

Flip Algorithm for Segment Triangulations

Mathieu Bréviliers, Nicolas Chevallier, and Dominique Schmitt

Laboratoire LMIA, Université de Haute-Alsace

4, rue des Frères Lumière, 68093 Mulhouse Cedex, France

{Mathieu.Brevilliers,Nicolas.Chevallier,Dominique.Schmitt}@uha.fr

Abstract. Given a set S of disjoint line segments in the plane, which we call sites, a segment triangulation of S is a partition of the convex hull of S into sites, edges, and faces. The set of faces is a maximal set of disjoint triangles such that the vertices of each triangle are on three distinct sites. The segment Delaunay triangulation of S is the segment triangulation of S whose faces are inscribable in circles whose interiors do not intersect S . It is dual to the segment Voronoi diagram. The aim of this paper is to show that any given segment triangulation can be transformed by a finite sequence of local improvements in a segment triangulation that has the same topology as the segment Delaunay triangulation. The main difference with the classical flip algorithm for point set triangulations is that local improvements have to be computed on non convex regions. We overcome this difficulty by using locally convex functions.

1 Introduction

In 1977, Lawson [16] has shown that any given triangulation of a planar point set can be transformed in a Delaunay triangulation (one whose triangles' circumcircles are empty of sites) by a sequence of local improvements. Every local improvement consisted in flipping a diagonal of a convex quadrilateral to the other diagonal. Since then, several extensions of flip algorithms have been proposed. On the one hand, they have been investigated in higher dimensions. The algorithm does not work as such in dimensions higher than two because flips should be applied to non convex polyhedrons, leading to geometrically unrealizable tetrahedrizations [13]. However, Joe [14] has shown that, once the Delaunay tetrahedrization of a point set in three dimensions is given, it can be updated by a sequence of flips, after the insertion of a new point. Cheng and Dey [7] have also proven that a surface triangulation that closely approximates a smooth surface with uniform density can be transformed to a Delaunay triangulation by a flip algorithm. On the other hand, flips have been studied for different types of triangulations such as constrained triangulations [9], weighted triangulations [10], pseudo-triangulations [2], pre-triangulations [1], ...

Independently of their efficiency when applied to a “not too bad” initial triangulation, flip algorithms have been implemented as subroutines for randomized algorithms [11]. They also enable to prove important properties of the manipulated triangulations. For example, they have been used for proving that, among

all triangulations of a point set in the plane, the Delaunay triangulation maximizes the minimum angle [16]. They also enable to structure and to enumerate triangulations as vertices of a graph in which two vertices are adjacent if they differ from each other by a flip [12], [5].

In this work we address the question of flip algorithm for the segment triangulations that have been introduced in [4]. Given a finite set S of disjoint line segments in the plane, a segment triangulation of S is a maximal set of disjoint triangles, each of them having its vertices on three distinct sites of S (see Figure 1). Segment triangulations form a very natural family of diagrams containing the dual of the segment Voronoi diagram. This dual diagram, called the segment Delaunay triangulation (or edge Delaunay triangulation), has been introduced much earlier by Chew and Kedem [8]. A topological dual of the segment Voronoi diagram has also been used to implement efficiently the construction of the segment Voronoi diagram in the CGAL Library [15]. In [4], we have given a local characterization of the segment Delaunay triangulation among the family of all segment triangulations of S as well as a local characterization of its topology.

An obstacle arises when trying to transform a segment triangulation into the segment Delaunay triangulation by a sequence of local improvements: As for three dimensional point sets, local transformations must be performed on non convex regions. We overcome this difficulty by allowing local improvements that not necessarily imply changes in the topology, as flips do. In order to characterize these local improvements and to prove that the constructed triangulations tend toward 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 [2], [3], ...).

Another difficulty comes out of segment triangulations: There are infinitely many segment triangulations of a given set of sites, 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 topology as the segment Delaunay triangulation. We shall show that such a triangulation is obtained in finitely many steps, thanks to geometrical estimates about the angles of the triangles arising during the algorithm. The segment Delaunay triangulation can then be deduced from this triangulation in linear time.

2 Segment Triangulations

In this section, we recall the main results about segment triangulations given in [4]. They generalize the concept of triangulation to a set of disjoint segments in the plane. Afterwards, we slightly extend these results.

Throughout this paper, S is a finite set of $n \geq 2$ disjoint closed segments in the plane, which we call sites. A closed segment may possibly be reduced to a single point. We shall denote by \mathbf{S} the set of points of the segments of S . 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.

Definition 1. A segment triangulation \mathcal{T} of S is a partition of the convex hull $\text{conv}(\mathbf{S})$ of \mathbf{S} in disjoint sites, edges, and faces such that:

1. Every face of \mathcal{T} is an open triangle whose vertices are in three distinct sites of S and whose open edges do not intersect \mathbf{S} ,
2. No face can be added without intersecting another one,
3. The edges of \mathcal{T} are the (possibly two-dimensional) connected components of $\text{conv}(\mathbf{S}) \setminus (F \cup \mathbf{S})$, where F is the union of faces of \mathcal{T} .

In the following, the word “triangle” will only be used for faces and never for edges, even if they have the shape of a triangle.




Fig. 1. A segment triangulation (a) (the sites appear in black, the faces in white, and the edges in gray) and its topology (b)

Using the fact that no triangle can be added to \mathcal{T} , it has been shown that the closure of an edge of a segment triangulation meets exactly two sites (see Figure 1). Thus, a planar combinatorial map M can be associated with \mathcal{T} in the following way:

- the vertices of M are the sites of S ,
- the arcs connecting two sites s and t in M are the edges of \mathcal{T} whose closures intersect s and t ,
- for every vertex s of M , 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} .

M represents the topology of \mathcal{T} . Using this topology, it has been shown that the number of faces of a segment triangulation of S depends only on S and is linear with the number of sites of S .

Definition 2. A segment triangulation of S is *Delaunay* if the circumcircle of each face does not contain any point of \mathbf{S} in its interior.

If all the sites are points, a segment Delaunay triangulation is a usual point set Delaunay triangulation. It has been shown that the segment Delaunay triangulation exists for any set S , is unique if S is in general position, and is dual to the segment Voronoi diagram.

As for point sets, the legal edge property has been defined for segment triangulations in [4]. A more intuitive formulation is:

Definition 3. *Let e be an edge adjacent to two triangles T_1 and T_2 in a segment triangulation and let r, t, u, v be the sites adjacent to T_1 and T_2 . The edge e is legal if there exists a segment triangulation \mathcal{T} of $\{r, t, u, v\}$ with the same topology as the segment Delaunay triangulation of $\{r, t, u, v\}$ and such that T_1 and T_2 are two triangles of \mathcal{T} .*

An edge adjacent to zero or one triangle is legal.

This led to a local characterization of the segment Delaunay triangulation:

Theorem 1. *A segment triangulation of S whose all edges are legal has the same topology as the Delaunay one.*

Since the segment Delaunay triangulation of four sites contains at most four triangles, it can be checked in constant time whether an edge is legal or not.

Note that the segment Delaunay triangulation can be easily computed once its topology is known. It suffices to put each triangle in its tangency position, which means that the interior of its circumcircle does not meet the three sites that contain its vertices. Thus, computing the segment Delaunay triangulation comes down to compute its topology. Therefore, the goal of our flip algorithm is to lead in finitely many “local” steps to a segment triangulation whose edges are all legal. To this aim, we shall need to constrain the segment triangulations in some subsets of the convex hull of \mathbf{S} . So, we need to extend slightly the above results.

Definition 4. *A subset U of $\text{conv}(\mathbf{S})$ is S -polygonal if U is closed and if the boundary of U is a finite union of disjoint segments of two kinds:*

- closed segments included in \mathbf{S} ,
- open segments $]p, q[$ such that $\mathbf{S} \cap [p, q] = \{p, q\}$.

Throughout this paper, U denotes an S -polygonal subset of $\text{conv}(\mathbf{S})$. Now, the definition of segment triangulations extends to U by replacing, in Definition 1, $\text{conv}(\mathbf{S})$ by U and \mathbf{S} by $U \cap \mathbf{S}$. Here again we can show that:

Theorem 2. *The number of faces of a segment triangulation of U depends only on the couple (U, S) .*

We say that a point q in U is *visible* from a point p in U if $]p, q[$ is included in $U \setminus \mathbf{S}$.

Definition 5. *1. A triangle t included in U with vertices in \mathbf{S} is a Delaunay triangle of U 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 triangles are Delaunay.

Theorem 3. *Every S -polygonal subset U admits a segment Delaunay triangulation.*

This result is a consequence of Theorem 7 of section 4. However, a segment Delaunay triangulation is not necessarily unique since four connected components of $U \cap \mathbf{S}$ may be cocircular even if S is in general position.

3 Description of the Flip Algorithm

The algorithm starts with a segment triangulation of S . The edges of the triangulation are stored in a queue.

One basic step of the algorithm goes as follows. The edge e at the head of the queue is popped. Let P be the closure of the union of e and of its at most two adjacent triangles: This region is called the input polygon of e (see Figure 2 (b) and (f)). Then, the algorithm computes a segment Delaunay triangulation of P . Since P meets at most four sites, the Delaunay triangles of P can be computed in constant time. The triangles adjacent to e are replaced with the Delaunay triangles of P . This gives rise to a new segment triangulation of S (it is a consequence of Theorems 2 and 3). Finally, the edge replacing e is pushed at the tail of the queue.

Beside this queue, the algorithm maintains the number of illegal edges in the current triangulation. The algorithm ends when all edges are legal.

If a basic step changes the topology of the current triangulation, we say that the processed edge is flipped.

In case of point set triangulations, when an illegal edge is processed by the flip algorithm, it is flipped, it becomes legal, and it will never reappear. Since there are finitely many 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 2):

- 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.

In the next two sections, we use the same approach to prove that our flip algorithm constructs a segment triangulation that has the same topology as the segment Delaunay triangulation. At first, for every S -polygonal subset 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 7). Then, we define the lifting of any segment triangulation that is not Delaunay (Definition 7) 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 8). In order to show the correctness of the algorithm, we prove that, after a basic step, the lifting of the resulting segment triangulation is lower than or equal to the lifting of the segment triangulation before the basic step (Theorem 9). This leads to prove that the sequence of basic steps builds a sequence of segment triangulations that converges to the segment Delaunay triangulation (Theorem 9). It remains to see that, after a finite number of basic steps, the segment triangulation constructed by the flip algorithm has the same topology as the segment Delaunay triangulation (Corollary 2). From Theorem 1, there is no more illegal edge in this triangulation and the algorithm stops.




Fig. 2. 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 of the edge e_3 . Then, in (c), it computes the segment Delaunay triangulation of the input polygon 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).

4 Locally Convex Functions and Segment Triangulations

Recall that, if V is a subset of \mathbf{R}^2 , a function $\phi : V \rightarrow \mathbf{R}$ is locally convex if the restriction of ϕ to each segment included in V is convex.

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 \mathbf{R}^3 . Note that it corresponds to this usual lower convex hull when the domain V is convex.

Definition 6. Let $L(V)$ be the set of functions $\phi : V \rightarrow \mathbf{R}$ 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 -polygonal domain U with the function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ defined by $f(x, y) = x^2 + y^2$. The convexity of f implies that $f_{U, \mathbf{S}} = f$ on $U \cap \mathbf{S}$. Using the geometrical assumptions on U , it can also be proven that $f_{U, \mathbf{S}}$ is continuous.

The main aim of this section is to explain that the function $f_{U, \mathbf{S}}$ determines a segment Delaunay triangulation of U (see Figure 3). Next theorem gives information about the value of the function $f_{U, \mathbf{S}}$ at a point p . For every point p in $U \setminus \mathbf{S}$, denote S_p the closure of the set of points in \mathbf{S} visible from p and V_p its convex hull (in general, V_p is not contained in U).




Fig. 3. 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 .

Theorem 4. Every point p of U belongs to a closed convex subset C of U whose extreme points are in \mathbf{S} and such that the function $f_{U, \mathbf{S}}$ is affine on C . Moreover $f_{U, \mathbf{S}}(p) = f_{V_p, \mathbf{S}}(p)$.

Proof. (1) We begin with the case $U = \text{conv}(\mathbf{S})$. In this case, the graph of $f_{U,\mathbf{S}}$ is the classical lower convex hull E of the lifting of S on the paraboloid. Every (0-, 1-, or 2-dimensional) face of E is contained in a non vertical supporting hyperplane of E , which implies that $f_{U,\mathbf{S}}$ is affine on the downward projection C of each face of E in the plane $z = 0$. Every point p of the plane belongs to such a set C . Moreover the extreme points of C are the projections of the extreme points of a face and thus they belong to S .

(2) We consider now a general S -polygonal subset U . The Theorem is more difficult to prove and we only give the steps of its proof.

If p is in \mathbf{S} or in the boundary of U , the Theorem is very easy to prove. So we may suppose that p is in the interior of U and not in \mathbf{S} . In this case, since U is S -polygonal, it can be shown that p is also in the interior of V_p . We use the result of (1) with V_p and S_p instead of $\text{conv}(\mathbf{S})$ and \mathbf{S} : There exists a convex set C containing p , included in V_p whose extreme points are in S_p and such that f_{V_p,S_p} is affine on the convex set C . Since U is an S -polygonal subset, we can see that for every point in $U \setminus \mathbf{S}$ there is a ball centered at this point whose intersection with U is convex. This property enables to prove that the convex C is entirely contained in U .

We know that the function f_{V_p,S_p} is affine on the convex set C . Since p is in the interior of V_p and since f_{V_p,S_p} is convex there exists an affine function $h : \mathbf{R}^2 \rightarrow \mathbf{R}$ equal to f_{V_p,S_p} on the convex set C and lower than or equal to f_{V_p,S_p} on V_p . This implies $h \leq f$ on S_p .

The main idea of the proof is to construct a locally convex function $g : U \rightarrow \mathbf{R}$ which is equal to h on the convex set C and such that $g \leq f$ on \mathbf{S} . Indeed, $f_{U,\mathbf{S}} \geq g$ by definition of the lower convex hull of a function and since $f_{U,\mathbf{S}}$ is convex on C , the function $f_{U,\mathbf{S}}$ must be lower than or equal to the function h on C . Therefore $f_{U,\mathbf{S}} = h$ on C .

The function g is defined as follows. Consider the open disk $A = \{q \in \mathbf{R}^2 : f(q) < h(q)\}$ and let W be the connected component of $A \cap U$ containing p . The function $g : U \rightarrow \mathbf{R}$ is defined by $g(q) = h(q)$ if $q \in W$ and $g(q) = f(q)$ otherwise. One can check that g is convex on each segment included in U , hence g is locally convex. Moreover, it is not difficult to see that C is included in $W \cup \{q \in \mathbf{R}^2 : h(q) = f(q)\}$. The last and most difficult thing to check is that $g \leq f$ on \mathbf{S} . It is enough to prove that W contains no point of \mathbf{S} . Since the function h is lower than or equal to f on all the points of \mathbf{S} visible from p and since $h < f$ on W , we know that W contains no point of \mathbf{S} visible from p . Suppose that there exists a point q in $W \cap \mathbf{S}$. Since W is connected, we can join the point p to the point q by a path γ in W and we can suppose that the length of γ is finite. The distance $\delta = d(\gamma, \partial A)$ from γ to the boundary of A is positive. It is not difficult to show that $W_\delta = W \cap \{x \in A : d(x, \partial A) \geq \delta\}$ is a closed set. It follows that there exists a shortest path from p to q in W_δ . Now, we know that for every point in $U \setminus \mathbf{S}$ there is a ball centered at this point whose intersection with U is convex. This shows that this shortest path from p to q is straight unless it meets some point in \mathbf{S} . Therefore, either q is visible from p or the shortest path contains a point of \mathbf{S} visible from p , which is impossible. \square

Corollary 1. *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 .*

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 $h : \mathbf{R}^2 \rightarrow \mathbf{R}$ 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 \mathbf{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 $U = \text{conv}(\mathbf{S})$. 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 $h_t : \mathbf{R}^2 \rightarrow \mathbf{R}$ 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, S_p) . Hence, by the convex case, f_{V_p, S_p} is affine on t . By the previous Theorem, we have $f_{U,\mathbf{S}}(p) = f_{V_p, S_p}(p)$. Since $f_{U,\mathbf{S}}$ is locally convex, we have $f_{U,\mathbf{S}} \leq f_{V_p, S_p}$ on t . Now, p is in the interior of t , therefore $f_{U,\mathbf{S}} = f_{V_p, S_p}$ 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. We are not going to prove this result, nevertheless we can explain why it is natural. On the one hand, the relative interiors of the faces of a closed convex set form a partition of this convex set (see [6]). In the case $U = \text{conv}(\mathbf{S})$, it follows that U is partitioned by the downward projections of the relative interiors of the lower faces of the convex hull of the lifting of \mathbf{S} on the paraboloid $\{z = x^2 + y^2\}$. On the other hand, in the case of an S -polygonal subset U , the maximal subsets of U where $f_{U,\mathbf{S}}$ is affine are intended to replace these downward projections. This leads to the Theorem:

Theorem 5. *For each point p in $U \setminus \mathbf{S}$ consider the set of all relatively open convex subsets of U containing p where $f_{U,\mathbf{S}}$ is affine. This set of convex subsets contains a maximal element C_p (maximal for the inclusion). Moreover, the extreme points of $\overline{C_p}$ are in \mathbf{S} and, if q is another point of $U \setminus \mathbf{S}$, either $C_p \cap C_q = \emptyset$, or $C_p = C_q$.*

The last statement of Theorem 5 means that the subsets C_p form a partition of $U \setminus \mathbf{S}$. Now we have to establish that the two-dimensional convex subsets among the C_p induce the triangles of a segment triangulation.

Theorem 6. *By decomposing the two-dimensional $(C_p)_{p \in U \setminus \mathbf{S}}$ into triangles we get the triangles of a segment triangulation of U , which we call a triangulation induced by $f_{U,\mathbf{S}}$.*

Proof. As before, by lack of place, we can only give some hints about the proof of this Theorem. Denote by C_{S_1, S_2} the set of all one-dimensional C_p with one endpoint in the site S_1 and the other in the site S_2 . Denote by U_{S_1, S_2} the union of all the segments of C_{S_1, S_2} . The continuity of $f_{U, \mathbf{S}}$ and the strict convexity of f allow to show that a point in $\mathbf{S} \setminus (S_1 \cup S_2)$ cannot be too close to a point in a segment $C_p \in C_{S_1, S_2}$. It follows that U_{S_1, S_2} is open in $V = U \setminus (\mathbf{S} \cup \{\text{the two-dimensional } C_p\})$. Now, by definition, a connected component A of V cannot meet more than one set U_{S_1, S_2} , $S_1 \in S$, $S_2 \in S$. Then, it is not difficult to prove that the closure of A meets exactly two sites. Therefore, it is impossible to add a triangle to the two-dimensional C_p without intersecting them. Thus we have a segment triangulation of U . \square

From Corollary 1 and Theorem 6 we deduce:

Theorem 7. *For any S -polygonal subset U , a segment triangulation of U is induced by $f_{U, \mathbf{S}}$ if and only if it is Delaunay.*

Using locally convex functions, we are able to lift any segment triangulation in the following way:

Definition 7. *Let \mathcal{T} be a segment triangulation of U . The function $f_{U, \mathbf{S}, \mathcal{T}} : U \rightarrow \mathbf{R}$ is equal to f on \mathbf{S} , to $f_{\bar{e}, \mathbf{S}}$ on every edge e of \mathcal{T} , and to $f_{\bar{t}, \mathbf{S}}$ on the interior of every triangle t of \mathcal{T} .*

The lifting of \mathcal{T} to \mathbf{R}^3 is the graph of the function $f_{U, \mathbf{S}, \mathcal{T}}$.

Since $f_{\bar{e}, \mathbf{S}} \geq f_{U, \mathbf{S}}$ on \bar{e} and $f_{\bar{t}, \mathbf{S}} \geq f_{U, \mathbf{S}}$ on t , we get:

Theorem 8. *If \mathcal{T} is a segment triangulation of U , then $f_{U, \mathbf{S}} \leq f_{U, \mathbf{S}, \mathcal{T}}$.*

5 Convergence of the Flip Algorithm

In case of point set triangulations, it is well known that a flip increases the smallest angle of the triangles. A weaker result holds for segment triangulations.

Given a segment triangulation \mathcal{T} of U , let the slope of \mathcal{T} be:

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

Denoting by $\theta(\mathcal{T})$ the minimum angle of the triangles of \mathcal{T} , we can prove:

Proposition 1. *There exists a positive constant c depending only on f , S , and U such that, for every segment triangulation \mathcal{T} of U , $\theta(\mathcal{T}) \geq c / (\max(1, \sigma(\mathcal{T})))$.*

It is not difficult to prove that $\sigma(\mathcal{T}) < +\infty$. Moreover, it is obvious that if \mathcal{T}' is a segment triangulation of U such that $f_{U, \mathbf{S}, \mathcal{T}} \leq f_{U, \mathbf{S}, \mathcal{T}'}$, then $\sigma(\mathcal{T}) \leq \sigma(\mathcal{T}')$.

Consider now our algorithm: It starts with a segment triangulation \mathcal{T}_0 of $\text{conv}(\mathbf{S})$ and computes a sequence $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n, \dots$ of triangulations.

Theorem 9. *The sequence $(f_n = f_{conv(\mathbf{S}),\mathbf{S},\mathcal{T}_n})_{n \in \mathbf{N}}$ decreases to $f_{conv(\mathbf{S}),\mathbf{S}}$ as n goes to infinity.*

Proof. At every stage n , we compute a Delaunay triangulation of the input polygon P_n of the edge at the head of the queue. Applying Theorem 8 to the S -polygonal subset U composed of P_n and of all the edges of \mathcal{T}_n adjacent to P_n , we get that $f_{n+1} \leq f_n$ on U , which implies that $f_{n+1} \leq f_n$ on $conv(\mathbf{S})$.

It follows that the sequence of functions $(f_n)_{n \in \mathbf{N}}$ decreases to a function $g : conv(\mathbf{S}) \rightarrow \mathbf{R}$. The only thing to show is that g is (locally) convex. Since $g = f$ on \mathbf{S} and $g \geq f$ on $conv(\mathbf{S})$, we are reduced to show that g is convex on any open segment $]p_0, p_1[$ included in the interior of $conv(\mathbf{S}) \setminus \mathbf{S}$. By Proposition 1 and since the sequence f_n decreases, the angles of the triangles generated by the algorithm are not too sharp. It allows to show that, for every point p in $]p_0, p_1[$, there exists $\varepsilon > 0$ such that the neighborhood $I_{p,\varepsilon}$ of p of radius ε in $]p_0, p_1[$ is included either in a triangle of \mathcal{T}_n or in the input polygon P_n treated at stage n , for infinitely many integers n . Thus, 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 $I_{p,\varepsilon}$, hence g is convex. Finally, since $f_n \geq f_{U,\mathbf{S}}$ for all n , we get $g \geq f_{U,\mathbf{S}}$. Moreover, $g = f_n = f$ on \mathbf{S} and g is convex, therefore $g \leq f_{U,\mathbf{S}}$. \square

Corollary 2. *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 $conv(\mathbf{S})$.*

Proof. Since the set of topologies of all the segment triangulations of S is finite, if the corollary were false, 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 representation t_{k_n} in \mathcal{T}_{k_n} converges to a triangle t_∞ when n goes to infinity (take subsequences of \mathcal{T}_{k_n}). Therefore, the function $f_{conv(\mathbf{S}),\mathbf{S}} = \lim_{n \rightarrow \infty} f_{conv(\mathbf{S}),\mathbf{S},\mathcal{T}_{k_n}}$ must be affine on each of these triangles t_∞ . Together with Theorem 7, this shows that the set of all triangles t_∞ defines the segment Delaunay triangulation of S and that all the segment triangulations \mathcal{T}_{k_n} have the same topology as the Delaunay one. \square

6 Conclusion

In this paper, we have shown that the segment Delaunay triangulation can be constructed by a flip algorithm in a finite number of steps.

The complexity of the flip algorithm seems difficult to estimate since the same edge is processed several times. Nevertheless, as for point set triangulations, we can expect that the practical complexity of the algorithm will be efficient if the input segment triangulation is not too bad. This practical complexity may be

improved by a better handling of the queue: It is not necessary to systematically insert all the edges in the queue and we could establish an ordering of these edges.

The proof of the convergence of the flip algorithm uses the control of the angles of the triangles during the algorithm. Moreover, as for point set triangulations, the segment Delaunay triangulation is the only segment triangulation whose three-dimensional lifting is convex. These are two strong hints that make us believe that the segment Delaunay triangulations should have some optimality properties.

At last, possible extensions of segment triangulations should be mentioned: Extension to three-dimensional space, to more general sites, to more general distance functions, ... 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. Aichholzer, O., Aurenhammer, F., Hackl, T.: Pre-triangulations and liftable complexes. In: Proc. 22th Annu. ACM Sympos. Comput. Geom., pp. 282–291 (2006)
2. Aichholzer, O., Aurenhammer, F., Krasser, H., Brass, P.: Pseudotriangulations from surfaces and a novel type of edge flip. *SIAM J. Comput.* 32(6), 1621–1653 (2003)
3. Aurenhammer, F., Krasser, H.: Pseudo-simplicial complexes from maximal locally convex functions. *Disc. Comput. Geom.* 35(2), 201–221 (2006)
4. Brévilliers, M., Chevallier, N., Schmit, D.: Triangulations of Line Segment Sets in the Plane. In: Arvind, V., Prasad, S. (eds.) *FSTTCS 2007*. LNCS, vol. 4855, pp. 388–399. Springer, Heidelberg (2007)
5. Brönnimann, H., Kettner, L., Pocchiola, M., Snoeyink, J.: Counting and enumerating pointed pseudotriangulations with the greedy flip algorithm. *SIAM J. Comput.* 36(3), 721–739 (2006)
6. Bronsted, A.: *An Introduction to Convex Polytopes*. Graduate Texts in Mathematics. Springer, New York (1983)
7. Cheng, S., Dey, T.K.: Delaunay edge flips in dense surface triangulations. In: *Proceedings of the 24th European Workshop on Computational Geometry*, pp. 1–4 (2008)
8. Chew, L.P., Kedem, K.: Placing the largest similar copy of a convex polygon among polygonal obstacles. In: Proc. 5th Annu. ACM Sympos. Comput. Geom., pp. 167–174 (1989)
9. Edelsbrunner, H.: Triangulations and meshes in computational geometry. In: *Acta Numerica*, pp. 133–213 (2000)
10. Edelsbrunner, H., Shah, N.R.: Incremental topological flipping works for regular triangulations. *Algorithmica* 15(3), 223–241 (1996)
11. Guibas, L.J., Knuth, D.E., Sharir, M.: Randomized incremental construction of delaunay and voronoi diagrams. *Algorithmica* 7(4), 381–413 (1992)
12. Hurtado, F., Noy, M.: Graph of triangulations of a convex polygon and tree of triangulations. *Comput. Geom. Theory Appl.* 13(3), 179–188 (1999)

13. Joe, B.: Three-dimensional triangulations from local transformations. *SIAM J. Sci. Stat. Comput.* 10(4), 718–741 (1989)
14. Joe, B.: Construction of three-dimensional delaunay triangulations from local transformations. *Comp. Aided Geom. Design.* 8(2), 123–142 (1991)
15. Karavelas, M.I.: A robust and efficient implementation for the segment Voronoi diagram. In: *Proceedings of the International Symposium on Voronoi Diagrams in Science and Engineering*, pp. 51–62 (2004)
16. Lawson, C.L.: Software for C^1 surface interpolation. In: Rice, J.R. (ed.) *Math. Software III*, pp. 161–194. Academic Press, New York, NY (1977)