

Voronoi Diagrams, Triangulations and Surfaces

Jean-Daniel Boissonnat

19 septembre 2006

Table des matières

1	Introduction	7
2	Convex hulls	9
2.1	Simplices and complexes	9
2.2	Convex hulls	10
2.2.1	The planar case	10
2.2.2	Convex hulls in \mathbb{R}^3	10
2.2.3	In higher dimensions	11
2.3	Space of lines and duality	12
2.4	Combinatorial bounds	13
2.4.1	Convex hulls in \mathbb{R}^3	13
2.4.2	Beyond the third dimension	15
2.5	Algorithms	15
2.5.1	An incremental algorithm	15
2.5.2	Can we do better?	16
2.5.3	Randomisation helps	16
2.6	Bibliographical notes	18
3	Voronoi diagrams and Delaunay triangulations	19
3.1	Voronoi diagrams	19
3.2	The space of spheres	20

3.3	Delaunay triangulation	22
3.4	Properties of the Delaunay triangulation	25
3.5	Bibliographical notes	27
4	Laguerre geometry	29
4.1	Laguerre diagrams	29
4.2	Properties of Laguerre diagrams	31
4.3	Affine diagrams	33
4.4	Bibliographical notes	34
5	Delaunay triangulation restricted to a surface	35
5.1	Approximation criteria	35
5.1.1	Topological equivalence	35
5.1.2	Distance between two sets	36
5.2	Restricted Delaunay triangulation	37
5.3	Sampling conditions	38
5.4	Approximation of smooth surfaces	40
5.4.1	Approximation of a curve	41
5.4.2	Approximation of a smooth surface : local lemmas	42
5.4.3	$\text{DelS}(E)$ is a closed triangulated surface	46
5.4.4	$\text{DelS}(E)$ and S are isotopic	47
5.4.5	Loose ε -samples	51
5.5	Bibliographical notes	54
6	Surface meshing	55
6.1	Algorithm	55
6.2	Termination	56
6.3	Optimality	57
6.4	Experimental results	58

6.5	Bibliographical notes	60
7	Complexity of Voronoi diagrams of surface samples	61
7.1	Introduction	61
7.2	Definitions and notations	62
7.2.1	Notations	62
7.2.2	Polyhedral surfaces	62
7.2.3	Sample	63
7.3	Preliminary results	63
7.4	Counting Delaunay edges	66
7.4.1	Delaunay edges with both endpoints in the ε -regular zone	67
7.4.2	Delaunay edges with both endpoints in the ε -singular zone	68
7.4.3	Delaunay edges joining the ε -regular and the ε -singular zones	69
7.4.4	Main result	72
7.5	Bibliographical notes	73
8	Local systems of coordinates	
	Scattered data interpolation	75
8.1	Voisins naturels et systèmes de coordonnées associés	75
8.1.1	Gradient des coordonnées naturelles	79
8.1.2	Support des coordonnées naturelles	81
8.2	Interpolation de données non structurées	82
8.2.1	Reconstruction exacte des fonctions affines	82
8.2.2	Reconstruction exacte des fonctions quadratiques	82
8.3	Systèmes de coordonnées sur une surface	84
8.3.1	Voisins naturels sur S	84
8.3.2	Coordonnées naturelles sur S	85
8.4	Notes bibliographiques	86

Chapitre 1

Introduction

These notes intend to provide an introduction to the computational geometry of surfaces. The main motivation is to provide effective algorithms to construct piece-wise (PL) approximations of surfaces. Instead of surveying the many approaches that have been proposed in the literature, we focus on provably correct methods and, among them, on methods based on a fundamental data structure that has been extensively studied in computational geometry, the Voronoi diagram of a finite set of points. Given n points (also called sites) of \mathbb{R}^d , the associated Voronoi diagram subdivides \mathbb{R}^d into regions, each region consisting of the points of \mathbb{R}^d that are closer to one of the sites.

In a first part (sections 2-4), we introduce Voronoi diagrams and their dual Delaunay triangulations. We establish a connexion between Voronoi diagrams and polytopes that allows to derive tight combinatorial bounds and efficient algorithms to construct Voronoi diagrams. This first part is a (very brief) introduction to computational geometry. The interested reader will find more material in the textbooks devoted to the subject [12, 13, 19, 24].

In a second part (sections 5-7.5), we show how Voronoi diagrams and Delaunay triangulations can be used to sample and approximate a surface. The main tool is the concept of Delaunay triangulation restricted to a surface. Under some sample conditions, we will show that the restricted Delaunay triangulation of a surface S is a good approximation of S , both in a topological and in a geometric sense. We will present and analyze an algorithm to compute such an approximation, and present some results.

Notations : We identify a point $x \in \mathbb{R}^d$ and the vector of its coordinates. We note $x \cdot y$ the dot product and $x^2 = x \cdot x = \|x\|^2$ the squared euclidean norm of x . $B(c, r)$ denotes the ball of center c and radius r , $\text{aff}(f)$ the affine hull of f .

Remark : The latest version of these notes can be downloaded

`ftp ://ftp-sop.inria.fr/geometrica/boissonnat/notes_de_cours.pdf`

All comments are welcome and can be addressed to `Jean-Daniel.Boissonnat@sophia.inria.fr`

Chapitre 2

Convex hulls

2.1 Simplices and complexes

A linear combination of $k + 1$ points $\{p_0, \dots, p_k\}$ is a weighted sum of points

$$\sum_{i=0}^k \lambda_i p_i, \quad \text{with} \quad \sum_{i=0}^k \lambda_i = 1.$$

We say that the $k + 1$ points are *affinely independent* if the space generated by the affine combinations of these points has dimension k .

We call *k-simplex* of \mathbb{R}^d the convex hull of $k + 1$ points that are affinely independent, i.e. the set of points satisfying

$$\sum_{i=0}^k \lambda_i p_i, \quad \text{with} \quad \sum_{i=0}^k \lambda_i = 1 \quad \text{and} \quad \lambda_i \geq 0.$$

A k -simplex s is a simplex such that the smallest affine subspace that contains s is of dimension k . A 0-simplex is a point, a 1-simplex is a line segment, a 2-simplex is a triangle, a 3-simplex is a tetrahedron. The *faces* of a simplex are simplices of lower dimensions. The k -faces of a l -simplex s , $l > k$, are obtained by taking the convex hull of all subsets of $k + 1$ vertices of s . We note $g \prec f$ if g is a face of f .

If the points of a finite point set E are in *general position*, meaning that no subset of $k + 2$ points belong to an affine space of dimension k , for $k = 0, \dots, d - 1$, all the faces of $\text{conv}(E)$ are simplices.

A *simplicial complex* K is a finite set of simplices such that

1. if $f \in K$ and $g \prec f$, then $g \in K$
2. if $f, g \in K$ and $f \cap g = h \neq \emptyset$, then $h \prec f$ and $h \prec g$.

The dimension of K is the maximal dimension of its faces. If the dimension of K is k , we say that K is a k -complex.

A *cell complex* C is defined as a simplicial complex except that the faces may be general convex polyhedra that are not restricted to be simplices.

2.2 Convex hulls

The convex hull $\text{conv}(E)$ of a finite set of points E is the smallest convex set (for the inclusion relation) that contains E . $\text{conv}(E)$ can alternatively be defined as the intersection of a finite number of half spaces. We say that a hyperplane h *supports* $\text{conv}(E)$ if $h \cap E \neq \emptyset$ and if E belongs to one of the two half spaces defined by h . We note h^- this half-space. $\text{conv}(E)$ is the intersection of all such half-spaces h^- . The boundary of $\text{conv}(E)$ is a cell complex whose faces are the convex hulls of $E \cap h$ for all supporting hyperplanes h . The vertices of $\text{conv}(E)$ are the points of E . When the points are in general position, all the faces of $\text{conv}(E)$ are simplices. This will be assumed in the sequel.

2.2.1 The planar case

$\text{conv}(E)$ will be represented by a circular list L whose elements are the vertices of $\text{conv}(E)$ sorted in counterclockwise order. The numerical operations that will be used are comparing two numbers, and computing the orientation of a triangle, which reduces to evaluating the sign of a wedge product

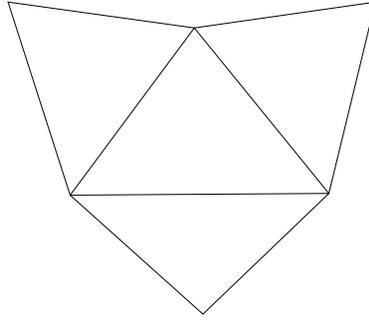
$$\text{orient}(p_i p_j p_k) = (p_i - p_j) \wedge (p_i - p_k) = \begin{vmatrix} 1 & 1 & 1 \\ p_i & p_j & p_k \end{vmatrix} = \begin{vmatrix} 1 & 1 & 1 \\ x_i & x_j & x_k \\ y_i & y_j & y_k \end{vmatrix}.$$

$\text{orient}(p_i p_j p_k)$ is 0 when p_i , p_j and p_k are colinear, positive if the triangle $p_i p_j p_k$ is oriented positively (counterclockwise), negative if $p_i p_j p_k$ is oriented clockwise.

Observe that the output of the numerical operations above belongs to the discrete set $\{-, 0, +\}$. We call them *predicates*. The orientation predicate is mandatory. Comparisons are not necessary but make things easier.

2.2.2 Convex hulls in \mathbb{R}^3

A vertex s of $\text{conv}(E)$ is represented by a pointer to the corresponding point $\text{point}(s)$. Each vertex points to one of its incident facets. Each facet of $\text{conv}(E)$ has *three* pointers to its 3 vertices p_1, \dots, p_3 and *three* pointers to its *three* incident facets, f_1, \dots, f_3 . By convention, f_i is the facet that has not p_i as a vertex.



The important predicate is still the orientation predicate

$$\text{orient}(p_i, p_j, p_k, p_l) = \begin{vmatrix} 1 & 1 & 1 & 1 \\ x_i & x_j & x_k & x_l \\ y_i & y_j & y_k & y_l \\ z_i & z_j & z_k & z_l \end{vmatrix} = ((p_j - p_i) \wedge (p_k - p_i)) \cdot (p_l - p_i).$$

If $\text{orient}(p_i, p_j, p_k, p_l) = 0$, the points p_i, p_j, p_k, p_l are coplanar.

We require, in the sequel, that the convex hull is *oriented*, meaning that all its facets are oriented negatively. A facet $f = p_i p_j p_k$ is oriented negatively if, for any $p \in E$, $\text{orient}(p_i, p_j, p_k, p) < 0$.

2.2.3 In higher dimensions

A simplex $p_{i_0} \dots p_{i_d}$ is oriented positively, negatively or is degenerate (the affine subspace generated by its vertices has dimension $< d$) whether the determinant of the $(d+1) \times (d+1)$ matrix is positive, negative or zero

$$\text{orient}(p_{i_0}, \dots, p_{i_d}) = \text{sign} \begin{vmatrix} 1 & \dots & 1 \\ p_{i_0} & \dots & p_{i_d} \end{vmatrix} = \text{sign} |p_{i_1} - p_{i_0} \ \dots \ p_{i_d} - p_{i_0}|$$

A geometric interpretation that will be useful is the following. Let us note h the hyperplane passing through $p_{i_0}, \dots, p_{i_{d-1}}$ and $\pi(p)$ the projection of point p onto $x_{i_d} = 0$. If h is not vertical and if

$$\text{orient}(\pi(p_{i_0}), \dots, \pi(p_{i_{d-1}})) > 0,$$

$\text{orient}(p_{i_0}, \dots, p_{i_d})$ is > 0 , < 0 or 0 whether the point p_{i_d} is above, below or in h .

2.3 Space of lines and duality

Let h be a non vertical line of \mathbb{R}^2 of equation $y = ax - b$. We associate to h the *dual* point $h^* = (a, b)$. Conversely, to point $p = (x, y)$, we associate the *dual* line $p^* = \{(a, b) : b = xa - y\}$. Note that p^* consists of the points dual to the lines passing through p .

The mapping $*$

- is *involutive* and thus is bijective : $h^{**} = h$ and $p^{**} = p$
- *preserves incidences* :

$$p = (x, y) \in h \iff y = ax - b \iff b = xa - y \iff h^* \in p^*.$$

Moreover, if $h^+ = \{(x, y) : y > ax - b\}$, the mapping $*$

- *reverses inclusions*

$$p \in h^+ \iff h^* \in p^{*+}.$$

Let us consider a set H of n lines h_1, \dots, h_n and note P the (unbounded) polygon that is the intersection of the half-planes h_i^+ . Each vertex s of P is the intersection point of two lines h_i and h_j , and s lies above all the other lines h_k , $k \neq i, j$. The dual point s^* of s is therefore the line $l_{ij} = (h_i^* h_j^*)$. Moreover, by the inclusion-reversing property of the mapping $*$, no point of $H^* = \{h_1^*, \dots, h_n^*\}$ lie below l_{ij} . It follows that l_{ij} supports the *lower convex hull* of H^* , i.e. the set of the faces of the convex hull of H^* whose supporting hyperplanes are not above any point of H^* .

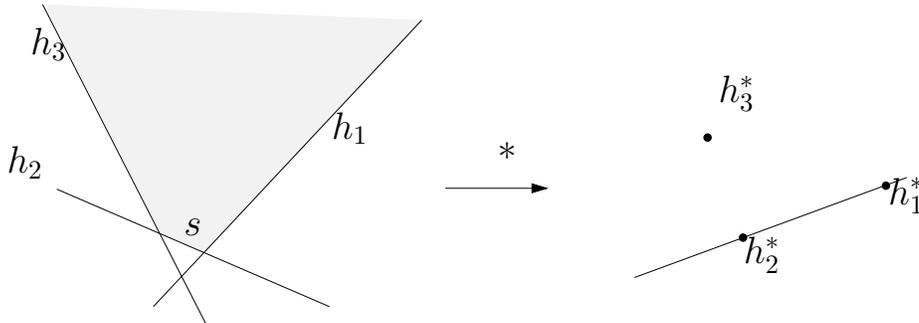


FIG. 2.1 – Point-line duality.

In this way, we reduce the construction of P to the construction of a lower convex hull.

In higher dimensions, we can proceed in a similar way and establish a dual mapping between points and hyperplanes, and therefore between convex hulls and intersections of half-planes.

Let h be a non vertical hyperplane of \mathbb{R}^d of equation $x_d = a \cdot x - b$, where $x = (x_1, \dots, x_{d-1}) \in \mathbb{R}^{d-1}$. We associate to h the *dual* point $h^* = (a, b)$ of \mathbb{R}^d . Conversely, to a point $p = (x, x_d) \in \mathbb{R}^{d-1} \times \mathbb{R}$, we associate the *dual* hyperplane $p^* = \{(a, b) \in \mathbb{R}^{d-1} \times \mathbb{R} : b = x \cdot a - x_d\}$.

Consider now a set of n hyperplanes h_1, \dots, h_n and note P the (unbounded) polyhedron that is the intersection of the half-planes h_i^+ . Let f be a face of P and assume without loss of generality that h_1, \dots, h_{k+1} are the $k+1$ hyperplanes that contain f . If p denotes a point of \mathbb{R}^d , we deduce from the discussion above

$$p \in h_1 \cap \dots \cap h_{k+1} \iff h_1^*, \dots, h_{k+1}^* \in p^*. \quad (2.1)$$

Moreover, if $p \in f$, $p \in h_i^+$ for all $i = 1, \dots, n$ and $p \in \text{int } h_i^+$ for all $k+1 < i \leq n$. Equivalently,

$$\forall i, k+1 < i \leq n, h_i^* \in \text{int } p^{*+}. \quad (2.2)$$

Let $H_{k+1}^* = \{h_i^*\}_{i=1, \dots, k+1}$. We deduce from (2.1) and (2.2) that, if $p \in f$, p^* is a hyperplane that supports the convex hull $\text{conv}(H_{k+1}^*)$ and that $f^* = \text{conv}(H_{k+1}^*)$ is a face of the lower convex hull P^* of H_{k+1}^* . We now associate to each face f of P the corresponding face f^* of P^* : this correspondence is involutive (and therefore bijective) and reverses inclusions. We say that polyhedron P^* is dual to polyhedron P .

Exercise 2.1

Extend the construction of the intersection H of n general half-spaces (not necessarily all above their bounding line) assuming we know a point o in H . (Hint: To the point p , we associate the dual line p^* $p^* = \{\forall x \in p^*, (x - o) \cdot (p - o) = 1\}$. To a line h not passing through o , we associate the dual point h^* : $(h^* - o) \cdot (x - o) = 1, \forall x \in h$. If we denote h^+ the half-space bounded by h not containing o , we can easily adapt the above discussion and get a bijection between the faces of the intersection of the half-spaces and the convex hull of the dual points.)

2.4 Combinatorial bounds

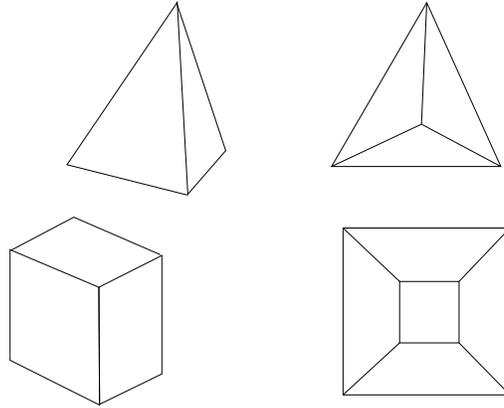
2.4.1 Convex hulls in \mathbb{R}^3

Euler's formula

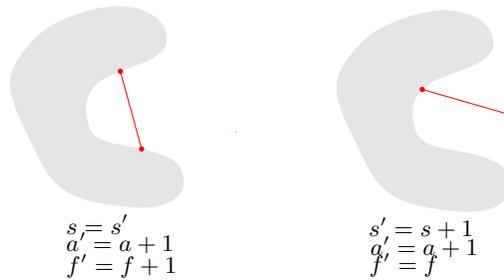
The number of vertices s , edges a and facets f are related by the celebrated *Euler's formula*:

$$s - a + f = 2.$$

We prove the formula for a bounded convex polyhedron P . The Schlegel diagram of P is a planar model obtained when P is seen in perspective from a position just outside the center of one facet. This facet appears as a large polygon with all the remaining faces filling its interior.



We prove Euler's formula for any connected planar map. We build the map from a single vertex by adding edges one-by-one. At each step, the edge to be inserted to the current map C has 1 or vertex in C . In the first case, s does not change while a and f both increase by 1. In the other case, f does not change while s and a both increase by 1. In both cases, $s - a + f$ remains invariant. At the beginning $s = f = 1$ and $a = 0$, and thus $s - a + f = 2$. This value 2 is maintained throughout the whole construction. The formula is therefore true for any connected planar map and, in particular, for any Schlegel diagram, and therefore for any bounded convex polyhedron of \mathbb{R}^3 .



Each edge is incident to two facets and each facet has at least three edges. We thus have

$$2a \geq 3f,$$

with an equality in the case where all facets are triangles. Using this inequality in Euler's formula, we get

$$\begin{aligned} a &\leq 3s - 6 \\ f &\leq 2s - 4 \end{aligned}$$

with an equality in the case where all facets are triangles.

2.4.2 Beyond the third dimension

The combinatorial complexity of a polyhedron defined as the intersection of n half-spaces of \mathbb{R}^d , which is given by the celebrated upper bound theorem..

Theorem 2.1 (Upper Bound Theorem) *The total number of faces of a polytope defined as the intersection of n half-spaces of \mathbb{R}^d is $\Theta\left(n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

Proof. We present a simple proof of the asymptotic version of the upper bound theorem due to Seidel [33]. Let P be the polytope and let p one of its vertices. We assume that all the vertices of P are incident to exactly d edges (which is the general situation). Otherwise, we can slightly perturb the hyperplanes bounding the half-spaces defining P so that the hyperplanes become in general position. During this process, the number of faces of P can only increase. Among the edges that are incident to p (there are d such edges), at least $\lfloor \frac{d}{2} \rfloor$ edges are in the half-space $x_d \geq x_d(p)$ or at least $\lceil \frac{d}{2} \rceil$ edges are in the half-space $x_d \leq x_d(p)$. If $k < d$ edges are incident to p , they belong to a same face of dimension k (k -face for short). Therefore, p is a vertex of extremal x_d -th coordinate for at least one face of dimension $\lfloor \frac{d}{2} \rfloor$. Since any face has at most one vertex of maximal x_d -th coordinate and one vertex of minimal x_d -th coordinate, the number of vertices of P is at most twice the number of $\lfloor \frac{d}{2} \rfloor$ -faces of P .

Under the general position assumption, a k -face is the intersection of $d - k$ of the hyperplanes that define P . We conclude that the number of k -faces is $\binom{n}{d - k} = O(n^{d-k})$, which is $O(n^{\lfloor \frac{d}{2} \rfloor})$ if $k = \lfloor \frac{d}{2} \rfloor$. From the above discussion, we then conclude that the number of vertices of P is $O(n^{\lfloor \frac{d}{2} \rfloor})$.

Under the general position assumption, the number of faces (of any dimension) that are incident to a vertex is bounded by a constant that depends on d but not on n . Hence, the upper bound holds also for the number of faces of any dimension. \square

2.5 Algorithms

2.5.1 An incremental algorithm

We start with the case of a set E of n points in \mathbb{R}^2 .

1. We sort the points by lexicographic order. We therefore have for any two indices $i < j$

$$x(p_i) < x(p_j) \quad \text{or} \quad \begin{cases} x(p_i) = x(p_j) \\ y(p_i) < y(p_j) \end{cases}$$

We represent $\text{conv}(E)$ by two lists that share the same first element p_1 and the same last element p_n . The first one, L^+ , represents the vertices of the *upper* hull, located above p_1p_2 , the other one, L^- , represents the

vertices of the *lower* hull, located below p_1p_2 . We describe only the construction of L^+ , the construction of L^- being entirely symmetrical.

2. L^+ is initialized with p_1 and p_2 .

3. **Loop :**

```

for  $i = 3, \dots, n$ 
  while  $\text{tail}(L^+) \neq \text{head}(L^+)$  and  $\text{orient}(p_i, \text{tail}(L^+), \text{pred}(\text{tail}(L^+))) \leq 0$ 
    remove  $\text{tail}(L^+)$  from  $L^+$ ;
    insert  $p_i$  in  $\text{tail}(L^+)$ .

```

The initial sorting can be done using $O(n \log n)$ comparisons. In step 3, a unique new element is inserted in L^+ . The other elements of L^+ that are considered are, all except the last one, removed from the list and won't be inserted again. The total cost of step 3 is therefore $O(n)$. Observe that the cost of a single iteration can be $O(n)$; nevertheless the total *amortized* cost is also linear.

The overall complexity of the algorithm is therefore $O(n \log n)$. Note that the algorithm executes $O(n \log n)$ comparisons of coordinates and only $O(n)$ orientation tests.

2.5.2 Can we do better ?

It is known that sorting n elements requires at least $\log_2(n!) = \Omega(n \log n)$ comparisons. We deduce that computing the convex hull of n points requires at least $\Omega(n \log n)$ comparisons and orientation tests.

Indeed, take n numbers x_1, \dots, x_n . We associate to each x_i the point $p_i = (x_i, x_i^2)$. The p_i lie on the parabola $y = x^2$. If we know the convex hull of the p_i , we can deduce in linear time the list of the x_i sorted by increasing values. Observe that the orientation test reduces in this case to three comparisons. Indeed,

$$\begin{vmatrix} x_i - x_j & x_i - x_k \\ x_i^2 - x_j^2 & x_i^2 - x_k^2 \end{vmatrix} = (x_i - x_j)(x_j - x_k)(x_k - x_i).$$

As any sorting algorithm must execute $\Omega(n \log n)$ comparisons, we conclude that any convex hull algorithm using comparisons and orientation tests only must execute $\Omega(n \log n)$ such operations.

2.5.3 Randomisation helps

The algorithm of section 2.5.1 can be extended to \mathbb{R}^3 but its complexity cannot remain $O(n \log n)$. In fact, no incremental algorithm can have an $O(n \log n)$ complexity in dimensions 3 and higher. Still, we can have an optimal algorithm provided we introduce some randomness in the algorithm.

The algorithm is incremental and introduces the points one by one. We first describe the case of point set E of \mathbb{R}^2 . Let o be a point in the interior of $\text{conv}(E)$, e.g. the barycenter of three non collinear points. We denote by E_i the subset of the i first points that have been inserted. The algorithm maintains two data structures. First, we maintain the current convex hull $\text{conv}(E_i)$ at step i . Second, we maintain a graph called *conflict graph*. The conflict graph is a bipartite graph that links any point p_j that has not been inserted yet (i.e. $j > i$) to an edge e of the current convex hull $\text{conv}(E_i)$ that is intersected by the line segment op_j . A point and an edge that are joined by an arc in the conflict graph are said to be in conflict.

Consider step $i + 1$. An edge of $\text{conv}(E_i)$ is said to be *red* if its supporting line intersects the line segment op_{i+1} . An edge that is not red is said to be *blue*. If $\text{conv}(E_i)$ is oriented positively, an edge ab is red if and only if $\text{orient}(a, b, p_{i+1}) < 0$. The union of the red edges is a connected portion of the boundary of $\text{conv}(E_i)$. Let s and t be its endpoints. Updating $\text{conv}(E_i)$ consists in replacing the red edges by the two new edges sp_{i+1} and $p_{i+1}t$.

Let e be an edge of $\text{conv}(E_i)$ in conflict with p_{i+1} . As e is intersected by op_{i+1} , e is red and must be removed. To find the other red edges, we walk on the boundary of $\text{conv}(E_i)$ (in both directions) until we meet a blue edge : we thus get s and t , and we can construct $\text{conv}(E_{i+1})$ in time proportional to the number of red edges.

It remains to update the conflict graph. For each point of index $j > i + 1$ whose conflict edge has just been removed, we look for the new edge that is intersected by the half-line $[op_j)$. Let $e' = a'b'$ be that edge. If $\text{orient}(a'b'p_j) \geq 0$, p_j is not a vertex of $\text{conv}(E)$ and is no longer considered. Otherwise, we add an arc in the conflict graph joining p_j and e' .

In higher dimensions, the algorithm is almost identical. The main difference comes from the fact that the boundary of the convex hull is a triangulation and not a polygonal curve. We identify the red facets by traversing the adjacency graph of the facets of $\text{conv}(E_i)$. Let R be the union of the red faces. To update $\text{conv}(E_{i+1})$, we create new faces $f_k = \text{conv}(p_{i+1}, g_k)$ for all faces g_k of the boundary of R , and remove the red faces. The adjacency graph is updated accordingly.

To analyze this algorithm, we assume that the points are inserted in random order.

Let us bound the expected number of facets created at step i . In \mathbb{R}^2 , this number is 2. Consider the situation in \mathbb{R}^d . Assume that E_i is given. A facet of $\text{conv}(E_i)$ has been created at step i if and only if one of its d vertices is p_i , which arises with probability $\frac{d}{i}$. Noting \mathcal{E} the expectation and n_i the number of facets of $\text{conv}(E_i)$ that are created at step i , we have, using the upper bound theorem

$$\mathcal{E}(n_i) = \sum_{f \in \text{conv}(E_i)} \frac{d}{i} = \frac{d}{i} O\left(i^{\lfloor \frac{d}{2} \rfloor}\right) = O(n^{\lfloor \frac{d}{2} \rfloor - 1}).$$

In particular, when $d = 2$ or 3 , $\mathcal{E}(n_i) = O(1)$. This bound does not depend on the choice of E_i and therefore remains valid when averaging on all subsets E_i of E of i elements. Therefore, the bound holds also for the expected number of facets created at step i . By summing over all i and using linearity of expectation, we obtain that the expected total number of created (and therefore also removed) facets is $O(n^{\lfloor \frac{d}{2} \rfloor})$.

We now bound the cost of updating the conflict graph. For any point p_j , $j > i + 1$, a new conflict (p_j, f_j) is created when inserting the $(i + 1)$ -th point if and only if the line segment op_j crosses one of the new facets. The expected number of new conflicts is therefore

$$\sum_{i+1 < j \leq n} \text{proba}(f_j \text{ is new}) = (n - i - 1) \frac{d}{i + 1}.$$

By summing over i and using linearity of expectation, we obtain that the total cost of updating the conflict graph is $O(n \log n + n^{\lfloor \frac{d}{2} \rfloor})$.

Theorem 2.2 *The convex hull of n points of \mathbb{R}^d can be constructed in time $\Theta\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

Using the duality introduced in section 2.3, we immediately deduce the following corollary

Corollary 2.1 *The intersection of n half-spaces of \mathbb{R}^d can be constructed in time $\Theta\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

Exercises 2.1

1. Show that any incremental algorithm that constructs the convex hull of n points of \mathbb{R}^3 takes $\Omega(n^2)$ time in the worst-case.
2. Show that we can remove a point in expected time $O(\log n)$.

2.6 Bibliographical notes

A modern introduction to the theory of polytopes can be found in Ziegler's book [35]. The original proof of the upper bound theorem has been established by McMullen in 1970. The simple asymptotic version given in Theorem 2.1 is due to Seidel [33]. Chazelle [16] has proposed a *deterministic* algorithm to compute the convex hull of a finite point set in any dimension. However, the algorithm is mostly of theoretical interest and no implementation is known.

An excellent introduction to algorithms can be found in the book of Cormen, Leiserson and Rivest [18]. The theory of randomized algorithms is well-developed and finds applications in many areas of computer science. See the book by Motwani and Raghavan for a broad perspective [27]. The CGAL library offers robust and efficient implementation of many geometric algorithms [8, 15].

Chapitre 3

Voronoi diagrams and Delaunay triangulations

Voronoi diagrams are fundamental data structures that have been extensively studied in Computational Geometry.

In this chapter, we introduce Euclidean Voronoi diagrams of points and establish a correspondence between those diagrams and polyhedra in a one dimension higher space. This allows to derive tight bounds on their combinatorial complexity.

3.1 Voronoi diagrams

Let $E = \{p_1, \dots, p_n\}$ be a set of points of \mathbb{R}^d . To each p_i , we associate its *Voronoi region* $V(p_i)$

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

$V(p_i)$ is the intersection of the $n - 1$ half-spaces bounded by the bisector hyperplanes of p_i and each of the other points of E . $V(p_i)$ is therefore a convex polyhedron, possibly unbounded. Observe that $V(p_i)$ contains p_i and therefore is not empty. The collection of the Voronoi regions and their faces, together with their incidence relations, constitute a *cell complex* called the Voronoi diagram of E . Since any point of \mathbb{R}^d belongs to at least one Voronoi region, the Voronoi diagram of E is a subdivision of \mathbb{R}^d .

Definition 3.1 *The Euclidean Voronoi diagram of E , noted $\text{Vor}(E)$, is the cell complex whose cells are the Voronoi regions and their faces.*

In the next subsection, we introduce a useful correspondence between Euclidean Voronoi diagrams of \mathbb{R}^d and a class of polyhedra of \mathbb{R}^{d+1} .

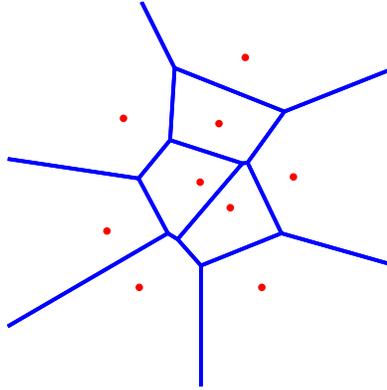


FIG. 3.1 – The Voronoi diagram of a set of 9 points.

3.2 The space of spheres

Let σ be the sphere of equation

$$\sigma(x) = (x - c)^2 - r^2 = x^2 - 2c \cdot x + s = 0,$$

where c is the center of σ , r its radius and $s = \sigma(0) = c^2 - r^2$.

Let

$$\phi : \sigma \in \mathbb{R}^d \longrightarrow \phi(\sigma) = (c, s) \in \mathbb{R}^{d+1}$$

be the bijection that maps a sphere of \mathbb{R}^d to a point of \mathbb{R}^{d+1} . We call \mathbb{R}^{d+1} *the space of spheres*. For convenience, we call *vertical* the $(d + 1)$ -th coordinate axis and $(x, s) \in \mathbb{R}^{d+1}$ is said to be above (x, s') if $s > s'$, and below (x, s') if $s < s'$. We call *vertical projection* the orthogonal projection onto the hyperplane $x_{d+1} = 0$.

Observe that the image by ϕ of a point, considered as a sphere of radius 0, is a point of the paraboloid \mathcal{Q} of \mathbb{R}^{d+1} of equation $x_{d+1} = x^2$.

Let us consider a sphere σ passing through a given point p . We denote by x its center, r its radius and $s = x^2 - r^2$. We have

$$\sigma(p) = p^2 - 2x \cdot p + s = 0.$$

Hence $\phi(\sigma) = (x, s = 2p \cdot x - p^2)$. It follows that the image by ϕ of the spheres that pass through p is the *hyperplane* h_p of \mathbb{R}^{d+1}

$$x_{d+1} = 2p \cdot x - p^2.$$



FIG. 3.2 – A face of $\mathcal{V}(E)$ projects vertically onto a face of $\text{Vor}(E)$.

The intersection of h_p and the paraboloid \mathcal{Q} is reduced to the point $\phi(p) = (p, p^2) : h_p$ is thus tangent to \mathcal{Q} in p .

Consider now a set $E = \{p_1, \dots, p_n\}$ of n points and the associated hyperplanes h_{p_1}, \dots, h_{p_n} tangent to \mathcal{Q} at p_1, \dots, p_n respectively.

If x is closer to p_i than p_j

$$(x - p_i)^2 < (x - p_j)^2 \iff 2p_i \cdot x - p_i^2 > 2p_j \cdot x - p_j^2.$$

In other words, the vertical line passing through x intersects h_{p_i} above h_{p_j} .

Consider now the polyhedron $\mathcal{V}(E) = h_{p_1}^+ \cap \dots \cap h_{p_n}^+$ of \mathbb{R}^{d+1} , where $h_{p_i}^+$ denotes the half-space above h_{p_i} . Each point (x, x_{d+1}) of the facet f_i of $\mathcal{V}(E)$ contained in $h_{p_i}^+$ is above all h_{p_j} . Hence, point x is closer to p_i than to any p_j , which proves that f_i projects vertically onto the Voronoi cell $V(p_i)$ (see figure 3.2).

The faces of the Voronoi diagram $\text{Vor}(E)$ of E are the vertical projections of the faces of the convex polyhedron $\mathcal{V}(E)$.

It follows from the above correspondence that the combinatorial complexity of the Voronoi diagram of n points of \mathbb{R}^d is at most the combinatorial complexity of a polyhedron defined as the intersection of n half-spaces of \mathbb{R}^{d+1} , which is $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ as shown in section 2.4. This bound is tight. In particular, the Voronoi

diagram of n points of \mathbb{R}^3 may be quadratic (see exercise 3.1.5 below).

Exercises 3.1

1. What are the preimages by ϕ of the points of \mathbb{R}^{d+1} that lie
 - (a) above \mathcal{Q} ?
 - (b) on the boundary of $\mathcal{V}(E)$? in the interior of $\mathcal{V}(E)$?
 - (c) on a line?
2. Consider the diagram obtained by projecting the faces of $h_{p_1}^- \cap \cdots \cap h_{p_n}^-$ vertically. Characterize the points that belong to a face of this diagram.
3. Show that if we take points on two non coplanar lines of \mathbb{R}^3 , say $n_1 + 1$ on one of the lines and $n_2 + 1$ on the other, their Voronoi diagram has $n_1 n_2$ vertices.

3.3 Delaunay triangulation

A *triangulation* of a finite set of points E of \mathbb{R}^d is a simplicial complex embedded in \mathbb{R}^d that cover the convex hull of E .

A given set of points admits, in general, many triangulations. The Delaunay triangulation of E , which is defined as the *dual* of the Voronoi diagram of E , is a canonical triangulation that has interesting properties and finds numerous applications.

Two cell complexes V and D are said to be dual if there exists an involutive correspondence between the faces of V and the faces of D that reverses inclusion relationships, i.e. for any two faces f and g of V , their dual faces f^* and g^* of D satisfy $f \subset g \implies g^* \subset f^*$.

Let f be a face of dimension k of the Voronoi diagram of E . All points in the interior of f have the same closest points in E . Let $E_f \subset E$ be the subset of those closest points. If the points are in *general position* (meaning no subset of $d+2$ points of E lie on a same sphere in this chapter), $|E_f| = d - k + 1$. The face dual to f is the convex hull of E_f . The *Delaunay triangulation of E* , noted $\text{Del}(E)$, is the complex consisting of all the dual faces. Under the general position assumption, all the faces of $\text{Del}(E)$ are simplices and $\text{Del}(E)$ is a simplicial complex. The fact that $\text{Del}(E)$ is indeed a triangulation, i.e. a simplicial complex embedded in \mathbb{R}^d and covering the convex hull of E , will follow from the well-known duality between polyhedra in the space of spheres.

In section 3.2, we have associated to the sphere σ of \mathbb{R}^d of equation $\sigma(x) = x^2 - 2c \cdot x + s = 0$ the point $\phi(\sigma) = (c, s)$ of \mathbb{R}^{d+1} . Similarly to what we did in section 2.3, we associate to this point $\phi(\sigma)$ the so-called *polar hyperplane* h_σ of \mathbb{R}^{d+1} of equation $x_{d+1} = 2c \cdot x - s$. Observe that if σ is reduced to a point c , h_σ is identical to the hyperplane h_c tangent to the paraboloid \mathcal{Q} that has been introduced in section 3.2. It should also be observed that the intersection of h_σ with \mathcal{Q} projects vertically onto σ .

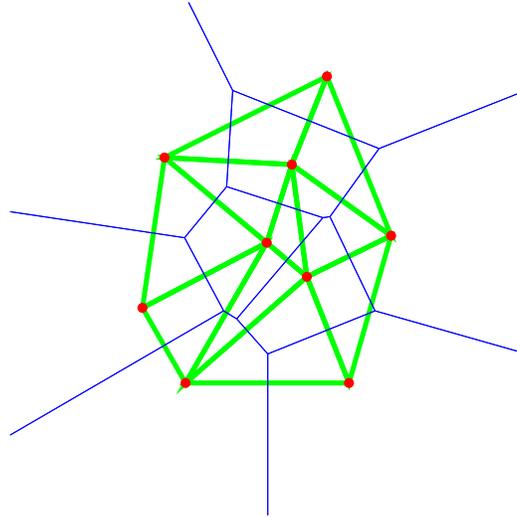


FIG. 3.3 – The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines).

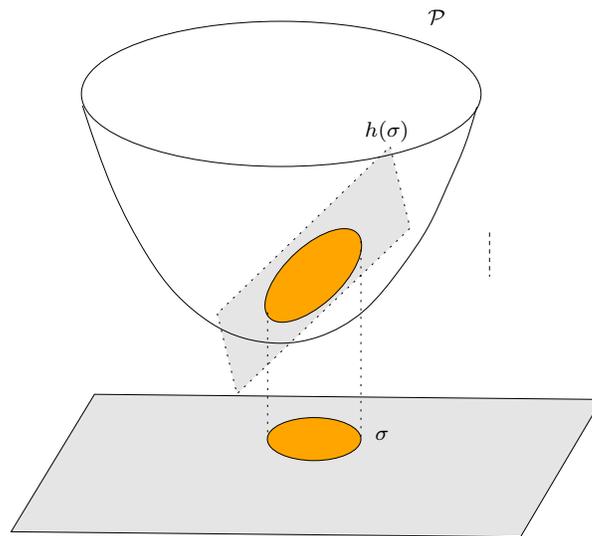


FIG. 3.4 – The polar hyperplane of a sphere.

We deduce the remarkable following property : $x \in \sigma$ if and only if $\phi(x) = (x, x^2) \in h_\sigma$ and σ encloses x if and only if $\phi(x)$ is below h_σ . Indeed

$$\begin{aligned} \sigma(x) = 0 &\iff x^2 = 2c \cdot x - s \iff \phi(x) \in h_\sigma \\ \sigma(x) < 0 &\iff x^2 < 2c \cdot x - s \iff \phi(x) \in \text{int } h_\sigma^-. \end{aligned}$$

Consider now a set $E = \{p_1, \dots, p_n\}$ of n points and let $\mathcal{V}(E)$ denote, as in section 3.2, the polyhedron intersection of the n halfspaces above the n polar hyperplanes h_{p_i} . We define $\mathcal{D}(E)$ as the lower convex hull of the points $\phi(p_i)$, $i = 1, \dots, n$. From section 2.3, we know that $\mathcal{D}(E)$ and $\mathcal{V}(E)$ are dual polyhedra of \mathbb{R}^{d+1} . More precisely, if f is a face of $\mathcal{V}(E)$ that is the intersection of k polar hyperplanes h_1, \dots, h_k , the face f^* dual to f is the face of $\mathcal{D}(E)$ that is the convex hull of h_1^*, \dots, h_k^* . As previously observed, $\pi(f)$ is the k -face of $\text{Vor}(E)$ that is common to the k Voronoi regions $V(p_1), \dots, V(p_k)$. Plainly, $\pi(f^*)$ is the convex hull of p_1, \dots, p_k .

The convex polyhedron $\mathcal{D}(E)$ projects vertically onto a triangulation of E in \mathbb{R}^d . The fact that $\text{Del}(E)$ is properly embedded in \mathbb{R}^d comes from the fact that the vertical projection of $\mathcal{D}(E)$ onto the hyperplane $x_{d+1} = 0$ is 1-1. Moreover, $\text{Del}(E)$ is a subdivision of the convex hull of E since projection preserves convexity.

We deduce that $\mathcal{D}(E)$ projects vertically onto the cell complex dual to $\text{Vor}(E)$, i.e. $\text{Del}(E)$.

$$\begin{array}{ccc} \mathcal{V}(E) = h_{p_1}^+ \cap \dots \cap h_{p_n}^+ & \longleftrightarrow & \mathcal{D}(E) = \text{conv}^-(\phi(E)) \\ \downarrow & & \downarrow \\ \text{Voronoi Diagram } \text{Vor}(E) & \longleftrightarrow & \text{Delaunay Triangulation } \text{Del}(E) \end{array}$$

It follows from the above correspondence that the combinatorial complexity of the Delaunay triangulation of n points is the same as the combinatorial complexity of the dual Voronoi diagram. Moreover, the Delaunay triangulation of n points of \mathbb{R}^d can be deduced from the dual Voronoi diagram or vice versa in time proportional to its size. The Delaunay triangulation of n points of \mathbb{R}^d can be constructed directly by computing the lower convex hull of n points of \mathbb{R}^{d+1} .

If the p_i are in general position, the points $\phi(p_i)$ are affinely independent and all the faces of $\mathcal{D}(E)$ are simplices. It follows that $\text{Del}(E)$ is a triangulation. Otherwise, it is always possible to triangulate the faces that are not simplices. We then obtain a triangulation of E we also call a Delaunay triangulation. Since there are several ways of triangulating the faces, such a Delaunay triangulation of E is no longer unique.

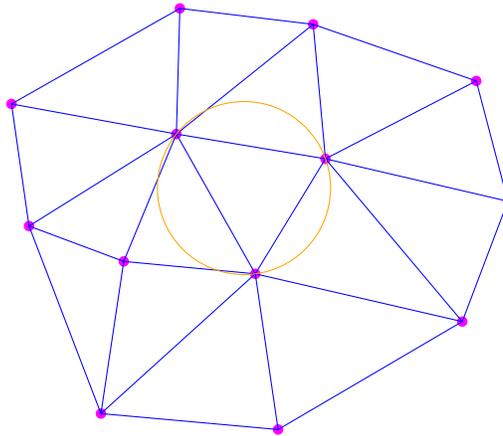
We deduce from theorem 2.2 and corollary ?? the following theorem.

Theorem 3.1 *The Voronoi diagram and the Delaunay triangulation of n points of \mathbb{R}^d can be constructed in time $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.*

3.4 Properties of the Delaunay triangulation

The following lemma can easily be proved either directly or as a consequence of the correspondence between $\text{Del}(E)$ and $\mathcal{D}(E)$.

Lemma 3.1 *A triangulation T of a finite set of points E such that any d -simplex of T has a $(d - 1)$ -sphere circumscribing that does not enclose any point of E is a Delaunay triangulation of E . Any k -simplex with vertices in E that can be circumscribed by a $(d - 1)$ -sphere that does not enclose any point of E is a face of a Delaunay triangulation of E .*



The following lemma provides a *local* characterization of Delaunay triangulations.

Lemma 3.1 *A triangulation $T(E)$ of a finite set of points E whose pairs are all regular is a Delaunay triangulation.*

Proof : Let $\hat{T}(E)$ be the triangulated surface in the space of spheres obtained by lifting the vertices of $T(E)$ by ϕ . The fact that a pair (t_1, t_2) is regular is equivalent to (see section 3.3)

$$\phi(q_2) \in h_{\sigma_1}^+ \quad \text{and} \quad \phi(q_1) \in h_{\sigma_2}^+.$$

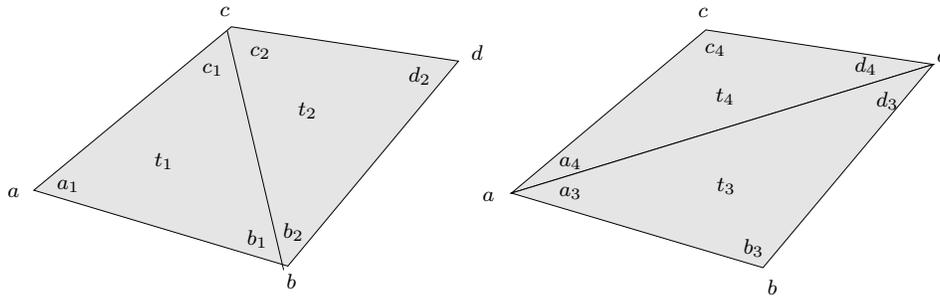
Considering $\hat{T}(E)$ as the graph of a function γ defined on $\text{conv}(E)$, γ is thus locally convex at each point of the interior of its domain of definition which is convex. Hence, γ is convex, $\hat{T}(E)$ is a convex polyhedron which is therefore $\mathcal{D}(E)$. It follows that $T(E)$ is a Delaunay triangulation. \square

The following lemma shows that the Delaunay triangulation of a given set E is the triangulation of E with the largest smallest angle, a useful property in scientific computing where meshes with small angles result

in convergence problems. We define the *angular vector* of a triangulation $T(E)$ of a finite set E of points of the plane as the vector $\text{ang}(T(E)) = (\alpha_1, \dots, \alpha_{3t})$ where the α_i are the angles of the t triangles of $T(E)$ sorted by increasing order.

Lemma 3.2 *Given a finite set E of points in the plane, the triangulation whose angular vector is maximal for the lexicographical order is a Delaunay triangulation of E .*

Proof. Let $T(E)$ be any triangulation of E . Let us consider two adjacent triangles $t_1 = abc$ and $t_2 = bcd$ of $T(E)$. Let Q be their union. If Q is not strictly convex, the pair is regular. If Q is strictly convex, we can retriangulate Q by creating the triangles $t_3 = abd$ and $t_4 = acd$. Let us prove that $\text{ang}(t_1, t_2) > \text{ang}(t_3, t_4)$ if and only if the pair (t_1, t_2) is regular.



If (t_1, t_2) is regular, d does not belong to the disk circumscribing abc and a does not belong to the disk circumscribing bcd . Hence we have, denoting by p_i the angle of vertex p in triangle of index i ,

$$d_3 \leq c_1, \quad d_4 \leq b_1, \quad a_3 \leq c_2, \quad a_4 \leq b_2.$$

As $a_1 = a_3 + a_4$ and $d_2 = d_3 + d_4$, $\text{ang}(t_1, t_2) \geq \text{ang}(t_3, t_4)$. If (t_1, t_2) is not a regular pair, we have

$$d_3 > c_1, \quad d_4 > b_1, \quad a_3 > b_2, \quad a_4 > c_2.$$

With $b_3 = b_1 + b_2$ and $c_4 = c_1 + c_2$, we conclude that $\text{ang}(t_1, t_2) < \text{ang}(t_3, t_4)$.

If the triangulation of E of maximal angular vector, noted $T_M(E)$, was not a Delaunay triangulation, there will be a non regular pair of triangles that could be regularized by the above procedure. As the above procedure increases the angular vector of the triangulation, this contradicts the fact that $T_M(E)$ is optimal. \square

Exercises 3.2

1. Let S be a sphere of \mathbb{R}^d passing through $d + 1$ points p_0, \dots, p_d . Show that a point p_{d+1} of \mathbb{R}^d lies on S , in the interior of the ball B_S bounded by S or outside B_S , depending whether the determinant of the $(d + 2) \times (d + 2)$ matrix

$$\text{in_sphere}(p_0, \dots, p_{d+1}) = \begin{vmatrix} 1 & \dots & 1 \\ p_0 & \dots & p_{d+1} \\ \|p_0\|^2 & \dots & \|p_{d+1}\|^2 \end{vmatrix}$$

is 0, < 0 or > 0 . Show that this predicate is the only numerical operation that is required to check if a triangulation is a Delaunay triangulation.

2. Project vertically the faces of the upper hull $\text{conv}^+(\{\phi(E)\})$. Show that we obtain a triangulation of the vertices of $\text{conv}(E)$ such that each ball circumscribing a simplex contains all the points of E . Define a dual and make a link with question 4 in section 3.2.
3. Prove lemma 3.1.
4. Consider the following algorithm for constructing a Delaunay triangulation of a finite set E of points in the plane. We first compute any triangulation of E and then, we regularize the non regular pairs as indicated in the proof of lemma 3.2. Show that the first step can be done in time $O(n \log n)$ and that the second step in time $O(n^2)$. (Hint : use the fact that regularizing a pair of triangles corresponds, in the space of spheres, to replacing two faces of a tetrahedron by its two other faces.)

3.5 Bibliographical notes

To know more about the space of spheres, one may read the books by Pedoe [29] and Berger [6]. An entire book is devoted to Voronoi diagrams [28]. One may also look at the survey by Aurenhammer and Klein [5] and the part of the textbook by Boissonnat and Yvinec [12, 13] devoted to Voronoi diagrams.

Chapitre 4

Laguerre geometry

We have seen that Euclidean Voronoi diagrams of \mathbb{R}^d are obtained by projecting vertically a polyhedron of \mathbb{R}^{d+1} which is the intersection of finitely many half-spaces. These half-spaces were bounded by hyperplanes tangent to the paraboloid \mathcal{Q} . If we relax the constraint that the hyperplanes are tangent to \mathcal{Q} , we still get a polyhedron whose faces project onto the faces of an affine diagram. Polarity will still provide a dual triangulation (assuming general position). We will see in this section that these diagrams, called *Laguerre diagrams*, have a geometric meaning in the so-called Laguerre geometry where E is replaced by a set of spheres and the Euclidean distance to an element of E is now the power to a sphere. We will also see that all affine diagrams are obtained this way and are in fact Laguerre diagrams.

4.1 Laguerre diagrams

We call *power* of a point to a sphere σ of center c and radius r the real number¹

$$\sigma(x) = (x - c)^2 - r^2.$$

Let $\mathcal{S} = \{\sigma_1, \dots, \sigma_n\}$ be a finite set of spheres of \mathbb{R}^d . We denote by c_i the center of σ_i and r_i its radius. To each σ_i , we associate the region $L(\sigma_i)$ consisting of the points of \mathbb{R}^d whose power to σ_i is not larger than the power to the other spheres of \mathcal{S} :

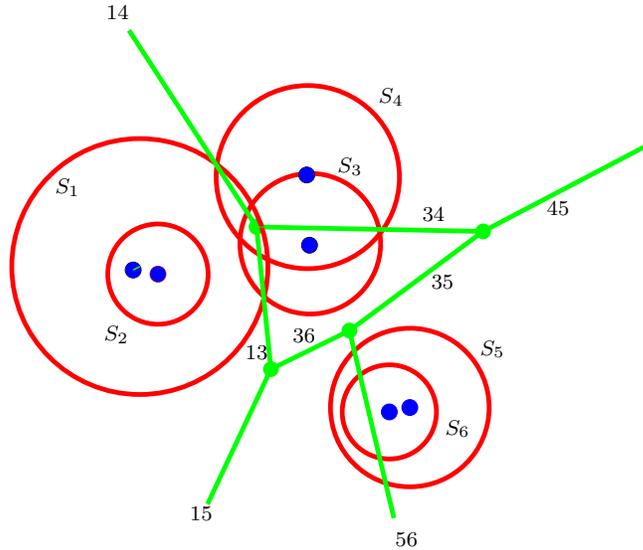
$$L(\sigma_i) = \{x \in \mathbb{R}^d : \sigma_i(x) \leq \sigma_j(x), 1 \leq j \leq n\}.$$

The set of points that have equal power to two spheres σ_i and σ_j is a hyperplane, noted π_{ij} , called *the radical hyperplane* of σ_i and σ_j . π_{ij} is orthogonal to the line joining the centers of σ_i and σ_j . We denote by

¹One may see some similarity with the distance between two events in the relativist space-time : $d = \sqrt{(x - x')^2 + (y - y')^2 + (z - z')^2 - (t - t')^2}$.

π_{ij}^i the half-space bounded by π_{ij} consisting of the points whose power to σ_i is smaller than their power to σ_j . $L(\sigma_i)$ is the intersection of all half-spaces $\pi_{ij}^i, j \neq i$. If this intersection is not empty, it is a convex polyhedron, possibly not bounded. We call *Laguerre regions* the non empty $L(\sigma_i)$.

Definition 4.1 We define the *Laguerre diagram* of \mathcal{S} , noted $\text{Lag}(\mathcal{S})$, as the cell complex whose cells are the Laguerre regions and their faces.



When all spheres have the same radius, their Laguerre diagram is identical to the Voronoi diagram of their centers.

Let's go back to the space of spheres and use the same notations as in section 3.3.

$$\sigma_i(x) < \sigma_j(x) \iff 2c_i \cdot x - s_i > 2c_j \cdot x - s_j$$

can be understood as : the vertical line passing through the point $(x,0)$ intersects hyperplane h_{σ_i} above hyperplane h_{σ_j} .

If we denote by $\mathcal{L}(\mathcal{S})$ the polyhedron $h_{\sigma_1}^+ \cap \dots \cap h_{\sigma_n}^+$, we get that *the faces of the Laguerre diagram are the vertical projections of the faces of $\mathcal{L}(\mathcal{S})$* . Hence, Laguerre diagrams are very similar to Voronoi diagrams : the only difference is that the hyperplanes supporting the faces of $\mathcal{L}(\mathcal{S})$ are not necessarily tangent to the paraboloid \mathcal{Q} and that some hyperplane may not contribute a face. In other words, some sphere σ_i may have an empty Laguerre region.

By proceeding as in section 3.3, we can define a polyhedron $\mathcal{R}(\mathcal{S})$ dual to $\mathcal{L}(\mathcal{S})$. The vertical projections of the faces of $\mathcal{R}(\mathcal{S})$ constitute the faces of a (in general) simplicial complex called a *regular triangulation*.

$$\begin{array}{ccc}
 \mathcal{L}(\mathcal{S}) = h_{\sigma_1}^+ \cap \dots \cap h_{\sigma_n}^+ & \longleftrightarrow & \mathcal{R}(\mathcal{S}) = \text{conv}^-(\phi(\mathcal{S})) \\
 \downarrow & & \downarrow \\
 \text{Laguerre diagram } \text{Lag}(\mathcal{S}) & \longleftrightarrow & \text{Regular triangulation}
 \end{array}$$

It follows, as for Euclidean Voronoi diagrams and Delaunay triangulations, that the combinatorial complexity of the Laguerre diagram of n spheres of \mathbb{R}^d (or equivalently its dual regular triangulation) is $\Theta\left(n^{\lceil \frac{d}{2} \rceil}\right)$ and that such a diagram can be computed in optimal time $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.

4.2 Properties of Laguerre diagrams

We say that a half-space \mathcal{E} is an *upper* half-space if it lies above its bounding hyperplane h and we note $\mathcal{E} = h^+$.

Lemma 4.1 *Let P be the intersection of a finite number of upper half-spaces of \mathbb{R}^{d+1} . If we project the faces of P vertically, we obtain a Laguerre diagram.*

Proof. Let $P = h_1^+ \cap \dots \cap h_n^+$ where hyperplane h_i is defined by $x_{d+1} = 2c_i \cdot x - s_i$. We associate to h_i the sphere of \mathbb{R}^d (possibly imaginary) whose center is c_i and the squared radius $c_i^2 - s_i$. As shown in section 4.1, the faces of the Laguerre diagram are the vertical projections of the faces of $\mathcal{L}(E) = h_1^+ \cap \dots \cap h_n^+$. \square

The following lemma gives a conservation law for flows entering a Laguerre region normally to the facets of the region. This property makes Voronoi and Laguerre diagrams useful when applying finite volume methods in fluid dynamics.

Lemma 4.2 *If $f_{ij}, j \in J$, are the facets of a Laguerre region $L(\sigma_i)$, we have*

$$\sum_{j \in J} \text{vol}(f_{ij}) \frac{c_j - c_i}{\|c_j - c_i\|} = 0.$$

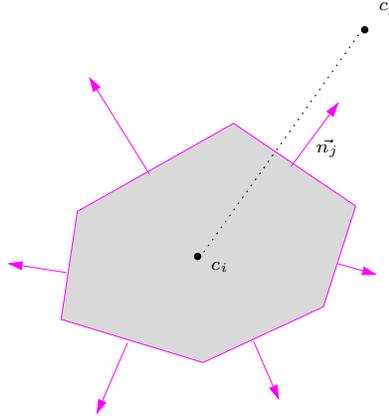
Proof. We first show that if P is a convex polyhedron and $f_j, j \in J$, are its facets, we have

$$\sum_{j \in J} \text{vol}(f_j) \vec{n}_j = 0,$$

where \vec{n}_j is the unit normal vector to f_j (oriented towards the outside of P). Indeed, the volume of P is the sum of the volumes of the pyramids obtained by joining a point x in the interior of P to the facets of P :

$$d \text{ vol}(P) = \sum_{j \in J} \text{vol}(f_j) (x - p_j) \cdot \vec{n}_j.$$

Since the volume of P does not depend on x , the gradient of the right term must be zero, which proves the lemma.



We apply Minkowski's theorem to the Laguerre region $L(\sigma_i)$ and we use the fact that f_{ij} is contained in the radical hyperplane of the spheres σ_i and σ_j , which is orthogonal to the line passing through the centers c_i and c_j of the two spheres. Hence we have $\vec{n}_i = \frac{c_i - c_j}{\|c_i - c_j\|}$. \square

Lemma 4.3 *The intersection of a Voronoi diagram with an affine subspace \mathbb{E} is a Laguerre diagram.*

Proof. Let p'_i denote the projection of p_i onto \mathbb{E} and let $d_i = \|p_i - p'_i\|$

$$\forall x \in \mathbb{E} : \|x - p_i\|^2 < \|x - p_j\|^2 \iff \|x - p'_i\|^2 + d_i^2 < \|x - p'_j\|^2 + d_j^2 \iff \sigma_i(x) < \sigma_j(x),$$

where σ_i is the (imaginary) sphere centered at p'_i whose squared radius is $r^2 = -d_i^2$ \square

We equip \mathbb{R}^d with a non Euclidean metric defined by

$$\|x - y\|_Q = (x - y) Q (x - y)^t,$$

where Q is a symmetric matrix, i.e. $Q = Q^t$. We can now adapt the definition of Euclidean Voronoi diagrams and define the Voronoi diagram of a finite set of points for this (non Euclidean) metric.

Lemma 4.4 *The Voronoi diagram of a set of n points $E = \{p_1, \dots, p_n\}$ for the distance $\|\cdot\|_Q$ is the Laguerre diagram of n spheres $\sigma_1, \dots, \sigma_n$ of \mathbb{R}^d . The center c_i of σ_i is the point $p_i Q$ and its squared radius is $r_i^2 = c_i \cdot c_i - p_i Q p_i^t$.*

Proof.

$$\|x - p_i\|_Q < \|x - p_j\|_Q \iff -2p_i Q x^t + p_i Q p_i^t < -2p_j Q x^t + p_j Q p_j^t \iff \sigma_i(x) < \sigma_j(x).$$

\square

4.3 Affine diagrams

Both Voronoi diagrams and Laguerre diagrams are cell complexes whose cells are convex polyhedra. Let us consider a more general setting where \mathcal{O} is a finite set of objects and δ is a continuous function between a point of \mathbb{R}^d and an object (we do not assume δ to be a distance and, in particular, we do not require that it satisfies the triangular inequality). Assume that the bisector between any two objects is a hyperplane. We can then define the diagram of \mathcal{O} under the distance δ in the way we defined Voronoi and Laguerre diagrams. Such a diagram is called an *affine diagram*. Voronoi and Laguerre diagrams are special cases of affine diagrams. Lemmas 4.3 and 4.4 have provided other examples of such diagrams. As stated in those lemmas, each of these diagrams is identical to some Laguerre diagram. In fact, the following theorem, due to Aurenhammer [4], states that all affine diagrams are Laguerre diagrams.

Theorem 4.1 *Let \mathcal{O} be a set of n objects o_1, \dots, o_n , and let $\delta(x, o_i)$ be a function that measures the distance from any point x of \mathbb{R}^d to object i , $i = 1, \dots, n$. Assume that the set of points of \mathbb{R}^d that are closer to o_i than to o_j , is a given halfspace of \mathbb{R}^d . The affine diagram defined in \mathbb{R}^d for \mathcal{O} under the distance function δ is identical to the Laguerre diagram of a set of n spheres of \mathbb{R}^d .*

Exercises 4.1

1. Show that if some spheres are imaginary (i.e. their squared radii are negative) does not lead to any additional difficulty.
2. Show that the intersection of a Laguerre diagram with an affine subspace is still a Laguerre diagram.
3. Show that the only numerical operation that is required to check if a triangulation is the regular triangulation of a set of spheres σ_i is the evaluation of the sign of the determinant of the $(d+2) \times (d+2)$ matrix

$$\text{power_test}(\sigma_0, \dots, \sigma_{d+1}) = \begin{vmatrix} 1 & \dots & 1 \\ c_0 & \dots & c_{d+1} \\ \|c_0\|^2 - r_0^2 & \dots & \|c_{d+1}\|^2 - r_{d+1}^2 \end{vmatrix}$$

where c_i and r_i are respectively the center and the radius of σ_i .

4.4 Bibliographical notes

To know more about Laguerre diagrams, also called power diagrams, one may look at the survey paper by Aurenhammer [4].

Chapitre 5

Delaunay triangulation restricted to a surface

In this section, we introduce the concept of restricted Delaunay triangulation. Given a surface \mathcal{S} of \mathbb{R}^3 and a sample E of \mathcal{S} , i.e. a finite set of points on \mathcal{S} , the Delaunay triangulation of E restricted to \mathcal{S} , noted $\text{Del}_{\mathcal{S}}(E)$, is a subcomplex of the Delaunay triangulation of E . The main result of this chapter (Theorem 5.3) states that, under some sampling conditions to be discussed in section 5.3, $\text{Del}_{\mathcal{S}}(E)$ is a good approximation of \mathcal{S} . First, we need to define what good approximation means.

5.1 Approximation criteria

There are many ways of measuring how close two objects are. We distinguish between topological and geometric criteria.

5.1.1 Topological equivalence

Definition 5.1 *Two subsets X and Y of \mathbb{R}^d are said to be homeomorphic if there exists a continuous, bijective map $f : X \rightarrow Y$ with continuous inverse f^{-1} .*

When X and Y are homeomorphic, they have the same number of handles. If X is the standard unit ball of \mathbb{R}^d (and therefore has no handle), Y is called a *topological ball*.

Definition 5.2 *Two subsets X and Y of \mathbb{R}^d are said to be isotopic if there exists a continuous map $f : X \times [0, 1] \rightarrow \mathbb{R}^d$ such that $f(\cdot, 0)$ is the identity of X , $f(X, 1) = Y$, and for each $t \in [0, 1]$, $f(\cdot, t)$ is a homeomorphism onto its image.*

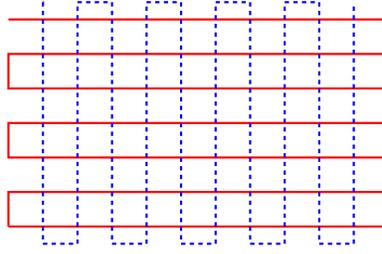


FIG. 5.1 – Two curves whose Hausdorff distance is small compared to their Fréchet distance.

5.1.2 Distance between two sets

Given a compact set X of \mathbb{R}^d , define $\text{tube}_\varepsilon(X)$ to be the tubular neighborhood of X of radius ε , i.e. the set of all points at distance at most ε from X

$$\text{tube}_\varepsilon(X) = \{y \in \mathbb{R}^d : \inf_{x \in X} \|x - y\| \leq \varepsilon\} = \bigcup_{x \in X} B(x, \varepsilon).$$

Definition 5.3 *The Hausdorff distance $d_H(X, Y)$ between two subsets X and Y of \mathbb{R}^d is the smallest ε such that $X \subset \text{tube}_\varepsilon(Y)$ and $Y \subset \text{tube}_\varepsilon(X)$. Equivalently,*

$$d_H(X, Y) = \max \left(\sup_{y \in Y} (\inf_{x \in X} \|x - y\|), \sup_{x \in X} (\inf_{y \in Y} \|x - y\|) \right).$$

The Hausdorff distance is not always a good measure of the similarity of two shapes. This is illustrated in Figure 5.1.2 where two curves are close for the Hausdorff distance but look quite different.

A more satisfactory measure of the similarity of two shapes is the so-called Fréchet distance.

Definition 5.4 *The Fréchet distance between two subsets X and Y of \mathbb{R}^d is*

$$d_F(X, Y) = \inf_h \sup_{p \in X} d(p, h(p)),$$

where h ranges over all homeomorphisms from X to Y .

If X and Y are isotopic, $\sup_{x \in X} \|f(x, 0) - f(x, 1)\|$ is an upper bound on the Fréchet distance between X and Y .

Other criteria are also useful when considering smooth curves and surfaces. Of utmost importance is the approximation of the normal field. As illustrated in Figure 5.1.2, a triangulated surface can be arbitrarily

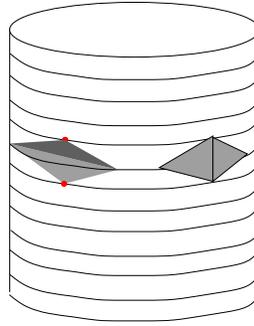


FIG. 5.2 – The normals of the two facets on the left side are almost perpendicular to the normal field of the cylinder. This is in sharp contrast with the normals of the two facets on the right side.

close to a cylinder for the Hausdorff or even the Fréchet distance, while the normal fields of the two surfaces are very different. As a consequence, it is impossible to estimate the area of the cylinder from the area of the triangulated surface.

5.2 Restricted Delaunay triangulation

In this section, X denotes a subset of \mathbb{R}^d and E a finite set of points of \mathbb{R}^d .

Definition 5.5 *If f is a face of the Voronoi diagram of E , we say that $f \cap X$ is the restriction of f to X . The subcomplex made of all non-empty restrictions of faces of $\text{Vor}(E)$ to X is called the restriction of $\text{Vor}(E)$ to X , and noted $\text{Vor}|_X(E)$.*

Definition 5.6 *We call restriction of a Delaunay triangulation $\text{Del}(E)$ to X the sub-complex of $\text{Del}(E)$, noted $\text{Del}_X(E)$, that consists of the faces of $\text{Del}(E)$ whose dual Voronoi faces intersect X . The regularized restriction of the Delaunay triangulation $\text{Del}(E)$ to X is the subcomplex of $\text{Del}_X(E)$ consisting of the faces of maximal dimension of $\text{Del}_X(E)$ and their subfaces.*

In this chapter, X will denote most of the time a surface of \mathbb{R}^3 . For pedagogical and illustration purposes, we will also consider the case of a curve in \mathbb{R}^2 .

Assume that X is a surface we denote by \mathcal{S} and that E is a *generic sample* meaning that no vertex of $\text{Vor}(E)$ lies on \mathcal{S} . It is always possible to slightly perturb E so that it is generic. Therefore $\text{Del}_{\mathcal{S}}(E)$ does not contain cf Edels. tetrahedra and is a 2-complex. The regularized $\text{Del}_{\mathcal{S}}(E)$ consists of the facets of $\text{Del}_{\mathcal{S}}(E)$, i.e. the Delaunay facets dual to the Voronoi edges that intersect \mathcal{S} .

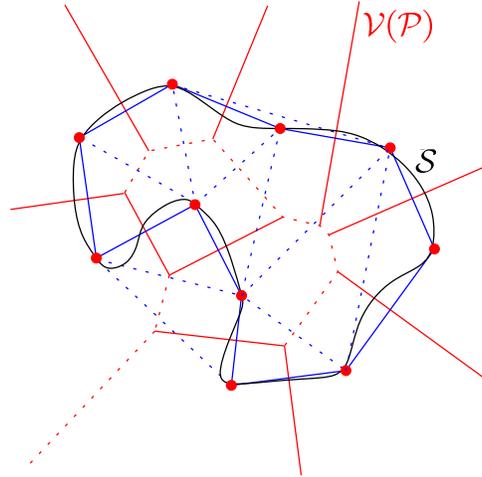


FIG. 5.3 – The Voronoi diagram and the Delaunay triangulation of a point set restricted to a planar closed curve. The edges of the restricted Voronoi diagram and of the Delaunay triangulation of the sample points are in bold lines.

We call *surface Delaunay ball* a ball that is centered on \mathcal{S} and that circumscribes a facet of $\text{Del}_{\mathcal{S}}(E)$. There may be several surface Delaunay balls associated to a given facet. In the sequel, we denote by f a facet of $\text{Del}_{\mathcal{S}}(E)$ and by $B_f = B(c_f, r_f)$ a surface Delaunay ball that circumscribes f .

An example is shown in Figure 5.3. As can be seen on this example, $\text{Del}_{\mathcal{S}}(E)$ is a simple closed polygon that correctly approximates \mathcal{S} . We will show, in section 5.4, that this is indeed the case when E is a sufficiently dense sample of \mathcal{S} .

5.3 Sampling conditions

Let \mathcal{O} be an open set of \mathbb{R}^d . The medial axis $\mathcal{M}(\mathcal{O})$ of \mathcal{O} will play a crucial role in the sequel. $\mathcal{M}(\mathcal{O})$ can be seen as a generalization of the notion of Voronoi diagram to infinite point sets.

Definition 5.7 (Medial axis) *The medial axis of \mathcal{O} is the closure of the set of points with at least two closest points on the boundary of \mathcal{O} .*

Observe that if \mathcal{O} is the complement $\mathbb{R}^d \setminus E$ of a finite set of points E , $\mathcal{M}(\mathcal{O})$ is the subcomplex of $\text{Vor}(E)$ obtained by removing from $\text{Vor}(E)$ the cells of full dimension d .

A ball that is centered on the medial axis, whose interior is contained in \mathcal{O} , and whose bounding sphere intersects the boundary of \mathcal{O} is called a *medial ball* (see Figure 5.4).

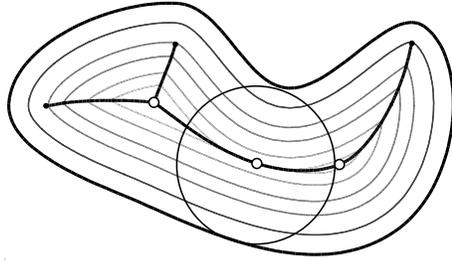


FIG. 5.4 – The medial axis of a planar domain. The thin curves are parallel to the boundary of the domain. The circle bounds a medial ball.

Definition 5.8 (Local feature size) *The local feature size at a point $x \in \mathcal{O}$, noted $\text{lfs}(x)$, is the distance of x to the medial axis of \mathcal{O} .*

It can easily be proved that lfs is a 1-Lipschitz function, i.e. $\forall x, y \in \mathcal{O}$, $\text{lfs}(x) \leq \text{lfs}(y) + \|x - y\|$.

We consider now the case where \mathcal{S} denotes a surface of \mathbb{R}^3 and $\mathcal{O} = \mathbb{R}^3 \setminus \mathcal{S}$. For convenience, we call *medial axis of \mathcal{S}* the medial axis of \mathcal{O} , and write $\mathcal{M}(\mathcal{S})$ for $\mathcal{M}(\mathcal{O})$.

We will restrict our attention to the class $C^{1,1}$ of surfaces that admits a normal at each point and whose normal field is Lipschitz. This class is larger than the class of C^2 surfaces and includes surfaces whose curvature may be discontinuous at some points. An example of a surface that is $C^{1,1}$ but not C^2 is the offset (i.e. the boundary of a tube) of a cube.

An important property of $C^{1,1}$ surfaces is the following lemma.

Lemma 5.1 *For a surface of class $C^{1,1}$, lfs is bounded away from 0.*

Lemma 5.2 *Let \mathcal{S} be a surface of \mathbb{R}^3 and B a ball centered at a point x of radius r that intersects \mathcal{S} . If $B \cap \mathcal{S}$ is not a topological disk, B contains a point of the medial axis of \mathcal{S} .*

Proof. The result is trivial when x belongs to the medial axis $\mathcal{M}(\mathcal{S})$ of \mathcal{S} . Therefore assume that $x \notin \mathcal{M}(\mathcal{S})$ and, for a contradiction, that $B \cap \mathcal{S}$ is not a topological disk.

We denote by B_x the largest ball centered at x whose interior does not intersect \mathcal{S} . Let r_x be its radius. B_x is tangent to \mathcal{S} in a unique point y since x does not belong to $\mathcal{M}(\mathcal{S})$. Hence $B_x \cap \mathcal{S} = \{y\}$ is a topological disk. Since $B \cap \mathcal{S}$ is not a topological disk, there exists a point z of \mathcal{S} at distance $r > r_c > r_x$ from x such that the ball $B_c = B(x, r_c)$ is tangent to \mathcal{S} in z . The medial ball tangent to \mathcal{S} in z and centered on the half-line $[zx)$ is contained in B_c (since it is tangent to \mathcal{S} in z and cannot contain y) and therefore in B . Its center belongs to $\mathcal{M}(\mathcal{S})$. \square

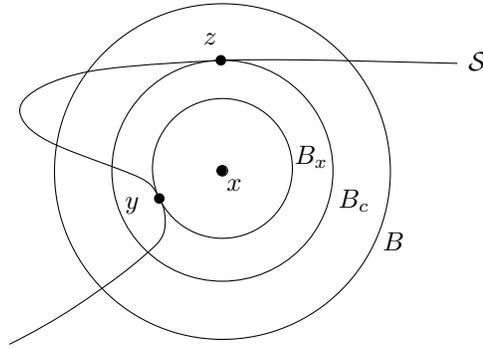


FIG. 5.5 – For the proof of Proposition 5.2.

Since $B(x, r)$ cannot intersect the medial axis of \mathcal{S} for any $r < \text{lfs}(x)$, Lemma 5.2 implies

Lemma 5.3 *For any x of \mathcal{S} , and any $r < \text{lfs}(x)$, the intersection of \mathcal{S} with the ball $B(x, r)$ centered at x of radius r is a topological disk.*

Definition 5.9 (ε -sample) *Let $\varepsilon < 1$ and \mathcal{S} be a surface of class $C^{1,1}$. We say that a finite point set $E \subset \mathcal{S}$ is an ε -sample of \mathcal{S} if any point of \mathcal{S} is at distance at most $\varepsilon \text{lfs}(x)$ from a point of E .*

In other words, E is a finite set of points on \mathcal{S} such that $d_H(E, \mathcal{S}) \leq \varepsilon$. Observe that ε -samples are non uniform samples which are denser where lfs is small (See Figure 5.6).

Let E be an ε -sample of \mathcal{S} , f a facet of $\text{Del}_{\mathcal{S}}(E)$ and v a vertex of f .

Lemma 5.4 *For any $p \in E$, $V(p) \cap \mathcal{S} \subset B(p, r)$ where $r \leq \frac{\varepsilon}{1-\varepsilon} \text{lfs}(p)$.*

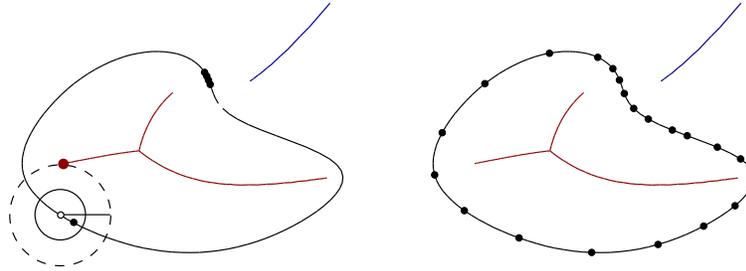
Proof. Let x be a point of $V(p) \cap \mathcal{S}$. Using the Lipschitz property of lfs , we have

$$\|p - x\| \leq \varepsilon \text{lfs}(x) \leq \varepsilon \text{lfs}(p) + \|p - x\| \leq \frac{\varepsilon}{1-\varepsilon} \text{lfs}(p).$$

□

5.4 Approximation of smooth surfaces

In this section, \mathcal{S} denotes an oriented closed surface of \mathbb{R}^3 of class $C^{1,1}$. Our goal is to prove that, when E is an ε -sample of \mathcal{S} for some sufficiently small ε , $\text{Del}_{\mathcal{S}}(E)$ and \mathcal{S} are isotopic and that the distance between a

FIG. 5.6 – An ε -sample.

point and its image by the isotopy is a decreasing function of the sampling density (Theorem 5.3). We will also prove that the normals of the facets of $\text{Del}_{\mathcal{S}}(E)$ are close to the normal field of \mathcal{S} (Lemma 5.7).

Before considering the case of a surface of \mathbb{R}^3 , we consider the simpler case of a curve in the next section.

The proof for surfaces is subdivided into three parts. We first prove some local results, then prove that $\text{Del}_{\mathcal{S}}(E)$ is a closed surface and, finally, we exhibit an isotopy between \mathcal{S} and $\text{Del}_{\mathcal{S}}(E)$.

5.4.1 Approximation of a curve

In this section, \mathcal{S} denotes a simple curve of \mathbb{R}^d and E a sample of \mathcal{S} . We assume that E contains at least three points on each closed connected component of \mathcal{S} (H1). If a component is an arc, we assume that the two endpoints of the arc belong to E (H2). \mathcal{S} is then the union of *elementary arcs* joining two points of E that are consecutive along \mathcal{S} . Note that two elementary arcs have at most one endpoint in common. In addition, we assume that no face of dimension less than $d - 1$ of $\text{Vor}(E)$ intersect \mathcal{S} (H3), which can always be ensured by slightly perturbing E .

Theorem 5.1 *If E be an ε -sample of \mathcal{S} , then $\text{Del}_{\mathcal{S}}(E)$ is a polygonal curve homeomorphic to \mathcal{S} and $d_F(\text{Del}_{\mathcal{S}}(E), \mathcal{S}) \leq 2\varepsilon \sup_{x \in \mathcal{S}} \text{lfs}(x)$.*

Proof. By (H3), $\text{Del}_{\mathcal{S}}(E)$ contains no face of dimension 2 or more, and therefore consists only of edges and vertices.

Let p_1 and p_2 be two points of E that are consecutive along \mathcal{S} . We prove that $p_1 p_2$ is an edge of $\text{Del}_{\mathcal{S}}(E)$. Let γ_{12} be the elementary arc of \mathcal{S} that joins p_1 to p_2 . We denote by c an intersection point between γ_{12} and the bisector of p_1 and p_2 , and by σ the largest sphere centered at c that does not enclose any point of E . If σ passes through p_1 and p_2 , then $p_1 p_2$ is a Delaunay edge and therefore an edge of $\text{Del}_{\mathcal{S}}(E)$. Otherwise, σ passes through a third point p_3 of E and the intersection of the ball bounded by σ with \mathcal{S} contains at

least two arcs. By Lemma 5.3 (adapted to the case where \mathcal{S} is a curve, see Exercise 5.11), the radius of σ is therefore greater than $\text{lfs}(c)$, which contradicts the assumption that E is an ε -sample of \mathcal{S} .

Conversely, we show that any Delaunay edge $p_i p_j$ whose Voronoi dual facet $p_i p_j^*$ intersects \mathcal{S} joins two points of E that are consecutive along \mathcal{S} . We denote by c an intersection point of $p_i p_j^*$ with \mathcal{S} . The sphere σ that is centered at c and passes through p_i and p_j does not enclose any point of E : hence, its radius is less than $\text{lfs}(c)$ since E is an ε -sample. By Lemma 5.3, the portion of \mathcal{S} that is contained in the ball bounded by σ is an arc : p_i and p_j are therefore consecutive along \mathcal{S} .

Two consecutive points defining a unique elementary arc (H1), we have a bijection between the elementary arcs of \mathcal{S} and the edges of $\text{Del}_{\mathcal{S}}(E)$. Each elementary arc is homeomorphic to its associated edge, and, since an elementary arc of \mathcal{S} and the associated edge of $\text{Del}_{\mathcal{S}}(E)$ have the same endpoints, we obtain a homeomorphism between \mathcal{S} and $\text{Del}_{\mathcal{S}}(E)$ by concatenating the elementary homeomorphisms between the elementary arcs and their associated edges of $\text{Del}_{\mathcal{S}}(E)$.

The last claim comes from the observation made above that an elementary arc γ_{pq} and its associated edge of $\text{Del}_{\mathcal{S}}(E)$ are both contained in the surface Delaunay ball¹ circumscribing pq . Hence, the distance between a point and its image is at most twice the radius of such a ball, which is less than $2\varepsilon \sup_{x \in \mathcal{S}} \text{lfs}(x)$. \square

5.4.2 Approximation of a smooth surface : local lemmas

Notations. For $x \in \mathcal{S}$, n_x denotes the unit vector normal to \mathcal{S} , oriented according to the orientation of \mathcal{S} , and T_x denotes the tangent plane at x . In this section, f denotes a facet of $\text{Del}_{\mathcal{S}}(E)$, B_f a surface Delaunay ball circumscribing f , c_f its center, r_f its radius, and n_f a unit vector perpendicular to f . Unless it is explicitly mentioned, we do not specify an orientation for n_f in this section. The modulus of the angle between two vectors n and n' , measured in $[0, \pi]$, is noted (n, n') . We define the angle between a vector n and a plane P , noted (n, P) , as the smallest angle between n and a vector of P . We define similarly the angle between a vector and a line.

For a real α , we write $\alpha^+ = \frac{\alpha}{1-\alpha}$.

Lemma 5.5 (Chord angle lemma) *For any two points x and y of \mathcal{S} with $\|x - y\| \leq \eta \text{lfs}(x)$, $\eta \leq 2$, the angle between xy and the tangent plane T_x at x is at most $\arcsin \frac{\eta}{2}$.*

Proof. \mathcal{S} and therefore y do not intersect the interior of the tangent balls of radius $\text{lfs}(x)$ at x . The modulus of the angle θ_x between xy and T_x is maximized when y lies on the boundary of one of the tangent balls. We then have $\sin \theta_x = \frac{\|x-y\|}{2\text{lfs}(x)} \leq \frac{\eta}{2}$. \square

Lemma 5.6 (Normal variation lemma) *Let x and y be two points of \mathcal{S} with $\|x-y\| \leq \eta \min(\text{lfs}(x), \text{lfs}(y))$, $\eta \leq 2$. The angle (n_x, n_y) between the normals n_x and n_y is at most $2 \arcsin \frac{\eta}{2}$.*

¹Curve Delaunay ball would be more appropriate.

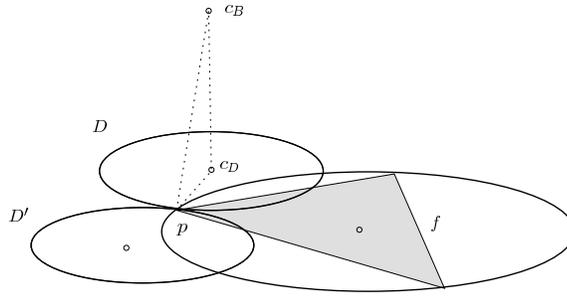


FIG. 5.7 – For the proof of Lemma 5.7.

Proof. To do. \square

Lemma 5.7 (Facet normal lemma) *Let $f = pqr$ be a facet of $\text{Del}(E)$ and assume that $\hat{p} \geq \frac{\pi}{3}$. If the circumradius ρ_f of f is at most $\eta \text{lfs}(p)$, then $(n_f, n_p) \leq \arcsin(\eta\sqrt{3})$.*

Proof. Let B and B' denote the two balls of radius $\text{lfs}(p)$ tangent to \mathcal{S} at p , $D = B \cap \text{aff}(f)$ and $D' = B' \cap \text{aff}(f)$ (see Figure 5.4.2). We call c_B and c_D the centers of B and D respectively, and ρ_D the radius of D and D' .

We first observe that $\rho_D \leq \rho_f\sqrt{3}$. Indeed, the interiors D and D' contain no vertex of f . ρ_D is then maximized with respect to ρ_f when f is equilateral with one vertex other than p , say q , on the boundary of D and r on the boundary of D' .

We then have

$$\sin(n_f, n_p) = \sin(\angle c_B p c_D) \leq \frac{\|p - c_D\|}{\|p - c_B\|} = \frac{\rho_D}{\text{lfs}(p)} \leq \eta\sqrt{3}$$

\square

Lemma 5.8 (Projection lemma) *Let E be an ε -sample of \mathcal{S} for $\varepsilon < 0.24$, and f and f' be two facets of $\text{Del}_{\mathcal{S}}(E)$ with a common vertex p . Then, the orthogonal projections of f and f' onto the tangent plane at p do not overlap, i.e. the projections of the relative interiors of f and f' do not intersect.*

Proof. Assume, for a contradiction, that there exists a line l parallel to n_p that intersects f and f' . We orient l as n_p and assume without loss of generality that f intersects l before f' . Let $T = t_1, \dots, t_s$ be the sequence of tetrahedra intersected by l between f and f' . Without loss of generality, we take a line l very close to p so that all the tetrahedra of T have p as a vertex. We denote by c_1, \dots, c_s the centers of the Delaunay balls circumscribing t_1, \dots, t_s , and γ the polygonal chain $\gamma = c_f, c_1, \dots, c_s, c_{f'}$. Clearly, γ is a path in $\text{Vor}(E)$. Moreover, γ is monotone with respect to l . Indeed, the directed line l intersects t_i before t_{i+1} , $i = 1, \dots, s-1$. Because t_i and t_{i+1} are Delaunay tetrahedra, the same is true for the line $(c_i c_{i+1})$ directed from c_i to c_{i+1} . Hence $(c_i - c_{i+1}) \cdot n_p > 0$ for $i = 1, \dots, s-1$. For essentially the same reason, we also have $(c_1 - c_f) \cdot n_p > 0$ and $(c_{f'} - c_s) \cdot n_p > 0$.

The proof is in parts. First, we provide an upper bound on the angle between $c_f c_{f'}$ and T_p , the tangent plane at p . Then, we will provide a lower bound on the same angle, leading to a contradiction.

1. Using the Lipschitz property of function lfs and the hypothesis of the lemma, we have

$$\begin{aligned} \|c_f - c_{f'}\| &\leq \|c_f - p\| + \|p - c_{f'}\| \\ &\leq \varepsilon (\text{lfs}(c_f) + \text{lfs}(c_{f'})) \\ &\leq \varepsilon (2 \max(\text{lfs}(c_f), \text{lfs}(c_{f'})) + \|c_f - c_{f'}\|) \\ &\leq 2\varepsilon^+ \min(\text{lfs}(c_f), \text{lfs}(c_{f'})). \end{aligned} \tag{5.1}$$

Then, by the Chord angle lemma, the angle ϕ between $c_f c_{f'}$ and the tangent plane T_{c_f} at c_f is at most $\arcsin \varepsilon^+$. Moreover, since $\|c_f - p\| \leq \varepsilon \text{lfs}(c_f) \leq \varepsilon^+ \text{lfs}(p)$, the Normal variation lemma 5.6 implies that $\psi = (n_{c_f}, n_p) \leq 2 \arcsin \frac{\varepsilon^+}{2}$. It follows that

$$(c_f c_{f'}, T_p) \leq \phi + \psi \leq \arcsin \varepsilon^+ + 2 \arcsin \frac{\varepsilon^+}{2} \tag{5.2}$$

2. We exhibit now a lower bound on $(c_f c_{f'}, T_p)$. We first need to prove that all the c_i are close to p . Let B_p be the ball centered at p of radius $\varepsilon^+ \text{lfs}(p)$, and m_p and m'_p be the two points on the line normal to \mathcal{S} at p , at distance $\text{lfs}(p)$ from p . We define W_p as the union of two solid half-cones C_p and C'_p . Refer to Figure 5.8. The boundary of C_p (resp. C'_p) is the union of the rays issued from m_p (resp. m'_p) that are tangent to B_p . Observe that, since the radius of B_p is $\varepsilon^+ \text{lfs}(p)$, the half-angle of C_p and C'_p is $\theta_p = \arcsin \varepsilon^+$. Let $V(p)$ be the Voronoi cell of p . By Lemma 5.4, m_p and m'_p belongs to $V(p)$ and $V(p) \cap \mathcal{S} \subset B_p$. We claim that $V(p)$ is contained in W_p . Indeed, otherwise, let $x \in V(p) \setminus W_p$. Assuming that x and m_p lie on distinct sides of \mathcal{S} (otherwise exchange the role of m_p and m'_p), the line segment $[x m_p]$ is not entirely contained in $V(p)$ since it must intersect \mathcal{S} but does not intersect $B_p \supseteq V(p) \cap \mathcal{S}$. This violates the convexity of $V(p)$ and proves the claim. Hence, the c_i , which are vertices of $V(p)$, belong to W_p . Moreover, by the monotonicity of γ , the c_i lie in the slab limited by the two planes perpendicular to n_p passing respectively through c_f and $c_{f'}$. By the hypothesis of the lemma, $\|c_f - p\| \leq \varepsilon \text{lfs}(c_f) \leq \varepsilon^+ \text{lfs}(p)$ and the same inequality holds for $\|c_{f'} - p\|$. The

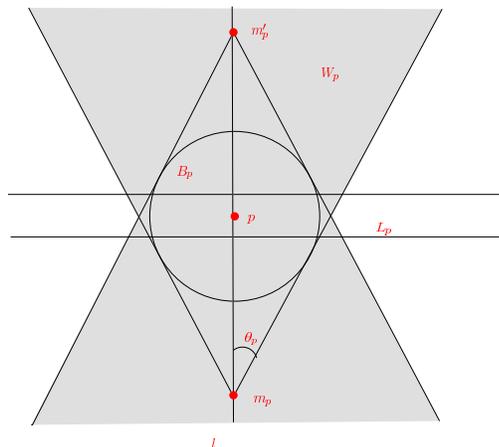


FIG. 5.8 – For the proof of the Projection lemma.

Chord angle lemma then implies that $|(c_f - p) \cdot n_p| \leq \frac{1}{2} \varepsilon^{+2} \text{lfs}(p) \stackrel{\text{def}}{=} h$, and also $|(c_{f'} - p) \cdot n_p| \leq h$. Hence the c_i lie in the slab L_p limited by the two planes perpendicular to n_p at distance h from p . The points of $W_p \cap L_p$ that are most distant from p lie on one of the two circles $W_p \cap \partial L_p$. Elementary computations then show that

$$\|c_i - p\| \leq \varepsilon^+ \frac{\sqrt{1 + \frac{5}{4}\varepsilon^{+2}}}{\sqrt{1 - \varepsilon^{+2}}} \text{lfs}(p) = \varepsilon^* \text{lfs}(p).$$

By the Facet normal lemma 5.7 and the Normal deviation lemma 5.6, we obtain

$$(n_{f_i}, n_p) \leq \arcsin(\varepsilon^* \sqrt{3}) + 2 \arcsin \frac{\varepsilon^*}{2(1 - \varepsilon^*)} \stackrel{\text{def}}{=} \chi.$$

Observing that $c_i - c_{i+1}$ is perpendicular to facet f_i , we get $(c_{i+1} - c_i) \cdot n_p \geq \|c_{i+1} - c_i\| \cos \chi$.

Using this last inequality (a) and the monotonicity of γ (b), we get

$$\begin{aligned} \|c_f - c_{f'}\| \cos \chi &\leq \sum_{i=0}^s \|c_{i+1} - c_i\| \cos \chi \\ &\stackrel{a}{\leq} \sum_{i=0}^s (c_{i+1} - c_i) \cdot n_p \\ &\stackrel{b}{=} (c_f - c_{f'}) \cdot n_p \\ &= \|c_f - c_{f'}\| \sin((c_f c_{f'}, T_p)). \end{aligned} \tag{5.3}$$

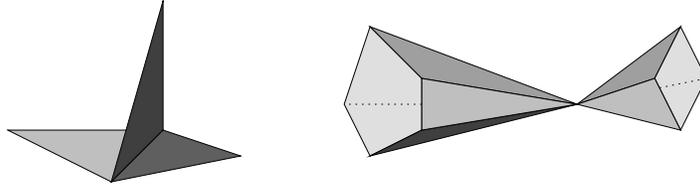


FIG. 5.9 – An edge with three incident facets and a vertex whose star is not a simple polygon.

For $\varepsilon < 0.24$, $\cos \chi > \sin \left(\arcsin \varepsilon^+ + 2 \arcsin \frac{\varepsilon^+}{2} \right)$ and we obtain a contradiction from inequalities (5.2) and (5.3).

□

Remark. Observe that the proof of the Projection lemma still holds if c_f and $c_{f'}$ are the centers of two surface Delaunay balls circumscribing the same facet $f' = f^2$. Hence, the Voronoi edge dual to any facet of $\text{Del}_{\mathcal{S}}(E)$ cannot intersect \mathcal{S} more than once if E is an ε -sample of \mathcal{S} with $\varepsilon < 0.24$.

5.4.3 $\text{Del}_{\mathcal{S}}(E)$ is a closed triangulated surface

In this subsection, we prove that $\text{Del}_{\mathcal{S}}(E)$ is a closed triangulated surface of \mathbb{R}^3 . Since $\text{Del}_{\mathcal{S}}(E)$ is a 2-complex, this is equivalent to proving that each edge of $\text{Del}_{\mathcal{S}}(E)$ is incident to exactly two facets, and that the boundary of the star of each vertex is a simple closed polygon. The *star* of a vertex p is the union of the facets of $\text{Del}_{\mathcal{S}}(E)$ that are incident to p .

Proposition 5.1 *If E is an ε -sample and $\varepsilon < 0.24$, $\text{Del}_{\mathcal{S}}(E)$ is a closed surface.*

Proof. 1. We first prove that every edge of $\text{Del}_{\mathcal{S}}(E)$ is incident to exactly two facets of $\text{Del}_{\mathcal{S}}(E)$. Let e be an edge of $\text{Del}_{\mathcal{S}}(E)$. We denote by e^* the Voronoi face dual to e . Since \mathcal{S} has no boundary, $\mathcal{S} \cap \text{aff}(e^*)$ is a union of simple closed curves, none of which intersects the boundary ∂e^* of e^* tangentially, by the Facet normal lemma 5.7. Hence, by the Jordan curve theorem, each component of $\mathcal{S} \cap \text{aff}(e^*)$ intersects ∂e^* an even number of times. It follows that \mathcal{S} intersects ∂e^* at an even number of points. Moreover, by the remark after the Projection lemma, each edge of ∂e^* cannot be intersected more than once by \mathcal{S} . Thus, \mathcal{S} intersects an even number of edges of ∂e^* . Equivalently, e is incident to an even number of restricted Delaunay facets.

²The main change is the fact that T is now empty, which makes the proof much simpler.

In addition, e cannot be incident to more than two facets of $\text{Del}_{\mathcal{S}}(E)$. Indeed, consider any two facets f and f' incident to e and let p be an endpoint of e . By the Projection lemma 5.8, the projections of f and f' onto $T(p)$ do not overlap. Hence, they must lie on different sides of the line supporting the projection of e , which proves the claim.

In conclusion, the number of facets of $\text{Del}_{\mathcal{S}}(E)$ that are incident to e is even, at least 1 and at most 2. It follows that e is incident to exactly two facets of $\text{Del}_{\mathcal{S}}(E)$.

2. Consider now a vertex p of $\text{Del}_{\mathcal{S}}(E)$ whose star boundary is not a simple closed polygon. By the discussion above, the star of p consists of several components we call *umbrellas*. Each umbrella is a triangulated topological disk. All umbrellas of p have p in common but two distinct umbrellas have distinct edges and facets.

We prove now that every vertex of $\text{Del}_{\mathcal{S}}(E)$ has exactly one umbrella. Let U be one of the umbrellas incident to p and \bar{U} its projection onto the tangent plane T_p at p .

We claim that p belongs to the interior of \bar{U} . Let us assume the contrary. Then U has a silhouette edge pq whose projection on T_p belongs to the boundary of \bar{U} . Since, by Proposition 5.1, pq is incident to two facets of U , these two facets must project onto the same side of (p, q) , and therefore overlap, which contradicts the Projection lemma 5.8.

As a consequence, p cannot have two umbrellas U and U' . Indeed, otherwise, there would be a facet of U and a facet of U' whose projections onto T_p overlap, which contradicts again the Projection lemma 5.8. \square

5.4.4 $\text{Del}_{\mathcal{S}}(E)$ and \mathcal{S} are isotopic

We assume that $\varepsilon < 0.24$. By Proposition 5.1, $\text{Del}_{\mathcal{S}}(E)$ is a closed surface we can therefore orient : specifically, we orient the normals of the facets of $\text{Del}_{\mathcal{S}}(E)$ so that they all point outside the bounded region limited by $\text{Del}_{\mathcal{S}}(E)$. In the sequel, we write n_f for the *oriented* normal of a facet f of $\text{Del}_{\mathcal{S}}(E)$.

We define the ε -thickening of \mathcal{S} as $\text{thick}_{\varepsilon}(\mathcal{S}) = \bigcup_{x \in \mathcal{S}} B(x, \varepsilon \text{lf}_s(x))$. For $\varepsilon < 1$, $\text{thick}_{\varepsilon}(\mathcal{S})$ does not intersect the medial axis of \mathcal{S} and its boundary has two components $\mathcal{S}_{\varepsilon}^{+}$ and $\mathcal{S}_{\varepsilon}^{-}$. We call *fiber* at x the line segment normal to \mathcal{S} at x with one endpoint on $\mathcal{S}_{\varepsilon}^{+}$ and the other on $\mathcal{S}_{\varepsilon}^{-}$.

The following result is a basic theorem in differential topology [25]. Let $\pi_{\mathcal{S}} : \mathbb{R}^3 \rightarrow \mathcal{S}$ map a point of \mathbb{R}^3 to its closest point on \mathcal{S} . The restriction of $\pi_{\mathcal{S}}$ to $\text{thick}_{\varepsilon}$ is a well-defined continuous function since $\text{thick}_{\varepsilon}(\mathcal{S})$ does not intersect the medial axis of \mathcal{S} .

Theorem 5.2 *Let \mathcal{S} be a $C^{1,1}$, compact, oriented, closed surface and T be a compact closed surface (not necessarily smooth) such that :*

1. $T \subset U = \text{thick}_{\varepsilon}(\mathcal{S})$,
2. *Any fiber intersects T in exactly one point.*

Then the restriction of $\pi_{\mathcal{S}}$ to T induces an isotopy that maps T to \mathcal{S} . The isotopy does not move the points

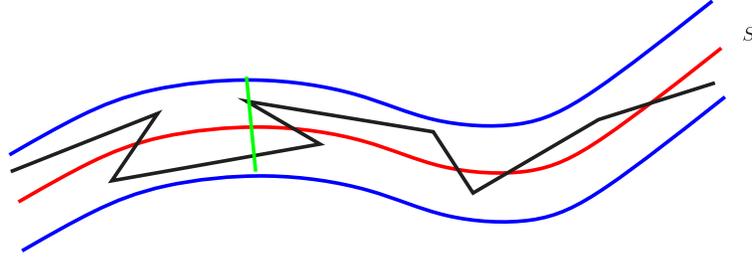


FIG. 5.10 – $\text{Del}_{\mathcal{S}}(E)$ is included in $\text{thick}_{\varepsilon}(\mathcal{S})$. A fiber is shown in light grey.

by more than $\varepsilon \sup_{x \in \mathcal{S}} \text{lfs}(x)$.

This theorem will be used to prove the following result :

Theorem 5.3 *If E is an ε -sample, $\varepsilon < 0.2$, then the restriction of $\pi_{\mathcal{S}}$ to $\text{Del}_{\mathcal{S}}(E)$ induces an isotopy that maps $\text{Del}_{\mathcal{S}}(E)$ to \mathcal{S} . The isotopy does not move the points by more than $\varepsilon \sup_{x \in \mathcal{S}} \text{lfs}(x)$.*

Before proving the theorem in its full generality, we prove that, for any $p \in E$, the fiber passing through p intersects $\text{Del}_{\mathcal{S}}(E)$ at point p only.

Lemma 5.9 *Let E be a ε -sample of \mathcal{S} , $\varepsilon < 1$, $p \in E$, and let m_p be the center of a medial ball B_p tangent to \mathcal{S} at p . The segment pm_p can only intersect $\text{Del}_{\mathcal{S}}(E)$ at point p .*

Proof. Let $p \in E$, m_p the center of a medial ball B_p tangent to \mathcal{S} at p . We prove that the open segment pm_p cannot intersect $\text{Del}_{\mathcal{S}}(E)$. Assume that pm_p intersects a facet f of $\text{Del}_{\mathcal{S}}(E)$. ∂B_p and ∂B_f necessarily intersect since the vertices of f do not belong to the interior of B_p and p does not belong to the interior of B_f . Let H be the plane containing the intersection of the spheres bounding B_p and B_f , H^+ be the open halfspace limited by H that contains $B_f \setminus B_p$, and H^- the other open halfspace which contains $B_p \setminus B_f$. Since p does not belong to the interior of B_f and is on the boundary of B_p , p must belong to $H^- \cup H$. On the other hand, m_p belongs to H^- since otherwise m_p would belong to B_f , which contradicts the assumption that $r_f < \text{lfs}(c_f)$. Hence the open segment pm_p is contained in the open half space H^- , and therefore cannot intersect f which is contained in $H^+ \cup H$. \square

We now proceed to the proof of Theorem 5.3.

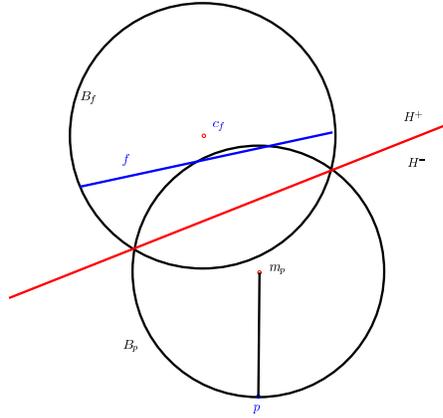


FIG. 5.11 – For the proof of Lemma 5.9.

Proof. We prove that the hypotheses of Theorem 5.2 are satisfied for \mathcal{S} , $T = \text{Del}_{\mathcal{S}}(E)$ and $U = \text{thick}_{\varepsilon}(\mathcal{S})$. Each facet of T is contained in the corresponding surface Delaunay ball. By definition of ε -samples, the radius of the surface Delaunay ball of any facet f of $\text{Del}_{\mathcal{S}}(E)$ is $\leq \varepsilon \text{lfs}(c_f)$ where c_f is the center of the ball. Hence the first condition of Theorem 5.2 is fulfilled.

Consider now the second condition. Let $\mathcal{S}_1 \subset \mathcal{S}$ be the subset of points of \mathcal{S} that are images by $\pi_{\mathcal{S}}$ of exactly one point of T . $\mathcal{S}_1 \neq \emptyset$ by Lemma 5.9. Moreover, a sufficiently small neighborhood of any vertex p of T projects 1-1 onto the tangent plane T_p to \mathcal{S} at p since the star of p in T is a closed polygon (Proposition 5.1) that projects 1-1 onto T_p (Projection lemma 5.8). The continuity of the normal field then implies that \mathcal{S}_1 contains an open neighborhood of each vertex of T .

If \mathcal{S}_1 has no boundary, then, since E intersects all the connected components of \mathcal{S} (exercise 5.12), $\mathcal{S}_1 = \mathcal{S}$ and the theorem is proved. Assume, for a contradiction, that $\partial\mathcal{S}_1 \neq \emptyset$, and let $x \in \partial\mathcal{S}_1$ and l be the fiber at x . l intersects T at a point t that necessarily belongs to an edge $e = pq$ of T . Let us assume that t is not an endpoint of e (otherwise move x slightly along $\partial\mathcal{S}_1$). By Proposition 5.1, there are exactly two facets of T , say f and f' , that are incident to e and, since $x \in \partial\mathcal{S}_1$, $n_f \cdot n_x$ and $n_{f'} \cdot n_x$ must have different signs.

Let us bound (n_f, n_x) . We first bound (n_f, n_p) using the Facet normal lemma. We then bound the distance between x and p and deduce from the Normal variation lemma a bound on (n_x, n_p) .

Let s be a vertex of f with greatest angle. By the Facet normal lemma 5.7, $(n_f, n_s) \leq \arcsin \varepsilon^+ \sqrt{3}$. Moreover, since $\|s-p\| \leq 2\varepsilon \text{lfs}(c_f) \leq 2\varepsilon^+ \min(\text{lfs}(p), \text{lfs}(s))$, the Normal variation lemma 5.6 gives $(n_s, n_p) \leq 2 \arcsin \varepsilon^+$. Hence

$$(n_f, n_p) \leq \arcsin \varepsilon^+ \sqrt{3} + 2 \arcsin \varepsilon^+ \tag{5.4}$$

We now bound the distance between p and x . Write y for the projection of x onto edge e and assume without

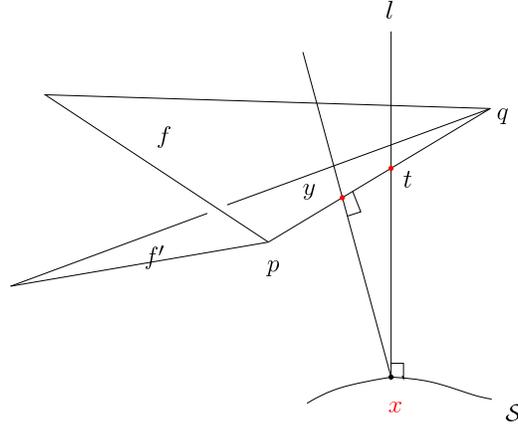


FIG. 5.12 – For the proof of Theorem 5.3.

loss of generality that $\|y - p\| \leq \|y - q\|$. We have $\|x - p\|^2 = \|x - y\|^2 + \|y - p\|^2 \leq \varepsilon^2 \text{lfs}^2(x) + \varepsilon^2 \text{lfs}^2(c_f)$. and, since

$$\begin{aligned} \text{lfs}(c_f) &\leq \text{lfs}(x) + \|x - c_f\| \\ &\leq \text{lfs}(x) + \|x - t\| + \|t - c_f\| \\ &\leq (1 + \varepsilon) \text{lfs}(x) + \varepsilon \text{lfs}(c_f) \\ &\leq \frac{1 + \varepsilon}{1 - \varepsilon} \text{lfs}(x) \end{aligned}$$

we obtain

$$\|x - p\| \leq \varepsilon^+ \sqrt{2(1 + \varepsilon^2)} \text{lfs}(x) \stackrel{\text{def}}{=} \eta \text{lfs}(x) \leq \eta^+ \text{lfs}(p).$$

By the Normal deviation lemma 5.6, we deduce $(n_p, n_x) \leq 2 \arcsin \frac{\eta^+}{2}$, which, together with equation (5.4), yields

$$(n_f, n_x) \leq \arcsin \varepsilon^+ \sqrt{3} + 2 \arcsin \varepsilon^+ + 2 \arcsin \frac{\eta^+}{2} \quad (5.5)$$

For $\varepsilon > 0.2$, $(n_f, n_x) \leq \frac{\pi}{2}$. The same inequality holds for $(n_{f'}, n_x)$. It follows that $n_f \cdot n_x$ and $n_{f'} \cdot n_x$ have the same sign, which contradicts our previous observation. It follows that $\partial \mathcal{S}_1 = \emptyset$ and therefore that $\mathcal{S}_1 = \mathcal{S}$, which proves that $\text{Del}_{\mathcal{S}}(E)$ and \mathcal{S} are isotopic.

The isotopy is induced by projecting $\text{Del}_{\mathcal{S}}(E)$ onto \mathcal{S} . Since, $\text{Del}_{\mathcal{S}}(E)$ and \mathcal{S} are both contained in $\text{thick}_{\varepsilon}(\mathcal{S})$, the isotopy does not move the points by more than $\varepsilon \sup_{x \in \mathcal{S}} \text{lfs}(x)$. \square

5.4.5 Loose ε -samples

The notion of ε -sample is well-suited for analysis but is difficult to use in practise. Indeed, the definition involves a condition for infinitely many points of \mathcal{S} . The following concept relaxes this condition to a finite set of points.

Definition 5.10 *Given a surface of class $C^{1,1}$ and $\varepsilon < 1$, we say that a finite set of points $E \subset \mathcal{S}$ is a loose ε -sample of \mathcal{S} if*

1. $E \subset \mathcal{S}$ and $\text{Del}_{\mathcal{S}}(E)$ has at least one vertex on each connected component of \mathcal{S} ,
2. for any facet f of $\text{Del}_{\mathcal{S}}(E)$ and any surface Delaunay ball $B_f = B(c_f, r_f)$ circumscribing f , we have $r_f \leq \varepsilon \text{lfs}(c_f)$.

Clearly, an ε -sample is a loose ε -sample. Interestingly, the converse is true asymptotically as stated in Theorem 5.4 below. We need first the following result :

Lemma 5.10 *Let E be a loose ε -sample of \mathcal{S} . Let $p \in E$ and $V(p)$ denote the Voronoi cell of p . $V(p) \cap \mathcal{S}$ is contained in the ball $B(p, r)$ centered at p of radius $r = \varepsilon(1 + O(\varepsilon^2)) \text{lfs}(p)$.*

Proof. We consider a slab L_p of width $h = O(\varepsilon^2)$ which is a thickening of the tangent plane T_p at p . We call $C(p)$ the portion of $V(p) \cap \mathcal{S}$ that is contained in L_p . The proof consists of two steps. We first prove that $C(p)$ is contained in $B(p, r)$. We then show that $V(p) \cap \mathcal{S}$ lies entirely in L_p and therefore is identical to $C(p)$, which proves the lemma.

1. The first part of the proof is similar to the proof of the Projection lemma 5.8 and one can refer to Figure 5.8. Let U_p be the star of p in $\text{Del}_{\mathcal{S}}(E)$ and write f_1, \dots, f_s for the facets of U_p . The centers c_i of the surface Delaunay balls circumscribing the f_i are at distance at most $\varepsilon \text{lfs}(c_i) \leq \varepsilon^+ \text{lfs}(p) \stackrel{\text{def}}{=} \rho$ from p . Moreover, by the Chord angle lemma 5.5, we have

$$|(c_i - p) \cdot n_p| \leq \frac{\varepsilon^+}{2} \text{lfs}(p) \stackrel{\text{def}}{=} h.$$

Hence, the c_i belong to $B(p, \rho) \cap L_p$ where L_p is the slab bounded by the two planes H et H' normal to n_p lying at distance h from p .

Let e_1, \dots, e_s be the edges of U_p that are incident to p and let f_1, \dots, f_s be the facets of U_p that are incident to p . We denote by f_i^* the Voronoi facet dual to f_i and by e_i^* the Voronoi facet dual to e_i . In addition,

write H_i^p for the halfspace bounded by $\text{aff}(e_i^*)$ that contains p . Write $Q(p) = \cap_{i=1,\dots,s} H_i^p$. Observe that the Voronoi cell $V(p)$ of p is included in $Q(p)$.

The intersection of $Q(p)$ with H (resp. H') is a convex polygon P (resp. P'). A vertex of P (resp. P') is the intersection point of H (resp. H') with the affine hull of a Voronoi edge dual to a facet of U_p . Any such Voronoi edge f_i^* contains the center c_i of the surface Delaunay ball circumscribing f_i . As observed above, the c_i belong to $B(p, \rho) \cap L_p$. Writing θ_p for the maximal angle between n_p and an edge f_i^* (or equivalently the normal to f_i), we obtain that all the vertices of P and P' lie in the intersection of the slab L_p with the cylinder W_p of axis (p, n_p) and radius

$$\rho^* \stackrel{\text{def}}{=} \rho + 2h \tan \theta_p = \rho(1 + \varepsilon^+ \tan \theta_p),$$

According to Lemmas 5.6 and 5.7, $\theta_p = \arcsin \varepsilon^+ \sqrt{3} + 2 \arcsin \varepsilon^+$. We therefore have $V(p) \cap L_p \subset W_p \cap L_p$.

Let $C(p) = V(p) \cap \mathcal{S} \cap L_p$. Since \mathcal{S} cannot intersect the two balls B_p and B'_p of radius $\text{lfs}(p)$ that are tangent to \mathcal{S} at p , we deduce from the above discussion that $C(p) \subset W_p \cap L_p$ which implies that, for any x in $C(p)$,

$$\|x - p\| \leq \sqrt{h^2 + \rho^{*2}} = \varepsilon(1 + O(\varepsilon^2)) \text{lfs}(p).$$

2. We prove now that $C(p) = V(p) \cap \mathcal{S}$, which, together with the first part of the proof, proves the lemma. Assume the contrary and let $C'(p)$ be the portion of $V(p) \cap \mathcal{S}$ that lie outside L_p . We claim that $C'(p)$ cannot intersect an edge of $V(p)$. Indeed, otherwise, since $C'(p)$ lies outside the two balls B_p and B'_p , such an intersection point c would be the center of a surface Delaunay ball whose radius is $\|c - p\| > \|c - m_p\| \geq \text{lfs}(c)$ (see Figure ??) where m_p denotes the center of B_p or B'_p closer to c . This would contradict the assumption that E is a loose ε -sample.

We show now that $C'(p) = \emptyset$. Consider the subdivision \mathcal{S}^\dagger of \mathcal{S} induced by $\text{Vor}(E)$. A facet of \mathcal{S}^\dagger is a connected component of the intersection of a cell of $\text{Vor}(E)$ with \mathcal{S} . We distinguish, among the facets of \mathcal{S}^\dagger , those that contain a sample point (we call them the *punctured facets*) and those that do not contain any point of E (called the *orphan facets*). A punctured facet cannot be incident to an orphan facet since any edge of a punctured facet has its two endpoints on edges of $V(p)$ while an orphan facet intersect no edge of $V(p)$. Hence, a connected component F of orphan facets has no boundary and is a connected component of \mathcal{S} . Since F contains no point of E , we get a contradiction with the assumption that E is a loose ε -sample and therefore has points on each connected component of \mathcal{S} . \square

Theorem 5.4 *Any loose ε -sample is an $\varepsilon(1 + O(\varepsilon^2))$ -sample.*

Proof. Let p be a vertex of $\text{Del}_{\mathcal{S}}(E)$. From Lemma 5.10, the intersection of \mathcal{S} with the Voronoi cell $V(p)$ of p is contained in a ball of center p and radius $\eta \text{lfs}(p)$ with $\eta = \varepsilon(1 + O(\varepsilon^2))$. Hence, for any point x in $V(p) \cap \mathcal{S}$, we have $\|x - p\| \leq \eta \text{lfs}(p) \leq \eta (\text{lfs}(p) + \|x - p\|) \leq \eta^+ \text{lfs}(x)$. The theorem follows. \square

In particular, one can check that if E is a loose η -sample of \mathcal{S} with $\eta \leq 0.13$, then it is an ε -sample of \mathcal{S} with $\varepsilon \leq 0.2$ and therefore Theorem 5.3 applies.

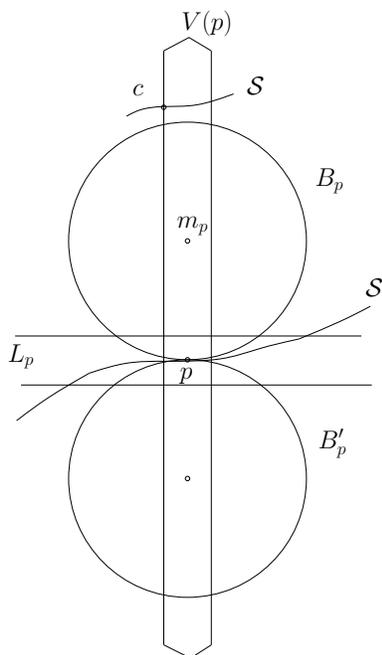


FIG. 5.13 – For the proof of Lemma 5.10.

Exercises 5.1

1. Extend Lemmas 5.2 and 5.3 to the case of manifolds of \mathbb{R}^d .
2. Let E be an ε -sample of a surface \mathcal{S} . Prove that E has points on all the connected components of \mathcal{S} .
3. Prove that the surface Delaunay balls of the restricted Delaunay triangulation $\text{Del}_{\mathcal{S}}(E)$ of an ε -sample E of \mathcal{S} intersect \mathcal{S} along topological disks.
4. Show that the Fréchet distance between $\text{Del}_{\mathcal{S}}(E)$ and \mathcal{S} is $O(\varepsilon^2)$ for a fixed surface \mathcal{S} and $\varepsilon \rightarrow 0$.
5. Show that Theorem 5.3 still holds if we replace $\text{Del}_{\mathcal{S}}(E)$ by the subcomplex of $\text{Del}_{\mathcal{S}}(E)$ consisting of all facets of $\text{Del}_{\mathcal{S}}(E)$ whose dual Voronoi edge intersects \mathcal{S} an odd number of times.
6. Improve the constants in the previous results when the smallest angle of the facets of $\text{Del}_{\mathcal{S}}(E)$ are bounded from below, say larger than $\frac{\pi}{6}$.
7. (Research problem) Extend the approximation results to piecewise smooth surfaces.

5.5 Bibliographical notes

Good introductions to algebraic topology and differential topology can be found in the books of Rotman [31] and Hirsch [25]. The concept of restricted Delaunay triangulation is related to the notion of *nerve* in algebraic topology [20]. Theorem 5.3 is due to Amenta and Bern [1]. For its extension to loose ε -samples and non smooth surfaces see [7] and [11]. The Facet normal lemma 5.7 provides a bound on the normal deviation between the surface \mathcal{S} and the restricted Delaunay triangulation $\text{Del}_{\mathcal{S}}(E)$. Cohen-Steiner and Morvan have further shown that one can estimate the tensor of curvatures from $\text{Del}_{\mathcal{S}}(E)$ [17].

Chapitre 6

Surface meshing

Let S be a surface of \mathbb{R}^3 . If we know a loose ε -sample E for $\varepsilon < 0.1$ of S , then, according to section 5.4.2, $\text{Del}_S(E)$ is a good approximation of S . In this section, we present an algorithm that can construct such a sample and the associated restricted Delaunay triangulation. We restrict the presentation to the case of surfaces that are compact, $C^{1,1}$ and closed. S being compact and of class $C^{1,1}$, for any $x \in S$, $\text{lfs}(x) \geq \text{lfs}_{\text{inf}} > 0$.

6.1 Algorithm

The algorithm is greedy. It inserts points one by one and maintains the current set E , the Delaunay triangulation $\text{Del}(E)$ and its restriction $\text{Del}_S(E)$ to S , and a list L of bad facets of $\text{Del}_S(E)$. Any point that is inserted is the center of the surface Delaunay ball of a bad facet of $\text{Del}_S(E)$. The algorithm stops when there are no more bad facets (which eventually happens as we will see).

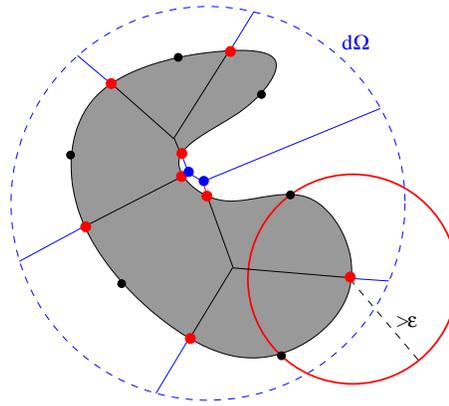
We define a *bad facet* as a facet f of $\text{Del}_S(E)$ that has a circumscribing surface Delaunay ball $B_f = B(c_f, r_f)$ satisfying $r_f > \psi(c_f)$, where ψ is a function defined over S and that satisfies

$$\forall x \in S, \quad \psi(x) \geq \psi_{\text{inf}} > 0.$$

The surface is only queried through an *oracle* that, given a line segment f^* (to be the edge of $\text{Vor}(E)$ dual to a facet f of $\text{Del}_S(E)$), determines whether f^* intersects S and, in the affirmative, returns an intersection point and the value of ψ at this point.

We initialize the construction with a (usually small) set of points $E_0 \subset S$. Three points per connected component of S are sufficient (see exercise 6.12). The algorithm then executes the following loop

while L is not empty {
 take an element f of L ;
 insert c_f in E and in $\text{Del}(E)$;
 update $\text{Del}_S(E)$, i.e.
 remove the facets that are no longer facets of $\text{Del}(E)$;
 add the new facets of $\text{Del}(E)$ whose dual Voronoi edge intersects S ;
 update L , i.e.
 remove the elements of L that are no longer facets of $\text{Del}_S(E)$;
 add the new facets of $\text{Del}_S(E)$ that are bad; }



6.2 Termination

Proposition 6.1 *The algorithm terminates after a finite number of steps.*

Proof. We call *radius of insertion* of a point p the distance of p to the current set E , just before inserting p . We note it ρ_p .

We show that $\rho_p \geq \psi_{\text{inf}}$ for all points $p \in E \setminus E_0$. Indeed, at each step, we insert the center of a surface Delaunay ball that circumscribes a bad facet f . If r_f denotes the radius of this ball and c_f its center, we have $\rho_{c_f} = r_f$. And, as f is a bad facet, $r_f > \psi(c_f) \geq \psi_{\text{inf}} > 0$. The insertion radius of the points of $E \setminus E_0$ is therefore always greater than ψ_{inf} . The balls of radius $\frac{\psi_{\text{inf}}}{2}$ centered at the points of $E \setminus E_0$ thus have disjoint interiors. The surface being compact, we can only insert a finite number of points, which proves that the algorithm terminates. \square

Upon termination, any facet f of $\text{Del}_S(E)$ has a circumscribing surface Delaunay ball B_f of center c_f and radius r_f with $r_f < \psi(c_f)$. To be able to apply theorem 5.3, we need 1. to take $\psi \leq \varepsilon$ lfs, for a sufficiently

small ε . 2. To ensure that $\text{Del}_S(E)$ has at least one vertex on each connected component of S . This can be done by taking in E_0 three points per component of S that are sufficiently close (see exercise 6.1).

6.3 Optimality

We bound the number of points inserted by the algorithm when ψ is 1-Lipschitz, i.e. $\psi(x) \leq \psi(y) + \|x - y\|$ for any $x, y \in S$.

Proposition 6.2 *If $\psi(x)$ is 1-Lipschitz and at most $\frac{1}{2} \text{lfs}(x)$, the number of points inserted by the algorithm after the initialization phase is $n = |E \setminus E_0| = O\left(\int_S \frac{dx}{\psi^2(x)}\right)$.*

Proof. Let $\tau(x) = \inf\{r : |B(x, r) \cap E| \geq 2\}$ and $B_p = B(p, \frac{\tau(p)}{2})$, $p \in E$. It is easy to see that τ is 1-Lipschitz.

$$\begin{aligned} \int_S \frac{dx}{\tau^2(x)} &\geq \sum_{p \in E \setminus E_0} \int_{(B_p \cap S)} \frac{dx}{\tau^2(x)} && \text{(the balls } B_p \text{ have disjoint interiors)} \\ &\geq \frac{4}{9} \sum_{p \in E \setminus E_0} \frac{\text{area}(B_p \cap S)}{\tau^2(p)} && (\tau(x) \leq \tau(p) + \|p - x\| \leq \frac{3}{2} \tau(p)) \\ &\geq \frac{4}{9} \sum_{p \in E \setminus E_0} \frac{3\pi}{16} = \frac{\pi}{12} n \end{aligned}$$

The last inequality comes from the fact that $\text{area}(B_p \cap S) \geq \frac{3\pi\tau^2(p)}{16}$. To prove the latter, we first observe that $\tau(p)$ is no more than twice the radius of the surface Delaunay balls incident to p . Hence, if f is a facet of $\text{Del}_S(E)$ incident to p , we have

$$\tau(p) \leq 2r_f \leq 2\psi(c_f) \leq 2\psi(p) + \|p - c_f\| = 2\psi(p) + r_f \leq 4\psi(p).$$

Now consider the two balls B and B' tangent to S at p of radius $\frac{1}{2}\tau(p)$. Since $\frac{\tau(p)}{2} \leq 2\psi(p) \leq \text{lfs}(p)$, $B_p \cap S$ is a topological disk by lemma 5.5. Moreover, this disk cannot intersect the interiors of the two balls B and B' , and its boundary lies on ∂B_p . Hence the area of $B_p \cap S$ is larger than the area of the disk bounded by the circle $\partial B \cap \partial B_p$. Denoting by ρ its radius, we have $\rho = \frac{\tau(p)\sqrt{3}}{4}$. The claim follows.

Using the Lipschitz property of τ and ψ , $\forall x \in B_p$, we have

$$\frac{1}{2} \tau(p) \leq \tau(p) - \|x - p\| \leq \tau(x) \tag{6.1}$$

$$\psi(x) \leq \psi(p) + \|p - x\| \leq \psi(p) + \frac{\tau(p)}{2} \tag{6.2}$$

Let q be a point of E closest to p : $\tau(p) = \|p - q\|$. According to the algorithm, $\|p - q\| > \psi(p)$ or $\|p - q\| > \psi(q)$ depending whether p has been inserted after or before q . In both cases, due to the Lipschitz property of ψ , we have

$$\tau(p) = \|p - q\| \geq \psi(p) - \|p - q\| > \frac{1}{2} \psi(p) \quad (6.3)$$

From inequalities (6.1), (6.2) and (6.3), we get $\psi(x) \leq 5\tau(x)$, from which we deduce

$$n \leq \frac{300}{\pi} \int_S \frac{dx}{\psi(x)^2}.$$

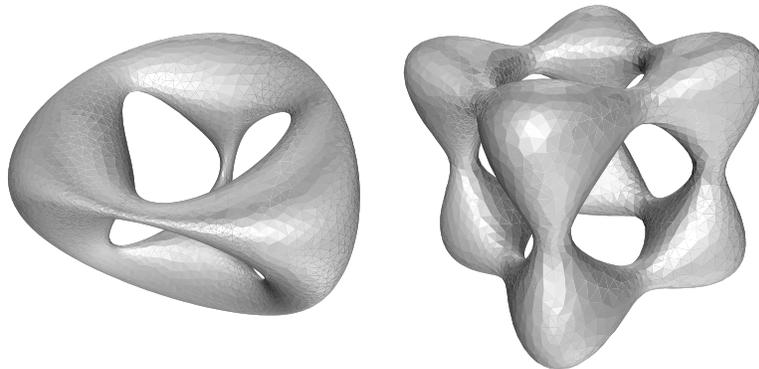
□

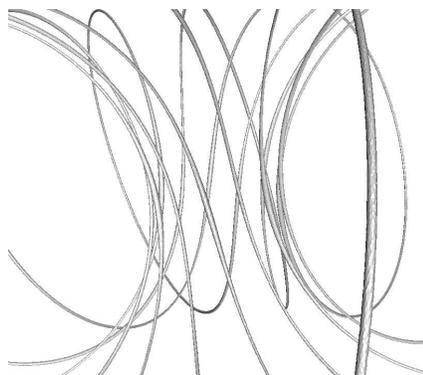
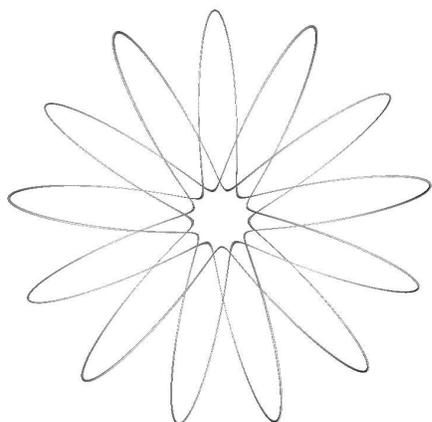
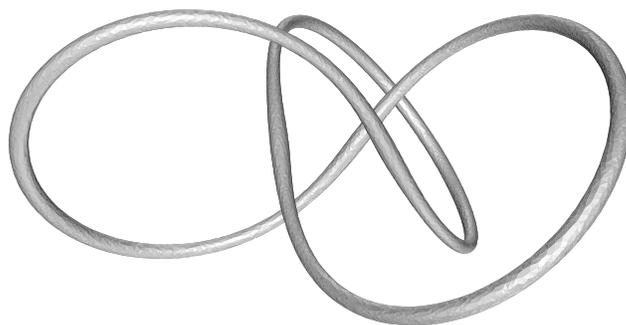
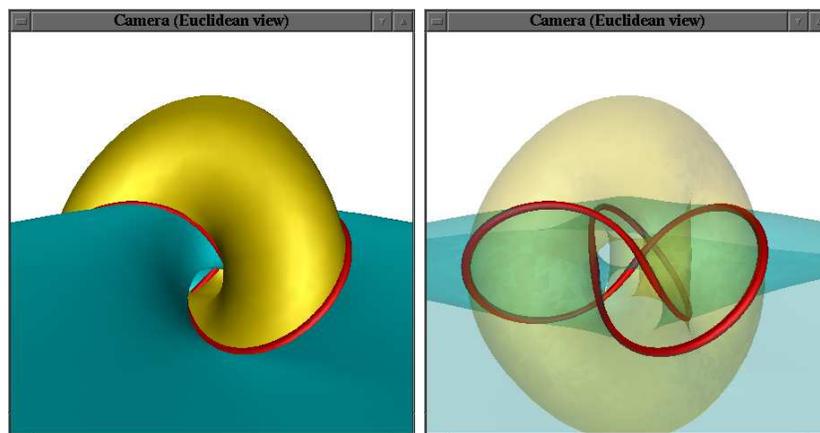
We sum up the results of this section in the following theorem. The proof of optimality is left as an exercise.

Theorem 6.1 *Given a compact oriented C^2 surface S and a positive function ψ on S , one can compute a loose ε -sample E of S for ψ . If ψ is 1-Lipschitz and at most $\frac{1}{2} \text{lfs}(x)$, the size of the sample is $O\left(\int_S \frac{dx}{\psi^2(x)}\right)$, which is optimal. If $\psi < 0.1 \text{lfs}$, $\text{Del}_S(E)$ is isotopic to S and the isotopy does not move the points by more than $\sup_{x \in S} \psi(x)$.*

6.4 Experimental results

The algorithm has been implemented by S. Oudot. We present some experimental results on algebraic surfaces. The first two surfaces have genus respectively 3 and 5. We then show the standard trefoil knot and a more intricate knot in "sausage" format. A thorough discussion of the implementation of the algorithm and other experimental results can be found in [10].





Exercises 6.1

1. Adapt the algorithm to the case where S is a curve of \mathbb{R}^d .

2. Assume that ψ is a Lipschitz function. Show that a facet f circumscribed by Delaunay ball B_f of center c_f and radius r_f with $r_f \leq \frac{1}{3}\psi(c_f)$ will remain a restricted Delaunay facet throughout the course of the algorithm.
3. We modify the definition of a bad facet so as to eliminate facets with small angles. A facet f of $\text{Del}_S(E)$ is now said to be bad if r_f is too large (as above) or if one of its angles is smaller than $\frac{\pi}{6}$. Show that the algorithm still terminates.
4. Show that the bound in proposition 6.2 is tight.
5. Since S is compact and lfs is continuous, there exists $x \in S$ such that $\text{lfs}(x) = \inf_{y \in S} \text{lfs}(y)$. Show that either $\text{lfs}(x)$ is the minimal radius of curvature at x (in which case, one of the two medial spheres passing through x is osculating S), or the line normal to S at x intersects S at another point z and is also the line normal to S at z , and the sphere with diameter xz is a medial ball. Show then how to compute $\inf_{y \in S} \text{lfs}(y)$.
6. Show that the algorithm remains valid if we replace the restricted Delaunay triangulation by the subcomplex of $\text{Del}_S(E)$ consisting of its bipolar facets (see exercise ??2).

6.5 Bibliographical notes

The meshing algorithm presented in this chapter is due to Boissonnat and Oudot [7]. The paradigm of Delaunay refinement has been first proposed by Ruppert for meshing planar domains [32].

Chapitre 7

Complexity of Voronoi diagrams of surface samples

7.1 Introduction

We have seen in chapter 3 that the complexity of the Delaunay triangulation of n points of \mathbb{R}^d can be quadratic in the worst-case. In this chapter, we will see that, if the points form a good sample of a fixed set of non intersecting polygons of \mathbb{R}^3 , the complexity of the Delaunay triangulation is linear.

In this chapter, we consider the case of points distributed on a fixed finite set of interior-disjoint planar regions whose total area is positive and whose total perimeter is finite. This includes the case of polyhedral surfaces. Under a mild uniform sampling condition (depending on a parameter κ), we show that the complexity of the Delaunay triangulation of the points is linear when κ is a constant. Our bound is deterministic. The constants are explicitly given and depend on κ and on the number of planar regions C_S , the total area A_S and the total perimeter L_S of the regions. More precisely, our main result states that the number of Delaunay edges is at most :

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

where κ is a constant characterizing the set of points. Our bound holds for any $n > 0$.

7.2 Definitions and notations

7.2.1 Notations

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

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

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

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

7.2.2 Polyhedral surfaces

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

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

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

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

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

Definition 7.1 *Let $\varepsilon \geq 0$. The ε -regular zone of a facet $F \subset \mathcal{S}$ is $F \ominus \varepsilon$. The ε -regular zone of \mathcal{S} is the union of the ε -regular zones of its facets. The ε -singular zone of F (resp. \mathcal{S}) is the set of points that do not belong to the ε -regular zone of F (resp. \mathcal{S}).*

Observe that the 0-singular zone of \mathcal{S} consists exactly of the edges of \mathcal{S} .

7.2.3 Sample

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

Definition 7.2 *Let \mathcal{S} be a polyhedral surface. $P \subset \mathcal{S}$ is said to be a (ε, κ) -sample of \mathcal{S} if and only if for every facet F of \mathcal{S} and every point $x \in F$:*

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

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

In the rest of the paper, P denotes a (ε, κ) -sample of \mathcal{S} and we provide asymptotic results when the sampling density increases, i.e. when ε tends to 0. As already mentioned, we consider κ and the surface \mathcal{S} (and, in particular, the three quantities $C_{\mathcal{S}}$, $A_{\mathcal{S}}$ and $L_{\mathcal{S}}$) to be fixed and not to depend on ε .

7.3 Preliminary results

\mathcal{S} designates a polyhedral surface and $P \subset \mathcal{S}$ a (ε, κ) -sample of \mathcal{S} . We denote by $\sharp(A)$ the number of elements of A . Let $n(R) = \sharp(P \cap R)$ be the number of sample points in the region $R \subset \mathcal{S}$. Let $n = \sharp(P)$ be the total number of sample points. We first establish two propositions relating $n(R)$ and n . We start with the following lemma :

Lemma 7.1

$$\frac{A_{\mathcal{S}}}{4\pi\varepsilon^2} \leq n$$

Proof. Let F be a facet of \mathcal{S} . Let $\{D(x_i, \varepsilon)\}_{i \in \{1, \dots, \lambda\}}$ be a maximal set of λ non-intersecting disks lying inside $F \oplus \varepsilon$. Because the set of disks is maximal, no other disk can be added without intersecting $\cup_{i=1}^{\lambda} D(x_i, \varepsilon)$. This implies that no point m of F is at distance greater than 2ε from a point x_i (see Figure 7.1). Therefore, $\{D(x_i, 2\varepsilon)\}_{i \in \{1, \dots, \lambda\}}$ is a covering of F . We have $\frac{\text{area}(F)}{4\pi\varepsilon^2} \leq \lambda$. Because of our sampling condition, every disk $D(x_i, \varepsilon)$ contains at least one sample point. Therefore, $\lambda \leq n(F)$ and

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

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

Lemma 7.2 *Let F be a facet of S . For any $R \subset F$, we have :*

$$n(R) \leq \frac{4\kappa \text{ area} \left(R \oplus \frac{\varepsilon}{2} \right)}{\pi \varepsilon^2}$$

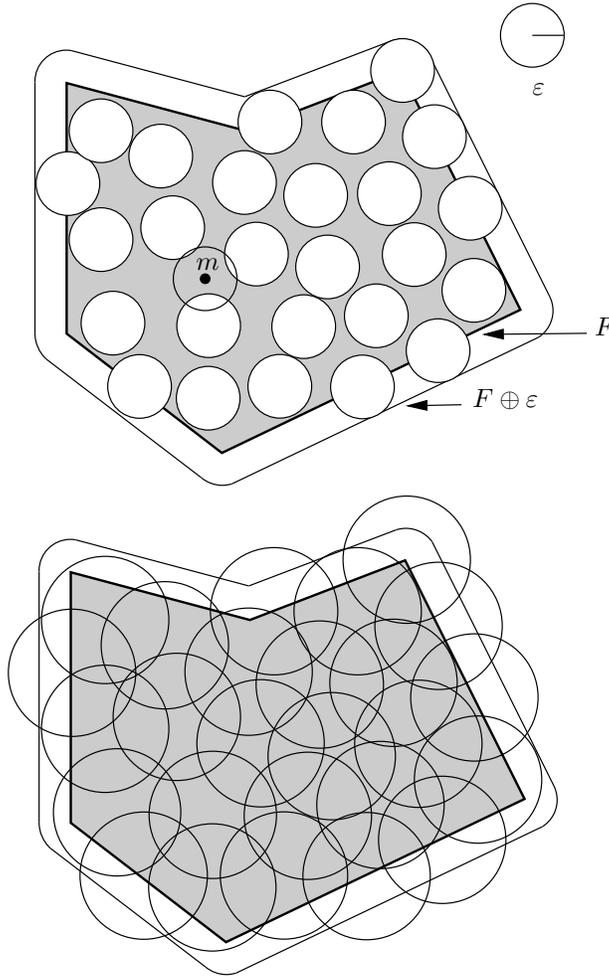


FIG. 7.1 – A maximal set of non-intersecting disks contained in $R \oplus \varepsilon$ and the corresponding covering of R obtained by doubling the radii of the disks.

Proof. Let $\{D(x_i, \frac{\varepsilon}{2})\}_{i \in \{1, \dots, \lambda\}}$ be a maximal set of λ non-intersecting disks lying inside $R \oplus \frac{\varepsilon}{2}$. Because the set of disks is maximal, no other disk can be added without intersecting $\cup_{i=1}^{\lambda} D(x_i, \frac{\varepsilon}{2})$. This implies that no

point m of R is at distance greater than ε from a point x_i . Therefore, $\{D(x_i, \varepsilon)\}_{i \in \{1, \dots, \lambda\}}$ is a covering of R . We have :

$$n(R) \leq \kappa \lambda \leq \kappa \times \frac{\text{area}(R \oplus \frac{\varepsilon}{2})}{\frac{\pi \varepsilon^2}{4}}$$

□

Proposition 7.1 *Let F be a facet of \mathcal{S} . For any $R \subset F$, we have :*

$$n(R) \leq 16\kappa \frac{\text{area}(R \oplus \frac{\varepsilon}{2})}{A_S} n$$

Proof. By Lemma 7.1, we have :

$$\frac{A_S}{4\pi\varepsilon^2} \leq n \tag{7.1}$$

We apply Lemma 7.2 to bound $n(R)$ from above.

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

Eliminating ε from the two inequalities yields the result. □

Proposition 7.2 *Let F be a facet of \mathcal{S} . Let $\Gamma \subset F$ be a curve contained in F . Let $a > 0$. We have :*

$$n(\Gamma \oplus a\varepsilon) \leq \frac{(4a+1)^2}{a} \kappa \frac{\text{length}(\Gamma)}{\varepsilon} \leq \frac{2(4a+1)^2}{a} \sqrt{\pi} \kappa \frac{\text{length}(\Gamma)}{\sqrt{A_S}} \sqrt{n}$$

Proof. Arguing as in the proof of Lemma 7.2, we see that the region $\Gamma \oplus a\varepsilon$ can be covered by $\frac{\text{length}(\Gamma)}{a\varepsilon}$ disks of radius $2a\varepsilon$ centered on Γ and contained in the supporting plane of F .

Applying Lemma 7.2 to a disk R with radius $2a\varepsilon$, we get :

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

Therefore, we have :

$$n(\Gamma \oplus a\varepsilon) \leq \kappa \frac{(4a+1)^2}{a} \frac{\text{length}(\Gamma)}{\varepsilon}$$

From Lemme 7.1, we get :

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

Combining the two inequalities leads to the result. □

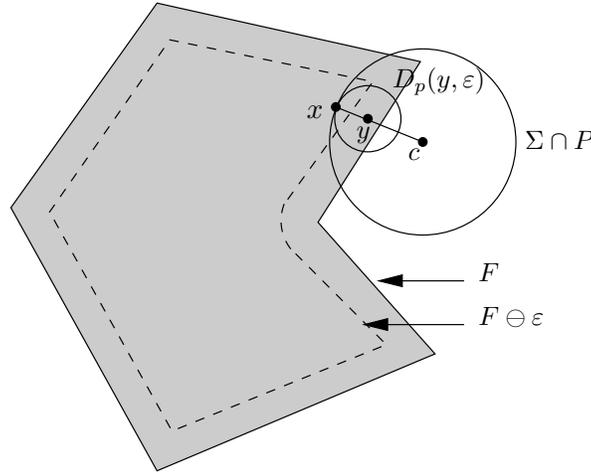


FIG. 7.2 – Assume Σ is an empty sphere passing through a point $x \in F \ominus \varepsilon$ and intersecting the supporting plane of F in a circle of radius greater than ε . Then, Σ contains an empty disk $D_H(y, \varepsilon)$ centered on F .

Lemma 7.3 *Let x be a sample point in the ε -regular zone of \mathcal{S} . Let H be the supporting plane of the facet through x . Any empty sphere passing through x intersects H in a circle whose radius is less than ε .*

Proof. The proof is by contradiction. Let H be the supporting plane of F . Consider an empty sphere Σ passing through x and intersecting H along a circle of radius greater than ε (see Figure 7.2). Let c be the center of this circle. Let y be the point on the segment $[xc]$ at distance ε from x . Because x belongs to the ε -regular zone of F , $y \in F$. The empty sphere Σ encloses the disk $D_H(y, \varepsilon)$. Therefore, $D_H(y, \varepsilon)$ is an empty disk of H , centered on F and of radius ε , which contradicts our sampling condition. \square

7.4 Counting Delaunay edges

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

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

$$t_p = 2e_p - 4$$

since the boundary of the union of those tetrahedra is a simplicial polyhedron of genus 0. Using the same

argument, if p lies on the boundary of the convex hull, we have :

$$t_p < 2e_p - 4$$

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

$$t < e - n.$$

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

We distinguish three types of Delaunay edges : those with both endpoints in the ε -regular zone, those with both endpoints in the ε -singular zone and those with an endpoint in the ε -regular zone and the other in the ε -singular zone. They are counted separately in the following subsections,

We denote by P_s the set of sample points in the ε -singular zone of \mathcal{S} .

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

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

Lemma 7.4 *Let x be a sample point in the ε -regular zone and F the facet that contains x . x has at most κ Delaunay neighbours in F .*

Proof. By Lemma 7.3, any empty sphere passing through x intersects F in a circle whose radius is less than ε . Therefore, the Delaunay neighbours of x on F are at distance at most 2ε from x . By assumption, the disk centered at x with radius 2ε contains at most κ points of P . \square

Lemma 7.5 *Let x be a sample point in the ε -regular zone of a facet F . Let $F' \neq F$ be another facet of \mathcal{S} . x has at most κ Delaunay neighbours in the ε -regular zone of facet F' .*

Proof. Refer to Figure 7.3. H and H' are the supporting planes of F and F' , y is a Delaunay neighbour of x in the ε -regular zone of F' and Sigma is an empty sphere passing through x and y . Let B be the closed ball whose boundary is Sigma . B intersects the planes H and H' along two disks whose radii are respectively r and r' . By Lemma 7.3, $r \leq \varepsilon$ and $r' \leq \varepsilon$.

Let M be the bisector plane of H and H' . Let x' and y' be the points symmetric to x and y with respect to M . Consider the sphere Sigma_0 centered on M and passing through the four points x, x', y and y' . Let B_0 be the closed ball whose boundary is Sigma_0 . B_0 intersects H and H' along two disks D_0 and D'_0 of the same radius r_0 . We claim that $r_0 \leq \max(r, r')$. Indeed, let v_0 be the center of Sigma_0 and v be the center of Sigma . Let M_{xy} (resp. $M_{x'y'}$) be the bisector plane of x and y (resp. of x' and y'). Observe that

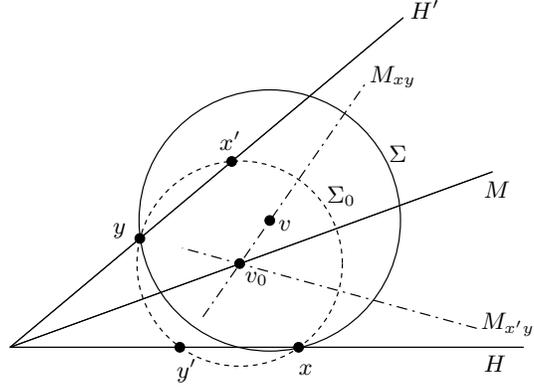


FIG. 7.3 – Any sphere passing through x and y intersects one of the two planes H or H' in a circle whose diameter is at least $\|x' - y\|$.

$v_0 \in M_{xy} \cap M_{x'y'}$ and $v \in M_{xy}$. If $v \in M_{xy} \cap M_{x'y'}$, $r_0 = r = r'$ and the claim is proved. Otherwise, v must belong to one of the two open halfspaces limited by $M_{x'y'}$. If v belongs to the halfspace that contains x' , B encloses D'_0 and therefore $r_0 \leq r'$ while in the second it encloses D_0 and $r_0 \leq r$.

We therefore have :

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

and consequently :

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

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

Proposition 7.3 *There are at most $\frac{1}{2}C_S\kappa n$ Delaunay edges with both endpoints in the ε -regular zone of \mathcal{S} .*

Proof. The surface S has C_S facets. Therefore, by Lemmas 7.4 and 7.5, a point x in the ε -regular zone of \mathcal{S} has at most $C_S\kappa$ Delaunay neighbours. \square

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

In this section, we count the Delaunay edges joining two points in the ε -singular zone (see Figure 7.4).

Proposition 7.4 *The number of Delaunay edges with both endpoints in the ε -singular zone is less than*

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

Proof. By Proposition 7.2, the number $\#(P_s)$ of sample points in the ε -singular zone is at most

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

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

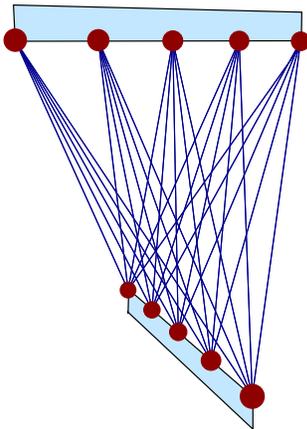


FIG. 7.4 – Example of a Delaunay triangulation of m points having a quadratic number of edges. Even if such a configuration can occur for a subset of the sample points, the number m of sample points involved in this configuration is $O(\sqrt{n})$. Therefore, the number of Delaunay edges involved in this configuration is $O(n)$.

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

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

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

Let H be a plane in \mathbb{R}^3 and $X \subseteq \mathbb{R}^3$ be a finite set of points. We assign to each point x of X the region $V(x) \subset H$ consisting of the points $h \in H$ for which the sphere tangent to H at h and passing through x encloses no point of X (see Figure 7.5). In other words, if $R(h, x)$ denotes the radius of the sphere tangent to H at h and passing through x , we have :

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

It is easy to see that the set of all $V(x)$, $x \in X$, is a subdivision of H which we denote $\mathcal{V}_H(X)$ (see Figure 7.8). The diagram $\mathcal{V}_H(X)$ is a multiplicatively-weighted power Voronoi diagram. Let \mathcal{P}_x be the paraboloid

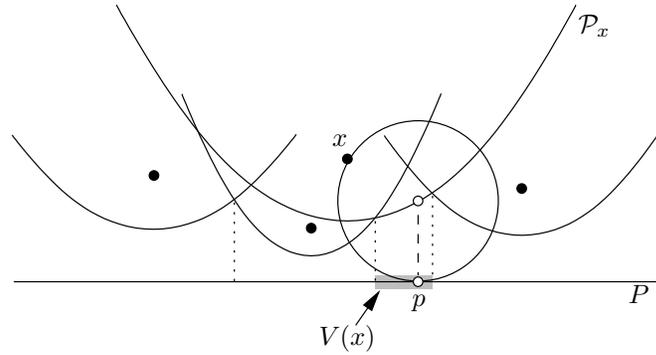


FIG. 7.5 – The cell $V(x)$ is the set of contact points between a plane H and a sphere passing through x and tangent to H . The part of the paraboloid \mathcal{P}_x on the lower envelope of the paraboloids projects to the cell $V(x)$.

of revolution with focus x and director plane H . The paraboloid \mathcal{P}_x consists of the centers of the spheres passing through x and tangent to H . Assume that the points X are all located above plane H . If not, we replace x by the point symmetric to x with respect to H , which does not change $\mathcal{V}_H(X)$. Let us consider the lower envelope of the collection of paraboloids $\{\mathcal{P}_x\}_{x \in X}$. Cell $V(x)$ is the projection of the portion of the lower envelope contributed by \mathcal{P}_x (see Figures 7.5 and 7.8).

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

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

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

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

It follows that the edges $E(x, y)$ of $V(x)$ are circle arcs that we call *convex* or *concave* with respect to x depending whether the disk $D(x, y)$ (whose boundary contains $E(x, y)$) is labelled $+$ or $-$ (see Figure 7.6). Observe that the convex edges of $V(x)$ are included in the boundary of the convex hull of $V(x)$.

Proposition 7.5 *The number of Delaunay edges with one endpoint in the ε -regular zone and the other in the ε -singular zone is at most :*

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

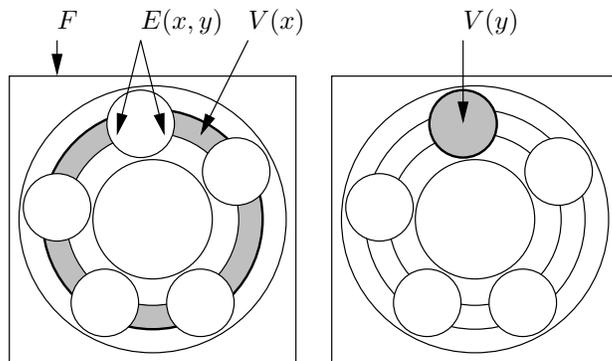


FIG. 7.6 – The bold edges are the convex edges of the shaded cells. The edge $E(x,y)$ which is concave with respect to x is convex with respect to y . The convex edges of a cell lie on the boundary of its convex hull.

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

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

We first show that the Delaunay neighbours of x that belong to the ε -regular zone of F belong to $V(x) \oplus 2\varepsilon$. Consider a Delaunay edge (xf) with $x \in P_s$, $x \notin H$ and $f \in P_s \cap (F \ominus \varepsilon)$. Let $\mathcal{S}igma$ be an empty sphere passing through x and f , v its center (see Figure 7.7). By Lemma 7.3, $\mathcal{S}igma$ intersects H in a circle whose radius r is less than ε . For a point c on the segment $[vx]$, we denote by $\mathcal{S}igma_c$ the sphere centered at c and passing through x . Because $\mathcal{S}igma$ encloses $\mathcal{S}igma_c$, $\mathcal{S}igma_c$ is an empty sphere. For $c = v$, $\mathcal{S}igma_c$ intersects H . For $c = x$, $\mathcal{S}igma_c$ does not intersect H . Consequently, there exists a position of c on $[vx]$ for which $\mathcal{S}igma_c$ is tangent to H . Let $p = \mathcal{S}igma_c \cap H$ for such a point c . We have $p \in V(x)$ and $\|p - f\| \leq 2r \leq 2\varepsilon$. Hence, $f \in V(x) \oplus 2\varepsilon$. Now, let us consider a Delaunay edge (xf) with $x, f \in P_s \cap H$. Applying Lemma 7.3 leads to $f \in V(x) \oplus 2\varepsilon$.

Let N_F be the number of Delaunay edges between P_s and $F \ominus \varepsilon$. We have, using the fact that \mathcal{V}_F is a subdivision of F and Proposition 7.2 :

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

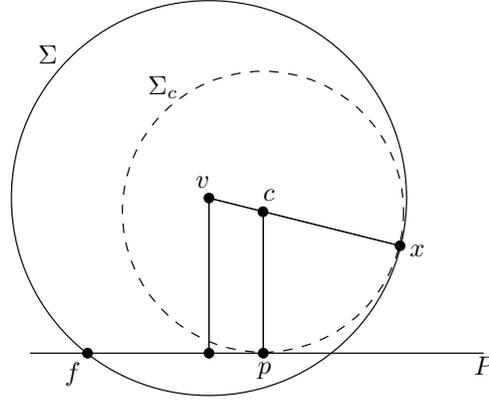


FIG. 7.7 – Every sphere Σ passing through x and $f \in H$ contains a sphere Σ_c passing through x and tangent to H .

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

The convex edges of x are contained in the boundary of the convex hull of $V(x)$. Since $V(x) \subset F$, the length of the boundary of the convex hull of $V(x)$ is at most the length of ∂F . Consequently :

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

Since, by Proposition 7.2, $\#(P_s) \leq 50\sqrt{\pi} \kappa \frac{L_S}{\sqrt{A_S}} \sqrt{n}$, we have :

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

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

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

□

7.4.4 Main result

We sum up our results in the following theorem :

Theorem 7.1 *Let \mathcal{S} be a polyhedral surface and P a (ε, κ) -sample of \mathcal{S} of size $\sharp(P) = n$. The number of edges in the Delaunay triangulation of P is at most :*

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

Notice that our bound holds for any $n > 0$. It should be observed also that the bound does not depend on the relative position of the facets (provided that their relative interiors do not intersect). In particular, it does not depend on the dihedral angles between the facets. Notice also that the bound is not meaningful when $A_S = 0$, which is the case of the quadratic example in Figure 7.4.

7.5 Bibliographical notes

This chapter is taken from a paper by Attali and Boissonnat [2]. The extension to the case of smooth surfaces is more difficult. Erickson [21] has exhibited a lower bound $\Omega(n\sqrt{n})$ while Attali, Boissonnat and Lieutier have provided a $O(n \log n)$ upper bound for *generic* surfaces [3].

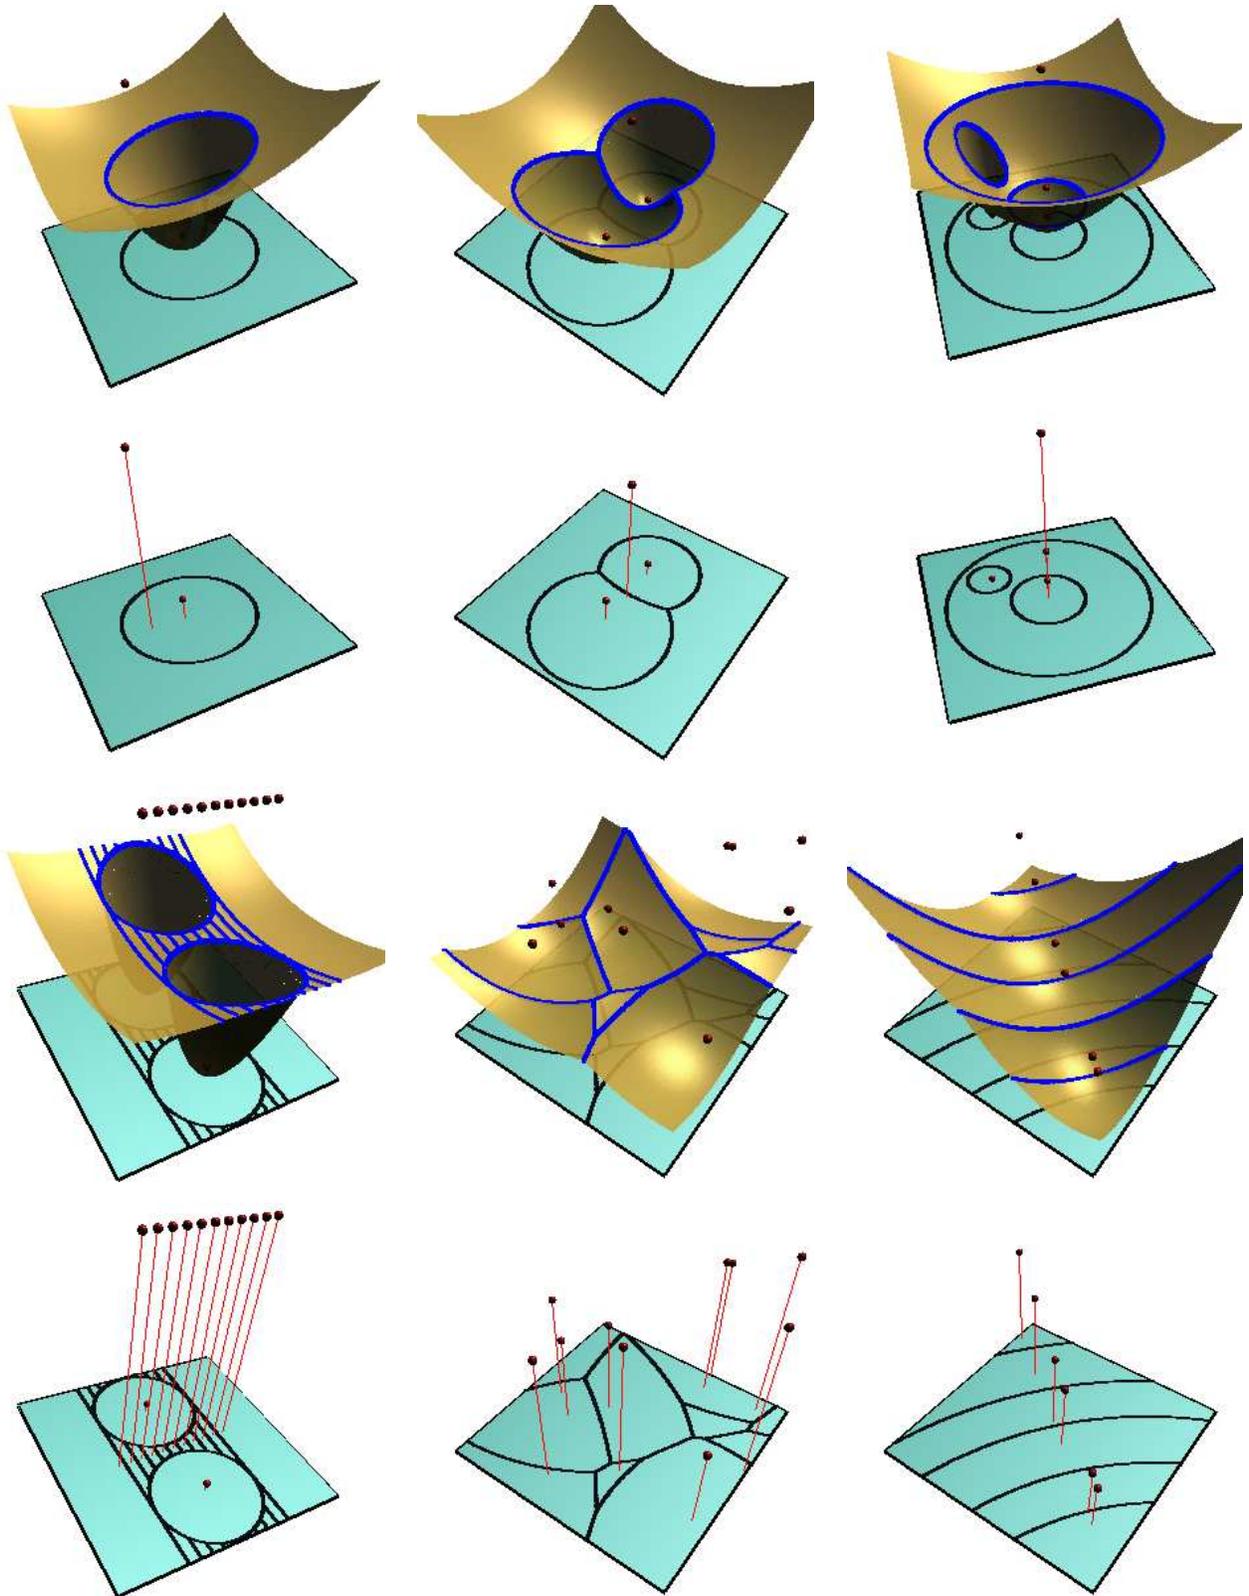


FIG. 7.8 – Decomposition of a facet F into cells for different set of points P_s . The lower envelope of the paraboloid $\{\mathcal{P}_x\}_{x \in P_s}$ has been represented. The red spheres represent the points of P_s and the red lines materialize the projection of the points of P_s on the plane H . The bisector of two points is a circle. The projection of x on H do not belong necessary to its cell. The decomposition of F can have a quadratic number of edges.

Chapitre 8

Local systems of coordinates Scattered data interpolation

Dans tout ce chapitre, E désigne un ensemble fini de points de \mathbb{R}^d .

Definition 8.1 On appelle système de coordonnées associé à E un ensemble de n fonctions continues λ_i , $i = 1, \dots, n$ vérifiant pour tout $x \in \text{conv}(E)$

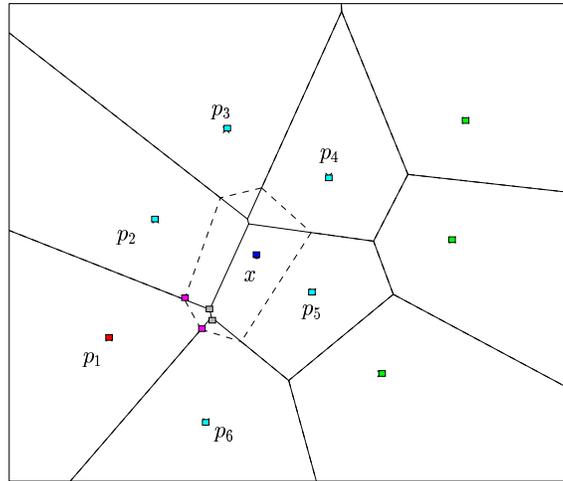
1. $x = \sum_i \lambda_i(x) p_i$
2. $\lambda_i(p_j) = 1$ si $j = i$ et 0 si $j \neq i$
3. $\lambda_i(x) \geq 0$, $\sum_i \lambda_i(x) = 1$

Une façon d'obtenir un système de coordonnées est de trianguler E , de localiser x dans la triangulation $T(E)$, c'est-à-dire d'identifier un simplexe t_x qui contient x et de définir $\lambda_i(x)$ comme les coordonnées barycentriques de x dans le simplexe t_x . On appelle ces coordonnées les coordonnées barycentriques de x dans $T(E)$. Deux remarques motivent les développements qui vont suivre. Tout d'abord, ces coordonnées dépendent du choix d'une triangulation et ne sont donc pas intrinsèques. D'autre part, elles sont continues mais pas continuellement différentiables aux bords des simplexes.

8.1 Voisins naturels et systèmes de coordonnées associés

Soit x un point de \mathbb{R}^d . On note $E^+ = E \cup \{x\}$ et $V^+(x)$ la région de Voronoï de x dans $\text{Vor}(E^+)$. On dit que p_i est un *voisin naturel* de x si $V^+(x)$ intersecte la région de Voronoï $V(p_i)$ de $\text{Vor}(E)$. Si p_i est un voisin naturel de x , on note $W_i(x)$ l'intersection de $V^+(x)$ avec $V(p_i)$ et $V_i(x)$ la facette commune à $V^+(x)$

et $V^+(p_i)$. $W_i(x)$ est l'ensemble des points de \mathbb{R}^d qui ont x comme plus proche voisin et p_i comme second plus proche voisin dans E^+ . $V_i(x)$ est la facette de $W_i(x)$ constituée des points à égale distance de x et p_i .



De manière équivalente, p_i est un voisin naturel de x si xp_i est une arête de $\text{Del}(E^+)$ (puisque'il existe un point c qui a x et p_i comme plus proches voisins, la sphère centrée en c et passant par x et p_i n'englobe aucun point de E). Les voisins naturels de x sont donc les points de E auxquels serait relié x si on l'insérait dans la triangulation de Delaunay de E . Ce sont aussi les sommets des tétraèdres de $\text{Del}(E)$ dont les sphères circonscrites englobent x .

Notons Nat la relation d'équivalence qui relie deux points de \mathbb{R}^d s'ils ont mêmes voisins naturels. Il découle de la discussion précédente que les classes d'équivalence de Nat sont les faces de l'arrangement des sphères de Delaunay¹.

Definition 8.2 Soit $v_i(x) = \text{vol}(V_i(x))$, $\bar{v}_i(x) = \frac{v_i(x)}{\|x-p_i\|}$ et $w_i(x) = \text{vol}(W_i(x))$. On appelle coordonnées de Laplace les n fonctions $\lambda_1, \dots, \lambda_n$ définies par $\lambda_i(x) = \frac{\bar{v}_i(x)}{\sum_{i=1}^n \bar{v}_i(x)}$ si $x \notin E$, $\lambda_i(p_i) = 1$ et $\lambda_i(p_j) = 0$ si $j \neq i$. On appelle coordonnées de Sibson les n fonctions $\varsigma_i(x) = \frac{w_i(x)}{\sum_{i=1}^n w_i(x)}$, $i = 1, \dots, n$.

Lemma 8.1 Les coordonnées de Laplace et les coordonnées de Sibson sont continues en tout point x de $\text{conv}(E)$. Les coordonnées de Sibson sont continuellement différentiables en tout point x de $\text{conv}(E)$ à l'exception des points de E .

Preuve :

¹C'est-à-dire la partition du plan induite par l'union des sphères.

1. Considérons tout d'abord la fonction $v_i(x)$. Le volume de $F^+(x, p_i)$ varie continuellement sauf aux points x où le plan médiateur de x et p_i contient une facette de $V(p_i)$, c'est-à-dire quand x est un des voisins naturels de p_i . Il s'ensuit que les v_i et donc aussi les λ_i sont continus en tout point $x \notin E$. Montrons que les λ_i sont également continus aux points de E . Quand x tend vers p_i , $\bar{v}_i(x)$ tend vers l'infini tandis que tous les autres $\bar{v}_j(x)$ restent bornés. Les $\lambda_j(x)$ ($j \neq i$) tendent donc vers 0 et $\lambda_i(x)$ vers 1. Comme par définition λ_i vaut 1 au point p_i et 0 aux autres points de E , les λ_i sont donc également continus aux points de E .

2a. A l'intérieur de chaque cellule de l'arrangement des sphères de Delaunay, les voisins naturels sont fixés et $w_i(x)$ est de classe C^∞ .

b. Considérons un point x sur une des sphères de Delaunay σ distinct des p_i . Au vu du point a, pour que $w_i(x)$ soit continue en x , il suffit que la restriction de w_i sur une droite δ coupant transversalement σ soit continue. De même, pour que la fonction $w_i(x)$ soit continuellement différentiable, il suffit qu'elle le soit sur δ .

c. Soit δ_i une demi-droite issue de p_i dont on exclut p_i . Considérons la restriction $w_{i|\delta_i}$ de w_i à δ_i . On paramètre δ_i par $t \in [0, +\infty[$. Soit $x = x(t_x)$ un point sur δ_i . Soit $t_m \geq 0$ la plus petite valeur de t telle que, $\forall t > t_m$, $v_i(x(t)) = 0$. Soit $I = [t_x, t_m[$. D'après le point 1, $v_i(x(t))$ est continue en tout point t de I (noter cependant que $v_i(x(t))$ n'est pas continu en t_m si $x(t_m)$ est un point de E dont la cellule de Voronoï est adjacente à celle de p_i). On observe que $w_{i|\delta_i}(x)$ est obtenue en intégrant $v_i(x(t))$ sur I . La continuité de $v_i(x(t))$ sur I entraîne que $w_{i|\delta_i}$ est continuellement différentiable en x . On conclut que $w_{i|\delta_i}$ est continuellement différentiable en tout point x de δ_i qui n'est pas un voisin naturel de p_i .

d. Il suffit maintenant de prendre comme droite δ (cf. point b) une droite passant par un des p_i appartenant à σ . On conclut que w_i est continuellement différentiable sauf aux points de E . Il en va clairement de même pour ς_i .

e. Reste à montrer la continuité de ς_i aux points de E . Celle-ci découle du fait que, quand x tend vers p_i , tous les $w_j(x)$ sauf $w_i(x)$ tendent vers 0. Il s'ensuit que les $\varsigma_i(x)$ sont continus aux points de E et qu'on a $\varsigma_i(p_i) = 1$ et $\varsigma_i(p_j) = 0$ si $j \neq i$. On notera que $w_i(x)$ n'est pas continu en p_i .

□

Lemma 8.2 $x = \sum_i \lambda_i(x) p_i$.

Preuve : C'est une conséquence directe du lemme 4.2 appliqué à la région de Voronoï $V^+(x)$:

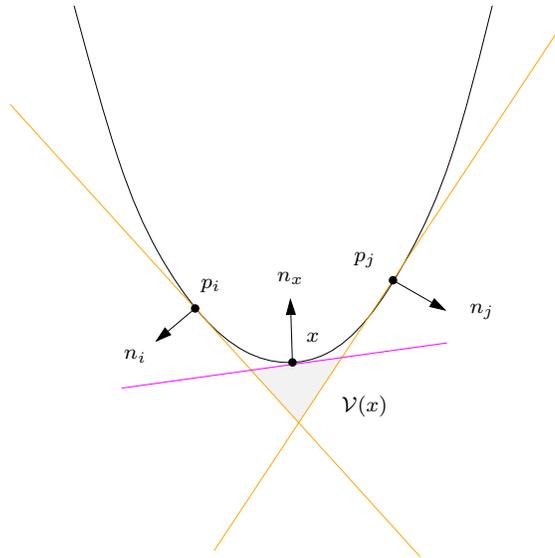
$$\sum_i v_i \frac{x - p_i}{\|x - p_i\|} = 0 \iff \left(\sum_i \frac{v_i}{\|x - p_i\|} \right) x = \sum_i \frac{v_i}{\|x - p_i\|} p_i.$$

□

Lemma 8.3 $x = \sum_i \varsigma_i(x) p_i$.

Preuve : Nous revenons à l'espace des sphères qu'on identifie à \mathbb{R}^{d+1} (cf. paragraphe 3.2). Pour éviter les confusions, on rebaptise p le point x du lemme. On note I l'ensemble des indices des voisins naturels de p . Considérons le polyèdre \mathcal{V}_p intersection des demi-espaces $h_{p_i}^+$, $i \in I$, d'équations $h_{p_i}^+ : x_{d+1} \geq 2p_i \cdot x - p_i \cdot p_i$, et du demi-espace h_p^- d'équation $h_p^- : x_{d+1} \leq 2p \cdot x - p \cdot p$. On note f_i la facette de \mathcal{V}_p portée par $h_{p_i}^+$ et f_p la facette de \mathcal{V}_p portée par h_p^- . Le vecteur $\vec{n}_i = (2p_i, -1)$ est normal à f_i et dirigé vers l'extérieur de \mathcal{V}_p et le vecteur $\vec{n}_p = (-2p, +1)$ est normal à f_p et dirigé vers l'extérieur de \mathcal{V}_p . Appliquons le théorème de Minkowski à \mathcal{V}_p :

$$\sum_i \text{vol}(f_i) \frac{\vec{n}_i}{\|\vec{n}_i\|} + \text{vol}(f_p) \frac{\vec{n}_p}{\|\vec{n}_p\|} = 0. \tag{8.1}$$



La facette f_i se projette dans $x_{d+1} = 0$ selon $V^+(p, p_i)$ et la facette f_p se projette selon $V^+(p)$, d'où, en notant $\vec{i}_{d+1} = (0, \dots, 0, 1)$ de \mathbb{R}^{d+1} le vecteur unitaire porté par le dernier axe de coordonnées,

$$\begin{aligned} w_i(p) = \text{vol}(V^+(p, p_i)) &= -\text{vol}(f_i) \frac{\vec{n}_i}{\|\vec{n}_i\|} \cdot \vec{i}_{d+1} = \frac{\text{vol}(f_i)}{\|\vec{n}_i\|} \\ w(p) = \text{vol}(V^+(p)) &= \frac{\text{vol}(f_p)}{\|\vec{n}_p\|}, \end{aligned}$$

et, avec (8.1),

$$\sum_i w_i(p) \vec{n}_i + w(p) \vec{n}_p = 0.$$

En utilisant les expressions des normales et en projetant dans $x_{d+1} = 0$, on obtient finalement

$$\sum_i w_i(p) p_i - w(p) p = 0.$$

□

On rassemble les résultats précédents dans le

Theorem 8.1 *Etant donné un ensemble fini de points E , les λ_i constituent un système de coordonnées associé à E et il en va de même pour les ς_i .*

- Exercices 8.1**
1. Montrer que si $d = 1$, les coordonnées $\varsigma_i(x)$ s'identifient aux coordonnées barycentriques de x dans la subdivision de la droite réelle induite par E .
 2. Définir des coordonnées de Laplace et de Sibson dans un diagramme de Laguerre. Vérifier que les lemmes 8.1, 8.3 et 8.4 restent valides.
 3. S'inspirer de l'insertion d'un point dans une triangulation de Delaunay pour calculer efficacement $\lambda_i(x)$ et $\varsigma_i(x)$ pour $i = 1, \dots, n$ et x un point de $\text{conv}(E)$.
 4. Montrer que

$$w_i(x) = w_i^{-p_j}(x) - \left(w_i^{-p_j}(p_j) - w_i^{-p_j+x}(p_j) \right)$$

où w_i^{-p+q} est la i -ième coordonnée non normalisée de Sibson calculée dans le diagramme de Voronoï de l'ensemble $E \setminus \{p\} \cup \{q\}$. En déduire que ς_i est, sauf aux points de E^+ , une fonction continuellement différentiable de p_j (x étant fixé).

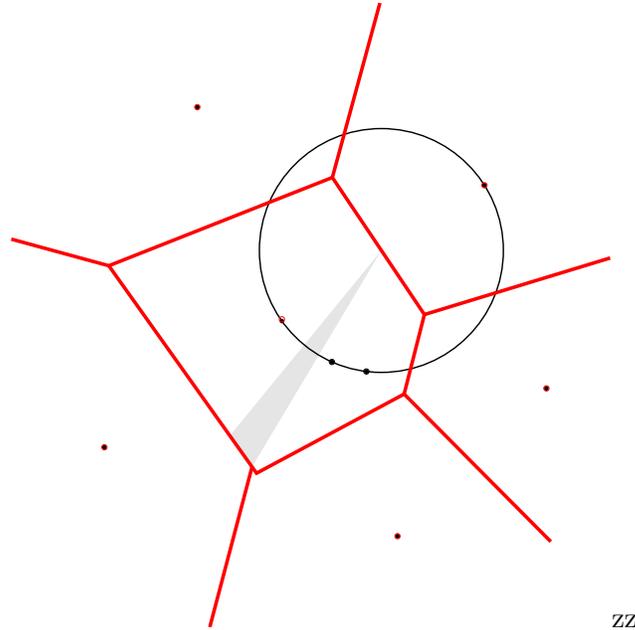
8.1.1 Gradient des coordonnées naturelles

On peut donner une formule explicite pour le gradient de ς_i due à Piper. En fait le lemme suivant donne le gradient de la coordonnée non normalisée w_i . Le gradient de ς_i s'en déduit par la relation

$$\nabla \varsigma_i(x) = \frac{1}{\sum_j w_j(x)} \left(\nabla w_i(x) - \varsigma_i(x) \sum_j \nabla w_j(x) \right)$$

Lemma 8.4 *En tout point $x \in \text{conv}(E) \setminus E$, $\nabla w_i(x) = \bar{v}_i(x)(c_i - x)$, où c_i est le centre de gravité de la facette $F^+(x, p_i)$.*

Proof. Nous donnons la preuve dans le cas plan. Soit c et c' les extrémités d'une arête de $V^+(x)$. Le point c est équidistant de x , p_i et p_j , et c' est équidistant de x , p_i et p_k . On note σ et σ' les cercles centrés respectivement en c et c' qui passent par x .



On calcule la projection du gradient selon deux directions. La première est la tangente au cercle σ en x . Soit $x(t)$ le point de σ image de $x = x(0)$ par une rotation de centre c et d'angle t . L'arête $F^+(x, p_i) = [cc']$ devient $F^+(x(t), p_i) = [cc'(t)]$. Si $x \notin E$ et pour t suffisamment petit, la différence d'aires $a(t) = w_i(x(t)) - w_i(x)$ est égale à l'aire du triangle $cc'c'(t)$, soit $\frac{1}{2} \|c - c'\| \times \|c - c'(t)\| \sin \theta$, en notant θ l'angle entre cc' et $cc'(t)$. Comme $x p_i$ et $x(t) p_i$ sont respectivement perpendiculaires à cc' et $cc'(t)$, on a $\theta = \angle c'cc'(t) = \angle x p_i x(t) = \frac{t}{2}$. On en déduit que

$$a'(0) = \lim_{t \rightarrow 0} \frac{a(t)}{t} = \frac{\|c - c'\|^2}{4}.$$

Comme $a(t) = w_i(x(t)) - w_i(x)$, on a

$$a'(0) = \nabla w_i(x(0)) x'(0) = \nabla w_i(x) \cdot (c - x)^\perp,$$

où v^\perp représente l'image du vecteur v par une rotation de $\frac{\pi}{2}$. On en déduit

$$\nabla w_i(x) \cdot (c - x)^\perp = \frac{\|c - c'\|^2}{4}. \tag{8.2}$$

Un argument analogue pour le cercle σ' conduit à

$$\nabla w_i(x) \cdot (c' - x)^\perp = \frac{-\|c - c'\|^2}{4}. \tag{8.3}$$

On vérifie maintenant que la formule du lemme conduit aux mêmes projections sur $(c - x)^\perp$ et $(c' - x)^\perp$. En effet, on déduit de la formule du lemme et de $c_i = \frac{1}{2}(c + c')$

$$\begin{aligned} \nabla w_i(x) \cdot (c - x)^\perp &= \frac{1}{2} \bar{v}_i(x) ((c - x) + (c' - x)) \cdot (c - x)^\perp \\ &= \bar{v}_i(x) \text{aire}(xcc') \\ &= \frac{1}{4} \bar{v}_i(x) \|x - p_i\| \times \|c - c'\| \quad (cc' \perp xp_i) \\ &= \frac{1}{4} \|c - c'\|^2. \end{aligned}$$

On retrouve (8.2). On montre de même que (8.3) est vérifiée. La formule du lemme étant vérifiée pour deux directions, elle est donc correcte. \square

Exercices 8.2

1. Etendre le lemme 8.4 au cas de coordonnées de Sibson définies sur un diagramme de puissance.
2. Montrer qu'il existe des constantes c et ε telles que si $\|x - p_j\| \leq \varepsilon$, on a pour tout i

$$|\zeta_i(x) - \zeta_i(p_j)| \leq c \|x - p_j\|.$$

3. Démontrer la formule du gradient en dimensions supérieures à 2.

8.1.2 Support des coordonnées naturelles

Le *support* d'une fonction définie sur \mathbb{R}^d est l'ensemble des $x \in \mathbb{R}^d$ pour lesquels cette fonction ne s'annule pas.

On déduit facilement de leur définition que le support de la coordonnée de Laplace λ_i et le support de la coordonnée de Sibson ζ_i est l'union des boules de Delaunay incidentes à p_i . On peut préciser quel est le diamètre du support de ces coordonnées si on suppose que E est un échantillon suffisamment dense.

Soit R un domaine borné de \mathbb{R}^d et $\varepsilon > 0$. On dit que E est un ε -échantillon de R si tout point de R est à distance $< \varepsilon$ d'un point de E . On note R^ε l'ensemble des points de R à distance au plus ε du bord de R qu'on appellera aussi *érodé* de R .

Lemma 8.5 *Si E est un ε -échantillon uniforme de R , les voisins naturels de tout point de R^ε sont contenus dans une boule centrée en x de rayon 2ε . Les supports des coordonnées de Laplace et Sibson λ_i et ζ_i associées à un point $p_i \in E \cap R^\varepsilon$ sont contenus dans une boule centrée en p_i de rayon 2ε .*

Proof. On montre que toute sphère passant par $x \in R^\varepsilon$ et n'englobant aucun point de E a un rayon au plus ε . Supposons le contraire. Soit σ une telle sphère, c son centre, et y le point du segment cx à distance ε de x . x étant dans l'érode de R , y est dans R . La boule B_c limitée par σ contient la boule B_y centrée en y et de rayon ε . L'intérieur de B_c ne contenant pas de point de E , il en va de même pour B_y , ce qui contredit l'hypothèse sur E . \square

8.2 Interpolation de données non structurées

8.2.1 Reconstruction exacte des fonctions affines

Soit un ensemble E fini de points p_i de \mathbb{R}^d . A chaque p_i est associée un réel f_i . On cherche une fonction $F : \text{conv}(E) \rightarrow \mathbb{R}$ qui interpole les données, c'est-à-dire telle que $F(p_i) = f_i$, $i = 1, \dots, n$.

En utilisant les coordonnées définies précédemment, on obtient deux telles fonctions

$$\begin{aligned} F_0(x) &= \sum_i \lambda_i(x) f_i \\ F_1(x) &= \sum_i \varsigma_i(x) f_i. \end{aligned}$$

Les propriétés suivantes découlent immédiatement des propriétés des coordonnées de Laplace et Sibson :

Interpolation : $F_0(p_i) = F_1(p_i) = f_i$, $i = 1, \dots, n$.

Continuité : F_0 et F_1 sont continues en tout $x \in \text{conv}(E)$. F_1 est continuellement différentiable en tout point $x \in \text{conv}(E) \setminus E$.

Précision : Si $F(x)$ est une fonction affine, i.e. $F(x) = a \cdot x + b$ où a et b sont deux vecteurs de \mathbb{R}^d , les lemmes 8.2 et 8.3 montrent que, en tout point de $\text{conv}(E)$, $F_0(x) = F_1(x) = a \cdot x + b$. Les interpolants F_0 et F_1 reconstruisent donc exactement les fonctions affines.

8.2.2 Reconstruction exacte des fonctions quadratiques

On suppose ici qu'à chaque point $p_i \in E$ est associé un réel $f_i = F(p_i)$ et un vecteur $n_i = \nabla F(p_i)$.

On note comme précédemment

$$F_1(x) = \sum_i \varsigma_i(x) f_i$$

et on définit

$$H_i(x) = f_i + n_i^t(x - p_i).$$

Observer que $H_i^{-1}(0)$ est le plan de $\mathbb{R}^d \times \mathbb{R}$ passant par le point (p_i, f_i) et perpendiculaire à n_i .

On utilise le lemme 8.3 pour calculer l'erreur d'approximation de F_1 et H_i dans le cas où les données proviennent d'une fonction quadratique $Q(x) = a + b^t x + x^t Q x$, où a, b sont des réels et Q une matrice symétrique.

$$\begin{aligned}
E_1(x) &= F_1(x) - Q(x) = \sum_i \varsigma_i(x) f_i - Q(x) \\
&= \sum_i \varsigma_i(x) (a + b^t p_i + p_i^t Q p_i) - (a + b^t x + x^t Q x) \\
&= \sum_i \varsigma_i(x) p_i^t Q p_i - x^t Q x \\
&= \sum_i \varsigma_i(x) (p_i^t Q p_i + x^t Q x - 2x^t Q p_i) \\
&= \sum_i \varsigma_i(x) (x - p_i)^t Q (x - p_i)
\end{aligned}$$

On considère une combinaison des H_i donnée par $H(x) = \sum_i \tau_i(x) H_i(x)$ avec $\sum_i \tau_i(x) = 1$. On obtient

$$E_2(x) = Q(x) - H(x) = \sum_i \tau_i(x) (Q(x) - H_i(x)) = \sum_i \tau_i(x) (x - p_i)^t Q (x - p_i)$$

On définit alors F_2 par $F_2(x) = \frac{E_1(x)H(x) + E_2(x)F_1(x)}{E_1(x) + E_2(x)}$. On vérifie que $F_2(x) = \frac{(F_1(x) - Q(x))H(x) + (Q(x) - H(x))F_1(x)}{(F_1(x) - Q(x)) + (Q(x) - H(x))} = Q(x)$.

Si on prend $\tau_i = \sigma_i$, $F_2(x)$ est calculable à partir de E , des f_i et des n_i : F_2 est une fonction interpolante qui reconstruit exactement les fonctions quadratiques.

Exercices 8.3

1. Montrer que F_2 reconstruit exactement les fonctions sphériques, c'est-à-dire les fonctions quadratiques pour lesquelles $Q = c\mathbb{I}$.
2. L'interpolant de Sibson F_s est obtenu en prenant pour $\tau_i(x) = \frac{\tau_i^*(x)}{\sum_i \tau_i^*(x)}$ où $\tau_i^*(x) = \frac{\varsigma_i(x)}{\|x - p_i\|}$. Par la question précédente, cet interpolant reconstruit les fonctions sphériques. Montrer de plus qu'il interpole les gradients n_i , c'est-à-dire que $\nabla F_s(p_i) = n_i$. En déduire que F_s est continuellement différentiable en tout point de $\text{conv}(E)$ (on pourra se servir de la question 8.2.2).

8.3 Systèmes de coordonnées sur une surface

Soit S une variété compacte de codimension 1 (on dira surface par commodité), orientable, de classe C^2 , plongée dans \mathbb{R}^d , et $E = \{p_1, \dots, p_n\}$ un ensemble de points de S . On aimerait définir un système de coordonnées associé à E et pouvoir interpoler des données sur S .

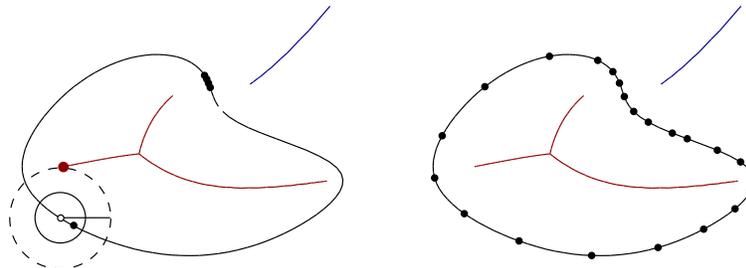
Les systèmes de coordonnées introduits au paragraphe 8.2.2 et définis sur \mathbb{R}^d présentent au moins deux difficultés. La première est que S n'est pas contenue dans $\text{conv}(E)$: on ne pourra donc pas interpoler sur tout S . La deuxième est que les voisins naturels d'un point de S ne sont pas nécessairement proches du point.

On peut pallier à ces difficultés en rajoutant des points sur une boîte de grande taille contenant S . On peut alors montrer que, si la boîte et S sont bien échantillonnés, les voisins naturels d'un point $x \in S$ qui sont loin de x ont des coordonnées petites. Cette solution est néanmoins coûteuse : le nombre de points de E a grossi et il faut gérer un grand nombre de voisins (alors que leur contribution est faible).

On présente au paragraphe 8.3.1 une autre façon de procéder qui consiste à définir un système de coordonnées sur S (et pas dans tout \mathbb{R}^d). Mais nous faisons auparavant un petit rappel pour définir un bon échantillon d'une surface.

On appelle $\text{lfs}(x)$ la distance de x à l'axe médian de S . On montre facilement que lfs est 1-Lipschitz : Pour tous $x, y \in S$, $\text{lfs}(x) \leq \text{lfs}(y) + \|x - y\|$.

Definition 8.3 On dit qu'un ensemble fini de points E de S est un ε -échantillon de S (pour lfs) si tout point x de S est à distance au plus $\varepsilon \text{lfs}(x)$, $\varepsilon < 1$, d'un point de E .



8.3.1 Voisins naturels sur S

Soit x un point de S . Le plan tangent T_x à S en x coupe $\text{Vor}(E)$ selon un diagramme de Laguerre $\text{Lag}(E')$ (lemme 4.3). Plus précisément, si on note p'_i le projeté de $p_i \in E$ sur T_x et $h_i = \|p_i - p'_i\|$, $\text{Lag}(E')$ est le diagramme de Laguerre des sphères (imaginaires) centrées aux p'_i et dont les carrés des rayons sont $-h_i^2$. On peut alors associer à x des voisins naturels dans $\text{Lag}(E')$ (comme au paragraphe 8.1 (question 8.1)) et

des voisins naturels sur S : un point $p_i \in E$ sera un *voisin naturel de x sur S* si son projeté x'_i est un voisin naturel de x dans $\text{Lag}(E')$.

Le lemme suivant affirme que les voisins naturels de x sur S sont proches de S .

Lemma 8.6 *Si E est un ε -échantillon avec $\varepsilon < \frac{1}{2}$, les voisins naturels d'un point $x \in S$ sont à distance au plus $\frac{2\varepsilon}{\sqrt{1-2\varepsilon}} \text{lfs}(x)$ de x .*

Proof. On note $E^+ = E \cup \{x\}$ et $E'^+ = E' \cup \{x\}$. Montrons tout d'abord que la plus grande distance de x à un de ses voisins naturels sur S n'est pas plus de deux fois la plus grande distance de x aux sommets de la cellule $L(x)$ de x dans $\text{Lag}(E'^+)$. En effet, si p_i est voisin naturel sur S de x , p_i est sur le bord d'une boule vide centrée en un sommet v de $L(x)$ qui passe également par x . On a donc $\|x - p_i\| \leq 2\|x - v\|$.

Majorons la distance de x aux sommets de $L(x)$. Soit B_x la boule de rayon $\text{lfs}(x)$ tangente à S en x telle que le segment joignant son centre c à v coupe S en un point q . On note $\alpha = \angle vcx = \angle vxq'$ où q' est le projeté de x sur le segment vc . La boule (ouverte) B_v de centre v et de rayon $\|x - v\|$ ne contient pas de point de E puisque $v \in V^+(x)$. Comme B_x ne contient pas non plus de point de E , x est le point de E^+ le plus proche de q . En effet la boule ouverte centrée en q et de rayon $\|q - x\|$ est contenue dans $B_x \cup B_v$ dont l'intersection avec E est vide. On a donc $\|x - q\| \leq \varepsilon \text{lfs}(q) \leq \frac{\varepsilon}{1-\varepsilon} \text{lfs}(x)$. Par ailleurs, $\|x - q\| \geq \|x - q'\| = \sin \alpha \text{lfs}(x)$. On en déduit que $\alpha \leq \arcsin \frac{\varepsilon}{1-\varepsilon}$. Le rayon de B_x étant $\text{lfs}(x)$, on conclut

$$\|x - v\| = \tan \alpha \text{lfs}(x) \leq \tan(\arcsin \frac{\varepsilon}{1-\varepsilon}) \text{lfs}(x) = \frac{\varepsilon}{\sqrt{1-2\varepsilon}} \text{lfs}(x).$$

□

8.3.2 Coordonnées naturelles sur S

Pour pouvoir affecter des coordonnées aux voisins naturels de x dans T_x , il suffit que la cellule de x dans $\text{Lag}(E' \cup \{x\})$ soit bornée. Le lemme suivant montre que ceci est vrai si E est un 1-échantillon de S , c'est-à-dire que tout point de S est plus près d'un point de E que du squelette de S .

Lemma 8.7 *Si E est un 1-échantillon de S , la cellule de x dans $\text{Lag}(E' \cup \{x\})$ est bornée.*

Proof. Pour prouver le lemme, il suffit de montrer que l'intersection de la cellule $V'(x)$ de x dans $\text{Vor}_{|T_x}(E \cup \{x\})$ est bornée. Supposons le contraire. L'intérieur de $V'(x)$ contient alors un point à l'infini p_∞ . Soit H le plan passant par x et perpendiculaire à (xp_∞) et H^+ le demi-espace limité par H qui contient p_∞ . Puisque x est parmi les points de $E \cup \{x\}$ celui qui est le plus de p_∞ , H^+ ne contient aucun point de E .

Soit $c \in H$ le centre de la sphère médiane passant par x et contenue dans la région bornée limitée par S . c appartient au squelette de s . Soit y un point d'intersection de S avec la demi-droite issue de c , perpendiculaire

à H et contenue dans H^+ . Un tel point existe puisque S est compacte et sans bord. $\|y - c\| \geq \text{lfs}(y)$ et, pour tout $p \in E$, $\|y - p\| > \|y - c\|$, ce qui contredit l'hypothèse que E est un 1-échantillon de S . \square

Lemma 8.8 *Si E est un ε -échantillon avec $\varepsilon < \sqrt{2} - 1$, le support Δ_i de τ_i est contenu dans une boule centrée en p_i et de rayon $r_i = \frac{2\varepsilon}{\sqrt{1-2\varepsilon}-\varepsilon}$.*

Proof. Du lemme 8.6, on déduit que pour tout $x \in \Delta_i$, $\|x - p_i\| \leq \eta \text{lfs}(x) \leq \frac{\eta}{1-\eta} \text{lfs}(p_i)$, avec $\eta = \frac{2\varepsilon}{\sqrt{1-2\varepsilon}}$. \square

Exercise 8.1 On peut montrer des résultats de continuité analogues à ceux obtenus au paragraphe 8.1 et utiliser les fonctions d'interpolation définies au paragraphe 8.2 pour interpoler des fonctions sur des surfaces.

8.4 Notes bibliographiques

Les coordonnées de Sibson ont été introduites par Sibson [34] et les coordonnées de Laplace par Sugihara [26]. Les propriétés de continuité de ces coordonnées énoncées par Sibson ont été prouvées par Farin [22] puis par Piper qui a donné la formule du gradient [30]. Le lemme 8.3 a été prouvé par Sibson. La preuve géométrique présentée ici reprend une idée de Sugihara pour le cas des coordonnées de Laplace. Plusieurs autres preuves sont connues [14]. L'extension des coordonnées naturelles aux diagrammes de Laguerre est étudiée en détail dans la thèse de Julia Flöttoto [23].

La notion de voisins naturels sur une surface a été introduite par Boissonnat et Flöttoto. Les résultats de la section 8.3.2 sont repris de [9].

On trouvera une implantation des coordonnées naturelles dans \mathbb{R}^2 et sur une surface dans la bibliothèque CGAL [15].

Bibliographie

- [1] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4) :481–504, 1999.
- [2] 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.
- [3] Dominique Attali, Jean-Daniel Boissonnat, and André Lieutier. Complexity of the Delaunay triangulation of points on surfaces : The smooth case. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 237–246, 2003.
- [4] F. Aurenhammer. Power diagrams : properties, algorithms and applications. *SIAM J. Comput.*, 16 :78–96, 1987.
- [5] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
- [6] M. Berger. *Géométrie (vols. 1-5)*. Fernand Nathan, Paris, 1977.
- [7] J-D. Boissonnat and S. Oudot. An effective condition for sampling surfaces with guarantees. In *Proc. 10th ACM Symposium on Solid Modeling and Applications*, 2004.
- [8] Jean-Daniel Boissonnat, Olivier Devillers, Monique Teillaud, and Mariette Yvinec. Triangulations in CGAL. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 11–18, 2000.
- [9] Jean-Daniel Boissonnat and Julia Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36 :161–174, 2004.
- [10] Jean-Daniel Boissonnat and Steve Oudot. An effective condition for sampling surfaces with guarantees. Technical Report 5064, INRIA, 2003. ECG-TR-304100-02.
- [11] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of lipschitz surfaces. In *Proc. 22th Annu. ACM Sympos. Comput. Geom.*, 2006.
- [12] Jean-Daniel Boissonnat and Mariette Yvinec. *Géométrie Algorithmique*. Ediscience international, Paris, 1995.
- [13] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

- [14] J. L. Brown. Systems of coordinates associated with points scattered in the plane. *Computer Aided Design*, 14 :547–559, 1997.
- [15] *The CGAL Library*. <http://www.cgal.org/>.
- [16] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10 :377–409, 1993.
- [17] David Cohen-Steiner and Jean-Marie Morvan. Restricted Delaunay triangulations and normal cycle. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 237–246, 2003.
- [18] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.
- [19] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry : Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.
- [20] Herbert Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, 2001.
- [21] Jeff Erickson. Nice point sets can have nasty Delaunay triangulations. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 96–105, 2001.
- [22] G. Farin. Surfaces over Dirichlet tessellations. *Comput. Aided Geom. Design*, 7 :281–292, 1990.
- [23] Julia Flötotto. *A coordinate system associated to a point cloud issued from a manifold : definition, properties and applications*. Thèse de doctorat en sciences, Université de Nice-Sophia Antipolis, France, 2003.
- [24] J. E. Goodman and J. O’Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press LLC, Boca Raton, FL, 1997.
- [25] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, NY, 1976.
- [26] Hisamoto Hiyoshi and Kokichi Sugihara. Voronoi-based interpolation with higher continuity. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 242–250, 2000.
- [27] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.
- [28] Atsuyuki Okabe, Barry Boots, and Kokichi Sugihara. *Spatial Tessellations : Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Chichester, UK, 1992.
- [29] D. Pedoe. *Geometry, a comprehensive course*. Dover Publications, New York, 1970.
- [30] B. Piper. Properties of local coordinates based on dirichlet tessellations. *Computing*, 8(227–239), 1993.
- [31] J. J. Rotman. *An Introduction to Algebraic Topology*. Springer-Verlag, New York, NY, 1988.
- [32] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18 :548–585, 1995.
- [33] R. Seidel. The upper bound theorem for polytopes : an easy proof of its asymptotic version. *Comput. Geom. Theory Appl.*, 5 :115–116, 1995.

- [34] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnet, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.
- [35] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, 1994.