the work on information-based computation [109, 110]. The probing cost measures the number of probes and indicates the amount of information that becomes available to our algorithm. The displacement cost accounts for the motion of the probing device. The combinatorial cost measures the arithmetic operations and comparisons required, as well as the maintenance cost of the data structures. As discussed later, it is not possible in general to optimize all costs simultaneously.

8.1.3 Overview of the chapter

Under assumptions A1-A3, we show that S can be approximated by a triangulated surface $\hat{S}$ within any desired accuracy. Moreover, $\hat{S}$ recovers the exact topology of S and the error on the normal deviation of the facets of $\hat{S}$ is also bounded.

The chapter is organized as follows. In Section 8.2 we describe the probing algorithm, present its main properties in Section 8.3, and analyze its complexity in Section 8.4.

8.2 The probing algorithm

If we except the moves of the probing device, our algorithm is quite similar to Chew's algorithm. The main difference concerns the oracle that is used to discover the surface S. In our case, to check whether a Voronoi edge e intersects S or not, we must first move our probing device to one of its endpoints. This requires two things: first, that at least one endpoint v of e be located in $\mathbb{R}^3 \setminus \mathcal{O}$; second, that we know a *free path* from $\mathbb{R}^3 \setminus \Omega$ (where the probing device can move freely) to v, *i.e.* a continuous curve included in $\mathbb{R}^3 \setminus \mathcal{O}$ that goes from $\mathbb{R}^3 \setminus \Omega$ to v.

Recall that, given a point set E, the *Voronoi graph of E*, or $\mathrm{VG}(E)$ for short, is the 1-skeleton graph of the Voronoi diagram $\mathrm{Vor}(E)$. Our basic intuition is to constrain the probing device to move along the

edges of $\mathrm{VG}(E) \setminus \mathcal{O}$, which are called the *free edges*[1]. A difficulty arises from the fact that, when a new point p is inserted in E, some of the current Voronoi vertices and edges may disappear. It follows that portions of $\mathrm{VG}(E) \setminus \mathcal{O}$ that could be reached by the probing device from $\mathbb{R}^3 \setminus \Omega$ before the insertion of p may no longer be reachable afterwards. See Figure 8.1 for an illustration.

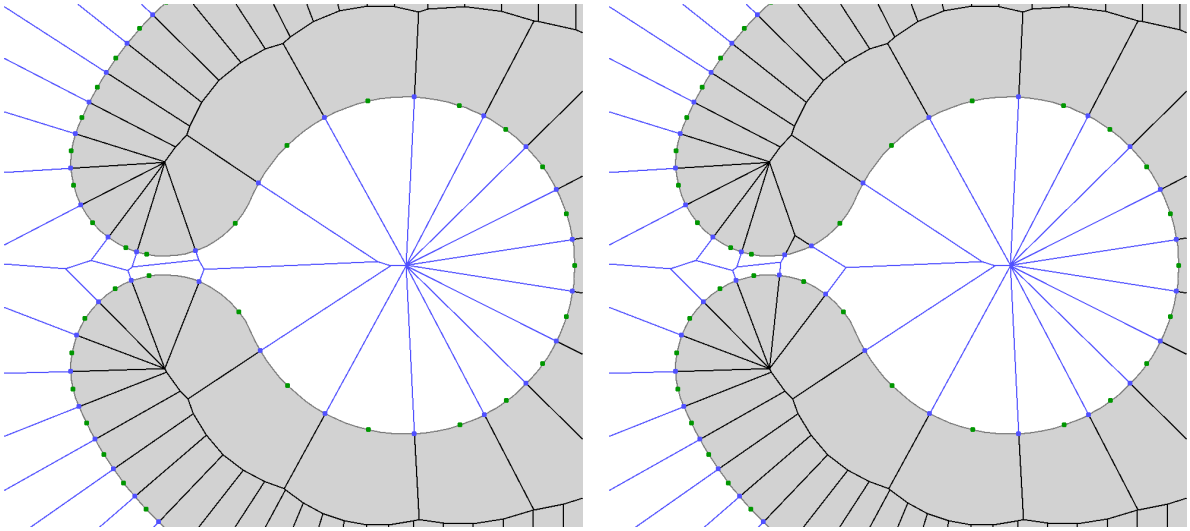

Figure 8.1: The insertion of a point splits $\mathrm{VG}(E) \setminus \mathcal{O}$ into two connected components, one of which is then no longer reachable from $\mathbb{R}^3 \setminus \Omega$.

To overcome this difficulty, once a free path $\pi(v)$ from $\mathbb{R}^3 \setminus \Omega$ to some Voronoi vertex v has been found, we store $\pi(v)$ in memory so that v will remain reachable by the probing device permanently. Hence our paths are made of two types of edges: edges that belong to the current Voronoi graph, and edges that do not but were edges in some former Voronoi diagram.

By moving the probing device along our free paths, and by probing from each visited Voronoi vertex towards its neighbor vertices in $\mathrm{Vor}(E)$, we can detect a subset of the points of $\mathrm{VG}(E) \cap S$ and construct a subcomplex of $\mathrm{Del}_{|S}(E)$ called the *visible restricted Delaunay triangulation of* E, or simply $\mathrm{Del}^v_{|S}(E)$. Every point of S that has been probed is the center of a Delaunay ball, called *ball of* $\mathrm{Del}^v_{|S}(E)$, that circumscribes a facet of $\mathrm{Del}^v_{|S}(E)$.

8.2.1 Data structure

We proceed as in Chew's algorithm, by storing $\mathrm{Del}^v_{|S}(E)$ as a subcomplex of $\mathrm{Del}(E)$. Inside every Delaunay tetrahedron, we mark each of the four facets as being or not being part of $\mathrm{Del}^v_{|S}(E)$. This way, every Delaunay facet is marked twice since it belongs to two Delaunay tetrahedra.

In order to store the paths for the probing device, every Voronoi vertex[2] v is given a pointer *prev* to the previous vertex on a path from $\mathbb{R}^3 \setminus \Omega$ to v. By convention, $v.prev = NULL$ means that we know no free path from $\mathbb{R}^3 \setminus \Omega$ to v. In such a case, v is said to be *inactive*. Otherwise, v is called *active*.

[1] More generally, any object (point, segment, curve etc.) that lies outside $\mathcal{O}$ is said to be *free*.

[2] In practice, it is its dual Delaunay tetrahedron that we consider. However, for simplicity, we will identify Delaunay tetrahedra with Voronoi vertices in the sequel.

If a newly created Voronoi vertex v belongs to $\mathbb{R}^3 \setminus \Omega$, then we set $v.prev \leftarrow v$ since v can be reached by the probing device. In particular, an infinite Voronoi vertex (*i.e.* the endpoint at infinity of an unbounded Voronoi edge) always lies outside Ω, which is compact. Thus, the *prev* field of an infinite vertex is never *NULL*. If v belongs to Ω, then we initialize $v.prev \leftarrow NULL$.

To construct and then update $\mathrm{Del}^v_{|S}(E)$, we use a routine named DETECT_ACCESS, introduced in Figure 8.2. Starting from an active vertex v_{start}, DETECT_ACCESS performs a depth-first traversal of $\mathrm{VG}(E) \setminus \mathcal{O}$ to see which previously inactive vertices can be reached by the probing device from v_{start} through free edges of the Voronoi graph.

```
DETECT_ACCESS (v_start):
 // Precondition: v_start is active
foreach neighbor v of v_start do
    PROBE edge [v_start, v];
    if ([v_start, v] ∩ S ≠ ∅) then
        add the dual of [v_start, v] to Del^v_|S(E);
    else if (v.prev = NULL) then
     // v becomes active
        if (v ∈ Ω) then
            v.prev ← v_start;
        end if
        MOVE the probing device from v_start to v;
        DETECT_ACCESS (v);
        MOVE the probing device from v to v_start;
    end if
end foreach
```

Figure 8.2: Routine DETECT_ACCESS

Initial construction Given an initial point set $E \subset S$, we compute $\mathrm{Del}^v_{|S}(E)$ by moving the probing device successively to all the vertices of $\mathrm{VG}(E)$ that lie outside Ω (including the infinite vertices[3]). For every such vertex v, we set $v.prev \leftarrow v$ and then we call DETECT_ACCESS on v.

After the initialization phase, every Voronoi vertex that can be reached from $\mathbb{R}^3 \setminus \Omega$ by walking along edges of $\mathrm{VG}(E) \setminus \mathcal{O}$ is active. Moreover, every active vertex is given a free path to $\mathbb{R}^3 \setminus \Omega$.

Update Each time a new point p is to be inserted in E, we update $\mathrm{Del}^v_{|S}(E)$ as follows:
• before the insertion, we look at the active vertices of $\mathrm{Vor}(E)$ that no longer exist in $\mathrm{Vor}(E \cup \{p\})$. By definition, they lie in $\mathrm{V}(p)$, the cell of p in $\mathrm{Vor}(E \cup \{p\})$. We keep these vertices in memory and we

[3]Processing the infinite vertices in the same manner as the other ones simplifies the presentation but is not quite satisfactory since it involves moving the probing device to infinity. However, this can be avoided easily by clipping $\mathrm{VG}(E)$ by $\partial\Omega$ and calling DETECT_ACCESS on all the intersection points of $\partial\Omega \cap \mathrm{VG}(E)$.

leave their *prev* pointers unchanged. This way, every active vertex will remain active in the sequel and will keep its path to $\mathbb{R}^3 \setminus \Omega$.

• after the insertion, we look at the new vertices of the Voronoi diagram (including the infinite ones), which by definition are the vertices of $\mathrm{V}(p)$. For any such vertex v, we need to determine whether v can be reached from $\mathbb{R}^3 \setminus \Omega$ through edges of $\mathrm{VG}(E) \setminus \mathcal{O}$:

- if $v \in \mathbb{R}^3 \setminus \Omega$, we set $v.prev \leftarrow v$ and move the probing device to v. Such a move is called a *positioning displacement*. Then, we call DETECT_ACCESS on v.
- otherwise, we look at the only neighbor v' of v that is not a vertex of $\mathrm{V}(p)$. If v' is active and if edge $[v, v']$ is free (which we can easily determine since $[v, v']$ is included in a former Voronoi edge that has been probed from v'), we perform a *positioning displacement* by moving the probing device to v'. Then, we call DETECT_ACCESS on v'.

8.2.2 The algorithm

The algorithm takes as input a user-defined value ε such that $0 < \varepsilon < 0.07\ \varepsilon_S$, which by A3 is less than $0.07\ \mathrm{rch}(S)$ if S has a positive reach and less than $0.07\ \mathrm{lr}_k(S)$ if S is k-Lipschitz. As will be shown in Section 8.3.3, controlling ε allows to bound the Hausdorff distance between S and the PL approximation built by the algorithm.

The algorithm starts by computing an initial point set E_I made of the three vertices of a λ-persistent facet, with $\lambda = 0.14\ \varepsilon_S$. To do so, it cannot use the method of Section 4.4 directly, since that method requires to compute the critical points of some height function. Instead, the algorithm places the probing device at a point p of $\partial\Omega$ and probes from p towards point o. Since $o \in \mathcal{O}$ and $p \in \mathbb{R}^3 \setminus \mathcal{O}$, the probing device finds a point $u \in S$, such that the segment $[p, u]$ is free. The algorithm then chooses two other directions, very close to direction $[p, u)$, so that the probing device will find two other points of S, namely v and w, such that the triangle (u, v, w) is circumscribed by a sphere centered on S of radius at most $\lambda/2$. According to Lemmas[4] 1.5 and 2.13, the normal of S in the vicinity of u is well approximated by the unit vector $\mathbf{n}_{uvw}$ orthogonal to $\mathrm{aff}(u, v, w)$ that has a non-negative inner product with $(\mathbf{p} - \mathbf{u})$. Then, the same construction as in Section 4.4 can be performed using $\mathbf{n}_{uvw}$ instead of $\mathbf{n}$, with the same theoretical guarantees.

After the construction of the initial point set E_I, the algorithm works as Chew's algorithm (see Section 4.1.3), using the probing device to answer the oracle. Specifically, the data structure is $\mathrm{Del}^v_{|S}(E)$, and the balls of $\mathrm{Del}^v_{|S}(E)$ whose radii are greater than ε are stored in a *priority queue* $\mathcal{Q}$. These balls are called the *bad* balls of $\mathrm{Del}^v_{|S}(E)$. While $\mathcal{Q}$ is not empty, the algorithm retrieves from $\mathcal{Q}$ a bad ball $B(c, r)$ of largest radius and inserts its center c in E. The algorithm then updates $\mathrm{Del}^v_{|S}(E)$ as described in Section 8.2.1, and updates $\mathcal{Q}$ as follows:

- the former bad balls that disappear because of the insertion of c are removed from $\mathcal{Q}$;
- the new bad balls that are created by the insertion of c are inserted in $\mathcal{Q}$.

The algorithm stops when $\mathcal{Q}$ is empty, that is, when no ball of $\mathrm{Del}^v_{|S}(E)$ is bad. The point set E is then renamed E_F and returned as well as $\mathrm{Del}^v_{|S}(E_F)$.

[4] In the nonsmooth case, the triangle (u, v, w) must not have too small inner angles, which can easily be ensured.

8.3 Correctness of the algorithm and quality of the approximation

In this section, we analyze the probing algorithm. We prove that it terminates in Section 8.3.1. In Section 8.3.2, we exhibit two invariants that are instrumental in proving the geometric properties of the output surface in Section 8.3.3. The analysis of the complexity of the algorithm is deferred to Section 8.4.

8.3.1 Termination

Since $\mathrm{Del}^v_{|S}(E)$ is a subcomplex of $\mathrm{Del}_{|S}(E)$, the proof of Lemma 4.7 holds when the algorithm is run with the probing device instead of the exhaustive oracle. Hence, the algorithm inserts finitely many points. This implies that the number of Voronoi edges created is finite. Since a Voronoi edge is probed only from its vertices, it is probed at most twice. Hence, the probing device performs finitely many probes. As a consequence, the algorithm terminates.

8.3.2 Invariants of the algorithm

Proposition 8.1 *The following assertions hold throughout the course of the algorithm:*
P1 *All active Voronoi vertices can be reached from $\mathbb{R}^3 \setminus \Omega$ by moving the probing device along current or former Voronoi edges.*
P2 *Any two Voronoi vertices that lie in the same connected component of $\mathrm{VG}(E) \setminus \mathcal{O}$ have the same status,* active *or* inactive.

Proof. We proceed by induction. Clearly, (P1) and (P2) are verified after the initialization phase. Let us now consider a step of the algorithm during which a new point (say p) is inserted in E and $\mathrm{Del}^v_{|S}(E)$ is updated. Our induction hypothesis is the following:

IH *Assertions (P1) and (P2) hold in set E before the insertion of p.*

We will prove successively that (P1) and (P2) are still verified after the insertion of p. In the sequel, E denotes the point sample before the insertion of p.

(P1) Let v be a vertex that is active after the insertion of p.
P1.1 If v existed and was already active before the insertion of p, then its path $\pi(v)$ to $\mathbb{R}^3 \setminus \Omega$ remains unchanged since all the vertices of $\pi(v)$ are kept in memory and DETECT_ACCESS does not change the status of active vertices. It follows that v is reachable by the probing device from $\mathbb{R}^3 \setminus \Omega$ after the insertion of p, since it was so before by (IH).
P1.2 If v did not exist or was not active before the insertion of p, then v is visited by DETECT_ACCESS during the update of $\mathrm{Del}^v_{|S}(E)$. Since we run DETECT_ACCESS only on new vertices lying in $\mathbb{R}^3 \setminus \Omega$ and on former active vertices, v is given a free path either to a new vertex lying in $\mathbb{R}^3 \setminus \Omega$, or to a former active vertex which, as explained in P1.1, remains reachable by the probing device after the insertion of p. In both cases, v is reachable by the probing device from $\mathbb{R}^3 \setminus \Omega$.

(P2) Let us prove that the vertices v and w of any free edge e of $\mathrm{VG}(E \cup \{p\})$ have the same status after the update of $\mathrm{Del}^v_{|S}(E)$. It will then follow, by transitivity, that (P2) still holds after the insertion of p.

P2.1 If a vertex of e (say v) is visited by DETECT_ACCESS during the update of $\mathrm{Del}^v_{|S}(E)$, then it becomes active if not so before, and DETECT_ACCESS visits also w if the latter is not active. Thus, v and w are both active afterwards.

P2.2 If neither v nor w is visited by DETECT_ACCESS, then they keep their status during the update of $\mathrm{Del}^v_{|S}(E)$. Hence, it suffices to prove that they have the same status right before. If neither v nor w is a vertex of $\mathrm{V}(p)$, then they are both old Voronoi vertices, and e is an old edge, which implies that v and w have the same status, by (IH). If one of them belongs to $\mathrm{V}(p)$, then none can be active, since otherwise, during the update of $\mathrm{Del}^v_{|S}(E)$, the algorithm would run DETECT_ACCESS on the one(s) that is (are) active, hereby contradicting the hypothesis of P2.2. □

8.3.3 Geometric properties of the output

In this section, we prove that $\mathrm{Del}^v_{|S}(E_F)$ approximates S topologically and gemetrically. Specifically, $\mathrm{Del}^v_{|S}(E_F)$ is isotopic to S and close to S for the Hausdorff distance. Moreover, E_F is a $O_S(\varepsilon)$-sample of S.

Theorem 8.2 *If S has a positive reach, then E_F and $\mathrm{Del}^v_{|S}(E_F)$ verify assertions (a)–(e) of Theorem 4.4. If S is k-Lipschitz, with $k \leq \tan\left(\frac{1}{2}\arctan\frac{1}{3\sqrt{3}}\right) \approx 0.1$, then E_F and $\mathrm{Del}^v_{|S}(E_F)$ verify assertions (â)–(ê) of Theorem 4.5, with $\varrho = 1$.*

For simplicity, we focus on the smooth case, the Lipschitz case being quite similar. From now on, we assume that the surface S has a positive reach.

Proof of the theorem. All we have to do is to show that $\mathrm{Del}^v_{|S}(E_F)$ verifies assertions I1–I3 of Section 1.2. The theorem will then follow.

Proof of I3. By definition, every facet of $\mathrm{Del}^v_{|S}(E)$ is circumscribed by a ball of $\mathrm{Del}^v_{|S}(E)$. Since the algorithm eliminates the balls of $\mathrm{Del}^v_{|S}(E)$ that have radii greater than ε, all the balls of $\mathrm{Del}^v_{|S}(E_F)$ have radii at most $\varepsilon < 0.07\ \varepsilon_S$, which is less than $\mu_0\ \mathrm{rch}(S)$, by hyp. A3 of Section 8.1.2. □

To prove I1, we need a technical result, which is a direct consequence of assertion (P2):

Claim 8.2.1 *Let ζ be a connected component of $\mathrm{VG}(E) \setminus \mathcal{O}$. Either all the points of $\partial\zeta \cap S$ are centers of balls of $\mathrm{Del}^v_{|S}(E)$, or none of them is.*

Proof. Let p and q be two points of $\partial\zeta \cap S$. By definition, p and q are centers of balls of $\mathrm{Del}_{|S}(E)$. If ζ contains no (finite or infinite) Voronoi vertex, then it is made of one piece of a Voronoi edge only. Therefore, p and q cannot be detected by the probing device, and none of them can be the center of a ball of $\mathrm{Del}^v_{|S}(E)$. If ζ contains some Voronoi vertices, then, by (P2), all the Voronoi vertices in ζ have the same status, *active* or *inactive*. In the first case, p and q are both centers of balls of $\mathrm{Del}^v_{|S}(E)$. In the second case, none of them is, which ends the proof of the claim. □

Using the above claim, we can now prove I1.

Proof of I1. Assertions M1 and M3 of Section 1.2 still hold in our context, since they do not involve restricted Delaunay facets. Moreover, since the balls of $\mathrm{Del}^v_{|S}(E)$ are surface Delaunay balls and since they have small radii upon termination of the algorithm, assertion M2 holds with $\mathrm{Del}_{|S}(E_F)$ replaced by $\mathrm{Del}^v_{|S}(E_F)$. As a consequence, the projection lemma 1.8 holds for the facets of $\mathrm{Del}^v_{|S}(E_F)$. It follows that every edge of $\mathrm{Del}^v_{|S}(E_F)$ is incident to at most two facets of $\mathrm{Del}^v_{|S}(E_F)$, by the same proof as in Lemma 1.10.

Let e be an edge of $\mathrm{Del}^v_{|S}(E_F)$ and $\mathrm{V}(e)$ be the dual Voronoi facet of e. Notice that $\partial\mathrm{V}(e) \cap S \neq \emptyset$, since e belongs to $\mathrm{Del}_{|S}(E_F)$. It follows that any connected component ξ of $\partial\mathrm{V}(e) \setminus \mathcal{O}$ is a simple polygonal arc, whose endpoints lie on S and are centers of balls of $\mathrm{Del}_{|S}(E_F)$. Moreover, ξ is included in a connected component of $\mathrm{VG}(E_F)\setminus\mathcal{O}$. Thus, by Claim 8.2.1, either both endpoints of ξ are centers of balls of $\mathrm{Del}^v_{|S}(E_F)$, or none of them is. It follows that the total number of centers of balls of $\mathrm{Del}^v_{|S}(E_F)$ that lie on $\partial\mathrm{V}(e)$ is even. In addition, by Remark 1.9, every edge of $\partial\mathrm{V}(e)$ contains at most one center of ball of $\mathrm{Del}^v_{|S}(E_F)$. Thus, the number of edges of $\partial\mathrm{V}(e)$ that contain centers of balls of $\mathrm{Del}^v_{|S}(E_F)$ is even. Equivalently, the number of facets of $\mathrm{Del}^v_{|S}(E_F)$ that are incident to e is even.

Hence, every edge of $\mathrm{Del}^v_{|S}(E_F)$ is incident to exactly two facets of $\mathrm{Del}^v_{|S}(E_F)$. In addition, the proof of Lemma 1.11 relies only on M1, M3, M5, M6, and on the properties of the umbrellas made of small restricted Delaunay facets, without any additional argument, thus it holds when we replace $\mathrm{Del}_{|S}(E_F)$ by $\mathrm{Del}^v_{|S}(E_F)$. It follows that every vertex of $\mathrm{Del}^v_{|S}(E_F)$ has one umbrella.

As a conclusion, $\mathrm{Del}^v_{|S}(E_F)$ is a compact surface without boundary. The fact that it is consistently oriented comes from Lemma 1.12, whose proof uses only the projection lemma 1.8. □

Proof of I2. Recall that E_I is made of the vertices of a λ-persistent facet f, with $\lambda = 0.014\ \varepsilon_S$. If f remains in $\mathrm{Del}_{|S}(E)$ until the end of the process, then $\mathrm{VG}(E_F) \cap S$ is not empty. Since $\mathrm{VG}(E_F)$ is connected, at least one point p of $\mathrm{VG}(E_F)\cap S$ belongs to the same connected component of $\mathrm{VG}(E_F)\setminus\mathcal{O}$ as some infinite Voronoi vertex. By (P2), p can be "seen" from an active Voronoi vertex. Hence, $\mathrm{Del}^v_{|S}(E_F)$ is not empty, which proves I2.

If f disappears from $\mathrm{Del}^v_{|S}(E)$ during the process, then, before f vanishes, I2 is verified by $\mathrm{Del}^v_{|S}(E)$, by the same argument as above. Therefore, by the same reasoning as in the proof of Theorem 4.4, $\mathrm{VG}(E_F) \cap S$ is not empty. Hence, $\mathrm{Del}^v_{|S}(E_F)$ verifies I2. □

Since $\mathrm{Del}^v_{|S}(E_F)$ verifies I1–I3, the various results of Chapter 1, applied with $\hat{S} = \mathrm{Del}^v_{|S}(E_F)$, show that assertions (a)–(e) of Theorem 4.4 hold with $\mathrm{Del}_{|S}(E_F)$ replaced by $\mathrm{Del}^v_{|S}(E_F)$. This concludes the proof of Theorem 8.2. □

In the smooth case, since $\varepsilon < 0.07\ \varepsilon_S < 0.07\ \mathrm{rch}(S)$, E_F is a $\mu_0\mathrm{rch}(S)$-sample, according to assertion (c) of Theorem 4.4. By Lemma 1.22, E_F is a loose $\mu_0\mathrm{rch}(S)$-sample. Thus, $\mathrm{Del}_{|S}(E_F)$ is isotopic to S, by Theorem 1.14. It follows that $\mathrm{Del}^v_{|S}(E_F)$ and $\mathrm{Del}_{|S}(E_F)$ are equal, since they are isotopic and since $\mathrm{Del}^v_{|S}(E_F)$ is a subcomplex of $\mathrm{Del}_{|S}(E_F)$. The same argument holds in the Lipschitz case, with Theorem 4.4 replaced by Theorem 4.5, Lemma 1.22 by Lemma 2.38, and Theorem 1.14 by Theorem 2.33.

8.4 Complexity of the algorithm

As mentioned in the introduction, the complexity of the algorithm has three components: the *combinatorial cost* that measures the memory space and time needed to store, construct and update the data structures; the *probing cost* that counts the number of probes performed by the probing device; the *displacement cost* that measures the effort spent in moving the probing device. Depending on the context, one can give emphasis to one type of cost or the other.

Notice that it is not possible in general to optimize all costs simultaneously. Take for instance the parabola $\mathcal{P}$ of Figure 4.2. Any Delaunay-based algorithm that optimizes the displacements of the probing device will somehow follow the curve $\mathcal{P}$, inserting the points of E more or less in their order along $\mathcal{P}$ (see Figure 4.2, left). This makes the overall complexity of the incremental Delaunay triangulation quadratic. Differently, our algorithm will insert the points in an order defined by the *largest empty ball* criterion (see Figure 4.2, right), which does not optimize the displacement cost but makes the combinatorial cost linear.

Open question 8.3 *What are the exact trade-offs between optimizing the combinatorial cost and the displacement cost?*

In the sequel, we analyze the combinatorial cost, probing cost and displacement cost separately. Since our algorithm enforces the probing device to move along the Voronoi edges, the size of the Voronoi diagram has a direct impact on all three costs. Let N be the size of the output point set E_F, and T be the overall number of Delaunay tetrahedra created during the course of the algorithm. By Theorem 8.2, we know that $N = O_S(\varepsilon^{-2})$.

8.4.1 Combinatorial cost

Lemma 8.4 *The space complexity of the algorithm is $O_S\left(N \log N\right) = O_S\left(\varepsilon^{-2} \log \frac{1}{\varepsilon}\right)$ if S is a smooth surface that satisfies the genericity condition of Lemma 4.8. Otherwise, the space complexity is $O_S(N^2) = O_S\left(\varepsilon^{-4}\right)$.*

Proof. The data structure stores the current Delaunay triangulation as well as some of the former Voronoi vertices. Since every vertex is stored at most once, the size of the data structure is at most the total number of Voronoi vertices created during the course of the algorithm, which is equal to T. Moreover, since a Voronoi edge is probed only from its vertices, at most two centers of balls of $\mathrm{Del}^v_{|S}(E)$ are stored in the priority queue $\mathcal{Q}$ per Voronoi edge. As a consequence, the space complexity of the algorithm is $O(T)$.

If S is a smooth surface that satisfies the genericity condition of Lemma 4.8, then $T = O_S(N \log N)$, by Theorem 4.20. Hence, the space complexity is $O_S(N \log N)$. Otherwise, $T = O(N^2)$, since each point insertion creates at most a linear number of Delaunay tetrahedra. Hence, the space complexity is $O(N^2)$. □

Lemma 8.5 *The time complexity of the algorithm is* $O_S\left(N \log^2 N\right) = O_S\left(\varepsilon^{-2} \log^2 \frac{1}{\varepsilon}\right)$ *if* S *is a smooth surface that satisfies the genericity condition of Lemma 4.8. Otherwise, the time complexity is* $O_S(N^2 \log N) = O_S\left(\varepsilon^{-4} \log \frac{1}{\varepsilon}\right)$.

Note that the constants in the O notations depend on the time spent by the probing device to perform a move or a probe. This amount of time is considered as a constant once the probing device is given. The proof of the lemma is very similar to that of Lemma 4.15, with an additional subtelty. We recall it here, for completeness.

Proof. - The cost of maintaining $\mathrm{Del}(E)$ is $O(T)$ since no point location is performed in our case.

- The cost of updating $\mathrm{Del}^v_{|S}(E)$ is also $O(T)$ since DETECT_ACCESS stops each time it reaches an active vertex and any vertex that becomes active remains so. Hence, the number of times a vertex is visited is at most the total number of incident Voronoi edges created by the algorithm.

- Since a Voronoi edge is probed from its vertices, it contains at most two centers of balls of $\mathrm{Del}^v_{|S}(E)$. Hence, the cost of maintaining the priority queue $\mathcal{Q}$ of bad balls of $\mathrm{Del}^v_{|S}(E)$ is $O(T \log T)$ since the total number of centers of balls of $\mathrm{Del}^v_{|S}(E)$ inserted in $\mathcal{Q}$ (and then retrieved from it) is at most twice the total number of Voronoi edges created during the process.

As a conclusion, the time complexity is $O(T \log T)$. If S is smooth and satisfies the genericity condition of Lemma 4.8, then $T = O_S(N \log N)$ by Theorem 4.20, and the time complexity is $O_S(N \log^2 N)$. Otherwise, $T = O(N^2)$ and the time complexity is $O(N^2 \log N)$. □

8.4.2 Probing cost

The algorithm probes only along the Voronoi edges and from their vertices. Since every Voronoi edge has two vertices, it is probed at most twice. Hence, the total number of probes is at most twice the total number of Voronoi edges created during the process, which is linear w.r.t. T. It follows that the probing cost is $O_S(N \log N) = O_S\left(\varepsilon^{-2} \log \frac{1}{\varepsilon}\right)$ if S is smooth and satisfies the genericity condition of Lemma 4.8, and $O(N^2) = O_S\left(\varepsilon^{-4}\right)$ otherwise.

8.4.3 Displacement cost

We bound the total number of Voronoi edges travelled by the probing device. During the update of $\mathrm{Del}^v_{|S}(E)$, two types of displacements are performed (see Section 8.2.1): *detection displacements* are performed inside the routine DETECT_ACCESS to locate the intersection points with the surface S; *positioning displacements* are performed during the update of $\mathrm{Del}^v_{|S}(E)$, when the probing device is moved from one place of $\mathrm{VG}(E)$ to another, before issuing a new sequence of probes.

Lemma 8.6 *The displacement cost of the algorithm is* $O_S\left(N^2 \log N\right) = O_S\left(\varepsilon^{-4} \log \frac{1}{\varepsilon}\right)$ *if* S *is a smooth surface that satisfies the genericity condition of Lemma 4.8, and* $O(N^4) = O_S\left(\varepsilon^{-8}\right)$ *otherwise.*

Proof. The overall cost of the detection displacements has been analyzed in the proof of Lemma 8.5 and shown to be $O_S\left(N \log N\right)$ if S is smooth and satisfies the genericity condition of Lemma 4.8, and $O(N^2)$ otherwise.

At each iteration, the point that is being inserted has $O(N)$ neighbors in $\mathrm{Del}(E)$. Hence, the probing device follows $O(N)$ paths from or to $\mathbb{R}^3 \setminus \Omega$ during the positioning displacements. The length of each path is at most the total number of Voronoi edges created, which is $O(N^2)$. Hence, the cost of the positioning displacements is $O(N^3)$ per iteration. This yields an overall cost of $O(N^4)$.

If S is smooth and verifies the genericity condition of Lemma 4.8, then we can achieve a better bound. According to Lemma 4.13, for every iteration $i > i_0$, $E(i)$ is a $2r(i)$-sample of S, with $2r(i) \leq \varepsilon_0$. Hence, by Theorem[5] 1.1, every cell of $\mathrm{Vor}(E(i))$ intersects S along a topological disk that divides the cell into two components: one lies in $\mathcal{O}$, the other lies in $\mathbb{R}^3 \setminus \mathcal{O}$. Therefore, if $p(i)$ is the point inserted in E at iteration i, then, right after its insertion, all the vertices of its Voronoi cell $\mathrm{V}(p(i))$ that can be reached by the probing device will be marked *active* during the first call to DETECT_ACCESS. As a consequence, the algorithm has to call DETECT_ACCESS on only one vertex of $\mathrm{V}(p(i))$ (or on its neighbor). Hence, at iteration i, two paths only are followed by the probing device during the positioning displacements. The lengths of these two paths are bounded by the overall number of Voronoi vertices created before iteration i. This number is $O_S(N \log N)$, by Theorem 4.20. Hence, the overall cost of the positioning displacements after iteration i_0 is $O_S\left(N^2 \log N\right)$. □

The bound for the generic smooth case is almost tight, since on some input objects with generic smooth boundaries the displacement cost of the algorithm is $\Omega\left(N^2\right)$. Figure 8.3 presents an example in the plane. The top image shows the object $\mathcal{O}$ and the initial point sample E_I, both symmetric with respect to the origin (marked by a point at the center of the object). The bottom image shows the point sample E and the balls of $\mathrm{Del}^v_{|S}(E)$ at some stage of the course of the algorithm. Since at each iteration the algorithm inserts the center of the largest ball of $\mathrm{Del}_{|S}(E)$, it is easily seen that E remains symmetric (or almost symmetric) with respect to the origin throughout the process. Hence, each time a point p lying inside a cavity is inserted in E, the iteration before or after the algorithm inserts in E the symmetric of p, which lies in the other cavity. Since the density of the output point sample is uniform, the number of points inserted inside the cavities is linear with respect to N. Therefore, the overall number of Voronoi edges travelled by the probing device is $\Omega\left(N^2\right)$.

8.5 Implementation and results

We have implemented the probing algorithm using the C++ library CGAL [33], which provided us with robust and flexible implementations of the Delaunay triangulation in 2D and in 3D. A video [22] is available online, which describes the algorithm and demonstrates its practicality. Results on a planar curve and on a surface are reported in Figures 8.4 and 8.5. The active part of the Voronoi graph is printed in blue, the inactive part in black. Moreover, the faces of $\mathrm{Del}^v_{|S}(E)$ are shown in green or in red, depending on whether they are circumscribed by a bad ball of $\mathrm{Del}^v_{|S}(E)$ or not. In the 2D example, the inactive part of the Voronoi graph is shown only in the first image, for clarity.

Open question 8.7 *In practice a physical scaffold has to be present around the object being sampled to support the probing device. Can we extend our theoretical results to the case of a probing device whose*

[5] In fact we refer to the proof of this theorem (see [6]), which shows that every d-face of $\mathrm{Del}_{|S}(E)$ intersects S along a topological d-ball, if ever it intersects S.

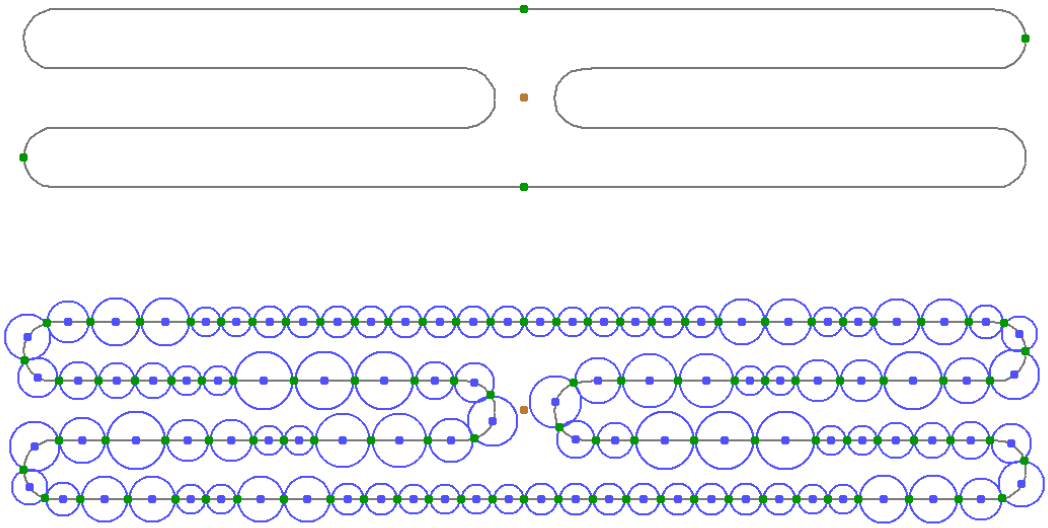

Figure 8.3: A quadratic example.

motions obey realistic constraints?

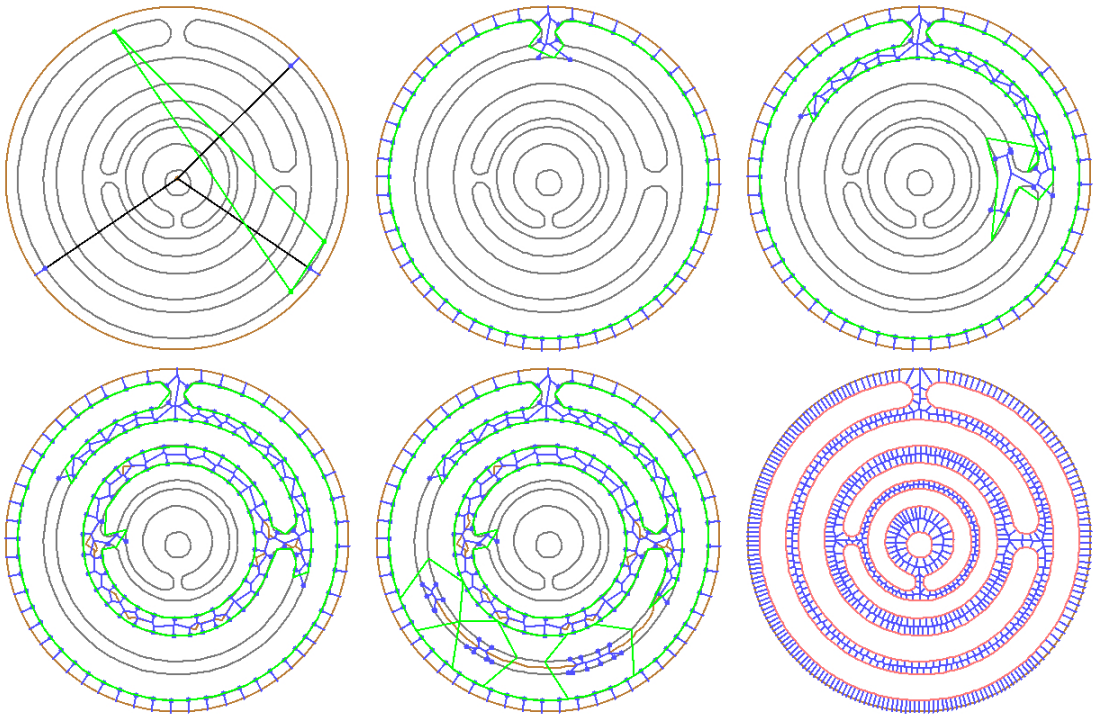

Figure 8.4: Course of the algorithm on a curve in $\mathbb{R}^2$

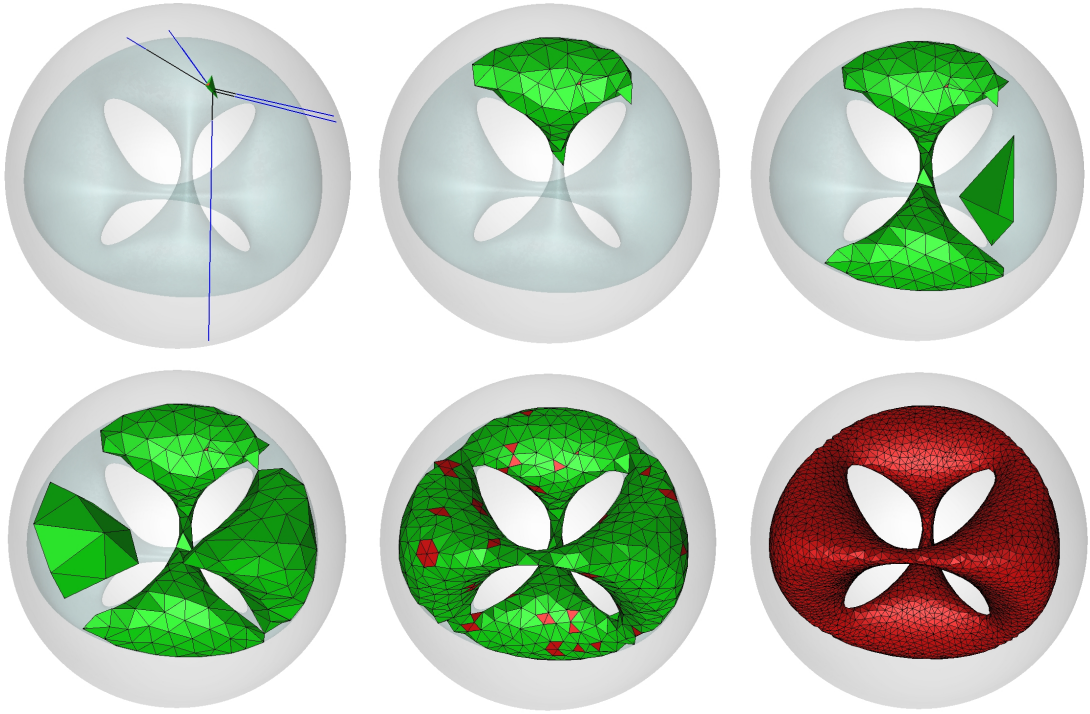

Figure 8.5: Course of the algorithm on a surface in $\mathbb{R}^3$

Chapter 9

Meshing Volumes with Curved Boundaries

9.1 Introduction

In this chapter, we focus on the problem of constructing tetrahedral meshes to approximate 3-dimensional objects with curved boundaries. Recall that a tetrahedral mesh is a simplicial mesh whose simplices have dimension at most 3.

Delaunay refinement, introduced in the plane by Ruppert [99] and then extended to 3-dimensional space by Shewchuk [102], has proven to be a powerful technique for generating provably good meshes. It allows the user to get an easy control on the sizes of the mesh elements, for instance through a (possibly non-uniform) sizing field. Moreover, it constructs meshes with a good grading, able to conform to quickly varying sizing fields.

However, existing Delaunay refinement algorithms deal exclusively with domains bounded by piecewise linear boundaries. In applications where the object to mesh has a curved boundary S, time is spent discretizing S into a polyhedron $\mathcal{P}$, before the interior of the object can be sampled. Then, the original boundary S is dropped away and replaced by its discretized version $\mathcal{P}$. As a result, the quality of the output mesh and the accuracy of the boundary approximation depend highly on the initial surface mesh $\mathcal{P}$.

Several methods have been proposed for meshing two-dimensional or three-dimensional domains with curved boundaries. Most of them deal only with specific types of boundaries (parametric, implicit etc.) [101], or they simply come with no guarantee regarding the topology of the ouput mesh, or the quality of its elements, or even the termination of the process [3, 52, 82]. One noticeable exception is [25], where the algorithm is able to handle any two-dimensional domain bounded by piecewise smooth curves, of any type, provided that a small number of geometric quantities can be estimated, such as the curvature of a given curve at a given point or the total variation of the unit tangent vector between two points on a given curve. The problem with this method is that it is designed exclusively for the two-dimensional case. Moreover, estimating the required geometric quantities can be time-consuming on certain types of curves.

9.1.1 Statement of the problem

Given a compact object $\mathcal{O}$ whose boundary S is a smooth or Lipschitz surface, our aim is to construct a 3-dimensional simplicial mesh whose boundary is an accurate approximation of S. The accuracy is measured by the Hausdorff distance. Moreover, the elements of the mesh (facets and tetrahedra) must have good aspect ratios and their circumradii are to be prescribed by a user-defined sizing field. Recall that the aspect ratio of a simplex is the ratio between its circumradius and the radius of its inscribed sphere.

9.1.2 Our approach

Recall that our surface mesher constructs iteratively a point set $E \subset S$ and maintains its 3-dimensional Delaunay triangulation $\mathrm{Del}(E)$. As shown in Chapter 4, upon termination, $\mathrm{Del}_{|S}(E)$ approximates S both topologically and geometrically. It can then be proved that the so-called Delaunay triangulation of E restricted to $\mathcal{O}$ is a good approximation of $\mathcal{O}$ – see Section 9.3.

Definition 9.1 *Given a finite point set E, the Delaunay triangulation of E restricted to $\mathcal{O}$, or $\mathrm{Del}_{|\mathcal{O}}(E)$ for short, is the subcomplex of $\mathrm{Del}(E)$ formed by the tetrahedra whose dual Voronoi vertices lie in $\mathcal{O}$.*

It suffices then to refine the tetrahedra of $\mathrm{Del}_{|\mathcal{O}}(E)$ by inserting their circumcenters in E, until they eventually satisfy some user-defined shape and size criteria. Note however that refining a tetrahedron of $\mathrm{Del}_{|\mathcal{O}}(E)$ means inserting a point in $\mathcal{O}$, which implies that E is no longer included in S afterwards. In order to apply the theoretical results of Part A, we have to make sure that the Delaunay triangluation of $E \cap S$ restricted to S is preserved. We will do this by inserting additional points of S whenever it is necessary.

9.1.3 Overview

We introduce our algorithm in Section 9.2. Section 9.3 deals with the accuracy of the approximation of the object by the output mesh. In Section 9.4, we prove that the algorithm terminates, and we bound the number of vertices of the output mesh. Section 9.5 addresses the practicality of the algorithm: it gives some details about the choice of the sizing field and it explains how to remove the so-called *slivers*. Finally, Section 9.6 provides a few examples and experimental results.

Note to the reader: *For clarity, the rest of the chapter is written entirely for the case where S is smooth. Our proofs can be easily extended to the Lipschitz setting, provided that the open question 2.20 is solved.*

9.2 Main algorithm

The algorithm takes as input the domain $\mathcal{O}$ to mesh, a sizing field σ, and two parameter values α and ϱ_τ. The domain is known through an exhaustive oracle that can tell whether a given segment intersects S and, in the affirmative, return all the points of intersection (which are finitely many, generically). Implementations of exhaustive oracles are given in Chapters 6 and 7, for specific applications. We

assume further that the oracle can tell whether a given point lies inside $\mathcal{O}$ or outside. The sizing field is a positive function $\sigma : \bar{\mathcal{O}} \to \mathbb{R}^+$ defined over $\bar{\mathcal{O}}$ and assumed to be 1-Lipschitz.

The algorithm first uses Chew's algorithm to construct an initial point set $E_I \subset S$ that verifies the following conditions:

- $\forall p \in S, \mathrm{d}(p, E_I) \leq 0.09\ \mathrm{d}_{\mathrm{M}}(p)$;
- $\forall p \in E_I, \mathrm{d}(p, E_I \setminus \{p\}) \geq 0.03\ \mathrm{d}_{\mathrm{M}}(p)$.

Once E_I is built, the algorithm constructs E iteratively, starting with $E = E_I$ and inserting one point in E per iteration. In the meantime, the restricted Delaunay triangulations $\mathrm{Del}_{|\mathcal{O}}(E)$ and $\mathrm{Del}_{|S}(E)$ are maintained, using the oracle.

At each iteration, one element of the mesh (a facet of $\mathrm{Del}_{|S}(E)$ or a tetrahedron of $\mathrm{Del}_{|\mathcal{O}}(E)$) is refined. To refine a tetrahedron, the algorithm inserts its circumcenter in E. A facet f of $\mathrm{Del}_{|S}(E)$ may be circumscribed by several surface Delaunay balls. Thus, to refine f, the algorithm inserts in E the center of the surface Delaunay ball $B(c, r)$ circumscribing f with largest ratio $r/\sigma(c)$. The choice of the next element to refine is driven by the following rules, considered in this order:

R1 if a facet f of $\mathrm{Del}_{|S}(E)$ does not have its three vertices on S, then refine f;

R2 if a facet f of $\mathrm{Del}_{|S}(E)$ has a surface Delaunay ball $B(c, r)$ with ratio $r/\sigma(c) > \alpha$, then refine f;

R3 if a tetrahedron t of $\mathrm{Del}_{|\mathcal{O}}(E)$ has a circumradius greater than $\sigma(c)$, where c is the circumcenter of t, or if t has a radius-edge ratio greater than ϱ_τ, then consider the circumcenter c of t:

R3.1 if c is not included in any surface Delaunay ball, then insert c in E;

R3.2 else, insert in E the center of a surface Delaunay ball containing c.

The algorithm terminates when the triggering conditions of Rules R1, R2 and R3 are no longer met. The point set E is then renamed E_F and returned as well as $\mathrm{Del}_{|\mathcal{O}}(E_F)$.

Every facet of $\mathrm{Del}_{|S}(E_F)$ has its three vertices on S (Rule R1) and every surface Delaunay ball $B(c, r)$ has a radius $r \leq \alpha\ \sigma(c)$ (Rule R2). Moreover (Rule R3), every tetrahedron t of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ has a circumradius $r \leq \min\{\sigma(c),\ \varrho_\tau\ l_{\min}\}$, where c is the circumcenter of t and $l_{\min}$ is the length of the shortest edge of t.

9.3 Approximation accuracy

In this section, we assume that the algorithm terminates. Termination is discussed in Section 9.4, which uses several results stated here. Let $E_{F|S} = E_F \cap S$.

Since E_I is a $0.09\mathrm{d}_{\mathrm{M}}$-sample of S, $E_{F|S}$ is also a $0.09\mathrm{d}_{\mathrm{M}}$-sample of S, since no point is deleted during the course of the algorithm. Hence, $\mathrm{Del}_{|S}(E_{F|S})$ is a closed 2-manifold isotopic to S, by Theorem 1.33. Therefore, to have topological guarantees on the output of the algorithm, it suffices to prove that the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is equal to $\mathrm{Del}_{|S}(E_{F|S})$.

There exists a strong relationship between the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ and $\mathrm{Del}_{|S}(E_F)$:

Lemma 9.2 *The boundary of* $\mathrm{Del}_{|\mathcal{O}}(E_F)$ *is a subcomplex of* $\mathrm{Del}_{|S}(E_F)$. *Moreover, if every edge of the Voronoi diagram* $\mathrm{Vor}(E_F)$ *intersects* S *at most once, and transversally, then the boundary of* $\mathrm{Del}_{|\mathcal{O}}(E_F)$ *is equal to* $\mathrm{Del}_{|S}(E_F)$.

Proof. Since $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is a union of Delaunay tetrahedra, its boundary is a union of Delaunay facets. Let f be a facet of the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$. By definition, it belongs to two Delaunay tetrahedra, one of which has its dual Voronoi vertex inside $\mathcal{O}$, whereas the other one has its dual Voronoi vertex outside $\mathcal{O}$. It follows that the Voronoi edge dual to f intersects S, which means that $f \in \mathrm{Del}_{|S}(E_F)$.

Let us now assume that every edge of $\mathrm{Vor}(E_F)$ intersects S at most once. Let f be a facet of $\mathrm{Del}_{|S}(E_F)$. By definition, the Voronoi edge dual to f intersects S. Since this edge intersects S only once, one of its vertices lies inside $\mathcal{O}$ whereas the other one (which may be at infinity) lies outside $\mathcal{O}$. It follows, by definition of $\mathrm{Del}_{|\mathcal{O}}(E_F)$, that one of the Delaunay tetrahedra incident to f belongs to $\mathrm{Del}_{|\mathcal{O}}(E_F)$, while the other one does not. Hence, f belongs to the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$. □

In our case, $\mathrm{Del}_{|S}(E_F)$ is precisely the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$, due to the following result:

Lemma 9.3 *Every edge of* $\mathrm{Vor}(E_F)$ *intersects* S *at most once, and transversally.*

Proof. Among the edges of $\mathrm{Vor}(E_F)$, only those whose dual Delaunay facets have their three vertices on S can intersect S, thanks to Rule R1. Let e be such an edge. It is included in an edge e' of $\mathrm{Vor}(E_{F|S})$. Since $E_{F|S}$ is a $0.09\mathrm{d_M}$-sample of S, Remark 1.9 tells that e' intersects S at most once, and transversally, which yields the lemma. □

Corollary 9.4 *The boundary of* $\mathrm{Del}_{|\mathcal{O}}(E_F)$ *is* $\mathrm{Del}_{|S}(E_F)$.

It follows from Corollary 9.4 that, if we can prove that $\mathrm{Del}_{|S}(E_F) = \mathrm{Del}_{|S}(E_{F|S})$, then the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ will be equal to $\mathrm{Del}_{|S}(E_{F|S})$, which is isotopic to S. We need an intermediate result.

Lemma 9.5 $\mathrm{Del}_{|S}(E_F)$ *has vertices on all the connected components of* S.

Proof. By Rule R1, every edge e of $\mathrm{Vor}(E_F)$ that intersects S has a dual Delaunay facet f whose three vertices are in $E_{F|S}$. Since $E_{F|S}$ is a $0.09\mathrm{d_M}$-sample of S, the point $c = e \cap S$ lies at distance at most $0.09\ \mathrm{d_M}(c)$ from the vertices of f. It follows that the Delaunay ball centered at c intersects S along a topological disk, by Lemma 1.26. Hence, c and the vertices of f lie on the same connected component of S. As a consequence, to prove the lemma, it suffices to show that every connected component of S is intersected by at least one Voronoi edge.

Notice that every connected component S_i of S is the fronteer between two connected components Ω_1 and Ω_2 of $\mathbb{R}^3 \setminus S$, so that every connected path from Ω_1 to Ω_2 crosses S_i. Therefore, to prove that S_i is intersected by a Voronoi edge, it suffices to prove that Ω_1 and Ω_2 both contain Voronoi vertices, since the Voronoi graph of E_F is connected.

Let us assume for a contradiction that some component Ω of $\mathbb{R}^3 \setminus S$ contains no Voronoi vertex. Since the Delaunay balls centered at the Voronoi vertices (including the ones at infinity) cover $\mathbb{R}^3$, at least one such ball (say $B(c, r)$) contains a point p of $\mathrm{M} \cap \Omega$, where M is the medial axis of S. Since c lies outside Ω while p lies inside, the line segment $[c, p]$ intersects the boundary of Ω (which is part of S). Let q be a point of intersection. The ball centered at q, of radius $\mathrm{d}(p, q)$, is contained in the interior of $B(c, r)$. Therefore, it contains no point of E_F. Now, its radius is $\mathrm{d}(p, q)$, which is at least the distance

from q to M since $p \in \mathrm{M} \cap \Omega$. Hence, q is farther from E_F than $\mathrm{d}_\mathrm{M}(q)$, which contradicts the fact that $E_{F|S}$ is a $0.09\mathrm{d}_\mathrm{M}$-sample of S. It follows that Ω contains at least one Voronoi vertex, which ends the proof of Lemma 9.5. □

We can now prove that $\mathrm{Del}_{|S}(E_F) = \mathrm{Del}_{|S}(E_{F|S})$, by using the fact that $\mathrm{Del}_{|S}(E_F)$ is the boundary of a three-dimensional object, namely $\mathrm{Del}_{|\mathcal{O}}(E_F)$ (Corollary 9.4).

Lemma 9.6 $\mathrm{Del}_{|S}(E_F) = \mathrm{Del}_{|S}(E_{F|S})$.

Proof. Thanks to Rule R1, all the facets of $\mathrm{Del}_{|S}(E_F)$ have their three vertices in $E_{F|S}$, hence their dual Voronoi edges are included in edges of $\mathrm{Vor}(E_{F|S})$. It follows that $\mathrm{Del}_{|S}(E_F)$ is a subcomplex of $\mathrm{Del}_{|S}(E_{F|S})$.

To prove the lemma, it suffices then to show that every facet of $\mathrm{Del}_{|S}(E_{F|S})$ is also a facet of $\mathrm{Del}_{|S}(E_F)$. Let us assume for a contradiction that there exists a facet f of $\mathrm{Del}_{|S}(E_{F|S})$ that is not a facet of $\mathrm{Del}_{|S}(E_F)$. Let C be the connected component of $\mathrm{Del}_{|S}(E_{F|S})$ to which f belongs. By Lemma 1.26, the vertices of C belong to a single component S_i of S. By Lemma 9.5, at least one vertex v of $\mathrm{Del}_{|S}(E_F)$ lies on S_i. Let C' be the connected component of $\mathrm{Del}_{|S}(E_F)$ that contains v. Since $\mathrm{Del}_{|S}(E_F)$ is a subset of $\mathrm{Del}_{|S}(E_{F|S})$, C' is included in C (which is a connected 2-manifold without boundary). Moreover, since f is not included in C' while v is, C' has a boundary. Now, by Corollary 9.4, C' is a connected component of the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$. Thus, C' cannot have a boundary, which raises a contradiction. □

It follows from the above results that the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is equal to $\mathrm{Del}_{|S}(E_{F|S})$, which is ambient isotopic to S, by Theorem 1.33. In addition to this topological result, we would like to give a bound on the Hausdorff distance between S and the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$, depending on the input sizing field σ. Let $\mu = \min\{0.09,\ \sup_{p\in S} \frac{\alpha\,\sigma(p)}{\mathrm{d}_\mathrm{M}(p)}\}$. Our bound will depend on μ. So far, we know that $E_{F|S}$ is a $0.09\mathrm{d}_\mathrm{M}$-sample of S.

Lemma 9.7 *The surface Delaunay balls of E_F and those of $E_{F|S}$ are the same.*

Proof. Since every edge of $\mathrm{Vor}(E_F)$ that intersects S is included in an edge of $\mathrm{Vor}(E_{F|S})$, the surface Delaunay balls of E_F are also surface Delaunay balls of $E_{F|S}$. Let us show that the converse is true. Let e be an edge of $\mathrm{Vor}(E_{F|S})$. If $e \cap S \neq \emptyset$, then $|e \cap S| = 1$, by Remark 1.9. Moreover, the Delaunay facet dual to e belongs to $\mathrm{Del}_{|S}(E_F)$, by Lemma 9.6. This means that e contains an edge e' of $\mathrm{Vor}(E_F)$, such that $|e' \cap S| \geq 1$. Hence $e \cap S = e' \cap S$. □

Thanks to Lemma 9.7, Rule R2 controls the radii of all the surface Delaunay balls of $\mathrm{Del}_{|S}(E_{F|S})$, which implies that, upon termination of the algorithm, $E_{F|S}$ is a loose $\mu\mathrm{d}_\mathrm{M}$-sample of S. From this observation we deduce the following

Theorem 9.8 $\mathrm{Del}_{|\mathcal{O}}(E_F)$ *is a 3-manifold ambient isotopic to $\bar{\mathcal{O}}$, at Hausdorff distance $O_S(\mu^2)$ from $\bar{\mathcal{O}}$, where $\mu = \min\{0.09,\ \sup_{p\in S} \frac{\alpha\,\sigma(p)}{\mathrm{d}_\mathrm{M}(p)}\}$. Moreover, the surface Delaunay balls of E_F cover S.*

Proof. By Corollary 9.4 and Lemma 9.6, the boundary of $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is equal to $\mathrm{Del}_{|S}(E_{F|S})$. Since $E_{F|S}$ is a loose $\mu\mathrm{d}_\mathrm{M}$-sample of S, we know by Theorem 1.33 that there exists an ambient isotopy $\phi : [0,1] \times \mathbb{R}^3 \to \mathbb{R}^3$ that maps S to $\mathrm{Del}_{|S}(E_{F|S})$. The map $\phi(1,.) : \mathbb{R}^3 \to \mathbb{R}^3$ is an ambient homeomorphism that maps the compact 3-manifold $\bar{\mathcal{O}}$ to a compact 3-manifold bounded by $\mathrm{Del}_{|S}(E_{F|S})$. Now, the only compact 3-manifold bounded by $\mathrm{Del}_{|S}(E_{F|S})$ is $\mathrm{Del}_{|\mathcal{O}}(E_F)$ itself[1]. Thus, we have $\phi(1, \bar{\mathcal{O}}) = \mathrm{Del}_{|\mathcal{O}}(E_F)$, which means that $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is ambient isotopic to $\bar{\mathcal{O}}$.

Since $\mathrm{Del}_{|\mathcal{O}}(E_F)$ and $\bar{\mathcal{O}}$ are both compact, their Hausdorff distance is achieved by a pair of points lying on their boundaries. Hence, we have $\mathrm{d}_\mathcal{H}(\mathrm{Del}_{|\mathcal{O}}(E_F),\ \bar{\mathcal{O}}) = \mathrm{d}_\mathcal{H}(\mathrm{Del}_{|S}(E_{F|S}),\ S)$, which by Theorem 1.33 is $O_S(\mu^2)$.

Finally, by Theorem 1.33 and Lemma 9.7, the surface Delaunay balls of E_F cover S. □

Observe that the results of this section do not rely on Rule R3. Hence, they hold not only upon termination, but also during the course of the algorithm, each time neither Rule R1 nor Rule R2 can be applied. In particular, Theorem 9.8 holds every time Rule R3 is triggered. This observation will be instrumental in proving Lemma 9.10 of Section 9.4.

9.4 Termination and size of the output

In this section, we provide conditions on parameters α and ϱ_τ to ensure that the algorithm terminates. We assume that the sizing field σ is 1-Lipschitz over $\bar{\mathcal{O}}$.

Our strategy is to prove an upper bound on the size of the point sample constructed by the algorithm. The termination of the algorithm results from this bound.

Definition 9.9 *Given a point p inserted in E by the algorithm, the* insertion radius *of p, or $r(p)$ for short, is the Euclidean distance from p to E right before its insertion*[2]*. The insertion radius of a point p of the initial point set E_I is the Euclidean distance from p to $E_I \setminus \{p\}$.*

Our first task is to provide a lower bound on the insertion radius of every point inserted by the algorithm. In fact, we will prove a stronger result, stated as Lemma 9.10. We define a sizing field σ_0 which can be considered as an extension of d_M over $\bar{\mathcal{O}}$:

$$\forall p \in \bar{\mathcal{O}},\ \sigma_0(p) = \inf\{\mathrm{d}(p,q) + \mathrm{d}_\mathrm{M}(q) \mid q \in S\} \tag{9.1}$$

It follows immediately from (9.1) that σ_0 is a positive 1-Lipschitz function, equal to d_M on S. As proved in [3, 90, 108], σ_0 is the pointwise maximal 1-Lipschitz function which is at most d_M on S. Let $\sigma'(p) = \min\{\alpha\ \sigma(p),\ 0.03\ \sigma_0(p)\}$, $\forall p \in \bar{\mathcal{O}}$. Notice that, since σ and σ_0 are 1-Lipschitz, σ' is γ-Lipschitz, where $\gamma = \max\{\alpha,\ 0.03\}$.

Lemma 9.10 *If $\alpha < \frac{1}{5}$ and $\varrho_\tau \geq \frac{4}{1-5\gamma}$, then, at any iteration of the algorithm, we have:*

C1 $\forall p \in E,\ r(p) \geq \sigma'(p)$

C2 $\forall p \in E \setminus E_{|S},\ \mathrm{d}(p,S) \geq \frac{1}{1-\gamma}\ \sigma'(p)$

[1] $\mathrm{Del}_{|\mathcal{O}}(E_F)$ is compact because it is a finite union of tetrahedra.

[2] Notice that it is also the length of the smallest Delaunay edge that is created when inserting p.

Proof. We prove the lemma by induction. Initially, we have $E = E_I$, and every point p of E_I verifies C1, since p is father than $0.03\mathrm{d}_\mathrm{M}(p)$ from $E_I \setminus \{p\}$. Moreover, the points of E_I belong to S, thus C2 is also verified. Let us now assume that C1 and C2 are verified by every point of E, up to a certain step where point c is inserted in E. We will prove that c also verifies C1 and C2.

• If Rule R1 is being applied, then c is the center of a surface Delaunay ball of E whose bounding sphere passes through a point $p \in E \setminus E_{|S}$. The insertion radius of c is the radius of the surface Delaunay ball, *i.e.* $r(c) = \mathrm{d}(c,p)$. Moreover, $\mathrm{d}(c,p) \geq \mathrm{d}(p,S)$, which is at is at least $\frac{1}{1-\gamma}\ \sigma'(p)$ by C2. Since σ' is γ-Lipschitz, we have $\sigma'(p) \geq (\sigma'(c) - \gamma\ \mathrm{d}(c,p))$, hence:

$$\mathrm{d}(c,p) \geq \frac{1}{1-\gamma}\left(\sigma'(c) - \gamma\ \mathrm{d}(c,p)\right) \ \Rightarrow\ r(c) = \mathrm{d}(c,p) \geq \sigma'(c)$$

It follows that C1 is verified for c. Moreover, C2 is also verified, since c belongs to S.

• If Rule R2 is applied, then c is the center of a surface Delaunay ball of radius greater than $\alpha\ \sigma(c) \geq \sigma'(c)$, thus the insertion radius of c is at least $\sigma'(c)$, which satisfies C1. Moreover, C2 is satisfied since c belongs to S.

• If Rule R3.1 is applied, then c is the center of a tetrahedron t, and the insertion radius $r(c)$ is the circumradius r of t. According to Rule R3.1, r is either greater than $\sigma(c)$ or greater than $\varrho_\tau\ l_{\min}$, where $l_{\min}$ is the length of the shortest edge of t. In the first case, we have $r > \sigma(c) > \alpha\ \sigma(c) \geq \sigma'(c)$, since $\alpha < 1$. In the second case, we have $r > \varrho_\tau\ l_{\min}$. Among the vertices of the shortest edge of t, let p be the one inserted last. We have $r(p) \leq l_{\min}$, thus $r > \varrho_\tau\ r(p)$. Moreover, by C1, we have $r(p) \geq \sigma'(p)$. Hence, $r \geq \varrho_\tau\ \sigma'(p)$. Since σ' is γ-Lipschitz, $\varrho_\tau\ \sigma'(p)$ is at least $\varrho_\tau\ (\sigma'(c) - \gamma\ \mathrm{d}(c,p)) \geq \varrho_\tau\ (\sigma'(c) - \gamma r)$. It follows that $r \geq \frac{\varrho_\tau}{1+\varrho_\tau\gamma}\ \sigma'(c)$, which means that C1 is verified for c if ϱ_τ satisfies:

$$\varrho_\tau \geq \frac{1}{1-\gamma} \tag{9.2}$$

To check C2, we notice that, in both cases ($r > \sigma(c)$ and $r > \varrho_\tau\ l_{\min}$), $r(c)$ is bounded from below by $\frac{\varrho_\tau}{1+\varrho_\tau\gamma}\ \sigma'(c)$. Let q be a nearest neighbor of c on S. We have $\mathrm{d}(c,S) = d(c,q) \geq \mathrm{d}(c,E_{|S}) - \mathrm{d}(q,E_{|S})$, where $\mathrm{d}(c,E_{|S}) \geq r(c) \geq \frac{\varrho_\tau}{1+\varrho_\tau\gamma}\ \sigma'(c)$.

Since Rule R3 is applied only when R1 and R2 are fulfilled, Theorem 9.8 holds right before c is inserted. Hence, the surface Delaunay balls of E cover S, and q belongs to a surface Delaunay ball B'', of center c'' and radius r''. Let $p \in E$ be a vertex of the facet of $\mathrm{Del}_{|S}(E)$ circumscribed by B''. $d(q,p)$ is at most $2r''$ because p and q both belong to B''. Due to Rule R2, r'' is at most $\alpha\ \sigma(c'')$, which is at most $\alpha(\sigma(q) + r'')$ since σ is 1-Lipschitz. It follows that $\mathrm{d}(q,E_{|S}) \leq \mathrm{d}(q,p) \leq \frac{2\alpha}{1-\alpha}\ \sigma(q)$, which is less than $3\alpha\ \sigma(q)$ because $\alpha < \frac{1}{3}$. Moreover, since $E_I \subseteq E_{|S}$ is a $0.09\mathrm{d}_\mathrm{M}$-sample of S, $\mathrm{d}(q,E_{|S})$ is also bounded by $0.09\ \mathrm{d}_\mathrm{M}(q) = 3\,(0.03\ \sigma_0(q))$. Thus, $\mathrm{d}(q,E_{|S}) \leq 3\ \sigma'(q)$.

Hence, $\mathrm{d}(c,S) = d(c,q) \geq \frac{\varrho_\tau}{1+\varrho_\tau\gamma}\ \sigma'(c) - 3\ \sigma'(q)$. Since σ' is γ-Lipschitz, we have $\sigma'(q) \leq \sigma'(c) + \gamma\ \mathrm{d}(c,q)$, thus $\mathrm{d}(c,S) = d(c,q) \geq \frac{1}{1+3\gamma}\left(\frac{\varrho_\tau}{1+\varrho_\tau\gamma} - 3\right)\sigma'(c)$. It follows that C2 is verified for c if ϱ_τ satisfies:

$$\frac{1}{1+3\gamma}\left(\frac{\varrho_\tau}{1+\varrho_\tau\gamma} - 3\right) \geq \frac{1}{1-\gamma}\ ,\ \textit{i.e.}\ \ \varrho_\tau \geq \frac{4}{1-5\gamma} \tag{9.3}$$

• If Rule R3.2 is applied, then c is the center of a surface Delaunay ball B, of radius $r = r(c)$, containing the circumcenter c' of a tetrahedron t' of circumradius $r' \geq \frac{\varrho_\tau}{1+\varrho_\tau\gamma}\ \sigma'(c')$ (see case R3.1). Since σ' is γ-Lipschitz, we have $\sigma'(c') \geq \sigma'(c) - \gamma\ r(c)$. Moreover, the circumsphere of t' is empty, thus $r' \leq \mathrm{d}(c', p)$, for any point p of E lying on the bounding sphere of B. Since B contains both p and c', $\mathrm{d}(c', p)$ is at most $2r(c)$. Hence,

$$2r(c) \geq \mathrm{d}(c', p) \geq r' \geq \frac{\varrho_\tau}{1 + \varrho_\tau\gamma} \left(\sigma'(c) - \gamma\ r(c)\right), \ \textit{i.e.}\ \ r(c) \geq \frac{\varrho_\tau}{2 + 3\varrho_\tau\gamma}\ \sigma'(c)$$

Therefore, C1 is verified for c if ϱ_τ satisfies:

$$\varrho_\tau \geq \frac{2}{1 - 3\gamma} \tag{9.4}$$

Moreover, C2 is verified because $c \in S$.

To conclude, Conditions C1 and C2 are verified if ϱ_τ and γ satisfy Eqs. (9.2), (9.3) and (9.4), which is granted if we choose $\gamma < \frac{1}{5}$ (and hence $\alpha < \frac{1}{5}$) and $\varrho_\tau \geq \frac{4}{1-5\gamma}$. □

From now on, we assume that $\alpha < \frac{1}{5}$ and that $\varrho_\tau \geq \frac{4}{1-5\gamma}$, where $\gamma = \max\{\alpha,\ 0.03\}$. Given a point p included in E_I or inserted by the algorithm, we define $B(p)$ as the open ball centered at p, of radius $\rho(p) = \frac{1}{2(1+\gamma)}\ \sigma'(p)$.

Lemma 9.11 *The balls $\{B(p)\}$ are pairwise disjoint.*

Proof. Given two points p and q inserted by the algorithm, we assume without loss of generality that q was inserted in E before p. The distance from p to q is then at least $r(p)$. By Lemma 9.10 (Condition C1), we have $r(p) \geq \sigma'(p)$, which is at least $\frac{1}{1+\gamma}\ \sigma'(q)$ since σ' is γ-Lipschitz. Thus, $\mathrm{d}(p, q) \geq \frac{1}{1+\gamma}\ \max\{\sigma'(p), \sigma'(q)\}$. It follows that $\frac{1}{2(1+\gamma)}\ \sigma'(p) + \frac{1}{2(1+\gamma)}\ \sigma'(q) \leq \mathrm{d}(p, q)$, which means that $B(p)$ and $B(q)$ are disjoint. The same argument holds if p or q belongs to E_I. □

Let $\rho^{\inf} > 0$ be the infimum of ρ over S. It follows from Lemma 9.11 that the distance between any two points inserted by the algorithm is at least $2\rho^{\inf}$. Then, by the same packing argument as in the proof of Lemma 4.7, the algorithm terminates.

To compute an upper bound on the size of E_F, we need another result which states that every ball $B(p)$ lies partly inside $\mathcal{O}$, and that the volume of the part included in $\mathcal{O}$ can be lower-bounded.

Lemma 9.12 *For any $p \in E_F$, $B(p) \cap \mathcal{O}$ contains a ball of radius $\frac{1}{2}\ \rho(p)$.*

Proof. We distinguish between two cases:
• If p lies inside $\mathcal{O}$, then, according to Lemma 9.10 (Condition C2), the distance from p to S is at least $\frac{1}{1-\gamma}\ \sigma'(p)$, which is greater than $\frac{1}{2}\ \rho(p)$. Hence, the ball centered at p, of radius $\frac{1}{2}\ \rho(p)$, is included in $\mathcal{O}$.
• Otherwise, p lies on S. There are two medial balls B^+ and B^- tangent to S at p. One of them (say B^+) is included in $\mathcal{O}$, whereas the other one is included in $\mathbb{R}^3 \setminus \mathcal{O}$. Since B^+ is a medial ball, its radius

is at least $\mathrm{d_M}(p) > \sigma'(p)$. Moreover, the radius of $B(p)$ is $\rho(p) < \sigma'(p)$. It follows that the intersection of $B(p)$ with B^+ contains a ball of radius $\frac{1}{2}\ \rho(p)$. □

Theorem 9.13 *If $\alpha < \frac{1}{5}$ and $\varrho_\tau \geq \frac{4}{1-5\gamma}$ (where $\gamma = \min\{\alpha,\ 0.03\}$), then*

$$|E_F| \ = \ O\left(\iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\sigma_0^3} \ + \ \frac{1}{\alpha^3} \iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\sigma^3}\right)$$

where σ_0 depends only on $\mathcal{O}$ (not on σ).

Proof. We use the same scheme as in Chapter 3. We will bound the integral of $1/\sigma'^3$ over $\mathcal{O}$, where σ' is the minimum of $\alpha\ \sigma$ and of $0.03\ \sigma_0$. Since $B(p) \cap \mathcal{O} \subseteq \mathcal{O}$ for any $p \in E_F$, we have

$$\iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\sigma'^3} \geq \iiint_{\bigcup_{p\in E_F}(B(p)\cap\mathcal{O})} \frac{d\mathcal{O}}{\sigma'^3}$$

Moreover, the balls $B(p)$ are pairwise disjoint, by Lemma 9.11, thus

$$\iiint_{\bigcup_{p\in E_F}(B(p)\cap\mathcal{O})} \frac{d\mathcal{O}}{\sigma'^3} = \sum_{p\in E_F} \iiint_{(B(p)\cap\mathcal{O})} \frac{d\mathcal{O}}{\sigma'^3}$$

In addition, since σ' is γ-Lipschitz, we have:

$$\forall q \in B(p),\ \sigma'(q) \leq \sigma'(p) + \gamma\ \mathrm{d}(q,p) \leq \sigma'(p) + \gamma\ \rho(p) = \left(1 + \frac{\gamma}{2(1+\gamma)}\right)\ \sigma'(p)$$

It follows that

$$\sum_{p\in E_F} \iiint_{(B(p)\cap\mathcal{O})} \frac{d\mathcal{O}}{\sigma'^3} \geq \sum_{p\in E_F} \frac{\mathrm{Vol}\,(B(p)\cap\mathcal{O})}{\left(1+\frac{\gamma}{2(1+\gamma)}\right)^3 \sigma'^3(p)}$$

Now, by Lemma 9.12, the volume of $B(p) \cap \mathcal{O}$ is at least $\frac{4}{3}\pi \frac{1}{64(1+\gamma)^3}\ \sigma'^3(p)$, which yields:

$$\sum_{p\in E_F} \frac{\mathrm{Vol}\,(B(p)\cap\mathcal{O})}{\left(1+\frac{\gamma}{2(1+\gamma)}\right)^3 \sigma'^3(p)} \geq \sum_{p\in E_F} \frac{\frac{\pi}{48(1+\gamma)^3}\ \sigma'^3(p)}{\left(1+\frac{\gamma}{2(1+\gamma)}\right)^3 \sigma'^3(p)} = \frac{\pi}{6\,(2+3\gamma)^3}\ |E_F|$$

which is at least $\frac{1}{34}\ |E_F|$ since $\gamma < \frac{1}{5}$. Hence, $|E_F|$ is at most $34\ \iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\sigma'^3}$. Now, $\sigma'(q)$ is defined as the minimum of $\alpha\ \sigma(p)$ and of $0.03\ \sigma_0(p)$, which are positive functions. It follows that $\iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\sigma'^3}$ is at most $\iiint_{\mathcal{O}} \frac{d\mathcal{O}}{\alpha^3\ \sigma^3} + \iiint_{\mathcal{O}} \frac{d\mathcal{O}}{0.03^3\ \sigma_0^3}$, which ends the proof of the theorem. □

9.5 Practicality of the algorithm

9.5.1 Sizing field

The meshing algorithm presented in the previous sections takes as input a sizing field $\sigma : \ \bar{\mathcal{O}} \to \mathbb{R}^+$ which, for the purpose of the analysis in Section 9.4, is assumed to be 1-Lipschitz. This section explains how to deal with user-defined sizing fields that are not Lipschitz or not defined everywhere in $\bar{\mathcal{O}}$.

Let us assume that the user wants a mesh whose grading conforms to a sizing field σ_u that is not 1-Lipschitz. Then we can use the technique of Miller, Talmor and Teng [90] to derive from σ_u a new sizing field σ'_u that is 1-Lipschitz:

$$\forall p \in \bar{\mathcal{O}},\ \sigma'_u(p) = \inf \{\mathrm{d}(p,q) + \sigma_u(q) \mid q \in \bar{\mathcal{O}}\}$$

Notice that $\sigma'_u(p) \leq \sigma_u(p)$, $\forall p \in \bar{\mathcal{O}}$. The field σ'_u is the best 1-Lipschitz approximation of σ_u [3], because any 1-Lipschitz function that is pointwise at most σ_u is also pointwise at most σ'_u.

The meshing algorithm can be run using the sizing field σ'_u, however it is not necessary to compute σ'_u inside $\mathcal{O}$. Indeed, the algorithm requires an evaluation of the sizing field at internal points only in Rule R3, when refining a tetrahedron. A tetrahedron t is refined either when its circumradius is greater than the value of σ at its circumcenter, or when its radius-edge ratio is greater than ϱ_τ. A careful look at the proof of termination shows that a positive 1-Lipschitz lower bound on the circumradii of tetrahedra is sufficient for the proof. Then, since $\sigma'_u(p) \leq \sigma_u(p)$ for any $p \in E$, the proofs still hold if rule R3 is applied when $r \geq \sigma_u(p)$. Besides saving some sizing field evaluations, this variant of the algorithm constructs sparser meshes whose densities conform to the user-defined sizing field, with a grading bounded only by the bound ϱ_τ on the radius-edge ratio.

In the case where the user has no particular sizing requirements, the 1-Lipschitz sizing field used in the analysis is the field σ_0 defined in Eq. (9.1), at the beginning of Section 9.4. Here again, the algorithm does not need to evaluate[3] σ_0 inside $\mathcal{O}$. It may simply skip the size test for tetrahedra and consider for refinement only the tetrahedra with a radius-edge ratio greater than ϱ_τ. Since the occasions of refining tetrahedra are fewer in this variant than in the original version of the algorithm, it is clear that this variant also terminates. Its output is a mesh whose sizing is a fraction of d_M on S and grows as fast as possible (regarding the bound on the radius-edge ratio) when moving towards the medial axis.

In any case, the algorithm needs to compute d_M at some points on S in order to test whether the precondition of Rule R2 is met. This issue has already been addressed in Section 5.4.

9.5.2 Sliver removal

Optimizing radius-edge ratios prevents our output mesh from containing any bad tetrahedra, except possibly *slivers*. Recall that a sliver is a tetrahedron whose vertices are close to a great circle of its circumsphere and equally spaced along this circle.

Cheng *et al.* [41], and later on Cheng and Dey [40], proposed to exude slivers from the mesh by assigning carefully chosen weights to the vertices, so that their weighted Delaunay triangulation contains as few slivers as possible. Li and Teng [84] proposed to avoid slivers by relaxing the choice of refinement vertices inside small areas around the circumcenters of the elements to be refined.

In our context, we use the sliver exudation algorithm of Cheng *et al.* [41] as a post-process. The output mesh is no longer a Delaunay triangulation, but a weighted Delaunay triangulation. Although the theoretically guaranteed bound on aspect ratios is known to be miserably low [41], the method is efficient in practice and generates almost sliver-free meshes [61].

[3] Note however that computing d_M on S is still necessary, in order to check the sizes of the facets of $\mathrm{Del}_{|S}(E)$.

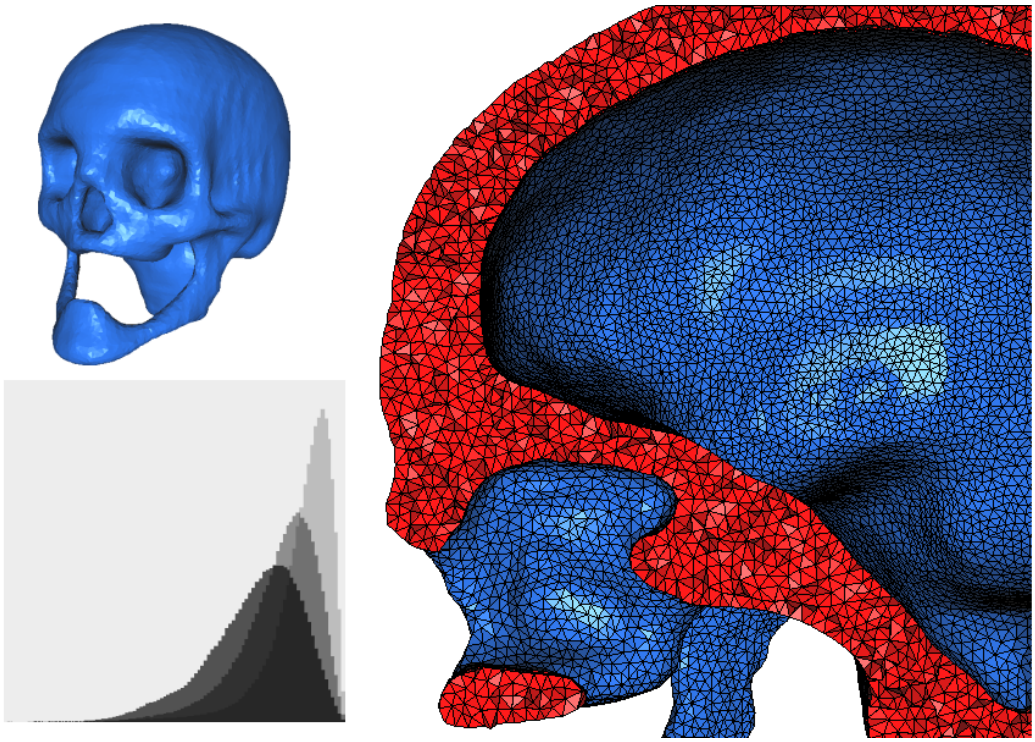

Figure 9.1: Skull model: $89,245$ vertices (among which $35,483$ lie on S) and $442,542$ tetrahedra.

9.6 Implementation and results

The algorithm has been implemented in C++, using the geometric library CGAL [33] which provided us with an efficient and flexible implementation of the three-dimensional Delaunay triangulation.

Figures 9.1 and 9.2 show two meshes generated by our algorithm coupled with the post-processing step described in Section 9.5.2. Each figure is composed of two views of the output mesh: one shows the boundary (top left), the other shows a zoom on the interior, cut by a plane[4] (right). The bottom-left corner of each figure shows the distribution of the inversed aspect ratios of the tetrahedra, represented on a linear scale ranging from 0 to $\frac{1}{3}$ (which corresponds to the inverse of the aspect ratio of a regular tetrahedron). The histograms are normalized with respect to area, so that we can make fair comparisons between meshes of different sizes.

In Figure 9.1, the boundary of the domain is a level set in a 3D grey-scaled image. Its diameter is about 280 millimeters, and its reach approximately 1 millimeter. Although our theoretical results require strict conditions on σ, α and ϱ_τ, in practice the algorithm works well under weaker conditions. For instance, in this example we used a uniform sizing field of 2 millimeters, with $\alpha = 1$ and $\varrho_\tau = 2$, which is far beyond the theoretical limits. Note that the topology of the domain has been captured, and that the boundary has been accurately approximated.

The inversed aspect ratios distribution of our algorithm (in medium grey) has been superimposed with those obtained by two other algorithms: the unit edge mesher of [70, 71] (in dark grey), and the variational mesher of [3] (in light grey). These two programs, run with our initial surface mesh

[4]The screenshots were obtained using Medit [89].

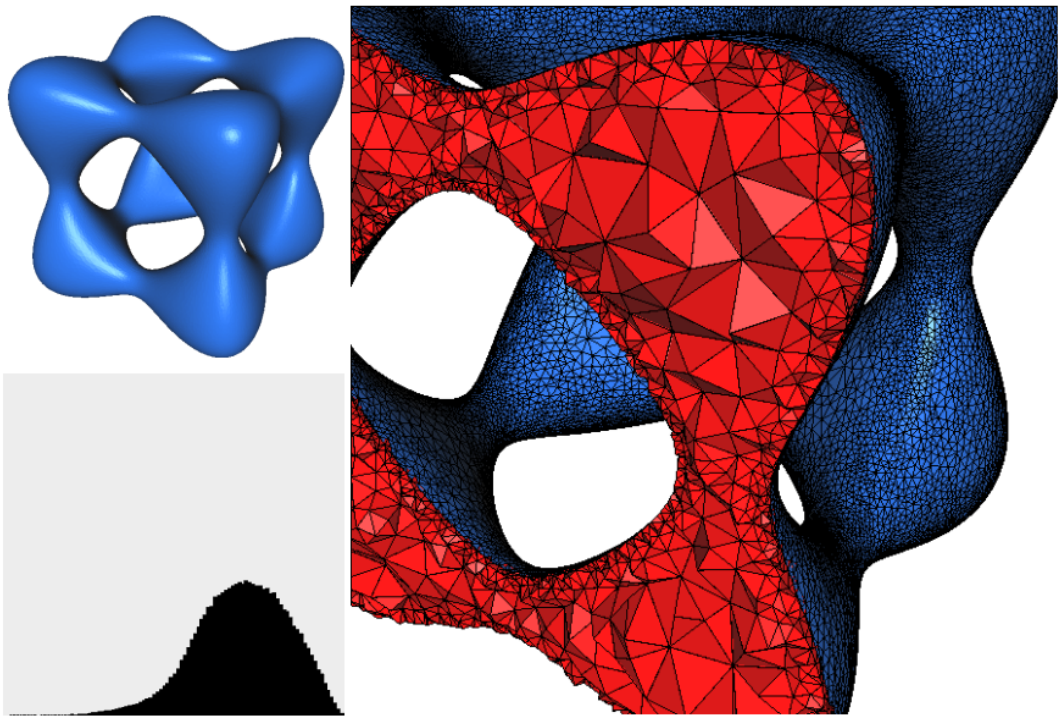

Figure 9.2: Tanglecube model: $57,293$ vertices and $226,010$ tetrahedra.

$\mathrm{Del}_{|S}(E_I)$ ($33,012$ vertices) as input, generated approximately the same number of vertices as our mesher. Their running times, on a Pentium IV at 1.7 GHz, are respectively 10 seconds and 10 minuts. The running time of our algorithm on this example is 20 seconds to insert the $53,762$ vertices lying inside $\mathcal{O}$ and the $2,471$ remaining vertices on S, and 1 minut to exude slivers from the mesh. Compared to the other meshers, our algorithm carries out a good compromise between running time and quality of the output.

In Figure 9.2, the boundary of the domain is the tanglecube surface, introduced in Chapter 6. We used no sizing field inside the domain and $\sigma_0 = 0.09\ \mathrm{d_M}$ on its boundary, as described in Section 9.5.1. The bound ϱ_τ on the radius-edge ratios was set to 2, which enforced the grading of the output mesh. Although the overall appearance of the inversed aspect ratios distribution is deteriorated due to the non-uniformity of the sizing field, the quality of the output mesh remains quite acceptable.

Conclusion

We have presented a simple meshing algorithm and provided evidence that it is generic enough to be of use in a wide variety of practical situations. To apply the algorithm to a new context, it suffices to implement an oracle capable of answering a few basic geometric questions. As an illustration, we have presented various implementations of the algorithm for sampling algebraic surfaces, extracting isosurfaces from 3D medical images, remeshing polyhedra, reconstructing smooth shapes, probing surfaces, and meshing 3-dimensional objects with curved boundaries. Our experiments illustrate the universality of the method and its practical efficiency. In particular, comparing our results and timings to those obtained by the Marching Cubes algorithm, we could emphasize on the advantages of our approach: topological guarantees, approximation of normals, facets with bounded aspect ratio, optimal size.

To prove the correctness of the algorithm, we have used the concept of Delaunay triangulation restricted to a surface or to a volume, and studied its properties under some sampling conditions. We have introduced the notion of loose ε-sample, which turned out to be as powerful as the classical notion of ε-sample, but which is easier to check and to construct. Maybe the most important contribution of this thesis is the concept of k-Lipschitz radius, defined on k-Lipschitz surfaces, which allowed us to prove that (loose) ε-samples provide theoretical guarantees in a more general setting than the traditional smooth case. As a result, the output of our mesher is certified both in the smooth and in the Lipschitz cases.

Prospects

Surface reconstruction We have proved that the restricted Delaunay triangulation of a (loose) ε-sample E of a Lipschitz surface S approximates S both topologically and geometrically. As mentioned in the introduction of the thesis, reconstructing S from E reduces then to finding which facets of the Delaunay triangulation belong to $\mathrm{Del}_{|S}(E)$. In the case where S is unknown, this question is not trivial. It has been elegantly solved in the smooth setting, as reported in [30, 31]. Can any of these methods be extended to the Lipschitz setting? This would lead to the first provably good surface reconstruction algorithm for nonsmooth surfaces.

Noisy data Our work uses a noise-free model, where every point sampled from a surface S lies precisely on S. Such a model is not always pertinent in practice, due to the imperfection of measuring devices. Dey and Goswami [56] introduced a more realistic noise model and proved that a variant of the reconstruction algorithm of [55] can reconstruct S with guarantees. As mentioned in Section 6.5,

the presence of noise leads to a wrong approximation of the normals of S by the Voronoi edges that intersect S. The idea of Dey and Goswami is to approximate the normals of S by Voronoi edges that lie farther from S, in areas where the noise in the input point set E has less impact on the structure of the Voronoi diagram of E. This idea has been carried further by Chazal and Lieutier [35], who proved that, under a more general noise model, a subset of the Voronoi diagram of E remains stable under Hausdorff perturbation. A nice extension of our work would be to combine these ideas with ours, to produce a theory of approximation that holds both in the smooth and Lipschitz settings, and both in the noise-free and noisy cases. Chazal and Lieutier [36] made a step in this direction, by showing that the homology groups of any object $\mathcal{O}$ with positive weak feature size can be retrieved from a sufficiently close Hausdorff approximation of $\mathcal{O}$.

Higher dimensions Our sampling condition can be used for curves in any dimension: extending the proofs is not difficult, and in fact the proofs are simpler. Further research is needed to extend this work to manifolds of dimension more than one embedded in spaces of higher dimensions. Note that most of our arguments extend to higher dimensions in a more or less straightforward manner. However, a problem remains: checking that a simplicial complex $\hat{S}$ is a manifold requires to check its faces of all dimensions. In 3D, this reduces to checking only the edges and vertices of $\hat{S}$. It is no longer the case in higher dimensions, and many pathological cases can occur. This is a serious bottleneck. There is also another issue: when S is a submanifold of $\mathbb{R}^d$ of codimension greater than one, it is a waste of time and space to construct the full-dimensional Delaunay triangulation of a given point sample of S. As a consequence, our meshing algorithm is no longer effective, and therefore other strategies should be considered [72].

Boundaries Our surfaces are assumed to have no boundary. The case of a surface with boundaries as some similarity with the case of a volume bounded by a smooth surface. The approach presented in Chapter 9 can be easily adapted to mesh surfaces with smooth boundaries. The proofs are more technical, due to the fact that the codimension is one and not zero, but they keep the same spirit. The major defect is the smoothness assumption. However, our work on Lipschitz curves/surfaces can probably be used to prove the correctness of the approach on nonsmooth surfaces with nonsmooth boundaries. This would yield a sampling theory for surfaces with boundaries, which would be simpler than existing theories in the smooth case [73], and which would also be relevant in the nonsmooth setting.

Singularities Our theoretical results in the nonsmooth setting assume that the surface S is k-Lipschitz, for some sufficiently small k. Our experiments show that the theoretical bound on k is far below the practicle bound, which seems to lie somewhere around $\sqrt{2}$. When k lies beyond this limit, pathological cases may occur since our sampling method is blind and does not detect singularities. The same observation has been made before in the context of conforming Delaunay meshes [43, 48], where the goal is to conform a Delaunay triangulation to a polyhedron with possibly small angles. The solution proposed then consisted in detecting and then protecting the sharp edges and corners. The same kind of approach could perhaps be used for meshing k-Lipschitz surfaces with large k, or at least piecewise smooth surfaces.

Non-Lipschitz case To conclude, let us emphasize that the non-Lipschitz case is also to be addressed in a near future. The sampling condition of Chazal and Lieutier holds for a larger class of shapes than our sampling condition does, but it provides weaker guarantees. As a consequence, what are the exact trade-offs between providing strong guarantees and providing guarantees for a large class of shapes?

Bibliography

[1] P. K. Agarwal and S. Suri. Surface approximation and geometric partitions. *SIAM J. Comput.*, 27:1016–1035, 1998.

[2] P. D. Alevizos, J.-D. Boissonnat, and M. Yvinec. Non-convex contour reconstruction. *J. Symbolic Comput.*, 10:225–252, 1990.

[3] P. Alliez, D. Cohen-Steiner, M. Desbrun, P. Schröder, and M. Yvinec. Variational tetrahedral meshing. In *Proceedings SIGGRAPH*, 2005.

[4] P. Alliez, D. Cohen-Steiner, O. Devillers, B. Levy, and M. Desbrun. Anisotropic polygonal remeshing. *ACM Transactions on Graphics*, 2003. SIGGRAPH '2003 Conference Proceedings.

[5] P. Alliez, E. Colin de Verdière, O. Devillers, and M. Isenburg. Isotropic surface remeshing. In *International Conference on Shape Modeling and applications*, pages 49–58, 2003.

[6] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.

[7] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. Journal of Comput. Geom. and Applications*, 12:125–141, 2002.

[8] N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory Appl.*, 19:127–153, 2001.

[9] N. Amenta, T. J. Peters, and A. C. Russel. Computational topology: ambient isotopic approximation of 2-manifolds. *Theor. Comput. Sci.*, 305(1–3):3–15, 2003.

[10] D. Attali and J.-D. Boissonnat. A linear bound on the complexity of the delaunay triangulation of points on polyhedral surfaces. *Discrete and Comp. Geometry*, 31:369–384, 2004.

[11] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of the medial axis — a state-of-the-art report. In T. Möller, B. Hamann, and B. Russell, editors, *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Springer-Verlag, 2004. Submitted.

[12] D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces: the smooth case. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 201–210, 2003.

[13] A. Axelsson and A. M^{c}Intosh. Hodge decompositions on weakly Lipschitz domains. *Advances in Analysis and Geometry: New Developments Using Clifford Algebras*, 30(3), 2004.

[14] M. Berger and B. Costiaux. *Differential Geometry: Manifolds, Curves, and Surfaces*, volume 115 of *Graduate Texts in Mathematics Series*. Springer, 1988. 474 pages.

[15] P. Bhattacharya. Efficient neighbor finding algorithms in quadtree and octree. Master's thesis, Indian Institute of Technology, Kampur, India, 2001.

[16] J. Bloomenthal. Polygonization of implicit surfaces. *Comput. Aided Geom. Design*, 5(4):341–355, 1988.

[17] J. Bloomenthal, editor. *Introduction to Implicit Surfaces*, volume 391. Morgan-Kaufmann, 1997.

[18] J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Comput. Geometry: Theory and Applications*, 19:155–173, 2001.

[19] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp. Geometry: Theory and Applications*, pages 185–203, 2002.

[20] J.-D. Boissonnat, D. Cohen-Steiner, and G. Vegter. Isotopic implicit surface meshing. In *Proc. 36th Annu. Sympos. on Theory of Computing*, pages 301–309, 2004.

[21] J.-D. Boissonnat, L. J. Guibas, and S. Oudot. Learning smooth objects by probing. In *Proc. 21st Annu. ACM. Sympos. Comput.Geom.*, pages 198–207, 2005.

[22] J.-D. Boissonnat, L. J. Guibas, and S. Oudot. Learning smooth objects by probing. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 364–365, 2005.
Video available at: `http://compgeom.poly.edu/acmvideos/socg05video/index.html`.

[23] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, September 2005.

[24] J.-D. Boissonnat and Mariette Yvinec. Probing a scene of non-convex polyhedra. *Algorithmica*, 8:321–342, 1992.

[25] C. Boivin and C. Ollivier-Gooch. Guaranteed-quality triangular mesh generation for domains with curved boundaries. *International Journal for Numerical Methods in Engineering*, 55(10):1185–1213, 2002.

[26] H. Borouchaki. Geometric surface mesh. In *Int. Conf. on Integrated and Manufacturing in Mechanical Engineering*, pages 343–350, 1998.

[27] H. Borouchaki, P. L. George, F. Hecht, P. Laug, and E. Saltel. Delaunay mesh generation governed by metric specifications. *Finite Elements in Analysis and Design*, 25:61–83, 1997.

[28] E. Brieskorn and H. Knörrer. *Plane Algebraic Curves*. Birkhaüser, Basel, Boston, Stuttgart, 1986.

[29] M. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, Englewood Cliffs, NJ, 1976.

[30] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction: ideas and algorithms. Research Report 5393, INRIA, November 2004.

[31] F. Cazals, J. Giesen, and M. Yvinec. Delaunay triangulation based surface reconstruction: a short survey. Research Report 5394, INRIA, November 2004.

[32] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Proc. 1st Symp. on Geometry Processing*, pages 177–187, 2003.

[33] The CGAL Library. Release 3.1 (`http://www.cgal.org`).

[34] F. Chazal and A. Lieutier. Weak feature size and persistent homology: Computing homology of solids in $\mathbb{R}^n$ from noisy data samples. Technical Report 378, Institut de Mathématiques de Bourgogne, 2004. Partially published in [36].

[35] F. Chazal and A. Lieutier. The λ-medial axis. *Graphical Models*, 67(4):304–331, July 2005.

[36] F. Chazal and A. Lieutier. Weak feature size and persistent homology: Computing homology of solids in $\mathbb{R}^n$ from noisy data samples. In *Proc. 21st Annual ACM Symposium on Computational Geometry*, pages 255–262, 2005.

[37] Y.-S. Chen and W.-H. Hsu. Parallel thinning algorithm for binary digital patterns. In *Handbook of Pattern Recognition and Computer Vision*, pages pp 458–461. World Scientific, 1993.

[38] H.-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. M. Sullivan. Dynamic skin triangulation. *Discrete and Computational Geometry*, 25:525–568, 2001.

[39] H.-L. Cheng and X. Shi. Guaranteed quality triangulation of molecular skin surfaces. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 481–488, Washington, DC, USA, 2004. IEEE Computer Society.

[40] S.-W. Cheng and T. K. Dey. Quality meshing with weighted delaunay refinement. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on discrete algorithms*, pages 137–146, 2002.

[41] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver exudation. *J. ACM*, 47(5):883–904, 2000.

[42] S.-W. Cheng, T. K. Dey, E. A. Ramos, and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. In *Proc. 20th Annu. ACM Sympos. on Computat. Geom.*, pages 280–289, 2004.

[43] S.-W. Cheng, T. K. Dey, E. A. Ramos, and T. Ray. Quality meshing for polyhedra with small angles. *International J. Comput. Geom. and Applications*, 15(4):421–461, 2005.

[44] E. Chernyaev. Marching cubes 33: Construction of topologically correct isosurfaces. Technical report CN 95-17, CERN, 1995.

[45] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.

[46] F. H. Clarke. Generalized gradients and applications. *Trans. Amer. Math. Soc.*, 205:247–262, 1975.

[47] F. H. Clarke. *Optimization and Nonsmooth Analysis*. Classics in applied mathematics. SIAM, 1990. Reprint.

[48] D. Cohen-Steiner, E. Colin de Verdière, and M. Yvinec. Conforming Delaunay triangulations in 3D. *Computational Geometry: Theory and Applications*, pages 217–233, 2004.

[49] D. Cohen-Steiner and J.-M. Morvan. Restricted delaunay triangulations and normal cycle. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 237–246, 2003.

[50] R. Cole and C. K. Yap. Shape from probing. *J. Algorithms*, 8(1):19–38, March 1987.

[51] H. L. de Cougny and M. S. Shephard. Surface meshing using vertex insertion. In *Proc. 5th International Meshing Roundtable*, pages 243–256, PO Box 5800, MS 0441, Albuquerque, NM, 87185-0441, 1996. Sandia National Laboratories. Also Sand. Report 96-2301.

[52] S. Dey, R. O'Bara, and M. S. Shephard. Curvilinear mesh generation in 3d. In *Proc. 8th Internat. Meshing Roundtable*, pages 407–417, 1999.

[53] T. K. Dey, S. Funke, and E. A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. In *Abstracts 17th European Workshop Comput. Geom.*, pages 129–132. Freie Universität Berlin, 2001.

[54] T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 257–263, 2001.

[55] T. K. Dey and S. Goswami. Tight cocone: a watertight surface reconstructor. *Journal of Computing Informat. Sci. Engin.*, 13:302–307, 2003.

[56] T. K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. In *Proc. 20th Annu. ACM Sympos. Comput. Geom.*, pages 330–339, 2004.

[57] T. K. Dey, G. Li, and T. Ray. Polygonal surface remeshing with Delaunay refinement. In *Proc. 14th Internat. Meshing Roundtable*, 2005.

[58] T. K. Dey and J. Sun. An adaptive MLS surface for reconstruction with guarantees. In *Proc. 3rd Symposium on Geometry Processing*, pages 43–52, 2005.

[59] T. K. Dey and W. Zhao. Approximating medial axis from the voronoi diagram with convergence guarantee. *Algorithmica*, 38(1):179–200, 2004.

[60] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge, 2001.

[61] H. Edelsbrunner and D. Guoy. An experimental study of sliver exudation. *Engineering With Computers, Special Issue on 'Mesh Generation' (10th IMR 2001)*, 18(3):229–240, 2002.

[62] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.

[63] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. on Comp. Geom.*, 7:365–378, 1997.

[64] J. Erickson. Dense point sets have sparse delaunay triangulations: or "...but not too nasty". In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 125–134, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[65] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Discrete and Computational Geometry*, 30(1):109–132, 2003.

[66] H. Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93:418–491, 1959.

[67] H. Federer. *Geometric Measure Theory*. Classics in Mathematics. Springer-Verlag, 1996. Reprint of the 1969 ed.

[68] P. Frey. About surface remeshing. In *Proc. 9th Internat. meshing roundtable*, pages 123–136, 2000.

[69] J. H. G. Fu. Tubular neighborhoods in Euclidean spaces. *Duke Mathematical Journal*, 52(4):1025–1046, 1985.

[70] P.-L. George, F. Hecht, and E. Saltel. Fully automatic mesh generator for 3d domains of any shape. *Impact of Computing in Science and Engineering*, 2:187–218, 1990.

[71] P.-L. George, F. Hecht, and E. Saltel. Automatic mesh generator with specified boundary. *Computer Methods in Applied Mechanics and Engineering*, 92:269–288, 1991.

[72] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds with high co-dimension. *Discrete and Computational Geometry*, 32:245–267, 2004.

[73] M. Gopi. *Theory and practice of sampling and reconstruction for manifolds with boundaries*. Doctoral thesis, University of North Carolina, at Chapel Hill, 2001.

[74] K. Grove. Critical point theory for distance functions. In *Proc. of Symposia in Pure Mathematics*, volume 54, 1993. Part 3.

[75] P. M. Gruber. Approximation of convex bodies. In Peter M. Gruber and J. M. Wills, editors, *Convexity and its Applications*, pages 131–162. Birkhäuser, Basel, Switzerland, 1983.

[76] J. C. Hart and B. T. Stander. Guaranteeing the topology of an implicit surface polygonization for interactive modeling. In *Proc. SIGGRAPH*, pages 279–286, 1997.

[77] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, NY, 1976.

[78] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 26(2):71–78, 1992. Proc. SIGGRAPH '92.

[79] R. Kolluri. Provably good moving least squares. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms*, pages 1008–1017, 2005.

[80] Nico Kruithof and Gert Vegter. Meshing skin surfaces with certified topology. Technical Report ECG-TR-364100-02, Rijksuniversiteit Groningen, 2004.

[81] P. Laug and H. Borouchaki. Molecular surface modeling and meshing. In *Proc. 10th International Meshing Roundtable*, pages 31–41, 2001.

[82] P. Laug, H. Bourouchaki, and P.-L. George. Maillage de courbes gouverné par une carte de métriques. Technical Report RR-2818, INRIA Rocquencourt, 1996.

[83] D. Levin. Mesh-independent surface interpolation. In G. Brunett, B. Hamann, K. Mueller, and L. Linsen, editors, *Geometric Modeling for Scientific Visualization*. Springer-Verlag, 2003.

[84] X.-Y. Li and S.-H. Teng. Generating well-shaped delaunay meshed in 3d. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 28–37. Society for Industrial and Applied Mathematics, 2001.

[85] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as it medial axis. *Computer-Aided Design*, 36(11):1029–1046, September 2004.

[86] M. Lindenbaum and A. M. Bruckstein. Blind approximation of planar convex sets. *IEEE Trans. Robot. Autom.*, 10(4):517–529, August 1994.

[87] W. Lorensen and H. Cline. Marching cubes: a high resolution 3d surface construction algorithm. *Comput. Graph.*, 21(4):163–170, 1987.

[88] D. Mavriplis. Cfd in aerospace in the new millenium. *Canadian Aeronautics and Space Journal*, 46(4):167–176, 2000.

[89] Medit, a scientific visualization tool (`http://www-rocq.inria.fr/gamma/medit/medit.html`).

[90] G. L. Miller, D. Talmor, and S.-H. Teng. Data generation for geometric algorithms on non uniform distributions. *International Journal of Computational Geomety and Applications*, 9(6):577–599, 1999.

[91] F. Morgan. *Geometric Measure Theory: a Beginner's Guide*. Academic Press, 1988.

[92] J.-M. Morvan and B. Thibert. On the approximation of the area of a surface. Rapport de recherche 4375, INRIA, 2002.

[93] F. Murat and J. Simon. Sur le contrôle par un domaine géométrique. Pré-publication du Laboratoire d'Analyse Numérique 74015, University of Paris VI, 1976.

[94] J. Nečas. *Les méthodes directes en théorie des équations elliptiques*. Masson, 1967.

[95] P. P. Pebay and T. J. Baker. Comparison of triangle quality measures. In *Proc. 10th Internat. Meshing Roundtable*, pages 327–340, 2001.

[96] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *Proc. 2nd Sympos. on Geometry Processing*, pages 251–260, 2004.

[97] T. J. Richardson. Approximation of planar convex sets from hyperplanes probes. *Discrete and Computational Geometry*, 18:151–177, 1997.

[98] G. Rote. The convergence rate of the Sandwich algorithm for approximating convex functions. *Computing*, 48:337–361, 1992.

[99] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18:548–585, 1995.

[100] V. V. Savchenko, A. A. Pasko, O. G. Okunev, and T. L. Kunii. Function representation of solids reconstructed from scattered surface points and contours. *Computer Graphics Forum*, 14(4):181–188, 1995.

[101] A. Scheffer and A. Ungor. Efficient adaptive meshing of parametric models. In *Proc. 6th ACM Sympos. Solid Modeling and Applications*, pages 59–70, 2001.

[102] J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. In *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, pages 86–95, 1998.

[103] J. R. Shewchuk. What is a good linear element? Interpolation, conditioning, and quality measures. In *Proc. 11th Internat. Meshing Roundtable*, pages 115–126, 2002.

[104] K. Shimada and D. C. Gossard. Automatic triangular mesh generation of trimmed parametric surfaces for finite element analysis. *Comput. Aided Geom. Des.*, 15(3), 1998.

[105] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnet, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.

[106] S. S. Skiena. Geometric reconstruction problems. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 26, pages 481–490. CRC Press LLC, Boca Raton, FL, 1997.

[107] The SYNAPS library. `http://www-sop.inria.fr/galaad/logiciels/synaps/index.html`.

[108] S.-H. Teng and C. W. Wong. Unstructured mesh generation: Theory, practice, and perspectives. *Int. J. Comput. Geometry Appl.*, 10(3):227–266, 2000.

[109] J. F. Traub, G. W. Wasilkowski, and H. Wozniakowski. *Information-based Complexity*. Academic Press, 1988.

[110] J. F. Traub and A. G. Werschulz. *Complexity and Information*. Cambridge University Press, 1998.

[111] P. Trébuchet. *Vers une résolution stable et rapide des équations algébriques*. Doctoral thesis, University of Paris VI, 2002.

[112] J. R. Tristano, S. J. Owen, and S. A. Canann. Advancing front surface mesh generation in parametric space using a Riemannian surface definition. In *Proc. 7th international meshing reoundtable*, pages 429–445, 1998.

[113] L. Velho. Simple and efficient polygonization of implicit surfaces. *Journal of Graphics Tools*, 1(2):5–24, 1996.

[114] L. Velho, J. Gomes, and L. H. de Figueiredo. *Implicit objects in Computer Graphics*. Springer-Verlag, 2002.

Abstract In the last decade, a great deal of work has been devoted to the elaboration of a sampling theory for smooth surfaces. The goal was to work out sampling conditions that ensure a good reconstruction of a given smooth surface S from a finite subset E of S. Among these conditions, a prominent one is the ε-sampling condition of Amenta and Bern, which states that every point p of S is closer to E than ε times $\mathrm{d}_{\mathrm{M}}(p)$, where $\mathrm{d}_{\mathrm{M}}(p)$ is the distance from p to the medial axis of S. Amenta and Bern proved that it is possible to extract from the Delaunay triangulation of E a PL surface that approximates S both topologically and geometrically.

Nevertheless, the important issues of checking whether a given point set is an ε-sample, and constructing ε-samples of a given smooth surface, have never been addressed. Moreover, the sampling conditions proposed so far offer guarantees only in the smooth setting, since d_{M} vanishes at points where the surface is not differentiable.

In this thesis, we introduce the concept of loose ε-sample, which can be viewed as a weak version of the notion of ε-sample. The main advantage of loose ε-samples over ε-samples is that they are easier to check and to construct. Indeed, checking that a finite set of points is a loose ε-sample reduces to checking whether a finite number of spheres have small enough radii. When the surface S is smooth, we prove that, for sufficiently small ε, ε-samples are loose ε-samples and vice-versa. As a consequence, loose ε-samples offer the same topological and geometric guarantees as ε-samples.

We further extend our results to the nonsmooth case by introducing a new measurable quantity, called the Lipschitz radius, which plays a role similar to that of d_{M} in the smooth setting, but which turns out to be well-defined and positive on a much larger class of shapes. Specifically, it characterizes the class of Lipschitz surfaces, which includes in particular all piecewise smooth surfaces such that the normal variation around singular points is not too large. Our main result is that, if S is a Lipschitz surface and E is a sample of S such that any point p of S is at distance less than a fraction of the Lipschitz radius of S, then we obtain similar guarantees as in the smooth setting. More precisely, we show that the Delaunay triangulation of E restricted to S is a 2-manifold isotopic to S lying at Hausdorff distance $O(\varepsilon)$ from S, provided that its facets are not too skinny. We also extend our previous results on loose samples. Furthermore, we are able to give tight bounds on the size of such samples.

To show the practicality of the concept of loose ε-sample, we present a simple algorithm that constructs provably good surface meshes. Given a compact Lipschitz surface S without boundary and a positive parameter ε, the algorithm generates a sparse loose ε-sample E and at the same time a triangular mesh extracted from the Delaunay triangulation of E. Taking advantage of our theoretical results on loose ε-samples, we can guarantee that this triangular mesh is a good topological and geometric approximation of S, under mild assumptions on the input parameter ε. A noticeable feature of the algorithm is that the input surface S needs only to be known through an oracle that, given a line segment, detects whether the segment intersects the surface and, in the affirmative, returns the intersection points. This makes the algorithm useful in a wide variety of contexts and for a large class of shapes. We illustrate the genericity of the approach through a series of applications: implicit surface meshing, polygonal surface remeshing, unknown surface probing, and volume meshing.

Résumé Cette dernière décennie a vu apparaître et se développer toute une théorie sur l'échantillonnage des surfaces lisses. L'objectif était de trouver des conditions d'échantillonnage qui assurent une bonne reconstruction d'une surface lisse S à partir d'un sous-ensemble fini E de points de S. Parmi ces conditions, l'une des plus importantes est sans conteste la condition d'ε-échantillonnage, introduite par Amenta et Bern, qui stipule que tout point p de S doit être à distance de E au plus ε fois $\mathrm{d_M}(p)$, où $\mathrm{d_M}(p)$ désigne la distance de p à l'axe médian de S. Amenta et Bern ont montré qu'il est possible d'extraire de la triangulation de Delaunay d'un ε-échantillon E une surface affine par morceaux qui approxime S du point de vue topologique (isotopie) et géométrique (distance de Hausdorff).

Néanmoins restaient ouvertes les questions cruciales de vérifier si un ensemble de points donné est un ε-échantillon d'une part, et de construire des ε-échantillons d'une surface lisse donnée d'autre part. De plus, les conditions d'échantillonnage proposées jusque là n'offraient des garanties que dans le cas lisse, puisque $\mathrm{d_M}$ s'annule aux points où la surface n'est pas différentiable.

Dans cette thèse, nous introduisons le concept d'ε-échantillon lâche, qui peut être vu comme une version faible de la notion d'ε-échantillon. L'avantage majeur des ε-échantillons lâches sur les ε-échantillons classiques est qu'ils sont plus faciles a vérifier et à construire. Plus précisément, vérifier si un ensemble fini de points est un ε-échantillon lâche revient à regarder si les rayons d'un nombre fini de boules sont suffisamment petits. Quand la surface S est lisse, nous montrons que les ε-échantillons sont des ε-échantillons lâches et réciproquement, à condition que ε soit suffisamment petit. Il s'ensuit que les ε-échantillons lâches offrent les mêmes garanties topologiques et géométriques que les ε-échantillons.

Nous étendons ensuite nos résultats au cas où la surface échantillonnée est non lisse en introduisant une nouvelle grandeur, appelée rayon Lipschitzien, qui joue un rôle similaire à $\mathrm{d_M}$ dans le cas lisse, mais qui s'avère être bien défini et positif sur une plus large classe d'objets. Plus précisément, il caractérise la classe des surfaces Lipschitziennes, qui inclut entre autres toutes les surfaces lisses par morceaux pour lesquelles la variation des normales aux abords des points singuliers n'est pas trop forte. Notre résultat principal est que, si S est une surface Lipschitzienne et E un ensemble fini de points de S tel que tout point de S est à distance de E au plus une fraction du rayon Lipschitzien de S, alors nous obtenons le même type de garanties que dans le cas lisse, à savoir : la triangulation de Delaunay de E restreinte à S est une variété isotope à S et à distance de Hausdorff $O(\varepsilon)$ de S, à condition que ses facettes ne soient pas trop aplaties. Nous étendons également ce résultat aux échantillons lâches. Enfin, nous donnons des bornes optimales sur la taille de ces échantillons.

Afin de montrer l'intérêt pratique des échantillons lâches, nous présentons ensuite un algorithme très simple capable de construire des maillages certifiés de surfaces. Étant donné une surface S compacte, Lipschitzienne et sans bord, et un paramètre positif ε, l'algorithme génère un ε-échantillon lâche E de S de taille optimale, ainsi qu'un maillage triangulaire extrait de la triangulation de Delaunay de E. Grâce à nos résultats théoriques, nous pouvons garantir que ce maillage triangulaire est une bonne approximation de S, tant sur le plan topologique que géométrique, et ce sous des hypothèses raisonnables sur le paramètre d'entrée ε. Un aspect remarquable de l'algorithme est que S n'a besoin d'être connue qu'à travers un oracle capable de détecter les points d'intersection de n'importe quel segment avec la surface. Ceci rend l'algorithme assez générique pour être utilisé dans de nombreux contextes pratiques et sur une large gamme de surfaces. Nous illustrons cette généricité à travers une série d'applications : maillage de surfaces implicites, remaillage de polyèdres, sondage de surfaces inconnues, maillage de volumes.