Locally Uniform Anisotropic Meshing

Jean-Daniel Boissonnat, Camille Wormser, Mariette Yvinec

Appeared in the 24th Annual Symposium on Computational Geometry, SoCG'08

Locally Uniform Anisotropic Meshing*

Jean-Daniel Boissonnat Camille Wormser Mariette Yvinec

INRIA, 2004 route des lucioles
06902 Sophia-Antipolis
FirstName.LastName@inria.fr

Abstract

Anisotropic meshes are triangulations of a given domain in the plane or in higher dimensions, with elements elongated along prescribed directions. Anisotropic triangulations have been shown to be particularly well suited for interpolation of functions or numerical modeling. We propose a new approach to anisotropic mesh generation, relying on the notion of locally uniform anisotropic mesh. A locally uniform anisotropic mesh is a mesh such that the star around each vertex v coincides with the star that v would have if the metric on the domain was uniform and equal to the metric at v. This definition allows to define a simple refinement algorithm which relies on elementary predicates, and provides, after completion, an anisotropic mesh in dimensions 2 and 3.

A practical implementation has been done in the 2D case.

1 Introduction

Anisotropic meshes are triangulations of a given domain in the plane or in higher dimensions, with elements elongated along prescribed directions. Anisotropic triangulations have been shown [10] to be particularly well suited for interpolation of functions or numerical modeling. They allow

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

to minimize the number of triangles in the mesh while retaining a good accuracy in computations. For such applications, the directions along which the elements should be elongated 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 the generation of anisotropic meshes have been proposed. Li et al. [8] and Shimada et al. [12] 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.

Recently, Labelle and Shewchuk [5] have settled the foundations for a rigorous approach based on the so-called anisotropic Voronoi diagram. They have used this geometric structure to compute anisotropic meshes in dimension 2. Their definition of anisotropic Voronoi diagram was used in [1] to provide a direct computation of the dual mesh. However, some kind of flat tetrahedra, called slivers, prevented the extension of these methods to dimension 3. Still, this approach was extended by Cheng et al.[4] for the anisotropic meshing of surfaces embedded in 3D, but this extension is not claimed to be practically implementable.

We propose a new approach for the generation of anisotropic meshes. Given a set of sites V, for each site $v \in V$, computing the Delaunay triangulation $\mathrm{Del}_v(V)$ for the metric M_v is simple, since it is just the image of a Euclidean Delaunay triangulation under a stretching transformation. We define the star S_v of a site v as the set of simplices incident to v in $\mathrm{Del}_v(V)$. With this notation, we can define a *locally uniform anisotropic mesh* as a mesh such that for each site v, the set of elements incident to v in the mesh is exactly its star S_v. Our algorithm allows to build such a *locally uniform anisotropic mesh.*

Initially, there are *inconsistencies* among the stars of the sites, in the sense that it is impossible to merge these stars into a mesh. Then, by adding new points in V at carefully chosen locations, we show how to remove all the inconsistencies. The data structure involved is similar to the one presented by Shewchuk[11], in the context of maintaining triangulations of moving points. Furthermore, the method for guaranteeing termination is inspired by the work of Li and Teng[7] on sliver removal.

Some notable advantages of this new method are:

- programming this algorithm is simple and straightforward, since it relies on the usual Delaunay predicates (applied to some stretched spaces);

- it is valid in 3D;
- in 3D, the termination of the algorithm relies on the sliver removal method of Li and Teng[7], adapted to avoid configurations unsuitable for the algorithm. Consequently, slivers tetrahedra, which are a typical problem for numerical computations, are also avoided without further expense.

In this abstract, we present the results in dimension 3. However most of these results are still true in higher dimensions (and dimension 2), with few or without modifications.

2 Preliminaries

2.1 Anisotropic Metric

We consider a domain $\Omega \subset \mathbb{R}^d$ and assume that each point $p \in \Omega$ is given a symmetric positive definite quadratic form represented by a $d \times d$ matrix M_p, called the metric at p. The distance between two points a and b, as measured by a metric M is defined as

$$d_M(a,b) = \sqrt{(a-b)^t M (a-b)}$$

and we use the notations $d_p = d_{M_p}$, $d_p(a) = d_p(p,a)$ and $d(a,b) = \min(d_a(b), d_b(a))$.

Given the positive definite quadratic form M_p of a point p, we denote by F_p any matrix such that $\det(F_p) > 0$ and $F_p^t F_p = M_p$. The Cholesky decomposition provides such a square root matrix F_p. Note however that F_p is not unique. The Cholesky decomposition provides an upper triangular F_p, while a symmetric F_p can be obtained by diagonalizing the quadratic form M_p and computing the quadratic form with the same eigenvectors and the square root of each eigenvalue.

The Delaunay triangulation $\mathrm{Del}_p(V)$ of a finite set of points V with metric M_p is simply obtained by computing the Euclidean Delaunay triangulation of the stretched image $F_p(V)$, and stretching it back with F_p^{-1}. In the sequel, the points of V, associated with their metrics, are called *sites*, and we refer to the elements of maximal dimension in the triangulation (tetrahedra in 3D) as *simplices*.

Definition 2.1 *Given some metric M, a sphere or a ball computed for M are called M-sphere and M-ball. In the same way, we define the M-circumsphere $\mathcal{C}_M(\tau)$, the M-circumball $\mathcal{B}_M(\tau)$ and the M-circumradius $R_M(\tau)$ of a simplex τ, and the M-volume of a domain.*

Given some metric M, the M-radius-edge ratio $\rho_M(\tau)$ of a simplex τ is the ratio $R_M(\tau)/d_M(\tau)$, where $d_M(\tau)$ denotes its shortest edge, as measured by M.

Note that if M and N are two metrics, an M-sphere is in general an ellipsoid for N. In particular, an M-sphere is an empty Euclidean ellipsoid, elongated along the eigenvectors of M_p. $\mathrm{Del}_p(V)$ is the triangulation of V such that each simplex has an empty M_p-circumsphere. By empty, we mean that the circumsphere contains no site of the triangulation.

2.2 Distortion

The definitions in this section are mostly the ones proposed by Labelle and Shewchuk[5].

Given two metrics M and N, and their square-roots F_M and F_N, the relative *distortion* between M and N is then defined as

$$\gamma(M,N) = \max\{\|F_M F_N^{-1}\|_2, \|F_N F_M^{-1}\|_2\},$$

where $\|\cdot\|_2$ denotes the operator norm associated to the Euclidean metric. Similarly, given two points p and q, the relative *distortion* between p and q is then defined as $\gamma(p,q) = \gamma(M_p, M_q)$.

A fundamental property of $\gamma(p,q)$ is that it bounds the difference between d_p and d_q: for any points x, y, we have $1/\gamma(p,q)\ d_q(x,y) \leq d_p(x,y) \leq \gamma(p,q)\ d_q(x,y)$. The *bounded distortion radius* $\mathrm{bdr}(p,\gamma)$ is the upper bound of numbers ℓ such that for all q and r in Ω, $\max(d_p(q), d_p(r)) \leq \ell \Rightarrow \gamma(q,r) \leq \gamma$. Furthermore, the *minimal bounded distortion radius* associated to γ is $\mathrm{bdr}_{\min}(\gamma) = \inf \mathrm{bdr}(p,\gamma)$, with the minimum taken over all points $p \in \Omega$. Note that this is not exactly the same definition as the one proposed by Labelle and Shewchuk (denoted bdr_{LS} here), but we have

Lemma 2.2 *The two notions of bounded distortion radius are related by the following inequalities:* $\mathrm{bdr}_{LS}(p,\sqrt{\gamma}) < \mathrm{bdr}(p,\gamma) < \mathrm{bdr}_{LS}(p,\gamma)$.

Proof $\mathrm{bdr}_{LS}(p,\gamma)$ is the upper bound of numbers ℓ such that for all q in Ω, $d_p(q) \leq \ell \Rightarrow \gamma(p,q) \leq \gamma$. In particular, if $\ell < \mathrm{bdr}_{LS}(p,\sqrt{\gamma})$, we have that $\max(d_p(q), d_p(r)) \leq \ell$ implies $\gamma(q,r) \leq \gamma(q,p)\gamma(p,r) \leq \gamma^2$. The lower bound follows: $\mathrm{bdr}_{LS}(p, \sqrt{\gamma}) < \mathrm{bdr}(p, \gamma)$. The other inequality is a direct consequence of the definition. $\square$

In dimension 3, each simplex $\tau = abcd$ has four circumspheres $\mathcal{C}_a(\tau)$, $\mathcal{C}_b(\tau)$, $\mathcal{C}_c(\tau)$ and $\mathcal{C}_d(\tau)$. We define the *total distortion* over τ as the maximal distortion between any pairs of points of Ω which are both inside $\mathcal{C}_a(\tau)$ or both inside $\mathcal{C}_b(\tau)$, or $\mathcal{C}_c(\tau)$ or $\mathcal{C}_d(\tau)$. This total distortion is denoted by $\gamma(\tau)$.

In the following, we assume that the domain Ω to be meshed is compact, and that the metric field is continuous over Ω. It follows that $\Gamma = \max_{x,y\in\Omega} \gamma(x, y)$ is finite.

3 Stars and Refinement

We now define the local structures that are built and refined by our algorithm. These definitions rely on the notion of restricted Delaunay triangulation.

Let Ω be a domain of $\mathbb{R}^3$, and let V be a finite set of points of Ω.

Definition 3.1 *The* restriction *to* Ω *of the Delaunay triangulation* $\mathrm{Del}(V)$ *of* V *is the sub-complex of* $\mathrm{Del}(V)$ *consisting of the simplices whose Voronoi dual belongs to* Ω*.*

3.1 Stars

Definition 3.2 *We define the* star S_v *of a site* v *as the set of simplices incident to* v *in* $\mathrm{Del}_v(V)$ *restricted to* Ω*.*

Definition 3.3 *Two stars* S_v *and* S_w *are said to be* inconsistent *if edge* $[vw]$ *appears in only one of the two stars* S_v *and* S_w*. Any simplex containing* $[vw]$ *is also said to be* inconsistent *(see Figure 1).*

Definition 3.4 *The* conflict zone *of a star* S_v *is the union of the balls* $\mathcal{B}_{M_v}(\tau)$ *circumscribing the simplices* τ *that compose* S_v*. We denote it by* Z_v*.*

The following result is a simple property of the Delaunay triangulation:

Lemma 3.5 *The* conflict zone *of a star* S_v *is non-increasing upon insertion of new sites.*

It follows that the star of a site v can be maintained by maintaining a local triangulation around v: to each site v is attached a triangulation $\mathcal{T}_v$, computed as the Delaunay triangulation for metric M_v, and a new site s is inserted into $\mathcal{T}_v$ only if s belongs to the conflict zone of S_v.

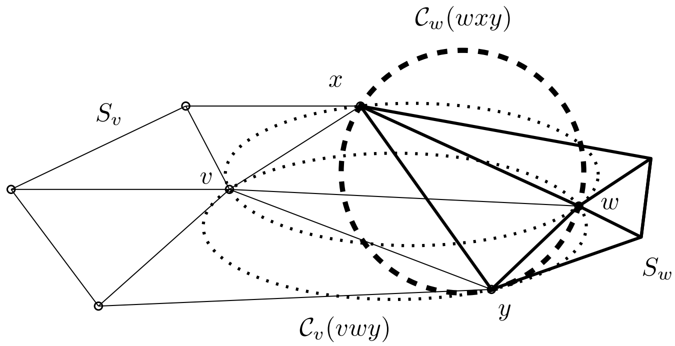

Figure 1: Example of inconsistent stars in 2D: stars S_v and S_w are inconsistent because edge $[vw]$ belongs to S_v but not to S_w.

3.2 Quasi-Cosphericity

Let $\gamma_0 > 1$ be a bound on the distortion. We introduce now the notion of γ_0-cosphericity and show its link with inconsistent simplices.

Definition 3.6 *Five sites a, b, c, d, e are said to be* γ_0-cospherical *for metric M if there exist two metrics N, N' such that*

- $\gamma(M, N) \leq \gamma_0$*,* $\gamma(M, N') \leq \gamma_0$*,* $\gamma(N, N') \leq \gamma_0$*;*
- *the triangulations* $\mathrm{Del}_N(\{a, b, c, d, e\})$ *and* $\mathrm{Del}_{N'}(\{a, b, c, d, e\})$ *are different.*

If γ_0 is implicit, we say that a, b, c, d, e are quasi-cospherical.

See Figure 2 for an illustration in 2D. Note that the five points a, b, c, d, e play symmetric roles in the definition of γ_0-cosphericity. We have the following simple fact:

Lemma 3.7 *Five points a, b, c, d, e are γ_0-cospherical for metric M if there exist two metrics N, N' such that*

- $\gamma(M, N) \leq \gamma_0$*,* $\gamma(M, N') \leq \gamma_0$*,* $\gamma(N, N') \leq \gamma_0$*;*
- *e is outside $\mathcal{C}_N(abcd)$;*
- *e is inside $\mathcal{C}_{N'}(abcd)$.*

Let us now show how the notion of γ_0-cosphericity is related to inconstitencies:

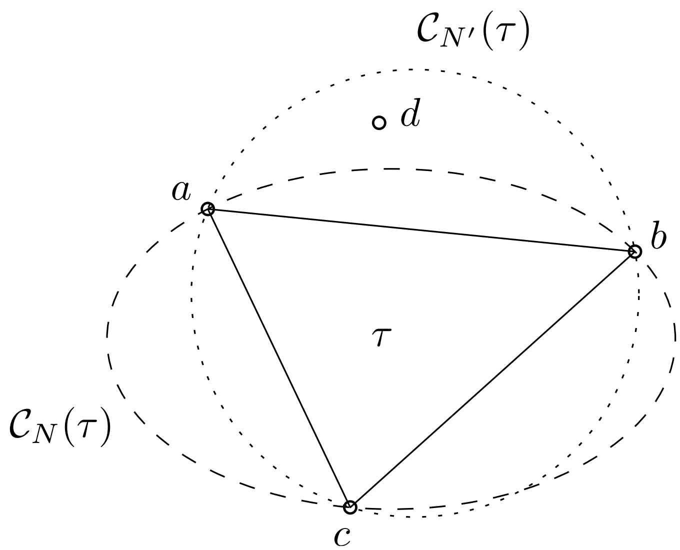

Figure 2: Example of quasi-cospherical points in 2D: a, b, c and d are quasi-cospherical because d is outside of $\mathcal{C}_N(abc)$ but inside $\mathcal{C}_{N'}(abc)$

Lemma 3.8 *Let $\tau = (v, w, x, y)$ be some inconsistent simplex with distortion $\gamma(\tau) < \gamma_0$, which appears in star S_v but not in star S_w. Then there exists a vertex p of S_w such that $\{v, w, x, y, p\}$ are γ_0-cospherical for metric M_v.*

Proof Since $\tau \in S_v$, $\mathcal{C}_v(vwxy)$ is empty. But since $\tau \notin S_w$, there exists some site p of S_w which is inside $\mathcal{C}_w(vwxy)$. It follows that v, w, x, y, p are γ_0-cospherical for metric M_v. $\square$

Definition 3.9 *Given some metric M and five points $x_1, \dots, x_5$ γ_0-cospherical for the metric M, the M-*radius *r of the quasi-cospherical configuration is the minimum of the M-circumradii of the simplices $x_i x_j x_k x_l$, for i, j, k, l distinct integers in $\{1, \dots, 5\}$.*

*The M-*radius-edge ratio *of the quasi-cospherical configuration is the ratio $r/d_{\min}$, where $d_{\min} = \min_{1 \le i \ne j \le 5} d(x_i, x_j)$.*

3.3 Picking Region

The refinement algorithm consists of refining the simplices which do not satisfy the required conditions in terms of size, shape, distortion radius or consistency by inserting a point in the empty circumscribing ball of each bad simplex (the circumscribing ball being computed for the metric of the star currently considered). In the usual Delaunay refinement, this point is simply the circumcenter of the simplex.

However, we cannot guarantee that the consistency problems will disappear if new sites are inserted exactly at the circumcenter of the simplices. As we have seen in the previous section, once the distortion radii of all elements are small, remaining inconsistencies are related to the occurrence of quasi-cospherical configurations. At this point, if the exact circumcenter is inserted, cascading configurations are possible: the refinement could create smaller and smaller inconsistent quasi-cospherical simplices. This is easily seen from the fact that the classical Delaunay refinement cannot get rid of almost flat and cocyclic tetrahedra, called *slivers*. We quantify this by measuring the shortest distance between sites:

Definition 3.10 *The* shortest interdistance $\ell(V)$ *of the set of sites* V *is the shortest distance between pairs of sites of* V *:*

$$\ell(V) = \min_{a,b \in V} d(a,b)$$

In order to prove the termination of the refinement procedure, we need to provide a positive lower bound on $\ell(V)$. In the same way as Li and Teng[7] did for avoiding slivers in 3D Delaunay refinement, we define for each simplex, face and edge (generically called *face* in the sequel) a picking region. Let $\delta < 1$ be a constant to be specified later. If c_τ and r_τ are the M-circumcenter and M-circumradius of a face τ, where M is the metric of some site, we define the M-picking region of τ as the intersection of the M-ball $\mathcal{D}_M(c_\tau, \delta r_\tau)$ with the affine subspace generated by τ. For this reason, δ is called the *picking ratio*.

To avoid cascading constructions, we need to insert a point which is not γ-cospherical with any of the existing simplices. Writing $W(\tau)$ for the set of points that are γ-cospherical with a given simplex τ, we therefore need to bound the M-volume of $W(\tau)$.

Lemma 3.11 *Let* M *be a metric, let* $\tau = vwxy$ *be some simplex with* M*-circumradius* R *and radius-edge ratio smaller than* ρ_0*, and let* $\gamma_0 > 1$ *be a distortion bound. The set* $W(\tau)$ *of points* z *such that* v, w, x, y, z *are* γ_0*-cospherical, with* M*-radius smaller than* βR*, is included in a region of* M*-volume* $V_M < R^3 f(\gamma_0)$*, where* f *is such that* $f(x)$ *tends to* 0 *when* x *tends to* 1*.*

Proof Denote by c_M the center of $\mathcal{C}_M(vwxy)$. Denote by N and N' the two metrics involved in the definition of γ_0-cosphericity. Assume that $vwxy$ is a

Delaunay simplex for metric N. By definition, z is outside $\mathcal{C}_N(vwxy)$ but inside $\mathcal{C}_{N'}(vwxy)$. Denote by c_N and $c_{N'}$ the centers of these circumscribed spheres, for metrics N and N' respectively.

We can assume, without loss of generality, that N is the Euclidean distance. Recall that the Euclidean circumcenter of $vwxy$ can be expressed as

$$c_N = f(v,w,x,y) = v + \frac{c(1,1,1)}{\det(w-v,x-v,y-v)}, \text{ with}$$

$$\begin{aligned} c(1,1,1) &= (y-v)^2(w-v)\times(x-v) \\ &+ (w-v)^2(x-v)\times(y-v) \\ &+ (x-v)^2(y-v)\times(w-v) \end{aligned}$$

Denote now by A a square root of N' (see Section 2.1 for a definition of F_p, the square root of M_p). We can assume that $A = \mathrm{Diag}(\lambda,\mu,\nu)$ with $0 < \lambda \leq \mu \leq \nu \leq \gamma_0$ and $\nu \geq 1/\lambda$ (by changing the frame of coordinates and exchanging N and N' if needed). We then have

$$\begin{aligned} c_{N'} &= A^{-1} f(Av,Aw,Ax,Ay) \\ &= v + \frac{A^{-1}\,\mathrm{Com}(A)}{\det(A)} \frac{c(\mu_1,\mu_2,\mu_3)}{\det(w-v,x-v,y-v)} \\ &= v + A^{-2}\frac{c(\mu_1,\mu_2,\mu_3)}{\det(w-v,x-v,y-v)}, \text{ with} \end{aligned}$$

$$\begin{aligned} c(\mu_1,\mu_2,\mu_3) &= \mu_1(y-v)^2(w-v)\times(x-v) \\ &+ \mu_2(w-v)^2(x-v)\times(y-v) \\ &+ \mu_3(x-v)^2(y-v)\times(w-v) \end{aligned}$$

with $\lambda \leq \mu_1,\mu_2,\mu_3 \leq \nu$ and $\mathrm{Com}(A) = \mathrm{Diag}(\mu\nu,\nu\lambda,\lambda\mu)$. Furthermore, we have

$$c(\mu_1,\mu_2,\mu_3)\cdot(y-v) = \mu_1(y-v)^2\det(w-v,x-v,y-v),$$

and the same formulas with cyclic permutations of w,x,y. It follows that

$$\left\| \frac{c(\mu_1,\mu_2,\mu_3)-c(1,1,1)}{\det(w-v,x-v,y-v)} \right\| \leq 3(\gamma_0-1)\kappa(\rho_0,\beta)2R_N,$$

where $\kappa(\rho_0,\beta) = \sqrt{(2\rho_0)^2 + |\beta^2\gamma_0^2 - 1|}$. The first part of $\kappa(\rho_0,\beta)$ accounts for the angle between the edges of the tetrahedron and the radii of the circumsphere, while the second part bounds the distance in the sliver case, i.e.

the case when the vertices are almost coplanar: in such a case, $c(\mu_1, \mu_2, \mu_3)$ moves in the direction orthogonal to the facets of the tetrahedron.

We define

$$\tilde{c_N} = v + \frac{1}{\det(w-v, x-v, y-v)} c(\mu_1, \mu_2, \mu_3).$$

The triangular inequality then shows that $d_N(c_N, c_{N'}) \leq d_N(c_N, \tilde{c_N}) + d_N(\tilde{c_N}, c_{N'}) \leq 3(\gamma_0 - 1)\kappa(\rho_0, \beta) 2R_N + \|A^{-2} - I\| R_N \leq 7(\gamma_0^2 - 1)\kappa(\rho_0, \beta) R_N$.

Finally, $d_N(c_N, c_{N'}) < C\gamma_0(\gamma_0^2 - 1)R$. Note that this inequality is valid for any metrics N, N' such that the distorsions $\gamma(N, N'), \gamma(M, N), \gamma(M, N')$ are smaller than γ_0. In particular, we may have $M = N$ or $M = N'$.

For metric M, $\mathcal{C}_N(vwxy)$ is an ellipsoid whose minor half-axis is bigger than R/γ_0. It follows from the upper bound of the distance between c_N and $c_{N'}$ that $\mathcal{C}_N(vwxy)$ contains the Euclidean sphere centered at c_M with radius $(1/\gamma_0 - C\gamma_0(\gamma_0^2 - 1))R > (2 - 2\gamma_0 - C\gamma_0(\gamma_0^2 - 1))R$.

Similarly, for metric M, $\mathcal{C}_{N'}(vwxy)$ is an ellipse whose major half-axis is smaller than $\gamma_0 R$. It follows from the upper bound of the distance between c_N and $c_{N'}$ that $\mathcal{C}_{N'}(vwxy)$ is contained in the Euclidean sphere centered at c_M with radius $(\gamma_0 + C\gamma_0(\gamma_0^2 - 1))R$.

Finally, the volume V_M is bounded by $4/3\pi R^3$ $((\gamma_0 + C\gamma_0(\gamma_0^2 - 1))^3 - (2 - 2\gamma_0 - C\gamma_0(\gamma_0^2 - 1))^3) = R^3 f(\gamma_0)$. $\square$

Similarly, we need to bound the M-area of the intersection of $W(\tau)$ with a plane and the M-length of the intersection of $W(\tau)$ with a line: in order to conform the mesh to the prescribed boundary, the algorithm may need to restrict the insertion of a point to a given triangle or segment.

Lemma 3.12 (Plane restriction) *Given a metric M, and a simplex $\tau = vwxy$ with M-circumradius R and radius-edge ratio smaller than ρ_0, and a bound $\gamma_0 > 1$, the set $W(\tau)$ of points z such that v, w, x, y, z are γ_0-cospherical, with M-radius smaller than βR, intersected with a plane π, is included in a region of M-area $V_M < R^2 g(\gamma_0)$, where g is such that $g(x)$ tends to 0 when x tends to 1.*

Lemma 3.13 (Line restriction) *Given a metric M, and a simplex $\tau = vwxy$ with M-circumradius R and radius-edge ratio smaller than ρ_0, and a bound $\gamma_0 > 1$, the set $W(\tau)$ of points z such that v, w, x, y, z are γ_0-cospherical, with M-radius smaller than βR, intersected with a line ℓ, is included in a region of M-length $V_M < Rh(\gamma_0)$, where h is such that $h(x)$ tends to 0 when x tends to 1.*

See Lemmas 3.2.5 and 3.2.6 of [6], for the detailed computations needed for proving Lemmas 3.12 and 3.13.

Lemma 3.14 *Let ρ_0, β be positive bounds, with $\beta < 1$, and let $\epsilon > 0$ be the shortest interdistance. There is at most a constant number $K(\rho_0, \beta)$ of possible new γ_0-cospherical configurations p, q, r, s, t with M_p-radius smaller than βr_τ if a point p is inserted in the picking region $\mathcal{D}(c_\tau, \delta r_\tau)$ of a face τ (see Figure 3).*

Proof Let q, r, s, t be four points such that p, q, r, s, t are γ_0-cospherical for metric M_p and with M_p-radius smaller than βr_τ. Since q, r, s, t are supposed to be at M_p-distance less than $2\beta r_\tau$ from p, a volume argument follows from the fact that all sites have an interdistance greater than ϵ. $\square$

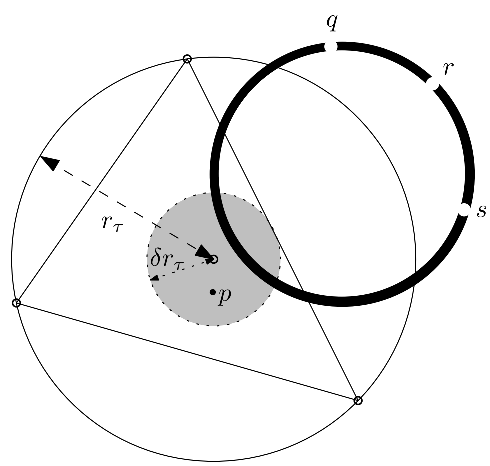

Figure 3: q, r, s define a forbidden region (black annulus) for p in the picking region (grey area)

Lemma 3.15 *If γ_0 is such that $K(\rho_0, \beta)\max(f(\gamma_0), g(\gamma_0), h(\gamma_0))\beta^3 < 4/3\pi\delta^3$, the set of points p that would create new γ_0-cospherical configurations, with radius smaller than βr_τ and radius-edge ratio smaller than ρ_0, does not cover the entire picking region.*

Proof The total area of the set of points that may create such γ_0-cospherical configurations upon insertion of p is smaller than $K(\rho_0, \beta) f(\gamma_0)(\beta r_\tau)^3$. If γ_0 is chosen so that this volume is smaller than the volume $4/3\pi\delta^3 r_\tau^3$ of the picking region, the picking region is not entirely covered.

The same proof remains valid in the case of restricted picking if one replaces f by g and h. $\square$

3.4 Encroachment and Star Initialization

Let us now present how the boundary of the domain is preserved during the refinement process. We assume that the domain Ω to be meshed is a polyhedral domain in dimension 3. By preserving the boundary $\partial\Omega$ of the domain, we mean that the vertices, edges and faces of $\partial\Omega$ appear as elements of the final mesh.

As in the usual Delaunay refinement algorithms, this goal is reached by protecting the boundary $\partial\Omega$ from encroachment by inserted points. Let us recall these notions precisely, in the Euclidean context. See [9] for the original and detailed presentation of this method.

Definition 3.16 *A point p is said to* encroach *a boundary edge or facet f if p is inside the smallest circumscribing sphere of f. This sphere is called the* diametral *sphere of an edge, and the* equatorial *sphere of a facet. This sphere being empty is called the* Gabriel *property for f.*

Maintaining the Gabriel empty property for each boundary edge and facet provides the protection needed for the boundary. Recall that maintaining the Gabriel property of boundary edges and facets upon insertion of a new site v means applying the insertion function `Insert_or_snap_e`(v) defined as follow:

- `GInsert_or_snap_e`(c):
  if c encroaches some boundary edge e, insert the circumcenter of e.
  Otherwise, `Insert_or_snap_f`(c).

- `GInsert_or_snap_f`(c):
  if c encroaches some boundary triangle f, insert the circumcenter of f.
  Otherwise, insert (c).

In this manner, all protected edges and facets do appear in the final mesh and no circumcenter is ever inserted outside the domain.

In our context, we do the same for each of the stars: all constraints are inserted in all stars, and the Gabriel property is maintained in each star for the corresponding metric.

Note that in practice, as soon as the conflict zone Z_v of S_v has an empty intersection with the union of the diametral balls of the constraints, updating S_v is done without taking the constraints into account anymore. This immediately follows from the fact that Z_v is a non increasing set.

This procedure guarantees that boundary facets and edges will be kept all along the algorithm.

4 Algorithm

4.1 Algorithm Outline

The refinement algorithm that we consider constructs the set of sites V in a greedy way while maintaining the set of stars $\{S_v\}_{v \in V}$ and the corresponding sets of constraints E_v whose diametral balls intersect Z_v.

The algorithm refines the simplices of inconsistent stars, until inconsistent stars disappear. Once all stars are consistent, they can be merged together to form a triangulation $\mathcal{T}$ of the domain, with the property that the 1-neighborhood of any vertex v in $\mathcal{T}$ is Delaunay for metric M_v. For this reason, we call the resulting triangulation a *locally uniform anisotropic mesh*.

As we have seen in Section 3.3, simply refining inconsistent simplices by inserting their circumcenter does not allow to maintain a positive insertion radius, which is the condition for the algorithm to terminate. In order to avoid this problem, we manage not to create a forbidden quasi-cospherical configuration by inserting the new site randomly in the picking region around the circumcenter of the simplex to be refined. If the picked point creates any γ-cospherical configuration with γ too small, it is discarded, and a new point is picked in the picking region.

Let $\gamma_0 > 1$, $\delta > 0$, $\rho_0 > 0$ and $\beta > 0$ be constants to be specified in Section 4.2. In order to describe precisely the algorithm, we define the insertion procedures to be used. Face τ is either a simplex, a triangle or an edge:

- `Pick_valid`(τ, M):
  denote by c and r the center and radius of $\mathcal{C}_M(\tau)$. Pick randomly a point x in the picking region $B_M(c, \delta r) \cap H$, where H is the affine subspace spanned by τ. If there exists points p, q, r, s such that $xpqr$ is a new simplex with $\gamma(xpqr) < \gamma_0$ and x, p, q, r, s are γ_0-cospherical with radius smaller than βr_τ and radius-edge ratio smaller than ρ_0, discard x and pick another random point x, until no such points p, q, r, s exist. Return x.

- `Refine`(τ): `Insert_or_snap_e(Pick_valid(`τ, M`))`,
  where M is the metric of the star that is being refined.

- `Insert_or_snap_e`(c):
  if c encroaches some boundary edge e, `Refine`(e).
  Otherwise, `Insert_or_snap_f`(c).

- `Insert_or_snap_f`(c):
  if c encroaches some boundary triangle f, `Refine`(f).
  Otherwise, insert c.

The algorithm consists of applying the following rules. Rule (i) is applied only if Rule (j) with $j < i$ cannot be applied:

Rule (1) **Encroachment:** Refine encroached elements (edges and then triangles) e by calling `Refine`(e).

Rule (2) **Distortion:** If a simplex τ is such that $\gamma(\tau) \geq \gamma_0$, `Refine`$(\tau)$;

Rule (3) **Radius-edge ratio:** If a simplex τ of S_v is such that $\rho_{M_v}(\tau) > \rho_0$, `Refine`$(\tau)$;

Rule (4) **Cosphericity:** If a simplex $\tau = vxyz$ of star S_v is such that there exists a site p such that v, x, y, z, p are γ_0-cospherical for M_v, `Refine`(τ).

Once the algorithm terminates, a simple sweep allows to merge all the stars into the final *locally uniform anisotropic mesh*.

4.2 Termination of the Algorithm and Quality of the Mesh

Let us now prove that the algorithm presented in the previous section does terminate, for suitable choices of distortion bound γ_0, picking ratio δ, radius-edge ratio ρ_0 and size ratio β. Let us consider the refinement rules, in their order of priority.

Lemma 4.1 *Assume that for any boundary edge e, the angle between the two boundary facets incident to e, computed for the metric of any point belonging to e, is greater than $90°$. Then Rule (1) is applied only a finite number of times during the algorithm.*

Proof Once the boundary is sufficiently refined, the diedral angle at any boundary edge, as computed for the metric at any point in the star of its vertices, is greater than $90°$, thanks to the continuity of the metric field. At this point, the usual proofs apply. $\square$

Denote by ϵ_1 the shortest interdistance between sites once Rule (1) cannot be applied anymore. Recall the definition $\Gamma = \max_{x,y\in\Omega} \gamma(x, y)$. Let us now consider the shortest interdistance created by Rule (2):

Lemma 4.2 *Let $\gamma_0 > 0$ be a distortion bound. Denote by r_0 the minimal bounded distortion radius associated to γ_0. Any simplex τ such that $\gamma(\tau) > \gamma_0$ can be refined, while creating no interdistance shorter than $(1-\delta)^3 r_0/(4\Gamma^3)$.*

Proof If a M_x-sphere $\mathcal{C}(x, r)$ has a radius r less than $r_0/2$, then $\gamma(p, q) < \gamma_0$ for any $p, q \in \mathcal{C}$. Let τ be a simplex such that $\gamma(\tau) > \gamma_0$, and denote by a a vertex of τ such that $\gamma(x, y) > \gamma_0$ for two points x, y which are inside $\mathcal{C}_a(\tau)$. It follows that $R_{M_a}(\tau) > r_0/2$. Denote by c_a the center of $\mathcal{C}_a(\tau)$. For any site $w \neq a$, and any point x in the picking region around c_a, we have $d_w(x) \geq d_a(w, x)/\Gamma \geq (d_a(w, c_a) - \delta R_{M_a})/\Gamma$. The Delaunay empty ball property then implies $(d_a(w, c_a) - \delta R_{M_a})/\Gamma \geq (1-\delta)R_{M_a}/\Gamma$ and by the high distortion condition $R_{M_a}(\tau) > r_0/2$, we finally have $(1-\delta)r_a/\Gamma \geq (1-\delta)r_0/(2\Gamma)$.

To summarize, we have proved that $d_w(x) \geq (1-\delta)r_0/(2\Gamma)$. The same lower bound is obviously also valid for $d_x(w)$.

In case boundary elements are encroached, the same proof can be applied to the boundary elements instead of τ, with a penalty of a factor at most $(1-\delta)^2/(2\Gamma^2)$: if a point x, chosen in the picking region of a simplex τ, encroaches a boundary facet f (for a metric M), the distance r_x from x to any site is at most $\sqrt{2}R_M(f)$. Furthermore, as we have seen in the first part of the proof, the point y picked in the picking region of f has a distance r_y to any site of at least $(1-\delta)R_M(f)/\Gamma$. It follows that $r_y \geq (1-\delta)R_M(f)/\Gamma \geq (1-\delta)r_x/(\sqrt{2}\Gamma)$.

Hence, the penalty for one encroachment is a factor of $(1-\delta)/(\sqrt{2}\Gamma)$. It follows that the penalty for two consecutive encroachments (of a face and then of an edge) is a factor of $(1-\delta)^2/(2\Gamma^2)$. This concludes the proof. $\square$

Denote by ϵ_2 the shortest interdistance obtained after Rule (1) and Rule (2) have been applied: $\epsilon_2 = \min(\epsilon_1, (1-\delta)^3 r_0/(4\Gamma^3))$. In the following, we can assume that all simplices have a distortion less than γ_0, and that the interdistance is greater than $\epsilon_2 > 0$. In case simplices with high distortion were to appear again later in the process, the previous lemma shows that we could again refine them and maintain the same bound ϵ_2. Let us now consider the case of simplices with high radius-edge ratio.

Lemma 4.3 *If $(1-\delta)^3\rho_0 > 2\gamma_0^3$, refining the simplices with a radius-edge ratio larger than ρ_0 does not decrease the shortest interdistance.* $\square$

Proof Denote by ϵ the shortest interdistance before the refinement of a simplex with a radius-edge ratio larger than ρ_0. In a way similar to the proof of Lemma 4.2, one computes easily that after the refinement, the shortest interdistance is still greater than $(1-\delta)^3\rho_0\epsilon/(2\gamma_0^3)$. The result follows. $\square$

This proves that applying Rule (3) does not decrease the shortest interdistance. Hence, ϵ_2 remains lower bound of the interdistance. Finally, we can compute how much the interdistance is decreased when Rule (4) is applied.

Lemma 4.4 *Let $\tau = vxyz$ be a simplex of star S_v with a site p such that v, x, y, z, p are γ_0-cospherical for M_v. Refining all such configurations does not create any interdistance shorter than $(1-\delta)^3\epsilon_2/2$ if $(1-\delta)^3\beta > 2\gamma_0^3$.*

Proof Denote by ϵ the current shortest interdistance. In a way similar to the proof of Lemma 4.2, one computes easily that the shortest interdistance after the refinement of such a γ_0-cospherical configuration stays bigger than $(1-\delta)^3\epsilon/(2\gamma_0^3)$.

Recall that, thanks to the definition of `Pick_valid`, no γ_0-cospherical configuration is ever created by the refinement of any simplex τ, except γ_0-cospherical configurations with radius bigger than βr_τ or radius-edge ratio bigger than ρ_0. If the radius-edge ratio is bigger than ρ_0, the configuration is to be refined by Rule (3). As we have just seen, if the radius is bigger than βr_τ, the shortest interdistance created to refine this new γ_0-cospherical configuration is at least $(1-\delta)^3\beta r_\tau/(2\gamma_0^3)$. Hence, if we choose β large enough, so that $(1-\delta)^3\beta > (2\gamma_0^3)$, refining this kind of new γ_0-cospherical configuration does not reduce the shortest interdistance.

It follows that $(1-\delta)^3\epsilon_2/2$ is a lower bound on the interdistance after applying Rule (4), under the condition that $(1-\delta)^3\beta > (2\gamma_0^3)$. $\square$

Lemma 4.1, 4.2, 4.3 and 4.4 show that the insertion radius admits a positive lower bound. This concludes the proof of the termination of the algorithm. Let us summarize this result in the following theorem, which also relies on Lemma 3.14:

Theorem 4.5 *Given a polyhedral domain Ω and a continuous metric field over Ω, and the following properties for the parameters of the algorithm,*

- *the angle at each boundary edge e, computed for the metric of any point of e, is greater than $90°$;*
- *ρ_0 is larger than 2;*
- *δ is small enough, so that $(1-\delta)^3\rho_0 > 2$;*
- *β is large enough, so that $(1-\delta)^3\beta > 2$;*

- *γ_0 is close enough to* 1*, so that* $K(\rho_0, \beta) \max(f(\gamma_0), g(\gamma_0), h(\gamma_0))\beta^2 < 4/3\pi\delta^2$ *and* $(1-\delta)^3\beta > 2\gamma_0^3$ *and* $(1-\delta)^3\rho_0 > 2\gamma_0^3$.

the refinement algorithm terminates, with a lower bound ρ_0 on the radius-edge ratio of the elements and an upper bound γ_0 on the distortion of the simplices. $\square$

Note that these bounds ρ_0 and γ_0 ensure that eventually all simplices are well-shaped for the metrics of their vertices. This guarantees the quality of the final mesh.

5 Conclusion

We have proposed a new definition for an anisotropic mesh and an algorithm to generate such a mesh. The algorithm is the first to offer guarantees in 3-space. Moreover, the algorithm is simple and has been implemented in the plane in C++ using CGAL.

Although the implementation has not been optimized, we had still a much more scalable algorithm than the one we proposed in [1]: our datastructure has asymptotically the same space complexity as a triangulation of the same pointset. Interestingly, the assumption that the metric field was continuous appeared crucial not only in theory, but also in practical tests: discontinuities typically prevent the algorithm from terminating, because the algorithm refines the locus of the discontinuity (usually a curve) indefinitely.

Figure 6 shows the output of the algorithm on a domain where the metric is stretched horizontally in the upper part and vertically in the lower part. In this example, we did not enforce any size bound, so that the variable density of the result clearly shows where more refinement was needed for removing inconsistencies. As expected, the higher densities are located along the line of high distorsion, where the eigenvectors exchange their eigenvalues.

Future directions of work include

- allowing more general constraints, in particular constraints with sharp edges, and using a protection scheme to avoid cascading insertions in the neighborhood of these edges;
- dealing with discontinuities by protecting points of discontinuity and by considering the curves of discontinuity as constraints of the triangulation;
- providing a 3D implementation;

- extending the results in dimension $d > 3$.

6 Acknowledgments

The authors would like to thank Tamal Dey for fruitful discussions.

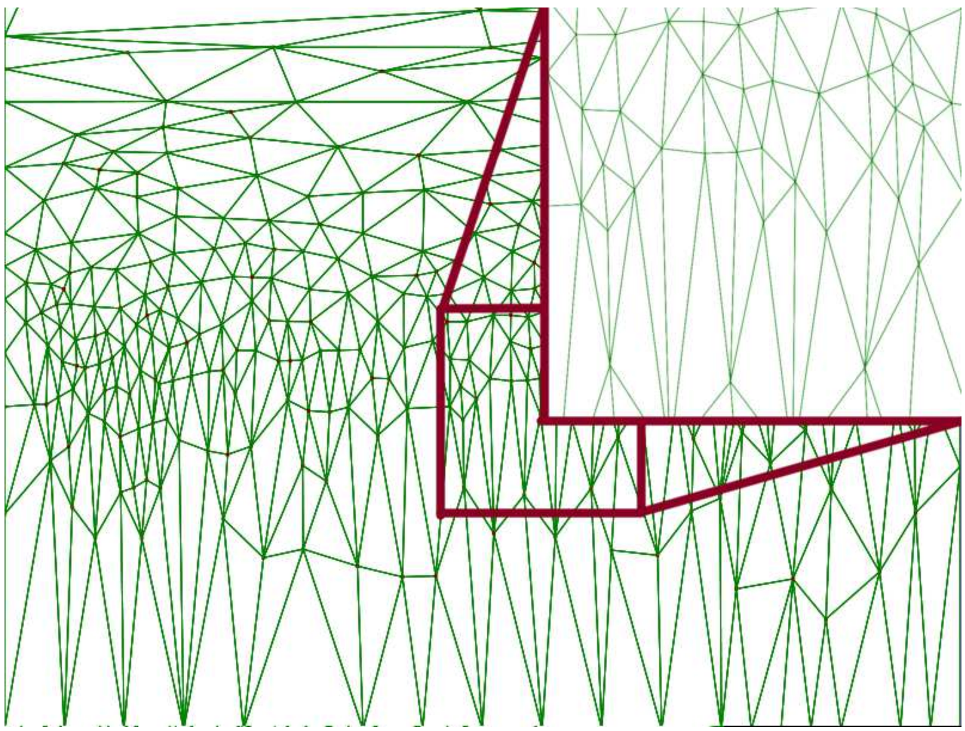

Figure 4: Simple example of the output of the algorithm. The metric changes vertically. The upper right square represents a zoom of the small square in the middle.

References

[1] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic diagrams: Labelle shewchuk approach revisited. *Theoretical Computer Science, to appear.*

[2] H. Borouchaki, P. L. George, F. Hecht, P. Laug, and E. Saltel. Delaunay mesh generation governed by metric specifications. part I algorithms. *Finite Elem. Anal. Des.*, 25(1-2):61–83, 1997.

[3] F. Bossen and P. Heckbert. A pliant method for anisotropic mesh generation. In *5th International Meshing Roundtable*, October 1996.

[4] S.-W. Cheng, T. K. Dey, E. A. Ramos, and R. Wenger. Anisotropic surface meshing. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 202–211, New York, NY, USA, 2006. ACM.

[5] F. Labelle and J. R. 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.

[6] X.-Y. Li. Sliver-free three dimensional delaunay mesh generation. PhD thesis, University of Illinois at Urbana-Champaign 2000.

[7] X.-Y. Li and S.-H. Teng. Generating well-shaped delaunay meshed in 3d. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 28–37, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.

[8] X.-Y. Li, S.-H. Teng, and A. Üngör. Biting ellipses to generate anisotropic mesh. In *8th International Meshing Roundtable*, October 1999.

[9] J. R. Shewchuk. Tetrahedral mesh generation by delaunay refinement. In *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry*, pages 86–95, New York, NY, USA, 1998. ACM.

[10] J. R. Shewchuk. What is a good linear finite element? Interpolation, conditioning, anisotropy, and quality measures. In *http://www.cs.cmu.edu/~jrs/jrspapers.html*, Manuscript 2002.

[11] R. Shewchuk. Star splaying: an algorithm for repairing delaunay triangulations and convex hulls. In *SCG '05: Proceedings of the twenty-first annual symposium on Computational geometry*, pages 237–246, New York, NY, USA, 2005. ACM.

[12] K. Shimada, A. Yamada, and T. Itoh. Anisotropic Triangulation of Parametric Surfaces via Close Packing of Ellipsoids. *Int. J. Comput. Geometry Appl*, 10(4):417–440, 2000.