

Anisotropic Diagrams: Labelle Shewchuk approach revisited

Jean-Daniel Boissonnat, Camille Wormser, Mariette Yvinec

► **To cite this version:**

Jean-Daniel Boissonnat, Camille Wormser, Mariette Yvinec. Anisotropic Diagrams: Labelle Shewchuk approach revisited. [Research Report] RR-5741, INRIA. 2006, pp.23. inria-00070277

HAL Id: inria-00070277

<https://hal.inria.fr/inria-00070277>

Submitted on 19 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Anisotropic Diagrams: Labelle Shewchuk approach
revisited*

Jean-Daniel Boissonnat — Camille Wormser — Mariette Yvinec

N° 5741

Novembre 2005

Thème SYM



*R*apport
de recherche



Anisotropic Diagrams: Labelle Shewchuk approach revisited

Jean-Daniel Boissonnat^{*}, Camille Wormser[†], Mariette Yvinec[‡]

Thème SYM — Systèmes symboliques
Projet Geometrica

Rapport de recherche n° 5741 — Novembre 2005 — 23 pages

Abstract: F. Labelle and J. Shewchuk [LS03] have proposed a discrete definition of anisotropic Voronoi diagrams. These diagrams are parametrized by a metric field. Under mild hypotheses on the metric field, such Voronoi diagrams can be refined so that their dual is a triangulation, with elements shaped according to the specified anisotropic metric field.

We propose an alternative view of the construction of these diagrams and a variant of Labelle and Shewchuk's meshing algorithm. This variant computes the Voronoi vertices using a higher dimensional power diagram and refines the diagram as long as dual triangles overlap. We see this variant as a first step toward a 3-dimensional anisotropic meshing algorithm.

Key-words: anisotropic Voronoi diagram, anisotropic meshing

Work partially supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-006413 (ACS - Algorithms for Complex Shapes)

Work partially supported by the European Network of Excellence AIM@shape (FP6 IST NoE 506766).

^{*} INRIA, BP 93 06902 Sophia Antipolis, France

[†] INRIA, BP 93 06902 Sophia Antipolis, France

[‡] INRIA, BP 93 06902 Sophia Antipolis, France

Maillages anisotropes : approche de Labelle et Shewchuk revisitée

Résumé : F. Labelle et J. Shewchuk [LS03] ont proposé une définition discrète des diagrammes de Voronoï anisotropes. Ces diagrammes sont paramétrés par un champ de métriques. Sous certaines hypothèses sur ce champ de métriques, ils peuvent être raffinés de telle sorte que leur graphe dual devienne une triangulation dont les éléments sont aussi réguliers que possible, au sens du champ de métriques considéré.

Nous proposons une autre façon de construire ces diagrammes ainsi qu'une variante de l'algorithme de maillage de Labelle et Shewchuk. Cette variante calcule les sommets de Voronoi par le biais d'un diagramme de puissance en plus grande dimension, puis raffine le diagramme tant que des triangles se chevauchent. Une telle approche doit être un premier pas vers un algorithme de maillage anisotrope en dimension 3.

Mots-clés : diagrammes de Voronoï anisotropes, maillages anisotropes

1 Introduction

Anisotropic meshes are triangulations of a given domain in the plane or in higher dimension, with elements elongated along prescribed directions. Anisotropic triangulations have been shown [She02] to be particularly well suited for interpolation of functions or numerical modeling. They allow to minimize the number of triangles in the mesh while retaining a good accuracy in computations. For such applications, the elongation directions are usually given as quadratic forms at each point. These directions may be related to the curvature of the function to be interpolated, or to some specific directions taken into account in the equations to be solved.

Various heuristic solutions for generating anisotropic meshes have been proposed. Li et al. [LTÜ99] and Shimada et al. [SYI97] use packing methods. Bossen and Heckbert [BH96] use a method consisting in centroidal smoothing, retriangulating and inserting or removing sites. Borouchaki et al. [BGH⁺97] adapt the classical Delaunay refinement algorithm to the case of an anisotropic metric.

Recently, Labelle and Shewchuk [LS03] have settled the foundations for a rigorous approach based on the so-called anisotropic Voronoi diagrams. We propose an alternative view of the construction of these diagrams and a variant of the meshing algorithm of Labelle and Shewchuk. This variant computes the Voronoi vertices using a higher dimensional power diagram and refines the diagram as long as dual triangles overlap.

2 Labelle and Shewchuk's approach

2.1 Anisotropic Voronoi diagram definition

Labelle and Shewchuk [LS03] have proposed a discrete definition of anisotropic Voronoi diagrams. This section presents the basis of their work. The diagram is defined over a domain $\Omega \subset \mathbb{R}^d$, and each point $p \in \Omega$ has an associated metric. More specifically, a point p is given a symmetric positive definite quadratic form represented by a $d \times d$ matrix M_p . The *distance* between two points x and y *as viewed by* p is defined as

$$d_p(x, y) = \sqrt{(x - y)^t M_p (x - y)},$$

and the distance between p and q is defined as $d(p, q) = \min(d_p(p, q), d_q(p, q))$. Note that d_p is a distance, whereas d is not, since it does not necessarily verifies the triangular inequality.

In a similar way, the *angle* $\angle_p xqy$ *as viewed by* p is defined as

$$\angle_p xqy = \arccos \frac{(x - q)^t M_p (y - q)}{d_p(x, q) d_p(y, q)}.$$

In order to compare the metric at points p and q , a transfer application is needed. Given the quadratic form M_p of a point p , we denote by F_p a matrix such that $\det(F_p) > 0$ and

$F_p^t F_p = M_p$. Then $d_p(x, y) = \|F_p(x - y)\|_2$ and the transfer application from p to q is

$$F_{p,q} = F_q F_p^{-1}.$$

This application $F_{p,q}$ is in fact an isometry between the metric spaces (\mathbb{R}^d, M_p) and (\mathbb{R}^d, M_q) . The *distortion* between p and q is then defined as $\gamma(p, q) = \gamma(q, p) = \max\{\|F_{p,q}\|_2, \|F_{q,p}\|_2\}$. For any points x, y , we have $1/\gamma(p, q) d_q(x, y) \leq d_p(x, y) \leq \gamma(p, q) d_q(x, y)$.

Labelle and Shewchuk [LS03] define the anisotropic Voronoi diagram in the following way:

Definition 2.1 *Let S be a set of points, called sites hereafter. The Voronoi cell of a site p in S is*

$$\text{Vor}(p) = \{x \in \mathbb{R}^d : d_p(p, x) \leq d_q(q, x) \text{ for all } q \in S\}.$$

Any subset $R \subset S$ induces a Voronoi face $\text{Vor}(R) = \bigcap_{q \in R} \text{Vor}(q)$ of points equally close to the sites in R and no closer to any others. If not empty, such a face has dimensionality $\dim(\text{Vor}(R)) \geq d + 1 - |R|$, achieving equality if the sites are in general position. The anisotropic Voronoi diagram of S is the arrangement of the Voronoi faces $\{\text{Vor}(R) : R \subset S, R \neq \emptyset, \text{Vor}(R) \neq \emptyset\}$.

It should be noted that

- each site is in the topological interior of its cell, which has dimensionality d ;
- the bisectors are quadric surfaces (conic curves in dimension 2);
- the Voronoi faces are not always connected.

For brevity, we use in the sequel the term *k-Vface* to name Voronoi faces that have dimensionality k . The *label* of a Vface $\text{Vor}(R)$ is the set R . As noted, faces are not necessarily connected. In particular, a 0-Vface is not necessarily a unique point, but may consist of several ones. We call each of these points a *Voronoi vertex*.

For any diagram D , and any domain Ω , we denote by D_Ω the diagram D restricted to Ω , i.e. the diagram obtained by intersecting the cells of D with Ω .

Definition 2.2 *The dual complex of the anisotropic Voronoi diagram of S is the simplicial complex whose set of vertices is the set S , with a simplex associated to each subset $R \subset S$ such that $\text{Vor}(R) \neq \emptyset$. We associate to each one of those simplices a geometric simplex, its canonical linear embedding in \mathbb{R}^d . The set of those geometric simplices, with their incidence relations, is called the geometric dual of the anisotropic Voronoi diagram of S .*

In two dimensions and with points in generic position, the geometric dual includes, for each Voronoi vertex v , a dual triangle whose vertices are the three sites that compose the label of v . There is no reason why these triangles should form a triangulation. The two issues to be considered are the combinatorial planarity of the graph, which depends on the connectivity of the cells, and the ability to stretch its edges without crossing, which depends on the curvature of the bisectors.

The goal of the meshing algorithm is to refine the anisotropic Voronoi diagram by inserting new sites, so that its geometric dual becomes a triangulation, with well-shaped triangles.

2.2 The wedge property

This section summarizes some results presented in [LS03]: Labelle and Shewchuk have introduced the wedge property and have proven the following results to ensure that their algorithm converges to a triangulation.

Definition 2.3 *The wedge between two sites p and q is the locus of points x such that the angle $\angle_p x p q$ and the angle $\angle_q x q p$ are less than $\pi/2$, or equivalently $d_p(x, q)^2 \leq d_p(p, x)^2 + d_p(p, q)^2$ and $d_q(x, p)^2 \leq d_q(q, x)^2 + d_q(p, q)^2$.*

A k -Vface f , with $k < d$, is said to be wedged if, for any pair p, q of distinct sites such that $f \subset \text{Vor}(p) \cap \text{Vor}(q)$, we have $f \subset \text{wedge}(p, q)$.

Theorem 2.4 *If every subface of a d -Vface $\text{Vor}(p)$ is wedged, then the d -Vface is star-shaped around p .*

The following lemma is only valid in the two-dimensional case.

Lemma 2.5 *Let v be a Voronoi vertex labeled by the sites p, q and r . If v is wedged, then the orientation of the triangle pqr matches the ordering of the cells $\text{Vor}(p), \text{Vor}(q), \text{Vor}(r)$ locally around v .*

Let Ω be a polygonal domain of the plane and S be a set of sites in Ω that includes every vertex of Ω . We denote by D the anisotropic Voronoi diagram of S and D_Ω its restriction to Ω . The following result is central to the proof of the correctness of Labelle and Shewchuk's algorithm.

Theorem 2.6 *Suppose that each 1-Vface that intersects the boundary $\partial\Omega$ intersects a single edge of $\partial\Omega$ and that each edge of $\partial\Omega$ is intersected exactly once. If all the 1-Vfaces and vertices of D_Ω are wedged, then the geometric dual graph of D_Ω is a triangulation of Ω if S is in general position.*

If S is not in general position, the geometric dual is a polygonalization of Ω with strictly convex polygons. Labelle and Shewchuk represent the Voronoi diagram as the lower envelop of a set of paraboloids. When inserting a new site, this lower envelop is updated in a lazy way, which amounts to computing only the connected component of the cell that contains the new site. Theorem 2.6 validates their lazy computation of the diagram¹.

Labelle and Shewchuk's algorithm consists in incrementally inserting points

- on non-wedged Voronoi edges;
- at the center of triangles that are badly shaped, or are too large, or do not have the same orientation as the three Voronoi cells around their dual Voronoi vertices.

¹In fact, there is a slight imprecision in their claim about the triangulation output by their algorithm: since the algorithm never checks the wedge property for Voronoi edges that have not been computed, it does not ensure that no disconnected cell remains in the complete diagram.

3 Power Diagram and Anisotropic Voronoi Diagram

In this section, we reduce the construction of an anisotropic Voronoi diagram in \mathbb{R}^d to the computation of a power diagram in \mathbb{R}^D where $D = d(d+3)/2$ and its restriction to a d -manifold.

Definition 3.1 *A power diagram is defined for a set of spheres. Given a sphere σ centered at y and of radius r , the power of a point x with respect to σ is defined as*

$$\pi_\sigma(x) = \|x - y\|^2 - r^2$$

where $\|\cdot\|$ is the euclidean distance. The power diagram of a set of hyperspheres Σ of \mathbb{R}^D is the subdivision induced by the power cells of the spheres in Σ , where the power cell $Pow(\sigma)$ of a sphere σ is the locus of points with a smaller power with respect to σ than to any other sphere in Σ :

$$Pow(\sigma) = \{x \in \mathbb{R}^D, \pi_\sigma(x) \leq \pi_\tau(x), \forall \tau \in \Sigma\}.$$

We define the power cell of a family of spheres $\{\sigma_i\}_i$ as

$$Pow(\{\sigma_i\}_i) = \cap_i Pow(\sigma_i).$$

Let $D = \frac{d(d+3)}{2}$. Associate to each point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$

- the point $\tilde{x} = (x_r, x_s, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}}$;
- the point $\hat{x} = (x, \tilde{x}) \in \mathcal{P} \subset \mathbb{R}^D$.

where \mathcal{P} denotes the d -manifold of \mathbb{R}^D $\{\hat{x} \in \mathbb{R}^D, x \in \mathbb{R}^d\}$. Let $S = \{p_1, \dots, p_n\}$ be a finite set of sites in \mathbb{R}^d . To each point p_i of S we attach a symmetric positive definite matrix M_i and we define

- the point $q_i = (q^{r,s}, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}}$ defined as
 - $q^{r,r} = -\frac{1}{2}M_i^{r,r}$, for $1 \leq r \leq d$;
 - $q^{r,s} = -M_i^{r,s}$, for $1 \leq r < s \leq d$.
- the point $\hat{p}_i = (M_i p_i, q_i)$;
- the sphere $\sigma(p_i)$ of center \hat{p}_i and radius $\sqrt{\|\hat{p}_i\|^2 - p_i^t M_i p_i}$, where $\|\cdot\|$ denotes the euclidean norm.

Let Π be the projection $(x, \tilde{x}) \in \mathbb{R}^D \mapsto x \in \mathbb{R}^d$. Let Σ be the set of spheres $\{\sigma(p), p \in S\}$.

Lemma 3.2 *The anisotropic Voronoi diagram of $S \subset \mathbb{R}^d$ is the image by Π of the restriction of the D -power diagram of Σ to the d -manifold \mathcal{P} .*

Proof We have the following inequalities:

$$\begin{aligned} d_{p_i}(x, p_i)^2 &= x^t M_i x - 2p_i^t M_i x + p_i^t M_i p_i \\ &= -2q_i^t \tilde{x} - 2p_i^t M_i x + p_i^t M_i p_i \\ &= -2\hat{p}_i^t \hat{x} + p_i^t M_i p_i \end{aligned}$$

This implies that $d_{p_i}(x, p_i) < d_{p_j}(x, p_j)$ if and only if

$$\|\hat{x} - \hat{p}_i\|^2 - (\|\hat{p}_i\|^2 - p_i^t M_i p_i) < \|\hat{x} - \hat{p}_j\|^2 - (\|\hat{p}_j\|^2 - p_j^t M_j p_j)$$

It follows that x is closer to p_i than to p_j if and only if the power of \hat{x} with respect to σ_i is smaller than its power with respect to σ_j . This proves that, for a point $z \in \mathcal{P}$, being in the power cell of σ_i is equivalent to $\Pi(z)$ being in the Voronoi cell of p_i . \square

The previous lemma gives a construction of the anisotropic Voronoi diagram. As is well-known, computing a power diagram in \mathbb{R}^D reduces to computing a lower convex hull in \mathbb{R}^{D+1} . Hence, in the two-dimensional case, the computation of a six-dimensional convex hull is needed. To get the anisotropic Voronoi diagram, it remains to compute the intersection of the power diagram with the manifold \mathcal{P} . We detail the computations required by our algorithm in the following section,

4 Required computations

4.1 Computations of Voronoi vertices

Computing the complete anisotropic Voronoi diagram explicitly is not easy. However, our meshing algorithm only requires to compute Voronoi vertices. We now explain how to compute these vertices, in the two-dimensional case. Recall that a 0-Vface of \mathbb{R}^2 may be seen as the projection of a finite subset of \mathbb{R}^5 . This set is obtained as the intersection of a linear subspace of codimension 2 (obtained as the intersection of three cells of the power diagram of Σ) with the 2-dimensional manifold \mathcal{P} (see lemma 3.2).

The computation of the Voronoi vertices whose label is $\{a, b, c\}$ consists of the following steps:

- (1) Consider three sites a, b and c such that $\sigma(a), \sigma(b), \sigma(c)$ correspond to a triangle in the regular complex dual to the power diagram, i.e. such that their cells have a common non-empty intersection.
- (2) Compute the hyperplane H_{ab} , which is the bisector of $\sigma(a)$ and $\sigma(b)$, and the hyperplane H_{bc} , which is the bisector of $\sigma(b)$ and $\sigma(c)$, and then their intersections D_{ab} and D_{bc} with \mathcal{P} . Practically, D_{ab} and D_{bc} are represented by their projections by Π , named

respectively C_{ab} and C_{bc} . The curves C_{ab} and C_{bc} are conics of \mathbb{R}^2 , and the equation of C_{ab} in \mathbb{R}^2 is:

$$(x^t M_a x - 2a^t M_a x + a^t M_a a) - (x^t M_b x - 2b^t M_b x + b^t M_b b) = 0$$

We denote this equation by $C_{ab}(x) = 0$. The equation of C_{bc} is obtained similarly.

- (3) Compute the intersection points of C_{ab} and C_{bc} . This intersection is the set of Voronoi vertices whose label is $\{a, b, c\}$ in the Voronoi diagram of $\{a, b, c\}$.
- (4) In the diagram of S (i.e. not only three sites), some of the elements of $C_{ab} \cap C_{bc}$ are not Voronoi vertices because they belong to the cell of a closer site. Accordingly, in dimension D , the linear subspace $H_{ab} \cap H_{bc}$ may intersect the power cells of some other sphere $\sigma(x)$ for $x \in S \setminus \{a, b, c\}$. The pre-image by Π of a point z of $C_{ab} \cap C_{bc}$ lies on $H_{ab} \cap H_{bc}$ in \mathbb{R}^D . It belongs to the power cell $Pow(\{\sigma(a), \sigma(b), \sigma(c)\})$ if and only if it has a smaller power to $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$ than to any other $\sigma(d)$ in Σ . We do not have to check this fact for all the other spheres $\sigma(d)$ with $d \in S$, but only for the spheres whose cell is incident to the cells of $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$, since the cells of a power diagram are always connected. We realize this computation after projecting onto the plane: among the points z of $C_{ab} \cap C_{bc}$, we keep the ones such that for each tetrahedron of the regular complex defined by $\sigma(a)\sigma(b)\sigma(c)\sigma(f)$, the inequalities $C_{af}(z) < 0$, $C_{bf}(z) < 0$ and $C_{cf}(z) < 0$ are verified (it should be noted that those three inequalities are equivalent, since z has the same power with respect to any of the three spheres $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$). The points that we keep are the Voronoi vertices labeled by $\{a, b, c\}$.

4.2 Computations of encroachment

Our algorithm takes as input a set of segments which are required to appear in the final triangulation. These segments are called *constrained segments*. They may be refined during the algorithm, by the insertion of sites located on them. In such a case, the different parts delimited by the sites inserted on the constrained segment are called *constrained subsegments*.

Most notably, among them are the boundaries of the domain we want to triangulate. We now present how to compute the classical property of *encroachment* of a constrained subsegment.

Definition 4.1 *A constrained subsegment $e = (a, b)$ is encroached by a point $p \notin \{a, b\}$ if $\text{Vor}(p) \cap [a, b] \neq \emptyset$ in the Voronoi diagram of $\{a, b, p\}$.*

During the algorithm, we need to compute whether a constrained subsegment $e = (a, b)$, that was not previously encroached, is encroached by a point p to be inserted.

First note that, when inserting a site p , we have a small set of potentially encroached edges: among the constrained subsegments, it is sufficient to consider the ones that would have both of their endpoints joined to p in the geometric dual, if p were inserted in the diagram. Indeed, if p encroaches e , the cell $\text{Vor}(p)$ is adjacent to the cells of both endpoints

of e . Let $e = [a, b]$ be such a constrained subsegment. Then, let E be the intersection $C_{pa} \cap [a, b]$ of the bisector of p and a with $[a, b]$. If any of the $z \in E$ verifies $C_{pb}(z) < 0$, we have in fact $z \in \text{Vor}(\{a, p\}) \cap [a, b]$ and $\text{Vor}(p)$ intersects $[a, b]$.

A constrained subsegment may also completely disappear from the dual when a site p is inserted. Such a segment is obviously encroached by p .

5 Description of the Algorithm

From now on, let Ω be a polygonal domain of the plane, whose boundary is denoted by $\partial\Omega$. We denote by C the set of constrained subsegments and by S a finite set of sites in Ω . The set C is updated during the course of the algorithm to reflect the fact that some constrained segments have been refined into constrained subsegments. At the beginning, we assume that the edges of $\partial\Omega$ belong to C and that the vertices of $\partial\Omega$ belong to S . Refining the Voronoi diagram consists in adding sites to the set S . We assume that the quadratic form associated to any point of Ω can be obtained.

5.1 Local Embedding

We have seen in the previous section how to compute the Voronoi vertices. If the label of a vertex v is $\{a, b, c\}$, the triangle abc is called the *dual triangle* of v . We now introduce some properties that will ensure that the dual triangles define a triangulation of the domain they cover.

We consider a set of non-degenerate triangles T such that

- (i) the set of vertices of the triangles in T is exactly S ;
- (ii) each edge on $\partial\Omega$ is the edge of exactly one triangle in T ;
- (iii) if e is the edge of some triangle in T and is not an edge on $\partial\Omega$, e belongs to exactly two triangles in T , which do not overlap².

We prove that under those assumptions, T is a triangulation of Ω .

Definition 5.1 Let $p \in S$ be one of the sites and T_p be the set of triangles incident to p . Two triangles are said to be adjacent if they share an edge. The equivalence classes for the transitive closure of the adjacency relation in T_p are called the *umbrellas* of p .

The link $\text{link}(p)$ of a site p is the set of edges opposite to p in all the triangles of T_p .

Lemma 5.2 If the finite set of triangles T verifies Rules (i), (ii) and (iii), we claim that:

- (a) all the triangles in T are inside Ω ;

²Since all triangles are non-degenerate, the overlapping is well-defined.




Figure 1: Two umbrellas (left) and one umbrella winding twice (right) around p

- (b) if p is an internal site, its umbrellas are combinatorial disks and p is inside each of its umbrellas;
- (c) if p is a vertex of $\partial\Omega$, p has a unique umbrella, and p is on the boundary of this umbrella. Furthermore, this umbrella is embedded, i.e. its triangles do not overlap one another.

Proof For the sake of simplicity, we prove the result under the hypothesis that Ω is simply connected. The result is still true without this hypothesis. However the proof would be more complicated.

- (a) We consider an edge e of the boundary of the union U of all the triangles. From Rules (ii) and (iii), e has to be an edge of $\partial\Omega$. Thus the boundary of U is included in the boundary $\partial\Omega$. Assume for a contradiction that U is not included in the simply connected domain Ω and consider $x \in U \setminus \Omega$. Let π be a path in ∂U such that x is in the bounded component delimited by π . Since $\partial U \subset \partial\Omega$, π is in $\partial\Omega$. However, π is not contractible in Ω . This contradicts Ω being simply connected.
- (b) If p is an internal site, Rule (iii) implies that $link(p)$ is a union of closed polygonal curves (non necessarily simple curves), since Rule (iii) prevent any vertex of degree different from 2 to appear on the link. An umbrella is then obtained by choosing one of those closed curves, and linking p to every vertex of it. This proves that an umbrella is a combinatorial disk, since it has a combinatorial circle as boundary.

Consider the union of the triangles of an umbrella. Assume for a contradiction that p is not in the interior of this union. Then it is on the boundary and each edge of this boundary incident to p belongs to two triangles of the umbrella which overlap, which contradicts Rule (iii).

- (c) If p is a vertex of $\partial\Omega$, $\ell(p)$ may a priori contain some closed curves and some curves joining the two neighbors of p on $\partial\Omega$. As seen in (b), the closed curves have to enclose p . Thanks to (a) and to the fact that p is on $\partial\Omega$, this is not possible. Therefore, the link of vertex p cannot include a closed curve. Rule (ii) then implies that all curves in

$\ell(p)$ have the same first and last segment and because rule (iii) prevents any branching vertex in $\ell(p)$, the link $\ell(p)$ is a single curve. The fact that the triangles of the unique umbrella do not overlap follows from (iii) too.

□

Theorem 5.3 *Under assumptions (i), (ii), (iii), T is a triangulation of Ω .*

Proof A priori, an internal site may have multiple umbrellas and each of those umbrellas may wind more than once around p . To prove that T is a triangulation, we now glue the triangles of T along their common edges and vertices to build a surface: we denote by $\mathcal{T} = \{(x, t) \in \Omega \times T \mid x \in t\}$ the set of points associated to the triangles they belong to, and we define on \mathcal{T} the equivalence relation \sim by setting $(x, t) \sim (x', t')$ if $x = x'$, $x \in \partial t$ and $x' \in \partial t'$, so that taking the quotient of the set \mathcal{T} by the equivalence relation \sim amounts to gluing the common edges and vertices. The final glued space is denoted by $\mathcal{G} = \mathcal{T} / \sim$.

Let $h : (x, t) \in \mathcal{G} \mapsto x$ be the first projection, mapping \mathcal{G} to Ω . The correctness of the triangulation is equivalent to h being a homeomorphism. Let Ω_p be the punctured space obtained by removing from Ω the vertices of the triangles of T , and let \mathcal{G}_p be $h^{-1}(\Omega_p)$.

From assumption (iii), the restriction h_p of h to \mathcal{G}_p is a local homeomorphism. Thus h_p is a cover of Ω_p . As the points close to $\partial\Omega$ have only one pre-image, from assumption (ii), $h_p : \mathcal{G}_p \rightarrow \Omega_p$ has only one sheet and is in fact a homeomorphism.

This shows that each site p has a unique umbrella, which is well embedded and that h_p may be extended to \mathcal{G} as an homeomorphism. Thus, Ω is triangulated by T . □

5.2 Form Criterion

Let v be a Voronoi vertex of an anisotropic Voronoi diagram. Its label consists of three sites that form a dual triangle $t_v = abc$. The *radius* of v is $r(v) = d_a(a, v) = d_b(b, v) = d_c(c, v)$ (we define the radius of the center instead of the radius of the triangle, because the triangle may have multiple centers). The length of an edge is the distance between its endpoints as defined in section 2.1. We denote the shortest edge of t_v by $\delta(t_v)$. The *radius-edge ratio* of v is $\beta(v) = r(v)/\delta(t_v)$.

For a given *shape bound* B , a vertex v is considered to be *badly-shaped* if $\beta(v) > B$.

5.3 Refinement Algorithm

In this section, we present our algorithm, which refines an anisotropic Voronoi diagram V until the set of triangles dual to the Voronoi vertices of the restriction V_Ω of V to Ω have a good shape and satisfy conditions (i), (ii) and (iii), which ensure that the triangles form a triangulation of Ω .

First recall that, thanks to the monotonicity of the distance function associated to each point, there is always a unique point on a line segment that is equidistant from both endpoints.

Definition 5.4 *Assume that a constrained subsegment $e = (p, q)$ is encroached. The breakpoint of the edge (p, q) is defined as the point of $[p, q] \setminus (\text{Vor}(p) \cup \text{Vor}(q))$ that is as close to the midpoint of $[p, q]$ as possible. By midpoint, we mean the intersection of $[p, q]$ with the bisector of p and q , i.e. the point z of $[p, q]$ such that $d_p(p, z) = d_q(q, z)$.*

We now present our refinement algorithm. We are given a shape bound B . At each step of the algorithm, we maintain the set T of dual triangles, obtained as the labels of the computed Voronoi vertices that are inside Ω (see section 4.1).

We define a procedure of conditional insertion, needed for the presentation of the algorithm:

CONDITIONALLY INSERT(x):

- if x encroaches no constrained subsegment, insert x ;
- if x encroaches some constrained subsegment e , insert a site on the breakpoint of e instead.

The algorithm inserts points iteratively, applying the following rules. Rule i is applied only if no rule j with $j < i$ applies:

- Rule (1) if some constrained subsegment $e \in C$ does not appear as the edge of a dual triangle because it is encroached, insert a site located at the breakpoint of edge e ;
- Rule (2) if some constrained subsegment $e \in C$ does not appear as the edge of a dual triangle, because its dual Vface is a complete ellipse, denote by Δ the support line of e . Then conditionally insert a site located at the intersection of $\Delta \setminus e$ with the ellipse;
- Rule (3) if a Voronoi vertex v is badly shaped (see section 5.2), conditionally insert a site located at that vertex;
- Rule (4) if a triangle abc is the dual of several Voronoi vertices, conditionally insert a site located at the vertex that is the furthest from a , b and c (consider the largest of the three distances);
- Rule (5) if two triangles sharing an edge overlap, conditionally insert a site at the dual Voronoi vertex of one of them: choose the triangle which contains the edge (x, y) such that $\gamma(x, y)$ is maximal.

We will now prove that if the algorithm terminates, Conditions (i), (ii) and (iii) of section 5.1 are verified. By Theorem 5.3, the geometric dual is therefore a triangulation, without any badly-shaped vertex.

Lemma 5.5 *Upon termination of the algorithm, the dual triangles in T form a triangulation of the domain Ω and all the constrained subsegments appear in this triangulation.*

Proof First, let us prove that each constrained subsegment is incident to least one triangle in T . Consider some constrained subsegment s with endpoints a and b .

- Thanks to Rule (1), s is not encroached and therefore lies in the union of the cells of its endpoints.
- Since each site lies in its own cell, s cannot be included in one cell only. This proves that the dual edge $\text{Vor}(\{a, b\})$ is not empty and intersects s and the domain Ω .
- If the bisector of a and b is an ellipse, Rule (2) implies that the Voronoi edge $\text{Vor}(\{a, b\})$ has endpoints within Ω . In all cases, observe that, owing to the monotonicity of the distance $d_a(a, x)$ along ab , $\text{Vor}(\{a, b\})$ intersects s in at most one point. Furthermore $\text{Vor}(\{a, b\})$ cannot intersect any other constrained subsegment because they are not encroached either. Therefore $\text{Vor}(\{a, b\})$ has at least one endpoint in Ω .

Therefore in any case the dual edge $\text{Vor}(\{a, b\})$ has endpoints in Ω , and the dual triangles of those endpoints are incident to s .

We still have to ensure that the three hypothesis (i), (ii) and (iii) of Theorem 5.3 are verified. (i) is obviously verified and (iii) is implied by Rule (5). Let us now prove (ii): consider a constrained subsegment s of $\partial\Omega$. From the first part of the proof, we know that the dual Voronoi edge e of s intersect $\partial\Omega$ in one point and therefore has an odd number of endpoints within Ω . If e had more than one endpoint, i.e. if s had more than one incident triangle, it would in fact have at least three, and s would have at least three incident triangles, contradicting Rule (5). This proves that s has exactly one incident triangle, as required by hypothesis (ii). All three hypothesis are verified. In case of termination, Theorem 5.3 shows that the set T is a triangulation of Ω . \square

Note that Rule 2 can be omitted if we assume that the graph consisting of all constrained segments of C has no vertex of degree 1. Indeed, in such a case, if no constrained subsegment is encroached, none of them can have an ellipse as dual Vface.

6 Termination of the Algorithm

We now consider the conditions needed to ensure the termination of the algorithm. These conditions depend on the shape bound B and on the geometry of the initial set of constrained segments C .

6.1 Distortion and overlapping

In this subsection, we prove that two well-shaped dual triangles cannot overlap if the relative distortion between adjacent sites is small enough. In the following, abc and abd are two adjacent triangles that are respectively dual to Voronoi vertices q_c and q_d . The points q_c and q_d lie inside Ω (see Figure 2), otherwise, their dual triangles would not be considered. We define γ as $\max(\gamma(x, y))$ (see 2.1) where the maximum is taken over all edges (x, y) of the two triangles, and $\delta = \max(\delta(abc), \delta(abd))$.




Figure 2: The two dual triangles and their Voronoi centers

If q_c and q_d verify $\beta(q_c) \leq B$ and $\beta(q_d) \leq B$, we have the following inequalities:

$$\begin{aligned}
 d_c(c, q_d) &\leq d_c(c, q_c) + d_c(q_c, a) + d_c(a, q_d) && \text{(triangular inequality)} \\
 &\leq d_c(c, q_c) + \gamma(a, c)d_a(q_c, a) + \gamma(a, c)d_a(a, q_d) && \text{(distortion between } a \text{ and } c) \\
 &\leq (1 + \gamma(a, c))B\delta(abc) + \gamma(a, c)B\delta(abd) \\
 &\leq (1 + 2\gamma)B\delta
 \end{aligned}$$

The same inequality holds when c and d are exchanged.

In the same way, we have

$$\begin{aligned}
 d_c(c, a) &\leq d_c(c, q_c) + d_c(q_c, a) \\
 &\leq d_c(c, q_c) + \gamma(a, c)d_a(a, q_c) \\
 &\leq (1 + \gamma)B\delta(abc) \\
 &\leq (1 + \gamma)B\delta
 \end{aligned}$$

and

$$\begin{aligned}
 d_c(c, d) &\leq d_c(c, q_d) + d_c(q_d, d) \\
 &\leq (1 + 2\gamma)B\delta + \gamma(c, d)d_d(d, q_d) \\
 &\leq (1 + 3\gamma)B\delta
 \end{aligned}$$

Let $r = (1 + 4\gamma)B\delta$. We consider the zones $Z_3 = B(a, r) \cap B(b, r) \cap B(c, r)$ and $Z_4 = B(a, r) \cap B(b, r) \cap B(c, r) \cap B(d, r)$, where $B(p, r) = \{x \in \mathbb{R}^2, d_p(p, x) \leq r\}$. As shown by the previous inequalities, the four sites a, b, c and d are in Z_4 , as well as the two centers q_c and q_d .

Lemma 6.1 *If a triangle abc is not badly shaped, any point $q \notin Z_3$, is far from each of the three sites a, b and c . More precisely, for any $x \in \{a, b, c\}$, we have*

$$d_x(x, q) > 3B\delta(abc)$$

Proof Assume that $q \notin B(b, r)$ for example. We then have

$$\begin{aligned}
d_a(a, q) &\geq d_b(a, q)/\gamma \\
&\geq (d_b(b, q) - d_b(a, b))/\gamma \\
&\geq (r - (1 + \gamma)B\delta(abc))/\gamma && \text{(by (*))} \\
&\geq ((1 + 4\gamma)B\delta(abc) - (1 + \gamma)B\delta(abc))/\gamma \\
&> 3B\delta(abc)
\end{aligned}$$

□

Let V_{Z_4} be the restriction of the Voronoi diagram to Z_4 . We now establish a sufficient condition on the bound B and on the distortion bound γ so that the vertices and the edges of the Voronoi diagram $V_{Z_4}(\{a, b, c, d\})$ are wedged.

Lemma 6.2 *We are given $B > 1$. If $B^4(\gamma^2 - 1)(1 + \gamma)^2(1 + 4\gamma)^2 \leq 1$, all the 0 and 1-Vfaces of the Voronoi diagram of $\{a, b, c, d\}$ restricted to Z_4 are wedged.*

Proof Let $x, y \in \{a, b, c\}$ with $x \neq y$. Let z be a point of Z_4 on the bisector between x and y . We want to ensure that $d_x(z, y)^2 \leq d_x(x, z)^2 + d_x(x, y)^2$. We have $d_x(z, y)^2 \leq \gamma^2 d_y(y, z)^2$ and since z is on the bisector between x and y , $d_x(x, z) = d_y(y, z)$ †.

Now, if $[x, y]$ is the common edge of the two triangles, we have $d_x(x, y) \geq \delta$. If $[x, y]$ is not the common edge, we only have $d_x(x, y) \geq \delta(abc) \geq d(a, b)/(B(1 + \gamma)) \geq \delta/(B(1 + \gamma))$ ‡.

Finally, by inequalities † and ‡, we have $d_x(x, z)^2 + d_x(x, y)^2 \geq d_y(y, z)^2 + \frac{\delta^2}{B^2(1 + \gamma)^2}$. If $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{B^2(1 + \gamma)^2}$, we then have $d_x(z, y)^2 \leq d_x(x, z)^2 + d_x(x, y)^2$.

Thus, a sufficient condition for z to be wedged is $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{B^2(1 + \gamma)^2}$ (and the condition obtained by swaping x and y). The domain Z_4 was chosen so that $d_y(y, z) < r = (1 + 4\gamma)B\delta$. Hence, a sufficient condition for the point z to be wedged is $(\gamma^2 - 1)(1 + \gamma)^2(1 + 4\gamma)^2 B^4 \leq 1$. □

In the following, the condition $B^4(\gamma^2 - 1)(1 + \gamma)^2(1 + 4\gamma)^2 \leq 1$ will be denoted by (H).

Under the condition stated in Lemma 6.2, the proof of theorem 2.4 may be easily adapted to show that every cell is connected in Z_4 , by showing that it is star-shaped around its site: the segment $[ay]$ is entirely included in Z_4 because Z_4 is convex, as an intersection of ellipses. In order to show that a point y is visible from site a , we only need to consider the Voronoi edges that are intersected by the segment $[ay]$ (see proof of theorem 4 [LS03]). Those intersections lie inside Z_4 , and we can use Lemma 6.2.

Lemma 6.3 *Consider three connected components of distinct 2-Vfaces, whose topological interior are denoted by \mathcal{A} , \mathcal{B} and \mathcal{C} . On the boundary of \mathcal{C} , we cannot find four points $\alpha, \beta, \alpha', \beta'$ in this order such that $\alpha, \alpha' \in \partial\mathcal{A}$ and $\beta, \beta' \in \partial\mathcal{B}$ (see Figure 6.1).*

Proof Assume for a contradiction that $\alpha, \beta, \alpha', \beta'$ exist. Then, since \mathcal{C} and \mathcal{A} are connected, there is a simple path $\pi_{\mathcal{C}}$ in \mathcal{C} and a simple path $\pi_{\mathcal{A}}$ in \mathcal{A} joining α and α' . We concatenate

those two paths into a closed curve π . As \mathcal{B} is connected, there is also a path $\pi_{\mathcal{B}}$ in \mathcal{B} joining β and β' . By Jordan theorem, β and β' should therefore be in the same connected component delimited by π . However, if we follow the boundary of \mathcal{C} from β to β' , we cross π exactly once. So β and β' do not belong to same connected component, which contradicts our hypothesis. \square




Figure 3: Impossible case described in Lemma 6.3

Lemma 6.4 *If (H) is verified and if all the Voronoi vertices in Z_4 are well shaped, all Voronoi vertices of $V_{Z_4}(\{a, b, c, d\})$ have different labels.*

Proof Assume for a contradiction that two Voronoi vertices of $V_{Z_4}(\{a, b, c, d\})$ have the same label $\{a, b, c\}$. By Lemma 2.5, the cells around them have the same orientation. This common orientation shows that we can find four points $\alpha, \beta, \alpha', \beta'$ on the boundary of $\text{Vor}(c) \cap Z_4$ which contradicts Lemma 6.3 (see Figure 3). \square

The same proof applies to the following lemma.

Lemma 6.5 *If $B^4(\gamma^2 - 1)(1 + \gamma)^2(1 + 4\gamma)^2 \leq 1$ and if all Voronoi vertices in Z_3 are well shaped, there is a unique Voronoi vertex with label $\{a, b, c\}$ in $V_{Z_3}(\{a, b, c\})$.*

The following lemma states that under low distortion of the metric, the cells are arranged along the border of Z_4 in the same order as the vertices of the convex hull of $\{a, b, c, d\}$. This topological property will help us prove that we have a triangulation in Z_4 .

Lemma 6.6 *Let p be a vertex of the convex hull of $\{a, b, c, d\}$. If (H) holds, the cell $\text{Vor}(p)$ contains a segment that joins p to a point on the boundary of Z_4 and does not intersect the convex hull of the sites.*

Proof Let us assume that $p = a$ in the following. Let y be one of the points b, c or d . As proved in Lemma 6.2, under condition (H), any point in Z_4 equidistant to y and p is in the wedge that y and p define. Therefore the cell of p in $V_{Z_4}(p, y)$ contains the intersection with Z_4 of a half-plane H^+ defined as follows. H^+ is the half-plane not containing y and bounded by the hyperplane H that goes through p and is normal to $[py]$, from the point of view of p . Denote by H_x^+ the half-plane corresponding to $y = x$, for $x = b, c, d$. Since p is on the boundary of the convex hull of a, \dots, d , the domain $H_b^+ \cap H_c^+ \cap H_d^+$ contains at least one half-line r with origin p : this half-line is any half-line contained in the cone orthogonal (in the sens of the metric of p) to the cone delimited by the tangents to the convex hull at point p . This ray r does not intersect the convex hull of the four sites, and it is inside the cell of site p in the four-sites-diagram. \square

Lemma 6.7 *The 1-Vface of the restricted diagram $V_{Z_4}(\{a, b, c, d\})$ labeled $\{a, b\}$ is connected.*

Proof If the dual Vface of (a, b) in $V_{Z_4}(\{a, b, c, d\})$ does not intersect the boundary of Z_4 or intersect it once, it has to be connected. Indeed, thanks to Lemma 6.4, this 1-Vface has at most two endpoints, labeled $\{a, b, c\}$ and $\{a, b, d\}$, within Z_4 .

We still have to consider the case where such a 1-Vface intersects at least twice the boundary of Z_4 . Let x and y be two of those intersection points.

The connectivity of the cells of a and b shows that there is a path π_a inside the cell of a from x to y and a path π_b inside the cell of b from x to y . Those two paths cannot cross each other (see Figure 4).

The part of Z_4 limited by π_a and which does not contain π_b is now denoted the a -part. We define the b -part symmetrically, and the third part is called the central part. The cells of c and d have to be in the central part. Indeed, if it was not the case with, for example, c in the a -part (recall that the cells are connected), the cell of c could not touch the one of b , and the vertex abc would not exist.

But the central part does not touch the boundary of Z_4 , and neither do the cells of c and d . This last fact contradicts Lemma 6.6. \square




Figure 4: Paths defined in the proof of Lemma 6.7

6.2 Proof of termination

We now consider the algorithm at some point during its execution. We have an arbitrary shape bound B and a distortion coefficient G , chosen so that $(G^2 - 1)(1 + G)^2(1 + 4G)^2 B^4 \leq 1$. This section aims at proving a lower bound on the insertion radius d_{\min}^w of the next inserted site w . By insertion radius, we mean the shortest distance between the new site and all the previously inserted sites. It may depend on the current shortest distance d_{\min} between sites, on the shape bound B , on the geometry of the constrained segments and on the metric field on Ω . The distortion coefficient G is used as a way to discriminate different configurations inside the proof. As we have seen, no such coefficient intervenes in the algorithm itself.

The following definitions are taken from [LS03]:

Definition 6.8 *The bounded distortion radius $\text{bdr}(p, \gamma)$ is the upper bound on the numbers ℓ such that for any point q , we have $d_p(p, q) \leq \ell \Rightarrow \gamma(p, q) \leq \gamma$, and $\text{bdr}_{\min}(\gamma) = \inf \text{bdr}(p, \gamma)$, with the minimum taken over all points $p \in \Omega$.*

Definition 6.9 *Two points q and q' that belong to the union of constrained subsegments $\cup C$ are said to be intertwined if they lie on a common segment of C or if they lie on two edges e and e' of C that share an endpoint b such that $\gamma(q, b) < G$ and $\gamma(q', b) < G$.*

For a set of constrained segments C , the local feature size $\text{lfs}_{\min}(C)$ is the upper bound on the distances r such that $x < r$ implies that for all p , $B(p, x)$ does not contain two non-intertwined points of $\cup C$.

The two following lemmas are Lemma 14 and Lemma 16 of [LS03]:

Lemma 6.10 *Let u and v be two sites of a Voronoi diagram D , and q a point on the bisector $\text{Cell}(u) \cap \text{Cell}(v)$. Assume that there exists some $\gamma > 1$ such that*

$$d_u(u, v) \geq \text{bdr}(u, \gamma).$$

Then for any site w of D ,

$$d(w, q) \geq \text{bdr}_{\min}(\gamma) / (\gamma^3 + \gamma).$$

Lemma 6.11 *Let x be a point in Ω . For any $\gamma > 1$ and for every site p ,*

$$d_x(x, p) \geq \min \left(\frac{d_p(p, x)}{\gamma}, \text{bdr}(x, \gamma) \right).$$

We study now the inter-sites distances created while inserting a new site w along the five rules of the algorithm.

Rule 1:

If Rule 1 applies, the inter-sites distances created by the insertion of the breakpoint of the encroached subsegment are treated in the exact same way as Lemma 17 of [LS03], which we recall:

Lemma 6.12 *Assume that any pair of incident subsegments of C forms an angle of at least $2 \arcsin(G^2/2)$, as measured from the perspective of the intersection point. Let $e = (a, b)$ be a subsegment of C . Let s be a site that encroaches e . Let w be a point in $\text{Vor}(s) \cap e$. Let $m = \min\{d_a(a, s), d_b(b, s)\}$.*

Then for any site p of the diagram, $d(p, w) \geq \min(m, \text{lfs}_{\min}(C)/G, \text{bdr}_{\min}(G))$.

In the following, we call *original refinement point* the point passed as parameter the conditional insertion procedure.

We now consider the cases of Rules 2, 3, 4 and 5 when point w is the original refinement point and not a point lying on an encroached edge.

Rule 2:

If Rule 2 applies, the inserted site w lies on an edge $\text{Vor}(p) \cap \text{Vor}(q)$ but outside of $\text{wedge}(p, q)$ so that we can apply Lemma 5 of [LS03] (stating that in this case, $d_p(p, w) \geq \frac{d_p(p, q)}{\sqrt{G^2-1}}$) and Lemma 6.10. Then, we have for every site p ,

$$d_p(p, w) \geq \min\left(\frac{d_{\min}}{\sqrt{G^2-1}}, \frac{\text{bdr}_{\min}}{G^3+G}\right)$$

and

$$d(p, w) \geq \min\left(\frac{d_{\min}}{G\sqrt{G^2-1}}, \frac{\text{bdr}_{\min}}{G^4+G^2}\right).$$

Rule 3:

If Rule 3 applies, w is located at a Voronoi vertex dual to the triangle abc and at distance $d_p(p, w) \geq r = d_a(a, w) > B\delta(abc)$ from any site p . Lemma 6.11 implies that for every site p and any coefficient G ,

$$d(p, w) \geq \min\left(\frac{B}{G}d_{\min}, \text{bdr}_{\min}(G)\right).$$

Rule 4:

If Rule 4 applies, no vertex is badly shaped, and w is one of the multiple vertices dual to triangle abc . Because w is located at the furthest vertex from a, b and c , Lemma 6.5, implies that either the distortion between the sites a, b and c is greater than G , or w is not in the zone Z_3 .

If the distortion is greater than G , we can use Lemma 6.10. If w is not in Z_3 , thanks to Lemma 6.1, for every site p ,

$$d_p(p, w) \geq B\delta(abc),$$

so that, using Lemma 6.11,

$$d(p, w) \geq \min\left(\frac{B}{G}\delta(abc), \text{bdr}_{\min}(G)\right).$$

In summary, if w is inserted because of Rule 3, we have

$$d(p, w) \geq \min\left(\frac{B}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right).$$

Rule 5:

Finally, if Rule 5 applies, w is located at the Voronoi vertex of abc with abc and abd overlapping one another. Rule 4 implies that this position is unique. The previous section proves that this is only possible if the bound G on the distortion is not respected in abc or abd , since none of them is badly shaped. We may assume that the bound G is not respected in abc , so that we have the bound given by Lemma 6.10:

$$d_p(p, w) \geq \frac{\text{bdr}_{\min}(G)}{G^3 + G}$$

and

$$d_p(p, w) \geq \frac{\text{bdr}_{\min}(G)}{G^4 + G^2}.$$

Summary for Rules 2,3,4,5:

We have proven that, if the original refinement point is inserted, the minimal distance d_{\min}^w after insertion of w verifies

$$d_{\min}^w \geq \min \left(\frac{d_{\min}}{G\sqrt{G^2-1}}, \frac{B}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^4 + G^2} \right)$$

Note that we have implicitly used the relation between B and G stated at the beginning of this section when we have used the property of zones Z_3 and Z_4 , while considering the cases of Rules 4 and 5.

Rules 2,3,4,5 with encroachment:

We now consider the case when the original refinement point s is rejected for encroaching some constrained subsegment $e = (a, b)$. In this case, we insert the corresponding breakpoint w on e . First recall the following fact, extracted from the proof of Lemma 23 in [LS03]:

If w belongs to $\text{Vor}(a)$ and if for some $G > 1$, we have $d_s(s, a) \geq \text{bdr}(s, G)$, then for any site p , $d(p, w) \geq \text{bdr}_{\min}(G)/(G^3 + G)$. Otherwise, $d_p(w) \geq d_a(a, s)/(G\sqrt{G^2 + 1})$.

We can now use the bounds established for Rules 2, 3 and 4 (where w is replaced by s):

$$d_p(p, s) \geq \min \left(\frac{d_{\min}}{\sqrt{G^2-1}}, Bd_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3 + G} \right).$$

Thus,

$$d_p(p, w) \geq \min \left(\frac{d_{\min}}{G\sqrt{G^4-1}}, \frac{Bd_{\min}}{G\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^4 + G^2)\sqrt{G^2+1}} \right).$$

and we have, by Lemma 6.11

$$d(p, w) \geq \min \left(\frac{d_{\min}}{G^2\sqrt{G^4-1}}, \frac{Bd_{\min}}{G^2\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^5 + G^3)\sqrt{G^2+1}} \right).$$

Termination

In order to handle the first two terms in the previous equation and to respect the condition of lemma 6.2, let $B > 1$ and $G > 1$ be such that

$$(G^2 - 1)(1 + G)^2(1 + 4G)^2 B^4 \leq 1$$

$$\begin{aligned} G^2\sqrt{G^4-1} &\leq 1 \\ G^2\sqrt{G^2+1} &\leq B. \end{aligned}$$

Note that for any $B > \sqrt{2}$, a suitable $G > 1$ may be found, since every condition is verified when $G \rightarrow 1^+$. We also demand that any pair of incident edges of C forms an angle of at least $2 \arcsin(G^2/2)$, so that it complies to the requirements of lemma 6.12. Under those conditions, the minimal inter-distance d'_{\min} after the insertion of a new site is bounded from below:

$$d_{\min} \geq \min \left(\frac{\text{bdr}_{\min}(G)}{(G^5 + G^3)\sqrt{G^2 + 1}}, \frac{\text{lfs}_{\min}(C)}{G} \right)$$

Finally, if we can find G satisfying the conditions and such that $\text{bdr}_{\min}(G) > 0$ the above bound is not trivial, and an easy induction shows that we indeed have a lower bound on the minimal inter-distance. It proves that the algorithm will not insert sites indefinitely, by a classical volume argument.

Moreover, because

$$(G^2 - 1)B^2 < 1,$$

the shape condition parametrized by B may be translated into a condition in terms of a lower bound on the angles of the triangles, as measured by any point inside the triangle (see Corollary 10 of [LS03]).

Theorem 6.13 *Let $B > \sqrt{2}$ be a constant, and let $G > 1$ be such that*

$$\begin{aligned} (G^2 - 1)(1 + G)^2(1 + 4G)^2 B^4 &\leq 1 \\ G^2\sqrt{G^4 - 1} &\leq 1 \\ G^2\sqrt{G^2 + 1} &\leq B. \end{aligned}$$

Let C be a set of constrained segments which delimits a polygonal domain of the plane such that incident segments always form an angle greater than $2 \arcsin(G^2/2)$. If $\text{bdr}_{\min}(G) > 0$, the algorithm presented in section 5.3 terminates and provides a triangulation whose dual Voronoi vertices respect the shape bound B .

7 Conclusion

The approach we have presented is built upon the work of Labelle and Shewchuk. Instead of using a lower envelop of paraboloids, we rely on a power diagram in higher dimension. Moreover, we present the algorithm by focusing on the overlapping condition, thus minimizing the dependence over the Voronoi diagram itself, apart from the computation of the Voronoi vertices. As an aside, we also rely only on the Voronoi vertices that are inside the domain Ω , while Labelle and Shewchuk compute the whole diagram.

This algorithm has been implemented using the Computational Geometry Algorithms Library [CGA05].

A similar algorithm can be considered in three dimensions. However, we currently cannot prove that this meshing algorithm terminates in three dimensions because sliver tetrahedra may overlap their neighbors, without inducing a large insertion distance for the new refining point. This may happen even in the case of low distortion of the metric field.

References

- [BGH⁺97] Housman Borouchaki, Paul Louis George, Frédéric Hecht, Patrick Laug, and Eric Saltel. Delaunay mesh generation governed by metric specifications. part i algorithms. *Finite Elem. Anal. Des.*, 25(1-2):61–83, 1997.
- [BH96] Frank Bossen and Paul Heckbert. A pliant method for anisotropic mesh generation. In *5th International Meshing Roundtable*, October 1996.
- [CGA05] *The CGAL Manual*, 2005. Release 3.0.
- [LS03] Francois Labelle and Jonathan Richard Shewchuk. Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, pages 191–200, New York, NY, USA, 2003. ACM Press.
- [LTÜ99] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. Biting ellipses to generate anisotropic mesh. In *8th International Meshing Roundtable*, October 1999.
- [She02] Jonathan Richard Shewchuk. What is a good linear finite element? interpolation, conditioning, anisotropy, and quality measures. Manuscript 2002.
- [SYI97] Kenji Shimada, Atsushi Yamada, and Takayuki Itoh. Anisotropic triangular meshing of parametric surfaces via close packing of ellipsoidal bubbles. In *6th International Meshing Roundtable*, October 1997.

Contents

1	Introduction	3
2	Labelle and Shewchuk's approach	3
2.1	Anisotropic Voronoi diagram definition	3
2.2	The wedge property	5
3	Power Diagram and Anisotropic Voronoi Diagram	6
4	Required computations	7
4.1	Computations of Voronoi vertices	7
4.2	Computations of encroachment	8
5	Description of the Algorithm	9
5.1	Local Embedding	9
5.2	Form Criterion	11
5.3	Refinement Algorithm	11
6	Termination of the Algorithm	13
6.1	Distortion and overlapping	13
6.2	Proof of termination	18
7	Conclusion	21



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399