Triangulations and Meshes

Mariette Yvinec

24 octobre 2006

Table des matières

Première partie

Introduction

Chapitre 1

Triangulations

1.1   Definitions

Simplexes

Assume that we work in the space $\mathbb{R}^d$ of dimension d. Let $k \leq d$. A *k-simplex* is a polytope which is the convex hull of $k+1$ independent points. The k simplexes are the smallest (with respect to the number of vertices) polytopes with dimension k. The simplexes of dimension 0, 1, 2, 3 are called respectively points, segments, triangle and tetrahedra.
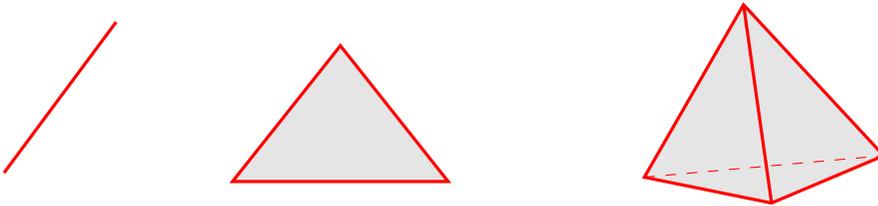
Let A be a subset of $k+1$ independent points, and let $\sigma(A)$ be the simplex that is the convex hull of A. Any subset $B \subset A$ with $l+1$ points in A is the subset of vertices of a l-simplex $\sigma(B)$ which is a face of $\sigma(A)$. Therefore the combinatorics of a simplex is completely fixed : A k-simplex $\sigma(A)$ has $\begin{pmatrix} k \\ l \end{pmatrix}$ faces of dimension l, and for any j, l such that $0 \leq j \leq l \leq k$, a j-face of $\sigma(A)$ is included in $\begin{pmatrix} k-j \\ k-l \end{pmatrix}$ l-faces of $\sigma(A)$.
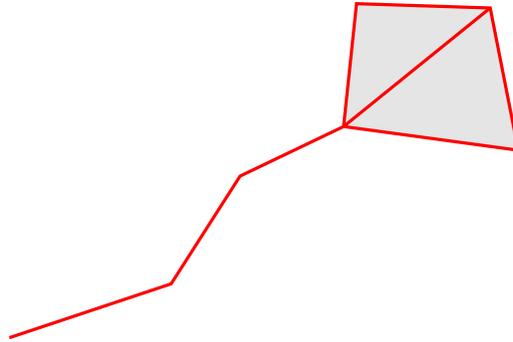
FIG. 1.1 – A complex.

Complexes

A *simplicial complex* $\mathcal{C}$ is a set of simplexes such that :
– any face of a simplex in $\mathcal{C}$ belongs to $\mathcal{C}$,
– two simplexes in $\mathcal{C}$ are either disjoint or share a lower dimensional face.
Any simplex in a complex simplicial $\mathcal{C}$ is called *a face* of $\mathcal{C}$. The *dimension* of a simplicial complex $\mathcal{C}$ is the maximal dimension of the simplexes in $\mathcal{C}$. The *domain* of a simplicial complex $\mathcal{C}$ is the union of the simplexes of $\mathcal{C}$.

Remark 1.1 *Similar definitions can be given to define* cellular complexes *whose elements are convex cells rather than simplexes.*

Remark 1.2 *The above definitions concern* embedded complexes *whose vertices are points embedded in a Cartesian space. On the opposite, an* abstract simplicial complexes *is defined from a set $\mathcal{S}$ as a set $\mathcal{C}$ of subsets of $\mathcal{S}$ such that :*
– any subset T' of a subset in T in $\mathcal{C}$ is in $\mathcal{C}$.

A complex $\mathcal{C}$ is said to be *pure* if any face of $\mathcal{C}$ is a face of a face of $\mathcal{C}$ with maximal dimension. A complex is *connected* if its domain is connected. For example, the complex in Figure 1.1 is connected but not pure because it has antenna edges which are not included in a triangle.

Let F be a k-face of a simplicial complex $\mathcal{C}$. The *star* of F in $\mathcal{C}$ is the simplicial complex formed by the subset of simplexes in $\mathcal{C}$ that contain F and their faces. The *link* of F in $\mathcal{C}$ is the simplicial complex formed by the subset of simplexes in the star of F that do not include F. A face F in a complex $\mathcal{C}$ is *singular* if the domain of its link is neither a topological ball nor a topological disk. See Figure 1.2.
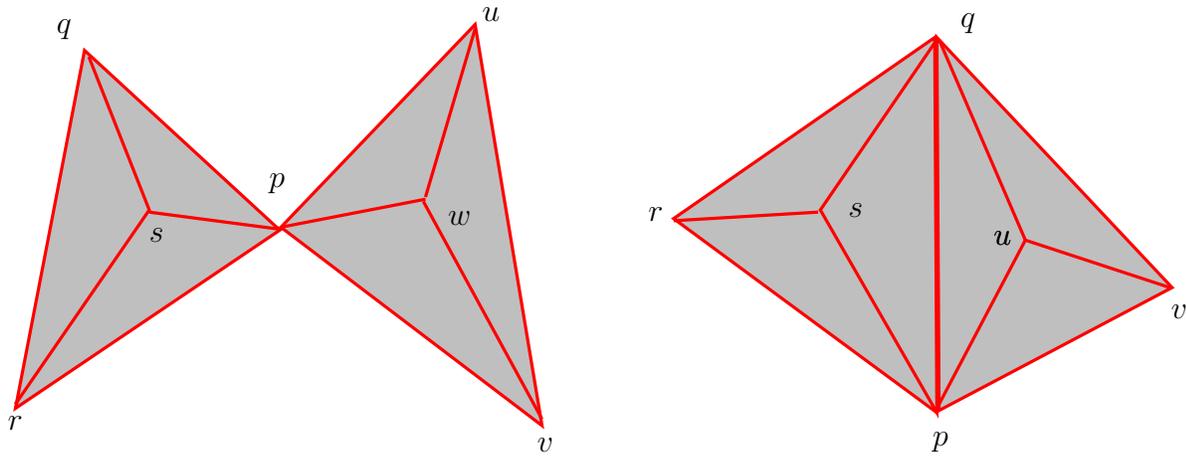
FIG. 1.2 – Singular faces. The complex on the left is made of two tetrahedra $pqrs$ and $puvw$, the vertex p is singular : its link is made of the two triangles qrs and uvw. The complex on the right side is made of the two tetrahedra $pqrs$ and $pquv$, the edge pq is singular : its link is made of the two triangles rs and uv.

Triangulations.

A triangulation is a simplicial complex that is pure, connected and without singular faces. The Figure 1.3 shows two examples of triangulations. The left one is a triangulation of a set of points. The triangulation of a set of points $\mathcal{P}$ is a triangulation having $\mathcal{P}$ as set of vertices and whose domain is the convex hull of $\mathcal{P}$. The domain of a triangulation needs to be connected. Note however that it does not need to be convex nor simply connected. The second example in Figure 1.3 is the triangulation of a polygonal region whose domain is not convex nor simply connected.

1.2 Existence and number of triangulations of point sets

The existence of triangulations of sets of points, for any set of points in any dimension, can be easily proved using for instance a constructive proof. We describe here a construction called the lazy incremental . triangulation of a set of points. For simplicity, we assume here that the set of points $\mathcal{P}$ is in general position, however the construction can be generalized to any set of points.

1. Compute the convex hull conv($\mathcal{P}$) of $\mathcal{P}$.

2. Initialize the triangulation $\mathcal{T}$ by a triangulation obtained by *staring* the polytope conv($\mathcal{P}$) from any of its vertices. This means the following : we choose a vertex p of conv($\mathcal{P}$) and consider the set of simplexes $\{conv(p, F)\}$ where F is any facet on the boundary of conv($\mathcal{P}$) that do not contain p.
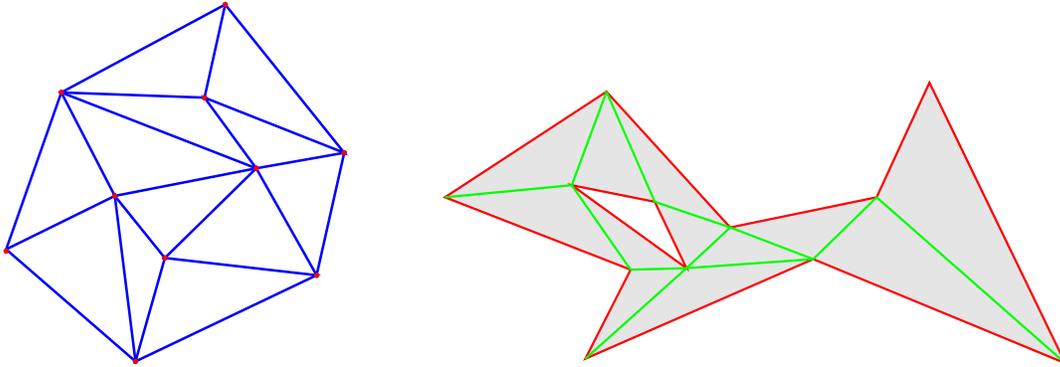
FIG. 1.3 – Examples of triangulations. Left : the triangulation of a set of points in 2D. Right : the triangulation of a 2D polygonal region.

3. Update incrementally the triangulation $\mathcal{T}$ as follows. For any point p of $\mathcal{P}$ that is not a vertex of $conv(\mathcal{P})$, we locate p in $\mathcal{T}$, The tetrahedron T of $\mathcal{T}$ containing p is replaced by the $(d+1)$ tetrahedra $conv(p, F_i)$ formed with the $(d+1)$ facets F_i of T.

Generally, point sets have several triangulations. Very few is known about the number of different triangulations of a point set. In the planar case, it is known that n points in convex positions have $C_n = \frac{1}{n+1} \begin{pmatrix} 2n \\ n \end{pmatrix} \approx \frac{1}{\sqrt{\pi}} \frac{4^n}{n^{\frac{3}{2}}}$ triangulations. Also the number of triangulations of a planar set of n point is known to be $\Omega\left((2+\varepsilon)^n\right)$, for some $\varepsilon > 0$,. This number is also $O(59^n/n^6)$.

1.3   Size of 2-dimensional triangulations.

1.3.1   Triangulations of a planar point sets

We consider here a triangulation of a set $\mathcal{P}$ of n points in dimension 2. We know that the number of vertices is n, and our goal is to determine or at least bound the number e of edges and the number f of facets (i. e. triangles) of such a triangulation.

Our starting point is the Euler characteristic of the triangulation. The homology theory says that any two complexes that have homeomorphic domains have the same Euler characteristic. For 2-dimensional triangulations of point sets the Euler characteristic is 1 (as can be checked on any example) and the Euler relation writes :

$$n - e + f = 1 \tag{1.1}$$

Then comes the edge/facet incidence relation. This relation is obtained from counting the number of inci-
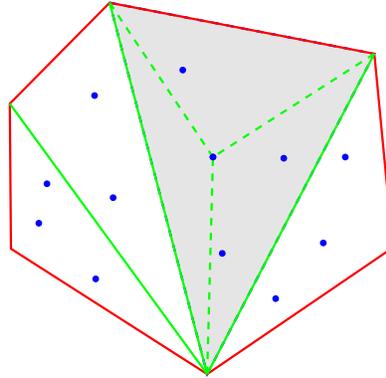
FIG. 1.4 – The lazy incremental construction of a triangulation of a set of points.

dence relations between edges and facets. From the point of view of facets, this number is $3f$, because each facet is a triangle and is therefore incident to three edges. From the point of view of edges this number is $2e - n_e$ where n_e is the number of edges (or vertices) on the convex hull $\mathrm{conv}(\mathcal{P})$ of the point set, because each edge is incident to two facets except the edges on the boundary of $\mathrm{conv}(\mathcal{P})$ which are incident to only one facet. Therefore,

$$3f = 2e - n_e \tag{1.2}$$

From equations 1.1 and 1.2 we get the following results.

Theorem 1.1 *Any triangulation of a set of $\mathcal{P}$ of n points in the plane with n_e vertices on the convex hull has the same number of edges (e) and facets (f) given by :*

$$
\begin{aligned}
f &= 2n - 2 - n_e \\
e &= 3n - 3 - n_e
\end{aligned}
$$

1.3.2   General planar triangulations

More generally, any planar triangulation with n vertices obeys the Euler relation

$$n - e + f = 1 - k \tag{1.3}$$

where k is the number of holes in the (connected) domain of the triangulation. The edge/facet incidence relation 1.2 is still valid if n_e is the number of edges (or vertices) on the boundary of the triangulation. This yields the following theorem.

Theorem 1.2 *For any planar triangulation with n vertices, whose domain include k holes and with n_e vertices on its boundaries, the number e of edges and the number f of facets are :*

$$\begin{aligned} f &= 2(n-1+k) - n_e \\ e &= 3(n-1+k) - n_e \end{aligned}$$

1.3.3   Size of triangulated surfaces.

Consider now triangulated surfaces, that is 2-dimensional triangulations embedded in the 3-dimensional space. For such a surface, the Euler relation writes :

$$n - e + f = 2 - 2h - k \tag{1.4}$$

where n, e, f, are respectively the number of vertices, edges and facets as before, h is the number of handles and k the number of boundaries. To justify this, let us consider first a surface with no handles, and no boundary. Such a surface is a topological sphere and its Euler characteristic is 2. To obtain a surface with handles and boundaries, we conceptually add the handles and boundaries, one at a time and then one boundary at a time. Adding a handle, amounts to glue the current surface with a topological sphere, putting two facets of one surface in correspondence with two facets of the other and removing the four glued facets. The resulting change in the Euler characteristic is -2 for each handle. Adding a boundary amounts to remove one facet which yields a change of -1 in the Euler characteristic.

The edge/facet incidence relation 1.2 is still valid if we note n_e the total number of vertices (or edges) on the surface boundaries, which together with equation 1.4, yields the following theorem.

Theorem 1.3 *Any triangulated surface with n vertices, h handles, k boundaries and n_e edges (or vertices) on the boundary has e edges and f facets such that :*

$$\begin{aligned} f &= 2(n-2+2h-k) - n_e \\ e &= 3(n-2+2h-k) - n_e. \end{aligned}$$

In particular, if a triangulated surface is a topological sphere with n vertices, it has $3n - 6$ edges and $2n - 4$ facets.

1.4   Size of 3 dimensional triangulations

Things get more complicated for 3-dimensional triangulations. First, all the triangulations of the same set of points do not have the same numbers of edges, facets and cells. Let us consider for example the two triangulations of a set of five points in Figure 1.5. The first one has two tetrahedra, 7 facets and 9 edges while the other has 3 tetrahedra, 9 facets and 10 edges.
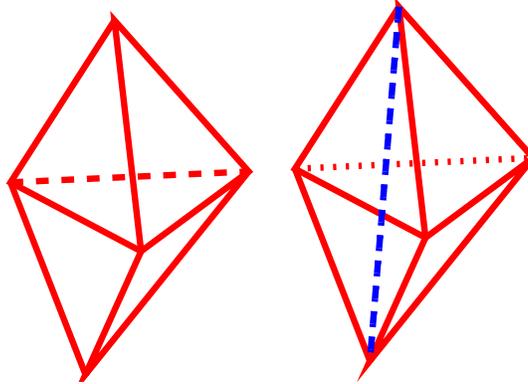
FIG. 1.5 – Two triangulations of a set of five points.

Even worse, a set of points may have triangulations with a linear number of faces and triangulation with a quadratic number of faces. Consider for example the set $\mathcal{P}$ of $2n+1$ points in Figure 1.6. There are n points on the edge pq which together with the edge rs form a triangulation with $n-1$ tetrahedra sharing the edge rs. We now add $n-2$ points on the edge rs. Each point added in rs splits in two the subedge containing that point and the $n-1$ tetrahedra sharing this subedge which yields $2(n-2)(n-1)$ tetrahedra. Last we add point t at the middle of rt splitting in two the single tetrahedra containing the edge rt, for a final number of $2(n-1)(n-2)+1$ tetrahedra. Now, re-starting from the $n-1$ tetrahedra sharing edge rs, we could have first inserted the point t, splitting in two the single tetrahedra containing edge rt and then the $n-2$ points on rs which results in $2(n-1)$ tetrahedra sharing the facet pst.

Some point sets admit only quadratic size triangulations. For instance, withdrawing the point t from the above example, yields a set of $\mathcal{P}'$ of $2n$ points with n points on edge pq and n points on edge rs. Such a set has a single triangulation which is of quadratic size. However, we know that sets in general position have linear size triangulations. Indeed, let us consider once more the lazy incremental construction presented in section 1.2. Assume that the set of n points has n_e points on the convex hull and therefore $n-n_e$ internal points. The convex hull has $2n-4$ facets and if we star it from a vertex with degree δ, the initial triangulation has $t = 2n_e-4-\delta$ tetrahedra. Each internal point splits a tetrahedra in four, and yield a gain of 3 tetrahedra. The total number of tetrahedra of the lazy triangulation is therefore $t = 2n_e-4-\delta+3(n-n_e) = 3n-n_e-\delta-4$ which is less than $3n-11$, because the number of vertices on the convex hull is at leat 4 and the degree of a vertex on the convex hull is at least 3.

Let us however try to bound the number of faces in the triangulation of a set of n point in dimension 3. We note n, e, f, t, the number of vertices, edges, facets and cells. The triangulation of a set of points is a topological ball, and the Euler relation for a topological ball in dimension 3 is

$$n - e + f - t = 1. \tag{1.5}$$
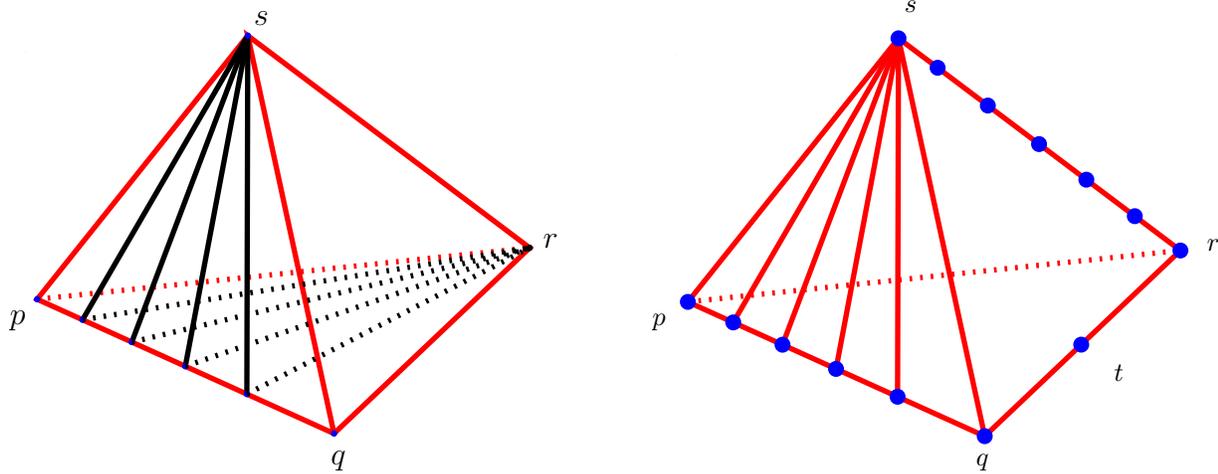
FIG. 1.6 – A point set with a linear size triangulation and a quadratic size one.

Let n_e and f_e be respectively the number of vertices and facets on the convex hull. Because, the boundary of the convex hull is a triangulated topological sphere, we get from theorem 1.3,

$$f_e = 2n_e - 4. \tag{1.6}$$

Then, counting the number of incidence relations between facets and tetrahedra, we get

$$4t = 2f - f_e \tag{1.7}$$

From equations 1.5, 1.6 and 1.7, we get that the number t of tetrahedra is :

$$t = e - n - n_e + 3 \tag{1.8}$$

The bound on the number of tetrahedra now comes from bound on the number of edges. On the one hand, the number of edges e is less than $\frac{n(n-1)}{2}$. On the other hand, the number of edges is $e = e_e + e_i$ where e_e is the number of edges on the convex hull and e_i is the number of internal edges. We know from theorem 1.3 that $e_e = 3n_e - 6$. The number e_i of internal edges is such that $2e_i \geq 4(n - n_e)$ because each internal vertex is incident to at least four internal edges. This altogether proves the following theorem :

Theorem 1.4 *In 3-dimensional space, the number t of tetrahedra in the triangulation of a point set with size n, is such that*

$$n - 3 \leq t \leq \frac{n^2}{2} - \frac{3n}{2} - n_e + 3 \tag{1.9}$$

where n_e is the number of vertices on the convex hull.

1.5 Bibliographical notes

The definition and combinatorics of triangulations can be found in any good textbook, e. g. [2]. The lower bound on the number of triangulations for planar point set is due to is due to O. Aichholzer, F. Hurtado, and M. Noy [1] and the upper one to F. Santos and R. Seidel [5].

Deuxième partie

Constrained triangulations

Chapitre 2

Constrained triangulations in dimension 2

Roughly speaking, in dimension 2 a constrained triangulation is a triangulation of a set of points which is required to include among its edges a given subset of non intersecting segments joining the points. Constrained triangulation are highly useful to discretize polygonal regions. Indeed those regions can be represented as a union of facets in a triangulation including as edges the segments forming the boundary of the regions.

2.1 Definition

Let us first introduce the notion of planar straight line graph which is useful to describe the input of a constrained triangulation problem. A *planar straight line graph*, later called *PSLG* for short, is a pair $(\mathcal{P}, \mathcal{S})$ where $\mathcal{P}$ is a set of points and $\mathcal{S}$ is a set of segments forming together a 1-dimensional simplicial complex. This means that :
– each endpoint of a segment in $\mathcal{S}$ is a point of $\mathcal{P}$,
– any two segments in $\mathcal{S}$ are either disjoint or share an endpoint.
A triangulation T is a *constrained triangulation* of the PSLG $(\mathcal{P}, \mathcal{S})$ if :
– the set of vertices of T is $\mathcal{P}$
– any segment of $\mathcal{S}$ is an edge of T
As an example, a constrained triangulation can be used to discretize a polygonal region. See Figure **??**. Assume we are given a polygonal region defined by a set of non intersecting polygons $\mathcal{Q}$. For instance, the first polygon describes the outer boundary of the region while the others polygons describe the boundaries of the holes piercing the region. We take as set $\mathcal{P}$ the set of vertices of the polygons in $\mathcal{Q}$ and as set $\mathcal{S}$ the set of edges of the polygons in $\mathcal{Q}$. Any constrained triangulation T of $(\mathcal{P}, \mathcal{S})$ allows to represent the polygonal region as the union of a subset of the facets of T.
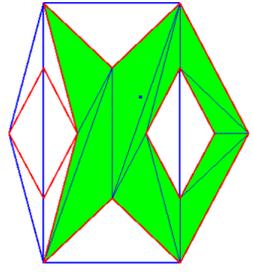
FIG. 2.1 – Triangulation of a polygonal region

2.2   Existence of constrained triangulations

The following theorem asserts the existence of constrained triangulations in dimension 2.

Theorem 2.1 *In dimension 2, any PSLG (P, S) admits a constrained triangulation.*

To prove this theorem, we first prove the following lemma which gives an alternative definition for a triangulation of a set of points in dimension 2.

Lemma 2.2 *Let $\mathcal{P}$ be a set of points in the 2-dimensional space. The set of edges of any triangulation of $\mathcal{P}$ is a maximal set of edges with endpoints in $\mathcal{P}$ and such that two edges do not intersect except may be at their endpoints. (The term maximal is refers to the set inclusion relation.)*

Proof.   For shortness, we say here that two edges do not intersect when then do not intersect except may be at a common endpoint. Let T be a triangulation of the set $\mathcal{P}$. We take as obvious that there is no edge with endpoints in $\mathcal{P}$ that is not in T or does not intersect some edge in T. Let us therefore prove the reverse proposal which says that any maximal set of non intersecting edges with endpoints in $\mathcal{P}$ is the set of edges of a triangulation. Let us consider a maximal set $\mathcal{E}$ of non intersecting edges with endpoints in $\mathcal{P}$. We first notice that set $\mathcal{E}$ includes the edges of the convex hull of $\mathcal{P}$ otherwise it would not be maximal. Next, we show that any bounded region in the planar map form by the edges in $\mathcal{E}$ is triangular. Assume for contradiction that there is a non triangular bounded region G in this map. See Figure 2.2. Let a be the vertex of G with minimum x-coordinate. Vertex a belongs to the outer boundary of G and is connected by edges of $\mathcal{E}$ to two other vertices b and c such that edges ba and ac are on the boundary of G. Edge bc does not belong to $\mathcal{E}$, otherwise G would be a triangle. Because $\mathcal{E}$ is maximal, there must be some edges of $\mathcal{E}$ intersecting bc and because those edges cannot intersect ab and ac, they have endpoints in triangle abc. Therefore there are some points of $\mathcal{P}$ in triangle abc. Let d be the point in triangle abc encountered first when sweeping a line parallel to bc from a toward bc. Then ad is a segment with endpoints in $\mathcal{P}$ and no intersection with segments in $\mathcal{E}$ which contradicts the fact that $\mathcal{E}$ is maximal. $\square$
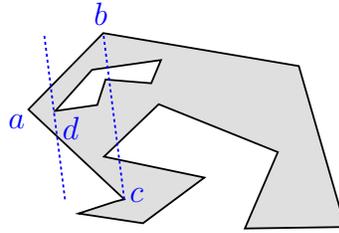
FIG. 2.2 – For the proof of lemma 2.1.
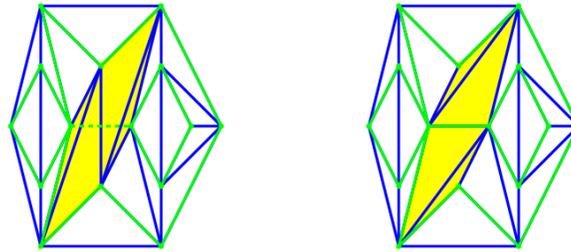
FIG. 2.3 – Incremental construction of constrained triangulations

Proof. [Proof of theorem 2.1] To achieve the proof theorem 2.1, it suffices to notice that the set of edges $\mathcal{S}$ of a PSLG $(\mathcal{P}, \mathcal{S})$ is a set of non intersecting edges with endpoints in $\mathcal{P}$. Therefore it can always be completed into a maximal set of non intersecting edges with endpoints in $\mathcal{P}$, which from lemma 2.2 is the set of edges of a triangulation of $\mathcal{P}$. □
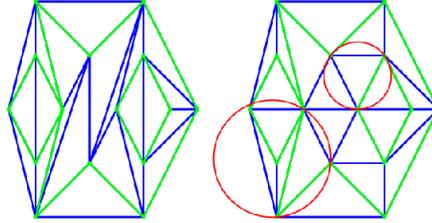
FIG. 2.4 – Constrained and constrained Delaunay triangulations

2.3   Algorithms for constrained triangulation

2.3.1   Incremental algorithm

2.3.2   A sweeping algorithm

2.3.3   Triangulation of a polygon

2.4   Constrained Delaunay triangulations

There are generally several possible constrained triangulations of a given PSLG $(\mathcal{P}, \mathcal{S})$. Most of the time, none of them is the Delaunay triangulation $\mathrm{Del}(\mathcal{P})$ of $\mathcal{P}$, because the set $\mathcal{S}$ include some edges which are not Delaunay edges. The goal of this section is to define the constrained Delaunay triangulations which are particular constrained triangulations.

Let $(\mathcal{P}, \mathcal{S})$ be a PSLG. The definition of constrained Delaunay triangulation is based on a notion of visibility which considers the segment of $\mathcal{S}$ as obstacles to the visibility. More precisely, we say that point p is *visible* from point q, or that p and q are mutually visible if the interior of the segment pq intersect no segment in $\mathcal{S}$.

The definition of constrained Delaunay triangulation is easier if we define first constrained Delaunay simplexes.

Definition 2.3 (Constrained Delaunay simplex) *Given a PSLG $(\mathcal{P}, \mathcal{S})$, a simplex e (segment or triangle) is* a constrained Delaunay simplex *(for $(\mathcal{P}, \mathcal{S})$) if and only if :*
– the vertices of e are points of $\mathcal{P}$,
– the (relative) interior of e intersects no segment in $\mathcal{S}$,
– there is a circumcircle of e that encloses no point of $\mathcal{P}$ visible from the interior of e (that is from any point in the relative interior of e).
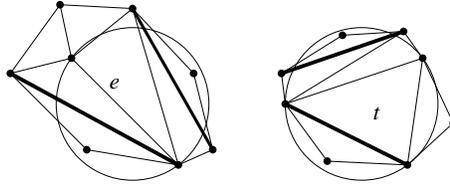
FIG. 2.5 – Constrained Delaunay simplexes. Edge e and triangle t are both constrained Delaunay simplexes for the PSLG whose edges are marked as bold edges.

Definition 2.4 (Constrained Delaunay triangulation.) *The constrained triangulation T of the PSLG $(\mathcal{P}, \mathcal{S})$ is a constrained Delaunay triangulation of $(\mathcal{P}, \mathcal{S})$ if and only if any triangle t in T is a constrained Delaunay triangle.*

Theorem 2.5 *The constrained triangulation T of the PSLG $(\mathcal{P}, \mathcal{S})$ is a constrained Delaunay triangulation of the PSLG $(\mathcal{P}, \mathcal{S})$ if and only if any edge e in T is either a segment of $\mathcal{S}$ or a constrained Delaunay segment.*

Proof. Let us first show that any edge of a constrained Delaunay triangulation T of a PSLG $(\mathcal{P}, \mathcal{S})$ is a constrained Delaunay segment for $(\mathcal{P}, \mathcal{S})$. First notice that any edge e of T is an edge of a triangle t of T and therefore the endpoints of e belong to $\mathcal{P}$ and the interior of e intersects no segment of $\mathcal{S}$. If e is a segment of $\mathcal{S}$ nothing else has to be proved. Otherwise, notice that any point enclosed by the circumcircle of t and visible from the interior of e is also visible from the interior of t. By hypothesis the circumcircle of t, encloses no point of $\mathcal{P}$ visible from the interior of t. Therefore it is a circumcircle of e that encloses no point of $\mathcal{P}$ visible from the interior of e. This proves that e is a constrained Delaunay segment.

Assume now that T is a constrained triangulation of the PSLG $(\mathcal{P}, \mathcal{S})$ whose edges are either segments in $\mathcal{S}$ or a constrained Delaunay segments and let us prove that T is a constrained Delaunay triangulation of $(\mathcal{P}, \mathcal{S})$. Let t be a triangle of T. Because T is a constrained triangulation of $(\mathcal{P}, \mathcal{S})$, t has points of $\mathcal{P}$ as vertices and its interior intersects no segment in $\mathcal{S}$. It remains to show that the circumcirle of t encloses no point of $\mathcal{P}$ visible from t. Let $b(t)$ be the disk bounded by the circumcircle of t. The subset $b(t)/t$ is formed of three regions each of which is separated from t by an edge of t. Consider the part of b_e of $b(t)/t$ separated from t by the edge e of t. Either e is a segment of $\mathcal{S}$ and no point in b_e is visible from the interior of t. Or, segment e is constrained Delaunay and therefore there is a circumcircle of e enclosing a disc d_e that contains no point of $\mathcal{P}$ visible from the interior of e. Observe that disk d_e contains no point of $\mathcal{P}$ visible from the interior of t (otherwise such a point would be visible from a point in the interior of e). Futhermore, disk d_e does not include the vertex of t opposite to e which is visible from e and therefore it covers b_e. Thus, no region of $b(t)/t$ contains a point in $\mathcal{P}$ visible from the interior of t and neither does $b(t)$, which proves that triangle t is constrained Delaunay. $\square$
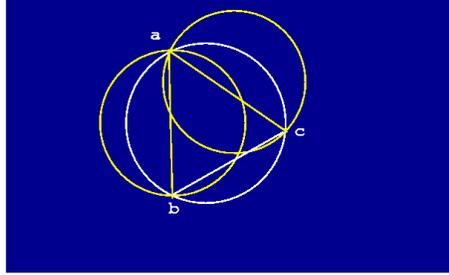
FIG. 2.6 – For the proof of theorem 2.5.

2.5   Existence and unicity of constrained Delaunay triangulations

Theorem 2.6 *Any PSLG $(\mathcal{P}, \mathcal{S})$ has a constrained Delaunay triangulation.*
If $\mathcal{P}$ includes no subset of four co-circular points, this triangulation is unique.

Proof.   assume first that $\mathcal{P}$ has no subset of four co-circular points. We consider the set $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ of constrained Delaunay segments and claim that $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$ is a maximal set of edges with endpoints in $\mathcal{P}$ and no intersection, except at a common endpoints.

Let us first prove that two edges in $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$ do not intersect except at maybe at endpoints. By definition two edges in $\mathcal{S}$, or one segment in $\mathcal{S}$ and one in $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ do not intersect. Let us consider two edges e and e' in $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ and assume for contradiction that they intersect in a common point x in the interior of e. See figure 2.7. Let C_e be the circumcircle of e that encloses no point of $\mathcal{P}$ visible the interior of e. The circle C_e encloses point x but do not enclose the endpoints or e' that are visible from a point on e close enough to x. We note a and b the two points where e' intersect C_e. Any circle circumscribing e' encloses a circle circumscribing segment ab and any circle circumscribing ab (except C_e) encloses one of the endpoint of e which are both visible from any point in e' close enough to x. Thus e' cannot be an edge of $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ which is a contradiction.

Let us show that the set $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S}))$ is maximal. $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ includes the edges of the convex hull of $\mathcal{P}$. If $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$ is not maximal, there is a bounded non triangular facet in the planar map formed by $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$. Let G be such a facet and let a be the vertex of G with minimum x-coordinate. Vertex a is incident to two edges ab and ac on the external boundary of G and edge bc is not in $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$. Therefore bc does not belong to $\mathcal{S}$, nor to $\mathcal{CD}(\mathcal{P}, \mathcal{S})$. This implies that either bc intersects segments in $\mathcal{S}$ which have endpoints in triangle $t = abc$ or that any circumcircle of bc, and in particular the circumcircle C_t of triangle t, enclose some points of $\mathcal{P}$ visible from bc. If the second case occurs, observe that the visible point are either in t or in the portion D_{bc} of the ball B_t bounded by C_t which is separated from t by bc. Indeed the edge ab is either belongs to $\mathcal{S}$ in which case D_{ab} is not visible from bc, or ab is in $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ in which case D_{ab} is enclosed by the circle of ab that encloses no point of $\mathcal{P}$ visible from ab. Therefore region D_{ab} contains
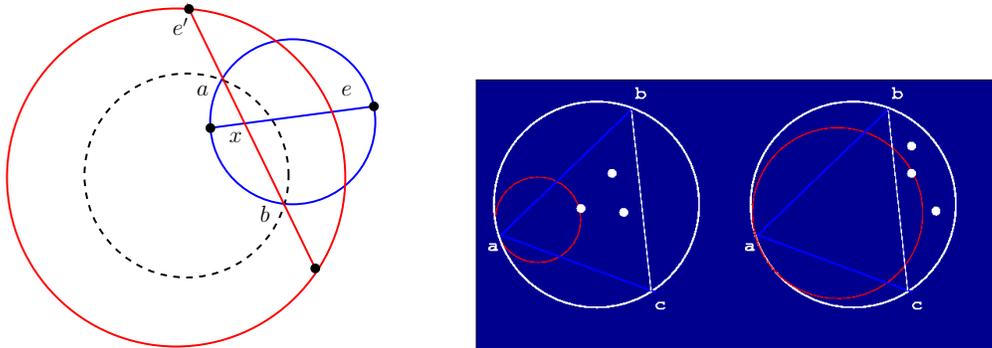
FIG. 2.7 – For the proof of theorem 2.6.

no point of $\mathcal{P}$ visible from t and the same holds for D_{ac}. Let's grow a circle C from point a keeping the tangency in a to C_t until C hits a point d of $\mathcal{P}$ in t or D_{bc}. Then, the segment ad should be in $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ which is contradiction.

From theorem 2.2, we know that $\mathcal{S} \cup \mathcal{CD}(\mathcal{P}, \mathcal{S})$ is the set of edges of a constrained triangulation of $(\mathcal{P}, \mathcal{S})$ and from theorem 2.5, we know that this triangulation is a constrained Delaunay triangulation of $(\mathcal{P}, \mathcal{S})$. This proves the existence and at the same time the unicity of such the constrained Delaunay triangulation when $\mathcal{P}$ has no subset of four cocircular points. When case of cocircularities occur among points of $\mathcal{P}$ we may have intersecting segments in $\mathcal{CD}(\mathcal{P}, \mathcal{S})$. However we can still extract from $\mathcal{CD}(\mathcal{P}, \mathcal{S})$ a maximal set of non intersecting segments with togeteher with $\mathcal{S}$ forms a constrained Delaunay triangulation. □

2.6 Constrained Delaunay triangulations and conforming Delaunay triangulations

Constrained Delaunay triangulation should not be confused with conforming Delaunay triangulation. The constrained Delaunay triangulation of a PSLG $(\mathcal{P}, \mathcal{S})$ includes segment in $\mathcal{S}$ as edges and has $\mathcal{P}$ as vertex set. It is not a Delaunay triangulation, because its simplexes are constrained Delaunay and not Delaunay : they satisfy only a weak version of the empty sphere property based on an the notion of visibility.

A conforming Delaunay triangulation of a PSLG $(\mathcal{P}, \mathcal{S})$ is a Delaunay triangulation in which the segments of $\mathcal{S}$ have been refined in subsegments by the introduction of additional points (called Steiner vertices) until all subsegments are Delaunay segments. Therefore, a conforming Delaunay triangulation is a Delaunay triangulation. The vertex set of a conforming triangulation includes $\mathcal{P}$ and additionnal vertices. The conforming triangulation do not necessarily include the segment of $\mathcal{S}$ as edges, but each segment in $\mathcal{S}$ is represented as a union of edges of the triangulation. Figure **??** show the differences between constrained and conforming
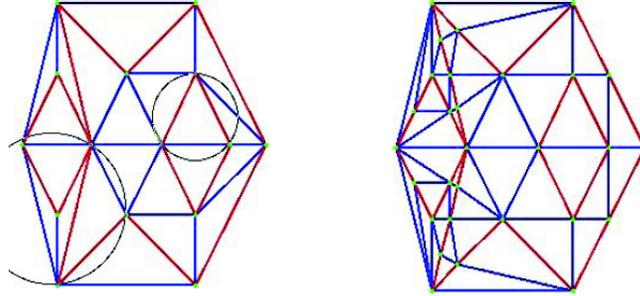
FIG. 2.8 – Constrained and conforming Delaunay triangulations
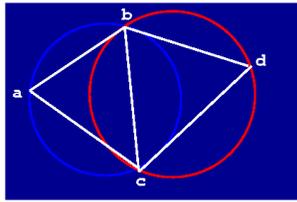
FIG. 2.9 – A regular or local Delaunay edge.

Delaunay triangulations for the same PSLG $(\mathcal{P}, \mathcal{S})$.

2.7   Local characterization of constrained Delaunay triangulations

Definition 2.7 (Regular edges) *An edge bc incident to two triangles abc and bcd is said to be* regular *or* locally Delaunay *iff it is an edge of the Delaunay triangulation of the four point set $\{a, b, c, d\}$. This condition is equivalent to say either that point a is not enclosed by in the circumcircle C_{bcd} of triangle bcd or that point d is not enclosed by the circumcircle C_{abc} of triangle abc.*

Theorem 2.8 *Any constrained triangulation of the PSLG $(\mathcal{P}, \mathcal{S})$ whose edges are either constrained edges (i. e. edges of $\mathcal{S}$) or regular edges, is the constrained Delaunay triangulation of $(\mathcal{P}, \mathcal{S})$.*

Proof.   Let T be a triangulation of the PSLG $(\mathcal{P}, \mathcal{S})$ whose edges are either constrained edges and regular edges. For any triangle t in T, we prove that the circumball C_t of t encloses no point of $\mathcal{P}$ visible from t.
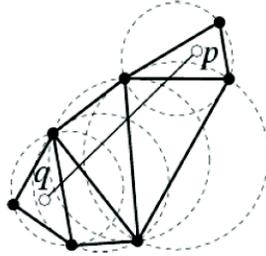
FIG. 2.10 – For the proof of theorem 2.8

Let t be a triangle of T and p a point of $\mathcal{P}$ (that is a vertex of T) visible from a point q inside t. Let $t_0, t_1, \ldots t_n = t$ be the sequence of triangles traversed by segment pq ordered from p to q. Let $\Pi(p, t_i)$ be the power of p with respect to the circumcircle $C(t_i)$ of t_i. Because p and q are mutually visible, each crossed edge is a regular edge and this implies that the power sequence $\{\Pi(p, t_i), i = 0, \ldots, n\}$ is a growing sequence. Indeed, triangle t_i is crossed before t_{i+1} because it is, with respect to the line l_i supporting line of the edge $e_i = t_i \cap t_{i+1}$ on the same side as p is. Line l_i is the radical axis of the circumcircles $C(t_i)$ and $C(t_{i+1})$ and the regularity of edge e_i imply that triangle t_i and therefore p is on the side of l_i where points have a smaller power to Ct_i than the power to $C(t_{i+1})$. Thus,

$$\Pi(p, t_0) \leq \Pi(p, t_1) \ldots \leq \Pi(p, t_n) = \Pi(p, t).$$

Furthermore, $\Pi(p, t_0) = 0$ because p is a vertex of t_0 and therefore $\Pi(p, t) \geq 0$ which implies that p is not enclosed in the circumcircle of t. $\square$

2.8   A flip algorithm

Let abc and bcd be two adjacent triangles in a triangulation T. If the quad Q which is the union of theses two triangles is convex, it can be re-triangulated with triangles adb and adc. The local change in the triangulation which consists in replacing edge bc by the edge ad and the two triangles abc and bcd by the two triangles adb and adc is called a flip.

Assume that the quad Q is convex. Because the sum of internal angles of Q is 2π, exactly one of the two edges bc and ad is regular. A flip is said to be a Delaunay flip when it replaces a non regular edge by a regular edge.

We know from the proof of Lemma 3.2 in chapter **??** that a Delaunay flip increases the angular sequence of the triangulation.
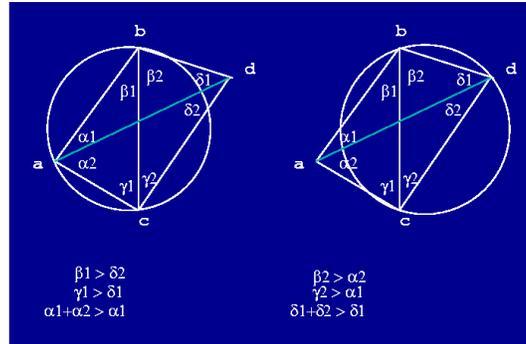
FIG. 2.11 – A Delaunay flip improves the angular sequence of the triangulation.

The flipping algorithm

The following algorithm transforms any constrained triangulation of the PSLG $(\mathcal{P}, \mathcal{S})$ into the constrained Delaunay triangulation.

1. Start from any constrained triangulation of $(\mathcal{P}, \mathcal{S})$

2. Initialize a stack with edges that are neither constrained edges nor regular edges

3. While the stack is not empty do
   – pop an edge ad from the stack
   – if ad is still not regular
        flip ad,
        re-test the regularity of the four wing edges ab, ac, db, dc
        update the stack.

Theorem 2.9 *The flip algorithm ends and yields the constrained Delaunay triangulation of a PSLG with n vertices in time $O(n^2)$.*

Proof. When the algorithm stops, every edge is either a constrained edge or a regular edge and the triangulation is constrained Delaunay by theorem 2.8. It remains to show that the algorithm cannot loop for ever. This is guaranteed by the fact that each performed flip is a Delaunay flip and therefore increases the angular sequence.

In the framework of the lifted map introduced in section **??** and **??** of chapter **??**, the current triangulation is lifted in a bumped surface in $\mathbb{R}^3$. Each Delaunay flip locally lowers the surface, and the updated surface hide the lifted version of the removed edge. Therefore, a flipped edge will never appear again in the triangulation. and the total number of flips performed is at most $\frac{n(n-1)}{2}$. $\square$

Chapitre 3

Optimal triangulations

3.1 Introduction

The quality of a given triangulation can be estimated for different criteria. The most common criteria are based on the shape of triangles and the quality of a triangulation is usually defined as the measure of its worst triangle or as the ordered sequence of measures of its element. Common measures on triangle shapes are for instance the minimum or maximum angle, the length of the minimum elevation, the radius of the circumcircle or the radius-edge ratio which is the ratio between the circumradius and the shortest edge length.

We have seen in chapter **??** that a Delaunay flip improves the minimum angle, which simply means that that the minimum angle of the two triangles produced by the flip is greater than the minimum angle of the two triangles removed. This fact implies that among all the triangulation of a set of points $\mathcal{P}$, the Delaunay triangulation maximizes the angular sequence or the minimum angle of the triangulation. For the same reason, the constrained Delaunay triangulation of a PSLG $(\mathcal{P}, \mathcal{S})$ is, among all the constrained triangulations of $(\mathcal{P}, \mathcal{S})$, the optimal triangulation for the angular sequence or the minimum angle. More generally, Delaunay and constrained Delaunay triangulations are optimal for all the quality measures which are improved by a Delaunay flip.

The main goal of this chapter is to prove the following theorem, summarizing the optimality properties of Delaunay and constrained Delaunay triangulations.

Theorem 3.1 (Optimality of Delaunay and constrained Delaunay triangulation.) *Among all the triangulation of a planar set of points, the Delaunay triangulation achieves the following optima :*

MaxMin angle : *the Delaunay triangulation maximizes the minimum angle of any triangle.*

MinMax circumradius : *the Delaunay triangulation minimizes the length of the largest cicumradius of any triangle*
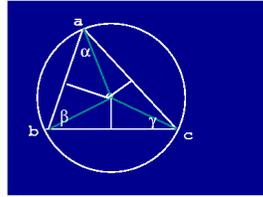
FIG. 3.1 – Circumradius, angles and edge lengths

MinMax smallest enclosing sphere : *we consider for each triangle its smallest enclosing sphere. The Delaunay triangulation minimizes the radius of the largest of all the smallest enclosing spheres.*

The constrained Delaunay triangulation of a PSLG $(\mathcal{P}, \mathcal{S})$ achieves the same optimality properties among all the constrained triangulations of $(\mathcal{P}, \mathcal{S})$.

It is to be noted that the third optimality properties (MinMax smallest enclosing sphere) is the only one that is still valid for triangulations of dimension 3 or higher.

The end of the chapter provides a way to optimize other measures which are not improved by Delaunay flips.

3.2   MinMax circumradius

Circumradius, angles and edge lengths

Let us first notice how the circumradius of a triangle is closely related to the lengths of its edges and to the values of its internal angles. Let $t = abc$ be a triangle. We note r the circumradius, l_a, l_b, l_c the length of the edges opposite to the vertices a, b, c respectively and α, β, γ the internal angles at vertices a, b, c respectively. Elementary triangle geometry, see figure 3.1 shows that :

$$r = \frac{l_a}{\sin \alpha} = \frac{l_b}{\sin \beta} = \frac{l_a}{\sin \gamma}$$

Optimality of Delaunay triangulations

To prove the optimality of Delaunay and constrained Delaunay triangulations with respect to the circumradius, we just have to prove the following theorem.

Theorem 3.2 *A Delaunay flip decreases the maximum circumradius.*
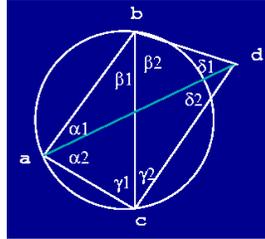
FIG. 3.2 – Delaunay flips decreases the maximum circumradius.

Proof. Lets assume that the Delaunay flip of edge ad, replaces the triangles abd and adc by abc and bcd. See figure 3.2. Because the flip is Delaunay, d is outside the circumcircle of abc and a is outside the circumcircle of bcd. Therefore we have the following angular comparisons :

$$\gamma_1 > \delta_1 \quad \text{and} \quad \beta_2 > \alpha_2,$$

which implies :

$$\text{circumradius}(abc) = \frac{ab}{2sin\gamma_1} \quad < \quad \text{circumradius}(abd) = \frac{ab}{2sin\delta_1}$$

$$\text{circumradius}(bcd) = \frac{cd}{2sin\beta_2} \quad < \quad \text{circumradius}(adc) = \frac{cd}{2sin\alpha_2}$$

□

3.3   Smallest enclosing sphere

The smallest enclosing sphere of a simplex is the enclosing sphere with smallest radius. It should not be confused with the circumsphere of the simplex which may or may not be the smallest enclosing sphere. (See Figure 3.3) The following theorem makes precise the relationship between the circumsphere and the smallest enclosing sphere of any simplex t. Let us stress that it is valid in a space of any dimension for a simplex of any dimension.

Theorem 3.3 *Let t be a k-simplex in a space of dimension d. The center x_{min} of the smallest enclosing sphere of t is :*
– the circumcenter x_c of t if x_c is included in t,
– the point of t closest to x_c otherwise.
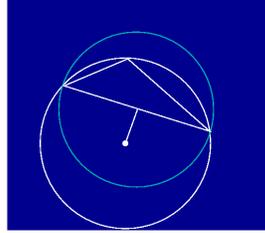
FIG. 3.3 – Smallest enclosing sphere.

Proof. The circumcenter x_c of t is the single point of the affine hull of t equidistant to the $k+1$ vertices of t. The center of the smallest enclosing sphere x_{min} is the point of the affine hull of t that minimizes its distance to the vertex of t that is the furthest. If x_c is included in t, any small displacement increases its distance to the furthest vertex, hence x_c realizes the minimum distance to the furthest vertex and coincides with x_{min}.

Assume now that x_c is not included in t and let $\varpi(x_c)$ be the point of t closest to x_c and let f be the single face of t such that $\varpi(x_c)$ belongs to the relative interior of f. The segment $x_c\varpi(x_c)$ is orthogonal to f, therefore we have :
– for any vertex v_i of f, $\|x_c - v_i\|^2 = \|x_c - \varpi(x_c)\|^2 + \|\varpi(x_c) - v_i\|^2$
– for any vertex v_i of $t \setminus f$ $\|x_c - v_i\|^2 \geq \|x_c - \varpi(x_c)\|^2 + \|\varpi(x_c) - v_i\|^2$.
Because x_c is equidistant to all vertices v_i, this shows that $\varpi(x_c)$ is equidistant to the vertices of f and closer to any vertex in $t \setminus f$ than to the vertices in f. This ensures first that $\varpi(x_c)$ is the center of the smallest sphere enclosing f. and second that the smallest sphere enclosing f encloses t as well. Therefore $\varpi(x_c)$ is the center x_{min} of the smallest sphere enclosing t. $\square$

To prove the optimality of Delaunay and constrained Delaunay triangulation with respect to the MinMax smallest enclosing sphere we need to prove that a Delaunay flip decreases the maximum of the radii of smallest enclosing spheres. However, we shall prove a more powerful result, stated in the theorem below, saying that in any dimension the Delaunay triangulation of a set of points minimizes the maximum radius of smallest enclosing spheres. Applied to the set of four points concerned by a flip, this provides the wanted result concerning the Delaunay flip in dimension 2 and thus also yields the optimality of 2 dimensional constrained Delaunay triangulations with respect to the MinMax smallest enclosing sphere.

Theorem 3.4 *In any dimension, the Delaunay triangulation minimizes the maximum radius of smallest enclosing spheres among all the triangulations of a set of points.*

Proof. The proof includes four main steps.

Step 1 : defines a function F.   Let $t = (p_0, p_1, \ldots p_d)$ be a d simplex. For any point $x \in R^d$, we define the barycentric coordinates of x with respect to t, $\lambda_i(x)$, $i = 0 \ldots d$, as

$$x = \sum_i \lambda_i(x) p_i, \quad \sum_i \lambda_i(x) = 1$$

Then we define the function $F(t, x)$ as :

$$F(t, x) = \sum_i \lambda_i(x)(p_i - x)^2 = \sum_i \lambda_i(x) p_i^2 - x^2$$

Step 2 : gives a second interpretation of F.   Let (x_c, r_c) be the center and radius of the circumsphere of t and (x_{min}, r_{min}) be the center and radius of the smallest enclosing sphere of t. We have :

$$
\begin{aligned}
F(t, x) &= \sum_i \lambda_i(x)(p_i - x)^2 \\
F(t, x) &= \sum_i \lambda_i(x) \left((p_i - x_c)^2 + 2(p_i - x_c)(x_c - x) + (x_c - x)^2 \right) \\
F(t, x) &= r_c^2 - (x - x_c)^2
\end{aligned}
$$

The last equation shows that $F(x, t)$ is just the power of x with respect to the circumsphere of t. Further more the maximum of $F(t, x)$ is r_c^2 and is obtained for $x = x_c$. But if we add the condition that x must stay within t, than the conditional maximum of $F(t, x)$ is reached for $x = x_{min}$ and is just r_{min}^2.

Step 3 : gives a third interpretation of F.   We can give another interpretation of F in the space of spheres. Recall that the lift map send points of $\mathbb{R}^d$ onto the unit paraboloid of $\mathbb{R}^{d+1}$.

$$
\begin{aligned}
p_i &\longrightarrow \phi(p_i) = (p_i, p_i^2) \\
x &\longrightarrow \phi(x) = (x, x^2)
\end{aligned}
$$

Let $\phi(t)$ be the d-simplex in $\mathbb{R}^{d+1}$ whose vertices are the lifted points $\phi(p_i)$. The expression $F(t, x) = \sum_i \lambda_i(x) p_i^2 - x^2$ shows that $F(t, x)$ is just the vertical distance $d(\phi(t), \phi(x))$ that is the distance between $\phi(x)$ and its vertical projection on the affine hull of $\phi(t)$.

If x is included in t (meaning that $\forall i, 0 \le \lambda_i(x) \le 1$) $\phi(t)$ is above x and he vertical distance $d(\phi(t), \phi(x)) = F(t, x)$ is always positive.

Let $\mathcal{P}$ be a set of points. Let us now consider the set $\mathcal{S}_P$ of all simplexes in $\mathbb{R}^d$ with vertices in $\mathcal{P}$ and the subset $\mathcal{S}_P(x)$ of those including x. Then, because the lifted transform of the Delaunay triangulation of $\mathcal{P}$: $\{\phi(t), \quad t \in \mathrm{Del}(\mathcal{P})\}$ is just the lower hull of its set of vertices, we know that the minimum of $F(t, x)$ for x fixed and t in $\mathcal{S}_P(x)$ is reached for when t is the simplex of the Delaunay triangulation $\mathrm{Del}(\mathcal{P})$ that contains x.
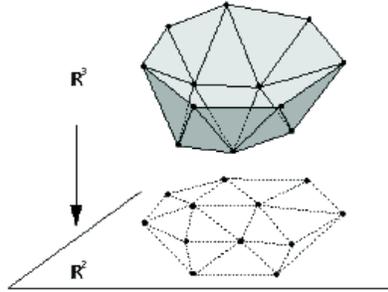
FIG. 3.4 – The lift map of a triangulation

Step 4 : concludes. Let $\mathcal{P}$ be a set of points. We consider a triangulation T of $\mathcal{P}$ and compare it with the Delaunay triangulation $\text{Del}(\mathcal{P})$ noted here DT for short. To each of the trianglations T and DT, w e associate a functions defined on the convex hull $\text{conv}(\mathcal{P})$ of $\mathcal{P}$ as follows :
– $F_T(x)$ is the value of $F(t,x)$ for the simplex t of T that contains x.
– $F_{DT}(x)$ is the value of $F(t,x)$ for the simplex t of DT that contains x.
$F_T(x)$ and $F_{DT}(x)$ are continuous, piecewise linear functions on $\text{conv}(\mathcal{P})$. Let x_T be the point of $\text{conv}(\mathcal{P})$ where $F_T(x)$ is maximum and x_{DT} the point $\text{conv}(\mathcal{P})$ where $F_{DT}(x)$ reaches its maximum. We have :

$$\max_{t \in T} r_{min}(t)^2 = F_T(x_T) \geq F_T(x_{DT}) \geq F_{DT}(x_{DT}) = \max_{t \in DT} r_{min}(t)^2,$$

This shows that the maximum radius of all the smallest enclosing spheres of the simplexes in the Delaunay triangulation DT is less than the maximum radius of all the smallest enclosing spheres of the simplexes in any other triangulation. $\square$

3.4   When Delaunay flips does not work

Some measures are not always improved by a Delaunay flip and therefore not optimized by the Delaunay triangulation. For instance Delaunay triangulation do not achieve :

the MinMax angle that is the minimum of the maximum of all the internal angles of the triangles. of the maximum angle of a triangle.

the MaxMin elevation that is the maximum of the minimum length of all the triangles elevations.

the minimum total edge length that is minimum of the total edge length.

One could have the idea to adapt the flip algorithms to those measures, performing only flips that locally improve the measure. However the example below shows that such a strategy may lead to a dead lock where
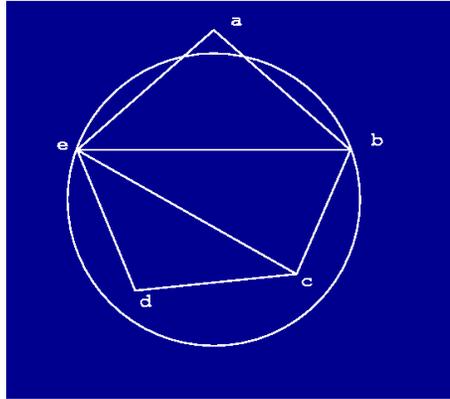
FIG. 3.5 – When Delaunay flip does not work

no improving flip is possible while the current triangulation is still not optimal. In the following example, see figure 3.5, we try to achieve the MinMax angle. The five point configuration $\{a, b, c, d, e\}$ is obtained as follows. We start from a regular pentagon, then slightly pull vertex a out of the enclosing circle, then pull d slightly closer to a along ad and then c a little bit closer to b along bd so that we get

$$\hat{a} < \hat{e} = \hat{b} < \hat{d} < \hat{c}.$$

The optimal triangulation would obviously use the internal edges ab and ac. Nevertheless, if we get the triangulation with internal edges eb and ec, no improving flip is possible.

Two approaches might be possible to avoid getting stuck in such a local minimum. The first one, called *simulated annealing*, consists in allowing flips which do not improve locally the measure. The second one consist in using elementary transformations of the triangulation which are more powerful than flips. The *edge insertion* which we present below is one of them. It allows to get triangulations that achieve the MinMax angle and the MaxMin elevation. More generally *edge insertion* allows to optimize a class of measure called anchored measures.

3.5 Optimal triangulations through edge insertion

Assume that the measure to be optimized is defined for a triangulation T as being the maximum (or the minimum) over all the triangles t in T of a measure $f(t)$ defined for each triangle. Wlog, we assume hereafter that $f(T) = \max_{t \in T} f(t)$ and that we wish to minimize $f(T)$.

Definition 3.5 (Anchored measures) *A triangle abc is said* to have an anchor *at vertex a iff any triangulation T whose measure is better than $f(abc)$ ($F(T) < f(abc)$) has an edge ad incident to vertex a and*
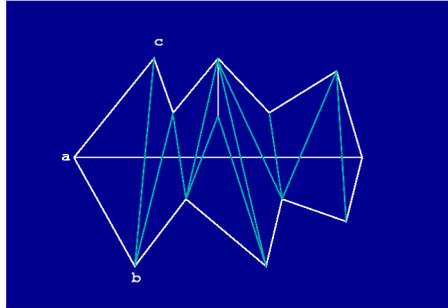
FIG. 3.6 – Edge insertion

intersecting bc.

A measure is anchored *if any triangle has an anchor*

The triangulation measure is defined as the maximum of all internal angles of the triangles in the triangulation is an example of anchored measure.

Definition 3.6 (Edge insertion) *Let T be a triangulation of a set of points $\mathcal{P}$. Given two points a and d of $\mathcal{P}$, the insertion of edge ad in T consist in (see Figure 3.6 :*
– remove from T all edges and triangles intersected by ad
– re-triangulate the two regions R_1 and R_2 formed when adding edge ad in the hole left over by removed triangles and edges.

Theorem 3.7 *Any anchored measure can be optimized through edge insertion*

Proof.   The proof consists in showing that while the current triangulation T is not optimal, there is always an edge insertion that improves the measure or at least does not make it worse. Let $t = abc$ be the worst triangle in T (meaning that $f(T) = f(t)$) and assume that t has an anchor at vertex a. Let ad be one of the edge incident to a and intersecting bc in the optimal triangulation T^*. We claim that the insertion of edge ad improves the measure or at least does not make it worse. To prove the claim, we show that the two regions R_1 and R_2 that have to be re-triangulate upon the insertion of ad may be re-triangulated with two sets of triangles $T(R_1)$ and $T(R_2)$ such that : $f(T(R_1) = \max_{t \in T(R_1)} f(t) \leq f(abc)$ and $f(T(R_2) = \max_{t \in T(R_2)} f(t) \leq f(abc)$.

Let us consider region R_1. Three successive vertices p, q, r on the boundary of R_1 form an *ear* if the angle $\widehat{pqr}$ of R_1 is less than π. An ear of R_1 is said to be chopped by a segment s if s intersects both edges pq and qr (intersections in p and r are allowed but not in q).

If p, q, r form an ear chopped by an edge T^* the triangle $t' = pqr$ can be used in the triangulation $T(R_1)$. Indeed, because edges pq and qr were incident in T to triangles intersecting edge ad, the segment qr does not intersect the boundary of R_1. Furthermore, we have $f(t') \le f(t)$. Indeed, if t' has an anchor at p, $f(t') \le f(T^*)$ because there is no edge in T^* breaking the anchor at p. Otherwise t has an anchor at p or r and $f(t') \le f(T)$ because there is no edge in $f(T)$ breaking the anchor at p or r.

Now the main point is that there is always an ear of region R_1 chopped by an edge of the triangulation T^*. To see this, think of the edge e^* of T^* that intersects R_1 and is "maximally far" to ad. Therefore region R_1 can be triangulated by repeatedly removing from R_1 an ear chopped by an edge of T^* and adding the corresponding triangle to $T(R_1)$.

The following algorithm ensures that the process of edge insertions will not loop. Therefore it yields the optimal triangulation. The algorithm repeatedly try to insert an edge breaking the anchor of the worst triangle in the current triangulation. On each insertion trial, the holes are tentatively re-triangulated by successive removal of ears whose triangles have a measure not worse than the current measure of the triangulation. If the retriangulation process succeeds, the insertion is successfull and the edge opposite to the anchor in the worst triangle is definitevily eliminated. Otherwise the tentativily inserted edge is eliminated.

Edge insertion algorithm

1. Initialize T with any constrained triangulation of $(\mathcal{P}, \mathcal{S})$
2. While there is a triangle t of T such that $f(t) = F(T)$ and a non eliminated edge ad breaking the anchor of t, do
   – insert edge ad if this improves the measure,
   – eliminate edge ad otherwise.

The algorithm performs at most $O(n^2)$ edge insertion. Each tentive edge insertion can be performed in $O(n)$ time (using a stack similar to the one used in a Graham scan or in the incremental construction of constrained triangulation) which gives a running time of $O(n^3)$. $\square$

3.6 MaxMin elevation

Let us first notice that the minimum elevation of a triangle $t = abc$ is always the elevation incident to the vertex with maximum angle. Indeed, let $h(a), h(b), h(c)$ be the lentgh of elevations incident to a, b, c repectively, and let α, β, γ be the angles in a, b, c. We have :

$$
\begin{aligned}
h(a) &= ab \sin \beta = ac \sin \gamma \\
h(b) &= bc \sin \gamma = ba \sin \alpha \\
h(c) &= ca \sin \alpha = cb \sin \beta,
\end{aligned}
$$

which shows that $h(a) \ge h(b) \ge h(c)$ if $\sin \alpha \le \sin \beta \le \sin \gamma$.
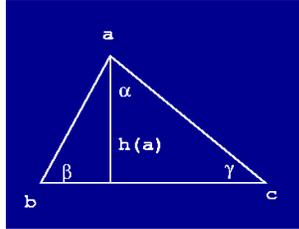
FIG. 3.7 – Elevation in a triangle
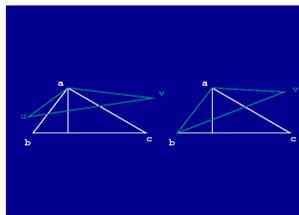
FIG. 3.8 – MaxMin elevation

The following theorem shows that a triangulation of a set of point $\mathcal{P}$ maximizing the minimum elevation length can be constructed by the edge insertion algorithm.

Theorem 3.8 *The minimum elevation length is an anchored measure.*

Proof. Let $t = abc$ be a triangle and assume that the smallest elevation of t is issued from vertex a and has length $h(a)$ We show that any triangulation T which has no edge breaking anchor a, i.e.incident to a and intersecting bc, is such that its smallest elevation length $h_{min}(t)$ is such that $h_{min}(t) \leq h(a)$. Indeed, either t is included in T or T has a triangle incident to a whose edge e opposite to a chops abc and intersects the elevation incident to a, see Figure 3.8. Therefore in both case $h_{min}(t) \leq h(a)$. $\square$

3.7   Dynamic programming

Optimizing the total edge length of the triangulation of a polygon.

3.8   Bibliographic notes

Most of this chapter is taken from the survey citebe-mgot-95 by M. Bern and D. Eppstein.

Chapitre 4

Constrained triangulations in dimension 3

4.1   Definitions

In the 3-dimensional space, the constraints of a triangulation problem are described has a piecewise linear complex.

Definition 4.1 (Piecewise linear complex (PLC)) *A piecewise linear complex $\mathcal{C}$, also called a PLC for short, is a set of faces of of dimension 0,1,2 such that :*
– the boundary of any face of C is the union of faces of C,
– the intersection of two faces of C is either empty
  or the union of faces of C.

The faces of a PLC are called vertices, edges and facets when their dimension is respectively 0,1 and 2.

Definition 4.2 (Constrained triangulations) *A constrained triangulation of the PLC $\mathcal{C}$ is a triangulation $T(\mathcal{C})$ such that*
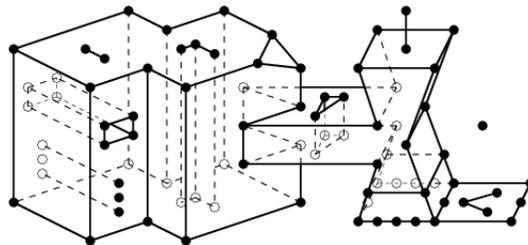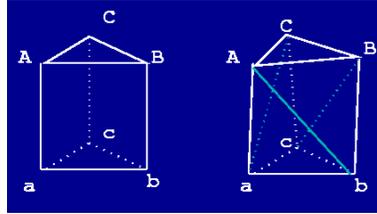
– *the vertex set of $T(\mathcal{C})$ is the vertex set of $\mathcal{C}$,*
– *any edge of $\mathcal{C}$ is an edge of $\mathcal{C}$,*
– *any facet of $\mathcal{C}$ is the union of faces of $T(\mathcal{C})$.*

4.2   The Schönhardt polyhedron

The main problem arising in dimension 3 is that the existence of a constrained triangulation is not guaranteed for any PLC. The most simple and famous example of PLC without any constrained triangulation is formed by the faces of a Schönhardt polyhedron.

A Schönhardt polyhedron can be described as follows : Let abc be three vertices forming an equilateral triangle in the plane $z = 0$. Let $a'b'c'$ be another equilateral triangle obtained from abc by a translation parallel to the vertical axis. We consider first the convex polyhedron obtained as the convex hull of $\{a, b, c, a', b', c'\}$. The Schönhardt polyhedron is obtained from this convex hull by slightly rotating the triangle $a'b'c'$ around its vertical axis. Then the lateral quadrilateral facets $abb'a'$, $bcc'b'$ and $caa'c'$ are non longer planar facets, each of them is split in two triangles by the additional edges ab', bc' and ca'.

To see that Schönhardt polyhedron admits no constrained triangulation, we first notice that, using the edges ab', bc' and ca', we have fold the quadrilateral facets toward the inside. This means that the other diagonal segments ba', cb' and ac' are outside the polyhedron and therefore cannot be used to triangulate the polyhedron. Then it just remains to notice, that any tetrahedron formed with four out of the six vertices of the Schönhardt polyhedron includes one of the forbidden diagonals. By symmetry, there are only three types of such tetrahedra whose prototypes are for example $abca'$, $aba'b'$ and $aba'c'$.

4.3   Triangulation of a polyhedron

To triangulate a polyhedron with no constrained triangulation one has to build a triangulation with additional (Steiner vertices). Of course, convex polyhedra can be triangulated without Steiner vertices, for instance by staring them from one of their vertices. A non convex polyhedron has *reflex* edges where the internal dihedral angle is more than π. The main result of this section is the following theorem (whose proof is only

sketched) stating precisely the intuitive fact that polyhedron with only few reflex edges can be triangulated with only few Steiner vertices.

Theorem 4.3 *Any polyhedron of genus 0 with n edges out of each only $r < n$ are reflex edges can be triangulated using $O(n + r^2)$ tetrahedra.*

The proof of this theorem is based on the notion of vertical decomposition of a polyhedron which has its own interest.

4.3.1   The vertical decomposition of a polyhedron

Let Q be a polyhedron with size $O(n)$, that is Q has $O(n)$ vertices, edges and facets. We obtain a vertical decomposition of the polyhedron Q as follows, see figure 4.1) :

Step 1 From each point p on an edge e of Q we draw a vertical upward segment and a vertical downward segment which extend until they first intersect the boundary of Q in a point different from the starting point p. Let us consider the sets of segments stemming from the same edge of Q and butting on the same facet. The maximal connected components of those sets form vertical trapezoids called 2-*walls of type 1*. The number of 2-walls of type 1 is $O(n^2)$ and these 2-walls decompose the polyhedron Q in cylindrical cells, each cell having only two non vertical facets called the floor and the ceiling of the cell. The floor and the ceiling of each cell may be not convex and even not simply connected.

Step 2 The cylindrical cells are then further subdivided as follows. Each floor or ceil facet is decomposed along its trapezoidal map obtained by drawing from each vertex two segments parallel to the y-axis and extended until they first meet the boundary of the cell in a point different from the starting point. These segments are 1-walls. From each point of a 1-wall we draw a vertical upward segment and a vertical downward segment, extending until they hit the ceiling or floor of the ceil. The vertical segments stemming from the same 1-wall form vertical trapezoids called 2-walls of type 2. There are $O(n^2)$ 1-wall and $O(n^2)$ 2-walls of type 2. The 2-walls of type 2 subdivide further the 2-walls of type 1 into $O(n^2)$ 2-walls of type 1'. The resulting set of 2-walls with type 2 or 1' subdivide the polyhedron into convex cells. Each convex cell has a cylindric shape with a trapezoid floor and ceiling. Each cell may have a large number of 2-walls of type 1' as vertical facets but the total complexity of the decomposition is still $O(n^2)$.

Triangulation of a polyhedron

A triangulation of the polyhedron into $O(n^2)$ tetrahedra can be obtained from the vertical decomposition. For instance, each trapezoid facet of the decomposition, vertical or not, can be split in two triangles. Then each cell can be starred from a Steiner point added in the interior of the cell. The resulting triangulation has $O(n^2)$ tetrahedra.
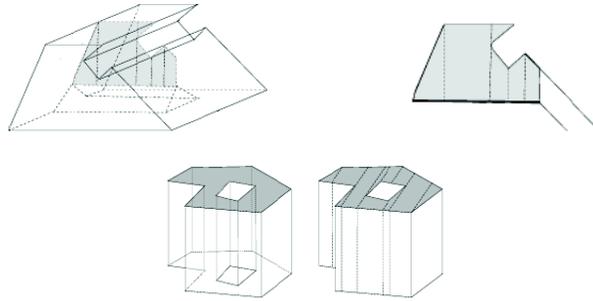
FIG. 4.1 – Vertical decomposition of a polyhedron : top left : the vertical segments stemming from a reflex edge, top right : the 2-walls of type 1, bottom : the cylindrical cells.

The main idea of proof of theorem 4.3 is to use the vertical decomposition but only after having simplified as much as possible the polyhedron. The simplification performed here concerns convex vertices. A vertex of a polyhedron is said to be *convex* when it is not incident to any reflex vertex. Let p be a convex vertex of a polyhedron Q. For each facet f of Q incident to p, we note $h(f)$ the hyperplane containing f and $h^+(f)$ the half space bounded by $h(f)$ and containing the edges of Q incident to p. The *cone* of p is the intersection of the closed half-space $h^+(f)$ for all the facets incident to p. We consider the set $\mathcal{V}(p)$ of vertices of Q included in the cone of p and note $\mathcal{V}'(p) = \mathcal{V}(p) \setminus \{p\}$. The *cup* of p is now define as the difference $\mathrm{conv}(\mathcal{V}(p)) \setminus \mathrm{conv}(\mathcal{V}'(p))$ between the convex hulls of $\mathcal{V}(p)$ and $\mathcal{V}'(p)$. See figure 4.2. The boundary of the cup of p include two parts : the *dome* formed by the facets of the cup which are facets of $\mathrm{conv}(\mathcal{V}'(p))$ and the *lateral boundary* formed by facets of $\mathrm{conv}(\mathcal{V}(p))$. The dome and the lateral boundary share a common boundary which is a closed polygonal line called the *crown*. The cup is star shaped with respect to p. If there is no vertex on Q that is on the dome without being the crown, the cup is said to be free. The simplification of the polyhedron consists in repeatedly removing free cups from the polyhedron. It has been shown than if this is performed as long as there exists convex vertex with a free cup, the residual polyhedron has size $O(r)$ where r is the number of reflex edges of the initial polygon. The polyhedron can be triangulated by staring each removed free cup from its apex and applying the vertical decomposition to the residual polyhedron. The resulting triangulation has size $O(n + r^2)$.

A lower bound for the decomposition of polyhedron

Theorem 4.4 *There are polyhedra with n reflex edges such that any of their triangulation has a size $\Theta(n^2)$.*

Proof.   The proof is the description of the polyhedron in Figure 4.3. This polyhedron is an axis parallel cube with unit size and n notches on the top and bottom faces. The reflex edges of the notches of the top face are parallel to the y-axis, equally spaced along to the x-direction, and they lie on the paraboloid
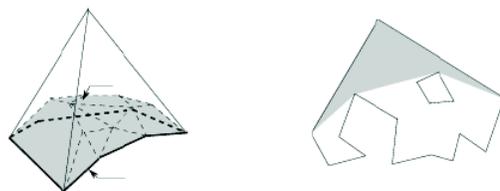
FIG. 4.2 – Reduction of a polyhedron : the left part shows a cup in dimension 3, the right image shows a cup in dimension 2
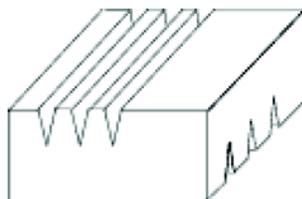
FIG. 4.3 – A polyhedron with only quadratic triangulations

$z = xy + \epsilon$ The reflex edges of the notches of the bottom face are parallel to the x axis, equally spaced along to the y-direction and they lie on the paraboloid with equation $z = xy$. It can be shown that any convex shape included in this polyhedron has a volume at most $O(1/n^2)$. Therefore any tetrahedral decomposition of such a polyhedron includes $O(n^2)$ tetrahedra. $\square$

4.4   Constrained Delaunay triangulations

The last section of this chapter presents a sufficient condition for a PLC $\mathcal{C}$ to admit a constrained triangulation. As it turns out, the triangulation produced to prove the sufficient condition is in fact a constrained Delaunay triangulation. Let us start with a few definitions.

Definition 4.5 (Delaunay and strongly Delaunay simplexes) *Let $\mathcal{P}$ be a set of points.*

A simplex s with vertices in $\mathcal{P}$ is a Delaunay simplex for the set $\mathcal{P}$ iff there is a circumsphere of s enclosing no point of $\mathcal{P}$.

A simplex s with vertices in $\mathcal{P}$ is a strongly Delaunay simplex for the set $\mathcal{P}$ iff there is a circumsphere of s enclosing no point of $\mathcal{P}$ and passing through no point of $\mathcal{P}$ except the vertices of s.

To define constrained Delaunay triangulations, we extend the notion of visibility introduced in section 2.4 for the 2-dimensional case. Given a PLC $\mathcal{C}$ the visibility obstacles are the facets of the PLC. More precisely, we say that point p is *visible* from point q, or that p and q are mutually visible if the interior of the segment pq intersects no facet of $\mathcal{C}$.

Definition 4.6 (Constrained Delaunay simplexes) *Let $\mathcal{C}$ be a PLC and let $\mathcal{P}$ be the set of vertices of $\mathcal{C}$. A simplex s with vertices in $\mathcal{P}$ is constrained Delaunay for $\mathcal{C}$ iff :*
– the interior of s intersects no face f of $\mathcal{C}$ except if s is included in f,
– there is a circumsphere of s enclosing no vertex of $\mathcal{C}$ visible from a point in the relative interior of s.

Then, the main result of this section reads as follows

Theorem 4.7 *Let $\mathcal{C}$ be a PLC and let $\mathcal{P}$ be the set of vertices of $\mathcal{C}$. If the edges in $\mathcal{C}$ are strongly Delaunay with respect to $\mathcal{P}$ and if $\mathcal{P}$ has no subset of five co-spherical points, the constrained Delaunay tetrahedra of $\mathcal{C}$ form together a constrained triangulation of $\mathcal{C}$.*
This triangulation is called the constrained triangulation of $\mathcal{C}$.

Remark. The edges of the PLC are required to be strongly Delaunay edges and not just Delaunay edges. Think for example at the Schönhart polyhedron. The six vertices of the Schönhart polyhedron lie on the

same sphere, so the edges of this polyhedron are Delaunay edges but not strongly Delaunay edges, and there is no constrained triangulation for this polyhedron.

Proof. Before going into the proof of theorem 4.7, we need one more definition relative to the facets of the PLC. Indeed the facets of $\mathcal{C}$ are not required to be triangular facet. They can be polygonal regions with arbitrary complexity and may be pierced with holes, slits or isolated vertices. In the constrained triangulation, those facets are going to be represented as a union of facets of the triangulation. In fact, we already know how facets of the PLC are going to be subdivided in the constrained Delaunay triangulation and we define the *subfacets* of the PLC facets as follows : Let $\mathcal{C}$ be a PLC whose edges are strongly Delaunay. Let f be a facet of $\mathcal{C}$ and h_f be the hyperplane including f. We call $\mathcal{V}_f$ the subset of vertices of $\mathcal{C}$ included in h_f and note $\mathrm{Del}(\mathcal{V}_f)$ the 2-dimensional Delaunay triangulation of the set $\mathcal{V}_f$ in h_f. The edges of $\mathcal{C}$ are strongly Delaunay and therefore the edges of $\mathcal{C}$ included in h_f are edges of $\mathrm{Del}(\mathcal{V}_f)$. Thus any triangle t of $\mathrm{Del}(\mathcal{V}_f)$ is either included in f or its interior is disjoint from f. The subfacets of a facet f are defined as the triangles t of $\mathrm{Del}(\mathcal{V}_f)$ that are included in f.

Now the proof of theorem 4.7 includes three steps which can be summarized as follows :

Step 1 proves that any point in the convex hull $\mathrm{conv}(\mathcal{C})$ of $\mathcal{C}$ is included in a constrained Delaunay tetrahedra.

Step 2 proves that constrained Delaunay tetrahedra form together a simplicial complex.

Step 3 proves that any subfacet of a facet in $\mathcal{C}$ is a facet of a constrained Delaunay tetrahedron.

□

4.4.1 Step 1 of the proof of theorem 4.7

Step 1 consist in proving the following lemma

Lemma 4.8 *Let $\mathcal{C}$ be a PLC whose edges are strongly Delaunay. Any point p in the convex hull $\mathrm{conv}(\mathcal{C})$ of $\mathcal{C}$ is included in some constrained Delaunay tetrahedron of $\mathcal{C}$*

Proof. The proof of this lemma is based on a process, called the growing sphere process, to build constrained Delaunay tetrahedra.

The growing sphere process The growing sphere process takes as input a k-dimensional constrained Delaunay simplex with its circumsphere enclosing no visible vertex and builds up a $(k+1)$-dimensional constrained Delaunay simplex. For $k \leq 2$, let s be a k-simplex that is a constrained Delaunay k-simplex of $\mathcal{C}$ and let S be the circumsphere of s enclosing no vertex of $\mathcal{C}$ visible from the interior of s. Let h be a hyperplane including s and let h^+ be one of the two half-spaces bounded by h. The sphere S belong to a pencil of spheres sharing the intersection $S \cap h$. Starting from S, we move a sphere in that pencil, such that
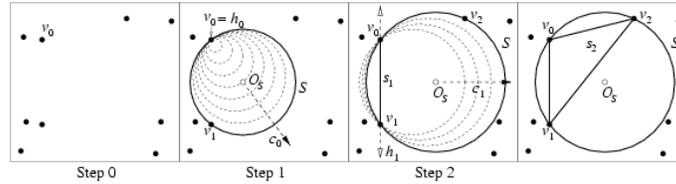
FIG. 4.4 – The growing sphere process.

the volume of h^+ enclosed by the sphere grows. The motion is stopped as soon as a vertex u of $\mathcal{C}$ visible from the interior of s is encountered by the growing sphere. We claim that the simplex $\mathrm{conv}(s, u)$ is constrained Delaunay. and this claim is proved by the growing sphere lemma (lemma 4.12) stated below.

The same growing sphere process can be performed from a 2-simplex s which is a constrained subfacet. In that case the pencil of spheres is defined by the circumcircle of s and the starting sphere is just the hyperplane including s. A variant of the growing sphere lemma (lemma 4.17) proves that the simplex $\mathrm{conv}(s, u)$ is constrained Delaunay.

Of course, there is a condition for the success of a growing sphere process which is the existence in half-space h^+ of a vertex of $\mathcal{C}$ visible from a point in the interior of s. The visibility lemma 4.9 stated below ensures that, because $\mathcal{C}$ has strongly Delaunay edges, if half-space h^+ contains a vertex of $\mathcal{C}$, it contains a vertex visible from a point in the interior of s.

Let us come back to the proof of lemma 4.8. Let p be a point of $\mathrm{conv}(\mathcal{C})$, we build a constrained Delaunay tetrahedron including p. This is done in two phases. The first phase builds a first constrained Delaunay tetrahedron for $\mathcal{C}$. It starts from any vertex in $\mathcal{C}$, and using the growing sphere process builds a constrained Delaunay edge, then a constrained Delaunay facet and at last a constrained Delaunay tetrahedron. At each substep, the hyperplane h^+ is chosen so that its includes vertices of $\mathcal{C}$ and the visibility lemma 4.9 below grants the success of the growing process. At the end of this phase we have one constrained Delaunay tetrahedron t. Either t includes p and we are done, or we choose a point q inside t such that pq intersect no edge of $\mathcal{C}$ and we process to the second phase. A each substep of this second phase, we build a new constrained tetrahedra as follows. We consider the facet f of the last built tetrahedron which is intersected by pq. This facet f is either a constrained Delaunay facet or a subfacet of $\mathcal{C}$. Let h_f^+ be the half-space bounded by the hyperplane containing f that contains p. We grow a new tetrahedron in the half-space h_f^+. Because h_f^+ contains the point p which is in $\mathrm{conv}(\mathcal{C})$, it contains some vertex of $\mathcal{C}$ and the visibility lemma 4.9 below grants the success of the growing process. The phase ends as soon as the new tetrahedron contains p which has to happen after a finite number of substeps because there is a finite number of possible tetrahedra and each new built tetrahedron covers a piece of segment pq with non zero length. This ends the proof of lemma 4.8, provided that we prove the instrumental lemmas 4.9, 4.12 and 4.17. $\square$
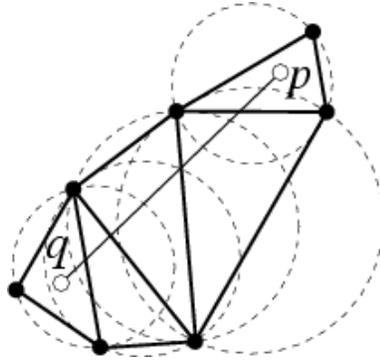
FIG. 4.5 – Building the constrained Delaunay tetrahedron that contains p .

4.4.2 Step 2 and step 3 of the proof of lemma 4.7

To perform Step 2, we first prove that two constrained Delaunay tetrahedra have no common interior point. Let t and s be two constrained Delaunay tetrahedra, and let S_t and S_s be their respective circumspheres. Assume for contradiction that s and t have some common interior point p. Because s and t are constrained Delaunay tetrahedra their vertices are visible from p. The circumsphere S_t and S_s are different because $\mathcal{C}$ has no subset of five co-circular points and they have to intersect because they both contains p. Let h be the radical hyperplane of S_t and S_s. The two halfspaces bounded by h are the power cells of the spheres S_s and S_t in their power diagram. Let h^s be the half-space of points with smallet power to S_s than to S_t and h^t the other half-space. Assume wlog that p belong to the $h \cup h^t$. Because p is in the interior of s, some vertex v of s has to be in h^t. Vertex v has power zero with respect to S_s and a smaller, thus negative, power with respect to S_t. Thus v is enclosed by S_t and visible from the interior point p of t this contradicts the fact that t is a constrained Delaunay tetrahedra.

Thus, any two constrained Delaunay tetrahedra have no common interior point. Assuming for now, the proof of lemmas 4.9, 4.12 and 4.17, we know from the previous subsection that, using the growing sphere process we can build a constrained Delaunay tetrahedron incident to any facet of a constrained Delaunay tetrahedron which is not on the convex hull of $\mathcal{C}$. Therefore constrained Delaunay tetrahedra meet neatly along their facets and because their union covers the convex hull of $\mathcal{P}$, they form collectively a triangulation.

To perform Step 3, it suffice to notice that the growing sphere process described in the proof of lemma 4.8 allows to build a constrained Delaunay tetrahedron incident to any subfacet f of $\mathcal{C}$. Therefore any subfacet of $\mathcal{C}$ is a facet of some constrained Delaunay tetrahedra.
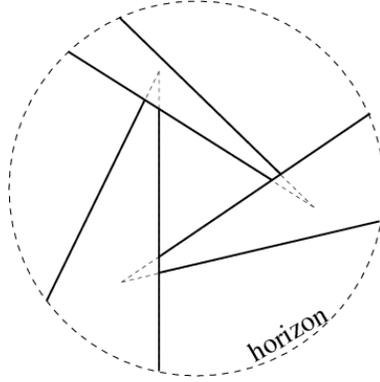
FIG. 4.6 – Overlapping facet cycle.

4.4.3   Instrumentals lemmas

Lemma 4.9 (Visibility lemma) *Let C be a PLC with strongly Delaunay edges. Let h be a hyperplane and p a point in h. If the open half-space h^+ bounded by h contains a vertex of C, it contains a vertex of C visible from p.*

Proof.   Figure 4.6 shows what could be seen by an observer located at point p : a cycle of overlapping constrained facet with no visible vertex. The proof of this lemma just consist in showing that such a cycle cannot occur if the constrained edges are strongly Delaunay.

Let s and t be two edges in C. We say that s covers t from the point p if there is a point p_s on s and a point p_t on t such that p_s does not belong to t and lies in the segment pp_t. Now the proof of lemma 4.9 derives from lemmas 4.10 and 4.11 below. Indeed, lemma 4.10 states that if there is some vertex of C in h^+ but no vertex visible from p, there must be, among edges of C, a cycle of covering edges from p and lemma 4.11 state that such a cycle cannot exist when the edges of C are strongly Delaunay.

Lemma 4.10 *If the PLC C has some vertices in h^+ and if none of them is visible from p, there must be among the edges of C a cycle of covering edges from point p.*

Proof.   Let u be a vertex of C in h^+. Either u is visible from p or there is a facet f of C occluding the visibility between p and u. Let q be the point of intersection $pv \cap f$. At least one part of the boundary of f lies in h^+. We move a point q' on f from point q towards a point on $\partial f \cup H^+$ keeping the visibility with p. The motion stop when either q' reaches a point m on an edge e of f visible from p or the visibility is blocked, which means that there is an edge e of C intersecting $q'p$. In the second case we set $m = e \cap q'p$. In both cases we have a point m of h^+ lying on an edge e of C and visible from p. Because $m \in e$, the edge e
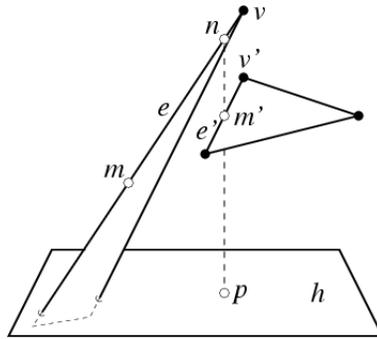
FIG. 4.7 – Some vertex in h^+ is visible from p.

has at least one of its vertex v in h^+. See Figure 4.7. If v is visible from p we are done. Otherwise we walk on e from m toward v until we reach the first point n were the visibility is blocked by an edge e' of $\mathcal{C}$. Edge e' covers e from p and has a point m' ($m' = pn \cap e'$) in h^+ visible from p. Therefore edge e' has a vertex v' in h^+. Either v' is visible from p and we are done or we continue the process walking on e' toward its vertex in h^+. Because $\mathcal{C}$ has a finite number of edges, the process either ends up with a vertex of $\mathcal{C}$ visible from p or $\mathcal{C}$ has a cycle of covering edges from p. $\square$

Lemma 4.11 *If the PLC $\mathcal{C}$ has strongly Delaunay edges, there is no cycle of covering edges from a point p.*

Let s and t be two strongly Delaunay edges of $\mathcal{C}$. Let S_s (resp. S_t) be the circumspheres of s (resp. of t) that do not enclose nor pass through vertices of $\mathcal{C}$. Let h be the radical hyperplane of the spheres S_s and S_t. The vertices of s have a null power with respect to S_s and a positive power with respect to S_t because t is strongly Delaunay. Therefore the vertices of s, and hence s itself, are in the half-space h^s bounded by h including points that have a smaller power to S_s than to S_t. For the same reason t is in the other half-space h^t. Now if s covers t from p, the segment pp_t issued from p intersect s before t, which implies that p belong to h^s. See Figure 4.8 Therefore the power $\Pi(p, S_s)$ of p to S_s is less than its power $\Pi(p, S_t)$ to S_t. This ordering prevents the existence of a cycle of covering edges from a point p. $\square$

Lemma 4.12 (Growing sphere lemma) *Let $\mathcal{C}$ be a PLC with strongly Delaunay edges. Let s be a constrained Delaunay simplex. Let u be a vertex of $\mathcal{C}$ visible from a point p in the interior of s and such that there is a circumsphere $S(s, u)$ of $\mathrm{conv}(s, u)$ that encloses no vertex of $\mathcal{C}$ visible from the interior of s. Then $\mathrm{conv}(s, u)$ is a constrained Delaunay simplex for $\mathcal{C}$.*

Proof. We first prove first that any point r in s is visible from u and then that the sphere $S(s, u)$ encloses no vertex of $\mathcal{C}$ visible from the interior of $\mathrm{conv}(s, u)$. The first claim ensures that the interior of simplex
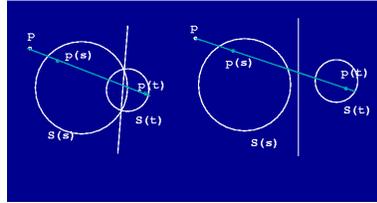
FIG. 4.8 – There is no cycle of covering edges
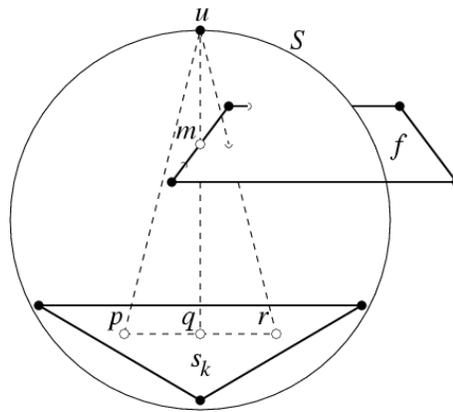
FIG. 4.9 – For the proof of claim 4.13

$conv(s, u)$ intersects no face of $\mathcal{C}$ and the second achieve to prove that $conv(s, u)$ is a constrained Delaunay simplex for $\mathcal{C}$.

Claim 4.13 *Any point r in s is visible from u.*

Assume for contradiction that there is a point r of s that is not visible from u. Because p is visible from u, there is a point q where the visibility is first blocked when moving from p to r. This means that qu intersects an edge e of $\mathcal{C}$ in a point m. See Figure 4.9. The point m is enclosed by $S(s, u)$ and visible from p. Then, we prove in the two lemma 4.15 and 4.16 below that that :

1. there is a vertex of e enclosed by $S(s, u)$ (lemma 4.15)
2. there is a vertex of $\mathcal{C}$ enclosed by $S(s, u)$ and visible from p (lemma 4.16),

which provides a contradiction with the definition of $S(s, u)$.

Claim 4.14 *The sphere $S(s, u)$ encloses no vertex of $\mathcal{C}$ visible from the interior of $conv(s, u)$*
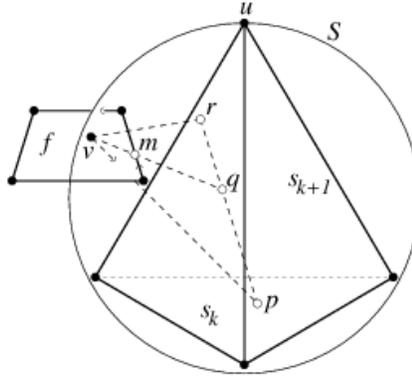
FIG. 4.10 – For the proof of Claim 4.14

Assume for contradiction that there is a vertex v of $\mathcal{C}$ visible from a point r in the interior of $\operatorname{conv}(s, u)$. Vertex v is visible from r but not from p. Let f be the facet occluding v from p. If there are several such facets choose the facet that intersect pv closest to p. We walk on segment pr from p towards r until we reach the first point q where either segment qv intersect an edge on the boundary of f (which has to occur before reaching r) or segment qv intersects an edge of another facet f'. See Figure 4.10. In both cases, let e be the intersected edge and m the intersection point. Point m is enclosed by $S(s, u)$ and visible from p. Then, lemma 4.15 asserts that there is a vertex of e in $S(s, u)$ and lemma 4.16 asserts that there is a vertex of $\mathcal{C}$ in in $S(s, u)$ visible from p, which contradicts the definition of $S(s, u)$.

This achieve the proof of the growing sphere lemma provided the lemmas 4.15 and 4.16 below. $\square$

Lemma 4.15 *Let $\mathcal{C}$ be a PLC, S a sphere and H_S the convex hull of the vertices of $\mathcal{C}$ enclosed by S or on S. If t is a strongly Delaunay simplex that intersects the interior of H_S, one of the vertices of t is enclosed by S.*

Proof. Let S_t be the circumsphere of t enclosing no vertices of $\mathcal{C}$. The strongly Delaunay simplex t intersects the interior of H_S, therefore H_S has vertices that do not belong to t and are not enclosed by S_t and the spheres S and S_t are distinct and intersecting spheres. Let h be the radical hyperplane of spheres S and S_t. Because t is strongly Delaunay, all the vertices of $\mathcal{C}$ in H_S are outside S_t and therefore lie on the half-space h^+ bounded by h of points with smaller power to S than to S_t. Hence, all vertices in H_S are in h^+ and because t intersects the interior of H_S, t has at least a vertex in h^+ which is therefore enclosed by S. See Figure 4.11 $\square$

Lemma 4.16 *Let $\mathcal{C}$ be a PLC with strongly Delaunay edges. Let S be a sphere and p a point in S. If there*
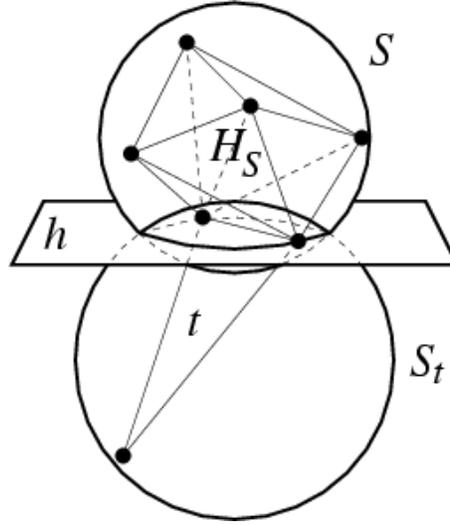
Fig. 4.11 – For the proof of lemma 4.15

is an edge e of $\mathcal{C}$ with a vertex v enclosed by S and a point m in e enclosed by S and visible from p, there is a vertex of $\mathcal{C}$ enclosed by S and visible from p.

Proof. If the endpoint v of e is visible from p, we are done. Otherwise we walk on e from m toward v until the we reach the point n where the visibility from p is first occluded. See Figure 4.12 Then segment pn intersects an edge e' of $\mathcal{C}$ in a point m' visible from p. Segment e' covers e from p. By lemma 4.15, one of the vertices of e' is enclosed by S and segment e' satisfies the hypothesis of lemma 4.16. Repeating the same process we find either a vertex of $\mathcal{C}$ enclosed by S and visible from p or a cycle of covering edges from p. The second possibility being ruled out by lemma 4.11, we are done. □

Lemma 4.17 (Growing sphere lemma bis) *Let $\mathcal{C}$ be a PLC with strongly Delaunay edges. Let s be a subfacet of a facet in $\mathcal{C}$ and let h be the hyperplane containing s. Assume that there is a vertex u of $\mathcal{C}$ in the halfspace h^+ such that u is visible from a point p in the interior of s and that the circumsphere $S(s,u)$ of $\mathrm{conv}(s,u)$ encloses no vertex of $\mathcal{C}$ in h^+ visible from the interior of s. Then $\mathrm{conv}(s,u)$ is a constrained Delaunay simplex for $\mathcal{C}$.*

Proof. Similarly to the proof of lemma 4.12, and establish the two following claims :
– any point r in s is visible from u.
– the sphere $S(s,u)$ encloses no vertex of $\mathcal{C}$ visible from the interior of $\mathrm{conv}(s,u)$
The proof of each claim is only slightly adapted to the fact that the sphere $\mathrm{conv}(s,u)$ may enclose vertices of $\mathcal{C}$ in the halspace h^-.
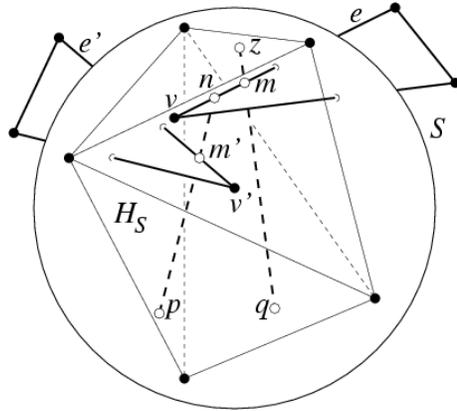
FIG. 4.12 – For the proof of lemma 4.16

Let p' be a point on pu, sufficently close to p so that ;
– p' cannot see any vertex of $\mathcal{C}$ in h^-,
– For any vertex v of $\mathcal{C}$ in h^+, line vp' intersects s,
– $\operatorname{conv}(s, p')$ intersect no facet of $\mathcal{C}$.
Because s is a subfacet of a PLC with strongly Delaunay edges, there is no vertices of $\mathcal{C}$ in $h \cap S$. Therefore any vertex of $\mathcal{C}$ visible from p' is in h^+.

Repeating the proof of lemma 4.12, with p' replacing p, we can prove that if there were a point r of s that is not visible from u, there would be a vertex w of $\mathcal{C}$ enclosed by $S(s, u)$ and visible from p'. Vertex w is in h^+ and visible from a point in s since wp' intersects the interior of s, wich provides a contradiction.

Consider the second claim and assume for contradiction that there is a vertex v of $\mathcal{C}$ enclosed by $S(s, u)$ and visible from some point r in the interior of $\operatorname{conv}(s, u)$. If v is in h^-, let $p'' = p'$. Otherwise, let p'' be a point on pv satisfying the same condition as p'. We repeat the argument in the proof of lemma 4.12 and deduce that there is a vertex w of $\mathcal{C}$ enclosed by $S(s, u)$ and visible from p''. As above, w must be in h^+ and visible from the interior of s, hence the contradiction.

$\square$

4.5 Bibliographical notes

The simplicial decomposition of a polyhedron of genus 0 is due to Chazelle and Palios [4] and the polyhedron with notches proving the worst case optimality of such a decomposition was proposed by Chazelle [3]. The sufficient condition to guarantee the existence of a constrained Delaunay triangulation of a piecewise linear

complex is due to Shewchuck[6].

Troisième partie

Mesh generation

Chapitre 5

Delaunay refinement in dimension 2

5.1 Introduction

A mesh is a cellular complex partitioning a given object or domain into elementary cells. The definition of a complex implies that the cells are convex, and that any two cells are disjoint or share a common lower dimensional face. The cells of a mesh are furthermore required to be elementary, which means that they admit a bounded size description (i.e. a bounded number of vertices and faces).

Meshes are ubiquitous in graphical applications, modelisations, simulations, scientific computing, CAD-CAM systems, reverse engineering, medical imaging processing ....

There are two families of meshes : the structured and unstructured meshes. In structured meshes, every vertex has the same environment, i. e. the same number of incident faces of any dimension. Structured meshes are highly appreciated in many applications, especially in computation using numerical methods like the finite element or finite difference methods. Indeed those meshes offer very economic storage and easy addressing of the different cells.

The unstructured meshes are mostly simplicial meshes. These meshes are appreciated for their flexibility to fit the geometry of the domain to be meshed.

The main challenges of a mesh generator can be summarized as follows :
– The mesh is required to respect or at least approximate accurately, the geometry of boundaries and internal constraints.
– The mesh cells should be well shaped. In case of simplicial meshes, most of the applications, in particular the finite element computations, required meshes whose cells have a shape closed to the shape of a regular simplex.
– The size of the cells should be adapted to the user requirements. These requirements can be defined using a sizing field which may be highly inhomogeneous.

Fig. 5.1 – An example of mesh.

Fig. 5.2 – Examples of structured meshes.
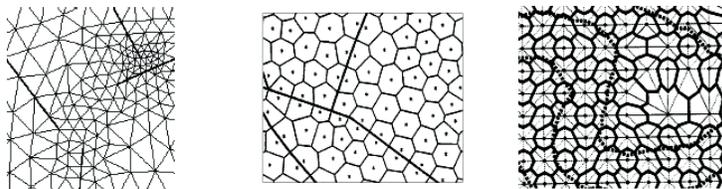
Fig. 5.3 – Examples of unstructured meshes.
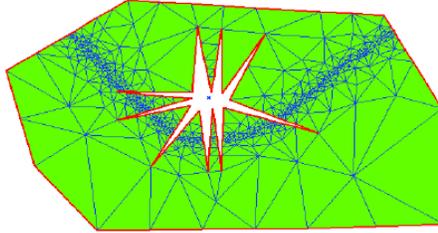
FIG. 5.4 – A mesh adapted to an inhomogeneous sizing field.

– The output mesh should fulfill the above goals with the smallest possible number of vertices. Most of the time, the efficiency of subsequent treatment is linearly related to the number of mesh elements.

The Delaunay triangulation is a fundamental tool for unstructured mesh generation. The following of these notes focuses on a Delaunay based meshing method which is called the Delaunay refinement.

5.2 Definition of a meshing problem

The input to a meshing problem is a PSLG (planar straight line graph) that describes the boundary of the domain to be meshed and also, may be, internal constraints whose geometry have to be respected. The domain to be mesh has to be bounded and its boundary has to be included in the edges of the input PSLG.

The input may also include shape and size criteria to be fulfilled by the mesh elements. We postpone to subsection 5.3 the description of shape criteria. The size criterion, if any, is an upper bound on some measure of mesh elements, for instance an upper bound on the length of longest edge or on the circumradius. The size bound may be inhomogeneous and given at each point of the domain by a sizing field. The sizing field is a function of the domain which may be itself described as a linear interpolation from a given mesh .....

Let $\mathcal{C} = (\mathcal{P}, \mathcal{S})$ be the input PSLG, whose set of vertices and edges are $\mathcal{P}$ and $\mathcal{S}$ respectively. We note Ω the domain to be meshed The searched mesh is a triangulation T such that :
– any vertex in $\mathcal{C}$ is a vertex of T,
– any edge in $\mathcal{C}$ is represented in T as a union of some edges of T,
– any triangle t of T included in the domain Ω satisfies the size and shape criteria.

5.3 The quality measures of a triangle

The Delaunay refinement process used in the algorithm described here is naturally designed to control the *radius-edge ratio* of mesh elements. The radius-edge ratio of a triangle t is the ratio ρ between the
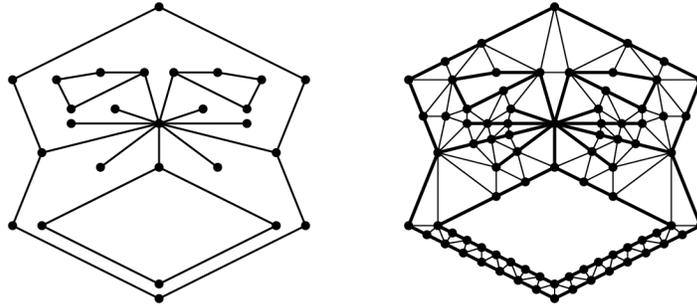
FIG. 5.5 – Left : the input PSLG. Right : the output mesh.

circumradius and the length of the shortest edge. Other measures can be used to quantify the quality of a triangle, however they are all more or less equivalent.

the smallest angle. Elementary geometry (see Figure 5.6) shows that the smallest angle α of triangle t is related to the radius-edge ratio by the relation

$$\rho = \frac{1}{2\sin\alpha} \tag{5.1}$$

maximum angle. If the smallest angle α, the maximum angle is at most $2\pi - \alpha$.

elevation ratio. The elevation-ratio of a triangle, noted ρ_h is the ratio of the length of the longest edge to the length of the shortest elevation. See Figure 5.6. The radius-edge ratio is related to the smallest angle by the relation

$$\frac{1}{\sin\alpha} \leq \rho_h \leq \frac{2}{\sin\alpha} \tag{5.2}$$

radius-radius ratio. The radius-radius ratio, ρ_i, is the ratio of the circumradius to the radius of the inscribed circle. Elementary geometry (see Figure 5.6) shows that the radius-radius ratio is related to the smallest angle by the relation

$$\frac{1}{\sin\alpha} \leq \rho_i \leq \frac{2}{2\sin^2\alpha} \tag{5.3}$$

Exercise 5.1 *Show that the shortest elevation of a triangle is always the elevation starting from the vertex with the maximum angle.*

Exercise 5.2 *Prove the relations 5.1, 5.2 and 5.3.*
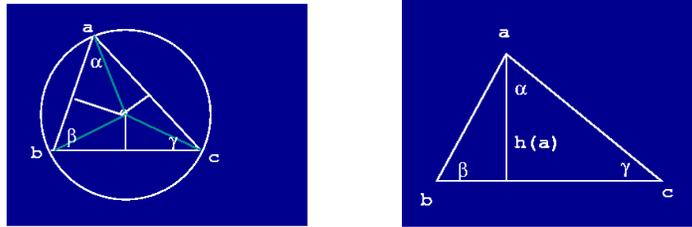
FIG. 5.6 – Elementary geometry of triangles
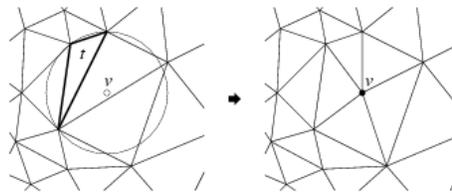
FIG. 5.7 – Removing a bad triangle by the insertion of its circumcenter

5.4 The Delaunay refinement algorithm

5.4.1 The basis of Delaunay refinement

Delaunay refinement based meshing algorithms maintain the Delaunay triangulation of a set points including the vertices of the input PSLG and the added Steiner vertices. In the following we say that a triangle is bad if it is included in the domain to be mesh and does not fulfill either the shape or the size requirement. The main idea of the Delaunay refinement is to iteratively remove the bad triangles from the current mesh by adding a new Steiner vertex at the circumcenter of this triangle. (see Figure 5.7).

However, the algorithm has to ensure that the edges of the input PSLG are represented in the final mesh as a union of mesh edges. this, the algorithm begins by refining the PSLG edges into subedges by adding Steiner vertices on these edges. In the following we called *constrained edges* the subedges of the PSLG edges formed by the Steiner vertices added on thoses edges. In fact, for a reason that will be clear later, the algorithm does not only refine constrained edges until they are Delaunay but a bit more, until they are *Gabriel edges*.

Definition 5.1 (Gabriel edges.) *An edge of a triangulation T is said to be a Gabriel edge if the smallest circumcircle of e encloses no vertex It T*

An edge e is said to be encroached *by a point p if the smallest circumcircle of e encloses p.*

5.4.2   The algorithm

The algorithm is more easily described with the help of two procedures called respectively `refine-edge(e)` and `conditionally-refine-facet(t)`.

`refine-edge(e)`. The procedure takes as input an edge e of the current triangulation and insert as vertex the midpoint of e.

`conditionally-refine-facet(t)`. The procedure takes as input a facet t of the current triangulation and computes the circumcenter c of t. Then,
  – if c encroaches some constrained edge e, `refine-edge(e)` is called
  – otherwise c is inserted as a vertex of the triangulation.

Now, the Delaunay refinement algorithm can be described as follows :

Initialization. Initialize the triangulation T with the Delaunay triangulation of the vertices in $\mathcal{C}$

Refinement. Apply one of the following refinement rules as long as one of them apply. There is a priority order on those rules, that is Rule R_2 is applied only if rule R_1 cannot be applied.

   Rule R_1. If there is a constrained edge e that is encroached by some vertex in T
       call `refine-edge(e)`

   Rule R_2. If there is a bad facet t in $T \cap \Omega$
       call `conditionnally-refine-facet(t)`

In summary the algorithm refines encroached constrained edges and bad facets with a priority given to the encroached facets. In particular, when the circumcenter of a facet is inserted in the triangulation there is no encroached constrained edge.

Note that an encroached constrained edge may be either a subedge that is missing in the triangulation because it is not a Delaunay edge, or an edge of the triangulation that is not a Gabriel edge. A variant of the algorithm consists in starting with the constrained Delaunay triangulation of $\mathcal{C}$ and maintaining the constrained Delaunay triangulation of the current set of subedges rather than the Delaunay triangulation of the current set of vertices. The final mesh is the same, but at any stage all constrained subedges are edges of the current triangulation.

5.4.3   The Delaunay refinement theorem

It is obvious that if it ends up, the Delaunay refinement algorithm yields a mesh whose triangles satisfy the size and shape criteria. The remaining of this section is devoted to the proof of the following theorem which gives conditions sufficient to ensure the Delaunay refinement algorithm terminates.

Theorem 5.2 (The Delaunay refinement theorem) *The Delaunay refinement algorithm terminates if the following conditions are satisfied*
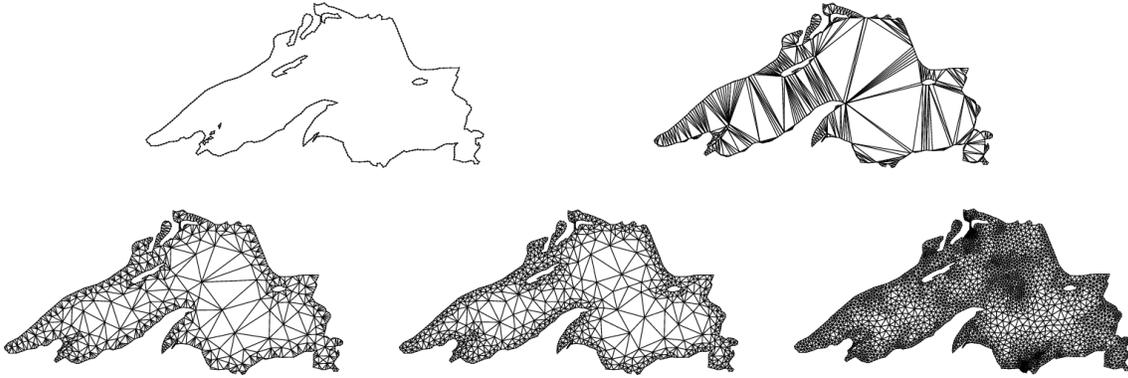
FIG. 5.8 – Top left : a PSLG. Top right : its constrained Delaunay triangulation. Bottom left to right : meshes obtained with the angular lower bound respectively equals to 15°, 25.6°, 34.2°.

– *There is no small input angles. This means that any two adjacent edges (i. e. edges sharing a vertex) of the input PSLG form angles that are at least* 60°.
– *The shape criteria is an upper bound b on the radius-edge ratio of the mesh triangles (i. e. of the form $\rho(t) \leq b$ where $\rho(t)$ is the radius-edge ratio of triangle t) and that the bound b is greater than sqrt(2).*
– *There is no size criteria or the size criteria is fulfilled by any triangle whose circumradius is less than a given sizing field.*

Note that the shape criteria implies a lower bound of $\arcsin \frac{1}{2b}$ for any angle of any triangle in the mesh. This lower bound on mesh angles is $20,7°$ for $b = sqrt(2)$. Thus $20,7°$ is the maximum lower bound guaranteed by the theorem. However, note that, the algorithms may terminates even if the condition of theorem 5.2 are not realized. For examples, the pictures in Figure 5.8 show that some $PSLG$ can be triangulated using a shape criterium that corresponds to a lower bound on mesh angles greater than $20,7°$.

In the following, we prove the Delaunay refinement theorem (theorem 5.2) in the case where there is no sizing field. The proof then can be easily modified to take into account a sizing field. The proof of the Delaunay refinement is based on a volume argument to bound the number of Steiner vertices. It needs a few definition and lemmas.

Lemma 5.3 (Steiner vertices lemma.) *Any Steiner vertex is inside or on the boundary of the domain Ω to be meshed.*

Proof. The lemma is trivial for Steiner vertices inserted on a constrained edge. Let t be a triangle in the current triangulation. Assume that t is included in Ω while c is not. Then there is a constrained edge e whose supporting line separates a vertex v of t from c. (See Figure **??**). Because t is a Delaunay triangle,
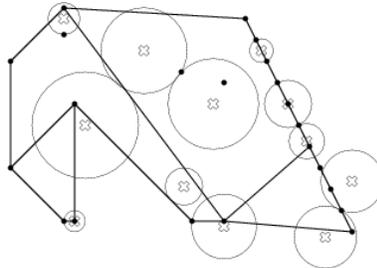
FIG. 5.9 – The local feature size.

the circumcircle of does not enclose the endpoints of e and e intersects this circumcircle in two points a and b. The smallest circumcircle of e encloses the smallest circumcircle of segment ab which encloses the vertex v of t. Therefore edge e is encroached. Because the algorithm maintains the fact that no constrained edge is encroached when the cirumcenter of a triangle is inserted, any triangle circumcenter inserted as a Steiner vertex is inside or on the boundary of the domain Ω. $\square$

Definition 5.4 (Local feature size.) *Given a PSLG $\mathcal{C}$ and a point p, the* local feature size *of $\mathcal{C}$ at point p, lfs(p), is the radius of the smallest disk centered at p and intersecting two disjoint faces of $\mathcal{C}$, i. e. two vertices, a vertex and a non incident edge, or two disjoint edges..*

It is trivial to show that, considered as a function of point p, the local feature size is a Lipschitz function, that is : for any two points, u and v

$$\text{lfs}(u) \leq \text{lfs}(v) + \|uv\|$$

The following definitions concern vertices that are currently in the triangulation or *rejected vertices* that are triangle circumcenters which at some stage have been considered for insertion in the mesh and rejected because they were encroaching some constrained edge.

Definition 5.5 (Insertion radius) *Let v be a vertex of the current triangulation T The* insertion radius *r_v of v if the length of the shortest edge incident to v, right after the insertion of v.*
The definition extends to rejected vertices as follows. The insertion radius *of a rejected vertex v is the length of the shortest edge that would have been incident to v right after its insertion if v had been inserted in the triangulation at the time it was rejected.*

Let us describe more precisely the insertion radius for each category of vertices.
– Vertices of the input PSLG, are considered to be inserted all at the same time. Therefore, if v is a vertex of the input PSLG, the insertion radius r_v is the distance from v to the vertex of $\mathcal{C}$ that is nearest to v.
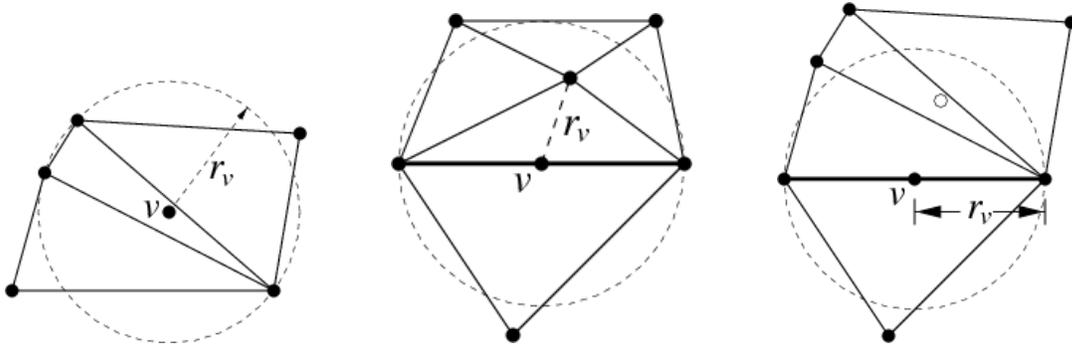
FIG. 5.10 – Insertion radius. Left : v is a triangle circumcenter. Middle : v is the midpoint of a constrained edge encroached by a rejected vertex. Right : v is the midpoint of a constrained edge encroached by triangulation vertices.
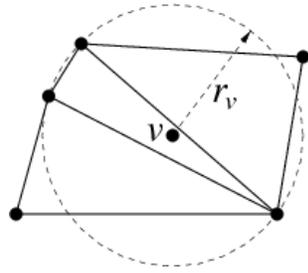
FIG. 5.11 – Parent. The parent p of a vertex v that is a triangle circumcenter.

– if v is the circumcircle of a triangle t, inserted in the mesh of rejected, the insertion radius r_v is the circumradius of t.
– Let v be a vertex inserted in a constrained edge e ,
  – $r_v = \|e\|/2$ if e is encroached by a rejected circumcenter
  – r_v is the distance from v to the closest vertex encroaching e otherwise.

Definition 5.6 (Parent vertex) *Each inserted or rejected vertex is associated to a parent vertex.*
– *If vertex v is the inserted or rejected circumcenter of a triangle t, the parent p of v is, among the two vertices of the shortest edge of t, the one that was inserted last.*
– *If vertex v is a vertex inserted in a constrained edge e, its parent p is the vertex encroaching e that is closest to v, p may be a vertex of the current mesh or a rejected vertex.*
A vertex v which is an input vertex of $\mathcal{C}$ has no parent.

Lemma 5.7 (The insertion radius lemma) *There is a constant c such that the insertion radius r_v of any vertex v of the mesh is such that*
- *either $r_v \geq \mathrm{lfs}(v)$*
- *or $r_v \geq cr_p$, where r_p is the insertion radius the the parent p of v.*

Proof.   More precisely, we show the following facts :

Fact 5.8 *If vertex v is a vertex of the input PLSG, $r_v \geq \mathrm{lfs}(v)$.*

Fact 5.9 *If vertex v is the circumcenter of a triangle, either inserted or rejected, $r_v \geq br_p$.*

Fact 5.10 *If v is a vertex inserted in an edge of the PSLG and if the parent p of v is rejected, $r_v \geq \frac{1}{\sqrt{2}} r_p$.*

Fact 5.11 *At last, if v is a vertex inserted on an edge e of the PSLG and if the parent p of v is on another edge e' of the PSLG, we have :*
- *$r_v \geq \mathrm{lfs}(v)$, if e and e' are disjoint,*
- *otherwise, let α be the smallest angle formed by e and e',*
   - *$r_v \geq \frac{1}{2\cos\alpha} r_p$, if $\alpha \in [45°, 90°]$*
   - *$r_v \geq \sin\alpha\, r_p$, if $\alpha \leq 45°$.*

Fact 5.8 is easy : the insertion radius of a vertex v of the input PSLG is the distance from v to its nearest input vertex which is at least $lfs(v)$ from the definition of the local feature size.

If the circumcenter v of triangle t is considered for insertion, rule R_2 is applied and the the insertion radius r_v, which is the circumradius of t, is at least $bl_{min}(t)$ where $l_{min}(t)$ is the length of the shortest edge of t. On the other hand, by definition of the parent p of v is the last inserted vertex of the shortest edge of t and the insertion radius r_p of p is at most $l_{min}(t)$. Therefore,

$$r_v \geq bl_{min}(t) \geq br_p$$

which proves Fact 5.9.

If the vertex v is a vertex inserted in a constrained edge e while the parent p of v is a rejected vertex, we have $r_v = \frac{\|e\|}{2}$. If a and b are the endpoints of e, the insertion radius r_p of p is less than $\min(\|pa\|, \|pb\|)$ which is less than $\frac{\|e\|}{\sqrt{2}}$ because p is enclosed by the smallest circumcircle of e. Therefore

$$r_v = \frac{\|e\|}{2} \geq \frac{1}{\sqrt{2}} r_p,$$

which proves Fact 5.10.
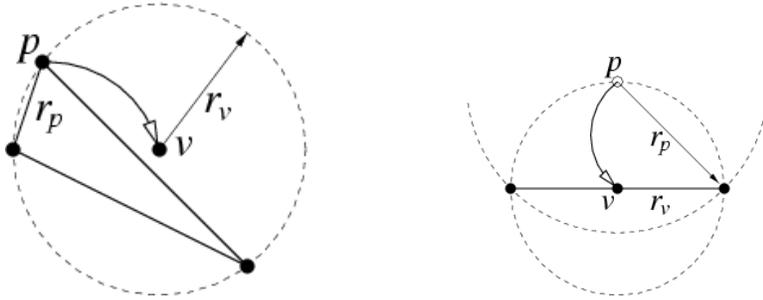
FIG. 5.12 – For the proof of facts 5.9 and 5.10
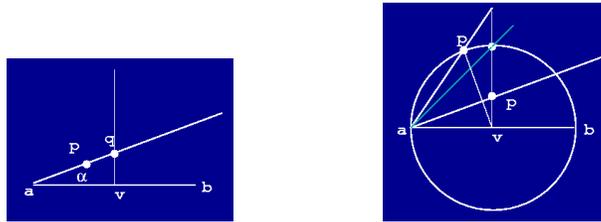
FIG. 5.13 – For the proof of Fact 5.11

At last consider the case of Fact 5.11. The insertion radius r_v of v is, in this case the distance $\|pv\|$ to its parent. If p and v belong to disjoint input edges this distance is more that lfs(v). Otherwise, let a and b be the endpoint of the edge e of $\mathcal{C}$ including v, and assume wlog that $\|pa\| \leq \|pb\|$. Then we have $r_p = \|pa\|$ and

$$
\begin{aligned}
\|pv\|^2 &= \|ap\|^2 + \|av\|^2 - 2\|ap\|\|av\|\cos\alpha \\
\frac{r_v^2}{r_p^2} = \frac{\|pv\|^2}{\|ap\|^2} &= 1 + \frac{\|av\|^2}{\|ap\|^2} - 2\frac{\|av\|}{\|ap\|}\cos\alpha
\end{aligned}
$$

Considered as a function of $x = \|ap\|$, the expression on the right hand side of the last equation is monotonously decreasing when p move away from a on e' and reaches its minimum, equal to $\sin^2\alpha$, when p reaches the position q which is the intersection of edge e' with the perpendicular to e from v. See Figure 5.13. However p is in the smallest circumcircle of e which encloses q if $\alpha \leq 45°$ and doest not enclose q if $\alpha \in [45°, 90°]$. Therefore, if $\alpha \leq 45°$ the minimum of $\frac{r_v}{r_p}$ is obtained when p is in q and is $\sin\alpha$. If $\alpha \in [45°, 90°]$, the minimum value of $\frac{r_v}{r_p}$ is obtained when p is at the intersection of aq with the smallest circumcircle of e, and is $\frac{1}{2\cos\alpha}$. This achieves the proof of Fact 5.11. $\square$

The last lemma proves a lower bound on the length of any edge in the mesh. Let lfs$_{min}$ be the minimum distance between two disjoint elements of C. Note that lfs$_{min}$ is the minimum of lfs(p) for any point p on a

face of the PSLG $\mathcal{C}$.

Lemma 5.12 (Lower bound on edge length) *If the input PSLG $\mathcal{C}$ has no pair of adjacent edges forming an angle less than $60°$, if the upper bound on radius-edge ratio is $b \geq \sqrt{2}$ and if there is no size criterion, the Delaunay refinement algorithm produces no edge shorter than lfs_{min}.*

Proof. The proof shows by induction that the insertion radius of each vertex in the mesh is at least lfs_{min}. This lower bound holds for the initial Delaunay triangulation. Assume that the bound holds until the insertion of vertex v. Let p be the parent of v and g be the parent of p if any. We note r_v, r_p and r_g the insertion radii of v, p and g respectively.

– If v is a triangle circumcenter, we know from Fact 5.9 that $r_v \geq br_p$ and that r_p is at least the length of the shortest edge of t, therefore $r_p \geq \mathrm{lfs}_{min}$ and

$$r_v \geq b\, r_p \geq b\, \mathrm{lfs}_{min} \geq \mathrm{lfs}_{min}.$$

– If v is a Steiner vertex inserted on an edge of $\mathcal{C}$ while its parent p is a rejected vertex, we know from Fact 5.10 that $r_v \geq \frac{r_p}{\sqrt{2}}$, and from Fact 5.9 that $r_p \geq br_g$ and by induction hypothesis we have $r_g \geq \mathrm{lfs}_{min}$. Therefore

$$r_v \geq \frac{r_p}{\sqrt{2}} \geq \frac{br_g}{\sqrt{2}} \geq r_g \geq \mathrm{lfs}_{min}.$$

– At last, let v be a Steiner vertex inserted on an edge e of $\mathcal{C}$ with a parent p on an edge e' of $\mathcal{C}$.
  – If the edges e and e' are disjoint, $r_v \geq \mathrm{lfs}(v) \geq \mathrm{lfs}_{min}$.
  – otherwise, because the edge e and e' forms an angle α of at least $60°$, we have from Fact 5.11

$$r_v \geq \frac{r_p}{2\cos\alpha} \geq r_p \geq \mathrm{lfs}_{min}.$$

Therefore in any case, the insertion radius r_v of vertex v is greater than lfs_{min}, which yields that the insertion of v does not produce edges with length smaller than lfs_{min}.
□

Proof. [Proof of the Delaunay refinement theorem 5.2.]

Recall that we assume that there is is no size criterion. Then, we know from the above lemma 5.12, that the refinement algorithm produces no edge shorter that lfs_{min}. Because, the triangulation is a Delaunay triangulation, it includes the edges connecting each vertex to its nearest neighbor. Therefore each vertex of the triangulation is at distance at least lfs_{min} from its nearestneighbor. We consider the set of disks centered at the mesh vertices and with radius $\frac{\mathrm{lfs}_{min}}{2}$. These disks have disjoint interior. Their centers lie inside Ω and because of the lower bound of $60°$ on the input angle it can be shown that at least $\frac{1}{6}$ of each disk is included in Ω, therefore we know that the area of Ω is such that

$$\mathrm{area}(\Omega) \geq \mathrm{n}\frac{1}{6}\pi\frac{\mathrm{lfs}_{min}^2}{4},$$
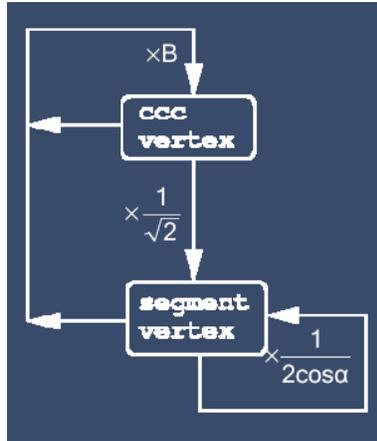
FIG. 5.14 – Flow chart for the Delaunay refinement showing the insertion radii multiplicative factors.

where n is the number of vertices of the final mesh. The above equation provides a bound on the number of vertices of the mesh and therefore proves that the algorithm terminates.

The proof can be easily modified to take into account a sizing criterion. See exercise 5.3. □

Exercise 5.3 *Assume that a sizing criterion is given by a sizing field $\sigma(x)$ defined on Ω. A triangle t with circumcenter c_t and circumradius r_t is over-sized and therefor bad if $r_t \geq \sigma(c_t)$. Modify the proof of the Delaunay refinement theorem 5.2 to take into account a sizing criterion.*

5.5 Optimality of the Delaunay refinement

The following theorem refines the bound on the number of vertices taking into account local variations of the local feature size. Then theorem 5.18 proves that the the Delaunay refinement yields optimal meshes with respect to the number of vertices.

Theorem 5.13 (Upper bound on the number of vertices) *If there is no size criterion, and if the input PSLG has no angle smaller than $60°$, the number n of vertices generated by the Delaunay refinement algorithm is at such that :*

$$n \leq \frac{6(3 + 2d_e)^2}{\pi} \int_\Omega \frac{dx}{\text{lfs}(x)^2}.$$

where Ω is the domain to be meshed and d_e is a constant depending on the radius-edge bound b and on the minimum angle of the input PLC.

To proof this theorem, we define *a weighted density* for mesh vertices and prove a bound for the weighted density of any vertex in the mesh.

Definition 5.14 (Weighted density) *The weighted density $d(v)$ of any vertex in the mesh is :*

$$d(v) = \frac{\text{lfs}(v)}{r_v}$$

where $\text{lfs}(v)$ *is the local feature size of the input PLC at v and r_v is the insertion radius of v.*

Lemma 5.15 (Weighted density lemma 1.) *For any vertex v, inserted in the mesh or rejected, if p is the parent of v and if the insertion radii of v of p are such that $r_v \geq cr_p$ for a given constant c, then $d(v) \leq 1 + \frac{d(p)}{c}$*

Proof. Because lfs() is a Lipschitz function,

$$\text{lfs}(v) \leq \text{lfs}(p) + \|pv\| \leq \text{lfs}(p) + r_v \leq d(p)r_p + r_v \leq \left(\frac{d(p)}{c} + 1\right) r_v.$$

$\square$

Lemma 5.16 *If there is no size criterion, and if the input PSLG has no angle smaller than $60°$, there are constant $d_e \leq d_f \geq 1$ such that :*
– for any circumcenter of a triangle, inserted or rejected, $d(v) \leq d_f$
– for any vertex v inserted on a PSLG edge $d(v) \leq d_e$.
Thus for any vertex in the mesh, $r_v \geq \frac{\text{lfs}(v)}{d_e}$.

Proof. First note that if $d_e \geq 1$ the result $r_v \geq \frac{\text{lfs}(v)}{d_e}$ holds for any vertex v of the PSLG. For Steiner vertices, the proof is be induction. Assume that the theorem holds up to the insertion of vertex v.
– If v is a triangle circumcenter, we have $r_v \geq br_p$ from Fact 5.9 and therefore $d(v) \leq 1 + \frac{d(p)}{b}$ from lemma 5.15. Thus, the lemma holds for v, if we have

$$1 + \frac{d_e}{b} \leq d_f. \tag{5.4}$$

– If v on an edge e of the PSLG
    – If the parent p of v is on an edge e' of the PSLG disjoint from e, $r_v = \text{lfs}(v)$ and $d(v) \leq 1$.
    – If the parent p of v is a rejected triangle circumcenter, $r_v \geq \frac{r_p}{\sqrt{2}}$ from Fact 5.10 and $d(v) \leq 1 + \sqrt{2}d_p$ from lemma 5.15. Thus, the theorem holds for v, if we have

$$1 + \sqrt{2}d_f \leq d_e. \tag{5.5}$$

– If the parent p of v is on an edge e' of the PSLG adjacent to e and forming an angle α with e, we have $r_v \geq \frac{r_p}{2\cos\alpha}$ from Fact 5.11 and $d(v) \leq 1 + 2\cos\alpha \, d_p$ from lemma 5.15. Thus, the theorem holds for v, if we have

$$1 + 2\cos\alpha \, d_e \leq d_e. \tag{5.6}$$

In summary, the theorem holds for the vertex v and is therefore proved if we can find d_e and d_f such that $d_e \geq d_f \geq 1$ and equations 5.4, 5.5 and 5.6 are satisfied for any input angle α. This is achieved if we choose

$$d_e \;\geq\; \max\left(\frac{(1+\sqrt{2})b}{b-\sqrt{2}}, \frac{1}{1 - 2\cos\alpha_{min}}\right)$$
$$d_f \;=\; 1 + \frac{d_e}{b}$$

where α_{min} is the smallest angle formed by adjacent edges of the input PSLG. $\square$

Lemma 5.17 (Local bound on edge length) *If there is no size criterion, and if the input PSLG has no angle smaller than $60°$, any edge in the mesh incident to vertex v has length $l(v) \geq \frac{\text{lfs}(v)}{d+1}$.*

Proof. Let vw be an edge incident to vertex v.
– If v was inserted after w, we have : $\|vw\| \geq r_v \geq \frac{\text{lfs}(v)}{d}$,
– otherwise we have : $\|vw\| \geq r_w \geq \frac{\text{lfs}(w)}{d} \geq \frac{\text{lfs}(v) - \|vw\|}{d}$ and the result follows.
$\square$

Proof. [Proof of the upper bound on the number of vertices theorem 5.13]
For any vertex v of the mesh, we consider the disc $\Sigma(v)$ with center v and radius $\frac{l(v)}{2}$. The discs have disjoint interior. Each disc has a fraction, with area at least $\frac{1}{6}$th of its total area, included in Ω. Therefore

$$\int_\Omega \frac{dx}{\text{lfs}(x)^2} \geq \sum_v \int_{\Sigma(v)\cap\Omega} \frac{dx}{\text{lfs}(x)^2} \geq \sum_v \frac{1}{6} \frac{\pi l(v)^2}{(4\text{lfs}(v) + l(v))^2} \geq \frac{n}{6} \frac{\pi}{(3 + 2d_e)^2}.$$

$\square$

Theorem 5.18 (A lower bound on the mesh size) *Any mesh with a minimum angle α_m of a domain Ω has a number n of vertices such that*

$$n \geq \frac{1}{3c^2\pi} \int_\Omega \frac{dx}{\text{lfs}(x)^2},$$

where the constant c depends on the minimum angle α_m.

The proof of this theorem includes a few lemma which aim to prove an upper bound on the longest edge of any triangle in the mesh. If $T(\Omega)$ is a mesh of the domain Ω, for any point p in Ω, we note $\mathrm{lm}(p)$ the length of the longest edge of the triangle t of $T(\Omega)$ that includes p.

Lemma 5.19 (Edge length ratio 1) *If T is a mesh with a minimum angle α_m,*
– the ratio between the lengths l_a and l_b of any two edges of the same triangle of T is upper bounded :

$$\frac{lb}{la} \leq \frac{1}{\sin \alpha_m}.$$

– the ratio between the lengths l_a and l_b of any two edges incident to the same vertex of T is upper bounded :

$$\frac{lb}{la} \leq \left(\frac{1}{\sin \alpha_m}\right)^{\frac{2\pi}{\alpha_m}}.$$

Proof. The first part comes from elementary geometry of the triangle, see for example Section 3.2. The second part is obtained by repeated application of the first part. Simply note that in a mesh with minimum angle α_m, each vertex is incident to at most $\left\lfloor \frac{2\pi}{\alpha_m} \right\rfloor$ triangles. $\square$

Lemma 5.20 (Longest edge lemma 1) *Let $T(\Omega)$ be a mesh with a minimum angle α_m of a domain Ω.*
– For any points p and q in adjacent triangles, we have

$$\mathrm{lm}(q) \leq \frac{1}{\sin \alpha_m} \mathrm{lm}(p).$$

– For any points p and q in Ω, we have

$$\mathrm{lm}(q) \leq c_1 lm(p) + c_2 \|pq\|$$

with

$$c_1 = \left(\frac{1}{\sin \alpha_m}\right)^{k+2} \quad \text{and} \ \ c_2 = 4 \left(\frac{1}{\sin \alpha_m}\right)^{k+2} \text{with} \ \ k = \left\lfloor \frac{\pi}{\alpha_m} \right\rfloor$$

Proof. The first part is obvious if one considers than $\mathrm{lm}(q)$ and $\mathrm{lm}(p)$ are both greater or equal to the length of the edge shared by the triangle containing p and q.

To prove the second part, we consider the set on triangles of $T(\Omega)$ intersecting the segment pq, ordered along pq. See Figure 5.15. Except for the first and last of these triangles, each of those triangles have two edges intersecting pq and one that does not. We call *lonely* the vertex of these triangles opposite to the edge that does not intersect pq, and we call fan the set of consecutive triangles along pq that share the same lonely vertex and we call *transition edge* an edge that separate two fans. A fan includes at most $K = \left\lfloor \frac{\pi}{\alpha_m} \right\rfloor$ triangles. Let l be the number of triangles intersected by pq.
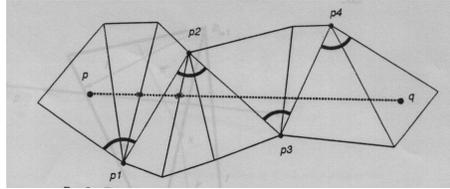
FIG. 5.15 – Triangles cross by the segment pq.
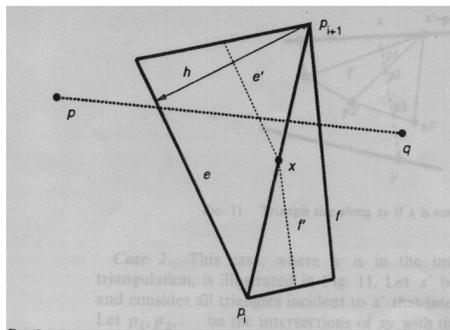
FIG. 5.16 – A transition edge crossed by pq.

– if $l \leq k + 3$, $\text{lm}(q) \leq \text{lm}(p) \left(\frac{1}{\sin \alpha} \right)^{k+2}$

– otherwise $l > k + 3$ and pq intersect at least one transition edge between two fans.

In this second case, let $p_i p_{i+1}$ be the last transition edge intersected by pq and let t_i and t_{i+1} be the two triangles of the mesh incident to $p_i p_{i+1}$. Let x be the midpoint of $p_i p_{i+1}$. Let e and f be the other edges of t_i and t_{i+1} respectively intersected by pq. Let e' and f' be the segment drawn from x in t_i and t_{i+1} respectively, parallel to e and f respectively. See Figure 5.16 Because pq intersects e, f and p_i, p_{i+1}, it has to intersect either e' or f'. Assume that pq intersects e', (the other case is similar). Then pq is greater than half the elevation h of t_i that is orthogonal to e. Let p' be a point in t_i. The elevation ratio of t_i (see Section 5.2) is at most $\frac{2}{\sin \alpha_m}$ and therefore :

$$\text{lm}(p') \leq \left(\frac{2}{\sin \alpha_m} \right) h \leq \frac{4}{\sin \alpha_m} \|pq\|$$

and

$$lm(q) \leq lm(p') \left(\frac{1}{\sin \alpha_m} \right)^{k+1} \leq \|pq\| \frac{4}{\sin \alpha_m} \left(\frac{1}{\sin \alpha_m} \right)^{k+1}$$

which achieve the proof of lemma 5.20

$\square$
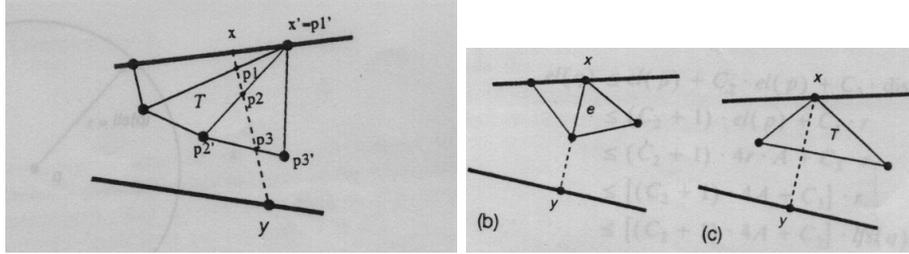
FIG. 5.17 – For the proof of lemma 5.21.

Lemma 5.21 (Longest edge lemma 2) *If x and y are two points on disjoint edges of the PSLG, there is a point q in segment xy such that $\mathrm{lm}(q) \geq \left(\frac{4}{\sin \alpha_m}\right) \|xy\|$*

Proof.   The proof is easy if either x, or y or both, are are vertices of the mesh. Let us therefore assume that neither x, nor y are vertices of the mesh. Then, xy has to cross a transition edge, (see Figure 5.17 and the proof is analog to the prove of the previous lemma (second case). $\square$

Lemma 5.22 (Longest edge lemma 3) *Let $T(\Omega)$ be a mesh with a minimum angle α_m of a domain Ω. For any point p in Ω, $\mathrm{lm}(p) \leq c_3 \mathrm{lfs}(p)$, for a given constant c_3 depending on the minimum angle.*

Proof.   Let p be a point in Ω. We consider the disc $\Sigma(p)$ centered in p with radius $\mathrm{lfs}(p)$. By definition of $\mathrm{lfs}(p)$, such a disk intersect two disjoint edges of the PSLG. Let x and y be two points in $\Sigma(p)$, one on each of the two disjoint edges intersecting $\Sigma(p)$. See Figure5.18. By lemma 5.21, there is a point q on xy such that :

$$\mathrm{lm}(q) \leq 4 \|xy\| \frac{1}{\sin \alpha_m} \leq 8 \, \mathrm{lfs}(p) \frac{1}{\sin \alpha_m}$$

From lemma 5.20

$$\mathrm{lm}(p) \leq c_1 \mathrm{lm}(q) + c_2 \|pq\|,$$

and therefore

$$\mathrm{lm}(p) \leq \left(\frac{8c_1}{\sin \alpha_m} + c_2\right) \mathrm{lfs}(p) \leq c_3 \mathrm{lfs}(p).$$

$\square$

Proof. [Proof of theorem 5.18.] The proof consists in summing $\frac{1}{\mathrm{lfs}(x)^2}$ on the domain Ω.

$$\int_\Omega \frac{dx}{\mathrm{lfs}(x)^2} \leq \sum_{t \in T(\Omega)} \int_t \frac{dx}{\mathrm{lfs}(x)^2} \leq c_3 \sum_{t \in T(\Omega)} \int_t \frac{dx}{\mathrm{lm}(x)^2} \leq c_3 n \frac{\sqrt{3}}{4}$$

$\square$
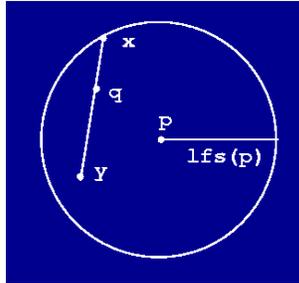
FIG. 5.18 – For the proof of lemma 5.22.
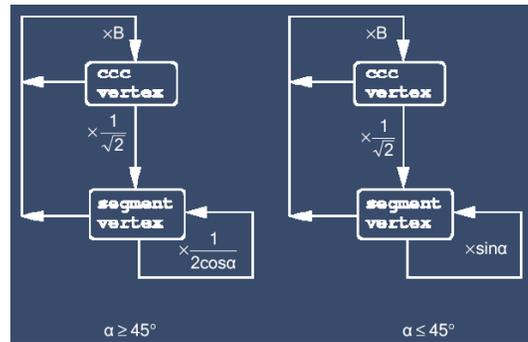
FIG. 5.19 – Flow chart of the Delaunay refinement algorithm and insertion radii.

5.6   Handling of small input angles

5.6.1   A negative result

The main problem of the Delaunay refinement algorithm when the input PSLG has small angles is that it might loop for ever. Indeed, as it is clear from the flow chart in figure 5.19, the algorithm may insert vertices whose insertion radii are smaller than their ancestor's insertion radius. Therefore, the lower bound on the length of mesh edges does not hold and neither the proof of termination.

If there are small input angles, i.e. if some pair of adjacent edges of the PSLG form small angles, there is no hope to build a mesh without small angles. Of course input angles are either still present in the final mesh or divide in smaller angles. However, there are still worse news. The following theorem simply state that there are PSLG that cannot be triangulated without adding small angles that are not already present in the PSLG.
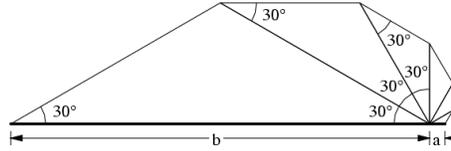
FIG. 5.20 – Edge lengths ratio.

Theorem 5.23 *For any bound θ on the mesh angles, there exist PSLG that cannot be meshed without creating angles smaller than θ.*

Proof. We first notice that any bound on the smallest angle of a mesh implies a limitation on the grading of triangle sizes. Indeed let l_a and l_b be the respective lengths of two constrained edges that are conscutive subedges of the same PSL edges. There is a bound for the ratio $\frac{la}{lb}$. Indeed, iterating the bound on the ratios between the lengths of edges in a triangle given in lemma 5.19 gives

$$\frac{lb}{la} \leq \left(\frac{1}{\sin \theta} \right)^{\frac{\pi}{\theta}} .$$

It can be shown that a tight bound for the ratio $\frac{la}{lb}$ is :

$$\frac{lb}{la} \leq b_1 \ \ \text{with} \ \ b_1 = (2\cos(\theta))^{\frac{\pi}{\theta}} . \tag{5.7}$$

This bound is reached by configuration of isosceles triangles of figure 5.20 :

An example of PSLG that cannot be meshed with a minimal angle bound of θ is the PSLG of figure 5.21. It includes three edges incident in a point o such that the two first edeges form a small angle ϕ, while the last two form a large angle, close to $\pi/2$ by example. Let e be the middle edge and assume that the PSLG also includes a vertex p that divide edges e in two subedges. Assume for contradiction that this PSLG has been meshed without angle smaller than θ. Then, the narrow wedge between the two first segment has been meshed, and vertex p is incident to a subedge pq of e such that o, p, q appear in this ordering and elementary geometry, see the left drawing on Figure 5.21, shows that the lengths $\|pq\|$ and $\|op\|$ are such that :

$$\frac{\|pq\|}{\|op\|} \leq b_2 \ \ \text{with} \ \ b_2 = \frac{\sin \phi}{\sin \theta} \left(\cos(\theta + \phi) + \frac{\sin(\theta + \phi)}{\tan \theta} \right) . \tag{5.8}$$

If the region above e is part of the domain, it has to be triangulated and the other subedge of e incident to p has a length l_b bounded by equation 5.7. Let us note b_1 and b_2 the two expressions on right hand side of equations 5.7 and 5.8 respectively. We have $l_b \leq b_1 \|pq\|$ and $\|pq\| \leq b_2 \|op\|$. Therefore, if $b_1 b_2 \leq 1$, op cannot be a mesh edge incident to p and there must be another vertex r between o and p on e. The same story can be restarted with r instead of p and we find that there is an infinite sequences of vertices on segment e. Therefore if $b_1 b_2 \leq 1$, such PSLG cannot be meshed ... unless with an infinite number of triangles ... For any bound θ, there is a choice of ϕ such that $b_1 b_2 < 1$. $\square$
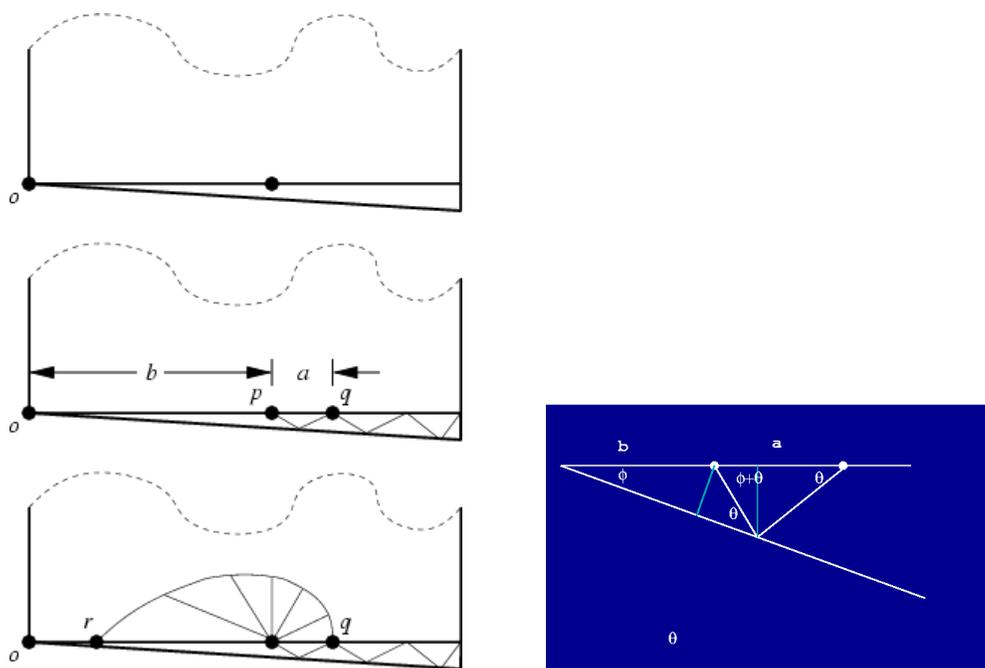
FIG. 5.21 – Right : a difficult PSLG. Left : the bound on $\frac{\|pq\|}{\|op\|}$.
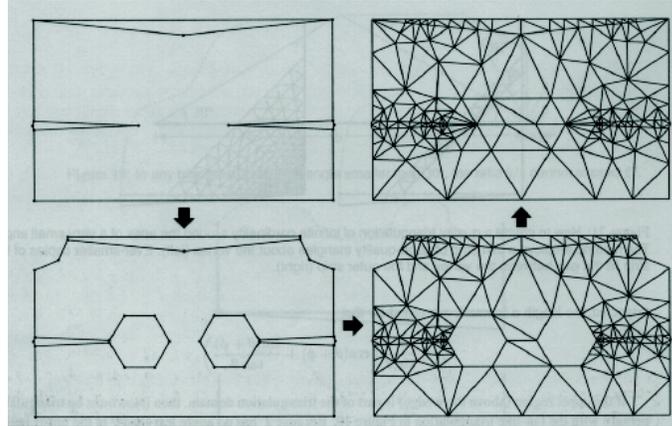
FIG. 5.22 – Corner looping

5.6.2   Corner looping

Corner looping is a first and quite natural solution to handle small input angles. In this method the input
PSLG $\mathcal{C}$ is replaced by a new one $\mathcal{C}'$ where small angle have been removed. To obtain $\mathcal{C}'$, we proceed as
follows for each apex v of a small angle :
– Vertex v is surrounded by a circle $\Sigma(v)$ of appropriate radius, 1/3 of lfs(v) is a reasonable choice for
  instance.
– Vertices are inserted at the intersection points of this circle with edges incident to v.
– Vertices inserted on $\Sigma(v)$ divide $\Sigma(v)$ into arcs. New vertices are further inserted on $\Sigma(v)$ if some arcs are
  greater than 60°
– vertices on $\Sigma(v)$ form a convex polygon $P(v)$ we obtain $\mathcal{C}'$ by adding the edges of this polygon and
  removing the subedges incident to v.
The PSLG $\mathcal{C}'$ has no small input angles and is meshed instead of $\mathcal{C}$. A mesh of $\mathcal{C}$ is then obtained, replacing
triangles inside the convex polygon $P(v)$ by a fan of triangles staring $P(v)$ from v.

Corner looping is conceptually simple and has the advantage that new small angles appear only at the apex
on input small angles. Furthermore it does not produce large angles : if θ is the minimal angle bound used
to mesh $\mathcal{C}'$ no angle larger than $2\pi - 2\theta$ appear in the mesh. The main drawback of this method is to reduce
the local feature size and to generally produce over refined meshes. The reminder of this section present an
alternative method, called terminator that, although conceptually more complicated is easy to implement
and tends to produce fewer triangles.

Exercise 5.4 *Show that if the radii of screening circles are chosen carefully, corner looping may achieve
that no new angle in the final mesh is smaller than the smallest input angle.*
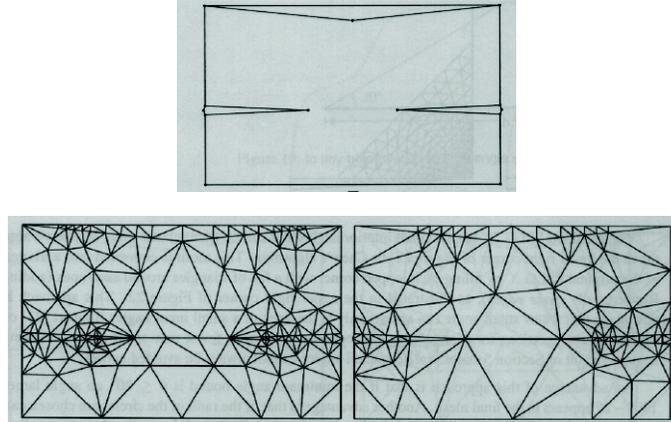
FIG. 5.23 – Up : the input PSLG. Left : a mesh produced using corner looping. Right : a mesh produced using Terminator.

5.6.3   Concentric shells refinement

A first problem caused by small input angles is the direct coupling between adjacent edges of the PSLG that form a small angle. This coupling is illustrated on Figure 5.24 and happens for input angles less than $45°$. Assume two edges of the input PSLG au and av form an angle θ less than $45°$ and that the shorter one, say av is :

– short enough to be enclosed the smallest circumcircle of au : $\|av\| \leq \|au\| \cos \theta$

– long enough for the midpoint w of au to be enclosed by the smallest circumcircle of av : $\frac{\|au\|}{2} \leq \|av\| \cos \theta$. Then point w has to be inserted in the triangulation because segment au is encroached by v, and the middle x of av is then inserted because av is encroached by w. Now, the vertices w and x reproduce exactly the same situation as u and v scaled by a factor 2 and an endless loop of refinement arises, see figure 5.24.

A possible way to avoid this endless loop is to use the concentric shells methods. This method is applied to refine the subsegments of a PSLG incident to the apexes of small input angles.

Definition 5.24 *Let $\mathcal{C}$ be a PSLG with small angles. A cluster of edges is a maximal subset of constrained edges incident to the same vertex and such that consecutive edges form an angle less than $60°$.*

We imagine that each apex of a small input angle is surrounded by concentric shells, i.e. by a set of circles centered at the apex and whose radii are all the integer power of 2 multiples of a given length l $\{2^i l : i \in \mathbb{Z}\}$. When an edge e of a cluster has to be refined, instead of using the midpoint of e as refinement point, we compute the intersection point between edge e and the concentric shell that is closest to the midpoint of e. Edges of the PSLG that are engaged in clusters at their both ends are split at their midpoints the first time.
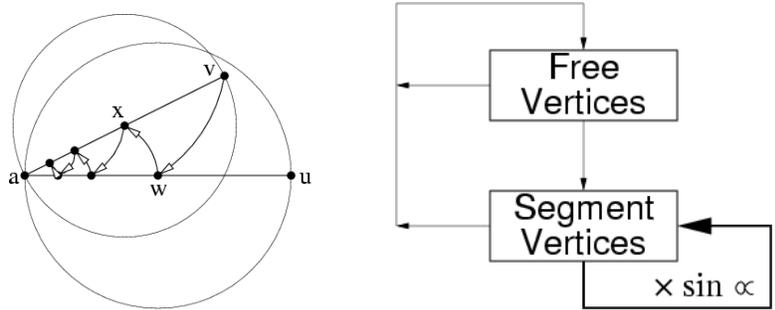
Fɪɢ. 5.24 – Endless refinement loop
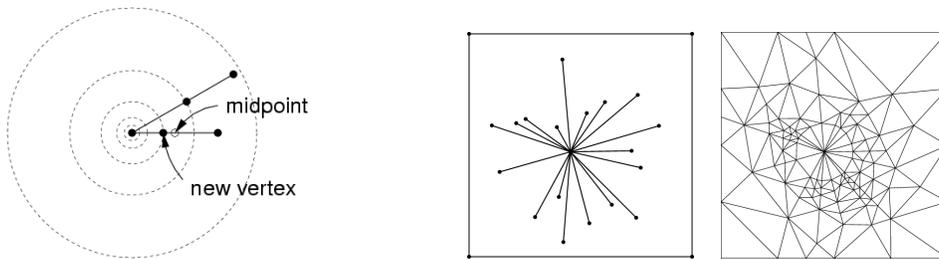
Fɪɢ. 5.25 – Concentric shells splitting. Left : the concentric shell and the choice of the refinement point. Right : A PSLG and is mesh generated using concentric shells splitting.

Concentric shells splitting solves the problem of direct coupling between adjacent edges because no mutual encroachment arises from adjacent constrained edges with equal lengths.

In fact, concentric shell splitting prevents endless loops due to the direct coupling of adjacent input segment forming a small angles. It does not prevent more complicated loops like the one appearing on Figure 5.26 where a constrained edge refinement triggers the refinement of another constrained edge which triggers triangle refinements and finally another constrained edge refinement.

Still, it is worth noting that concentric shell splitting is enough to solve the problem of meshing a polygon or polygonal region with holes, that is cases where every PSLG edges is on the boundary of the domain to be meshed. Indeed in the case of polygonal regions the creation of a short edge opposite to a small angle, does not further react on constrained edges incident to the apex of the small angles.
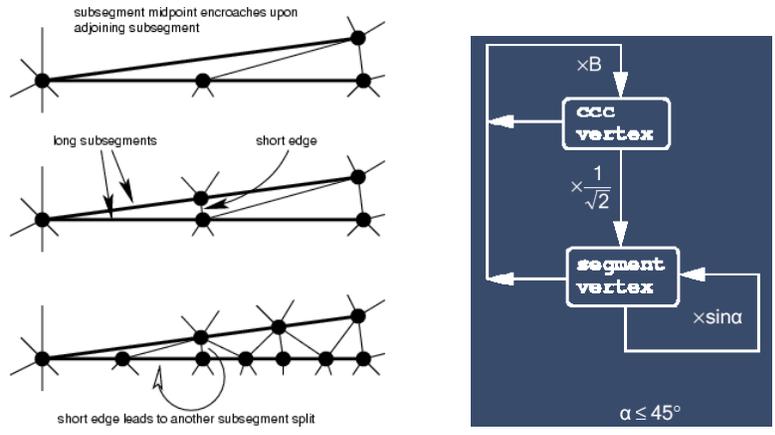
FIG. 5.26 – A complex indirect coupling and its flow chart

FIG. 5.27 – A simple strategy tend to produce large angle

5.6.4 Terminator

A simple way to ensure the termination of the Delaunay refinement process is to refuse the insertion of a vertex with insertion radius shorter than lfs$_{min}$. One way to implement this strategy is to refuse the insertion of any vertex whose insertion radius is smaller than the insertion radius of its most recently inserted ancestor (i. e. its parent if the parent is inserted, its grandparent if the parent is rejected) unless the parent is an input vertex or lies on a non incident segment. This simple strategy has two drawbacks. First the final mesh may have encroached edges and therefore the algorithm has to use a constrained Delaunay triangulation and the final mesh is not guaranteed to be Delaunay. Another drawback is that such a strategy tends to produce large angles : indeed, the algorithm will often allow the refinement of one of the edges incident to a small angle and refuse the refinement of the other, see Figure 5.27.

The terminator algorithm is a variant of the Delaunay refinement algorithm appliying the two following rules :
– Concentric shell splitting is used to refine constrained edges that are in clusters.
– The refinement induced by a bad triangle is refused if it reduces the insertion radius.
More precisely, let t be a bad triangle whose circumcenter p encroaches a subedge e in a cluster. Let v be the refinement point of e (according to the concentric shell splitting strategy) and let $r_{min}(v)$ be the

smallest insertion radius that will arise from the refinements of cluster edges that would be triggered by the insertion of v. Let g be the parent of p. The refinement of the cluster is agreed in one of the following cases

1. $r_{min}(v) \geq r_g$

2. t does not satisfy the size criteria if any

3. The cluster is not yet reduced, i.e. all the subedges of this cluster do not yet have the same lentgh

4. There is no ancestor of v in the PSLG edge including e. (This last case being optional).

In any other cases the circumcenter p is rejected without refinement of the cluster and the triangle t is kept in the mesh.

Theorem 5.25 (Terminator theorem) *The terminator algorithms ends and provides a Delaunay mesh with no encroached constrained edges. Small angles in the mesh occur only close to small input angles. If ϕ is the smallest input angle, there is no angles smallest than $\frac{\phi}{2\sqrt{2}}$.*

Proof. Let us call *diminishing vertex* a vertex inserted in the mesh and whose insertion radius is smaller than the insertion radii of all its ancestor. Diminishing vertices arise only in cases : **??** and 4 and only a finite number of vertices can be inserted for these cases. Therefore only a finite number of diminishing vertices are ever inserted in the mesh and when this process is over, the standard volume argument can still be applied to prove that the algorithm terminates.

The algorithm does not prevent the refinement of a constrained edge encroached by a mesh vertex, therefore the final mesh is Delaunay and has no encroached vertices.

Small angles in the mesh occur in bad triangles that are left in the mesh. Let t be a bad triangle left in the mesh. The circumcenter of t encroaches some constrained edge in a cluster and therefore cannot be far from a small input angle. Furthermore, we show below that the radius-edge ratio ρ of t is less than $\frac{1}{\sqrt{2}\sin(\phi/2)}$ which implies that the mesh has no angle smaller $\arcsin\left[\frac{\sin(\phi/2)}{2}\right] \approx \frac{\phi}{2\sqrt{2}}$.

We note p the circumcenter of t, l the lentgh of the shortest edges of t, r_v the half-length of the edges in the cluster encroached by p, r_{min} the smallest insertion radius of edges in the cluster, and r_p, r_g the insertion radii of p and its parent respectively. See Figure 5.28. We have $\rho = \frac{r_p}{l}$, with $r_p \leq \sqrt{2}r_v$, and $2r_v \sin\left(\frac{\phi}{2}\right) \leq r_{min}$. Also, $l \geq r_g$ and, because the refinement of the cluster has been refused, $r_g \geq r_{min}$. Therefore $\rho \leq \frac{1}{\sqrt{2}\sin(\phi/2)}$.
$\square$
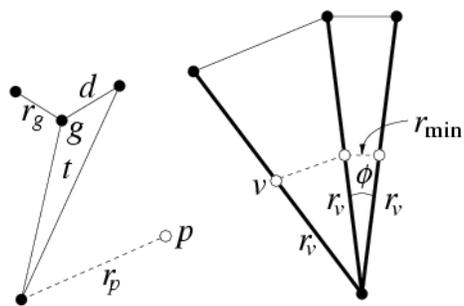
FIG. 5.28 – A bad triangle left in the mesh.

Chapitre 6

Delaunay refinement in dimension 3

6.1 Meshing problems in dimension 3

6.1.1 The constrained PLC

In dimension three, the geometry of the domain to be meshed and may be its internal structure are described through a piecewise linear complex (also called a PLC for short) $\mathcal{C}$. The domain Ω to be meshed is the union of some of the bounded three dimensional cells of $\mathcal{C}$. This imply in particular that the domain Ω is bounded and that its boundaries are included in the union of the facets of $\mathcal{C}$. A meshing problem is fully described by the data of the domain Ω together with some sizing and shape criteria for the mesh cells.

The output of the meshing problem is a mesh of domain Ω, i. e. is a triangulation T such that :
– Any vertex of $\mathcal{C}$ is a vertex of T.
– Any edge and facet of $\mathcal{C}$ is a union of faces in T
– The tetrahedra of T that are included in Ω satisfy the size and shape criteria.
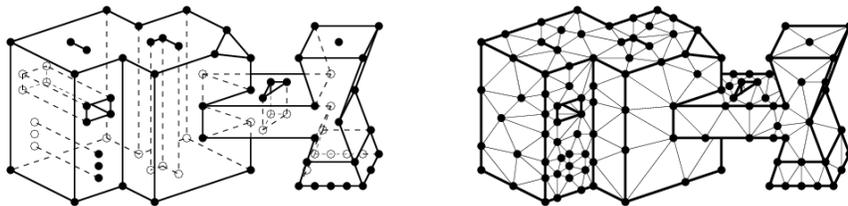
FIG. 6.1 – An input PLC (left) and the corresponding mesh (right).

6.1.2   Size and shape criteria

We assume that the size criterium is an upper bound on the circumradius of the tetrahedra in the mesh. This upper bound is allowed to be spatially inhomogeneous and can be for example given as a sizing field, i. e. as a scalar function whose defintion domain includes the domain Ω.

The shape criteria is, for now, assumed to be an upper bound on the radius-edge ratio of the tetrahedra in the mesh. The radius-edge ratio of a tetrahedra is the ratio $\rho = \frac{r}{l}$ between the circumradius r of the tetrahedra and the length l of its shortest edge. A bound on the radius-edge ratio od a tetrahedra implies the same bound for the facets of this tetrahedra and therefore o lower bound on the angles of those facets. However a bound on the radius-edge ratio do not imply a lower bound on the tetrahedra dihedral angle. We shall come back to this point latter.

6.2   The Delaunay refinement algorithm

6.2.1   Constrained edges and facets

The edges and facets of the input PLC $\mathcal{C}$ are going to be split into subedges and subfacets which are edges and facets of the final mesh. As in the two dimensional case, we call *constrained edges* the subedges of the PLC edges formed by Steiner vertices added on those edges.

The Delaunay refinement algorithm will refine constrained edges until they are Delaunay edges (and even Gabriel edges) for the current set of vertices.

Constrained facets are defined considering two dimensional Delaunay triangulations, as follows. Let F be a facet of the input PLC, we consider the hyperplan h_F spanned by F and the Delaunay triangulation T_F of the current subset of vertices lying in h_F. Once granted that constrained edges are Delaunay edges, the boundary edges of F are guaranteed to appear as edges in the Delaunay triangulation T_F. Hence the triangulation T_F is made of triangles that are either included in F or disjoint from the interior of F. We called *constrained facets* the facets of the Delaunay triangulation T_F that are included in F. To keep trace of the set of constrained facets, the refinement algorithm maintains, for each facet F of the PLC, a two dimensional Delaunay triangulation of the current set of vertices lying in the facet hyperplane. See Figure 6.2

6.2.2   The refinement algorithm

The algorithm maintains the Delaunay triangulation of the current set of vertices. A tetrahedron in this triangulation is said to be *bad* if it does not satisfy either the size or shape criteria. Rougly speaking, the Delaunay refinement algorithm will refine bad tetrahedra by circumcenter insertions until there are no longer bad tetrahedra. To ensure that the mesh respect the boundaries and constrained the algorithm maintains the fact that constrained edges are *Gabriel edges* for the current vertex set and the fact that constrained
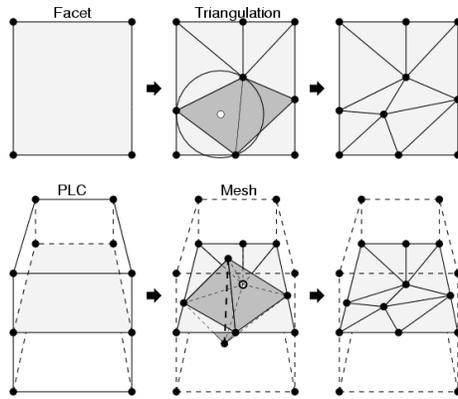
Fig. 6.2 – The constrained facets included in a PLC facet.
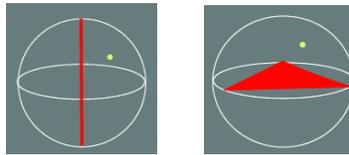
Fig. 6.3 – Encroached edge (left) and facet (right) .

facets are *Gabriel facets* where Gabriel edges and facets are defined as follows.

A constrained edge is said to be a *Gabriel edge* if its smallest circumsphere encloses no vertex of the triangulation. See Figure 6.3. A constrained facet is said to be a *Gabriel facet* if its smallest circumsphere encloses no vertex of the triangulation. Observe that a Gabriel edges and Gabriel facets are faces of the Delaunay triangulation. A constrained edge or a constrained facet whose smallest circumsphere encloses some vertex p is said to be encroached by p.

Note that an encroached constrained face may be either a constrained face that is missing in the triangulation because it is not a Delaunay, or a Delaunay face that is not a Gabriel face. Hereafter we call *circumcenter* a segment or of a triangle the center of its smallest circumsphere. The algorithm refines encroached constrained edges, and encroached constrained facets by insertion of their circumcenter. The refinement of encroached constrained edges is given priority over the refinement of encroached constrained facets which has a higher priority than the refinement of bad tetrahedra. Therefore when a constrained facet is refined there is no encroached constrained edge and when a bad tetrahedra is refined there is no encroached constrained edge nor encroached constrained facet.

We need to underline one more point before giving a precise description of the refinement algorithm. When a constrained facet f is considered for refinement, being encroached by a point p, there may be several
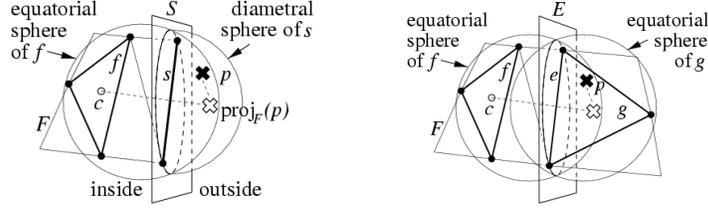
FIG. 6.4 – For the proof of lemma 6.1 Left : the projection of p is included in F. Right : p encroaches the facet $g \in F$ that contains its projection.

subfacets of the PLC facet F including f encroached by the same point p. The projection lemma below (lemma 6.1) states that the projection p_F of p on the hyperplane spanned by F is included in F and that the subfacet g of F containing p_F is also encroached by p. For reason that will be clear in the analysis of the algorithm, to optimize the bound on the radius-edge ratio of the mesh tetrahedra, the algorithm always refine the constrained subfacet g including the projection of the encroaching point.

Lemma 6.1 (Projection lemma) *Assume that a point p encroaches some constrained facet f included in a PLC facet F while no constrained edge is encroached :*
– the orthogonal projection p_F of p on the supporting hyperplan h_F of F belongs to F,
– p encroaches the constrained facet g in F that contains p_F.

Proof. Because no constrained edge is encroached, it can be shown as in lemma 5.3 that the circumcenter c of f is included in F. Therefore if we assume for contradiction that p_F is not included in F, the segment cp_F intersects at least one constrained edge e on the boundary of F. Let us consider the smallest circumspheres S_f of f and S_e of e. The radical hyperplan Π of the two spheres S_f and S_e is orthogonal to H_F, intersects H_F along the supporting line l_e of e and divides the space in two halfspaces. Let Pi^+ be the halspace including c and Pi^- the halfspace including p_F. Because pp_F is orthogonal to H_F, p belongs to the same halfspace Pi^- as p_F. Thus p belongs to the part of H_F^- enclosed by S_f. But because c is in H_F^+, the part of H_F^- enclosed by S_f is included in the part of H_F^- enclosed by S_e. Therefore p is enclosed by S_e which contradicts the fact that it does not encroaches e and proves the first part of the lemma.

Let us now show that p encroaches the subfacet g in F that contains p_F. If $f = g$, we are done. Otherwise let S_f and S_g be the smallest circumspheres of f and g respectively and let Π' be the radical hyperplane of S_f and S_g. Because the triangulation is Delaunay, f (resp. g) is included in the halfspace Π'^f (resp. Π'^g) that is locus of points whose power to S_f (resp. S_g) is smaller than their power to S_g (resp S_f). Thus p_F belongs to Π'^g. Because Π' is orthogonal to H_F, p and p_F are the same side of Π', and p belongs also to Π'^g. Then p belongs to the part of Π'^g enclosed be S_f which is included in the part of Π'^g enclosed be S_g. Therefore p is enclosed by S_g. $\square$

The algorithm is more easily described with the help of three procedures called respectively *refine-edge*, *conditionally-refine-facet* and *conditionally-refine-tetrahedron*.

`refine-edge(e).` The procedure takes as input an edge e of the current triangulation and insert the circumcenter of e as a vertex of the triangulation.

`conditionally-refine-facet(f).` The procedure takes as input a facet f of the current triangulation and performs the following :
 – compute the circumcenter c of f
 – if c encroaches some constrained edge e, `refine-edge(e)`
 – otherwise insert c as a vertex of of the triangulation.

`conditionally-refine-tetrahedron(t).` The procedure takes as input a tetrahedra t of the current triangulation and performs the following :
 – compute the circumcenter c of t
 – if c encroaches some constrained edge e, `refine-edge(e)`
 – else, if c encroaches some constrained facet f
     compute the projection c_F of c on the facet F of C including f
     find the subfacet g of F including c_F
     `conditionally-refine-facet(g))`
 – else, insert c as a vertex of the triangulation.

Now, the Delaunay refinement algorithm can be described as follows :

Initialization. Initialize the triangulation T with the Delaunay triangulation of the vertices in $\mathcal{C}$

Refinement. Apply one of the following refinement rules as long as one of them apply. There is a priority order on those rules, that is Rule R_i is applied only if rule R_j with $j < i$ cannot be applied.

 Rule R_1. If there is a constrained edge e that is encroached by some vertex in T
     call `refine-edge(e)`.

 Rule R_2. If there is a constrained facet f encroached by some vertex p,
     compute the projection p_F of p on the facet F of $\mathcal{C}$ including f,
     find the subfacet g of F including p_F,
     call `conditionally-refine-facet(g))`.

 Rule R_3. If there is a bad tetrahedra t, call `conditionally-refine-tetrahedron(t)`.

6.3 The Delaunay refinement theorem

If it ends up, the Delaunay refinement algorithm obviously yields a mesh whose elements respect the geometry of the domains and satisfy the size and shape criteria. The remaining of this section is devoted to the proof of the following theorem which gives conditions sufficient to ensure that the Delaunay refinement algorithm terminates.

Theorem 6.2 (The Delaunay refinement theorem) *The Delaunay refinement algorithm terminates if the following conditions are satisfied*
– All angles of the input PLC are greater than $90°$*.*
– The shape criteria is an upper bound b on the radius-edge ratio of the mesh tetrahedra and $b > 2$
– There is no size criteria or the size criteria is an upper bound on the circumradii of the mesh tetrahedra.

The angles of the input PLC, concerned by the $90°$ bound of the first hypothesis are :
– the dihedral angles formed by two facets of the PLC incident to the same edge,
– the edge angles formed by any two PLC edges incident to the same vertex
– the edge-facet angles formed by any pair made from a facet and an edge of the PLC sharing a vertex. The angle formed by the edge e and the facet f sharing a vertex v is defined as the minimum angle (e, vu) for any point u in f.
We give the proof of theorem 6.2 in the particular case where there is no sizing field. It can be easily modified to take into account a sizing field. As in the two dimensional case, the proof is based on a volume argument showing that the number of added vertices (Steiner vertices) is finite. A few definitions and lemma are necessary.

Lemma 6.3 (Steiner vertices lemma.) *Any Steiner vertex is inside or on the boundary of the domain.*

Proof. The claim is trivial for Steiner vertices added on the PLC edges. For Steiner vertices that are facets circumcenters or tetrahedra circumcenters the proof is analog to the proof of lemma 5.3, using the fact that when a facet circumcenter is inserted there is no encroached edge and when a tetrahedron circumcenter is inserted there are no encroached facet nor encroached edge. □

As in dimension 2, we define the *local feature size* for any point in the domain Ω to be meshed. We call *rejected vertex* a facet or tetrahedron circumcenter that has been considered for insertion and rejected for encroachement. We define the *insertion radius* and the *parent vertex* for any vertex in the mesh and any *rejected vertex*.

Definition 6.4 (Local feature size.) *Given a PLC* $\mathcal{C}$ *and a point p, the* local feature size *of* $\mathcal{C}$ *at point p,* $\mathrm{lfs}(p)$*, is the radius of the smallest disk centered at p and intersecting two disjoint faces of* $\mathcal{C}$*.*

Definition 6.5 (Insertion radius) *If a vertex v is inserted in the mesh, the insertion radius* $r(v)$ *is the lentgh of the shortest edge incident to v right after its insertion. If v is a rejected vertex, the insertion radius* $r(v)$ *is the lentgh of the shortest edge that would be incident to v right after its insertion if v was inserted in the mesh at the time it is rejected. If v is a PLC vertex is the insertion radius* $r(v)$ *is the distance from v to the vertex in* $\mathcal{C}$ *nearest to v.*

Definition 6.6 (Parent vertex) *The parent of a vertex v is defined as follows*

– *If v is a vertex of $\mathcal{C}$, it has no parent.*
– *If vertex v is the inserted or rejected circumcenter of a tetrahedra t, the parent p of v is one of the vertex of the shortest edge of t, the one that was inserted last.*
– *If vertex v is a vertex inserted in a constrained edge e, or on a constrained facet, its parent p is the encroaching vertex that is closest to v, p may be a vertex of the current mesh or a rejected circumcenter.*

Lemma 6.7 (The insertion radius lemma) *The insertion radius $r(v)$ of any vertex v of the mesh is such that*
– *$r(v) \geq \mathrm{lfs}(v)$ if v is a vertex of the input PLC or the circumcenter of a constrained face, (either an edge or a facet), encroached by a vertex of the triangulation.*
– *$r(v) \geq cr(p)$, where $r(p)$ is the insertion radius the parent p of v and $c = b$ if vertex v is a tetrahedron circumcenter and $c = \sqrt{2}$ if v is the circumcenter of a constrained face (either an edge or a facet) encroached by a rejected vertex.*

Proof. Let v be a vertex of the input PLC. The insertion radius of v is the distance fron v to the closest other vertex of the input PLC, hence at least $\mathrm{lfs}(v)$.

Let v be the center of a constrained edge or facet s encroached by a vertex of the triangulation. The parent p of v is the vertex encroaching s that is closest to v and the insertion radius of v is $\|pv\|$. Because of the angular hypothesis on the PLC, p necessary belongs to a constrained edge or facet s' disjoint from s and $\|pv\|$ is at least $\mathrm{lfs}(v)$.

Let v be the circumcenter of a tetrahedron t. The vertex v is considered for insertion when applying Rule 3 and the radius-edge ratio of t is more than b. Therefore the insertion radius of v which is the circumradius of t is more than $bl_{min}(t)$ where $l_{min}(t)$ is the length of the shortest edge of t. The parent p of v is the last inserted vertex of the shortest edge of t and its insertion radius $r(p)$ is less than $l_{min}(t)$. Thus $r(v) \geq br(p)$.

Let v be the midpoint of a constrained edge e encroached by the circumcenter p of facet or a tetrahedron. See Figure 6.5. If a and b are the endpoint of e, the insertion radius $r(p)$ of p is less than $\min(\|pa\|, \|pb\|)$ which is less than $\frac{\|e\|}{\sqrt{2}}$ because p is enclosed by the smallest circumcircle of e. Therefore $r(v) \leq \frac{r(p)}{\sqrt{2}}$.

At last, let v be the circumcenter of a constrained facet encroched by the circumcenter p of a tetrahedron. See Figure 6.5. The insertion radius $r(v)$ is the circumradius of f and the insertion radius $r(p)$ of p is less than $\|pa\|$ where a is the vertex of f closest to p. Let p_f be the projection of p on the hyperplan spanned by f. We have $\|pa\| = \sqrt{\|p_f a\|^2 + \|pp_F\|}$, $\|pp_F\| \leq r(v)$ and, because p_f is known to be in f $\|p_f a\| \leq r(v)$ and therefore $r(v) \leq \frac{r(p)}{\sqrt{2}}$. Figure 6.6 shows that we would get only $r(v) \leq \frac{r(p)}{2}$ if p_f was not restricted to lie in f. $\square$

We define the density $d(v)$ of a any vertex v, inserted in the mesh or rejected, as

$$d(v) = \frac{\mathrm{lfs}(v)}{r(v)}$$
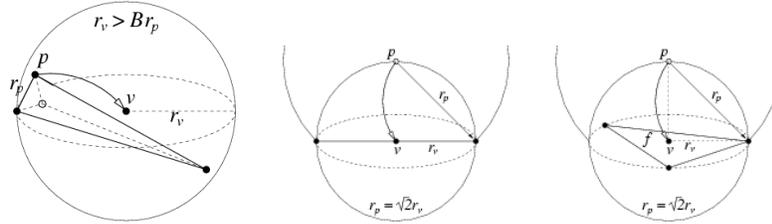
FIG. 6.5 – The insertion radii of a vertex v and its parent p. Left : v is the circumcenter of a tetrahedron. Middle : v is the midpoint of an edge encroached by a circumcenter. Left :v is the midpoint of a facet encroached by a tetrahedron circumcenter.
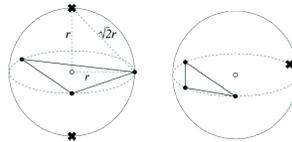
FIG. 6.6 – The insertion radius of the circumcenter v of a facet encroached by p compared to the insertion radius of $p : r(p)$ could be has big as $2r(v)$ of p is not restricted to project on f.
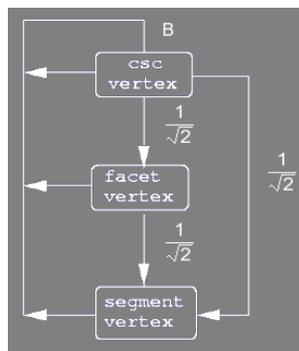
FIG. 6.7 – Flow diagram showing the insertion radii of Steiner vertices.

where lfs(v) is the local feature size of the input PLC at v and $r(v)$ is the insertion radius of v. The following lemma is identical to lemma 5.15.

Lemma 6.8 *For any vertex v, inserted in the mesh or rejected, if p is the parent of v and if the insertion radii of v of p are such that $r(v) \geq cr(p)$ for a given constant c, then $d(v) \leq 1 + \frac{d(p)}{c}$.*

Proof. See the proof of lemma 5.15. $\square$

Lemma 6.9 *If the hypothesis of theorem 6.2 are met and if there is no size criteria, there are constants $d_e \geq d_f \geq d_t \geq 1$ such that :*
– for any circumcenter v of a tetrahedron, inserted or rejected, $d(v) \leq d_t$
– for any circumcenter v of a constrained facet, inserted or rejected, $d(v) \leq d_f$
– for any any vertex v inserted in a constrained edge, $d(v) \leq d_e$.
Thus for any vertex in the mesh, $r(v) \geq \frac{\text{lfs}(v)}{d_e}$.

Proof. For any vertex v of the input PLC, $r(v) = \text{lfs}(v)$ and if $d_e \geq 1$ the result $r(v) \geq \frac{\text{lfs}(v)}{d_e}$ holds. For Steiner vertices, the proof is by induction. Assume that the lemma holds up to the insertion of vertex v.
– If v is a tetrahedron circumcenter, $r(v) \geq br(p)$ and from lemma 6.8, $d(v) \leq 1 + \frac{d(p)}{b}$. Thus lemma 6.9 holds for v if

$$1 + \frac{d_e}{b} \leq d_t. \tag{6.1}$$

– If v is the circumcenter of a facet f :
  – either the parent p of v is a vertex on a PLC face disjoint from the face including f, in which case $r(v) \geq \text{lfs}(v)$ and lemma 6.9 holds if $d_e \geq 1$,
  – or the parent p is a tetrahedron circumcenter, in which case we have $r(v) \geq \frac{r(p)}{\sqrt{2}}$ from lemma 6.7 and $d(v) \leq 1 + \sqrt{2}d_p$ from lemma 6.8. Then, lemma 6.9 holds for v if

$$1 + \sqrt{2}d_t \leq d_f. \tag{6.2}$$

– If v is a vertex inserted in an edge e of the PLC :
  – either the parent p of v is a vertex of the mesh on a constrained edge e' of the PLC disjoint from e, in which case $r(v) \geq \text{lfs}(v)$ and lemma 6.9 holds if $d_e \geq 1$,
  – or p is a rejected circumcenter of a facet or a tetrahedron, in which case we have $r(v) \geq \frac{r(p)}{\sqrt{2}}$ from lemma 6.7 and $d(v) \leq 1 + \sqrt{2}d(p)$ from lemma 6.8. Then, lemma 6.9 holds for v if

$$1 + \sqrt{2}d_f \leq d_e \tag{6.3}$$

The following choice ensures that equations 6.1, 6.2 and 6.3 are satisfied with $d_e \geq d_f \geq d_t \geq 1$:

$$
\begin{aligned}
d_e &= \left(3 + \sqrt{2}\right) \frac{b}{b-2} \\
d_f &= \frac{(1 + \sqrt{2}b) + \sqrt{2}}{b-2} \\
d_t &= \frac{b + 1 + \sqrt{2}}{b-2}
\end{aligned}
$$

□

Lemma 6.10 (Local lower bound on edge length) *If the hypothesis of the Delaunay refinement theorem (theorem 6.2) are met with no size criteria, any edge in the mesh incident to a vertex v has a length at least $\frac{\text{lfs}(v)}{d+1}$.*

Proof.   See the proof of 5.17. □

Proof. of the Delaunay refinement theorem (theorem 6.2) Once a lower bound proportionnal to lfs(v) on the length of the mesh edges incident to a vertex v is settled (lemma 6.10), an upper bound on the number of mesh vertices can be obtain from a computation of the integral $\int_\Omega \frac{dx}{\text{lfs}(x)^3}$, similarly to what is done in the 2D case (see the proof of theorem 5.13). □

6.4   Meshing domain with small angles

The angular condition stated in the Delaunay refinement theorem is highly restrictive. The main problem arising when this condition is not satisfied is that Steiner vertices inserted on a constrained facets or edges may encroached constrained facets or edges lying on non disjoint PLC faces, thus triggering the insertion of mesh vertices with diminishing insertion radii. Nevertheless it is possible to modify the Delaunay refinement algorithm so that it is guaranteed to terminate whatever may be the input PLC angles. Basically, the modified algorithm relaxes the guarantee concerning the tetrahedra shape and allows some bad tetrahedra to stay in the mesh in the neighboorhood of small input angles.

The modified algorithm can be described as follows :
– The constrained edge of the PLC forming small edge angles are grouped into clusters and the edges of a cluster are refined along concentric shells as in the 2 dimensional case. (See subsection **??**. This strategy forces constrained edges sharing a vertex to have the same lentgh and thus avoid mutual encroachement.
– When a constrained edge is encroached by a mesh vertex , it is always refined. Therefore the algorithm achieves the fact that constrained edges are refined into Gabriel edges.

dagger   blade

Fig. 6.8 – Nearly degenerate triangles

– When a constrained edge e is encroached by a facet or a tetrahedron circumcenter, the refinement is not always accepted. Let v be the midpoint of e. The algorithm determines the set of all constrained edges that will have to be refined if edge e is refined, and the refinement of e is accepted only if the smallest of the corresponding insertion radii is at least as large as the insertion radius of an ancestor of v. If the refinement of e is refused, the bad tetrahedron or the encroached facet will stay in the mesh and will never be further considered for refinement.
– If a constrained facet f is encroached and if the circumcenter v of f encroaches no constrained edge, the refinement of f is accepted only if the insertion radius of v is at least as large as the insertion radius of one of its ancester.

Not only some bad tetrahedra will not be refined mesh but also encroached constrained facets. Therefore it is not guaranteed that all constrained facets appear as facet in the Delaunay triangulation of the current set of vertices. To ensure the respect of the constraints, the algorithm maintains the constrained Delaunay triangulation of the PLC $\mathcal{C}'$ formed by the current set of vertices, the constrained edges and the facets of $\mathcal{C}$. Happily, because constrained edges are Gabriel edges, this constrained Delaunay triangulation exists.

6.5   Slivers

A nearly degenerate triangle has its three vertices closed to a single axis. According to the distribution of those vertices along the axis, the nearly degenerate triangle is of blade or dagger type, see Figure 6.8. In any case, the radius-edge rato of a nearly degenerate triangle is large and those triangle are eliminated by a Delaunay refinement algorithm.

The nearly degenerate tetrahedra have either their four vertices closed to a single line, in the case of thin tetrahedra, or closed to a single plane in the case of flat tetrahedra. See Figure 6.9. All nearly degenerate tetrahedra have a large radius-edge ratio except those that are of the sliver type. Roughly speaking, a sliver is a nearly flat tetrahedra obtained when the four vertices are closed to a circle and roughly equally spaced along that circle. A sliver has a radius-edge ratio of the order of $\frac{1}{\sqrt{2}}$ and therefore is not captured by the Delaunay refinement algorithm. Still, because they are nearly flat and have a too small volume, slivers are undesirable in most applications. Sliver elimination is one of the most outstanding problem in 3D mesh generation.

A sliver is a tetrahedra that has a reasonable radius edge ratio while its volume is small compared to the lentghs of its edges. In the following definition we characterize the slivers using two measures over the shape
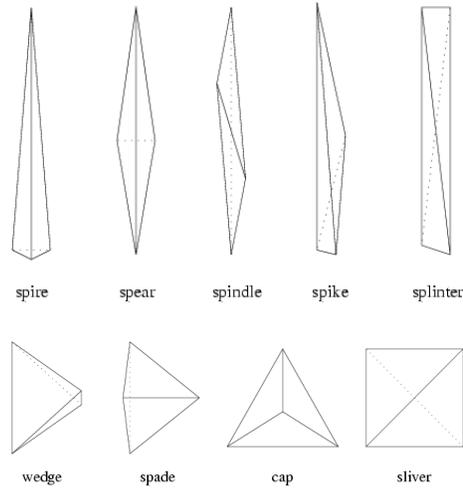
FIG. 6.9 – Nearly degenerate tetrahedra

of tetrahedra : the radius-edge ratio ρ and the volume-edge ratio $\sigma = \frac{V}{l^3}$ which is the ratio of the volume V of the tetrahedra to l^3 where l is the length of the shortest edge. A sliver are then defined by choosing bounds ρ_0 and σ_0 for this to measures.

Definition 6.11 (Sliver) *A sliver is a tetrahedra t that has a reasonable radius-edge ratio ρ and a volume-edge ratio σ such that $\rho \leq \rho_0$ and $\sigma =\leq \sigma_0$*

A fair measure for the quality of tetrahedra is a measure that tends to 0 for any type or nearly degenerate tetrahedra. Thus the radius-edge ratio is not a fair measure while the radius-radius ratio, id est the ratio of the circumradius to the radius of the inscribed sphere, is one.

In particular, one can show that any tetrahedra which has a bounded edge-ratio and is not a sliver has a bounded radius-radius ratio. Indeed, if $\rho \leq \rho_0$ and $\sigma \leq \sigma_0$, $\frac{r_{circ}}{r_{insc}} \leq \frac{\sqrt{3}\rho_0^3}{\sigma_0}$ where r_{circ} is the circumradius of t and r_{insc} the radius of its inscribed sphere.

In the next section, we show that bounded-edge ratio meshes, even if they include slivers, still have some nice properties that turns out to be usefull in the strategies of sliver elimination.

6.6   Delaunay meshes with bounded radius-edge ratio

Theorem 6.12 *Any Delaunay mesh with bounded radius-edge ratio tetrahedra is such that :*
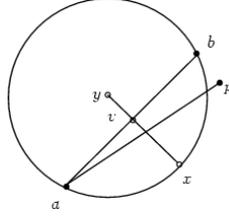
FIG. 6.10 – For the proof of lemma 6.13

– the ratio between the length of the longest edge incident to a mesh vertex and the length of the shortest one is bounded,
– the number of edges, facets and tetrahedra incident to a given vertex is bounded.

Lemma 6.13 *In a Delaunay mesh with bounded radius-edge ratio tetrahedra ($\rho \le \rho_0$ for any tetrahedra), two edges ab and ap incident to the same vertex a and forming an angle $(ab, ac) \le \eta_0 = \arctan\left[2\left(\rho_0 - \sqrt{\rho_0^2 - 1/4}\right)\right]$ are such that $\frac{\|ab\|}{2} \le \|ap\| \le 2\|ab\|$*

Proof. Let $\Sigma(y, r_y)$ be the circle that is the intersection of the plane spanned by ap and ab with the circumsphere of a tetrahedron incident to edge ab. The v be the projection of the center y on ab, and x the middle point of the arc ab. See Figure 6.10. We have,

$$\|xv\| = r_y - \sqrt{r_y^2 - \|ab\|^2/4}$$

$$\|xv\| \ge \left(\rho_0 - \sqrt{\rho_0^2 - 1/4}\right)\|ab\|$$

$$\widehat{(ab, ax)} \ge \arctan\left(\frac{2\|xv\|}{\|ab\|}\right) \ge \eta_0$$

Therefore, because $\widehat{(ab, ap)} \le \eta_0$, we have $\|ap\| \ge \|ax\| \ge \frac{\|ab\|}{2}$. $\square$

Proof. of theorem 6.12 Let ρ_0 be the bound on the radius-edge ratio of the tetrahedra, and let η_0, m_0 and ν_0 be defined as follows :

$$\eta_0 = \arctan\left[2\left(\rho_0 - \sqrt{\rho_0^2 - 1/4}\right)\right],$$

$$m_0 = \frac{2}{(1 - \cos(\eta_0/4))}$$

$$\nu_0 = 2^{2m_0 - 1}\rho_0^{m_0 - 1}.$$

We first prove that two mesh edges ab and ap incident to a are such that : $\frac{\|ab\|}{\nu_0} \leq \|ap\| \leq \nu_0\|ab\|$. Assume wlog that the shortest edge incident to a has unit length and let $\Sigma(a,1)$ be the sphere with unit radius centered on a. We consider a maximal packing on Σ of spherical caps with angle $\eta_0/4$. There are at most m_0 caps in this packing. Doubling the cap angles, we obtain a covering of Σ. Let G be the graph formed on Σ by the traces of mesh edges and mesh facets incident to a. The graph G is connected and there is a path in G from the trace of ab to the trace of ap. We then simplify the path by removing from it all the detour performed between two visits of the same cap. The simplified path visits at most m_0 caps and crosses at most m_0 boundary. From lemma 6.13 the length ratio is at multiplied by at most 2 within each cap. And when the path crosses a boundary the ratio is multiplied by at most $2\rho_0$. Indeed the path crosses a boundary between the trace of two edges incident to the same mesh facet. Any mesh facet is a triangle with radius-edge ratio at most ρ_0 and therefore a bound of at most $2\rho_0$ between the length of two of its edges.

We then prove that the number of edges incident to a given mesh vertex is at most $\delta_0 = (2\nu_0^2 + 1)^3$. Let ap be shortest edge incident to a and wlog let $\|ap\| = 1$. Let ab be the longest edge incident to a. From lemma 6.13, we have $\|ab\| \leq \nu_0$. For any vertex c adjacent to a, we have $1 \leq \|ac\| \leq \nu_0$, and for any vertex d adjacent to c, $\|cd\| \geq \frac{1}{\nu_0}$. Since the mesh is Delaunay, any vertex is adjacent to its nearest neighbor, and we may conclude that c is at distance at least $\frac{1}{\nu_0}$ from any other vertex in the mesh. Therefore the set of spheres $\Sigma_c(c, \frac{1}{2\nu_0})$, centered on the vertices adjacent to a, are disjoint and included in the sphere $\Sigma_a(a, \nu_0 + \frac{1}{2\nu_0})$. Thus the number of vertices adjacent to a is at most the ratio of the volume of sphere Σ_a to the volume of any sphere Σ_c. This ratio is $\delta_0 = (2\nu_0^2 + 1)^3$. The number of facets and tetrahedra incident to a mesh vertex are then respectively bounded by δ_0^2 and δ_0^3. $\square$

6.7   Sliver elimination

One of the method to eliminate the slivers from a mesh is to run as a post process a sliver elimination phase which is still a Delaunay refinement process but in which the choice of the refinement points is a bit relaxed. During this phase, encroached constrained edges, encroached constrained facets bad tetrahedra and slivers are refined by the insertion of a point which is not their circumcenter but a point carefully chosen in a small region surrounding this circumcenter. In the following we call *refinement region* the region in which the refinement point is chosen. The refinement region of a k-simplex τ with circumcenter c and circumradius r is a k-dimensional ball, included in the affine space spanned by τ, centered at c and with radius δr, where δ is a small constant to be precised later. Thus the refinement region of an encroached edge is a small subsegment, while the refinement region of facets and tetrahedra are small 2D and 3D balls respectively. See Figure 6.11

The main idea is that a carefull choice of the inserted vertex inside the refinement region avoid the creation of new slivers. In fact, it is not possible to avoid completely the creation of new slivers, but it is possible to avoid the creation of slivers with small circumradii. The remaining of this section proves these facts and show that this is enough to generate a sliver free mesh.
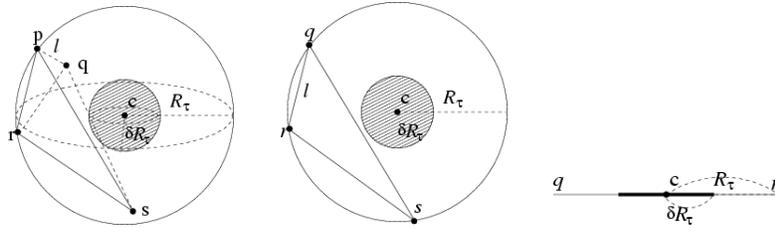
FIG. 6.11 – The refinement region of a tetrahedron(left), a facet(center) and an edge (right).

We first prove a few lemma on slivers.

Lemma 6.14 *Let p, q, r, s be the vertices of a tetrahedron t. with volume-edge ratio σ. Let r_y be the circumradius of the triangle qrs and d the distance from p to the hyperplane spanned by q, r, s. If $\sigma \le \sigma_0$, $d \le 12\sigma_0 r_y$*

Proof. Let V be the volume of the tetrahedra. From the definition of σ, $V = \sigma l^3$ where l is the shortest edge length. On the other hand, $V = \frac{1}{3}Sd$ where S, the area of triangle pqr, is at most $\frac{1}{2}l^2\frac{l}{2r_y}$. Therefore,

$$\sigma l^3 = V = \frac{1}{3}S\ d \ge \frac{1}{3}\left(\frac{1}{2}l^2\frac{l}{2r_y}\right)d = \frac{l^3}{12r_y}d.$$

□

Lemma 6.15 (Sliver lemma) *Let p, q, r, s be the vertices of a tetrahedron t and $\Sigma(y, r_y)$ be the circumcircle of triangle qrs. If the tetrahedron t is a sliver, the distance from p to the circumcircle $\Sigma(y, r_y)$ is at most $\gamma_2 r_y$ with $\gamma_2 = 48\sigma_0\rho_0$.*

For any triangle qrs, we called *forbidden region* the locus of points p such that the tetrahedron $pqrs$ is a sliver. Form the Sliver lemma, we know that the forbidden region is included in a torus of points around the circumcircle of triangle $pqrs$. The next lemma bounds the volume of the forbidden region, the area of its intersection with any hyperplane and the length of its intersection with any line.

Lemma 6.16 *For any base triangle qrs with circumradius r_y,*
– the volume of the forbidden region is at most $\gamma_3 r_y^3$ with $\gamma_3 = 2\pi^2\left(48\sigma_0\rho_0\right)^2$
– the area of the intersection of the forbidden region with any hyperplane is at most $\gamma_4 r_y^2$ with $\gamma_4 = 192\left(\pi\sigma_0\rho_0\right)$,
– the length the intersection of the forbidden region with any line is at most $\gamma_5 r_y$ with $\gamma_5 = 16\sqrt{3\sigma_0\rho_0}$.
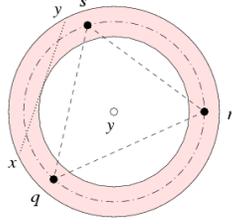
FIG. 6.12 – The torus including the forbidden region of triangle pqr.

Proof. The first point results from the computation of the volume of a torus which is the set of points at distance at most $48\sigma_0\rho_0$ from a circle with radius r_y. Figure 6.12 shows the hyperplane and the line that maximize the measure of their intersection with the torus. $\square$

The two next lemmas shows that, when refining an element τ of the mesh, there is only a finite number of forbidden region to take into account.

Lemma 6.17 *If the point p is to be chosen within the refinement region $(c_\tau, \delta r_\tau)$ of an element τ of the current mesh, there is at most a constant number of mesh facets qrs such that the tetrahedron $pqrs$ would be a new sliver with circumradius smaller than cr_τ for a given constant c.*

Proof. Let qrs be a mesh facet such that the tetrahedron $pqrs$ is sliver with circumradius smaller than cr_τ. Let $\mathcal{Q}$ be the set of vertices of mesh facets forming with vertex p a sliver with circumradius smaller than cr_τ.

If the circumradius of the tetrahedron $pqrs$ is less than cr_τ, $\|pq\| < 2cr_\tau$ and $\|qc\| < (2c + \delta)r_\tau$. Therefore any vertex in $\mathcal{Q}$ is enclosed in the sphere $\Sigma_1\left(c_\tau, (2C + \delta)r_\tau\right)$. centered at c_τ with radius $r_1 = (2c + \delta)r_\tau$.

Because p belongs to the refinement region $(c_\tau, \delta r_\tau)$ of a mesh element τ whose circumsphere encloses no vertex, $\|pq\| \geq (1 - \delta)r_\tau$, and the circumradius of $r(pqrs)$ of any tetrahedron $pqrs$ is such that $r(pqrs) \geq \frac{(1-\delta)r_\tau}{2}$. If the tetrahedra $pqrs$ is a sliver, its radius-edge ratio is less than ρ_0, and the edge qr of the mesh is such that $\|qr\| \geq \frac{r(pqrs)}{\rho(pqrs)} \geq \frac{(1-\delta)r_\tau}{2\rho_0}$. Because any tetrahedra in the mesh has a radius-edge ratio bounded by ρ_0 when a sliver is refined, from lemma 6.13, any edge in the mesh incident to q has length at least $\frac{(1-\delta)r_\tau}{2\rho_0\nu_0}$. Let us consider the set of spheres with radius $r_2 = \frac{(1-\delta)r_\tau}{4\rho_0\nu_0}$ centered at the vertices $\mathcal{Q}$. These spheres are disjoint. Therefore, the number of vertices in $\mathcal{Q}$ is at most

$$w = \left(\frac{r_1 + r_2}{r_2}\right)^3 = \left(\frac{4\rho_0\nu_0(2c + \delta) + (1 - \delta)}{1 - \delta}\right)^3,$$

and the number of mesh facets qrs likely to form a small sliver with p is at most w^3. $\square$

6.7.1 The sliver elimination phase.

We are now ready to describe the sliver elimination phase. This phase starts from a bounded radius-edge ratio mesh obtained through a standard Delaunay refinement algorithm. Then the sliver elimination phase applies in turn one of the following rule until none applies. Rule R_i is applied only if rule R_j with $j < i$ cannot be applied.

Rule R_1. If there is a constrained edge e that is encroached by some vertex in T
    call `sliver-free-refine-edge(e)`.

Rule R_2. If there is a constrained facet f encroached by some vertex p,
    compute the projection p_F of p on the facet F of C including f,
    find the subfacet g of F including p_F,
    call `sliver-free-conditionally-refine-facet(g))`.

Rule R_3. If there is a bad tetrahedra t call `sliver-free-conditionally-refine-tetrahedron(t)`.

Rule R_3. If there is a sliver t call `sliver-free-conditionally-refine-tetrahedron(t)`.

Thus the sliver elimination phase works as the usual Delaunay refinement phase except that it calls *sliver free* versions of the refinement procedures which avoid to create small slivers, i. e. slivers with circumradius less than cr where r is the circumradius of the tetrahedron being refined. Those sliver free refinement procedures do not use circumcenters as refinement point but choose a *sliver free* point in the refinement region. that is a point whose insertion does not create small slivers. For instance, the procedure `sliver--free-conditionally-refine-tet(t)` is given below.

`sliver-free-conditionally-refine-tetrahedron(`t`)`
– pick a sliver free point c in the refinement region
– if c encroaches some constrained edge e,
      sliver-free-refine-edge(e).
– else if c encroaches some constrained facet f,
      compute the projection p_F of p on the facet F of C including f
      find the subfacet g of F including p_F
      sliver-free-conditionally-refine-facet(f).
– else insert(c)
To choose a sliver free point in the refinement version, the algorithm picks a random a point of this region and check if the insertion of this point would create slivers with small circumradius. If this case occurs, the point is discarded and another random point is picked until a sliver free point is found. This is bound to occur because, for each refinement region, there is a finite number (w^3) of small slivers to avoid and from lemma 6.16, the constants σ_0, can be chosen so that the total volume of forbidden region is less than than the volume of the refinement region.

Theorem 6.18 *If the hypothesis of theorem 6.2 are met, and if the constants δ and c are chosen in such that $\frac{(1-\delta)^3 \rho_0}{2} \geq 1$ and $\frac{(1-\delta)^3 c}{4} \geq 1$, the sliver elimination phase terminates, yielding a sliver free bounded*

edge-ratio mesh, that is a mesh in which any tetrahedra t has a radius-edge ratio ρ and a volume edge ratio σ such that $\rho \leq \rho_0$ and $\sigma \leq \sigma_0$.

Proof. The two following lemma shows that if l_1 is the length of the shortest edge of the mesh before the sliver elimination phase, this phase will not create edges shorter than $l_2 = \frac{(1-\delta)^3 \, l_1}{4}$. $\square$

In the following, we call *original mesh* the bounded-edge ratio mesh before the sliver elimination phase. An *original sliver* is a sliver present in the original mesh and *secondary sliver* a sliver created during the sliver elimination phase.

Lemma 6.19 *Any point p whose insertion in the mesh is triggered by an original sliver has an insertion radius which is at least $\frac{(1-\delta)^3 \, l_1}{4}$.*

Proof. Let t be the sliver that triggers the insertion of p. The point p may be inserted in the refinement region of t, or in the refinement region of a constrained facet, or of a constrained edge.

In the first case, the insertion radius $r(p)$ of p is at least $(1 - \delta)r$ where r is the circumradius of t.

If p is inserted in the refinement region of a constrained facet f encroached by the refinement point q of t, the insertion radius $r(p)$ is at least $(1 - \delta)r(v)$, where $r(v)$ is the insertion radius of the circumcenter of the encroached facet. See Figure 6.13. Because q encroaches f, its insertion radius $r(q)$ is at most $\sqrt{2}r(v)$. Also, because q is in the refinement region of t, $r(q)$ is at least $(1 - \delta)r$. Therefore $\sqrt{2}r(v) \geq (1 - \delta)r$, and $r(p) \geq (1 - \delta)r(v) \geq \frac{(1-\delta)^2 r}{\sqrt{2}}$. The same bound holds if p is inserted in the refinement region of a constrained edge encroached by the refinement point of t.

At last, similarly, if p is inserted in the refinement region of a constrained edge encroached by the refinement point of a facet encroached by the refinement of t, $r(p) \geq \frac{(1-\delta)^3 r}{2}$. Therefore, in any case $r(p)$ is at least $\frac{(1-\delta)^3 r}{2}$ and the results follows because the circumradius r cannot be shorter than $\frac{l_1}{2}$. $\square$

Lemma 6.20 *If the constants δ and c are chosen so that $\frac{(1-\delta)^3 \rho_0}{2} \geq 1$ and $\frac{(1-\delta)^3 c}{4} \geq 1$ any vertex inserted during the sliver elimination phase has an insertion radius at least $l_2 = \frac{(1-\delta)^3 l_1}{4}$.*

Proof. The proof is by induction. Assume that the result holds until the insertion of p. Let t be the tetrahedron that triggers the insertion of p. From the previous lemma, the results holds if t is an original sliver.

Otherwise the tetrahedron t is either a bad tetrahedra or a sliver. The proof of the preceding lemma shows that the insertion radius $r(p)$ of p is at least $\frac{(1-\delta)^3 r}{2}$ where r is the circumradius of t. If t is a bad tetrahedra, $r \geq \rho_0 l_2$ and therefore $r(p) \geq \frac{(1-\delta)^3 \rho_0}{2} l_2$.
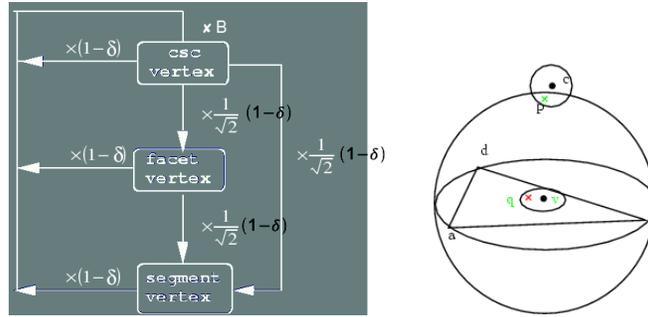
FIG. 6.13 – Flow diagram of the sliver elimination phase

If t is a secondary sliver the circumradius r of t is at least cr', where r' is the circumradius of the tetrahedra t' whose refinement creates t. By induction hypothesis, the length of t' edges is at least l_2 and r' is at least $\frac{l_2}{2}$ and $r(p) \geq \frac{(1-\delta)^3 c}{4} l_2$.

$\square$

Bibliographie

[1] O. Aichholzer, F. Hurtado, and M. Noy. On the number of triangulations every planar point set must have. In *13th Canadian Conf. on Computational Geometry*, pages 13–16, 2001.

[2] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

[3] Bernard Chazelle. Convex partitions of polyhedra : a lower bound and worst-case optimal algorithm. *SIAM J. Comput.*, 13 :488–507, 1984.

[4] Bernard Chazelle and L. Palios. Triangulating a non-convex polytope. *Discrete Comput. Geom.*, 5 :505–526, 1990.

[5] Francisco Santos and Raimund Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Comb. Theory, Ser. A*, 102(1) :186–193, 2003.

[6] Jonathan R. Shewchuk. A condition guaranteeing the existence of higher-dimensional constrained Delaunay triangulations. In *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, pages 76–85, 1998.