

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. Theoretical Computer Science, Elsevier, 2008, pp.163-173. <10.1016/j.tcs.2008.08.006>. <inria-00336798>

HAL Id: inria-00336798

<https://hal.inria.fr/inria-00336798>

Submitted on 5 Nov 2008

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.

Anisotropic Diagrams: Labelle Shewchuk approach revisited

Jean-Daniel Boissonnat, Camille Wormser*, Mariette Yvinec

INRIA, Geometrica, BP 93 06902 Sophia Antipolis, France

Abstract

F. Labelle and J. Shewchuk 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

1991 MSC: 65D18, 68U05

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 [8] 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

* Corresponding author.

Email addresses: jean-daniel.boissonnat@sophia.inria.fr (Jean-Daniel Boissonnat), camille.wormser@sophia.inria.fr (Camille Wormser), mariette.yvinec@sophia.inria.fr (Mariette Yvinec).

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. [7] and Shimada et al. [9] use packing methods. Bossen and Heckbert [3] use a method consisting in centroidal smoothing, retriangulating and inserting or removing sites. Borouchaki et al. [2] adapt the classical Delaunay refinement algorithm to the case of an anisotropic metric. In terms of applications, the question of tailoring anisotropic meshes to the specific needs of partial differential equations solvers has been studied by Simpson [10]. An example of strategy used to adapt anisotropic meshes thanks to a posteriori computations of the error in finite elements computations has been presented by Apel et al. [1], and typical examples of applications to fluid dynamics computations have been investigated by Frey and Alauzet [5].

Recently, Labelle and Shewchuk [6] have settled the foundations for a rigorous approach based on the so-called anisotropic Voronoi diagrams. We present these ideas in the first two sections, and we propose an alternative view of the construction of these diagrams in Section 3. After detailing in Section 4 the computations that we need, we expose a variant of the meshing algorithm of Labelle and Shewchuk in Section 5. This variant computes the Voronoi vertices using a higher dimensional power diagram and refines the diagram as long as dual triangles overlap. The last sections prove the correctness of this approach.

An extension of Labelle and Shewchuk results to the 2-manifold case was proposed by Cheng et al. [4], where a 3D anisotropic Voronoi diagram is considered to build an anisotropic mesh of the closed 2-manifold embedded in 3D.

2 Labelle and Shewchuk's Approach

Labelle and Shewchuk [6] 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 we write $d(p, q) = \min(d_p(p, q), d_q(p, q))$. Note that d_p is a distance, whereas d is not, since it does not necessarily verifies the triangular inequality.

In a similar way, $\angle_p xqy$ is defined as $\arccos \frac{(x-q)^t M_p (y-q)}{d_p(x,q)d_p(y,q)}$.

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 [6] define the anisotropic Voronoi diagram in the following way (and provide some examples):

Definition 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 $\text{Vor}(R) = \bigcap_{q \in R} \text{Vor}(q)$ which is the locus of points equally close to the sites in R and no closer to any other site. 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 *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 $\text{Vor}(R) \neq \emptyset$.*

In two dimensions and with points in generic position, the dual complex 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 straighten 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.

In order to prove the correctness of their algorithm, Labelle and Shewchuk [6] have defined the wedge property and have proven the following results to ensure that their algorithm converges to a triangulation.

Definition 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 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, i.e. $d = 2$.

Lemma 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 correctness of Labelle and Shewchuk's algorithm.

Theorem 6 *Suppose that each 1-Vface of D 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 dual complex of D_Ω is a triangulation of Ω if S is in general position, i.e. if all Voronoi vertices have degree 3.*

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 6 validates their lazy computation of the diagram¹.

¹ In 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.

Labelle and Shewchuk's algorithm consists in incrementally inserting points

- on edges of $\partial\Omega$ until these segments appear in D_Ω ;
- 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.

3 Relation to Power Diagrams

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. In the following, $\|\cdot\|$ denotes the Euclidean distance.

Definition 7 *A power diagram is defined for a set of spheres. Given a sphere σ centered at y and of radius r , the power distance of a point x with respect to σ is defined as $\pi_\sigma(x) = \|x - y\|^2 - r^2$.*

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 distance 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 set of spheres $\{\sigma_i\}_i$ as $Pow(\{\sigma_i\}_i) = \cap_i Pow(\sigma_i)$. The dual of the power diagram of Σ is called the regular complex of Σ .

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

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

where \mathcal{P} denotes the d -manifold of $\mathbb{R}^D \left\{ \hat{x} \in \mathbb{R}^D : x \in \mathbb{R}^d \right\}$.

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_{p_i} , whose elements are denoted by $(M_{p_i}^{r,s})_{1 \leq r, s \leq d}$, and we define

- the point $q_i = (q_i^{r,s}, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}}$ defined as
 - $q_i^{r,r} = -\frac{1}{2}M_{p_i}^{r,r}$, for $1 \leq r \leq d$;
 - $q_i^{r,s} = -M_{p_i}^{r,s}$, for $1 \leq r < s \leq d$.
- the point $\hat{p}_i = (M_{p_i} p_i, q_i) \in \mathbb{R}^D$;

- the sphere $\sigma(p_i) \subset \mathbb{R}^D$ of center \hat{p}_i and radius $\sqrt{\|\hat{p}_i\|^2 - p_i^t M_{p_i} p_i}$.

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 8 *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 equalities:

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

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

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

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 cell of p_i in the anisotropic diagram of S .

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 Basic Operations and Primitives

Computing the complete anisotropic Voronoi diagram explicitly is not easy. Our meshing algorithm only requires computing 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 8).

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

- (1) Compute the power diagram of Σ and consider three sites a , b and c such that $(\sigma(a), \sigma(b), \sigma(c))$ corresponds to a triangle in the regular complex of Σ (see Definition 7), which means 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 previous steps, we have only considered the bisectors of the spheres $\sigma(a), \sigma(b), \sigma(c)$ corresponding to the three sites involved, or, equivalently, the Voronoi diagram of $\{a, b, c\}$ alone. We now consider the bisectors of a, b, c in the Voronoi diagram of S , or, equivalently, the bisectors of the spheres $\sigma(a), \sigma(b), \sigma(c)$ in the power diagram of Σ . In the Voronoi diagram of S , 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 \mathbb{R}^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 $H_{ab} \cap H_{bc}$. It belongs to the power cell $Pow(\{\sigma(a), \sigma(b), \sigma(c)\})$ if and only if its power to $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$ is smaller than its power to any other $\sigma(c')$ in Σ . We do not have to check this fact for all the other spheres $\sigma(c')$ with $c' \in S$, but only for the spheres whose cells are incident to $Pow(\{\sigma(a), \sigma(b), \sigma(c)\})$, 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. Note that those three inequalities are equivalent, since z has the same power with respect to the three spheres $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$. The points kept are in fact the Voronoi vertices labeled by $\{a, b, c\}$.

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

Definition 9 *A constraint 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 constraint 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 constraint subsegments, it is sufficient to consider the ones that would have at least one of their endpoints joined to p in the dual complex, if p were inserted in the diagram. Indeed, if p encroaches e , the cell $\text{Vor}(p)$ is adjacent to the cells of at least one of the endpoints of e : before the insertion of p , $e = [a, b]$ was not encroached and was covered by $\text{Vor}(a)$ and $\text{Vor}(b)$. After the insertion of p , $\text{Vor}(p)$ covers a part of e , while $\text{Vor}(a)$ and $\text{Vor}(b)$ cover the rest of it.

Practically, let $e = [a, b]$ be such a constraint subsegment. Then, let E be the intersection $C_{pa} \cap [a, b]$ of the bisector of p and a with $[a, b]$. If some $z \in E$ 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 constraint 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

As above, let Ω be a polygonal domain of the plane, whose boundary is denoted by $\partial\Omega$. We denote by C the set of constraint 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 constraint segments have been refined into constraint 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.

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 (that is, triangles with non collinear vertices) 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 10 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 $link(p)$ of a site p is the set of edges opposite to p in all the triangles of T_p .




Fig. 1. Two umbrellas (left) and one umbrella winding twice (right) around p

Lemma 11 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 ;
- (b) if p is an internal site, its umbrellas are combinatorial disks and p is inside each of its embedded 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, the triangles of the umbrella do not overlap.

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 $\partial\Omega$. Since Ω is a simply connected polygon,

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

$\partial\Omega$ is a topological circle embedded in the plane. The set U is closed, and so is its boundary ∂U . It follows that ∂U is a closed non empty subset of the topological circle $\partial\Omega$, which implies that $\partial\Omega = \partial U$. Finally, both U and Ω are bounded domains, with the same circle as boundary, hence $U = \Omega$.

- (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 an embedded umbrella, i.e. the union U of the triangles of an umbrella. Assume for a contradiction that p is not in the interior of this union. Then p is on the boundary ∂U and both edges of this boundary that are incident to p belong to two triangles of the umbrella which have to overlap. This contradicts Rule (iii). In other words, we have proved that if there is a closed curve in the link of p , p is enclosed by it.

- (c) If p is a vertex of $\partial\Omega$, $link(p)$ may a priori contain some closed curves and some curves joining the two neighbors of p on $\partial\Omega$. As seen in the proof of (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 $link(p)$ have the same first and last segment and because Rule (iii) prevents any branching vertex in $link(p)$, the link $link(p)$ is a single curve. The fact that the triangles of the unique umbrella do not overlap follows from (iii) too.

Theorem 12 *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. Using the fact that \mathcal{G}_p is a separated space, that h_p is a proper map, and that Ω_p is connected, it follows that h_p is a covering 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 a homeomorphism. Thus, Ω is triangulated by T .

In order to present the refinement algorithm, we need to define a shape criterion. Let v be a Voronoi vertex of an anisotropic Voronoi diagram. The label of v 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 (a, b) is $d(a, b) = \min(d_a(a, b), d_b(a, b))$. 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 or the associated triangle are said to be *badly-shaped* if $\beta(v) > B$. Otherwise, they are said to be *well-shaped*.

Let us now present the algorithm, which refines an anisotropic Voronoi diagram V until the set of triangles dual to the Voronoi vertices of V_Ω , the restriction of V to Ω , have a good shape and satisfy conditions (i), (ii) and (iii) (stated in Section 5), and therefore form a triangulation of Ω , by Theorem 12.

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 of its endpoints.

Definition 13 *Assume that a constraint subsegment $e = (p, q)$ is encroached. The breakpoint of the edge (p, q) is defined as the point of $[p, q] \setminus (\text{Vor}(p) \cup \text{Vor}(q))$ closest to the midpoint of $[p, q]$ (this point is independent of the considered metric). 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). We define a procedure of conditional insertion, needed for the presentation of the algorithm:

CONDITIONALLY INSERT(x): if x encroaches some constraint subsegment e , insert a site at the breakpoint of e . Otherwise, insert x .




Fig. 2. Edge e is a constraint segment, with the cell of a being completely included in the cell of b . Voronoi bisectors are represented by dashed curves.

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 constraint subsegment $e \in C$ does not appear as the edge of a dual triangle because it is encroached, insert a site at the breakpoint of edge e ;
- Rule (2) if some constraint subsegment $e \in C$ does not appear as the edge of a dual triangle, because its dual Vface is a complete ellipse (it can happen if the constraint subsegment has a free endpoint, i.e. an endpoint which is not incident to any other constraint subsegment, see Figure 2 for an example), 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), 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 ;
- Rule (5) if two triangles share an edge and 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 ($\gamma(x, y)$ is the distortion between x and y defined in Section 2).

We will now prove that if the algorithm terminates, Conditions (i), (ii) and (iii) of Section 5 are verified. By Theorem 12, the dual complex is therefore a triangulation, without any badly-shaped vertex.

Lemma 14 *Upon termination of the algorithm, the dual triangles in T form a triangulation of the domain Ω and all the constraint subsegments appear in this triangulation.*

PROOF. First, let us prove that each constraint subsegment is incident to at least one triangle in T . Consider some constraint 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 $\text{Vor}(\{a, b\})$ is a union of curved segments, with an even number of endpoints. Furthermore, owing to the monotonicity of the distance $d_a(a, x)$ along ab , $\text{Vor}(\{a, b\})$ intersects s in at most one point (and at least once, thanks to Rule (1)). Consider the curved segment ℓ of $\text{Vor}(\{a, b\})$ which intersects s . One of the two endpoints of ℓ has to be inside Ω because $\text{Vor}(\{a, b\})$ cannot intersect any other constraint subsegment, since the other constraint subsegments are not encroached either. It follows that $\text{Vor}(\{a, b\})$ 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 hypotheses (i), (ii) and (iii) of Theorem 12 are verified. (i) is obviously verified and (iii) is implied by Rule (5). Let us now prove (ii): consider a constraint 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 12 shows that the set T is a triangulation of Ω .

Note that Rule 2 can be omitted if we assume that the graph consisting of all constraint segments of C has no vertex of degree 1. Indeed, in such a case, if no constraint subsegment is encroached, none of them can have an ellipse as a 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 K and on the geometry of the initial set of constraint segments C .

Let us prove that two well-shaped dual triangles (as defined in Section 5)

cannot overlap if the relative distortion between adjacent sites is small enough. In the following, abc and abc' are two adjacent triangles that are respectively dual to Voronoi vertices q_c and $q_{c'}$. The points q_c and $q_{c'}$ lie inside Ω , otherwise, their dual triangles would not be considered. We define γ as the maximum of the distortion $\gamma(x, y)$ (see Section 2) where the maximum is taken over all edges (x, y) of the two triangles, and $\delta = \max(\delta(abc), \delta(abc'))$ (as defined in Section 5).

If q_c and $q_{c'}$ are well-shaped, i.e. $\beta(q_c) \leq K$ and $\beta(q_{c'}) \leq K$, we have the following inequalities:

$$\begin{aligned} d_c(c, q_{c'}) &\leq d_c(c, q_c) + d_c(q_c, a) + d_c(a, q_{c'}) && \text{(triangular ineq.)} \\ &\leq d_c(c, q_c) + \gamma(a, c)d_a(q_c, a) + \gamma(a, c)d_a(a, q_{c'}) && \text{(distortion)} \\ &\leq (1 + \gamma(a, c))K\delta(abc) + \gamma(a, c)K\delta(abc') \\ &\leq (1 + 2\gamma)K\delta \end{aligned}$$

The same inequality holds when c and c' are exchanged.

In the same way, we have

$$\begin{aligned} d_c(c, a) &\leq d_c(c, q_c) + d_c(q_c, a) \leq d_c(c, q_c) + \gamma(a, c)d_a(a, q_c) \\ &\leq (1 + \gamma)K\delta(abc) \leq (1 + \gamma)K\delta \\ &\text{and} \\ d_b(b, a) &\leq (1 + \gamma)K\delta(abc) \quad (*) \\ &\text{and} \\ d_c(c, c') &\leq d_c(c, q_{c'}) + d_c(q_{c'}, c') \leq (1 + 2\gamma)K\delta + \gamma(c, c')d_{c'}(c', q_{c'}) \\ &\leq (1 + \gamma)^2K\delta \end{aligned}$$

Let $r = (1 + \gamma)^2K\delta$. We consider the zones $Z_3 = B(a, r) \cap B(b, r) \cap B(c, r)$, $Z'_3 = B(a, r) \cap B(b, r) \cap B(c', r)$ and $Z_4 = B(a, r) \cap B(b, r) \cap B(c, r) \cap B(c', r)$, 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 c' are in Z_4 , as well as the two centers q_c and $q_{c'}$.

Lemma 15 *If a triangle abc is well-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 $d_x(x, q) > 2K\delta(abc)$.*

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

$$\begin{aligned} d_a(a, q) &\geq d_b(a, q)/\gamma \geq (d_b(b, q) - d_b(a, b))/\gamma \\ &\geq (r - (1 + \gamma)K\delta(abc))/\gamma && \text{(by (*))} \\ &\geq ((1 + \gamma)^2K\delta(abc) - (1 + \gamma)K\delta(abc))/\gamma > 2K\delta(abc) \end{aligned}$$

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

Definition 16 *The three following conditions are called condition (H):*

- (i) $K > 1$ and $K^4(\gamma^2 - 1)(1 + \gamma)^6 \leq 1$
- (ii) *the triangles are well-shaped (for the bound K);*
- (iii) γ *is an upper bound on the distortion between the considered sites.*

Lemma 17 *Under condition (H), all the 0 and 1-Vfaces of the Voronoi diagram of $V_{Z_3}(\{a, b, c\})$ and all the 0 and 1-Vfaces of the Voronoi diagram of $V_{Z'_3}(\{a, b, c'\})$ 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 of 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 $d_x(x, y) \geq \delta$. Otherwise, we have by (*) $d_x(x, y) \geq \delta(abc) \geq d(a, b)/(K(1 + \gamma)) \geq \delta/(K(1 + \gamma))$. ‡

Finally, by inequalities † and ‡, $d_x(x, z)^2 + d_x(x, y)^2 \geq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$. And if $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$, we 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 $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$ (and the condition obtained by swapping x and y). The domain Z_4 was chosen so that $d_y(y, z) < r = (1 + \gamma)^2 K \delta$. Hence, a sufficient condition for the point z to be wedged is $(\gamma^2 - 1)(1 + \gamma)^2(1 + \gamma)^4 K^4 \leq 1$, i.e. $(\gamma^2 - 1)(1 + \gamma)^6 K^4 \leq 1$.

Lemma 18 *Under condition (H), the cells of a , b and c in $V_{Z_3}(\{a, b, c\})$ are connected.*

PROOF. Under condition (H), the proof of Theorem 4 (Theorem 4 in [6]) can easily be adapted to show that every cell is connected in Z_3 , by showing that it is star-shaped around its site: let y be some point of the cell of a in $V_{Z_3}(\{a, b, c\})$. The segment $[ay]$ is entirely included in Z_3 because Z_3 is convex, as an intersection of ellipses. In order to show that y is visible from site a , we only need to consider the Voronoi edges that are intersected by the segment $[ay]$. Those intersection points lie inside Z_3 . From Lemma 17, the intersection

points are wedged, and the proof of Theorem 4 [6] shows that the cell of a in $V_{Z_3}(\{a, b, c\})$ is star-shaped.

Lemma 19 *Consider three connected components of distinct 2-Vfaces, whose topological interiors 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 3).*

PROOF. Assume to the contrary that $\alpha, \beta, \alpha', \beta'$ exist. Then, since \mathcal{C} and \mathcal{A} are connected, there exists a simple path $\pi_{\mathcal{C}}$ in \mathcal{C} and a simple path $\pi_{\mathcal{A}}$ in \mathcal{A} joining α and α' . The union of those two paths forms 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.




Fig. 3. Impossible case described in Lemma 19

Lemma 20 *If (H) is verified and if all the Voronoi vertices in Z_3 are well-shaped, there is a unique Voronoi vertex with label $\{a, b, c\}$ in $V_{Z_3}(\{a, b, c\})$.*

PROOF. Assume for a contradiction that two Voronoi vertices v and v' of $V_{Z_3}(\{a, b, c\})$ have the same label $\{a, b, c\}$. By Lemma 5, the cells around v and v' have the same cyclic order.

By Lemma 18, the three cells $\text{Vor}(a) \cap Z_3$, $\text{Vor}(b) \cap Z_3$ and $\text{Vor}(c) \cap Z_3$ are connected. By considering the neighborhoods of v and v' , we can find four points $\alpha, \beta, \alpha', \beta'$ in this order on the boundary of $\text{Vor}(c) \cap Z_3$ such that α, α' belong to the boundary of the cell of a and β, β' belong to the boundary of the cell of b . This contradicts Lemma 19 (see Figure 3).

Lemma 21 *If (H) is verified and if all the Voronoi vertices labeled by $\{a, b, c\}$ and $\{a, b, c'\}$ in Z_4 are well-shaped, there is exactly one Voronoi vertex labeled by $\{a, b, c\}$ and one Voronoi vertex labeled by $\{a, b, c'\}$ in $V_{Z_4}(\{a, b, c, c'\})$.*

PROOF. The result follows from Lemma 20, because $Z_4 \subset Z_3$ and any vertex labeled by $\{a, b, c\}$ in Z_4 is also a vertex 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 $\{a, b, c, c'\}$. This topological property will help us prove that we have a triangulation in Z_4 .

Lemma 22 *Let x be a, b or c . If (H) holds, the cell $\text{Vor}(x)$ in $V_{Z_3}(\{a, b, c\})$ contains a segment that joins x to a point on the boundary of Z_3 and does not intersect the convex hull of the sites.*

PROOF. Let us assume that $x = a$ in the following. As proved in Lemma 17, under condition (H), any point in Z_3 equidistant to b and a is in the wedge defined by b and a . Therefore the cell of a in $V_{Z_3}(a, b)$ contains the intersection of Z_3 with a half-plane H_b^+ defined as follows. H_b^+ is the half-plane not containing b and bounded by the hyperplane H_b that goes through a and is normal to $[ab]$, from the point of view of a . Since a is on the boundary of the convex hull of a, b, c , the domain $H_b^+ \cap H_c^+$ contains at least one half-line r with origin a : this half-line is any half-line contained in the cone orthogonal (in the sense of the metric of a) to the cone delimited by the tangents to the convex hull at point a . This ray r does not intersect the convex hull of the three sites, and it is inside the cell of site a in the three-sites-diagram.

Lemma 23 *If (H) is verified and if all the Voronoi vertices labeled by $\{a, b, c\}$ and $\{a, b, c'\}$ in Z_3 and Z_3' respectively are well-shaped, the 1-Vface of the restricted diagram $V_{Z_4}(\{a, b, c, c'\})$ labeled $\{a, b\}$ is connected.*

PROOF. Let e be the dual 1-Vface of (a, b) in $V_{Z_4}(\{a, b, c, c'\})$. If e does not intersect the boundary of Z_4 or intersects it once, e has to be connected. Indeed, thanks to Lemma 21, e has at most two endpoints, labeled $\{a, b, c\}$ and $\{a, b, c'\}$, within Z_4 .

We now prove that e does not touch the boundary of Z_4 . From Lemma 18 and 20, the 1-Vface A labeled by $\{a, b\}$ in $V_{Z_3}(\{a, b, c\})$ is connected. Since a vertex labeled by $\{a, b, c'\}$ exists in $V_{Z_4}(\{a, b, c, c'\})$, it has to belong to A . Consider the arc $\ell \subset A$ of the bisector of $\{a, b\}$ which links the vertex q_c labeled by $\{a, b, c\}$ and the vertex $q_{c'}$ labeled $\{a, b, c'\}$. Let us prove that ℓ is

entirely included in Z_4 . Assume to the contrary that ℓ is not entirely included in Z_4 . The boundary of Z'_3 has to intersect it twice, because $Z_4 = Z_3 \cap Z'_3$. It follows that the 1-Vface A' labeled by $\{a, b\}$ in $V_{Z_3}(\{a, b, c'\})$, which contains $\ell \cap Z_4$, intersects the boundary of Z'_3 twice. Since there is only one vertex labeled $\{a, b, c'\}$ in Z'_3 , there is a sub-arc ℓ' of A' without any vertex on it. ℓ' cuts Z_3 into two parts (called the two sides of ℓ' in the following). The cell of c' is connected, and is on one side of ℓ' . The other side of ℓ' belongs to another cell. Let us assume, without loss of generality, that it belongs to the cell of a . This part of the cell of a is not adjacent to the cell of c' , which implies that the cell of a is disconnected. This contradicts the fact that the cells of $V_{Z_3}(\{a, b, c'\})$ are connected. We have proved that ℓ is entirely included in Z_4 . Then, since there is only one vertex labeled by $\{a, b, c'\}$ in Z_4 , ℓ is exactly the 1-VFace labeled by $\{a, b\}$ in $V_{Z_4}(\{a, b, c, c'\})$. This concludes the proof.

Lemma 24 *If (H) is verified and if all the Voronoi vertices in Z_4 are well-shaped, the two triangles abc and abc' do not overlap each other.*

PROOF. From Lemma 23, the 1-Vface labeled $\{a, b\}$ in the restricted diagram $V_{Z_4}(\{a, b, c, c'\})$ is connected. The two endpoints of the 1-Vface labeled $\{a, b\}$ are the Voronoi vertices q_c and $q_{c'}$. It follows that the cells of a, b and c around q_c and the cells of a, b and c' around $q_{c'}$ have opposite cyclic orders. Lemma 5 applied to $V_{Z_3}(\{a, b, c\})$ and $V_{Z'_3}(\{a, b, c'\})$ then implies that the triangles abc and abc' do not overlap each other.

We now consider the algorithm at some point during its execution. The proof makes use of an arbitrary shape bound K and a distortion coefficient G , chosen so that the following condition (C) is satisfied: any pair of adjacent segments of C forms an angle of at least $2 \arcsin(G^2/2)$ and $(G^2 - 1)(1 + G)^6 K^4 \leq 1$. (C)

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 Euclidean distance between the new site and all the previously inserted sites. It may depend on the current shortest anisotropic distance d_{\min} between the sites, on the shape bound K , on the geometry of the constraint 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 [6]:

Definition 25 *The bounded distortion radius $\text{bdr}(p, \gamma)$ is defined as $\sup\{\ell : d_p(p, q) \leq \ell \Rightarrow \gamma(p, q) \leq \gamma\}$ and $\text{bdr}_{\min}(\gamma)$ is the lower bound of these radii: $\text{bdr}_{\min}(\gamma) = \inf\{\text{bdr}(p, \gamma) : p \in \Omega\}$.*

Definition 26 Given some bound $G > 0$, two points q and q' that belong to constraint segments in C are said to be G -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 and are such that $\gamma(q, b) < G$ and $\gamma(q', b) < G$. For a set of constraint segments C , the local feature size $\text{lfs}_{\min}^G(C)$ is the upper bound on the distances r such that $x < r$ implies that for all $p \in \Omega$, $B(p, x)$ does not contain two non- G -intertwined points of $\cup C$.

The following four lemmas are Lemma 5, 14, 16 and 17 of [6]:

Lemma 27 Let w be a point on the bisector of a and b that lies outside the wedge of a and b , on the side of b . Let $G \geq 1$ be a constant for which $\gamma(a, b) \leq G$. Then the proximity of w to a and b is bounded by $d_a(a, w) = d_b(b, w) \geq d_b(b, a) / \sqrt{G^2 - 1}$.

Lemma 28 Let a and b be two sites of a Voronoi diagram D , and w a point on the bisector $\text{Vor}(\{a, b\})$ in D . Assume that there exists some $G > 1$ such that $d_a(a, b) \geq \text{bdr}(a, G)$. Then for any site x of D , $d(x, w) \geq \text{bdr}_{\min}(G) / (G^3 + G)$.

Lemma 29 Let p be a point in Ω . For any $G > 1$ and for every site x , $d_p(p, x) \geq \min\left(\frac{d_x(x, p)}{G}, \text{bdr}(p, G)\right)$.

Lemma 30 Assume that any pair of adjacent segments of C forms an angle of at least $2 \arcsin(G^2/2)$, as measured by the common endpoint. 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 x of the diagram, $d(x, w) \geq \min(m, \text{lfs}_{\min}^G(C)/G, \text{bdr}_{\min}(G))$.

We study now the inter-site distances created while inserting a new site w along the five rules of the algorithm, as presented in Section 5. Recall that G is assumed to satisfy Condition (C).

Rule 1: If Rule 1 applies, the inter-site distances created by the insertion of the breakpoint of the encroached subsegment are bounded by Lemma 30: for any site x of the diagram, $d(x, w) \geq \min(m, \text{lfs}_{\min}^G(C)/G, \text{bdr}_{\min}(G))$.

We call *original refinement point* the point passed as argument to the conditional insertion procedure. We now consider the cases of Rules 2, 3, 4 and 5 when the inserted 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 $\text{Vor}(a) \cap \text{Vor}(b)$ but outside wedge(a, b). We have two cases to consider. If $\gamma(a, b) \leq G$, we can apply Lemma 27 so that for every site x , $d_x(x, w) \geq d_a(a, w) = d_b(b, w) \geq \frac{d_{\min}}{\sqrt{G^2 - 1}}$, and Lemma 29 then implies $d(x, w) \geq \min\left(\frac{d_{\min}}{G\sqrt{G^2 - 1}}, \text{bdr}_{\min}(G)\right)$. If $\gamma(a, b) > G$, Lemma 28 implies $d(x, w) \geq \frac{\text{bdr}_{\min}(G)}{G^3 + G}$.

Rule 3: If Rule 3 applies, w is located at a Voronoi vertex dual to the triangle abc and at distance $d_x(x, w) \geq r = d_a(a, w) > K\delta(abc)$ from any site x . Lemma 29 implies that for every site x and any coefficient G , $d(x, w) \geq \min\left(\frac{K}{G}d_{\min}, \text{bdr}_{\min}(G)\right)$.

Rule 4: If Rule 4 applies, no vertex is badly shaped, and w is one of the vertices dual to triangle abc . Because w is located at the furthest vertex from a , b and c , Lemma 20, implies that either the distortion between the sites a , b and c is greater than G , or w does not belong to the zone Z_3 . If the distortion is greater than G , we can use Lemma 28. If w is not in Z_3 , thanks to Lemma 15, for every site x , $d_x(x, w) \geq K\delta(abc) \geq Kd_{\min}$, so that, using Lemma 29, $d(x, w) \geq \min\left(\frac{K}{G}\delta(abc), \text{bdr}_{\min}(G)\right)$. In summary, if w is inserted by Rule 4, $d(x, w) \geq \min\left(\frac{K}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right)$.

Rule 5: Finally, if Rule 5 applies, w is located at the Voronoi vertex of a triangle abc overlapping another triangle abc' . Rule 4 implies that abc has a unique dual vertex. Lemma 24 proves that this is only possible if γ , the maximum of the distortion $\gamma(x, y)$ where the maximum is taken over all edges (x, y) of the two triangles abc and abc' is greater than G , since both abc and abc' are well shaped. Then we have the bound given by Lemma 28 for every site x : $d(x, w) \geq \frac{\text{bdr}_{\min}(G)}{G^3+G}$.

Summary for Rules 2,3,4,5 without encroachment:

We have proven that, if the original refinement point is inserted, the minimal distance d_{\min}^w after insertion of w verifies $d_{\min}^w \geq \min\left(\frac{d_{\min}}{G\sqrt{G^2-1}}, \frac{K}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right)$ where K is the shape bound and G is any value satisfying Condition (C), as stated at the beginning of Section 6.

Rules 2,3,4,5 with encroachment: Denote by $e = (a, b)$ the constraint subsegment encroached by s . Since s encroaches e , we insert the corresponding breakpoint w on e . First recall the following fact, extracted from the proof of Lemma 23 in [6]: 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 x , $d(x, w) \geq \text{bdr}_{\min}(G)/(G^3 + G)$. Otherwise, we have $d_x(x, w) \geq d_a(a, s)/(G\sqrt{G^2 + 1})$. We use now the bounds established for Rules 2, 3 and 4 (with w replaced by s):

$$d_x(x, s) \geq \min\left(\frac{d_{\min}}{\sqrt{G^2-1}}, Kd_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right), \text{ and } d_x(x, w) \geq \min\left(\frac{d_{\min}}{G\sqrt{G^4-1}}, \frac{Kd_{\min}}{G\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^4+G^2)\sqrt{G^2+1}}\right). \text{ Lemma 29 then implies that } d(x, w) \geq \min\left(\frac{d_{\min}}{G^2\sqrt{G^4-1}}, \frac{Kd_{\min}}{G^2\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^5+G^3)\sqrt{G^2+1}}\right).$$

Termination

In order to handle the first two terms in the previous equation and to respect the condition of Lemma 17, let assume that $K > 1$ and $G > 1$ satisfy (C) and the two additional conditions $G^2\sqrt{G^4 - 1} \leq 1$ and $G^2\sqrt{G^2 + 1} \leq K$.

Note that for any $K > \sqrt{2}$, a suitable $G > 1$ may be found, since all conditions are 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 30. Under those conditions, the minimal inter-distance d'_{\min} after the insertion of a new site is bounded from below: $d_{\min} \geq \min\left(\frac{\text{bdr}_{\min}(G)}{(G^5+G^3)\sqrt{G^2+1}}, \frac{\text{ifs}_{\min}^G(C)}{G}\right)$. 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. This proves that the algorithm will not insert sites indefinitely, by a classical volume argument. Moreover, because $(G^2 - 1)K^2 < 1$, the shape condition parametrized by K 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 [6]).

Theorem 31 *Let $K > \sqrt{2}$ be a constant, and let C be a set of constraint segments which bounds a polygonal domain of the plane such that incident segments always form an angle greater than 60° . Under these assumptions, the algorithm presented in Section 5 terminates and provides a triangulation whose dual Voronoi vertices respect the shape bound K .*

PROOF. Let $G > 1$ be such that $(G^2 - 1)(1 + G)^6 K^4 \leq 1$ and $G^2 \sqrt{G^4 - 1} \leq 1$ and $G^2 \sqrt{G^2 + 1} \leq K$. We can also assume that G is close enough to 1, so that incident segments of C always form an angle greater than $2 \arcsin(G^2/2)$. We have seen that such a G can always be found. And if $\text{bdr}_{\min}(G) > 0$, which is always the case if the field of metrics is continuous, the algorithm terminates.

7 Conclusion and Future Work

The approach that we have presented is built upon the work of Labelle and Shewchuk. Instead of using a lower envelop of paraboloids, computed in a greedy way, we rely on a power diagram in higher dimension. As we have shown, we do not need all the combinatorial informations given by such a diagram, but only the zero-dimensional intersections of it with a 2-manifold. Indeed, we present the algorithm by focusing on the overlapping condition on dual triangles, 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.

The simplicity of the structure of our algorithm makes it a good candidate

for an extension to the 3-dimensional case, especially because of the absence of topological considerations. 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. The extension to the 3-dimensional case, while relying on a simple framework, raises interesting issues in terms of complexity of the computation of the restriction of a high dimensional power diagram, and in terms of termination conditions and proper embedding of a three-dimensional triangulation.

References

- [1] Apel, T. and Grosman, S. and Jimack, P.K. and Meyer, A., A new methodology for anisotropic mesh refinement based upon error gradients, *Applied Numerical Mathematics*, 2004, 329–341.
- [2] Borouchaki, H. and George, P.L. and Hecht, F. and Laug, P. and Saltel, E., Delaunay mesh generation governed by metric specifications. Part I algorithms, *Finite Elem. Anal. Des.*, 1997, 61–83.
- [3] Bossen, F. and Heckbert, P., A Pliant Method for Anisotropic Mesh Generation, 5th International Meshing Roundtable, 1996.
- [4] Cheng, S.W. and Dey, T.K. and Ramos, E.A. and Wenger, R., Anisotropic surface meshing, *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithms*, 2006, 202–211.
- [5] Frey, P. and Alauzet, F., Anisotropic mesh adaptation for CFD computations, *Computer Methods in Applied Mechanics and Engineering*, 2005, 5068–5082.
- [6] Labelle, F. and Shewchuk, J.R., Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation, *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, 2003, 191–200.
- [7] Li, X.Y. and Teng, S.H. and Üngör, A., Biting ellipses to generate anisotropic meshes, 8th International Meshing Roundtable, 1999.
- [8] Shewchuk, J.R., What Is a Good Linear Finite Elements Interpolation, Conditioning, Anisotropy, and Quality Measures, *Manuscript*, 2002.
- [9] Shimada, K. and Yamada, A. and Itoh, T., Anisotropic triangular meshing of parametric surfaces via close packing of ellipsoidal bubbles, 6th International Meshing Roundtable, 1997.
- [10] Simpson, R.B., Anisotropic mesh transformations and optimal error control, *Proceedings of the third ARO workshop on Adaptive methods for partial differential equations*, 1994, 183–198.