Constructing the segment Delaunay triangulation by a flip algorithm

Mathieu Brévilliers[a], Nicolas Chevallier[b], Dominique Schmitt[b,*]

[a]GIPSA-lab, Institut National Polytechnique, Grenoble, France
[b]Laboratoire LMIA, Université de Haute Alsace, Mulhouse, France

Abstract

Given a set S of line segments in the plane, we introduce a new family of partitions of the convex hull of S called segment triangulations of S. The set of faces of such a triangulation is a maximal set of disjoint triangles that cut S at, and only at, their vertices. A segment triangulation is Delaunay if its faces are inscribable in circles whose interiors do not intersect S. The main result of this paper is that any given segment triangulation can be transformed by a finite sequence of local improvements in a segment triangulation that has the same topological structure as the segment Delaunay triangulation. The main difference with the classical flip algorithm for point set triangulations is that some local improvements have to be performed on non convex regions. We overcome this difficulty by using locally convex functions.

Key words: Delaunay triangulation, Segment triangulation, Segment Voronoi diagram, Edge legality, Flip algorithm, Locally convex function

1. Introduction

Given a set S of points in the plane, a Delaunay triangulation of S is a triangulation of S whose triangles' circumcircles contain no point of S in their interiors. In 1977, Lawson [16] has shown that any triangulation of S can be transformed in a Delaunay triangulation by a sequence of local improvements: Every improvement consists in flipping a diagonal of a convex quadrilateral to the other diagonal. Since then, many extensions of flip algorithms have been proposed. In particular, they have been studied for different types of triangulations such as constrained triangulations [11], weighted triangulations [12], pseudo-triangulations [1], pre-triangulations [2], ... For a recent survey on flip algorithms, see [5].

In this paper we address the question of a flip algorithm to construct the segment Delaunay triangulation (or edge Delaunay triangulation). This triangulation has been introduced by Chew and Kedem [8] as the dual of the segment Voronoi diagram. Recall that, if S is a set of line segments in the plane, the segment Voronoi diagram of S is a partition of the plane whose regions are the points closer to one particular segment of S than to any others.

At first, we need to define a new family of diagrams, which we call segment triangulations, that contains the segment Delaunay triangulation. A segment triangulation of S is a partition of the convex hull of S whose set of faces is a maximal set of disjoint triangles such that the vertices of each triangle (and only its vertices) belong to three distinct sites of S (see Figure 1). The edges of the segment triangulation are the (possibly two-dimensional) connected components of the convex hull of S when the sites and open faces are removed. We show that these triangulations retain different geometrical and topological properties of point set triangulations and that they are intimately related to some generalized constrained triangulations.

The segment Delaunay triangulation is the one whose faces are inscribable in "empty" circles. In case of point set triangulations, the Delaunay triangulation can also be locally characterized: It is the only triangulation such that any two triangles sharing a common edge are Delaunay with respect to the four points defining the triangles [16]. We

*Corresponding author.
Email addresses: Mathieu.Brevilliers@gipsa-lab.inpg.fr (Mathieu Brévilliers), Nicolas.Chevallier@uha.fr (Nicolas Chevallier), Dominique.Schmitt@uha.fr (Dominique Schmitt)

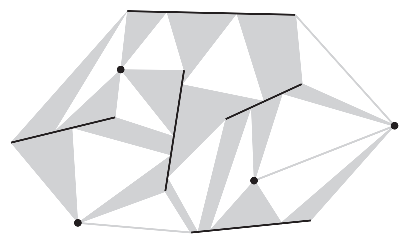

Figure 1: Example of segment triangulation. The sites of S are in black, the faces in white and the edges in grey.

show that the segment Delaunay triangulation can be characterized in the same way. We also give a local property that characterizes the set of segment triangulations having the same topological structure as the segment Delaunay triangulation.

An obstacle arises when we try to transform a segment triangulation into the segment Delaunay triangulation by a sequence of local improvements; some of these local transformations have to be performed on non convex regions. In order to characterize these transformations and to prove that the constructed triangulations tend to the segment Delaunay triangulation, we use a lifting on the three-dimensional paraboloid together with locally convex functions. The usefulness of locally convex functions in the context of flip algorithms has been already noticed by several authors (see [1], [3], ...). It is also worth noting that our algorithm is close to Perron's method for solving partial differential equations [19].

Another difficulty is that there are infinitely many segment triangulations of a given set, while the number of triangulations usually handled by flip algorithms is finite. So, a flip algorithm that aims to construct a segment Delaunay triangulation explicitly, might need infinitely many steps. Fortunately, this drawback can be circumvented by stopping the algorithm when it reaches a segment triangulation that has the same topological structure as the segment Delaunay triangulation. We shall show that such a triangulation is obtained in finitely many steps.

2. Segment triangulations

2.1. Definition and existence

Let S be a finite set of $n \geq 2$ disjoint closed segments in the plane, which we call sites. Throughout this paper, a closed segment may possibly be reduced to a single point. We say that a circle is tangent to a site s if s meets the circle but not its interior. The sites of S are supposed to be in general position, that is, we suppose that no three segment endpoints are collinear and that no circle is tangent to four sites (when a site is a point, we consider that this point is the only endpoint of the site). In the following, the word *triangle* will always denote a non degenerate triangle, that is a triangle whose vertices are not collinear. We will denote by $\mathbf{S}$ the union of the sites of S, that is the set of points of the sites of S. If U is a subset of $\mathbb{R}^2$, we will denote by $\overline{U}$ the closure of U, by U^o the interior of U, by $relint(U)$ the relative interior of U, and by ∂U the boundary of U.

Definition 1. *A segment triangulation $\mathcal{T}$ of S is a partition of the convex hull conv*$(\mathbf{S})$ *of* $\mathbf{S}$ *in disjoint sites, edges and faces such that:*
(i) Every face of $\mathcal{T}$ is an open triangle whose vertices are in three distinct sites of S and whose open sides do not intersect $\mathbf{S}$*,*
(ii) No face can be added without intersecting another one,
(iii) The edges of $\mathcal{T}$ are the (possibly two-dimensional) connected components of conv$(\mathbf{S}) \setminus (F \cup \mathbf{S})$*, where F is the union of the faces of $\mathcal{T}$.*

We will use indifferently the terms *triangle* and *face* to designate the faces of a segment triangulation.

Property 1. *Every set of sites admits a segment triangulation.*

Proof. To prove the existence of a segment triangulation of a set of sites S, it is enough to show that only a bounded number of disjoint open triangles can intersect $\mathbf{S}$ at, and only at, their vertices.

Obviously, if S contains only two sites, no such triangle exists. Otherwise, by considering the different kinds of convex hulls of two disjoint triangles in the plane, it is not difficult to see that, if the two triangles have their vertices (and only their vertices) in the same three sites, then the sites are encountered in two distinct orders when the triangles are traversed in counter clockwise direction. As a consequence, at most two disjoint triangles can have their vertices in the same three sites. □

2.2. Segment triangulations and weakly constrained triangulations

The segment triangulations defined here are not "real" triangulations since the convex hull is not decomposed in triangles. To this aim the two dimensional edges of a segment triangulation should be decomposed in triangles. A well-known triangulation defined on a set of points and line segments in the plane is the constrained triangulation. It is a triangulation of the set of points and line segment endpoints such that every given line segment is an edge of the triangulation. We show now that segment triangulations are intimately related to a kind of constrained triangulations with weaker constraints. This result will allow us to specify the shape of the edges of a segment triangulation and later on to give an algorithm to construct a particular segment triangulation.

Definition 2. *1. Given a set S of sites, we call S-polygon (possibly whith holes) any closed region A included in conv($\mathbf{S}$) and such that the boundary of A is composed of a finite number of disjoint segments that are of the two following forms:*
– on the one hand, closed segments included in $\mathbf{S}$ *(possibly reduced to points),*
– on the other hand, open segments that do not intersect $\mathbf{S}$ *and whose endpoints are in* $\mathbf{S}$.
2. An S-polygon A is said to be non degenerate if A is equal to the closure of its interior and if $A \setminus \mathbf{S}$ is connected.
3. We call weakly constrained triangulation of A (with respect to S), any partition of A into triangles whose vertices are in $\mathbf{S}$*, whose interiors do not intersect* $\mathbf{S}$*, and whose open sides either do not intersect* $\mathbf{S}$ *or are included in* $\mathbf{S}$.
When $A = $ conv($\mathbf{S}$), such a triangulation is also called a weakly constrained triangulation of S.

Note that a (classical) constrained triangulation is a particular case of weakly constrained triangulation.

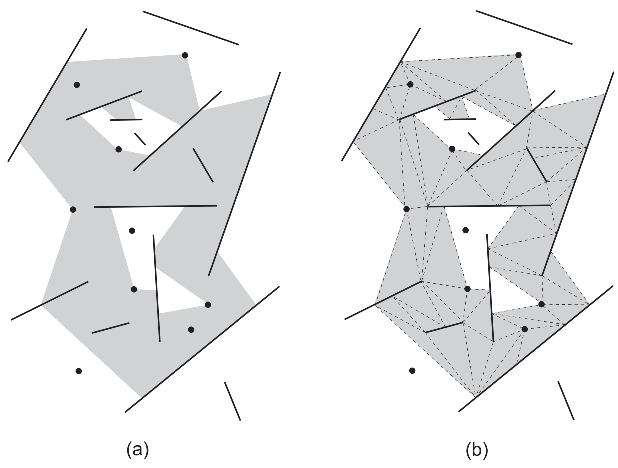

Figure 2: An S-polygon in gray (a) and a weakly constrained triangulation (dotted lines) of this S-polygon (b).

Lemma 1. *If A is a non degenerate S-polygon that intersects at least three sites of S, then any weakly constrained triangulation of A contains at least one triangle having its vertices in three distinct sites of S.*

Proof. Given a weakly constrained triangulation T of A, let $\Delta_T(A)$ be the (possibly empty) set of triangles of T having one side in $\mathbf{S}$. We show, by induction on the number $|\Delta_T(A)|$ of triangles of $\Delta_T(A)$, that T contains at least one triangle that has its vertices in three distinct sites of S.

Obviously, if $\Delta_T(A) = \emptyset$, every triangle of T has its vertices in three distinct sites of S. Suppose the result is true for any weakly constrained triangulation T of any non degenerate S-polygon A that intersects at least three sites of S and such that $|\Delta_T(A)| < k$, $k \geq 1$.

Let now T be a weakly constrained triangulation of such an S-polygon A but with $|\Delta_T(A)| = k$. Let t be a triangle of $\Delta_T(A)$. Since $A \setminus S$ is connected, at least one of the two sides of t that are not in $\mathbf{S}$ is not in the boundary of A. It follows that this side is also in the boundary of $\overline{A \setminus t}$. Hence, $\overline{A \setminus t}$ intersects the sites cut by t. Moreover, since A is equal to the closure of its interior, $\overline{A \setminus t} = \overline{A \setminus (t \cup \mathbf{S})}$ and it follows that $\overline{A \setminus (t \cup \mathbf{S})}$ cuts exactly the same sites as A. Now, the closure of every connected component of $A \setminus (t \cup \mathbf{S})$ cuts the two sites containing the vertices of t. Thus, the closure A' of at least one of these connected components cuts at least three sites.

By construction, A' is an S-polygon. Moreover, if T' is the restriction of T to A', then $|\Delta_{T'}(A')| < |\Delta_T(A)|$. Thus, by induction hypothesis, T' contains at least one triangle having its vertices in three distinct sites of S. It is the same for T. □

Theorem 1. *Every weakly constrained triangulation of S is a refinement of a segment triangulation of S, that is, a segment triangulation whose edges are decomposed in triangles.*

Proof. Obviously, every triangle of a weakly constrained triangulation T having its vertices in three distinct sites of S is a face of a segment triangulation of S. If F is the union of these open triangles, the closure of every connected component e of $conv(\mathbf{S}) \setminus (F \cup \mathbf{S})$ is an S-polygon. Since $\overline{e}$ contains no triangle of T having its vertices in three distinct sites, $\overline{e}$ intersects at most two sites, from Lemma 1. Hence, it is impossible to add in $\overline{e}$ a triangle having its vertices in three distinct sites. It follows that F is the union of the faces of a segment triangulation of S and that T is a refinement of this segment triangulation. □

Theorem 2. *(i) The closure of every edge of a segment triangulation $\mathcal{T}$ of S intersects exactly two sites of S.*
(ii) Every edge of $\mathcal{T}$ contains
– either two sides of two triangles of $\mathcal{T}$,
– or one side of one triangle of $\mathcal{T}$ and one side of $conv(\mathbf{S})$ that is not a site,
– or two such sides of $conv(\mathbf{S})$.

Proof. By definition, every edge e of $\mathcal{T}$ is a connected component of $conv(\mathbf{S}) \setminus (F \cup \mathbf{S})$ where F is the union of the faces of $\mathcal{T}$. As in proof of Theorem 1, $\overline{e}$ is an S-polygon. Actually, $\overline{e}$ is either a line segment or a non degenerate S-polygon.

In the first case, $\overline{e}$ connects two points of $\mathbf{S}$ and is either a common side of two triangles of $\mathcal{T}$ or a side of $conv(\mathbf{S})$ that is not a site.

In the second case, $\overline{e}$ is a "polygon" with at least three vertices and these vertices are in $\mathbf{S}$. Hence, $\overline{e}$ intersects at least two sites of S. But $\overline{e}$ cannot intersect more than two sites otherwise, by Lemma 1, a face of a segment triangulation could be placed in $\overline{e}$, contradicting the fact that $\mathcal{T}$ is already a segment triangulation. It follows that $\overline{e}$ is either a triangle with one vertex and its opposite side in $\mathbf{S}$ or a (possibly non convex) quadrilateral with two opposite sides in $\mathbf{S}$ (see Figure 3). In both cases, $\overline{e}$ admits two sides that are not in $\mathbf{S}$ and each of them is either a side of a triangle of $\mathcal{T}$ or a side of $conv(\mathbf{S})$ that is not a site. □

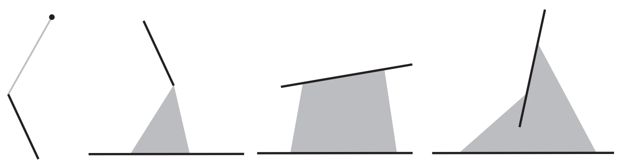

Figure 3: Examples of edges connecting two sites in a segment triangulation.

Theorem 2 shows that, as in point set triangulations, every edge of a segment triangulation "connects" two sites. This justifies the term "edge" for two dimensional regions of the segment triangulation.

2.3. Topological properties

Since every edge of a segment triangulation "connects" two sites, we can associate a combinatorial map with this triangulation in the following way:

Definition 3. *The combinatorial map associated with a segment triangulation $\mathcal{T}$ of S is such that:*

1. *The vertices of the map are the sites of S,*
2. *The arcs connecting two sites s and t in the map are the edges of $\mathcal{T}$ whose closures intersect s and t,*
3. *For every vertex s of the map, the cyclic ordering of the arcs out of s agrees with the counter-clockwise ordering of the associated edges around the site s in $\mathcal{T}$.*

Definition 4. *We say that two segment triangulations of S have the same topology if they have the same associated combinatorial map.*

Proposition 1. *The combinatorial map associated with a segment triangulation $\mathcal{T}$ of S is planar. Its faces match the faces of $\mathcal{T}$ and the outer face of $\mathcal{T}$ (that is, the complement of conv*$(\mathbf{S})$*).*

Proof. For every site s of S, let γ_s be a closed convex Jordan curve such that:
– s is inside γ_s (that is, in the region of the plane bounded by γ_s),
– for every site $s' \neq s$, $\gamma_{s'}$ is outside γ_s,
– the interior of γ_s only intersects the edges of $\mathcal{T}$ whose closures meet s.

Clearly, if an edge e intersects γ_s, $e \cap \gamma_s$ is connected (see Figure 4). Replace now every edge of $\mathcal{T}$ by one of the triangle or convex hull sides included int the edge. Then replace every site s by a point p_s inside γ_s and, for every reduced edge e that intersects γ_s, replace the part of e included in γ_s by a line segment connecting the points $e \cap \gamma_s$ and p_s. The order of the new edges around p_s is the same as the order of the corresponding initial edges around s in $\mathcal{T}$. Moreover, since the new edges do not intersect (except at their endpoints), we get a planar representation of the map whose faces correspond to the triangles of $\mathcal{T}$ and to the complement of *conv*$(\mathbf{S})$. □

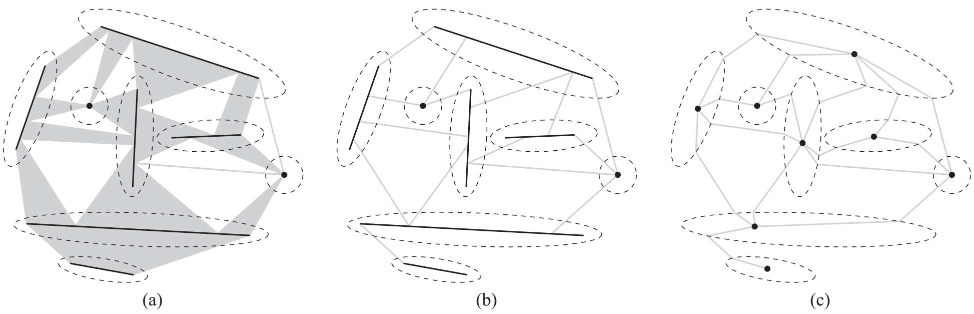

Figure 4: The sites are isolated inside convex curves (a), then the edges are replaced with line segments (b), and finally the sites are replaced with points (c).

This result allows to characterize the number of faces and edges in a segment triangulation:

Theorem 3. *Every segment triangulation $\mathcal{T}$ of a set S of n sites contains $3n - n' - 3$ edges and $2n - n' - 2$ faces, where n' is the number of sides of conv*$(\mathbf{S})$ *that are not sites.*

Proof. Let f be the number of faces of $\mathcal{T}$ and let e_0, e_1 and e_2 be the numbers of edges of $\mathcal{T}$ that contain respectively 0, 1 and 2 sides of faces of $\mathcal{T}$.

Since the sides of any face of $\mathcal{T}$ are contained in exactly three distinct edges of $\mathcal{T}$,

$$3f = 2e_2 + e_1. \tag{1}$$

Since the edges counted in e_1 and e_0 contain respectively 1 and 2 sides of *conv*(**S**),

$$n' = e_1 + 2e_0. \tag{2}$$

Since $\mathcal{T}$ has the same number of faces, of edges, and of vertices as its associated combinatorial map and since this map is planar, Euler's relation says that

$$f - (e_2 + e_1 + e_0) + n = 1. \tag{3}$$

The result follows from (1), (2), and (3). □

An interesting consequence of this theorem is that the size of a segment triangulation is linear with the number of sites. Moreover, it shows that the number of triangles of the triangulation is an invariant of the set of sites. This extends classical properties of point set triangulations.

2.4. Storage and construction

The combinatorial map associated to a segment triangulation $\mathcal{T}$ can be used as data structure to store the topology of $\mathcal{T}$. To maintain the geometrical informations of $\mathcal{T}$, it suffices to add the coordinates of the triangle vertices in the structure. These vertices can be associated with the edges of the map: one vertex per oriented edge. A segment triangulation of a set S of n sites can thus be stored using $O(n)$ space.

From Theorem 1, every constrained triangulation of S is a refinement of a segment triangulation of S. Now, a constrained triangulation T of S can be constructed in $O(n \log n)$ time [11]. A segment triangulation $\mathcal{T}$ can then be deduced from T in linear time by merging the triangles of T that are adjacent to the same two sites into a unique edge of $\mathcal{T}$ (see Figure 5).

Algorithms that construct the constrained triangulation T can also be adapted to directly construct the segment triangulation $\mathcal{T}$ in $O(n \log n)$ time [6].

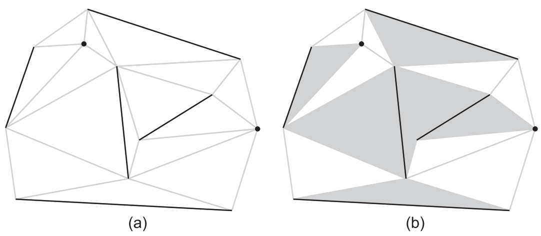

Figure 5: A constrained triangulation (a) and the corresponding segment triangulation (b).

2.5. Segment Delaunay triangulations

The classical Delaunay triangulation of a set of points can be easily extended to a set S of line segments in the following way:

Definition 5. *A segment triangulation of S is Delaunay if the circumcircle of each face does not contain any point of* **S** *in its interior.*

A way to show that every set S admits a segment Delaunay triangulation is to show that it is dual to the segment Voronoi diagram of S. Recall that the segment Voronoi diagram of S is a partition of the plane in regions. Each region contains the set of points closer to one site of S than to any other site. An edge of the segment Voronoi diagram is the set of points closest to and equidistant from two sites and it consists of line and parabola segments. When S is in general position, the vertices of its Voronoi diagram are the points of the plane closest to and equidistant from exactly three sites. Thus, each of these vertices is the center of a circle whose interior does not cut S and which circumscribes a triangle whose vertices are in three distinct sites of S. It is easy to see that these triangles are pairwise disjoint and, thus, that they are faces of a segment triangulation of S.

Moreover, if F is the union of these triangles, Chew and Kedem [8, 9] pointed out that the connected components of $conv(\mathbf{S}) \setminus (F \cup \mathbf{S})$ are dual to the edges of the segment Voronoi diagram of S and that each of these components is adjacent to exactly two sites of S. It follows that the triangles of F are all the faces of a segment triangulation of S and that this triangulation is dual to the segment Voronoi diagram of S. This extends a classical result of point set triangulations:

Theorem 4. *Every set S of sites in general position admits one and only one segment Delaunay triangulation. This triangulation is dual to the segment Voronoi diagram of S.*

The segment Delaunay triangulation of n sites can be constructed in $O(n \log n)$ time either by first building the segment Voronoi diagram [14, 4], or by adapting segment Voronoi diagram constructions [6].

3. Edge legality

3.1. Geometric legality of an edge

An interesting property of the Delaunay triangulation of a planar point set is the legal edge property. Consider an edge of a point set triangulation and its two adjacent triangles. The edge is illegal if a vertex of one of these triangles lies inside the circumcircle of the other triangle. It is well-known that the Delaunay triangulation of a point set is the unique triangulation of this point set without illegal edge. In the following, we are going to prove a similar property for segment triangulations.

Definition 6. *Let e be an edge of a segment triangulation $\mathcal{T}$ of S and let S' be the set of sites that contain the vertices of the zero, one, or two faces of $\mathcal{T}$ adjacent to e. The edge e is geometrically legal if:*
1. either e is not adjacent to any face of $\mathcal{T}$,
2. or the interior of the circumcircles of the faces adjacent to e do not cut $\mathbf{S}'$.

This directly extends the definition of edge legality in point set triangulations (see Figure 6).

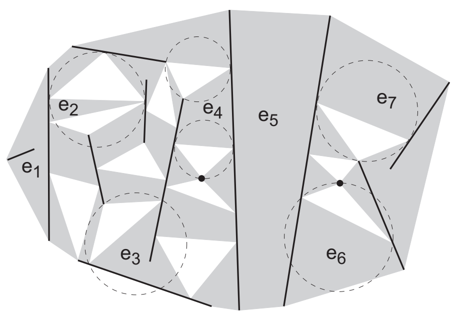

Figure 6: The edges e_1, e_4, e_5, and e_7 are geometrically legal whereas the edges e_2, e_3, and e_6 are not.

Theorem 5. *The segment Delaunay triangulation of S is the unique segment triangulation of S whose edges are all geometrically legal.*

Proof. By definition, the segment Delaunay triangulation of S has no geometrically illegal edge. Now, let $\mathcal{T}(S)$ be a segment triangulation of S that is not Delaunay and let us show that $\mathcal{T}(S)$ contains at least one geometrically illegal edge. Since $\mathcal{T}(S)$ is not Delaunay, it contains at least one face f such that the open disk $\mathcal{D}_f$ circumscribed to f cuts $\mathbf{S}$. Let x be a point in f and let p be a point in $\mathcal{D}_f \cap \mathbf{S}$. We can always choose p such that the open segment $]x, p[$ does not cut $\mathbf{S}$. Since p can not be a vertex of f (because it is in $\mathcal{D}_f$), $]x, p[$ cuts an edge e of $\mathcal{T}(S)$ adjacent to f. If $]x, p[$ cuts no other edge of $\mathcal{T}(S)$, either p is in one of the sites that cut the closure of e, or p is a vertex of the other face of $\mathcal{T}(S)$ adjacent to e. In both cases, the edge e is geometrically illegal (see Figure 7(a)).

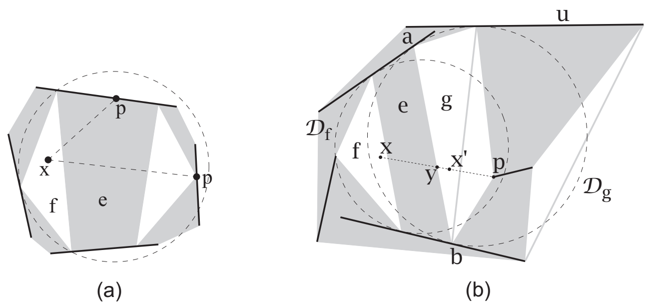

Figure 7: Illustrations of the proof of Theorem 5.

Now, we can state the following induction hypothesis: if $]x, p[$ cuts k edges, with $k \geq 1$, then at least one of them is geometrically illegal. Assume that $]x, p[$ cuts $k+1$ edges, with $k+1 \geq 2$, and let us show that $\mathcal{T}(S)$ contains at least one geometrically illegal edge. In this case, $]x, p[$ crosses e and also cuts the other face g adjacent to e. Let $\mathcal{D}_g$ be the open disk circumscribed to g and let $[a, b]$ be the side of ∂g included in e (see Figure 7(b)). If e is geometrically illegal, we are done. If e is geometrically legal, then $\mathcal{D}_f$ must be tangent to the sites that contain the vertices of f. Since two of these sites contain a and b, the points a and b can not be in $\mathcal{D}_f$. Since $]x, p[$ is included in $\mathcal{D}_f$ and crosses e, the point $y = [a, b] \cap]x, p[$ is in $\mathcal{D}_f$. It means that $[a, b]$ split $\mathcal{D}_f$ into two parts: one of them contains f and the other contains p. Let $\mathcal{D}_1$ be the part that contains f and $\mathcal{D}_2$ the part that contains p. Since the circumcircle of g passes through a and b, the disk $\mathcal{D}_g$ contains at least $\mathcal{D}_1$ or $\mathcal{D}_2$ (the two circles that bound $\mathcal{D}_f$ and $\mathcal{D}_g$ can not have four intersection points). If $\mathcal{D}_g$ contains $\mathcal{D}_1$, then it also contains f. Now, since e is geometrically legal, the vertices of f must lie on the boundary of $\mathcal{D}_g$, which means that $\mathcal{D}_g = \mathcal{D}_f$. As a consequence, $\mathcal{D}_g$ contains $\mathcal{D}_2$ and, thus, also contains p. Moreover, for any point x' of $[y, p] \cap g$, the segment $]x', p[$ cuts exactly k edges. From the induction hypothesis, we can conclude that $\mathcal{T}(S)$ contains at least one geometrically illegal edge. □

This theorem allows to know whether a given segment triangulation is Delaunay by checking the geometric legality of its edges. From Theorem 3, the number of edges in a segment triangulation of n sites is $O(n)$. Moreover, the legality test can be done in constant time since it is enough to compute the intersection of at most two circles and at most four segments. It means that we can know in linear time whether a given segment triangulation is Delaunay.

Due to precision errors, an algorithm could construct a segment triangulation that is almost Delaunay, in the sense that it has the same topology as the Delaunay one but the vertices of its faces are at distance $\varepsilon > 0$ from the vertices of the segment Delaunay triangulation faces. Clearly, geometric legality will not help to recognize such a triangulation. Furthermore, if a given segment triangulation has the same topology as the Delaunay one, then the latter can be computed by just moving the vertices of each face until its circumcircle becomes tangent to the three sites that contain the vertices of the face. As a consequence, an algorithm that constructs the segment Delaunay triangulation only needs to compute its topology. For this reason, we define in the following the concept of topologic legality of an edge, which will allow us to know whether a given segment triangulation has the same topology as the Delaunay one.

3.2. Topologic legality of an edge

In order to not consider the geometric position of the faces, we associate with each face a so called tangency triangle:

Definition 7. *Let f be a face of a segment triangulation of S. The tangency triangle of f is an open triangle such that*
– its vertices are on the same three sites as the vertices of f,
– its circumcircle is tangent to these three sites,
– these three sites are encountered in the same order when f and its tangency triangle are traversed in counter-clockwise direction.

The existence of the tangency triangle associated with a face f can be deduced from the existence of the Delaunay triangulation of the set of sites adjacent to f.

Definition 8. *Let e be an edge of a segment triangulation of S. The edge e is topologically legal in the two following cases:*

1. *e is adjacent to at most one face of the segment triangulation.*
2. *e is adjacent to two faces f_1 and f_2 and the following property holds. Denote t, r, u, v the sites such that t, r, u are incident to f_1 and r, t, v are incident to f_2 in counter-clockwise direction. Let $t_1r_1u_1$ and $r_2t_2v_2$ be the tangency triangles of f_1 and f_2 with $t_i \in t$, $r_i \in r$, $u_1 \in u$, and $v_2 \in v$. Then,*
   – the polygon $t_1t_2r_2r_1$ is either reduced to a segment or is a counter-clockwise oriented simple polygon (with three or four sides),
   – the circumcircles' interiors of $t_1r_1u_1$ and $r_2t_2v_2$ do not intersect the sites v and u respectively.

Case 2 of this definition can be stated in a more intuitive manner: The edge e is topologically legal if f_1 and f_2 are two faces of a segment triangulation of $\{r, t, u, v\}$ with the same topology as the segment Delaunay triangulation of $\{r, t, u, v\}$.

The two conditions of case 2 can be simultaneously false, as shown in Figure 8. Figure 9 gives an example where only the first condition is false. In the case where $t_1 = t_2$ and $r_1 = r_2$, the polygon $t_1t_2r_2r_1$ is reduced to a segment and only the second condition determines whether the edge is topologically legal or not.

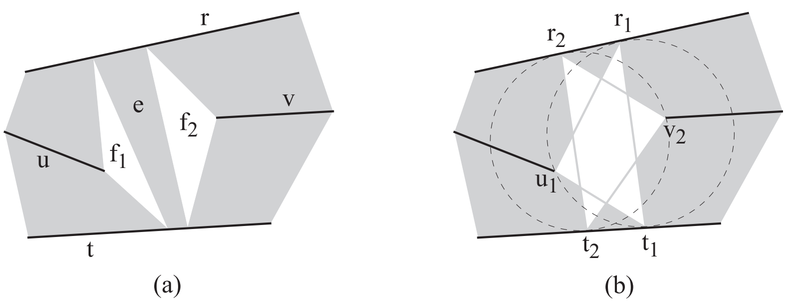

Figure 8: The edge e of the segment triangulation in (a) is topologically illegal because the tangency triangles of f_1 and f_2 (b) are such that neither $t_1t_2r_2r_1$ is counter-clockwise oriented nor their circumcircles are "empty".

In Theorem 6, we use the topologic legality of an edge in order to characterize the segment triangulations of S that have the same topology as the Delaunay one. At first, we give a preliminary result that states an important argument used in the proof of Theorem 6.

Lemma 2. *Let s_1, s_2, s_3 be three disjoint segments and let $a_i, a'_i \in s_i$, $i = 1, 2, 3$. Consider the oriented triangles $T = a_1a_2a_3$ and $T' = a'_1a'_2a'_3$.*
If $T \cap s_i = \{a_i\}$ and $T' \cap s_i = \{a'_i\}$, for $i = 1, 2, 3$, and if T and T' have the same orientation, then the curve γ formed by the segments $[a_2, a_3]$, $[a_3, a'_3]$, $[a'_3, a'_2]$, and $[a'_2, a_2]$ does not enclose any point of s_1.

Proof. We suppose that γ encloses a point O of s_1 and we show that it leads to a contradiction. We can distinguish four cases, which are illustrated in Figure 10. If we consider the case of Figure 10(a), since $O \in s_1$, since $T \cap s_1 = \{a_1\}$, and since $T' \cap s_1 = \{a'_1\}$, s_1 is entirely enclosed by the curve γ. It means that a'_1 is necessarily on the left of $\overrightarrow{a'_3a'_2}$ and that a_1 is on the left of $\overrightarrow{a_2a_3}$. As a result, T and T' do not have the same orientation, which contradicts the hypotheses. We can use the same argument in all other cases given in Figure 10. □

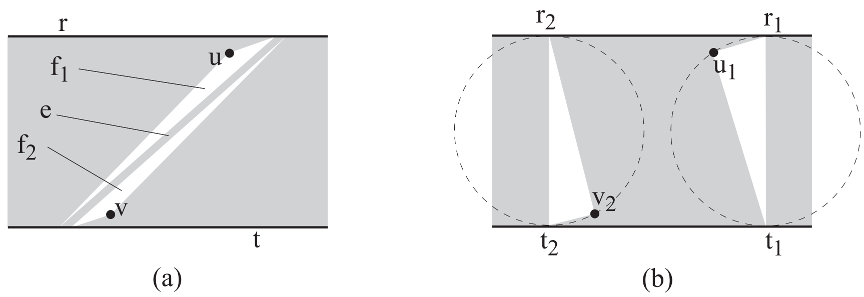

Figure 9: The edge e in (a) is topologically illegal even if the circumcircles of the tangency triangles of f_1 and f_2 are "empty" (b). Indeed, the polygon $t_1t_2r_2r_1$ is not counter-clockwise oriented.

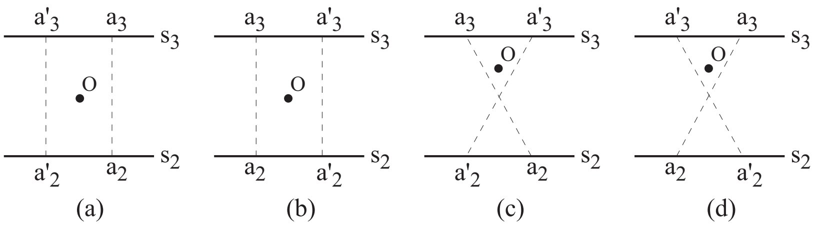

Figure 10: Illustration of the proof of Lemma 2.

Theorem 6. *A segment triangulation of S whose all edges are topologically legal has the same topology as the segment Delaunay triangulation of S.*

Proof. Let $\mathcal{T}(S)$ be a segment triangulation of S whose edges are all topologically legal. The Theorem will be proved if we can show that the tangency triangles of the faces of $\mathcal{T}(S)$ are the faces of a segment triangulation of S whose topology is the same as the topology of $\mathcal{T}(S)$. Indeed, from Theorem 5, this new triangulation is the segment Delaunay triangulation of S. The first goal is to prove that the tangency triangles are pairwise disjoint and do not cut any site of S.

The main idea of the proof is to use a result of Devillers et al. [10] which asserts that a representation of a combinatorial map by smooth curves in the plane is a planar graph if:
– All the circuits of the map are represented by simple closed curves,
– The ordering at each vertex s of the map is given by the geometric ordering of the curves emanating from the point representing s.
Actually, the result of Devillers et al. is stated with segments instead of smooth curves but an approximation argument leads to the same result for smooth curves. In the remaining of the proof, we will refer to this result as the planar representation lemma.

In the following, we define a planar geometric graph Γ from $\mathcal{T}(S)$ and we call C the map of which Γ is a geometric representation in the plane (Step 1). Then, we define another geometric representation Γ' of C with the tangency triangles of $\mathcal{T}(S)$ (Step 2). Finally, we prove that Γ' satisfies the conditions needed to use the planar representation lemma and we conclude that the tangency triangles are pairwise disjoint and do not cut any site of S (Step 3).

Step 1. At first, we construct a geometric graph Γ from the segment triangulation $\mathcal{T}(S)$.

Let ε be a strictly positive real number. For every site $s \in S$, let $\gamma_{s,\varepsilon}$ be the simple closed curve formed by the set of points at a distance ε from s. The curves $\gamma_{s,\varepsilon}$ are oriented in counter-clockwise direction. We choose ε small enough so that the curves $\gamma_{s,\varepsilon}$ are pairwise disjoint.

Let T be a triangle of $\mathcal{T}(S)$ whose vertices are in three sites s, t, and u in counter-clockwise direction. Let $p_{T,s}$, $p_{T,t}$, and $p_{T,u}$ be three points inside T such that $p_{T,s} \in \gamma_{s,\varepsilon}$, $p_{T,t} \in \gamma_{t,\varepsilon}$, and $p_{T,u} \in \gamma_{u,\varepsilon}$. It is easy to see that we can always choose the points $p_{T,s}$, $p_{T,t}$, and $p_{T,u}$ such that they can be joined with three disjoint curves $\gamma_{T,s,t}$, $\gamma_{T,t,u}$, and $\gamma_{T,u,s}$ that are inside T at a distance less than ε from the boundary of T and that meet $\gamma_{s,\varepsilon}$, $\gamma_{t,\varepsilon}$, and $\gamma_{u,\varepsilon}$ only at their endpoints. We do the same with the outer face of $\mathcal{T}(S)$, i.e., $\mathbb{R}^2 \setminus conv(\mathbf{S})$. In the following, $\gamma_{T,s,t}$ denotes the curve that goes from $p_{T,s}$ to $p_{T,t}$, and $\gamma_{T,t,s}$ denotes the curve that goes from $p_{T,t}$ to $p_{T,s}$, i.e., it is the same geometric curve, but with the reverse orientation.

Let s be a site and let $T_0, ..., T_{k-1}$ be the faces of $\mathcal{T}(S)$ incident to s in counter-clockwise direction around s (the outer face can be one of them). The curve $\gamma_{s,\varepsilon}$ is split into k disjoint simple curves $\gamma_{s,T_i,T_{i+1}}$ that go from the points $p_{T_i,s}$ to $p_{T_{i+1},s}$ ($i = i \bmod k$). As before, γ_{s,T_{i+1},T_i} and $\gamma_{s,T_i,T_{i+1}}$ are the same geometric curve, but with reverse orientations.

From the choice of ε, the geometric graph Γ is planar (see Figure 11(a)) and it defines a combinatorial map, which we call C.

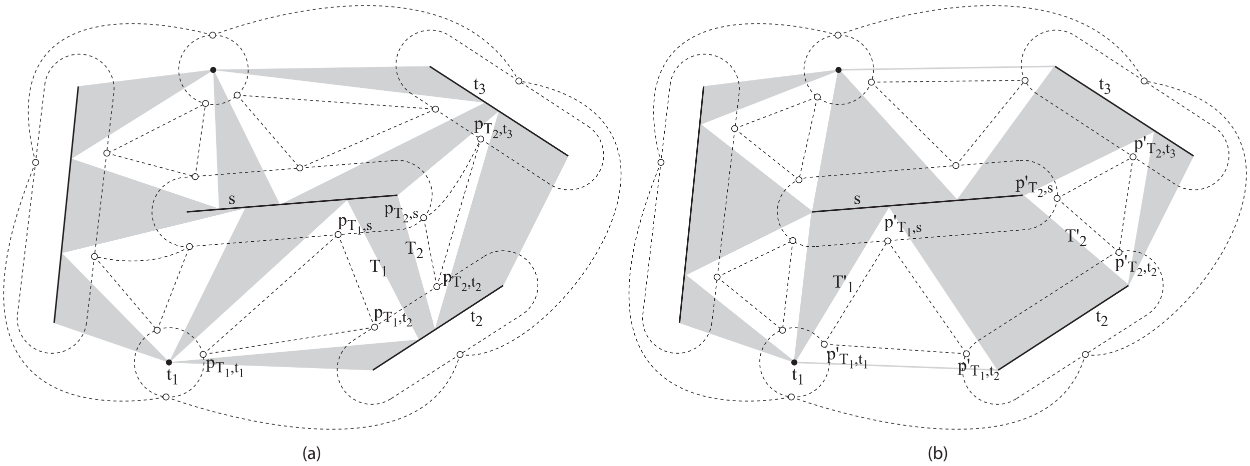

Figure 11: Illustration of Step 1 (a) and Step 2 (b) in the proof of Theorem 6.

Step 2. Now we construct another geometric representation Γ' of the map C. We use the same method as in Step 1, but with the tangency triangles of $\mathcal{T}(S)$ instead of its faces. For every triangle T of $\mathcal{T}(S)$, we denote by T' the corresponding tangency triangle.

In every tangency triangle T', we choose three points $p'_{T,s}$, $p'_{T,t}$, and $p'_{T,u}$ that are respectively on the curves $\gamma_{s,\varepsilon}$, $\gamma_{t,\varepsilon}$, and $\gamma_{u,\varepsilon}$. Consider now two triangles T_1 and T_2 that are consecutively incident to the same site s and their tangency triangles T'_1 and T'_2. Using the topologic legality, it easy to prove that T'_1 and T'_2 are disjoint. As in Step 1, the points $p'_{T,s}$, $p'_{T,t}$, and $p'_{T,u}$ can be joined with three disjoint simple curves $\gamma'_{T,s,t}$, $\gamma'_{T,t,u}$, and $\gamma'_{T,u,s}$ that are inside T' and at a distance less than ε from the boundary of T'.

In relation to Step 1, the only difference is the definition of the curves around the sites. Let s be a site and let $T_0, ..., T_{k-1}$ be the faces of $\mathcal{T}(S)$ incident to s in counter-clockwise direction around s. For every $i \in \{0, ..., k-1\}$, we consider the curve $\gamma'_{s,T_i,T_{i+1}}$ that goes from $p'_{T_i,s}$ to $p'_{T_{i+1},s}$ in counter-clockwise direction on $\gamma_{s,\varepsilon}$.

The set of curves $\gamma'_{T,s,t}$, $\gamma'_{s,T_i,T_{i+1}}$ defines a new geometric representation Γ' of the map C (see Figure 11(b)).

Step 3. The third step consists of proving that the geometric representation Γ' of C is planar. Indeed, if Γ' is planar for each $\varepsilon > 0$ small enough, then, letting ε go to 0, we see that the tangency triangles T' form a segment triangulation of S (note that the "outer" curves of Γ' are choosen at a distance less than or equal to 2ε from $conv(\mathcal{S})$).

We prove that Γ' is planar with the planar representation lemma. To this aim we prove that Γ' satisfies the two conditions of this lemma. Since each triangle T' has the same orientation as the corresponding triangle T in $\mathcal{T}(S)$ and since the curves $\gamma_{s,T_i,T_{i+1}}$ and $\gamma'_{s,T_i,T_{i+1}}$ are counter-clockwise oriented on $\gamma_{s,\varepsilon}$, the geometric order of the curves of Γ' around a vertex is the same as the geometric order of the curves of Γ around the corresponding vertex. Thus the second condition of the planar representation lemma holds. It remains to show that the geometric representations (in Γ') of the circuits of C are simple closed curves. The map C contains three types of circuits :

1. circuits with three arcs that correspond to the faces of $\mathcal{T}(S)$ and the outer circuit that corresponds to the boundary of $conv(\mathbf{S})$;
2. circuits with four arcs that correspond to the edges of $\mathcal{T}(S)$;
3. circuits around each site of $\mathcal{T}(S)$.

By construction, geometric representations of first-type circuits are simple curves. Geometric representations of second-type circuits are also simple curves because all edges of $\mathcal{T}(S)$ are topologically legal and all tangency triangles are pairwise disjoint. The result is not so obvious for third-type circuits.

Let s be a site and let $T_0, ..., T_{k-1}$ be the faces of $\mathcal{T}(S)$ that are incident to s in counter clockwise direction around s. For the sake of simplicity, we suppose that these faces are all internal faces. We begin by arbitrarily choosing a non zero vector $\overrightarrow{U}$ and an origin O in s. For every oriented curve $\gamma : [a, b] \to \mathbb{R}^2$ that does not contain O, we denote by $var(\gamma)$ the angle variation $\measuredangle(\overrightarrow{U}, \overrightarrow{O\gamma(t)})$ along γ. Since all the curves $\gamma'_{s,T_i,T_{i+1}}$ are oriented in counter clockwise direction, it is enough to show that:

$$var(\gamma'_{s,T_0,T_1}) + var(\gamma'_{s,T_1,T_2}) + ... + var(\gamma'_{s,T_{k-1},T_0}) = 2\pi. \tag{4}$$

For every $i \in \{0, ..., k-1\}$, let $t_0, ..., t_{k-1}$ be the different sites of s such that T_{i-1} and T_i are incident to t_i. The geometric representation α'_i of the circuit corresponding to a face T_i is a closed curve that begins at the point $p'_{T_i,s}$ and that is formed by the three curves γ'_{T_i,s,t_i}, $\gamma'_{T_i,t_i,t_{i+1}}$ and $\gamma'_{T_i,t_{i+1},s}$. It is obvious that the curves α'_i do not enclose the origin O, thus:

$$var(\alpha'_i) = var(\gamma'_{T_i,s,t_i}) + var(\gamma'_{T_i,t_i,t_{i+1}}) + var(\gamma'_{T_i,t_{i+1},s}) = 0.$$

Let β'_i be the geometric representation of the circuit corresponding to an edge incident to T_{i-1} and T_i. It is a closed curve that begins at the point $p'_{T_{i-1},s}$ and that is formed by the four curves γ'_{T_{i-1},s,t_i}, $\gamma'_{t_i,T_{i-1},T_i}$, $\gamma'_{T_i,t_i,s}$ and $\gamma'_{s,T_i,T_{i-1}}$. The orientation condition of the topologic legality implies that the curves β'_i do not enclose the origin O, thus :

$$var(\beta'_i) = var(\gamma'_{T_{i-1},s,t_i}) + var(\gamma'_{t_i,T_{i-1},T_i}) + var(\gamma'_{T_i,t_i,s}) + var(\gamma'_{s,T_i,T_{i-1}}) = 0.$$

Now, if we compute $var(\beta'_i) + var(\alpha'_i)$ for every $i \in \{0, ..., k-1\}$, then all the terms $var(\gamma'_{T_{i-1},s,t_i})$ and $var(\gamma'_{T_{i-1},t_i,s})$ cancel out. It is the same for the terms $var(\gamma'_{T_i,s,t_i})$ and $var(\gamma'_{T_i,t_i,s})$. Indeed, these terms are the angle variations along the same geometric curve, but in reverse orientations. We get:

$$\sum_{i=0}^{k-1}(var(\alpha'_i) + var(\beta'_i)) = \sum_{i=0}^{k-1}(var(\gamma'_{T_i,t_i,t_{i+1}}) + var(\gamma'_{t_i,T_{i-1},T_i}) + var(\gamma'_{s,T_i,T_{i-1}})) = 0.$$

We deduce that:

$$\sum_{i=0}^{k-1} var(\gamma'_{s,T_{i-1},T_i}) = \sum_{i=0}^{k-1}(var(\gamma'_{t_i,T_{i-1},T_i}) + var(\gamma'_{T_i,t_i,t_{i+1}})). \tag{5}$$

Since Γ is planar by construction, the same calculation with Γ gives rise to the following result:

$$\sum_{i=0}^{k-1} var(\gamma_{s,T_{i-1},T_i}) = \sum_{i=0}^{k-1}(var(\gamma_{t_i,T_{i-1},T_i}) + var(\gamma_{T_i,t_i,t_{i+1}})) = 2\pi. \tag{6}$$

The last thing to see is that the sums (5) and (6) are equal. To this aim, on each curve $\gamma_{t_i,\varepsilon}$, it is enough to choose a curve δ_i that joins $p_{t_i,T_{i-1}}$ to $p'_{t_i,T_{i-1}}$. Since the triangles T_i and T'_i have the same orientation, from Lemma 2, the successive curves

$$\gamma_{t_i,T_{i-1},T_i},\ \gamma_{T_i,t_i,t_{i+1}},\ \delta_{i+1},\ \gamma'_{T_i,t_{i+1},t_i},\ \gamma'_{t_i,T_i,T_{i-1}},\ -\delta_i$$

form a closed curve that do not enclose the origin O (see Figure 12). Summing over i the angle variations, we get:

$$\sum_{i=0}^{k-1}(var(\gamma_{t_i,T_{i-1},T_i}) + var(\gamma_{T_i,t_i,t_{i+1}}) + var(\delta_{i+1}) + var(\gamma'_{T_i,t_{i+1},t_i}) + var(\gamma'_{t_i,T_i,T_{i-1}}) + var(-\delta_i)) = 0$$

$$\sum_{i=0}^{k-1}(var(\gamma_{t_i,T_{i-1},T_i}) + var(\gamma_{T_i,t_i,t_{i+1}}) + var(\gamma'_{T_i,t_{i+1},t_i}) + var(\gamma'_{t_i,T_i,T_{i-1}})) = 0$$

$$\sum_{i=0}^{k-1}(var(\gamma_{t_i,T_{i-1},T_i}) + var(\gamma_{T_i,t_i,t_{i+1}})) + \sum_{i=0}^{k-1}(var(\gamma'_{T_i,t_{i+1},t_i}) + var(\gamma'_{t_i,T_i,T_{i-1}})) = 0$$

$$\sum_{i=0}^{k-1}(var(\gamma_{t_i,T_{i-1},T_i}) + var(\gamma_{T_i,t_i,t_{i+1}})) = \sum_{i=0}^{k-1}(var(\gamma'_{T_i,t_i,t_{i+1}}) + var(\gamma'_{t_i,T_{i-1},T_i})),$$

hence

$$2\pi = \sum_{i=0}^{k-1} var(\gamma_{s,T_{i-1},T_i}) = \sum_{i=0}^{k-1} var(\gamma'_{s,T_{i-1},T_i}).$$

It follows that the circuits of Γ' are simple curves. Thus, the tangency triangles are the faces of a segment triangulation of S. It remains to prove that this latter triangulation has the same topology as $\mathcal{T}(S)$. Since the adjacency relations of a triangle and of its tangency triangle are the same, the only thing to prove is that the cyclic ordering at each vertex is unchanged. But it is an easy consequence of (4). □

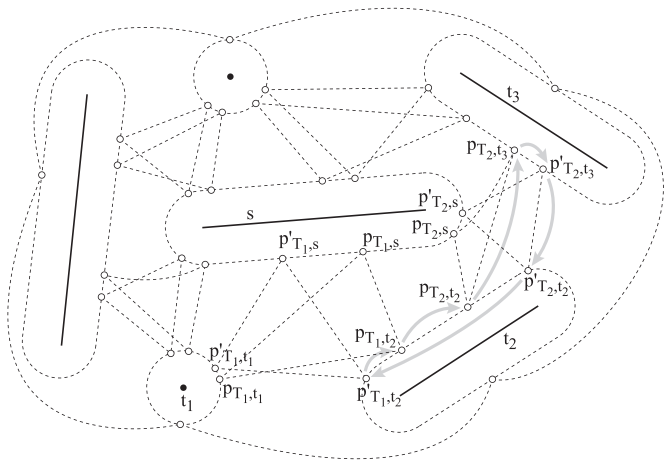

Figure 12: Illustration of Step 3 of the proof of Theorem 6.

Theorem 6 enables to test whether a segment triangulation has the topology of the segment Delaunay triangulation by checking the topologic legality of its edges. On the one hand, the topologic legality test of a given edge can be done in constant time since the only operations needed are computing the tangency points of a circle with three sites, checking the orientation of a polygon with (at most) four sides, and testing whether a line segment meets the interior of a circle. On the other hand, from Theorem 3, the number of edges in a segment triangulation is linear with the number of sites. Hence:

Corollary 1. *It can be checked in linear time whether a given segment triangulation has the same topology as the segment Delaunay triangulation.*

4. Flip algorithm

4.1. Segment triangulations of S-polygons

Each step of the flip algorithm that will be presented in Section 4.2 performs local modifications inside a subset of the current segment triangulation that is an S-polygon. So we first generalise the concept of segment triangulation to S-polygons (see Figures 13(a) and 13(b)).

Actually, when the intersection of an S-polygon and $\mathbf{S}$ is a finite point set, a segment triangulation of this polygon is nothing more than a classical triangulation of the polygon.

In the following, U always refers to an S-polygon, and we denote by S' the set of connected components of $\mathbf{S}' = U \cap \mathbf{S}$.

Definition 9. *A segment triangulation $\mathcal{T}$ of U (with respect to S) is a partition of U in disjoint sites, edges, and faces such that:*
(i) Every face of $\mathcal{T}$ is an open triangle whose vertices are in three distinct sites of S' and whose open sides do not intersect $\mathbf{S}'$,
(ii) No face can be added without intersecting another one,
(iii) The edges of $\mathcal{T}$ are the (possibly two-dimensional) connected components of $U \setminus (F \cup \mathbf{S}')$, where F is the union of the faces of $\mathcal{T}$.

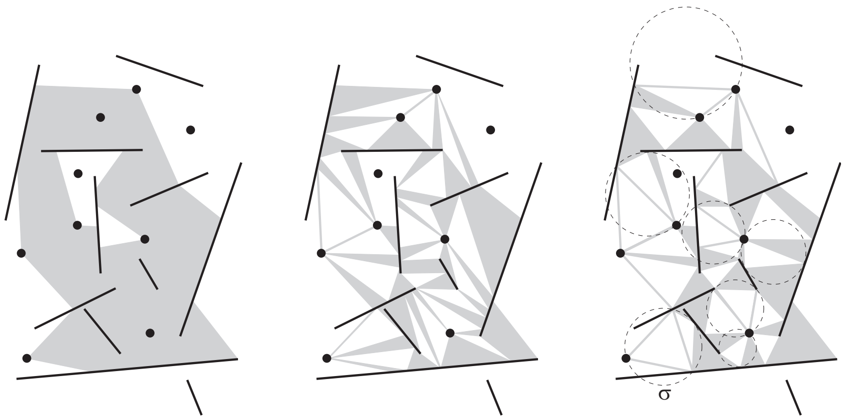

(a) (b) (c)

Figure 13: An S-polygon U (a), a segment triangulation of U (b), and a segment Delaunay triangulation of U (c). The circle σ is tangent to four connected components of $U \cap \mathbf{S}$.

Theorem 7. *The number of faces of a segment triangulation of U depends only on the couple (U, S).*

Proof. Let $\mathcal{T}$ be a segment triangulation of U. As in Definition 3, a combinatorial map M can be associated with $\mathcal{T}$. Moreover, using the same method as in proof of Proposition 1, M is planar. The faces of M match the faces of $\mathcal{T}$ together with the connected components of the complement of U in $\mathbb{R}^2$. Since M is planar, making use of Euler's relation, the result can be easily proved in the same way as Theorem 3. □

We can now define the segment Delaunay triangulation of an S-polygon U (see Figure 13(c)). In the following, we say that a point $q \in U$ is *visible* (relatively to U and S) from a point $p \in U$ if the open segment $]p, q[$ is included in $U \setminus \mathbf{S}$.

Definition 10. *1. Let t be a closed triangle included in U with vertices in $\mathbf{S}$. The interior of t is a Delaunay triangle of U (with respect to $\mathbf{S}$) if there exists a point p in the interior of t such that the interior of the circumcircle of t contains no point of $\mathbf{S}$ visible from p.*
2. A segment triangulation of U is Delaunay if all its faces are Delaunay triangles.

It should be noted that the circumcircle of a Delaunay triangle t contains no point of $\mathbf{S}$ visible from any point in t. This can be proved using Corollary 2.

We can also remark that the concept of Delaunay triangle used in this definition is very close to the one used in the classical definition of a constrained Delaunay triangulation [17, 7].

Theorem 8. *Every S-polygon admits a segment Delaunay triangulation.*

This result will be a consequence of Theorems 11 and 12 of Subsection 4.3. Note that a segment Delaunay triangulation of an S-polygon U is not necessarily unique since four connected components of $U \cap \mathbf{S}$ may be cocircular even if S is in general position (see Figure 13(c)).

4.2. Description of the flip algorithm

In this section, we give a flip algorithm that transforms any segment triangulation of S in a segment triangulation that has the same topology as the segment Delaunay triangulation of S.

The inputs of the flip algorithm are a segment triangulation of S and a queue that contains all edges of the triangulation.

One step of the algorithm goes as follows. The edge e at the head of the queue is popped. Let P_e be the closure of the union of e and of its at most two adjacent triangles: P_e is called the input polygon of e (see Figure 14 (b) and (f)). Clearly, P_e is an S-polygon and since it meets at most four sites, the Delaunay triangles of P_e can be computed in constant time. The triangulation of P_e is then replaced with a Delaunay triangulation of P_e. This gives rise to a new segment triangulation of S (it is a consequence of Theorems 7 and 8). Finally, the edge replacing e is pushed at the tail of the queue.

If this step changes the topology of the current segment triangulation, we say that the processed edge e has been *flipped*.

Beside the queue, the algorithm maintains the number of topologically illegal edges in the current triangulation. Notice that after the flip of an edge e, only the legality of the new edge and of the at most four edges adjacent to P_e has to be checked.

The algorithm ends when all edges are topologically legal. From Theorem 6, it means that the resulting segment triangulation of S has the same topology as the segment Delaunay triangulation of S.

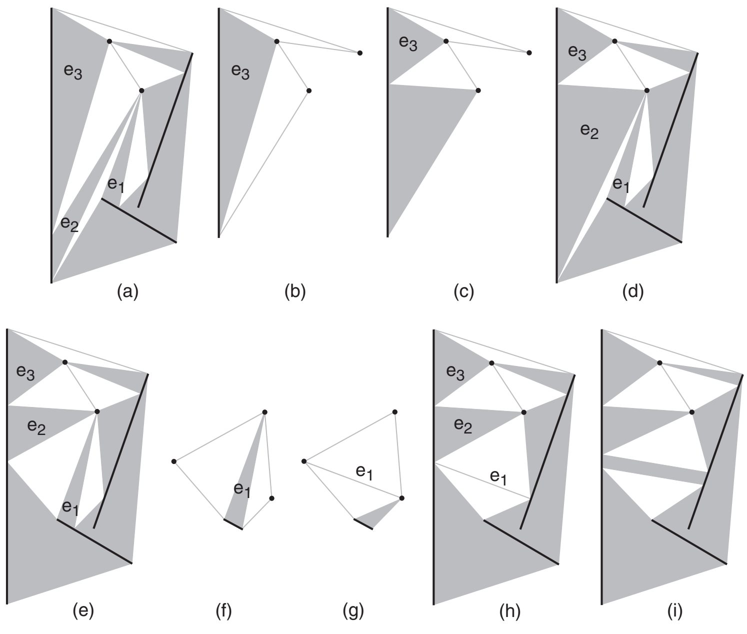

Figure 14: The flip algorithm transforms the given segment triangulation (a) in a segment triangulation (h) that has the same topology as the segment Delaunay triangulation (i).
The topology in (a) and the topology in (h) differ only by the flip of e_1, which is the only illegal edge of (a). However, the edge e_1 of (a) cannot be immediately flipped because its input polygon is not convex. So, the legal edges e_3 and e_2 have to be processed before e_1 becomes flippable.
In (b), the algorithm considers the input polygon P_{e_3} of the edge e_3. Then, in (c), it computes the segment Delaunay triangulation of P_{e_3} and this gives rise to a new segment triangulation in (d). In the same way, the processing of the edge e_2 leads to (e). Finally, the edge e_1 can be flipped (f, g), which leads to (h).

In case of point set triangulations, when an illegal edge is processed by the flip algorithm, it is flipped to a new legal edge, and the illegal edge will never reappear. Since there are finitely many possible edges, the flip algorithm reaches the Delaunay triangulation after a finite number of steps. Our flip algorithm looks very close to this classical flip algorithm, but we can not use the same idea to prove its convergence because of some important differences (see Figure 14):

– Even if an edge is not flipped, its geometry may change,
– some illegal edges cannot be flipped,
– a new constructed edge is not necessarily legal.

For point set triangulations, another way to prove the convergence of the flip algorithm to the Delaunay triangulation, is to lift the point set on the three-dimensional paraboloid $z = x^2 + y^2$. It is well known that the downward projection of the lower convex hull of the lifting is the Delaunay triangulation of the point set. Conversely, every other triangulation lifts to a non convex polyhedral surface above the lower convex hull. Now, it is enough to notice that an edge flip brings down the polyhedral surface.

We will use the same approach to prove that our flip algorithm always reaches a segment triangulation that has the same topology as the segment Delaunay triangulation. At first, for every S-polygon U, the lower convex hull of the lifting of $U \cap \mathbf{S}$ on the paraboloid is defined with the help of locally convex functions and we show that it projects down to the segment Delaunay triangulation of U (Theorem 12). Then, we define the lifting of any segment triangulation that is not Delaunay (Definition 12) and we show that the lifting of the segment Delaunay triangulation is lower than or equal to the lifting of any other segment triangulation (Theorem 13). In order to show the correctness of the algorithm, we prove that, after a step of the algorithm, the lifting of the resulting segment triangulation is lower than or equal to the lifting of the segment triangulation before this step (Theorem 14). This leads to prove that the sequence of steps builds a sequence of segment triangulations that converges to the segment Delaunay triangulation (Theorem 14). It remains to see that, after a finite number of steps, the segment triangulation constructed by the flip algorithm has the same topology as the segment Delaunay triangulation (Corollary 3).

4.3. Locally convex functions and segment triangulations

By using locally convex functions, we define and characterize the lifting on a paraboloid in $\mathbb{R}^3$ of a segment triangulation and of a segment Delaunay triangulation.

Recall that, if V is a subset of $\mathbb{R}^2$, a function $\phi : V \to \mathbb{R}$ is locally convex if the restriction of ϕ to each segment included in V is convex (see for example [3]). We define now the lower convex hull of a function, which we shall use instead of the usual lower convex hull of a subset in $\mathbb{R}^3$. Note that it corresponds to this usual lower convex hull when the domain V is convex.

Definition 11. *Let $L(V)$ be the set of functions $\phi : V \to \mathbb{R}$ that are locally convex on V. Given a real-valued function f defined on $V \cap \mathbf{S}$, the lower convex hull of f on $(V, \mathbf{S})$ is the function $f_{V,\mathbf{S}}$ defined on V by*

$$f_{V,\mathbf{S}}(x) = \sup\{\phi(x) : \phi \in L(V),\ \forall y \in V \cap \mathbf{S},\ \phi(y) \leq f(y)\}.$$

In the following, the above definition will be used on an S-polygon U with the function $f : \mathbb{R}^2 \to \mathbb{R}$ defined by $f(x, y) = x^2 + y^2$. The main aim of this subsection is to explain that the function $f_{U,\mathbf{S}}$ determines a segment Delaunay triangulation of U (see Figure 15). Next theorem gives information about the value of the function $f_{U,\mathbf{S}}$ at a point p. For every point p in U, we denote by vis_p the closure of the set of points of $\mathbf{S}$ visible from p, that is, the closure of the set of points q in $\mathbf{S}$ such that the open segment $]p, q[$ is included in $U \setminus \mathbf{S}$. The convex hull of vis_p is denoted by V_p.

For the sake of readability, intermediate lemmas and long proofs are postponed to Subsection 4.4.

Theorem 9. *Every point p of U belongs to a closed convex subset of U where the function $f_{U,\mathbf{S}}$ is affine and whose extreme points are one, two, or three points of $\mathbf{S}$. Moreover $f_{U,\mathbf{S}}(p) = f_{V_p,vis_p}(p)$.*

Corollary 2. *Let t be a triangle included in U with vertices in $\mathbf{S}$. $f_{U,\mathbf{S}}$ is affine on t if and only if t is a Delaunay triangle of U.*

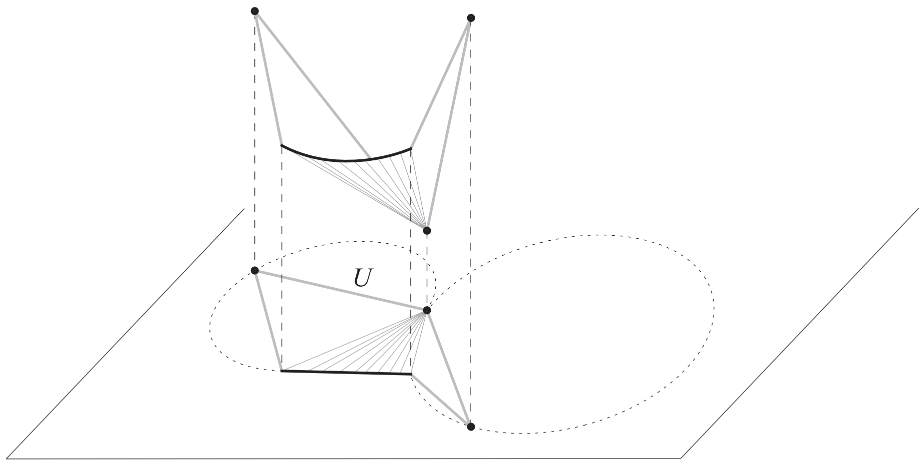

Figure 15: An S-polygon U and the graph of $f_{U,\mathbf{S}}$. U is decomposed into two triangles and infinitely many line segments where $f_{U,\mathbf{S}}$ is affine. The triangles are Delaunay triangles of U and the union of the segments forms the five edges of the segment Delaunay triangulation of U.

Proof. Suppose that $f_{U,\mathbf{S}}$ is affine on t. Let p be any point in the interior of t and $q \in S_p$. Denote $h : \mathbb{R}^2 \to \mathbb{R}$ the affine function equal to $f_{U,\mathbf{S}}$ on t. The function $f_{U,\mathbf{S}}$ is convex on $[p, q]$ and is equal to h on a neighborhood of p. Therefore $f_{U,\mathbf{S}} \geq h$ on $[p, q]$. Since $f_{U,\mathbf{S}} = f$ on $\mathbf{S}$, $f(q) = f_{U,\mathbf{S}}(q) \geq h(q)$. Hence q is not in the region of $\mathbb{R}^2$ where $f < h$, which is precisely the interior of the circumcircle of the triangle t.

Conversely, suppose that t is a Delaunay triangle. We begin by the case $U = conv(\mathbf{S})$. There exists a point p in the interior of t such that the interior of the circumcircle of t contains no point of $\mathbf{S}$ visible from p. Consider the affine function $h_t : \mathbb{R}^2 \to \mathbb{R}$ which is equal to f on the vertices of the triangle t. Since U is convex, the interior of the circumcircle contains no point of $\mathbf{S}$. Therefore $h_t \leq f$ on $\mathbf{S}$. It follows that $f_{U,\mathbf{S}} \geq h_t$ on the entire set U. On the other hand, $f_{U,\mathbf{S}} = f = h_t$ on the vertices of t. Thus, by convexity, $f_{U,\mathbf{S}} \leq h_t$ on t. It follows that $f_{U,\mathbf{S}} = h_t$ on t.

In the general case, if t is a Delaunay triangle of U then, by definition, it is also a Delaunay triangle of V_p with respect to vis_p. Hence, by the convex case, f_{V_p,vis_p} is affine on t. By the previous Theorem, we have $f_{U,\mathbf{S}}(p) = f_{V_p,vis_p}(p)$. Since $f_{U,\mathbf{S}}$ is locally convex, we have $f_{U,\mathbf{S}} \leq f_{V_p,vis_p}$ on t. Now, p is in the interior of t, therefore $f_{U,\mathbf{S}} = f_{V_p,vis_p}$ on t. □

The next step consists in showing that U can be partitioned into maximal convex subsets where the function $f_{U,\mathbf{S}}$ is affine (see Figure 15). Actually, when $U = conv(\mathbf{S})$, these subsets are the downward projections of the relative interiors of the lower faces of the convex hull of the lifting of $\mathbf{S}$ on the paraboloid $\{z = x^2 + y^2\}$. In the general case:

Theorem 10. *Let p be a point of U and denote by $\mathscr{C}_p$ the set of all relatively open convex subsets of U containing p where $f_{U,S}$ is affine.*
1. There is a maximal element C_p in $\mathscr{C}_p$, i.e., $C_p \in \mathscr{C}_p$ and, for every $C \in \mathscr{C}_p$, $C \subset C_p$.
2. $\overline{C_p}$ has finitely many extreme points, which are all in $\mathbf{S}$, and $f_{U,S}$ is affine on $\overline{C_p}$.
3. The collection of all sets C_p, $p \in U$, forms a partition of U.

Theorem 11. *Let T be a set of triangles that decompose all the two-dimensional C_p, $p \in U \setminus \mathbf{S}$, and whose vertices are extremal points of C_p. Then the interiors of the triangles of T are the faces of a segment triangulation of U.*

Such a triangulation is said to be *induced* by $f_{U,S}$. Then, from Corollary 2,

Theorem 12. *A segment triangulation of U is induced by $f_{U,\mathbf{S}}$ if and only if all its triangles are Delaunay.*

As a consequence, the graph of the function $f_{U,\mathbf{S}}$ is the lifting of a segment Delaunay triangulation on the paraboloid. We define now the lifting of any segment triangulation and show that it is above the lifting of a segment Delaunay triangulation.

Definition 12. *Let $\mathcal{T}$ be a segment triangulation of U. The function $f_{U,\mathbf{S},\mathcal{T}} : U \to \mathbb{R}$ is defined in the following way:*
– $f_{U,\mathbf{S},\mathcal{T}}(p) = f(p)$ if p is a point of $\mathbf{S}$,
– $f_{U,\mathbf{S},\mathcal{T}}(p) = f_{\bar{e},\mathbf{S}}(p)$ if p is in an edge e of $\mathcal{T}$,
– $f_{U,\mathbf{S},\mathcal{T}}(p) = f_{\bar{t},\mathbf{S}}(p)$ if p is in a face t of $\mathcal{T}$.

Theorem 13. *For every segment triangulation $\mathcal{T}$ of U, $f_{U,\mathbf{S}} \leq f_{U,\mathbf{S},\mathcal{T}}$.*

Proof. Let V be a closed subset of U. The inclusion $V \subset U$ implies both $V \cap \mathbf{S} \subset U \cap \mathbf{S}$ and $f_{U,\mathbf{S}}$ is locally convex on V. It follows that $f_{U,\mathbf{S}} \leq f$ on $V \cap \mathbf{S}$ and therefore $f_{U,\mathbf{S}} \leq f_{V,\mathbf{S}}$ on V. Using this last inequality with V the closure of a face or of an edge of $\mathcal{T}$, we get $f_{U,\mathbf{S}} \leq f_{U,\mathbf{S},\mathcal{T}}$. □

4.4. Proofs of the main results

4.4.1. Proof of Theorem 9

Lemma 3. *Let p, q, r be three points in U such that:*
– There exists a ball $B(p, \varepsilon)$ centered at p with radius $\varepsilon > 0$ such that $B(p, \varepsilon) \cap conv(\{p, q, r\})$ is included in U,
– The two open segments $]p, q[$ and $]p, r[$ are included in $U \backslash \mathbf{S}$,
– The interior of the triangle $t = conv(\{p, q, r\})$ contains no point u of $\mathbf{S}$ such that the segment $[p, u]$ is included in U.
Then, the triangle t is included in U.

Proof. Let E be the set of points x in $t^o \cap \partial U$ with $[p, x] \subset U$. The only thing to prove is that $E = \emptyset$. Suppose on the contrary that $E \neq \emptyset$. By the third hypothesis, $E \cap \mathbf{S} = \emptyset$. Furthermore, since $E \subset \partial U$, each point x in E must be in a boundary segment $]s_1, s_2[$ of U with s_1, s_2 in $\mathbf{S}$ and $]s_1, s_2[\cap \mathbf{S} = \emptyset$. Now, it is easy to see that a segment included in U cannot cross such a boundary segment; therefore $]s_1, s_2[$ crosses neither $]p, q[$ nor $]p, r[$. It follows that s_1 or s_2 is in t^o. Since there are only finitely many such boundary segments, there is an endpoint $s \in \mathbf{S} \cap t^o$ of such a segment where the affine function φ defined by $\varphi(p) = 0$ and $\varphi(q) = \varphi(r) = 1$ is minimal. By hypothesis, the segment $[p, s]$ is not included in U and therefore it must contain a point x in E. Finally, the inequalities $\varphi(s_i) \leq \varphi(x) < \varphi(s)$, which hold for one of the two points s_1 or s_2 associated with x, contradict the definition of s. □

Lemma 4. *Let p be in $U^o \backslash \mathbf{S}$ and let H be a closed half space such that $p \in \partial H$. Then there exists a point $q \in \mathbf{S} \backslash H$ such that $[p, q[$ is included in $U \backslash \mathbf{S}$.*

Proof. Moving on a half line from the point p in the open half plane $\mathbb{R}^2 \backslash H$, we can find a point $u \in \mathbf{S} \cup \partial U$ such that $u \in \mathbb{R}^2 \backslash H$ and $[p, u[\subset U \backslash \mathbf{S}$. If $u \in \mathbf{S}$, we are done. Otherwise, there exists a boundary segment $]s_1, s_2[$ containing u and such that $s_1, s_2 \in \mathbf{S}$ and $]s_1, s_2[\cap \mathbf{S} = \emptyset$. At least one of the two points s_1 and s_2 is in $\mathbb{R}^2 \backslash H$. Suppose it is $s = s_1$. Let v be the point of $\mathbf{S}$ in the triangle $t = conv(\{p, u, s\})$ such that the angle $\widehat{upv}$ is minimal. The triangle $t' = conv(\{p, u, v\})$ contains no point of $\mathbf{S}$ except on the line (p, v). If the segment $[p, v[$ is included in U, we are done. Otherwise there is a point r of ∂U in the segment $[p, v]$ such that $[p, r[$ is included in U^o. As before, if $[p, r]$ meets $\mathbf{S}$, we are done. Otherwise there exists a boundary segment $]r_1, r_2[$ with $r_1, r_2 \in \mathbf{S}$ and $r \in]r_1, r_2[$. One of the two points r_1, r_2 must be inside the triangle t', which contradicts the definition of v. □

Lemma 5. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a convex function, let U be a subset of $\mathbb{R}^n$, let A be a closed convex subset of $\mathbb{R}^n$, let $\phi : A \to \mathbb{R}$ be a convex function, and let W be a connected component of $A^o \cap U$. Assume that $\phi \geq f$ on A and that $\phi = f$ on ∂A. Then the function $g : U \to \mathbb{R}$ defined by*

$$g = \begin{cases} f \text{ on } U \backslash W \\ \phi \text{ on } W \end{cases}$$

is locally convex on U.

Proof. Let p and q be two points of U such that the segment $[p, q]$ is included in U. Let us show that the restriction $g_{|[p,q]}$ of g to $[p, q]$ is convex. If $[p, q]$ does not meet W, then $g_{|[p,q]} = f$ on $[p, q]$ and $g_{|[p,q]}$ is convex on $[p, q]$. Note that $I = [p, q] \cap W$ is an interval. Indeed, let r, s be two points of $W \cap [p, q]$. Since A^o is convex, $[r, s]$ is included in A^o. Moreover $[r, s] \subset [p, q] \subset U$, thus $[r, s] \subset U \cap A^o$. By definition of a connected component, $[r, s]$ is included in W.

Denote r and s the endpoints of I.
1. Suppose that $r, s \in \partial A$. By definition of the function, $g_{|[p,q]} = f$ on $[p, r]$, ϕ on $[r, s]$, and f on $[s, q]$. If $r = s$, then $g_{|[p,q]} = f$ on $[p, q]$ and $g_{|[p,q]}$ is convex on $[p, q]$. If $r \neq s$, on the one hand, the function ϕ is convex and $\phi \geq f$ on $[r, s]$, on the other hand, $f(r) = \phi(r)$, therefore the right derivative of $g_{|[p,q]}$ is non decreasing on $[p, s[$. It follows that $g_{|[p,q]}$ is convex on $[p, s[$. In the same way, using the left derivative, we show that $g_{|[p,q]}$ is convex on $]r, q]$ which shows that $g_{|[p,q]}$ is convex on $[p, q]$.
2. Suppose that $r, s \in A^o$. Let us show that $r = p$ and $s = q$. Suppose on the contrary that $[p, r[\neq \emptyset$. Since $r \in A^o$, there exists $r' \in [p, r[$ such that $[r', r] \subset A^o$. Now, $[r', r] \subset [p, q] \subset U$, therefore $[r, r'] \subset U \cap A^o$. It follows that r and r' are in the same connected component of $U \cap A^o$ and r cannot be an endpoint of I. By the same way of reasoning, we prove that $s = q$. Therefore $g_{|[p,q]} = \phi$ on $[p, q]$, and thus $g_{|[p,q]}$ is convex.
3. Suppose that $r \in \partial A$ and $s \in A^o$. In this case, we show as in second case that $s = q$ and the result follows as in the first case. □

Proof of Theorem 9.

Convex case. Assume that $U = conv(\mathbf{S})$. Consider the lifting $E = \{(x, f(x)) : x \in \mathbf{S}\}$ of $\mathbf{S}$ and the convex hull $\mathbf{K} = conv(E)$. The set of all lower points of $\mathbf{K}$ is the graph of a convex function $\varphi : U \to \mathbb{R}$. The convexity of f implies that $\varphi = f$ on $\mathbf{S}$, thus $\varphi \leq f_{U,\mathbf{S}}$. It is not difficult to prove that φ is affine on any downward projection G of a lower face F of $\mathbf{K}$. Moreover, since f is strictly convex, E is the set of extreme points of $\mathbf{K}$. It follows that the set $ext(G)$ of extreme points of G which are the projection of the extreme points of F, is included in $\mathbf{S}$. Now, we know that $f_{U,S} = \varphi$ on $ext(G)$, $f_{U,S}$ is convex and φ is affine on G, therefore $f_{U,\mathbf{S}} = \varphi$ on G. It follows that for any point $p \in U$, there is a closed convex subset G of U containing p such that $f_{U,\mathbf{S}}$ is affine on G and $ext(G) \subset \mathbf{S}$.
In the case of a point p in the interior of U, it is possible to strengthen the last conclusion: togetherwith the set G there exists an affine function $h : \mathbb{R}^2 \to \mathbb{R}$ equal to $f_{U,\mathbf{S}}$ on G and lower than or equal to f on $\mathbf{S}$. This can be easily seen using a supporting plane of $\mathbf{K}$ containing $(p, \varphi(p))$. We shall use this fact later in the proof.

General case. If p is in $\mathbf{S}$ or in ∂U, the proof of the Theorem is easy and we leave it to the reader. Let $p \in U^o \backslash \mathbf{S}$. Consider the convex set $V_p = conv(vis_p)$ instead of U. We proceed in five steps.

Step 1. We prove that p is in the interior of $V_p = conv(vis_p)$.
If $p \notin V_p^o$, there exists a supporting closed half plane H such that $p \in \partial H$ and $V_p \subset H$. By Lemma 4, there exists a point $q \in \mathbf{S}\backslash H$ such that $[p, q[\subset U\backslash\mathbf{S}$. Therefore $q \in vis_p$ and we have both $q \in vis_p \subset H$ and $q \in \mathbf{S}\backslash H$, which is impossible.

Step 2. There exists an affine function h_p on $\mathbb{R}^2$ and a subset S_p of vis_p such that $h_p \leq f$ on vis_p, $f = h_p$ on S_p and $p \in conv(S_p)$ (hence $f_{V_p,\mathbf{S}} = h_p$ on $conv(S_p)$).
The existence of S_p follows from the convex case of the theorem used with $V_p = conv(vis_p)$ instead of U and with vis_p instead of $\mathbf{S}$.

Step 3. We prove that $conv(S_p) \subset U$.
Denote by C the set of points $q \in conv(S_p)$ such that $[p, q] \subset U$. We would like to show that $C = conv(S_p)$. Since each point q in S_p is also in C, it is enough to prove that C is convex. Let q and q' be in C. If q, q' and p are on the same line then $[q, q'] \subset [p, q] \cup [p, q'] \subset U$, therefore $[q, q'] \in C$. It remains to study the case where q, q' and p are not on the same line. Since p is not in $\mathbf{S}$, p is not in S_p, hence by strict convexity of f, $f(p) < h_p(p)$. Again, the strict convexity of f implies that $(conv(\{p, q, q'\})\backslash\{q, q'\}$ is included in $\{x \in \mathbb{R}^2 : f(x) < h_p(x)\}$. Moreover $h_p \leq f$ on vis_p, thus $(conv(\{p, q, q'\})\backslash\{q, q'\}$ cannot contain a point $s \in \mathbf{S}$ such that $[p, s[\subset U\backslash\mathbf{S}$. Making use of Lemma 3, we get that $conv(\{p, q, q'\}) \subset U$; it follows that $[q, q'] \in C$ and $C = conv(S_p)$.

Step 4. We prove that $f_{U,\mathbf{S}} \leq h_p$ on $conv(S_p)$.
It simply comes from the convexity of $f_{U,\mathbf{S}}$ on $conv(S_p)$ which is included in U and from the equality $f_{U,\mathbf{S}} = f = h_p$ on S_p.

Step 5. We prove that $f_{U,\mathbf{S}} = h_p$ on $conv(S_p)$.
This is the main step of the proof and it needs Lemma 5. Let A be the convex subset of $\mathbb{R}^2$ where $f \leq h_p$. Since p is not in $\mathbf{S}$, p is not in S_p, and by strict convexity of f, we have $f(p) < h_p(p)$. It follows that p is in the interior A^o of A. Let W be the connected component of $U \cap A^o$ that contains p. By Lemma 5, the function $g : U \to \mathbb{R}$ defined by

$$g = \begin{cases} f \text{ on } U \backslash W \\ h_p \text{ on } W \end{cases}$$

is locally convex on U. If we can prove that $g \leq f$ on $\mathbf{S} \cap U$ then $f_{U,\mathbf{S}} \geq g$ and since $S_p \subset \{h_p = f\}$ and $conv(S_p) \backslash S_p \subset W$, $f_{U,\mathbf{S}} \geq g = h_p$ on $conv(S_p)$. We now prove that $g \leq f$ on $\mathbf{S} \cap U$.
It is enough to show that W does not contain any point of $\mathbf{S}$. Suppose on the contrary that there exists a point $q \in \mathbf{S} \cap W$. Since W is arcwise connected, there exists a path γ joining p to q in W. By continuity, the distance $\delta = d(\gamma, \partial A)$ is strictly positive.
Let $A_\delta = \{x \in A : d(x, \partial A) \geq \delta\}$. Obviously A_δ is closed and it is easy to see that it is convex. Let us show that $A_\delta \cap W$ is closed. Indeed, if $x \in \overline{A_\delta \cap W}$, then $x \in U$. It follows that there exists a ball $B(x, r)$ with $r > 0$ such that $U_x = B(x, r) \cap U$ is star shaped from x and thus connected. Therefore $W_x = U_x \cap A^o$ is also star shaped from x. Since $B(x, r) \cap W \cap A_\delta$ is not empty and is included in W_x, W_x contains at least one point of W. Therefore by definition of connected components, $W_x \subset W$ and $x \in W$ which implies that $A_\delta \cap W$ is closed.
It follows that there exists a shortest path γ_δ joining p to q in $A_\delta \cap W$. For all x in $U \backslash \mathbf{S}$, there exists $r > 0$ such that $U \cap B(x, r)$ is convex, therefore for all x in $A_\delta \cap W \backslash \mathbf{S}$, there exists $r > 0$ such that $A_\delta \cap W \cap B(x, r)$ is convex. It follows that the path γ_δ is straight on the parts where it does not meet $\mathbf{S}$. Let q' be the first point of $\mathbf{S}$ encountered by γ_δ. The segment $[p, q']$ is in $A_\delta \cap W$ and $[p, q'[$ does not contain any point of $\mathbf{S}$, it follows that $q' \in vis_p$. But $q' \in A_\delta \subset A^o$, thus by strict convexity of f, $f(q') < h_p(q')$ which contradicts the definition of h_p.

At last, by Caratheodory theorem, S_p can be chosen with two or three points. □

Note that, in case of $U = conv(\mathbf{S})$, the first step of the proof shows that $f_{U,\mathbf{S}}$ is lower semi continuous. Moreover, the main geometric fact about U, used in the proof of the Theorem, is that for all points of p in $U \backslash \mathbf{S}$, there exists a ball $B(p, r)$ such that $U \cap B(p, r)$ is convex. This is the reason why the Theorem may be hard to extend in higher dimensions.

4.4.2. Proof of Theorem 10

Lemma 6. *Let p_0, p_1, q_0, q_1 be four points of U such that:*
– The segments $[p_0, p_1]$ and $[q_0, q_1]$ are included in U,
– The intersection $]p_0, p_1[\cap]q_0, q_1[$ contains at least one point p,
– The function $f_{U,\mathbf{S}}$ is affine on both segments $[p_0, p_1]$ and $[q_0, q_1]$.
Then the quadrilateral $Q = conv(\{p_0, p_1, q_0, q_1\})$ is included in U and $f_{U,\mathbf{S}}$ is affine on Q.

Proof. If the four points p_0, p_1, q_0, and q_1 are on the same line, then the result is obvious. Otherwise, from the hypotheses, the point p is in the interior of Q. Let h_p be the affine function equal to $f_{U,\mathbf{S}}$ at p_0, p_1, and q_0. Since $f_{U,\mathbf{S}}$ and h_p are affine on $[p_0, p_1]$ and since they are equal at p and q_0, it means that they are equal on the entire segment $[q_0, q_1]$. Consequently, h_p and $f_{U,\mathbf{S}}$ are equal on $[p_0, p_1]$ and on $[q_0, q_1]$.

The application f is strictly convex, $f_{U,\mathbf{S}}(p_0) \geq f(p_0)$, $f_{U,\mathbf{S}}(p_1) \geq f(p_1)$ and $f_{U,\mathbf{S}}$ is affine on $[p_0, p_1]$, therefore $f_{U,\mathbf{S}}(p) > f(p)$ and $p \notin \mathbf{S}$.

Let us show that the interior ω of Q dose not contain a point of $\mathbf{S}$ visible from p. Let q be a point of ω such that $[p, q] \subset U$. Since $p \notin \mathbf{S}$, there exists a convex neighbourhood of p in U. Since $p \in]p_0, p_1[\cap]q_0, q_1[$ and since p_0, p_1, q_0 are not on the same line, the convex neighbourhood of p in U necessarily contains two points of $[p_0, p_1]$ and two points of $[q_0, q_1]$, which are the vertices of a quadrilateral that contains p in its interior and that is included in the convex neighbourhood of p in U. Thus, there exists $\varepsilon > 0$ such that the ball $B(p, \varepsilon)$ is included in U. Consequently, there exists a point $q' \in B(p, \varepsilon)$ such that $p \in]q', q[$ and such that q' is the barycenter of two points p_2, q_2 that are respectively in the segments $[p_0, p_1] \cap B(p, \varepsilon)$ and $[q_0, q_1] \cap B(p, \varepsilon)$. Since $q' \in [p_2, q_2] \subset U$, since $f_{U,\mathbf{S}}$ is convex on $[p_2, q_2]$, and since h_p is affine on $[p_2, q_2]$, we have $f_{U,\mathbf{S}}(q') \leq h_p(q')$. Moreover, the equality $f_{U,\mathbf{S}}(p) = h_p(p)$ and the convexity of $f_{U,\mathbf{S}}$ on $[q', q]$ implies that $f_{U,\mathbf{S}}(q) \geq h_p(q)$. Now, $q \in \omega$, $h_p \geq f$ on p_0, p_1, q_0, q_1, and f is strictly convex, thus $h_p(q) > f(q)$. As a consequence, $f_{U,\mathbf{S}}(q) > f(q)$ and $q \notin \mathbf{S}$.

Since ω contains no point of $\mathbf{S}$ visible from p, Lemma 3 shows that the four triangles $conv(\{p, p_0, p_1\})$, $conv(\{p, p_0, q_0\})$, $conv(\{p, p_1, q_0\})$ and $conv(\{p, q_0, q_1\})$ are included in U, therefore Q is included in U. The convexity of $f_{U,\mathbf{S}}$ on Q and the equality of $f_{U,\mathbf{S}}$ and h_p on $[p_0, p_1] \cup [q_0, q_1]$ show that $f_{U,\mathbf{S}} = h_p$ on Q. $\square$

Using standard arguments, one can show:

Lemma 7. *Let C_1 and C_2 be two relatively open convex subsets of $\mathbb{R}^n$. If the intersection $C_1 \cap C_2$ is not empty, then $conv(C_1 \cup C_2)$ is relatively open.*

Proof of Theorem 10.

Assume first that $p \notin \mathbf{S}$. Observe that $\mathscr{C}_p$ is non empty since $\{p\} \in \mathscr{C}_p$.

Step 1. Let us show that, if C_0 and C_1 are in $\mathscr{C}_p$, then $conv(C_0 \cup C_1)$ is in $\mathscr{C}_p$. We can assume that C_0 and C_1 are not reduce to $\{p\}$. From Lemma 7, we know that $conv(C_0 \cup C_1)$ is relatively open.
First, we prove that, for all p_0 in C_0 and all p_1 in C_1, $[p_0, p_1] \subset U$, and that $f_{U,\mathbf{S}}$ is affine on $[p_0, p_1]$. If p_0 or p_1 is equal to p, the result is obvious, thus we can assume that $p_0 \neq p$ and $p_1 \neq p$. Since C_0 and C_1 are relatively open, there exists $q_0 \in C_0$ and $q_1 \in C_1$ such that p is in $]p_0, q_0[$ and in $]p_1, q_1[$. By Lemma 6, the quadrilateral $Q = conv(\{p_0, p_1, q_0, q_1\})$ is included in U and $f_{U,\mathbf{S}}$ is affine on Q. Hence $[p_0, p_1] \subset U$ and $f_{U,\mathbf{S}}$ is affine on $[p_0, p_1]$. It remains to show that $f_{U,\mathbf{S}}$ is affine on each segment $[a, b] \subset conv(C_0 \cup C_1)$. The point a is in a segment $[a_0, a_1]$ and the point b is in a segment $[b_0, b_1]$ with a_0, b_0 in C_0 and a_1, b_1 in C_1. The segment $[a, b]$ is included in the quadrilateral $\mathcal{R} = conv(\{a_0, b_0, a_1, b_1\})$. We have just proved that all the segments $[x, y]$ with $x \in [a_0, b_0]$ and $y \in [a_1, b_1]$ are included in U and that $f_{U,\mathbf{S}}$ is affine on each of them. Furthermore, by definition, $f_{U,\mathbf{S}}$ is affine on $[a_0, b_0]$ and $[a_1, b_1]$. This implies that $f_{U,\mathbf{S}}$ is affine on the four triangles $conv(\{a_0, b_0, a_1\})$, $conv(\{a_0, b_0, b_1\})$, $conv(\{a_0, a_1, b_1\})$, and $conv(\{b_0, a_1, b_1\})$, and thus $f_{U,\mathbf{S}}$ is affine on $\mathcal{R}$ and $[a, b]$.

Step 2. Let us show that $C_p = conv(\cup_{C \in \mathscr{C}_p} C)$ is relatively open. By step 1, C_p is included in U and $f_{U,\mathbf{S}}$ is affine on C_p. So the only thing to show is that C_p is relatively open. If C_p is contained in a straight line, then $C_p = \cup_{C \in \mathscr{C}_p} C$ and it is clear that C_p is relatively open.
Suppose that C_p is two-dimensional. Fix C_0 and C_1 which are not included in the same straight line. For each point q in C_p, there exists at most three elements C_2, C_3 and C_4 in $\mathscr{C}_p$ such that $q \in conv(C_2 \cup C_3 \cup C_4)$. Therefore $q \in conv(\cup_{i=0}^4 C_i)$ which is open in $\mathbb{R}^2$ by Lemma 7.

Step 3. C_p contains no point of $\mathbf{S}$. Indeed, if q is in $C_p \cap \mathbf{S}$, then $f(q) = f_{U,\mathbf{S}}(q)$. The strict convexity of f implies that, on any segment $]r, s[$ containing q, f cannot be lower than or equal to an affine function h such that $h(q) = f(q)$. Since $f \leq f_{U,\mathbf{S}}$ and since $f_{U,\mathbf{S}}$ is affine on C_p, this leads to a contradiction.

Step 4. Let us show that the extreme points of $\overline{C_p}$ are all in $\mathbf{S}$. Let q be an extreme point of $\overline{C_p}$. The point q is the limit of a sequence $(q_n)_{n\geq 0}$ of points in C_p. By Theorem 9, for each n, there exists a_n, b_n and c_n in $\mathbf{S}$ such that $t_n = conv(\{a_n, b_n, c_n\})$ is included in U, $q_n \in t_n$, and $f_{U,\mathbf{S}}$ is affine on t_n. By step 3, q_n is neither equal to a_n, to b_n nor to c_n. Hence, we can suppose that q_n is in $relint(t_n)$, by removing, if necessary one of the points a_n, b_n or c_n. Making use of step 1 with q_n instead of p, we see that $C'_p = conv(C_p \cup relint(t_n))$ is included in U and that $f_{U,\mathbf{S}}$ is affine on C'_p. Hence C'_p is in $\mathscr{C}_p$and therefore $relint(t_n)$ is included in C_p and t_n is included in $\overline{C_p}$.
Now, Since the points a_n, b_n and c_n are in $\mathbf{S} \cap \overline{C_p}$, extracting a subsequence, we can assume that the sequences $(a_n)_{n\geq 0}$, $(b_n)_{n\geq 0}$ and $(c_n)_{n\geq 0}$ converge to a, b and c in $\mathbf{S} \cap \overline{C_p}$. Extracting once more a subsequence, we can suppose that the sequences of barycentric coefficients of q_n converge. It follows that q is a barycenter of three points a, b and c, which belong to $\mathbf{S} \cap \overline{C_p}$. Since q is an extreme point of $\overline{C_p}$, we have $q = a = b = c$, which shows that q is in $\mathbf{S}$.

Step 5. Let $h : \mathbb{R}^2 \to \mathbb{R}$ be an affine function such that $h = f_{U,\mathbf{S}}$ on C_p. Let us show that $f_{U,\mathbf{S}} = h$ on $\overline{C_p}$.
Let q be in $\overline{C_p}$ and let $(q_n)_{n\geq 0}$ be a sequence of point in C_p converging to q. As in step 4, there exists three converging sequences $(a_n)_{n\geq 0}, (b_n)_{n\geq 0}$ and $(c_n)_{n\geq 0}$ in $\mathbf{S}$ such that for all n, $t_n = conv(\{a_n, b_n, c_n\})$ is included in U, $q_n = \alpha_n a_n + \beta_n b_n + \gamma_n c_n \in t_n$, and $f_{U,\mathbf{S}}$ is affine on t_n. As in step 4, we can also suppose that the sequences $(\alpha_n)_{n\geq 0}, (\beta_n)_{n\geq 0}$ and $(\gamma_n)_{n\geq 0}$ converge to α, β and γ. Thus $q = \alpha a + \beta b + \gamma c$. By convexity,

$$\begin{aligned} f_{U,\mathbf{S}}(q) &\leq \alpha f_{U,\mathbf{S}}(a) + \beta f_{U,\mathbf{S}}(b) + \gamma f_{U,\mathbf{S}}(c) \\ &= \alpha f(a) + \beta f(b) + \gamma f(c). \end{aligned}$$

Since $f_{U,\mathbf{S}}$ is affine on t_n, $f_{U,\mathbf{S}}(q_n) = \alpha_n f(a_n) + \beta_n f(b_n) + \gamma_n f(c_n)$, hence

$$\lim_{n\to\infty} f_{U,\mathbf{S}}(q_n) = \alpha f(a) + \beta f(\beta) + \gamma f(c).$$

Since h is continuous, $h(q) = \lim_{n\to\infty} h(q_n) = \alpha f(a) + \beta f(b) + \gamma f(c) \geq f_{U,\mathbf{S}}(q)$. Moreover, $f_{U,\mathbf{S}}$ is convex on $\overline{C_p}$ and $f_{U,\mathbf{S}} = h$ on C_p, hence $f_{U,S} \geq h$ on $\overline{C_p}$ and $f_{U,\mathbf{S}}(q) \geq h(q)$.

Step 6. Since f is strictly convex, each segment $s \in S$ can contain one point of $\overline{C_p}$ at most. Thus, by step 4, there are finitely many extreme points.
To finish the proof, we have to consider the case $p \in \mathbf{S}$. Making use of the strict convexity of f, we show as in step 4, that the set $\mathscr{C}_p$ contains only one element : $\{p\}$.
At last, we have $C_x = C_p$ for all $x \in C_p$, which shows that, for all p and q in U, we have either $C_p = C_q$ or $C_p \cap C_q = \emptyset$. □

4.4.3. Proof of Theorem 11

Lemma 8. *$f_{U,\mathbf{S}}$ is continuous on $U \cap \mathbf{S}$.*

Proof. Let p be in $U \cap \mathbf{S}$. Since $f_{U,\mathbf{S}}(p) = f(p)$ and since $f \leq f_{U,\mathbf{S}}$, it is enough to prove that there exists $r > 0$ and a continuous function g defined on the neighborhood $V = U \cap B(p,r)$ such that $g(p) = f(p)$ and $f_{U,\mathbf{S}} \leq g$ on V. Since U is an S-polygon, we can find $r > 0$ such that $B(p,r) \cap U$ is an union of radius of the ball $B(p,r)$. The function f is bounded from above on U by a constant M. By Theorem 9, the function $f_{U,\mathbf{S}}$ is also bounded from above by M on U. Let g be the function defined on $\mathbb{R}^2$ by

$$g(x) = f_{U,\mathbf{S}}(p) + \frac{M - f_{U,\mathbf{S}}(p)}{r} \|x - p\|.$$

The three following properties hold:
– g is affine on each half line whose endpoint is p,
– $g = M \geq f_{U,\mathbf{S}}$ on the boundary of $B(p,r)$,
– $g(p) = f_{U,\mathbf{S}}(p)$.
Since $f_{U,\mathbf{S}}$ is locally convex, $f_{U,\mathbf{S}} \leq g$ on each radius of $B(p,r)$ included in U. Hence $g \geq f_{U,\mathbf{S}}$ on $B(p,r) \cap U$. □

Lemma 9. *For all $\varepsilon > 0$, there exists $\delta > 0$ such that for all $p' \in \mathbf{S}$ and all segment $[q,r]$ such that :*
– $d(p',[q,r]) \leq \delta$,
– $q, r \in \mathbf{S}$, $]q,r[\subset U \backslash \mathbf{S}$,
– $f_{U,\mathbf{S}}$ is affine on $[q,r]$,
we have $d(p',q)$ or $d(p',r) \leq \varepsilon$.

Proof. Suppose on the contrary that there exists $\varepsilon > 0$, a sequence $(p_n)_{n\in\mathbb{N}}$ of points of $U\backslash\mathbf{S}$, a sequence $(p'_n)_{n\in\mathbb{N}}$ of points of $\mathbf{S}$ and a sequence of segments $[q_n, r_n]$ such that:
– $d(p_n, p'_n) \to 0$ when $n \to \infty$,
$\forall n \in \mathbb{N}$,
– $q_n, r_n \in \mathbf{S}$, $]q_n, r_n[\subset U\backslash\mathbf{S}$, $p_n \in]q_n, r_n[$,
– $f_{U,\mathbf{S}}$ is affine on $[q_n, r_n]$,
– $d(p'_n, q_n)$ and $d(p'_n, r_n) \geq \varepsilon$.
We can assume that $d(p_n, p'_n) \leq \frac{\varepsilon}{2}$. For all $n \in \mathbb{N}$, $p_n = (1-\lambda_n)q_n + \lambda_n r_n$ where λ_n is in $[0,1]$. Moreover, for all n in $\mathbb{N}$, $d(p_n, q_n)$ and $d(p_n, r_n) \geq \frac{\varepsilon}{2}$. Now, $p_n - q_n = \lambda_n(r_n - q_n)$, thus

$$\frac{\varepsilon}{2} \leq d(p_n, q_n) = \lambda_n d(r_n, q_n) \leq \lambda_n \operatorname{diam} U$$

and $\lambda_n \geq \frac{\varepsilon}{2\operatorname{diam} U}$. In the same way, $1-\lambda_n \geq \frac{\varepsilon}{2\operatorname{diam} U}$. Extracting subsequences, we can suppose that the sequences $(p_n)_{n\in\mathbb{N}}$, $(p'_n)_{n\in\mathbb{N}}$, $(q_n)_{n\in\mathbb{N}}$ and $(r_n)_{n\in\mathbb{N}}$ converge to the points p, p', q and r. Since the sequence $d(p_n, p'_n)$ goes to 0, we have $p = p' \in \mathbf{S}$. By assumption, the function f is strictly convex and continuous, thus

$$\alpha = \inf\{(1-\lambda)f(x) + tf(y) - f((1-\lambda)x + ty) : x, y \in U,\ d(x,y) \geq \varepsilon,\ \lambda, 1-\lambda \geq \frac{\varepsilon}{2\operatorname{diam} U}\}$$

is strictly positive. Since $p \in \mathbf{S}$, by the previous lemma, the function $f_{U,\mathbf{S}}$ is continuous at p. At the same time, f is continuous at p, therefore there exists $\beta > 0$ such that $\forall x \in U,\ d(x,p) \leq \beta \Longrightarrow \left|f_{U,\mathbf{S}}(x) - f_{U,\mathbf{S}}(p)\right| \leq \frac{\alpha}{4}$ and $|f(x) - f(p)| \leq \frac{\alpha}{4}$. For n sufficiently large, $d(p_n, p) \leq \beta$, hence

$$\left|f_{U,\mathbf{S}}(p_n) - f_{U,\mathbf{S}}(p)\right| \leq \frac{\alpha}{4},$$
$$|f(p_n) - f(p)| \leq \frac{\alpha}{4}.$$

On the one hand, since $f_{U,\mathbf{S}}(p) = f(p)$, we get $\left|f_{U,\mathbf{S}}(p_n) - f(p_n)\right| \leq \frac{\alpha}{2}$. On the other hand,

$$\begin{aligned} f(p_n) = f((1-\lambda_n)q_n + \lambda_n r_n) &\leq (1-\lambda_n)f(q_n) + \lambda_n f(r_n) - \alpha \\ &= (1-\lambda_n)f_{U,\mathbf{S}}(q_n) + \lambda_n f_{U,\mathbf{S}}(r_n) - \alpha \\ &= f_{U,\mathbf{S}}((1-\lambda_n)q_n + \lambda_n r_n) - \alpha \\ &= f_{U,\mathbf{S}}(p_n) - \alpha, \end{aligned}$$

thus $f_{U,\mathbf{S}}(p_n) - f(p_n) \geq \alpha$, which contradicts the inequality $\left|f_{U,\mathbf{S}}(p_n) - f(p_n)\right| \leq \frac{\alpha}{2}$. □

Proof of Theorem 11.

Let T be the set of triangles induced by $f_{U,\mathbf{S}}$ and A be a connected component of $(U\backslash\mathbf{S})\backslash\cup_{t\in T}t^o$. We have to prove that $V = \overline{A}$ meets at most two sites of S. Suppose on the contrary that $V\backslash A$ is the union of $k \geq 3$ disjoints segments $S_1, S_2, S_3, ..., S_k$ included in $\mathbf{S}$. For $1 \leq i, j \leq k$, denote A_{ij} the set of $a \in A$ such that $C_a =]p'_i, p'_j[$ with $p'_i \in S_i$ and $p'_j \in S_j$ where C_a is defined in Theorem 10. Since $V\backslash\mathbf{S} = A$ is connected, it is enough to prove that for all couples (i, j), A_{ij} is open in A. Let a be in A_{ij}. Let $\rho = \min\{d(S_l, S_m) : l \neq m\}$. Consider $\varepsilon = \frac{1}{4}\min\{\rho, d(a, p'_i), d(a, p'_j)\}$ and $\delta > 0$ associated with ε by the previous lemma. We can assume that $\delta < \varepsilon/2$.

Let I be the set of points in C_a whose distances to the endpoints p'_i and p'_j are $\geq \varepsilon$. The set I is a segment $[a_i, a_j]$ which does not meet $\mathbf{S}$, hence there exists $\alpha > 0$ such that $I_\alpha = \{x \in \mathbb{R}^2 : d(x, I) \leq \alpha\}$ does not meet $\mathbf{S}$. Since the distances from a to the endpoints of I are greater than ε, there exists $\beta > 0$ such that:
a segment $[q, r]$ whose endpoints are in $\mathbf{S}$, that does not meet C_a, and that is at a distance $\leq \beta$ from a, contains a point q_i at a distance $\leq \delta/2$ from a_i and a point q_j at a distance $\leq \delta/2$ from a_j (see Figure 16).

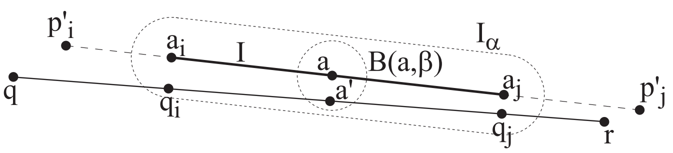

Figure 16: Illustration of the existence of β.

To see it, note that the endpoints of the segment $[q, r]$ cannot lie in I_α and if $[q, r]$ is very close to a, then $[q, r]$ is almost parallel to C_a.

Let a' be a point of A such that $d(a, a') \leq \beta$. Let us show that $C_{a'} =]q, r[$ where $q \in S_i$ and $r \in S_j$. By the choice of β, $C_{a'} =]q, r[$ must contain a point q_i at a distance $\leq \delta/2$ from a_i and a point q_j at a distance $\leq \delta/2$ from a_j.

Consider the segment $[q, q_i]$. This segment is included in $[q, r]$, which is almost parallel to $[p'_i, p'_j]$, therefore either it contains a point p at a distance $\leq \delta$ from p'_i or it is included in the ball $B(p'_i, \varepsilon + \delta)$. In the first case, we use the previous lemma with $[q, r]$ and $p' = p'_i$, it follows that C_p, which is equal to $C_{a'}$, has one of its endpoint at a distance $\leq \varepsilon$ from p'_i and therefore, $q \in S_i$. In the second, case $d(p'_i, q) \leq \rho/2$, hence $q \in S_i$. In the same way we show that $r \in S_j$. Finally, $A \cap B(a, \beta) \subset A_{ij}$ and A_{ij} is open in A. □

Lemma 10. *Let t be a triangle of a segment triangulation $\mathcal{T}$ of U and let h be the affine function equal to f at the vertices of t. The function $f_{U,\mathbf{S},\mathcal{T}} = h$ on ∂t.*

Proof. Let a and b be two vertices of t, and let e be the edge of $\mathcal{T}$ such that $]a,b[\subset e$. Since $[a,b] \subset \overline{e}$, $f_{\overline{e},\mathbf{S}} \leq f_{[a,b],\mathbf{S}}$ on $[a,b]$. On the one hand, by definition, $f_{U,\mathbf{S},\mathcal{T}} = f_{\overline{e},\mathbf{S}}$ on e, and thus also on $]a,b[$. On the other hand, clearly, $f_{[a,b],\mathbf{S}} = h$ on $]a,b[$. It follows that $f_{U,\mathbf{S},\mathcal{T}} \leq h$ on $]a,b[$. Hence, $f_{U,\mathbf{S},\mathcal{T}} \leq h$ sur ∂t.

Let $f'_{U,\mathbf{S},\mathcal{T}}$ be the function equal to h on ∂t and equal to $f_{U,\mathbf{S},\mathcal{T}}$ everywhere else in U. From the last inequality, we can see that the function $f'_{U,\mathbf{S},\mathcal{T}}$ is locally convex on the closures of the edges of $\mathcal{T}$. Since $f'_{U,\mathbf{S},\mathcal{T}} \leq f$ on $\mathbf{S}$, it follows that $f'_{U,\mathbf{S},\mathcal{T}} \leq f_{U,\mathbf{S},\mathcal{T}}$ on the edges. Hence, $f'_{U,\mathbf{S},\mathcal{T}} = f_{U,\mathbf{S},\mathcal{T}}$ on ∂t. □

Lemma 11. *The function $f_{U,\mathbf{S},\mathcal{T}}$ is continuous.*

Proof. Let p be a point of $U \cap \mathbf{S}$. There is a finite number of edges $e_1, ..., e_m$ of $\mathcal{T}$ whose closure contains p. In the same way, there is a finite number of triangles $t_1, ..., t_n$ of $\mathcal{T}$ whose closure contains p. Since the closures of the other edges and triangles of $\mathcal{T}$ do not contain p, there exists a real number $r > 0$ such that all these triangles and edges are at distance $\geq r$ from p. Consequently,

$$U \cap B(p,r) \subset \mathbf{S} \cup \overline{e_1} \cup ... \cup \overline{e_m} \cup \overline{t_1} ... \cup \overline{t_n}.$$

Now, from Lemma 8, the functions $f_{\overline{e_i},\mathbf{S}}$ and the functions $f_{\overline{t_i},\mathbf{S}}$ are continuous at p, thus $f_{U,\mathbf{S},\mathcal{T}}$ is continuous at p.

Moreover, $f_{U,\mathbf{S},\mathcal{T}}$ is continuous on each face of $\mathcal{T}$ and it is easy to see that $f_{U,\mathbf{S},\mathcal{T}}$ is also continuous on each edge. By Lemma 10, we see that $f_{U,\mathbf{S},\mathcal{T}}$ is continuous on every side of a triangle of $\mathcal{T}$. It follows that $f_{U,\mathbf{S},\mathcal{T}}$ is continuous. □

Definition 13. *Let $\mathcal{T}$ be a segment triangulation of S and let $U = conv(\mathbf{S})$. The slope of $\mathcal{T}$ is defined as follows :*

$$\sigma(\mathcal{T}) = \sup\{\frac{f_{U,\mathbf{S},\mathcal{T}}(p) - f_{U,\mathbf{S},\mathcal{T}}(q)}{d(p,q)} : p \in U \setminus \mathbf{S},\ q \in U \cap \mathbf{S},\ [p,q] \subset U\}.$$

Proposition 2. *If $\mathcal{T}$ is a segment triangulation of S, then $\sigma(\mathcal{T}) < +\infty$.*

Proof. We first show that $f_{U,\mathbf{S},\mathcal{T}}$ is piecewise C^1 and that the partial derivatives of $f_{U,\mathbf{S},\mathcal{T}}$ are bounded on each piece.

The pieces are given by the triangulation $\mathcal{T}$. We know that $f_{U,\mathbf{S},\mathcal{T}}$ is affine on each face of $\mathcal{T}$, this means that the partial derivatives of $f_{U,\mathbf{S},\mathcal{T}}$ exist and are constant on each face. The case of edges is less obvious, they need to be decomposed. The closure of each edge is an union of finitely many triangles and trapezoids. Furthermore, these triangles and theses trapezoids are union of segments where $f_{U,\mathbf{S},\mathcal{T}}$ is affine (see Figure 17). We study separately the cases of triangles and of trapezoids.

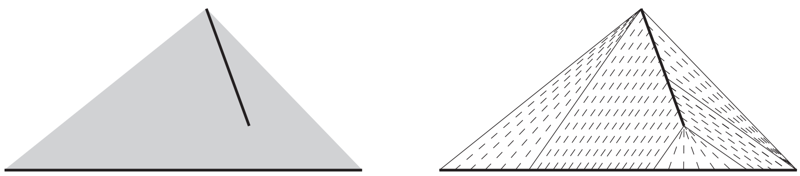

(a) (b)

Figure 17: An edge of a segment triangulation (a) is a finite union of triangles and trapezoids, which are union of segments where $f_{U,\mathbf{S},\mathcal{T}}$ is affine.

Cases of triangles. Consider a triangle $t = conv(O, A, B)$ included in the closure of an edge such that O is in a site $S_1 \in S$, A and B are in the same site $S_2 \in S$, and $f_{U,\mathbf{S},\mathcal{T}}$ is affine on each segment $[O,q]$ with q in $[A,B]$. In the frame $(O, \overrightarrow{i} = \overrightarrow{OA}, \overrightarrow{j} = \overrightarrow{OB})$, a point p of coordinates (x,y) is in t if, and only if, $x, y \geq 0$ and $x + y \leq 1$. Since the point q of coordinates $\frac{1}{x+y}(x,y)$ is in $[A,B]$, $f_{U,\mathbf{S},\mathcal{T}}$ is affine on $[O,q]$, thus

$$\begin{aligned} f_{U,\mathbf{S},\mathcal{T}}(p) &= (1-(x+y)) f_{U,\mathbf{S},\mathcal{T}}(O) + (x+y) f_{U,\mathbf{S},\mathcal{T}}(q) \\ &= (1-(x+y)) f_{U,\mathbf{S},\mathcal{T}}(O) + (x+y) f(q). \end{aligned}$$

In the frame $(O, \overrightarrow{i}, \overrightarrow{j})$, we get

$$f_{U,S,\mathcal{T}}(x, y) = (1 - (x + y)) f_{U,S,\mathcal{T}}(0, 0) + (x + y) f(\frac{x}{x + y}, \frac{y}{x + y}).$$

Therefore, the function $f_{U,S,\mathcal{T}}$ is C^1 on $t \backslash \{O\}$. In the frame $(O, \overrightarrow{i}, \overrightarrow{j})$, the function f is of the form

$$f(x, y) = Ax^2 + 2Bxy + Cy^2 + ax + by + c$$

where A, B, C, a, b and c are real numbers. Simple computations show that for $(x, y) \neq (0, 0)$ in t,

$$f_{U,S,\mathcal{T}}(x, y) = ax + by + \frac{Ax^2 + 2Bxy + Cy^2}{(x + y)}$$

and

$$\frac{\partial f_{U,S,\mathcal{T}}}{\partial x} = a - \frac{Ax^2 + 2Bxy + Cy^2}{(x + y)^2} + \frac{2Ax + 2By}{(x + y)}.$$

Since x and y are non negative, we get

$$\left| \frac{\partial f_{U,S,\mathcal{T}}}{\partial x}(x, y) \right| \leq |a| + 3 \max(|A|, |B|, |C|).$$

In the same way, for all $(x, y) \in t$, $(x, y) \neq (0, 0)$,

$$\left| \frac{\partial f_{U,S,\mathcal{T}}}{\partial y}(x, y) \right| \leq |b| + 3 \max(|A|, |B|, |C|).$$

Case of trapezoids. We can choose the frame $(O, \overrightarrow{i}, \overrightarrow{j})$ such that a point of coordinates (x, y) is in the trapezoid t iff $x, y \geq 0$ and $a \leq x + y \leq b$ where $b \geq a > 0$. A point in t of coordinates (x, y) lies in the segment whose end points are $(x + y, 0)$, $(0, x + y)$ and $f_{U,S,\mathcal{T}}$ is affine on this segment. Therefore,

$$\begin{aligned} f_{U,S,\mathcal{T}}(x, y) &= \frac{x}{x + y} f_{U,S,\mathcal{T}}(x + y, 0) + \frac{y}{x + y} f_{U,S,\mathcal{T}}(0, x + y) \\ &= \frac{x}{x + y} f(x + y, 0) + \frac{y}{x + y} f(0, x + y). \end{aligned}$$

It follows that $f_{U,S,\mathcal{T}}$ is C^1 on t.

Since there are finitely many such triangles and trapezoids, we can find a constant M in $\mathbb{R}$ such that

$$\left| \frac{\partial f_{U,S,\mathcal{T}}}{\partial x}(p) \right|, \left| \frac{\partial f_{U,S,\mathcal{T}}}{\partial y}(p) \right| \leq M$$

for all p in a trapezoid or in a triangle (except its vertices) associated with $\mathcal{T}$. Finally, by the mean value theorem, the previous inequalities and the continuity of $f_{U,S,\mathcal{T}}$ imply that, for all $p, q \in U$ such that the segment $[p, q]$ is included in U,

$$\frac{f_{U,S,\mathcal{T}}(p) - f_{U,S,\mathcal{T}}(q)}{|p - q|} \leq M.$$

$\square$

In the following, we denote by $\theta(\mathcal{T})$ the minimum angle of the triangles of $\mathcal{T}$.

Proposition 3. *There exists a positive constant c depending only on f, S and U, and such that, for every segment triangulation $\mathcal{T}$ of U,*

$$\theta(\mathcal{T}) \geq \frac{c}{\max(1, \sigma(\mathcal{T}))}.$$

Proof. Let $\Delta = conv(\{p_1, p_2, p_3\})$ be a triangle of $\mathbb{R}^2$. For $1 \le i \le 3$, denote h_i the length of the altitude through p_i, and a_i the lenght of the opposite side. We define the flatness of the triangle Δ by $flat(\Delta) = \max(\frac{a_1}{h_1}, \frac{a_2}{h_2}, \frac{a_3}{h_3})$. The flatness of a segment triangulation $\mathcal{T}$ of U is $flat(\mathcal{T}) = \max\{flat(\Delta) : \Delta \in \mathcal{T}\}$. It is not difficult to prove that $\theta(\Delta) \ge \frac{1}{\frac{3\sqrt{3}}{\pi} flat(\Delta)}$. Consequently, it is enough to show that there exists a positive constant C depending only on f, S and U, and such that, for every segment triangulation $\mathcal{T}$ of U,

$$flat(\mathcal{T}) \le C \max(1, \sigma(\mathcal{T})).$$

Let Δ be a triangle of a segment triangulation of U, and let h be the only affine function equal to f at the vertices of Δ. Set

$$\sigma(\Delta) = \sup\{\frac{h(p) - f(q)}{d(p,q)} : p \in \Delta\backslash\mathbf{S},\ q \in \mathbf{S} \cap \Delta\}.$$

Note that, if $\Delta \in \mathcal{T}$, then $\sigma(\Delta) \le \sigma(\mathcal{T})$. Thus, it is enough to prove that, for every triangle Δ of a segment triangulation of U,

$$flat(\Delta) \le C \max(1, \sigma(\Delta)).$$

Let $\Delta = conv(p, q, r)$ be a triangle of a segment triangulation of U. Since p, q and r are in distinct segments of S, there exists $l_0 > 0$ depending only on S such that $d(p,q),\ d(q,r),\ d(r,p) \ge l_0$. Suppose that the angle at the vertex r of Δ is maximal and denote s the foot of the altitude through r.

If two angles of Δ are greater than $\frac{\pi}{3}$, it means that one of the angles at p or at q is greater than $\frac{\pi}{3}$. It follows that $d(r,s) \ge (\sin\frac{\pi}{3}) l_0 \ge \frac{\sqrt{3}}{2} l_0$. Now, if $D = diam(U)$, we have $d(p,q) \le D$. Thus,

$$flat(\Delta) = \frac{d(p,q)}{d(r,s)} \le \frac{D}{\frac{\sqrt{3}}{2} l_0} \le \frac{2D}{l_0}.$$

If only one angle of Δ is greater than $\frac{\pi}{3}$, it must be the angle at r. In this case, we have $s \in [p,q]$ et $d(p,s), d(q,s) \ge (\cos\frac{\pi}{3}) l_0 = \frac{1}{2} l_0$. Set $s = (1-t)p + tq$. Since $d(p,q) \le D$, we have

$$\frac{l_0}{2D} \le t \le 1 - \frac{l_0}{2D}.$$

Denote

$$m = \min\{(1-t)f(p) + tf(q) - f((1-t)p + tq) : \\ p, q \in U,\ d(p,q) \ge l_0,\ \frac{l_0}{2D} \le t \le 1 - \frac{l_0}{2D}\}.$$

Since f is strictly convex, m is strictly positive. Furthermore, since f is uniformly continuous on U, there exists $\delta > 0$ such that $d(x,y) \le \delta$ implies $|f(x) - f(y)| \le \frac{m}{2}$. Since $flat(\Delta) = \frac{d(p,q)}{d(r,s)}$ and $d(p,q) \le D$, we have $flat(\Delta) \le \frac{D}{d(r,s)}$.

If $d(r,s) \ge \delta$, then

$$flat(\Delta) \le \frac{D}{d(r,s)} \le \frac{D}{\delta}.$$

Otherwise, when $d(r,s) \le \delta$, we have

$$\begin{aligned} f(r) &\le f(s) + \frac{m}{2} &(7)\\ &= f((1-t)p + tq) + \frac{m}{2} &(8)\\ &\le (1-t)f(p) + tf(q) - m + \frac{m}{2} &(9)\\ &\le h(s) - \frac{m}{2} &(10) \end{aligned}$$

hence

$$\sigma(\Delta) \ge \frac{h(s) - f(r)}{d(s,r)} \ge \frac{m}{2d(r,s)} \ge \frac{m}{2D} flat(\Delta).$$

Finally, $flat(\Delta) \le C \max(1, \sigma(\Delta))$ with $C = D\max(\frac{2}{l_0}, \frac{1}{\delta}, \frac{2}{m})$. □

The algorithm starts with a segment triangulation $\mathcal{T}_0$ of $U = conv(\mathbf{S})$ and computes a sequence $\mathcal{T}_1, \mathcal{T}_2, ..., \mathcal{T}_n, ...$ of new triangulations.

Theorem 14. *The sequence of functions* $(f_{conv(\mathbf{S}),\mathbf{S},\mathcal{T}_n})_{n\in\mathbb{N}}$ *decreases to* $f_{conv(\mathbf{S}),\mathbf{S}}$ *as n goes to infinity.*

Corollary 3. *There exists an integer N such that, for all integers* $n \geq N$*, the triangulation* $\mathcal{T}_n$ *has the same topology as the segment Delaunay triangulation of S.*

Proof of Theorem 14.

Set $f_n = f_{conv(\mathbf{S}),\mathbf{S},\mathcal{T}_n}$. The first thing to prove is that the sequence $(f_n)_{n\in\mathbf{N}}$ is decreasing. At n-th stage, the algorithm is performed with the edge e_n at the head of the queue. Denote t_1, t_2 the two triangles adjacent to e_n (or t_1 the unique triangle adjacent to e_n) and $a_1, ..., a_m$, $m \leq 4$, the other edges adjacent to the triangles t_i. The input polygon is $P_n = \overline{e_n \cup t_1 \cup t_2}$ (or $\overline{e_n \cup t_1}$). The algorithm computes a triangulation Q_n of P_n induced by $f_{P_n,\mathbf{S}}$. Let us denote the elements of Q_n by:

– t'_1 and t'_2 the new triangles (or t'_1),
– b_i the edge of Q_n adjacent to a_i, $1 \leq i \leq m$,
– e_{n+1} the new edge that replaces e_n.

The new edges of $\mathcal{T}_{n+1}$ are e_{n+1} and the edges $a'_i = a_i \cup b_i$, $1 \leq i \leq m$. By Theorem 13, since t_1, t_2, and e_n are included in P_n, we have $f_n \geq f_{P_n,\mathbf{S}}$ on P_n. Now, $f_{P_n,\mathbf{S}} = f_{P_n,\mathbf{S},Q_n}$, therefore $f_{n+1} = f_{P_n,\mathbf{S},Q_n} = f_{P_n,\mathbf{S}} \leq f_n$ on $t'^{\circ}_1 \cup t'^{\circ}_2 \cup e_{n+1}$. It remains to show that $f_{n+1} \leq f_n$ on a'_i, $i = 1, ..., m$. On the one hand, since $a_i \subset a'_i$, $f_{n+1} = f_{\overline{a'_i},\mathbf{S}} \leq f_{\overline{a_i},\mathbf{S}} = f_n$ on a_i. On the other hand, since $b_i \subset a'_i$, $f_{n+1} = f_{\overline{a'_i},\mathbf{S}} \leq f_{\overline{b_i},\mathbf{S}} = f_{P_n,\mathbf{S}} \leq f_n$ on b_i. Therefore, $f_{n+1} \leq f_n$ on $t'^{\circ}_1 \cup t'^{\circ}_2 \cup e_{n+1} \cup a'_1 \cup ... \cup a'_m$, which implies $f_{n+1} \leq f_n$ on $conv(\mathbf{S})$. By the way, we have proved the inequalities $f_{n+1} \leq f_{P_n,\mathbf{S}} \leq f_n$ on P_n.

It follows that the sequence of functions $(f_n)_{n\in\mathbb{N}}$ decreases to a function $g : conv(\mathbf{S}) \to \mathbb{R}$. The only thing to show is that g is locally convex.

Since $g \geq f$ on $conv(\mathbf{S})$ and $g = f$ on $\mathbf{S}$, it suffices to show that g is convex on any open segment $]p_0, p_1[$ included in the interior of $conv(\mathbf{S})$ and that does not meet $\mathbf{S}$. Let p be a point of such a segment $]p_0, p_1[$. The theorem will be proved if we can show that there exists a ball $B(p, \varepsilon)$ centered at p with radius $\varepsilon > 0$ and infinitely many integers n such that $I_{p,\varepsilon} =]p_0, p_1[\cap B(p, \varepsilon)$ is included either in a triangle of $\mathcal{T}_n$ or in the input polygon P_n. Indeed, for these integers n, either f_n or $f_{P_n,\mathbf{S}}$ is convex on $I_{p,\varepsilon}$, and since $f_{n+1} \leq f_{P_n,\mathbf{S}} \leq f_n$ on P_n, the function g is a limit of a sequence of convex functions on $I_{p,\varepsilon}$.

The inequality $f_{n+1} \leq f_n$ implies that $\sigma(\mathcal{T}_{n+1}) \leq \sigma(\mathcal{T}_n)$, therefore

$$\theta(\mathcal{T}_n) \geq \frac{c}{\max(1, \sigma(\mathcal{T}_n))} \geq \frac{c}{\max(1, \sigma(\mathcal{T}_0))}$$

for all integers n. Thus the angles of any triangle t generated by the algorithm are bounded from below by a constant $c_0 > 0$. Now, it is not difficult to see that there is a positive real number ε such that, if t is a triangle generated by the algorithm and if t meets the segment $I_{p,\varepsilon}$, then the length of the segment $]p_0, p_1[\cap t$ is greater than 2ε.

Case 1: Suppose that, at the step n_0, p lies in an edge e of the segment triangulation $\mathcal{T}_{n_0}$. While the topological edge e is not at the head of the queue, the geometrical edge e is increasing. At a later stage, the algorithm will be performed with the segment triangulation $\mathcal{T}_{n_1}$ and the edge e, which still contains p. By the choice of ε, the only triangles of $\mathcal{T}_{n_1}$ that can meet $I_{p,\varepsilon}$ are the triangles t_1 and t_2 adjacent to the edge e. Since the length of the segment $]p_0, p_1[\cap t_i$ is greater than 2ε, $I_{p,\varepsilon}$ is included in $e \cup t_1 \cup t_2$. This means that $I_{p,\varepsilon} \subset P_{n_1}$.

Case 2: Suppose that, at the step n_0, p lies in a triangle t of the segment triangulation $\mathcal{T}_{n_0}$. If $I_{p,\varepsilon}$ is included in t, we are done. Otherwise, $I_{p,\varepsilon}$ meets an edge e adjacent to t. As in Case 1, $I_{p,\varepsilon}$ is included in the input polygon P_{n_1} at the stage $n_1 \geq n_0$ when e is treated.

In all cases, we have shown that there exists infinitely many integers n such that $I_{p,\varepsilon}$ is included in P_n or in a triangle of $\mathcal{T}_n$. □

Proof of Corollary 3.

The set of topologies of all the segment triangulations of $\mathbf{S}$ is finite. Hence, if the corollary does not hold, then a non Delaunay topology would appear infinitely many times. Therefore, it is enough to prove that, if for an increasing sequence of integers $(k_n)_{n\in\mathbf{N}}$, the triangulations $\mathcal{T}_{k_n}$ have the same topology, then it is the topology of the segment Delaunay triangulation.

We can always suppose that, given a topological triangle t, its geometrical representations t_{k_n} in $\mathcal{T}_{k_n}$ converge to a triangle t_∞ when n goes to infinity (just take subsequences of $\mathcal{T}_{k_n}$). The set of all these triangles t_∞ defines a segment triangulation $\mathcal{T}_\infty$. It is clear that the triangulation $\mathcal{T}_\infty$ has the same topology as the $\mathcal{T}_{k_n}$. The function $f_{conv(\mathbf{S}),\mathbf{S}} = \lim_{n\to\infty} f_{conv(\mathbf{S}),\mathbf{S},\mathcal{T}_{k_n}}$ must be affine on each of these triangles t_∞. Therefore, all the triangles of $\mathcal{T}_\infty$ are Delaunay triangles and $\mathcal{T}_\infty$ is the segment Delaunay triangulation of **S**. $\square$

5. Conclusion

The aim of this paper was to show that the segment Delaunay triangulation can be constructed by a flip algorithm in finitely many steps. The precise complexity of the algorithm seems difficult to estimate since we do not know of any quantitative measure of the improvement of the triangulation after a step of the algorithm. However, we have applied our algorithm to triangualtions of 1,000 to 40,000 randomly generated segments. The initial triangulations have been obtained from constraint triangulations build by a sweep algorithm [11]. On these examples, the number of steps of the algorithm is nearly linear with the number of sites (about 340 steps per site). About 3/4 of the steps seems useless in the sense that they do not modify the current triangulation (neither by a flip, nor by a triangle shift). This is because all the edges are systematically processed by the main loop of the algorithm. If one wants to improve the practical performances of the algorithm, one should establish a priority ordering of the edges.

In case of point sets, the Delaunay triangulation is the one that maximizes the smallest angle of its triangles [21]. The proof of the convergence of our flip algorithm also uses the control of the angles of the triangles. Moreover, the three-dimensional lifting of the segment Delaunay triangulation is below the lifting of any other segment triangulation. These are two strong hints that make us believe that the segment Delaunay triangulation should have some optimal angular properties.

In recent years, particular attention has been paid to the study of the Voronoi diagram of a set of line segments in three dimensions [18], [20], [15], ... However, the topology of this diagram is really known only for a set of three lines [13]. The definition of segment triangulation extends to three dimensions: Its three-dimensional regions are tetrahedrons having their vertices on four distinct segments. The right knowledge of these triangulations will fairly facilitate the investigations about the three-dimensional segment Voronoi diagram, since it is dual to such a triangulation.

The three-dimensional extension is certainly a difficult problem; it will be easier to consider first more general convex sites in the plane. We believe that some of the results given in this paper can be extended to this more general setting.

References

[1] O. Aichholzer, F. Aurenhammer, H. Krasser, and P. Brass. Pseudotriangulations from surfaces and a novel type of edge flip. *SIAM J. Comput.*, 32(6):1621–1653, 2003.

[2] O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. In *Proc. 22th Annu. ACM Sympos. Comput. Geom.*, pages 282–291, 2006.

[3] F. Aurenhammer and H. Krasser. Pseudo-simplicial complexes from maximal locally convex functions. *Discrete and Computational Geometry*, 35(2):201–221, 2006.

[4] J.-D. Boissonnat, O. Devillers, R. Schott, M. Teillaud, and M. Yvinec. Applications of random sampling to on-line algorithms in computational geometry. *Discrete Comput. Geom.*, 8(1):51–71, 1992.

[5] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom. Theory Appl.*, 42(1):60–80, 2009.

[6] M. Brévilliers. *Construction de la triangulation de Delaunay de segment par un algorithme de flip*. PhD thesis, Université de Haute Alsace, Mulhouse, France, 2008.

[7] L. P. Chew. Constrained Delaunay triangulations. *Algorithmica*, 4:97–108, 1989.

[8] L. P. Chew and K. Kedem. Placing the largest similar copy of a convex polygon among polygonal obstacles. In *Proc. 5th Annu. ACM Sympos. Comput. Geom.*, pages 167–174, 1989.

[9] L. P. Chew and K. Kedem. A convex polygon among polygonal obstacles: placement and high-clearance motion. *Comput. Geom. Theory Appl.*, 3:59–89, 1993.

[10] O. Devillers, G. Liotta, F. P. Preparata, and R. Tamassia. Checking the convexity of polytopes and the planarity of subdivisions. *Comput. Geom. Theory Appl.*, 11:187–208, 1998.

[11] H. Edelsbrunner. Triangulations and meshes in computational geometry. *Acta Numerica*, pages 133–213, 2000.

[12] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15(3):223–241, 1996.

[13] H. Everett, S. Lazard, D. Lazard, and M. Safey El Din. The voronoi diagram of three lines. In *SCG '07: Proceedings of the twenty-third annual symposium on Computational geometry*, pages 255–264, New York, NY, USA, 2007. ACM Press.

[14] S. J. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.

[15] V. Koltum and M. Sharir. Three dimensional euclidean voronoi diagrams of lines with a fixed number of orientations. *SIAM J. Comput.*, 32 (3):616–642, 2003.

[16] C. L. Lawson. Software for C^1 surface interpolation. In J. R. Rice, editor, *Math. Software III*, pages 161–194. Academic Press, New York, NY, 1977.

[17] D. T. Lee and A. K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete Comput. Geom.*, 1:201–217, 1986.

[18] B. Mourrain, J.-P. Técourt, and M. Teillaud. On the computation of an arrangement of quadrics in 3d. *Comput. Geom. Theory Appl.*, 30(2): 145–164, 2005.

[19] O. Perron. Eine neue Behandlung der ersten Randwertaufgabe für $\delta u = 0$. *Math. Z.*, 18(1):42–54, 1923.

[20] E. Schömer and N. Wolpert. An exact and efficient approach for computing a cell in an arrangement of quadrics. *Comput. Geom. Theory Appl.*, 33(1–2):65–97, 2006.

[21] R. Sibson. Locally equiangular triangulations. *Comput. J.*, 21(3):243–245, 1978.