Voronoi diagrams on piecewise flat surfaces and an application to biological growth

C. Indermitte[a, *], Th.M. Liebling[a], M. Troyanov[a], H. Clémençon[b]

[a]*Departement d'Informatique, Ecole Polytechnique Federale de Lausanne, 105, Lausanne, Switzerland*
[b]*IBSG, Batiment de Biologie, Université de Lausanne, Switzerland*

Abstract

This paper introduces the notion of Voronoi diagrams and Delaunay triangulations generated by the vertices of a piecewise flat, triangulated surface. Based on properties of such structures, a generalized flip algorithm to construct the Delaunay triangulation and Voronoi diagram is presented. An application to biological membrane growth modeling is then given. A Voronoi partition of the membrane into cells is maintained during the growth process, which is driven by the creation of new cells and by restitutive forces of the elastic membrane. © 2001 Published by Elsevier Science B.V.

Keywords: Triangulation; Voronoi diagram; Delaunay triangulation; Flip algorithm

1. Introduction

The Delaunay triangulation of the convex hull of a finite set $V \subset \mathbb{R}^2$ can be computed with the well-known *flip algorithm*. The goal of this paper is to generalize this algorithm to the case of a piecewise flat surface where the set of singular points V is also the set of vertices of the triangulation. We will also show how to construct the associated Voronoi partition.

In the last section of this paper, we give a short description of a model of biological membrane growth and show the role of the Voronoi diagram therein.

2. Review of the planar case

Let V be a finite set in the plane. The *Voronoi cell* associated with a point $v \in V$ is the subset $\{x \in \mathbb{R}^2 \mid dist(x, v) \leqslant dist(x, v') \ \forall v' \in V\}$ of $\mathbb{R}^2$ and the *Voronoi diagram* is the complex defined by all Voronoi cells.
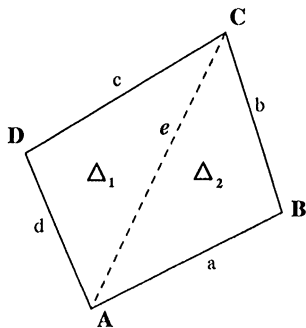
---

* Corresponding author.

Fig. 1. The hinge of an edge.

If the set V is in general position (no three points are aligned and no four points lie on a circle) then the Voronoi diagram is the dual of a triangulation of the convex hull of V called the *Delaunay triangulation*.

It is known that the Delaunay triangulation is characterized by the condition that each edge is *legal*, i.e. either it belongs to the boundary of $Conv(V)$ or the circumscribing circle to one of the two triangles incident with that edge contains no other point of V. See [1] or [3] for more on these notions.

Locality of the circle test. Any interior edge e is incident with two triangles. This pair of triangles will be called the *hinge* of edge e and denoted by $\Diamond_{e'}$ (Fig. 1). To verify the legality of an interior edge e, it suffices to look at the hinge $\Diamond_e$ and check whether the opposite vertex of one triangle is not contained in the circumscribing circle of the other. If this condition is satisfied, we say that two triangles are *confined*.

A *flip* is a local modification of a triangulation that consists in replacing the diagonal e in $\Diamond_e$ by the other diagonal e' and thus creating a new hinge $\Diamond_{e'}$, provided the associated quadrilateron is convex.

The flip algorithm: Starting with any initial triangulation of $Conv(V)$, a sequence of flips is performed on its illegal edges until there are none such left. At the end of this algorithm, all edges are legal and the resulting triangulation is a Delaunay triangulation. This algorithm stops after at most $O(|V|^2)$ flips [6].

3. Piecewise flat surfaces

Consider a surface $\mathscr{S}$ which is the boundary of a compact polyhedron in $\mathbb{R}^3$. The *intrinsic* or *geodesic* distance between two points $x, y \in \mathscr{S}$ is the length of the shortest path joining them. For the intrinsic geometry, a point lying on an edge of the polyhedron is not singular (because a neighbourhood of this point can be unfolded like a hinge). Thus, every point of $\mathscr{S}$ which is not a vertex admits a neighbourhood which is isometric to a disk in the euclidean plane. On the other hand, a vertex admits a
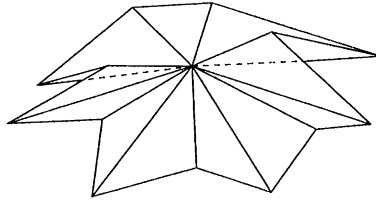
Fig. 2. A point with cone angle greater than 2π.

neighbourhood which is isometric to a piece of an Euclidean cone (the total angle of which may be greater than 2π, see Fig. 2).

Definition 1 (*Piecewise flat surface*). A compact metric space $\mathscr{S}$ is a Euclidean surface with conical singularities, or piecewise flat surface for short, if every point admits a neighbourhood which is isometric to a disk in the Euclidean plane or in a Euclidean cone.

These surfaces have been classified from the point of view of Riemannian geometry in [7].

Let us fix a compact euclidean surface with conical singularities $\mathscr{S}$. We will denote by $V = V_S$ the set of singular points of $\mathscr{S}$. This is a finite set. A *geodesic* on $\mathscr{S}$ is a curve $\gamma : [0, 1] \to S$ which locally minimizes the distances. We will call a *clean arc* a geodesic which is simple (i.e. has no self-intersection) and does not meet the singular set V away from its end points.

Proposition 1. *For any pair of points $p, q \in \mathscr{S}$ and any $L > 0$, the number of geodesic arcs of length $\leqslant L$ joining p and q is finite.*

Proof. Since $\mathscr{S}$ is a flat surface (away from its singularities), two geodesic arcs connecting p to q either must be nonhomotopic or they must together bound a region containing at least one singularity (see also the proof of Corollary 2 in [7]). In particular, if α is a geodesic arc joining p to q, we can find a simply connected neighbourhood U of α in $\mathscr{S}$ such that α is the only geodesic segment from p to q contained in U.

Now let $\mathscr{A}_{pq}(L)$ be the collection of all geodesic arcs $\alpha : [0, 1] \to \mathscr{S}$ parametrized at constant speed joining p to q and of length $\leqslant L$. In particular, the family $\mathscr{A}_{pq}(L)$ is uniformly Lipschitz (with Lipschitz constant L) and by Arzela–Ascoli theorem, it is a compact set (in the uniform topology). But the previous argument shows that each α is an isolated point in $\mathscr{A}_{pq}(L)$, hence $\mathscr{A}_{pq}(L)$ is finite. $\square$

Observe that this argument provides no bound on the number of arcs of length $\leqslant L$ joining p and q.

Corollary 1. *The set of lengths of all geodesics joining pairs of points in V is a discrete subset of $\mathbb{R}$.*

Definition 2. By a geodesic triangulation of the piecewise flat surface $\mathscr{S}$, we mean a finite collection of subsets $T_i \subset \mathscr{S}$ called triangles such that

1. the interiors of the T_i are pairwise disjoint and $\mathscr{S} = \bigcup_i T_i$,
2. the interior of each triangle is isometric to a triangle in $\mathbb{R}^2$,
3. all vertices belong to V_S,
4. V_S does not meet a triangle T_i off its vertices.

Observe that in a geodesic triangulation, each edge is a clean arc.

This definition is a slight abuse of language, for instance, two triangles are allowed to share more than one edge or a triangle may be glued to itself along an edge (such events are classically forbidden, see [4]). However, such triangulations are projections of genuine triangulations on some branch covering of the surface. Existence of geodesic triangulations is proved in [7].

In practice, surface $\mathscr{S}$ is usually given as a collection of triangles in $\mathbb{R}^3$ and is thus naturally equipped with a geodesic triangulation. However, we will have to consider many geodesic triangulations on $\mathscr{S}$. The edges of these triangulations will appear as polygonal lines in $\mathbb{R}^3$ even though they are clean arcs.

Definition 3. The hinge $\Diamond_e$ of an edge e of a geodesic triangulation of $\mathscr{S}$ is the unique pair of triangles incident with e. In case e belongs to only one triangle T_i, the hinge $\Diamond_e$ is the quadrilateron obtained by two isometric copies of T_i glued along the edge e.

Proposition 2. *In an arbitrary geodesic triangulation, every hinge is isometric to a quadrilateron in $\mathbb{R}^2$.*

Proof. A hinge is formed by two triangles T_1 and T_2 with a common edge e. By definition, there exists an isometric embedding $\varphi_i : T_i \rightarrow \mathbb{R}^2$ $(i = 1, 2)$. By composing φ_2 with an isometry of $\mathbb{R}^2$, we may construct an isometric embedding $\varphi : T_1 \cup T_2 \rightarrow \mathbb{R}^2$. $\square$

The image $\varphi(T_1 \cup T_2) \subset \mathbb{R}^2$ is called an *unfolding* or a *development* of the hinge (see Fig. 1).

Definition 4. An edge of a triangulation is legal if its unfolded hinge has two confined triangles (in the sense defined for the plane case).

Definition 5 (*Delaunay triangulation*). A geodesic triangulation of $\mathscr{S}$ is a Delaunay triangulation if all its edges are legal.

As in the planar case, one can show using the notion of universal branch covering, that this local definition implies that no other point of V_S is in the circumscribing circle of a given triangle when folded back to $\mathscr{S}$.
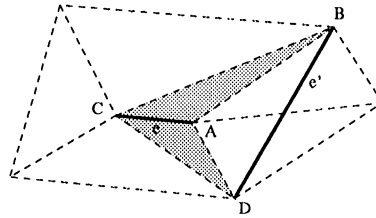
Fig. 3. An edge in a non-convex hinge is always legal. In this case, edge e will not flip.

4. The flip algorithm on piecewise flat surfaces

Let e be an edge of a geodesic triangulation $\mathcal{T}$ of the piecewise flat surface $\mathcal{S}$. Suppose that e is incident with two different triangles of $\mathcal{T}$ and that the development of the hinge $\Diamond_e$, results in a convex quadrilateron $ABCD$. Edge e is a diagonal of $ABCD$ and by replacing it with the other diagonal, one defines a new triangulation $\mathcal{T}'$ of $\mathcal{S}$.

Observe that a flip replace an edge e by another edge e' which is still a clean arc and two triangles ABC and CDA by two other triangles ABC and BCD which are still isometric to triangles in the plane. In short, a flip transforms a geodesic triangulation into a new geodesic triangulation.

Definition 6. This transformation is called the flip of $\mathcal{T}$ at the edge e and we write $\mathcal{T}' := \mathrm{flip}(\mathcal{T}, e)$.

Remark. It would be geometrically problematic to try and flip a triangulation at an edge e whose hinge $\Diamond_e$ is not convex (see Fig. 3). However, we need not worry about such edges as we can prove, as in the planar case, that they are always legal (see also [5]).

The following algorithm takes as input any geodesic triangulation on a piece wise flat surface $\mathcal{S}$ and produces a Delaunay triangulation.

Algorithm (The flip algorithm).
   Input: A geodesic triangulation $\mathcal{T}_0$ on $\mathcal{S}$.
   Output: A Delaunay triangulation $\mathcal{D}$ of $\mathcal{S}$.
   /* Initialisation */
        Set $\mathcal{D} = \mathcal{T}_0$.
        Create a heap H containing the set of edges of $\mathcal{D}$.
   While $H \neq \emptyset$
        Take $e \in H$
        If e is legal
            Remove e from H

Else

> Replace $\mathscr{D}$ by *flip*$(\mathscr{D}, e)$.
> Update H.

End.

The geometric data needed to run the algorithm is the length of all the edges of the initial triangulation $\mathscr{T}_I$, whereas the combinatorial data is the list of vertices, edges and faces.

Given these data, we can actually compute the development of a hinge $\lozenge_e$. If this edge must be flipped, the length of the new edge e' is easily computed and the data are modified accordingly.

The algorithm stops when H is the empty set.

Theorem 1. *This algorithm stops in a finite number of iterations.*

Proof. Let $\mathscr{T}$ be a geodesic triangulation of $\mathscr{S}$. For each triangle T of $\mathscr{T}$, we note $A(T)$ the area of the circumscribing circle of an isometric copy of T in the plane $\mathbb{R}^2$. We can compute $A(T)$ from the three sides a, b, c of T with the classical formula

$$16\pi A = \frac{(abc)^2}{s(s-a)(s-b)(s-c)},$$

where $s = \frac{1}{2}(a+b+c)$ is the semiperimeter of T. Next we define a function F on the set of all geodesic triangulation by

$$F(\mathscr{T}) = \sum_{\mathscr{T}} \mathscr{A}(\mathscr{T}).$$

By Corollary 1, the range of the function F is a discrete subset of $\mathbb{R}$. Now, we can prove as in [5, Theorem 6.1] that if $\mathscr{T}_I, \mathscr{T}_\infty, \mathscr{T}_\in, \dots$ is a sequence triangulations produced by the flip algorithm, then $F(\mathscr{T}_\rangle)$ is a strictly decreasing sequence. Hence it will reach its minimum value in finitely many steps. $\square$

In the planar case, there exists a quadratic bound for this algorithm, see [6]. Here, there is no polynomial bound, since exponentially many clean arcs can exist between two given points and any of them could a priori be legal, hence at some point it may have to be tested by the algorithm.

5. Computing the Delaunay edges and Voronoi cells

To draw the Delaunay triangulation on $\mathscr{S}$, we need to keep track of the ordered list of triangles from the initial triangulation $\mathscr{T}_I$ crossed by an edge. For instance in Fig. 4, the list of triangles crossed by e' is $\{T_1, \dots, T_5\}$.
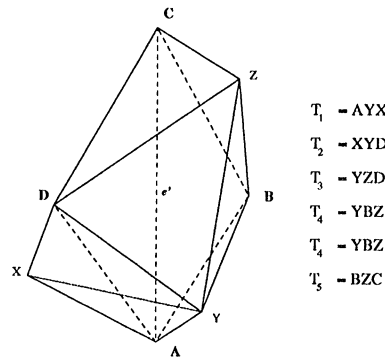
Fig. 4. An edge e crosses a triangle of $\mathcal{S}$ if this triangle crosses the hinge $\Diamond_e$.

This list may be computed from the informations contained in the list of the other edges of $\Diamond_{e'}$. Indeed it is enough to compare the concatenation C_{ABC} of lists of edges $\overline{AB}$ and $\overline{BC}$ with C_{ADC} of lists of edges $\overline{AD}$ and $\overline{DC}$.

We study the different situations where a triangle crosses an edge e'. Any triangle appearing twice in a list will be considered at its first occurrence only. An edge e' crosses a triangle T_i of $\mathcal{T}_I$ if:

1. ($T_i \in C_{ABC}$ or $B \in T_i$) and ($T_i \in C_{ADC}$ or $D \in T_i$),
2. $\overline{AB} \subset T_i$ and $\overline{AC} \subset T_i$,
3. $\overline{BC} \subset T_i$ and $\overline{DC} \subset T_i$.

Note that if e' is an edge of $\mathcal{T}_I$, its list is empty.

Along with the list of triangles transvered by a Delaunay edge, it is useful to determine the points where these edges enter and leave each triangle. This enables us to actually draw the Delaunay triangulation on $\mathcal{S}$ and will also be useful to construct the Voronoi diagram. Each intersection is computed on the development of $\Diamond_{e'}$ and the intersection points are then mapped on the surface.

Definition 7 (*Voronoi cell*). The Voronoi cell of the surface $\mathcal{S}$ associated with the vertex $v \in V_s$ is the set $\{x \in S \mid d(x,v) \leqslant d(x,v') \ \forall v' \in V\}$ where $d(x,v)$ is the geodesic distance between x and v.

Finally let us see how to construct the Voronoi cell associated with a vertex $v \in V$. For each Delaunay edge e incident with v, construct the corresponding Voronoi edge e_v

First unfold the hinge $\Diamond_e$ and then find the intersection I_1 of the perpendicular bisectors of AB, AC and the intersection I_2 of the perpendicular bisectors AC and AD. The line segment $I_1 I_2$ is the Voronoi edge e_v (Fig. 5). By finding the set of all triangles of the original triangulation $\mathcal{T}$ crossed by e_v, we can map e_v on $\mathcal{S}$ as we did before to draw the Delaunay edges.
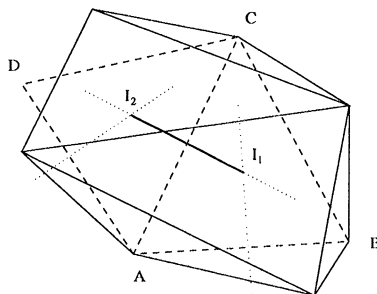
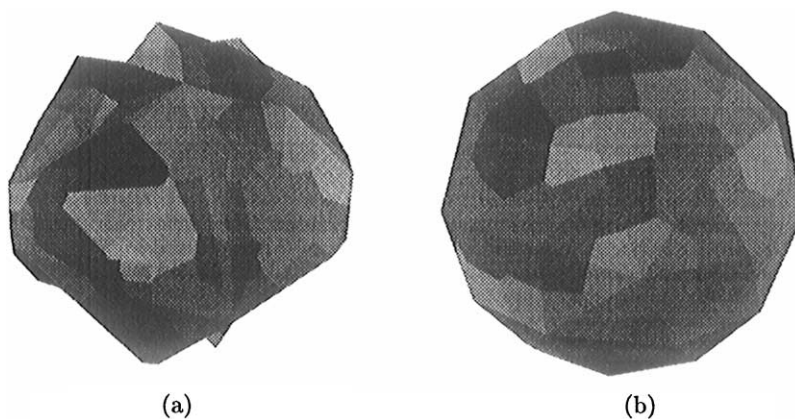Fig. 5. Voronoi edge I_1I_2 of Delaunay edge AC.

Fig. 6. Voronoi diagram on a piecewise flat surface. (a) Non-smoothed surface. (b) Smoothed surface.

6. Application to biological growth

The following section is meant as an illustration of possible modeling applications of the material introduced in this paper. Readers in quest of a more detailed description are invited to consult [2].

We begin by giving a very succinct description of a particular type of biological membrane growth process, namely that of hyphal walls of some mycelia (Fig. 6). It is followed by a brief presentation of the associated mathematical model, which was built using the structures of the preceding sections.

The biological tissues that constitute the hyphal walls of some mycelia are very thin surfaces that can be thought of as being made up of individual cell-like patches. Such cells are created whenever one of the numerous wall material vesicles floating within the hypha hits the membrane and makes room for itself there. All along this process, not unlike an inflating balloon, the expanding membrane is subject to surface tension and internal pressure with a convexifying effect, the fundamental difference being that a balloon has to do it with a constant amount of mass whereas membrane mass can
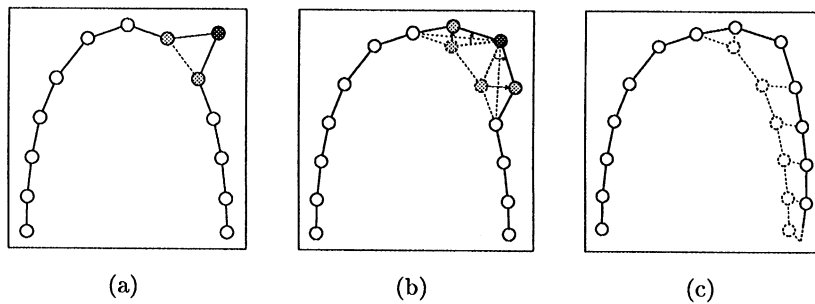
Fig. 7. Insertion of new cells in a wall modeled by a chain of cell and motion of all the cells to keep the convexity. (a) A new cell c is added between two other members of the chain. (b) Internal pressure pushes outside the two neighbours of c. (c) By recurrence, all the cells move away and the wall becomes convex again.

increase. An important element of hyphal growth the *Spitzenkörper*, which apparently plays a role in attracting the vesicles to the hyphal tip. In fact, hyphal branching coincides with the creation of a new *Spitzenkörper*. The process has been modeled in the plane [2], where the hyphal wall is a (closed) broken line. When a new vesicle arrives, it places itself between two existing ones and the membrane deforms itself under the effect of the forces mentioned before, expanding to take on a locally convex shape (Fig. 7). Our model generalizes this to three dimensions.

At any given time, the membrane is thought of as a piecewise flat closed surface, with triangular facets and homeomorphic to the sphere. Its cells are identified with those of the Voronoi partition of the surface induced by its vertices. The surface area is locally minimal, a property that is brought about, if necessary, by a number of flip operations of the triangular facets. These operations change the topography of the surface, as well as the associated Voronoi partition, which therefore has to be kept being updated. This property maintains a rather "smooth" surface and thus a closer representation of its biological model (Figs. 6(a) and (b)). When a new vesicle hits the membrane at some triangular facet, a new vertex of the surface is created along with four new facets and a new Voronoi cell, and the process continues. What remains to be given is a description of the motion of the vertices of the surface.

6.1. Surface deformation, motion of cells

The cells move and deform themselves to minimize surface tension. A cell models a physical surface element with a given area and an approximately circular shape. A cell can be stretched in all directions but will always try to return to its previous shape. Assuming that the "ideal" cell has a given shape (for example, circular) and a given area, one can find the corresponding "ideal" distance δ between the "centers" (i.e. the generators) of two cells. If the distance between two cells is smaller than δ, the cells will try to move away from each other. If the distance is longer, they will try to move closer. To model the driving forces, one can imagine a set of edges behaving
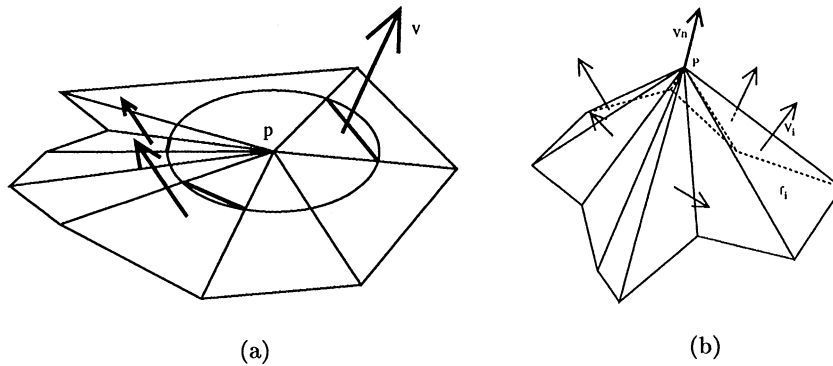
Fig. 8. Computation of a normal vector. (a) The weight of a "normal" vector $\vec{v}$ depends on the angle α of the two edges of the facet adjacent to P. (b) The "normal" vector of the surface at a point p_i is computed by adding the normal vector $\vec{v}_i$ of every facet adjacent to p_i.

like springs between the cells. The springs are given with their normal length δ. Each length modification induces a force on the two extremities of a spring and the resultant of the forces gives the motion direction of each cell.

Since a cell directly interacts only with its neighbours, it is reasonable to assume our springs to follow the edges of the Delaunay triangulation. Using the flip algorithm, we are able to efficiently locate the springs (i.e. the local cell interactions). Motion of these generators is split into two components:

- A "tangential" motion generated by tensions between the cells. The corresponding driving force is the projection on a supporting plane of the resultant of the spring forces acting on the generator. The normal vector to the supporting is a weighted sum of normal vectors to the faces incident with the generator. The weighing takes into account the angles at the generator of the various triangle (Fig. 8).
- A "normal" expansive motion, essentially due to inner pressure. The driving force has the direction of the normal vector to the supporting plane described above and is composed of the projection of the resultant of spring forces plus a contribution from internal pressure.

The brief description did not get into any detail in how the vesicles are generated, nor on branching, the creation of new *Spitzenkörper*. All these elements were taken into account in a model that was implemented using a Silicon Graphics work station. It produces images like those in Figs. 6 and 10.

Note that, not surprisingly, the filament structures obtained with this model have a very similar morphogenesis to that of the early development stages of real mycelia, as can be seen by comparing Figs. 9(a) (*Mucor spinosus*) and (b) (simulation). A closer look at the tip of a hypha tends to show a similar organization of the vesicles in both the model (Fig. 10) and the real mycelium. That could be one of the factors influencing morphogenesis, but further investigation on real mycelia is needed to confirm this.
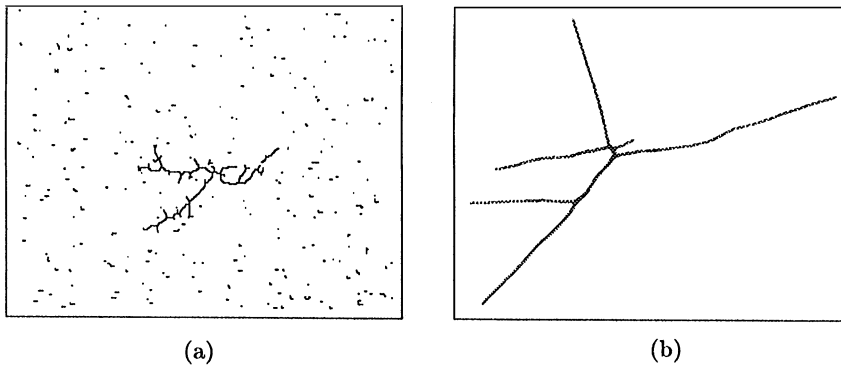
Fig. 9. Real and modeled growth of a mycelium. (a) Picture of a real mycelium (size: 6.5 mm × 4.5 mm) (b) Resultant of a membrance growth by insertion of new cells to several apices at regulated time interval.
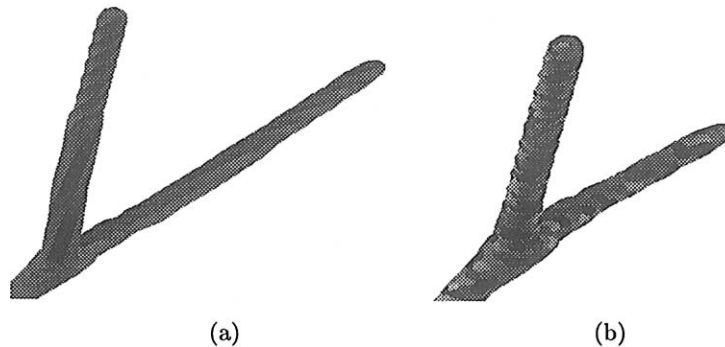

Fig. 10. Representation of the extremitites of two hyphae. (a) This apex was generated by the insertion of a set of new cells in the same region of the surface. (b) Modeling of the cellular membrane with Voronoi cells applied on $\mathscr{S}$.

Acknowledgements

References

[1] H. Edelsbrunner, Algorithms in Combinatorial Geometry, Springer, Berlin, 1987.
[2] C. Indermitte, Modélisation et Simulation de la Croissance d'un Mycélium, Ph.D. Thesis, Swiss Federal Institute of Technology, 1995.

[3] A. Okabe, B. Boots, K. Sugihara, Spatial Tesselations Concepts and Applications of Voronoi Diagrams, Wiley, New York, 1992.

[4] J. Stillwell, Classical Topology and Combinatorial Group Theory, Springer, New York, 1980.

[5] H. Telley, Modélisation et Simulation Bidimensionnelle de la Croissance des Polycristaux, Ph.D. Thesis, Swiss Federal Institute of Technology, 1989.

[6] H. Telley, Delaunay triangulation in the flat torus, Technical Report, Swiss Federal Institute of Technology, September 1992.

[7] M. Troyanov, Les surfaces euclidiennes à singularité coniques, Ens. Math. (1986) 79–94.