

Constructing the Segment Delaunay Triangulation by Flip

Mathieu Bréviliers*

Nicolas Chevallier*

Dominique Schmitt*

Abstract

Using locally convex functions, we show that the dual of the segment Voronoi diagram in the plane can be computed by a flip algorithm.

1 Introduction

The flip algorithm is a classical method to construct the Delaunay triangulation of a set of points in the plane, starting with any given triangulation [6]. In recent years, the method has been extended to generalized triangulations of point sets such as pseudo-triangulations or pre-triangulations [2], [3], [1], ...

In this paper, we propose a flip algorithm to construct the dual of the Voronoi diagram of a set of segments in the plane. This diagram, called segment Delaunay triangulation, has been introduced by Chew and Kedem [5]. In [4], we have already defined a family of diagrams containing the segment Delaunay triangulation: the segment triangulations (see Figure 1). The faces of such a triangulation form a maximal set of disjoint triangles resting on three distinct segments.

A classical method to study flip algorithms consists in lifting the triangulations to three-dimensional space. The problem here is that lifting has to be performed on non-convex regions in the plane. As in [2] and [3], we overcome this problem with the help of locally convex functions.

Another difficulty comes out of the fact that there are infinitely many segment triangulations of a given segment set. Thus, we give a flip algorithm that constructs, in a finite number of steps, a segment triangulation that has the same topology as the segment Delaunay triangulation.

2 Segment triangulations

In this section, we recall the main results about segment triangulations given in [4].

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 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




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

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 $\text{conv}(S)$ of S in disjoint sites, edges, and faces such that:

1. Every face of \mathcal{T} is an open triangle whose vertices belong to three distinct sites of S and whose open edges do not intersect 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 $\text{conv}(S) \setminus (F \cup S)$, where F is the set 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.

An edge of such a triangulation is adjacent to exactly two sites (see Figure 1). Moreover, the set of sites and edges defines a planar graph and thus a combinatorial map which represents the topology of the segment triangulation. 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 S in its interior.

The segment Delaunay triangulation of S always exists. Moreover, since S is in general position, it is unique and dual to the segment Voronoi diagram of S . Note that the geometry of the segment Delaunay triangulation is easy to compute once its topology is

*Laboratoire MIA, Université de Haute-Alsace, Mulhouse, France {Mathieu.Brevilliers, Nicolas.Chevallier, Dominique.Schmitt}@uha.fr

known. Indeed, it suffices to put every triangle t in tangency position on the three sites on which it rests, i.e., its circumcircle is tangent to these three sites and meets them in the same order as t .

As for point sets, a segment Delaunay triangulation can be recognized with local tests using edge legality. An edge of a segment triangulation is said to be (topologically) legal if the triangles adjacent to the edge, moved to their tangency positions, are Delaunay with respect to the sites adjacent to the triangles and if they retain locally the original topology. Hence:

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

It is easy to see that the legality of an edge can be checked in constant time. Thus, there is a linear time algorithm that checks whether a given segment triangulation has the same topology as the segment Delaunay triangulation.

In this paper, we shall need to constrain the segment triangulations in some subsets of the convex hull of S . Thus, we extend slightly the above results.

Definition 3 *A subset U of $\text{conv}(S)$ 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 S ,
- open segments $]p, q[$ such that $S \cap [p, q] = \{p, q\}$.

Now, the definition of segment triangulations extends to an S -polygonal subset U of $\text{conv}(S)$ by replacing, in Definition 1, $\text{conv}(S)$ by U and S by $U \cap S$. Here again we can show that the number of faces of a segment triangulation of U depends only on the couple (U, S) .

Definition 4 *A segment triangulation \mathcal{T} of U is Delaunay if the interior of the circumcircle of each triangle t of \mathcal{T} contains no point of S that is visible from an interior point of t , i.e., the open segment connecting these two points is not included in $U \setminus S$.*

Theorem 5 of section 4, shows that a segment Delaunay triangulation of U always exists. However, it is not necessarily unique since four connected components of $U \cap S$ 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. The edge e at the head of the queue is popped and a Delaunay triangulation of the S -polygonal subset P , union of e and of its adjacent triangles, is constructed (P is called the input polygon of e ; see Figure 2). This gives rise to a new segment triangulation.




Figure 2: Input polygons of some edges.




Figure 3: The flip algorithm transforms the given segment triangulation (a) in a segment triangulation (d) that has the same topology as the Delaunay one (e). The edge e_1 of (a) is treated but remains illegal in (b) because it cannot be flipped. The legal edge e_2 has to be processed before the flip of e_1 .

The edge replacing e is pushed at the tail of the queue. Beside this queue, a list of illegal edges is maintained. The algorithm ends when all edges are legal.

Studying the different cases, we can show that a Delaunay triangulation of P can be computed in constant time. If this triangulation admits two triangles and if the edge between them does not connect the same two sites as the edge used to determine P , then the edge is said to be flipped.

Even if the algorithm looks very close to the classical flip algorithm, there are important differences in their convergences. In case of segment sets:

- some illegal edges cannot be flipped (see Figure 3),
- a new constructed edge is not necessarily legal,
- a removed topological edge can reappear (Figure 4).

This shows that neither the legality of the edges nor the flip count suffices to prove the convergence of the algorithm. Another way to prove the convergence of the point set 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.

The lower convex hull of a set S of segments lifted on the paraboloid, also projects downward onto the segment Delaunay triangulation of S (see Theorem 5). The main difficulty is to lift the other segment

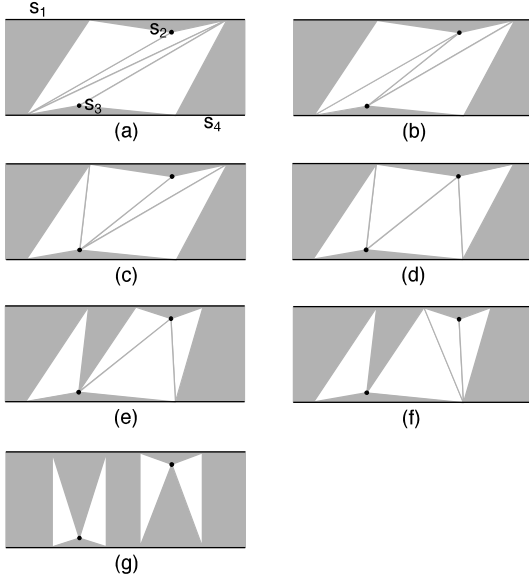


Figure 4: The edge flipped between (a) and (b) remains illegal. The edge connecting s_2 and s_4 in (b) disappears in (c) and reappears in (d).

triangulations and especially their non convex edges. To this aim we use locally convex functions.

4 Locally convex functions and segment triangulations

Recall that a real-valued function ϕ on the line segment s is convex if $\phi(tx + (1-t)y) \leq t\phi(x) + (1-t)\phi(y)$, for all t in $[0, 1]$ and all $x, y \in s$. More generally, if V is a subset of \mathbf{R}^2 and $\phi : V \rightarrow \mathbf{R}$ is a function, we say that ϕ 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 .

Definition 5 Given a real-valued function f defined on $V \cap S$, the lower convex hull of f on (V, S) is $f_{V,S} = \sup\{\phi : V \rightarrow \mathbf{R} : \phi \in L(V), \forall x \in V \cap S, \phi(x) \leq f(x)\}$ where $L(V)$ is the set of functions $\phi : V \rightarrow \mathbf{R}$ that are locally convex on V .

In the following, U denotes an S -polygonal subset of $\text{conv}(S)$ and the above definition will be used 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 $f_{U,S} = f$ on $U \cap S$. It can also be proven that $f_{U,S}$ is continuous.

The main aim of this section is to explain that the function $f_{U,S}$ determines a segment Delaunay triangulation of U . Next theorem gives information about the value of the function $f_{U,S}$ at a point p . It begins by the simplest case where U is convex. Then it shows how to reduce the general case to the convex case. For

every point p in $U \setminus S$, denote S_p the closure of the set of points in S visible from p and V_p its convex hull (in general, V_p is not contained in U). The theorem asserts that $f_{U,S}(p)$ depends only on the lower convex hull of f on (V_p, S_p) .

Theorem 2

1. If U is convex, then every point of $U \setminus S$ belongs to a closed convex subset C of U whose extremal points are in S and such that $f_{U,S}$ is affine on C .
2. In case of a (non convex) S -polygonal subset U , let p be a point of $U \setminus S$. If C is a closed convex subset of V_p , containing p , whose extremal points are in S_p , and such that f_{V_p, S_p} is affine on C , then C is included in U and $f_{U,S} = f_{V_p, S_p}$ on C .

The next step consists in showing that U can be partitioned into maximal convex subsets where the function $f_{U,S}$ is affine.

Theorem 3 Every point p in $U \setminus S$ belongs to a convex subset C_p that is maximal for the inclusion among the relatively open convex subsets of U where $f_{U,S}$ is affine. Moreover, the extremal points of C_p are in S and, if q is another point of $U \setminus S$, either $C_p \cap C_q = \emptyset$, or $C_p = C_q$.

The last statement of Theorem 3 means that the subsets C_p form a partition of $U \setminus S$. Now we have to establish that the two-dimensional convex subsets among the C_p are the faces of a segment triangulation.

Theorem 4 By decomposing the two-dimensional $(C_p)_{p \in U \setminus S}$ into triangles we get the faces of a segment triangulation \mathcal{T} of U , which we call a triangulation induced by $f_{U,S}$.

Suppose now that $U = \text{conv}(S)$ and let t be a triangle of \mathcal{T} and h the affine function that is equal to $f_{U,S}$ on t . The graph of h is a plane and its intersection with the graph of f is an ellipse whose downward projection is the circumcircle of t . Since U is convex, the function $f_{U,S}$ is convex. Therefore, $h \leq f_{U,S}$ on U . It follows that $h \leq f$ on $S \cap U$. We deduce that the circumcircle of t does not contain any point of S in its interior. By definition of a face of a segment Delaunay triangulation, we conclude that:

Theorem 5

1. If $U = \text{conv}(S)$, the segment triangulation induced by $f_{U,S}$ is the segment Delaunay triangulation of S .
2. For any S -polygonal subset U , a segment triangulation of U is induced by $f_{U,S}$ if and only if it is Delaunay.

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

Definition 6 Let \mathcal{T} be a segment triangulation of U . The function $f_{U,S,\mathcal{T}} : U \rightarrow \mathbf{R}$ is equal to f on S , to $f_{\bar{e},S}$ on any edge e of \mathcal{T} , and to $f_{t,S}$ on the interior of any triangle t of \mathcal{T} .

The lifting of \mathcal{T} to \mathbf{R}^3 is the graph of the function $f_{U,S,\mathcal{T}}$. Using the previous results we have then:

Theorem 6

1. If \mathcal{T} is a segment triangulation of U , then $f_{U,S} \leq f_{U,S,\mathcal{T}}$. Moreover $f_{U,S} = f_{U,S,\mathcal{T}}$ if and only if \mathcal{T} is induced by $f_{U,S}$.
2. If $U = \text{conv}(S)$, then \mathcal{T} is the segment Delaunay triangulation of S if and only if $f_{U,S} = f_{U,S,\mathcal{T}}$.

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:

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

Denoting by $\theta(\mathcal{T})$ the minimal angle of the triangles of \mathcal{T} , we get then:

Proposition 7 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 $\sigma(\mathcal{T}) < +\infty$ and, if \mathcal{T}' is a segment triangulation of U such that $f_{U,S,\mathcal{T}} \leq f_{U,S,\mathcal{T}'}$, then $\sigma(\mathcal{T}) \leq \sigma(\mathcal{T}')$.

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

Theorem 8 The sequence $(f_n = f_{\text{conv}(S),S,\mathcal{T}_n})_{n \in \mathbf{N}}$ decreases to $f_{\text{conv}(S),S}$ 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 6 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 $\text{conv}(S)$.

It follows that the sequence of functions $(f_n)_{n \in \mathbf{N}}$ decreases to a function $g : \text{conv}(S) \rightarrow \mathbf{R}$. The only thing to show is that g is locally convex, i.e., g is convex on any open segment $]p_0, p_1[$ included in the interior of $\text{conv}(S) \setminus S$. Since the angles of the triangles generated by the algorithm are not too sharp, it can be shown that, for every point p of $]p_0, p_1[$, there exists $\varepsilon > 0$ such that the neighbourhood $I_{p,\varepsilon}$ of p of length 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 $f_{P_n,S}$ is convex on $I_{p,\varepsilon}$, and since $f_{n+1} \leq f_{P_n,S} \leq f_n$ on P_n , the function g is a limit of a sequence of convex functions on $I_{p,\varepsilon}$. \square

Now, note that the set of topologies of all the segment triangulations of S is finite. We can also show that the only topology that appears infinitely many times in the sequence $(\mathcal{T}_n)_{n \in \mathbf{N}}$, is the topology of the segment Delaunay triangulation. Thus:

Corollary 9 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 $\text{conv}(S)$.

6 Conclusion

The aim of this paper was to show that the dual of the segment Voronoi diagram can be constructed by a flip algorithm in a finite number of steps. The remaining computational problems concern the implementation of the algorithm: robustness, time complexity, ... The algorithm has also to be compared with standard methods for computing segment Voronoi diagrams.

From the theoretical point of view, the fact that the angles of the triangles that appear during the algorithm cannot be too sharp, makes us believe that the segment Delaunay triangulation 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, ...

References

- [1] O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. In *Proc. 22th Annu. ACM Sympos. Comput. Geom.*, pages 282–291, 2006.
- [2] 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.
- [3] F. Aurenhammer and H. Krasser. Pseudo-simplicial complexes from maximal locally convex functions. *Discrete and Computational Geometry*, 35(2):201–221, 2006.
- [4] M. Bréviliers, N. Chevallier, and D. Schmitt. Triangulations of line segment sets in the plane. In *FSTTCS*, pages 388–399, 2007.
- [5] 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.
- [6] 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.