Anisotropic Delaunay Mesh Generation

Jean-Daniel Boissonnat, Camille Wormser, Mariette Yvinec

► To cite this version:

Jean-Daniel Boissonnat, Camille Wormser, Mariette Yvinec. Anisotropic Delaunay Mesh Generation. [Research Report] RR-7712, 2011, pp.37. inria-00615486v1

HAL Id: inria-00615486

<https://hal.inria.fr/inria-00615486v1>

Submitted on 19 Aug 2011 (v1), last revised 19 Dec 2014 (v2)

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

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

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Anisotropic Delaunay Mesh Generation

Jean-Daniel Boissonnat — Camille Wormser — Mariette Yvinec

N° 7712

Août 2011

*apport
de recherche*

Anisotropic Delaunay Mesh Generation

Jean-Daniel Boissonnat , Camille Wormser , Mariette Yvinec

Theme :
Équipe-Projet Geometrica

Rapport de recherche n° 7712 — Août 2011 — 34 pages

Abstract: Anisotropic meshes are triangulations of a given domain in the plane or in higher dimensions, with elements elongated along prescribed directions. Anisotropic triangulations are known to be well suited for interpolation of functions or solving PDEs. Assuming that the anisotropic shape requirements for mesh elements are given through a metric field varying over the domain, we propose a new approach to anisotropic mesh generation, relying on the notion of anisotropic Delaunay meshes. An anisotropic Delaunay mesh is defined as a mesh in which the star of each vertex v consists of simplices that are Delaunay for the metric associated to vertex v . This definition works in any dimension and allows to define a simple refinement algorithm. The algorithm takes as input a domain and a metric field and provides, after completion, an anisotropic mesh whose elements are shaped according to the metric field.

Key-words: mesh generation, anisotropic meshes, Delaunay triangulation

This research has been partially supported by the Agence Nationale de la Recherche (project GIGA).

Génération de maillages de Delaunay anisotropes

Résumé : Les maillages anisotropes sont des triangulations d'un domaine donné du plan ou d'un espace de plus grande dimension dont les éléments sont étirés selon des directions prescrites. Les maillages anisotropes sont utiles pour interpoler des fonctions ou résoudre des EDP. Dans cet article, nous supposons que l'anisotropie est prescrite par un champ de métrique défini sur le domaine à mailler. Nous proposons une nouvelle approche de génération de maillages anisotropes qui s'appuie sur la notion de maillage de Delaunay anisotrope. Un tel maillage est défini comme un maillage dont l'étoile de chaque sommet v est formée de simplexes qui sont de Delaunay pour la métrique de v . Cette définition est valide en toutes dimensions et un tel maillage peut être construit par un algorithme simple de raffinement.

Mots-clés : Génération de maillages, maillages anisotropes, triangulation de Delaunay

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 to be particularly well suited for interpolation of functions [15, 31] and for solving PDEs [5]. They allow to minimize the number of elements in the mesh while retaining a good accuracy in computations.

The required anisotropy is generally described through a metric field defined over the domain to be meshed. The directions along which the elements should be elongated are usually given, at each point of the domain, as a quadratic form. The eigenvectors and eigenvalues of the quadratic form describe the preferred directions and their anisotropic ratios.

Two main issues arise in this context. The first is to define the metric field. The second one is to generate a mesh whose elements are shaped according to the chosen metric field.

Defining good metric fields and error estimates is still an active research area. Alauzet et al. introduced the notion of continuous metrics and continuous meshes to minimize interpolation error [3, 26, 2]. Loseille et al. [27] applied this notion to a posteriori error estimates in order to minimize the approximation error during the solution process of some PDEs. Chen et al [11] considered anisotropic finite element approximation of functions in the L^p norm. Their result reveals that the accuracy of the approximation is governed by a quantity that depends non linearly on the hessian of the function. In his thesis, Mirebeau extends this result to finite elements of arbitrary degree and to Sobolev norms, and provides sharp asymptotic error estimates for the approximation of functions of two variables. The Sobolev norms $W^{1,p}$ measure the error on the gradient of the function.

Various methods have been proposed to generate anisotropic meshes whose elements are shaped according to a given metric field. In their early work on 2D meshes, Bossen and Heckbert [10] proposed to adapt in the anisotropic setting their *pliant* method for mesh generation. Starting from a constraint Delaunay triangulation, the pliant method performs local optimization operations including centroidal smoothing and retriangulation, with possible insertion or removal of vertices. Li et al. [25] and Shimada et al. [33] have proposed to place the mesh vertices close to the centers of ellipsoid bubbles optimally packed in the domain. Borouchaki et al. [9] proposed to adapt the standard Delaunay incremental construction to the anisotropic context. This construction is then combined with an anisotropic version of the *unit mesh* approach that aims at producing meshes whose edges have unit length. Lengths, in the anisotropic case, are measured in the Riemannian metric provided by the metric field. The efficiency of the method has been demonstrated in various contexts [20, 17].

Following a different line of research, some attempts have been done recently to define anisotropic Delaunay triangulation and meshes as the duals of some Voronoi diagrams derived from the metric field. Leibon and Letscher [22] introduced the Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. This approach requires to compute geodesic paths and intrinsic balls, which may be quite complicated in practice. A simpler approach has been proposed by Labelle and Shewchuk [21]. They define an anisotropic mesh as the dual of the so-called anisotropic Voronoi diagram. The sites of this diagram are the mesh vertices and the distance to a site is computed with respect to the metric attached to this site. In the 2-dimensional case, Labelle and Shewchuk have proposed a refinement algorithm that can provably produce anisotropic meshes. Their approach has somehow been simplified in [7], leading to a direct computation of the dual mesh, and extended by Cheng et al.[12] to produce anisotropic meshes of surfaces embedded in 3D. Extending Labelle and Shewchuk's approach to higher dimensions seems however difficult due to the presence of flat

tetrahedra called slivers [30]. Du and Wang [18] have proposed to use a definition of anisotropic Voronoi diagrams which is somehow symmetric to Labelle and Shewchuk's one. The Voronoi regions are based on distances from points to sites that are computed with respect to the metric of the point. Du and Wang compute centroidal Voronoi diagrams using this definition and show experimentally that the dual structures are generally anisotropic meshes of high quality. However they could not provide theoretical guarantees nor conditions that ensure that the dual structure is a valid triangulation.

In this paper, we introduce a new notion of anisotropic mesh which extends nicely in any dimension and is simple to compute. As in the previous approaches, we assume that the anisotropy is prescribed by a metric field that associates to each point p of the domain a symmetric positive definite square matrix M_p , describing the metric at point p . Given a set of points V called *sites*, we consider, for each site $v \in V$, the Delaunay triangulation $\text{Del}_v(V)$ of V , computed for the metric M_v attached to location of v . The triangulation $\text{Del}_v(V)$ is easy to compute: it is just the image of a standard Euclidean Delaunay triangulation under a stretching transformation. For each site $v \in V$, we keep the *star* S_v of v in $\text{Del}_v(V)$, i.e. the set of simplices of $\text{Del}_v(V)$ that are incident to v . The collection of stars is called the *star set* of V . In general, there are *inconsistencies* among the stars : a simplex τ may appear in the stars of some of its vertices without appearing in the stars of all of them. As a result, the simplices in the star set of V do not form a triangulation of V . However, we show in this paper that, given a compact domain of $\mathbb{R}^d$ and a smooth metric field, one can insert new sites in V at carefully chosen locations so that all inconsistencies are removed. The simplices in the star set then form a d -triangulation of V that we call an *anisotropic Delaunay mesh*.

The idea of maintaining independent stars for each vertex of a mesh has been first proposed by Shewchuk [32] for maintaining triangulations of moving points. The star set was even used [30] to build the dual of an anisotropic Voronoi diagram as defined by Labelle and Shewchuk. The method we used to ensure the star consistency is inspired by the work of Li and Teng [24, 23] for removing slivers in isotropic meshes. In our context, the method is extended so as to take into account the metric distortion between neighboring stars and also to avoid, in addition to slivers, more general quasi-cospherical configurations that may prevent the termination of the algorithm.

In addition to conforming to the given anisotropic metric field, this mesh generation method has several notable advantages over previous methods.

- It is not limited to the plane and works in any dimension;
- It is easy to implement. Through a stretching transform, the star of each vertex in the mesh can be computed as part of an Euclidean Delaunay triangulation. Therefore the algorithm relies only on the usual Delaunay predicates (applied in some stretched spaces);
- The star of each vertex in the output mesh is formed with simplices that are Delaunay with respect to the metric of the central vertex. This provides a neat characterization of the output mesh from its set of vertices.
- The algorithm provides theoretical guarantees about the size and shape of the output mesh elements. Each element is guaranteed to be well shaped according to the metric of all its vertices.

A preliminary version of this work, limited to the 3-dimensional case, has been presented at the Symposium on Computational Geometry [8].

2 Preliminaries

2.1 Anisotropic Metric

An anisotropic metric in $\mathbb{R}^d$ is defined by a symmetric positive definite quadratic form represented, in some vector basis, by a $d \times d$ matrix M . The distance between two points a and b , as measured by metric M is defined as

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

This definition provides a definition for M -lengths and, by integration, for higher dimensional M -volume measures.

In the following, we often use the same notation, M , for a metric and the associated matrix in a given basis. Given the symmetric positive definite matrix M , we denote by F_M any matrix such that $\det(F_M) > 0$ and $F_M^t F_M = M$. Note however that F_M is not unique. The Cholesky decomposition provides an upper triangular F_M , while a symmetric F_M can be obtained by diagonalizing the quadratic form M and computing the quadratic form with the same eigenvectors and the square root of each eigenvalue.

Note that

$$d_M(a, b) = \sqrt{(a - b)^t F_M^t F_M (a - b)} = \|F_M(a - b)\| \quad (1)$$

where the notation $\|\cdot\|$ stands for the Euclidean norm. Equation (1) proves that distance d_M enjoys the standard triangular inequality. In the following we call $F_M p$ the *stretching transform* of p .

Given some metric M , an M -sphere $\mathcal{C}_M(c, r)$, with center c and radius r , is defined as the set of points p such that $d_M(c, p) = r$, and likewise an M -ball $\mathcal{B}_M(c, r)$, is defined as the set of points p such that $d_M(c, p) \leq r$. Note that an M -sphere is in general an Euclidean ellipsoid, with its axes aligned along the eigenvectors of M .

Given a k -simplex τ in $\mathbb{R}^d$ and a metric M , we define the M -circumsphere $\mathcal{C}_M(\tau)$ as the circumscribing M -sphere of τ with smallest radius. The M -circumball $\mathcal{B}_M(\tau)$ is the M -ball bounded by $\mathcal{C}_M(\tau)$ and the M -circumradius $r_M(\tau)$ of a simplex τ is the radius of $\mathcal{C}_M(\tau)$. Equation (1) shows that $\mathcal{C}_M(\tau)$ is the reciprocal image $F_M^{-1}(\mathcal{C}(F_M(\tau)))$ of the Euclidean circumscribing sphere of simplexe $F_M(\tau)$.

Let M be a metric and V be a set of points, called *sites*. The Delaunay triangulation of V for metric M , denoted $\text{Del}_M(V)$, is the triangulation of V such that the interior of the M -circumball of each d -simplex is *empty*, i.e. contains no site of V . Owing to equation (1), the Delaunay triangulation $\text{Del}_M(V)$ of a finite set of points V for metric M is simply obtained by computing the Euclidean Delaunay triangulation of the stretched image $F(V) = \{F_M v, v \in V\}$, and stretching the result back with F_M^{-1} . The Delaunay triangulation $\text{Del}_M(V)$ is thus viewed as the dual of a stretched Voronoi diagram. Alternatively, $\text{Del}_M(V)$ can be computed as a weighted Delaunay triangulation. Indeed,

$$\begin{aligned} d_M(x, a) \leq d_M(x, b) &\Leftrightarrow x^2 - 2a^t M x + a^t M a \leq x^2 - 2b^t M x + b^t M b \\ &\Leftrightarrow \|x - a'\|^2 - w_{a'} \leq \|x - b'\|^2 - w_{b'} \end{aligned}$$

where $a' = Ma$, $b' = Mb$, $w_{a'} = a^t(M^2 - M)a$ and $w_{b'} = b^t(M^2 - M)b$. It follows that $\text{Del}_M(V)$ is the weighted Delaunay triangulation dual to the weighted Voronoi diagram of V' , where $V' = \{v' = Mv, v \in V\}$ with weights $w_{v'} = v^t(M^2 - M)v$.

2.2 Metric Field and Distortion

In the rest of the paper, we consider a compact domain $\Omega \subset \mathbb{R}^d$ and assume that we are given a metric field defined over Ω , i.e. a metric M_x is given at each point $x \in \Omega$.

In the following, to avoid double subscripts, we replace subscript M_x by x and simply write Y_x for Y_{M_x} . Hence, we will write for instance F_x for F_{M_x} and $d_x(a, b)$ for $d_{M_x}(a, b)$.

We recall some definitions due to Labelle and Shewchuk [21].

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

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

where $\|\cdot\|$ is the matrix norm operator associated with the Euclidean norm, i.e. for a $d \times d$ square matrix A , $\|A\| = \sup_{x \in \mathbb{R}^d} \frac{\|Ax\|}{\|x\|}$. In the context of a metric field, the relative *distortion* between two points p and q of the domain Ω is defined as $\gamma(p, q) = \gamma(M_p, M_q)$. Observe that $\gamma \geq 1$ and is equal to 1 iff $M_p = M_q$.

A fundamental property of $\gamma(p, q)$ is that it bounds the ratio between d_p and d_q :

$$\forall x, y, 1/\gamma(p, q) d_q(x, y) \leq d_p(x, y) \leq \gamma(p, q) d_q(x, y).$$

A d -simplex $\tau = p_0 p_1 \dots p_d$ has $d+1$ circumballs $\mathcal{B}_i(\tau)$, $i = 0, \dots, d$, where $\mathcal{B}_i(\tau)$ is the circumball of τ in the metric M_i attached to vertex p_i . The *distorsion distortion* $\gamma(B)$ of a ball B is defined as maximal distortion between any pairs of points of $B \cap \Omega$. We define the *distortion* $\gamma(\tau)$ of a simplex τ as the maximum of the distortion of its circumballs:

$$\gamma(\tau) = \max\{\gamma(\mathcal{B}_i(\tau)), i = 1, \dots, d+1\}.$$

2.3 Sizing field

In this paper, we will assume that the metric field is smooth over the domain Ω . The *distorsion* $\gamma(p, q)$ is then a continuous function and the maximum distortion over Ω , $\Gamma = \sup_{x, y \in \Omega} \gamma(x, y)$, is finite since Ω is compact.

We now consider a local view of the *distorsion*. Given a constant $\gamma_0 > 1$, called the *distortion bound*, we define for each point $p \in \Omega$ the *bounded distortion radius*, $\text{bdr}(p, \gamma_0)$, as the upper bound on distances ℓ such that for all q and r in Ω , $\max(d_p(p, q), d_p(p, r)) \leq \ell \Rightarrow \gamma(q, r) \leq \gamma_0$.

Lemma 2.1 (The bounded distortion radius lemma) *The bounded distortion radius $\text{bdr}(p, \gamma_0)$ enjoys the following property for any p, q in Ω :*

$$\frac{1}{\gamma(p, q)} [\text{bdr}(p, \gamma_0) - d_p(p, q)] \leq \text{bdr}(q, \gamma_0) \leq \gamma(p, q) [\text{bdr}(p, \gamma_0) + d_p(p, q)].$$

Proof Let x, y be any two points in Ω so that:

$$d_q(q, x) \leq \frac{1}{\gamma(p, q)} (\text{bdr}(p, \gamma_0) - d_p(p, q)), \quad (2)$$

$$d_q(q, y) \leq \frac{1}{\gamma(p, q)} (\text{bdr}(p, \gamma_0) - d_p(p, q)). \quad (3)$$

Then, we have, using the triangular inequality,

$$d_p(p, x) \leq d_p(p, q) + d_p(q, x) \leq d_p(p, q) + \gamma(p, q) d_q(q, x) \leq \text{bdr}(p, \gamma_0)$$

and, similarly,

$$d_p(p, y) \leq \text{bdr}(p, \gamma_0).$$

Then, by definition of the bounded distorsion radius, $\gamma(x, y) \leq \gamma_0$. Because the last inequality is true for any pair of points x, y satisfying inequalities (2) and (3), we conclude that

$$\frac{1}{\gamma(p, q)} [\text{bdr}(p, \gamma_0) - d_p(p, q)] \leq \text{bdr}(q, \gamma_0). \quad (4)$$

To prove the second inequality of Lemma 2.1, we simply write inequality (4) for the pair (q, p) , which yields:

$$\frac{1}{\gamma(p, q)} [\text{bdr}(q, \gamma_0) - d_q(p, q)] \leq \text{bdr}(p, \gamma_0)$$

from which we deduce

$$\begin{aligned} \text{bdr}(q, \gamma_0) &\leq \gamma(p, q) \text{bdr}(p, \gamma_0) + d_q(p, q) \\ &\leq \gamma(p, q) [\text{bdr}(p, \gamma_0) + d_q(p, q)]. \end{aligned}$$

□

We will further assume that the bounded distorsion radius has a strictly positive lower bound on domain Ω : $\min_{p \in \omega} \text{bdr}(p, \gamma_0) > 0$.

In our algorithm, we will use $\text{bdr}(p, \gamma_0)$ as a sizing field to adapt the mesh density to the variation of the anisotropic metric. In fact, as we will see, our algorithm may use more general sizing fields and to possibly take into account other sizing criteria.

Definition 2.2 (Smooth sizing field) *Let $\gamma_0 \geq 1$ be a given distorsion bound. We call sizing field and denote by $\text{sf}(p, \gamma_0)$ (or $\text{sf}(p)$ for short if γ_0 is understood), any function defined over the domain Ω , that satisfies the three following conditions:*

$$\text{positiveness} \quad \forall x \in \Omega, \quad \exists \text{sf}_0 > 0, \quad \text{sf}(x, \gamma_0) \leq \text{sf}_0 \quad (5)$$

$$\text{distorsion} \quad \forall x \in \Omega, \quad \text{sf}(x, \gamma_0) \leq \text{bdr}(x, \gamma_0) \quad (6)$$

$$\text{smoothness} \quad \forall x, y \in \Omega, \quad (7)$$

$$\frac{1}{\gamma(x, y)} [\text{sf}(x, \gamma_0) - d_x(x, y)] \leq \text{sf}(y, \gamma_0) \leq \gamma(x, y) [\text{sf}(x, \gamma_0) + d_x(x, y)]$$

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 $\mathbb{R}^d$ and let V be a finite set of points of Ω that are called hereafter *sites* or *vertices*.

The *restriction* to Ω of the Delaunay triangulation $\text{Del}((V))$ of V is the subcomplex of $\text{Del}(V)$ whose maximal faces are the d -simplices of $\text{Del}(V)$ that have their dual Voronoi vertices inside the

domain Ω . The natural extension to anisotropic Delaunay triangulations $\text{Del}_M(V)$ would be to define the restriction of $\text{Del}_M(V)$ to Ω as the subcomplex of $\text{Del}_M(V)$ whose maximal faces are the d -simplices τ of $\text{Del}(V)$ that have their M -circumcenter inside Ω .

However, for technical reasons, in the current framework of anisotropic metric fields, we need to be more restrictive. We define the *restriction* to Ω of a Delaunay triangulation $\text{Del}_M(V)$ of V as the subcomplex of $\text{Del}_M(V)$ whose maximal faces are the d -simplices τ of $\text{Del}_M(V)$ such that, for each vertex p_i of τ the dual circumcenter $c_i(\tau)$ for the metric M_i of p_i is inside Ω .

3.1 Stars and Inconsistencies

For each site v in V , we consider the Delaunay triangulation $\text{Del}_v(V)$ of V for the metric M_v . We define the *star* S_v of site v as the set of d -simplices incident to v in the restriction of $\text{Del}_v(V)$ to Ω .

The collection of all the stars S_v , $v \in V$, is called the *star set* of V .

Two stars S_v and S_w are said to be *inconsistent* if some simplex incident to v and w appears in only one of the two stars S_v and S_w . Any simplex that appear in the stars of some but not all of its vertices is also said to be *inconsistent* (see Figure 1).

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 .

Our algorithm incrementally inserts new sites in V and updates the star set $\{S_v, v \in V\}$ until there are no more inconsistencies. As shown below, when the mesh is dense enough with respect to the variation of the metric field, inconsistencies are related to the occurrence of special configurations of subsets of sites that are called *quasi-cospherical*. The algorithm will therefore aim at avoiding those *quasi-cospherical* configurations. As it turns out, it is possible to avoid *quasi-cospherical* configurations only when even more special configurations called *slivers* are avoided first. Both notions are now defined.

3.2 Slivers

The following definitions are taken from Li [23] and are easily extended to anisotropic metrics.

Let τ be a k -simplex. We denote by $C_p(\tau)$ the M_p -circumsphere of τ , by $r_p(\tau)$ the M_p -circumradius of τ and by $e_p(\tau)$ the M_p -length of the shortest edge of τ for the metric M_p .

We define two quality measures of τ for metric M_p . The M_p -radius-edge ratio is defined as the ratio $\rho_p(\tau) = r_p(\tau)/e_p(\tau)$. The *sliverity ratio* $\sigma_p(\tau)$ is the ratio $(\text{Vol}_p(\tau)/e_p^k(\tau))^{\frac{1}{k}}$ where $\text{Vol}_p(\tau)$ is the M_p -volume of τ .

Definition 3.1 (Sliver) Let ρ_0 and σ_0 be two positive constants and let M_p be a metric. A k -simplex τ is said to be

- well-shaped for M_p , if $\rho_p(\tau) \leq \rho_0$ and $\sigma_p(\tau) \geq \sigma_0$
- a sliver for M_p , if $\rho_p(\tau) \leq \rho_0$, $\sigma_p(\tau) < \sigma_0$
- a k -sliver for M_p , if it is a sliver and all its $(k-1)$ -dimensional faces are well-shaped.

It is easily shown that any k -dimensional simplex that is a sliver is either a k -sliver or include as a subspace a k' -sliver for $k' < k$.

The following lemma is known for slivers in dimension 3, see e.g. [19]. It has been extended to higher dimensions [23] and extends naturally to anisotropic metrics as proved in the appendix.

Lemma 3.2 (Sliver lemma) Let τ be a k -simplex and M a metric. If v is a vertex of τ , we denote by $\tau(v)$ the $(k-1)$ -face of τ opposite to vertex v , by $\text{aff}(\tau(v))$ the affine hull of $\tau(v)$, i.e. the $(k-1)$ -flat spanned by the vertices of $\tau(v)$, by $\mathcal{C}(v)$ the M -circumsphere of $\tau(v)$, and by $r(v)$ the M -radius of $\mathcal{C}(v)$.

If τ is a k -sliver with respect to M , the M -distance from v to $\text{aff}(\tau(v))$ is at most $2k\sigma_0 r(v)$ and the M -distance from v to $\mathcal{C}(v) \cap \text{aff}(v)$ is at most $4\pi k\rho_0\sigma_0 r(v)$.

3.3 Quasi-Cosphericity

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

Definition 3.3 (Quasi-cospherical configuration) A subset U of $d+2$ sites $\{p_0, p_1, \dots, p_{d+1}\}$ is said to be a (γ_0, M) -cospherical configuration if there exist two metrics N and N' such that :

- $\gamma(M, N) \leq \gamma_0$, $\gamma(M, N') \leq \gamma_0$ and $\gamma(N, N') \leq \gamma_0$;
- the triangulations $\text{Del}_N(U)$ and $\text{Del}_{N'}(U)$ are different.

Metrics N and N' are said to witness U . If M is clear from the context, we simply say that U is a γ_0 -cospherical configuration and if both M and γ_0 are understood, we say that U is a quasi-cospherical configuration.

See Figure 2 for an illustration in the plane. Note that the $d+2$ points in U play symmetric roles in the above definition. In the sequel, U will often consist of the set of vertices of a d -simplex τ belonging to the star S_v of some site $v \in V$, together with an additional site p of V . In such a case, we write $U = (\tau, p)$.

From Radon theorem, there are only two distinct triangulations of $U = (\tau, p)$ and any d -simplex with vertices in U belongs to exactly one of them. Therefore, we have the following easy lemma.

Lemma 3.4 A configuration (τ, p) is (γ_0, M) -cospherical iff there exist two metrics N and N' such that

- $\gamma(M, N) \leq \gamma_0$, $\gamma(M, N') \leq \gamma_0$ and $\gamma(N, N') \leq \gamma_0$;
- p belongs to the interior of exactly one of the two circumballs $\mathcal{B}_N(\tau)$ and $\mathcal{B}_{N'}(\tau)$.

Figure 2: $(\tau = abc, d)$ is a quasi-cospherical configuration because d is outside $\mathcal{C}_N(\tau)$ but inside $\mathcal{C}_{N'}(\tau)$

The following lemma relates quasi-cospherical configurations and inconsistencies.

Lemma 3.5 *Let τ be an inconsistent simplex that appears in star S_v but not in star S_w , where v and w are two vertices of τ . If $\gamma(\tau) < \gamma_0$, then there exists a site $q \in V$ such that the configuration (τ, q) is (γ_0, M_v) -cospherical.*

Proof Take $N = M_v$ and $N' = M_w$. Because the distortion of τ is less than γ_0 , we have $\gamma(v, w) = \gamma(M_v, M_w) \leq \gamma_0$. Since τ is a d -simplex in S_v but not in S_w , it belongs to $\text{Del}_N(V)$ and not to $\text{Del}_{N'}(V)$. Hence, there is a site $q \in V$ such that q is inside $\mathcal{B}_{N'}(\tau)$ and not inside $\mathcal{B}_N(\tau)$. It then follows from Lemma 3.4 that (τ, q) is a (γ_0, M_v) -cospherical configuration. $\square$

Given a metric M and a (γ_0, M) -cospherical configuration U , the M -radius $r_M(U)$ of U is defined as the minimum of the M -circumradii of all the d -simplices with vertices in U .

Definition 3.6 (Well-shaped configuration) *Let U be a (γ_0, M) -cospherical configuration witnessed by two metrics N and N' . U is said to be well-shaped if the simplices of the two triangulations $\text{Del}_N(U)$ and $\text{Del}_{N'}(U)$ are well-shaped for their respective metrics.*

4 Meshing Algorithm

4.1 Algorithm Outline

To mesh a given compact domain Ω , the algorithm constructs the set of sites V by inserting new sites in a greedy way. More precisely, the algorithm maintains the star set and, while there remain bad simplices in the star set, the algorithm selects one bad simplex and kills this simplex by inserting

a new site. Bad simplices are d -simplices that have a high distortion, those that are badly shaped (high radius-edge ratio or small sliverity ratio), and those that are inconsistent. To kill a bad simplex τ appearing in a star S_v , a new site p , called the *refinement point*, is inserted in the M_v -circumscribing ball of τ . Upon each insertion, the algorithm maintains the star set by calling the following **Insert** procedure.

Algorithm 1 $\text{Insert}(p)$

1. insert p in all the stars S_w that contain a simplex in conflict with p ;
 2. create the new star S_p .
-

Note that the new site p will be inserted not only in star S_v but also in all the other stars S_w that contain a d -simplex in conflict with p . In addition, a new star S_p will be created for p .

As noticed in Section 3.3, once the set of vertices is dense enough with respect to the variation of the metric field so that all simplices in the star set have a distortion smaller than γ_0 , inconsistencies arise only from quasi-cospherical configurations. The refinement algorithm therefore aims at avoiding those configurations. However, as will be clear from the proof of Lemma 4.3, it is not possible to avoid quasi-cospherical configurations involving slivers. The algorithm thus needs to remove slivers before removing inconsistent simplices.

The refinement algorithm consists in applying in turn the following rules. The rules are applied with a priority order : rule (i) is applied only if no rule (j) with $j < i$ can be applied. The algorithm ends when no rule applies any more. The algorithm relies on two procedures: procedure **Insert** inserts a new site in the data structures, and procedure **Pick_valid** chooses the location of the new site (see the next section).

The algorithm depends on constants α_0 , ρ_0 , σ_0 , γ_0 and on parameters β and δ . The values of constants α_0 , ρ_0 , σ_0 , γ_0 control the quality of the mesh elements and their adaptation to the metric field. Parameters β and δ influence the behaviour of the algorithm and their values are chosen in Section 5 in order to ensure the termination of the refinement algorithm.

If τ is a d -simplex in some star S_v , we write $B_v(\tau)$ or $B_v(c_v(\tau), r_v(\tau))$ for the M_v -circumball of τ with center $c_v(\tau)$ and radius $r_v(\tau)$, $\rho_v(\tau)$ for the M_v -radius-edge ratio of τ and $\sigma_v(\tau)$ for its M_v -sliverity ratio.

Remark. Parameter α_0 is always chosen less than 1. Therefore, when rule (1) does not apply anymore, the distortion of any d -simplex in any star is bounded by γ_0 .

Sections 5 and 6 will prove that the algorithm terminates. Before that, Subsection 4.2 describes the procedure **Pick_valid** while Subsection 4.3 analyses the properties of the resulting mesh.

4.2 Picking Region and Hitting Sets

In this section, we describe in more detail procedure **Pick_valid**. The simplest idea to kill a simplex would be to insert a refinement point at its circumcenter. However, with this simple strategy, the algorithm may loop, creating cascading configurations of slivers and quasi-cospherical configurations and is not guaranteed to terminate. To avoid slivers and quasi-cospherical configurations, the algorithm resorts to a strategy analog to the one used by Li and Teng [24, 23] to avoid slivers in isotropic meshes. The basic idea is to relax the position of the refinement point of a bad simplex. Instead of using the circumcenter of a bad simplex, the refinement point is picked from a small region

Algorithm 2 Refinement algorithm

Rule (1) Size:

If $\exists$ a d -simplex τ in star S_v such that $r_v(\tau) \geq \alpha_0 \text{sf}(c_v(\tau))$,
 Insert($c_v(\tau)$);

Rule (2) Radius-edge ratio:

If $\exists$ a d -simplex τ in star S_v such that $\rho_v(\tau) > \rho_0$,
 Insert(Pick_valid(τ, M_v));

Rule (3) Sliver removal:

If a d -simplex τ in star S_v is a M_v sliver (i.e. $\rho_v(\tau) \leq \rho_0$ and $\sigma_v(\tau) < \sigma_0$),
 Insert(Pick_valid(τ, M_v));

Rule (4) Inconsistency:

If a d -simplex τ in some star S_v is inconsistent,
 Insert(Pick_valid(τ, M_v));

around the circumcenter, called the *picking region*. The refinement point is carefully chosen in the picking region so as to avoid the formation of new slivers and new quasi-cospherical configurations.

Definition 4.1 (Picking region) Let $\delta < 1$ be a constant called the picking ratio. If τ is a bad simplex in star S_v , with M_v -circumball $\mathcal{B}_v(c_v(\tau), r_v(\tau))$, the M_v -picking region of τ , noted $P_v(\tau)$, is the M_v -ball $\mathcal{B}_v(c_v(\tau), \delta r_v(\tau))$.

In fact, it is not possible, when choosing a refinement point in the picking region $P_v(\tau)$ of a simplex τ of S_v to completely avoid the formation of new slivers and new quasi-cospherical configurations. The Pick_valid procedure will only avoid the creation of *small* slivers and *small* quasi-cospherical configurations where the meaning of *small*, precisely defined below, is relative to the radius $r_v(\tau)$ and controlled by a parameter β .

Definition 4.2 (Hitting set) Let p be a point in the M_v -picking region of a simplex τ . Let $r_v(\tau)$ be the M_v -circumradius of τ and β be a constant. A subset σ of the current set of sites V is said to hit p if one of the two following conditions is satisfied:

- σ consists of k sites and the k -simplex $\tau' = (\sigma, p)$ is, for some metric M such that $\gamma(M_p, M) \leq \gamma_0$, a k -sliver with M -circumradius $r_M(\tau') \leq \beta r_v(\tau)$.
- σ consists of $d + 1$ sites and $U = (\sigma, p)$ is for some metric M such that $\gamma(M_p, M) \leq \gamma_0$, a well-shaped (γ_0, M) -cospherical configuration with M -radius $r_M(U) \leq \beta r_v(\tau)$.

As usual, $r_v(\tau)$ denotes the M_v -circumradius of τ .

A point p in $P_v(\tau)$ is said to be a *valid refinement point* if it is not hit by any subset of V . A subset σ of sites in V that hits some point in the picking region $P_v(\tau)$ is said to be a *hitting set* for $P_v(\tau)$. Each hitting set σ of a picking region induces a *forbidden region* where the refinement point should not lie.

Note that the definition of valid refinement points depends on the constants δ and β : δ defines the size of the picking regions and β bounds from below the size of acceptable new slivers and new

quasi-cospherical configurations, with respect to the circumradius of the simplex being refined. The definition of valid refinement points also depends on the constants ρ_0 and σ_0 that define well-shaped simplices and slivers, and on the constant γ_0 that defines quasi-cospherical configurations. We prove in Section 6 the following *picking lemma*, which is fundamental for our algorithm and stated here for future reference.

Figure 3: $\{q, r, s\}$ is a hitting set for the picking region $P_v(\tau)$. It defines a forbidden region (grey area) to be avoided by the refinement point p of simplex τ .

Lemma 4.3 (Picking lemma) *For any values of parameters β and δ , and of constants α_0 and ρ_0 , it is possible to choose σ_0 small enough and the distortion bound γ_0 close enough to 1, so that valid refinement points exist in the M_v -picking region of any bad simplex τ in star S_v .*

To find a valid refinement point in the M_v -picking region $P_v(\tau)$ of some bad d -simplex τ , the insertion algorithm calls the following `Pick_valid` procedure. This procedure randomly chooses a point in the picking region $P_v(\tau)$ until it finds one that avoids forbidden regions. This procedure depends on constants ρ_0 , σ_0 , γ_0 , δ and β , to be fixed later in Section 5.

Algorithm 3 `Pick_valid`(τ, M_v)

Step 1 Pick randomly a point p in the picking region $P_v(\tau)$

Step 2 *Avoid small slivers*

For $k = 3$ to d ,

if, there exists a subset of k sites in V that hits p ,
 then discard p and go back to step 1.

Step 3 *Avoid small quasi-cospherical configurations*

If, there exists a subset of $d + 1$ sites in V that hits p ,
 then discard p and go back to step 1.

Step 4 Return p .

4.3 Quality of the final mesh

Upon termination of the algorithm, all stars are consistent and they can be merged together to form a triangulation $\mathcal{T}$ of the domain, with the property that all the simplices of $\mathcal{T}$ that are incident on a vertex v are Delaunay for the metric M_v and are well-shaped with respect to M_v .

Moreover, each simplex τ in $\mathcal{T}$ is well-shaped with respect to the metric of all its vertices, i.e. $\rho_p(\tau) \leq \rho_0$ and $\sigma_p(\tau) \geq \sigma_0$ for any vertex p of τ , and τ has a small distortion $\gamma(\tau) \leq \gamma_0$.

It now remains to prove that the algorithm terminates, which will be done in the two next sections.

5 Termination of the Algorithm

In this section, we prove that, if the picking lemma is true, the algorithm presented in the previous section does terminate, for suitable choices of constants $\alpha_0, \gamma_0, \rho_0, \sigma_0$ that respectively control the size, the distortion, the radius-edge ratio and the sliverity ratio of a simplex, and of parameters δ and β that bound the picking and the size ratios. The proof of the picking lemma is deferred to Section 6.

We prove that the algorithm terminates by bounding the number of mesh vertices using a volume argument.

Definition 5.1 We define the separation $s(a, b)$ between two sites a and b of V as :

$$s(a, b) = \min(d_a(a, b), d_b(a, b))$$

Let $V(p)$ be the set of inserted vertices that have been inserted before p . We define the *insertion radius* $r(p)$ of p as

$$r(p) = \min_{q \in V(p)} s(p, q).$$

To prove that the number of vertices is bounded, we will bound from below the insertion radius as a function of the sizing field $\text{sf}(p)$ in p .

We will consider in turn each of the refinement rules. We begin with a technical lemma relating the circumradius of a simplex with the insertion radius of its refinement point. As before, $r_v(\tau)$ and $c_v(\tau)$ denote respectively the M_v -circumradius and the M_v -circumcenter of τ .

Lemma 5.2 (Insertion radius lemma) Let τ be a d -simplex of star S_v to be refined by one of the algorithm rules and let p be the refinement point of τ .

- If rule (1) is applied, the refinement point p of τ is the circumcenter $c_v(\tau)$, and its insertion radius $r(p)$ is at least

$$r(p) \geq \frac{r_v(\tau)}{\Gamma}. \quad (8)$$

where Γ is the maximal distortion over Ω : $\Gamma = \max_{x, y \in \Omega} \gamma(x, y)$.

- If one of the rule (2), (3) or (4) is applied, the refinement point p is taken from the picking region $P_v(\tau)$ and the insertion radius $r(p)$ is at least:

$$r(p) \geq \frac{1 - \delta}{\Gamma} r_v(\tau). \quad (9)$$

Proof In the first case, $p = c_v(\tau)$, and therefore

$$\begin{aligned} \min_{q \in V(p)} d_v(p, q) &\geq r_v(\tau) \\ r(p) = \min_{q \in V(p)} s(p, q) &\geq \frac{r_v(\tau)}{\Gamma}. \end{aligned}$$

In the second case, p belongs to the picking region $P_v(\tau)$, and therefore

$$\begin{aligned} \min_{q \in V(p)} d_v(p, q) &\geq (1 - \delta)r_v(\tau) \\ r(p) = \min_{q \in V(p)} s(p, q) &\geq \frac{1 - \delta}{\Gamma} r_v(\tau). \end{aligned}$$

□

Lemma 5.3 *When rule (1) is applied, the insertion radius $r(p)$ of the inserted site p is at least:*

$$r(p) \geq C_1 \text{sf}(p) \text{ with } C_1 = \frac{\alpha_0}{\Gamma}. \quad (10)$$

Proof rule (1) is applied to a simplex τ in star S_v when the M_v -circumradius $r_v(\tau)$ of τ is greater than $\alpha_0 \text{sf}(c_v(\tau))$. The refinement point p is then $c_v(\tau)$ and we get from lemma 5.2

$$r(p) \geq \frac{r_v(\tau)}{\Gamma} \geq \frac{\alpha_0 \text{sf}(c_v(\tau))}{\Gamma} = \frac{\alpha_0}{\Gamma} \text{sf}(p).$$

□

Lemma 5.3 proves that any vertex p introduced by application of rule (1) has an insertion radius bounded from below by $C_1 \text{sf}(p)$ where $C_1 = \frac{\alpha_0}{\Gamma}$. The next lemmas aim at finding a constant C , and some conditions on α_0 , ρ_0 , γ_0 , β and δ so that rules (2)-(4) will maintain the invariant that the insertion radius of any inserted point is at least $C \text{sf}(p)$.

Lemma 5.4 *Let $\tau \in S(v)$ be a simplex to be refined by application of rule (2) and let p be the refinement point of τ . If, for any vertex q inserted before p , $r(q) \geq C \text{sf}(q)$ then we have $r(p) \geq C \text{sf}(p)$, provided that the following conditions hold*

$$\frac{(1 - \delta)\rho_0}{\Gamma\gamma_0^2} \geq 2, \quad (11)$$

$$\frac{(1 + \delta)\rho_0}{\gamma_0} C \leq 1. \quad (12)$$

Proof First, observe that $\gamma(\tau) \leq \gamma_0$ since rule (1) does not apply.

Now, because p is inserted by application of rule (2), the M_v -circumradius, $r_v(\tau)$, of τ is such that $r_v(\tau) \geq \rho_0 e_v(\tau)$, where $e_v(\tau)$ is the M_v -length of the shortest edge of τ according to metric M_v . Therefore, if q is the last inserted vertex of the shortest edge of τ , we have

$$\begin{aligned} r_v(\tau) &\geq \rho_0 e_v(\tau) \\ &\geq \frac{\rho_0}{\gamma_0} r(q) \quad (\text{since } v, q \in \tau) \end{aligned}$$

and, using the induction hypothesis and the smoothness assumption on the sizing field

$$\begin{aligned}
r_v(\tau) &\geq \frac{\rho_0}{\gamma_0} C \text{sf}(q) \\
&\geq \frac{\rho_0}{\gamma_0} C \frac{[\text{sf}(p) - d_p(p, q)]}{\gamma(p, q)} \\
&\geq \frac{\rho_0}{\gamma_0^2} C [\text{sf}(p) - d_p(p, q)] \\
&\geq \frac{\rho_0}{\gamma_0^2} C [\text{sf}(p) - \gamma_0 d_v(p, q)] \quad (\text{since } v, p \in B_v(\tau))
\end{aligned} \tag{13}$$

Now, because q is a vertex of τ and p is chosen in the picking region $P_v(\tau)$, $d_v(p, q) \leq (1 + \delta)r_v(\tau)$ which, together with inequality (13), gives:

$$\begin{aligned}
r_v(\tau) &\geq \frac{\rho_0}{\gamma_0^2} C [\text{sf}(p) - \gamma_0(1 + \delta)r_v(\tau)] \\
r_v(\tau) &\geq \frac{\frac{\rho_0}{\gamma_0^2} C \text{sf}(p)}{1 + \frac{\rho_0}{\gamma_0}(1 + \delta) C}.
\end{aligned} \tag{14}$$

Then, using the insertion radius Lemma 5.2, we get:

$$r(p) \geq \frac{\frac{(1-\delta)\rho_0}{\Gamma\gamma_0^2} C \text{sf}(p)}{1 + \frac{\rho_0}{\gamma_0}(1 + \delta) C},$$

which proves that $r(p) \geq C \text{sf}(p)$ when conditions (11) and (12) are fulfilled. $\square$

Lemma 5.5 *Let $\tau \in S(v)$ be a simplex to be refined by application of rule (3) or rule (4), and let p be the refinement point of τ . If, for any vertex q inserted before p , $r(q) \geq C \text{sf}(q)$ then we have $r(p) \geq C \text{sf}(p)$, provided that the following conditions hold*

$$\frac{\beta(1 - \delta)}{\Gamma\gamma_0^2(1 + \delta)} \geq 2 \tag{15}$$

$$C \leq \frac{\frac{(1-\delta)C_1}{2\Gamma\gamma_0^2}}{1 + \frac{1+\delta}{2\gamma_0} C_1} \tag{16}$$

$$\frac{\beta C}{\gamma_0} \leq 1. \tag{17}$$

Proof 1. Assume first that τ was created by application of rule (1). Then, if q is the last inserted vertex of τ , we have $r(q) \geq C_1 \text{sf}(q)$ by Lemma 5.3. Furthermore, $r_v(\tau)$ is plainly at least half the M_v -length of any edge of τ and, in particular, of any edge of τ that is incident to q . Therefore, using the fact that v and q belong to τ , we get

$$r_v(\tau) \geq \min_{q' \in \tau} \frac{d_v(q, q')}{2} \geq \frac{r(q)}{2\gamma_0} \geq \frac{C_1}{2\gamma_0} \text{sf}(q)$$

Now, using calculations similar to what has been done to deduce inequality (14) from (13), we obtain

$$\begin{aligned} r_v(\tau) &\geq \frac{C_1}{2\gamma_0^2} [\text{sf}(p) - \gamma_0(1+\delta)r_v(\tau)] \\ &\geq \frac{\frac{C_1}{2\gamma_0^2} \text{sf}(p)}{1 + \frac{C_1(1+\delta)}{2\gamma_0}}. \end{aligned}$$

Then, using the insertion radius Lemma 5.2, we get:

$$r(p) \geq \frac{\frac{(1-\delta)C_1}{2\gamma_0^2} \text{sf}(p)}{1 + \frac{1+\delta}{2\gamma_0} C_1}.$$

It follows that the bound $r(p) \geq C \text{sf}(p)$ holds, provided condition (16) is satisfied.

2. Now consider the case where τ was created by application of rule (2), (3) or (4), which means that τ is either a sliver or belong to a quasi-cospherical configuration. Assume that τ has been created when inserting the refinement point q of a simplex τ' in some star S_w (see Figure 4). The refinement point q was chosen by the procedure `Pick_valid`(τ' , M_w) and therefore, $r_v(\tau) \geq \beta r_w(\tau')$. Let us bound from below $r_w(\tau')$. Vertex q is the last inserted vertex of τ . It has been chosen in the picking region of τ' and therefore the vertices of τ' are at most at M_w -distance $(1+\delta)r_w(\tau')$ from q . Hence, since q and w belong to τ' , $\gamma(\tau') \leq \gamma_0$ and $r(q) \leq \gamma_0(1+\delta)r_w(\tau')$.

Figure 4: For the proof of Lemma 5.5.

Therefore:

$$\begin{aligned} r_v(\tau) &\geq \beta r_w(\tau') \geq \frac{\beta}{\gamma_0(1+\delta)} r(q). \\ r_v(\tau) &\geq \frac{\beta C}{\gamma_0(1+\delta)} \text{sf}(q). \end{aligned}$$

Using calculations similar to what has been done to deduce inequality (14) from (13), we obtain

$$r_v(\tau) \geq \frac{\beta C}{\gamma_0^2(1+\delta)} [\text{sf}(p) - \gamma_0(1+\delta)r_v(\tau)]$$

Hence,

$$r_v(\tau) \geq \frac{\frac{\beta C}{\gamma_0^2(1+\delta)} \text{sf}(p)}{1 + \frac{\beta C}{\gamma_0}},$$

and, from the insertion radius Lemma 5.2,

$$r(p) \geq \frac{\frac{\beta(1-\delta)C}{\Gamma\gamma_0^2(1+\delta)}}{1 + \frac{\beta C}{\gamma_0}} \text{sf}(p).$$

Conditions (17) and (15) ensure that $r(p) \geq C \text{sf}(p)$. $\square$

We can now give the main theorem of this section that states that the algorithm terminates and therefore, as mentionned above, that the resulting star set is indeed an anisotropic Delaunay mesh.

Theorem 5.6 *Given a compact domain Ω and a sizing field defined over Ω that satisfies the positiveness (5) and smoothness conditions (7), the refinement algorithm terminates provided that the constants ρ_0 , γ_0 , β and δ are chosen in such a way that:*

$$\frac{(1-\delta)\rho_0}{\Gamma\gamma_0^2} \geq 2 \tag{18}$$

$$\frac{(1-\delta)\beta}{\Gamma\gamma_0^2(1+\delta)} \geq 2 \tag{19}$$

Proof Observe that the inequalities in the theorem are just conditions (11) and (15) of Lemma 5.4 and 5.5. Assume that these inequalities hold. Let C be a constant small enough so that $C \leq C_1 = \frac{\alpha_0}{\Gamma}$ and condition (12) of Lemma 5.4, and conditions (16) and (17) of Lemma 5.5 hold. Then, the results of Lemmas 5.3, 5.4, and 5.5, and an easy induction shows that any vertex p in the mesh has an insertion radius $r(p) \geq C \text{sf}(p)$.

Let V be the set of vertices at some step of the algorithm and let p be a vertex of V . The minimum distance between p and any other vertex in V is bounded from below as follows:

$$\min_{q \in V \setminus \{p\}} s(p, q) \geq \frac{C}{(1+C)\Gamma} \text{sf}(p). \tag{20}$$

Indeed, let (p, q) be any pair of vertices of V . Either p were inserted after q , in which case we have $s(p, q) \geq r(p) \geq C \text{sf}(p)$, or q were inserted after p , in which case we have

$$s(p, q) \geq r(q) \geq C \text{sf}(q) \geq \frac{C}{\Gamma} [\text{sf}(p) - d_p(p, q)] \geq \frac{C}{\Gamma} [\text{sf}(p) - \Gamma s(p, q)]$$

$$\text{Hence } s(p, q) \geq \frac{C}{(1+C)\Gamma} \text{sf}(p).$$

Then, because $\text{sf}(x) \geq \text{sf}_0 > 0$ for any x in Ω , the minimum separation between any two vertices of the constructed mesh is at least $\frac{C}{(1+C)\Gamma} \text{sf}_0$. Then, for the metric M_y of any point y in Ω , the minimum M_y -distance between any two vertices is at least $\frac{C}{(1+C)\Gamma^2} \text{sf}_0$. Since Ω is a compact domain and has therefore a bounded M_y -volume, this proves that the algorithm can only insert a finite number of vertices and therefore terminates. $\square$

6 Proof of the picking lemma 4.3

To complete the proof of termination of the algorithm, it remains to prove the picking Lemma 4.3, which is the goal of this section.

Let us recall briefly the context. Assume that the algorithm needs to refine a simplex τ in star S_v , with circumball $B_v(c_v(\tau), r_v(\tau))$. The picking lemma states that it always finds a valid refinement point provided that the bound on the sliverity ratio σ_0 is small enough and that the bound on the distortion γ_0 is sufficiently close to 1. The refinement point is searched in the picking region $P_v(\tau)$, a M_v -ball with radius $\delta r_v(\tau)$ centered at the circumcenter $c_v(\tau)$. The refinement point is valid when it does not belong to the so-called forbidden regions. Each forbidden region is associated to a hitting set and consists of the points in the picking region that form with the hitting set either a *small* sliver or a *small* well-shaped quasi-cospherical configuration. *Small* is here relative to the circumradius $r_v(\tau)$ and controlled by parameter β , (see Definition 4.2).

The proof consists in showing that the union of the forbidden regions does not cover the picking region. In a first step, we show that the volume of each forbidden region is bounded and in fact can be made as small as required with a good choice of the parameters σ_0 and γ_0 (Lemmas 6.2 and 6.4). In a second step, we show that the number of hitting sets, or equivalently of forbidden regions to be avoided, is bounded (Lemma 6.5).

We begin with a technical lemma that bounds the difference between the two circumspheres of a well-shaped simplex (see Definition 3.1) with respect to two metrics with a bounded distortion.

Lemma 6.1 (Circumsphere lemma) *Let M_v and M_w be two metrics with a distortion $\gamma(M_v, M_w) \leq \gamma_0$ for some $\gamma_0 > 1$. Let τ be a k -simplex that is well shaped with respect to metric M_v . We write c_v and r_v for the M_v -circumcenter and M_v -circumradius of τ . and c_w, r_w for its M_w -circumcenter and M_w -circumradius.*

- The M_v -distance $d_v(c_v, c_w)$ between the M_v and M_w -circumcenters of τ satisfies

$$d_v(c_v, c_w) \leq f_k(\rho_0, \sigma_0, \gamma_0) r_v \quad (21)$$

where

$$f_k(\rho_0, \sigma_0, \gamma_0) = \left[1 + \frac{2^k}{k} \frac{\gamma_0^2 \rho_0^k}{\sigma_0^k} \right] (\gamma_0^2 - 1).$$

- The circumradius r_w is bounded as follows

$$r_w \in [h_k^-(\rho_0, \sigma_0, \gamma_0) r_v, h_k^+(\rho_0, \sigma_0, \gamma_0) r_v]$$

where

$$\begin{aligned} h_k^-(\rho_0, \sigma_0, \gamma_0) &= \frac{1}{\gamma_0} (1 - f_k(\rho_0, \sigma_0, \gamma_0)), \\ h_k^+(\rho_0, \sigma_0, \gamma_0) &= \gamma_0 (1 + f_k(\rho_0, \sigma_0, \gamma_0)). \end{aligned}$$

Remark. Observe that $f_k(\sigma_0, \gamma_0)$ tends to zero when σ_0 tends to 0, γ_0 tends to 1 and $(\gamma_0 - 1)/\sigma_0^k$ tends to 0.

Proof The proof is given in the appendix. $\square$

Avoiding slivers

Let τ be a k -simplex of a star $S(v)$. We again write r_v for its M_v -circumradius. Consider a refinement point p to be taken from the M_v -picking region $P_v(\tau)$ of τ . Point p is required to lie outside any forbidden region. We first consider the case of a forbidden region $Y_v(\tau')$ associated to a hitting set τ' of k' sites that form with p a small sliver for metric M close to M_p . More precisely (see Definition 4.2), by a metric M close to M_p , we mean a metric M such that $\gamma(M, M_p) \leq \gamma_0$, and by a small sliver we mean a sliver whose M -circumradius is smaller than βr_v . For convenience, τ' will denote both the subset of sites and the simplex formed by its convex hull. Note that τ' is not required to be a simplex appearing in some current star.

Lemma 6.2 (Forbidden regions due to slivers) *The M_v -volume of the forbidden region $Y_v(\tau')$ is bounded from above as follows*

$$\text{Vol}_v(Y_v(\tau')) \leq \mu_{k'}(\rho_0, \sigma_0, \gamma_0) \beta^d r_v^d,$$

where $\mu_{k'}(\rho_0, \sigma_0, \gamma_0)$ is a function that tends to 0 when σ_0 tends to 0 and γ_0 tends to 1 in such a way that $\frac{\gamma_0 - 1}{\sigma_0^{k'-1}}$ tends to 0.

Proof By definition of a hitting set, there is a metric M satisfying $\gamma(M_p, M) \leq \gamma_0$ such that τ' is a well-shaped $(k' - 1)$ -simplex with respect to M with a M -circumradius r' smaller than βr_v . Now, since p belongs to the M_v -picking region $P_v(\tau)$ of τ , we have $\gamma(M_v, M_p) \leq \gamma_0$. It follows that $\gamma(M_v, M) \leq \gamma_0^2$.

From the sliver Lemma 3.2, we know that $\text{conv}(p, \tau)$ is a k' -sliver for metric M if p is at at M -distance at most $4\pi k' \rho_0 \sigma_0 r'$ from $\mathcal{C}(c', r') \cap \text{aff}(\tau')$ where $\mathcal{C}(c', r')$ denotes the M -circumscribing sphere of τ' . Let $\mathcal{C}_v(c'_v, r'_v)$ denote the M_v -circumsphere of τ' . We further write $\eta_{k'}(\rho_0) = 4\pi k' \rho_0$. Applying the circumsphere Lemma 6.1 to the well-shaped $(k' - 1)$ -simplex τ' , we get

$$\begin{aligned} d_v(c'_v, p) &\leq d_v(c'_v, c') + d_v(c', p) \\ &\leq \gamma_0^2 d_M(c'_v, c') + \gamma_0^2 d_M(c', p) \\ &\leq \gamma_0^2 f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2) r' + \gamma_0^2 [1 + \eta_{k'}(\rho_0) \sigma_0] r' \\ &\leq \gamma_0^2 [1 + f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2) + \eta_{k'}(\rho_0) \sigma_0] r' \stackrel{\text{def}}{=} \lambda^+ r', \end{aligned}$$

writing $\lambda^+ = \gamma_0^2 [1 + f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2) + \eta_{k'}(\rho_0) \sigma_0]$. In the same way, we have:

$$\begin{aligned} d_v(c'_v, p) &\geq d_v(c', p) - d_v(c'_v, c') \\ &\geq \frac{1}{\gamma_0^2} d_M(c', p) - \gamma_0^2 d_M(c'_v, c') \\ &\geq \frac{1}{\gamma_0^2} [1 - \eta_{k'}(\rho_0) \sigma_0] r' - \gamma_0^2 f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2) r' \stackrel{\text{def}}{=} \lambda^- r', \end{aligned}$$

writing $\lambda^- = \frac{1}{\gamma_0^2} [1 - \eta_{k'}(\rho_0) \sigma_0] - \gamma_0^2 f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2)$.

By definition of a hitting set, the quasi-cospherical configuration (τ', p) is required to be small so that $r' \leq \beta r_v$. It follows that the forbidden region $Y_v(\tau')$ associated to τ' is included in the M_v -spherical shell $\mathcal{S}_v(c'_v, r_v^+, r_v^-)$ enclosed by two M_v -spheres centered at c'_v and with respective radii $r_v^+ = \lambda^+ \beta r_v$ and $r_v^- = \lambda^- \beta r_v$.

The M_v -volume of the spherical shell $\mathcal{S}_v(c'_v, r_v^+, r_v^-)$ and thus the M_v -volume of $Y_v(\tau')$ is at most

$$\phi_d \left(\frac{r_v^+ + r_v^-}{2} \right)^{d-1} (r_v^+ - r_v^-)$$

where ϕ_d is a constant that depends only on the dimension d , and

$$\begin{aligned} r_v^+ + r_v^- &= \frac{\lambda^+ + \lambda^-}{2} \beta r_v \\ &= \left[\left(\gamma_0^2 + \frac{1}{\gamma_0^2} \right) + \left(\gamma_0^2 - \frac{1}{\gamma_0^2} \right) \eta_{k'}(\rho_0) \sigma_0 \right] \beta r_v \end{aligned} \quad (22)$$

$$\begin{aligned} r_v^+ - r_v^- &= \frac{\lambda^+ - \lambda^-}{2} \beta r_v \\ &= \left[\left(\gamma_0^2 - \frac{1}{\gamma_0^2} \right) + 2\gamma_0^2 f_{k'-1}(\rho_0, \sigma_0, \gamma_0^2) + \left(\gamma_0^2 + \frac{1}{\gamma_0^2} \right) \eta_{k'}(\rho_0) \sigma_0 \right] \beta r_v \end{aligned} \quad (23)$$

Hence, we can write $\text{Vol}_v(Y_v(\tau')) \leq \mu_{k'}(\sigma_0, \gamma_0) \beta^d r_v^d$. Moreover, it follows from the remark after the circumsphere Lemma 6.1 that $r_v^+ - r_v^-$, and therefore $\mu_{k'}(\sigma_0, \gamma_0)$, tends to 0 when σ_0 tends to 0 and γ_0 tends to 1 in such a way that $\frac{\gamma_0 - 1}{\sigma_0^{k'-1}}$ tends to 0. $\square$

Avoiding cospherical configurations

In Lemma 6.2, we bounded the volume of a forbidden region associated to a sliver. We will now bound the volume of a forbidden region associated to a cospherical configuration. We first prove a technical lemma.

Lemma 6.3 (Cospherical lemma) *Given are*

1. *a metric M and a distortion bound $\gamma_0 > 1$,*
2. *a d -simplex τ that is well shaped with respect to M . We denote by c and r the M -circumcenter and the M -circumradius of τ .*
3. *a point p such that the configuration (p, τ) is a (γ_0, M) -cospherical configuration.*

Then p belongs to the M -spherical shell $\mathcal{S}_M(c, g_d^-(\rho_0, \sigma_0, \gamma_0) r, g_d^+(\rho_0, \sigma_0, \gamma_0) r)$ enclosed between two M -spheres centered at c , with respective radii $g_d^-(\rho_0, \sigma_0, \gamma_0) r$ and $g_d^+(\rho_0, \sigma_0, \gamma_0) r$, where:

$$\begin{aligned} g_d^+(\rho_0, \sigma_0, \gamma_0) &= \left[\gamma_0^2 + (1 + \gamma_0^2) f_d(\rho_0, \sigma_0, \gamma_0) \right] \\ g_d^-(\rho_0, \sigma_0, \gamma_0) &= \left[\frac{1}{\gamma_0^2} - (1 + \frac{1}{\gamma_0^2}) f_d(\rho_0, \sigma_0, \gamma_0) \right]. \end{aligned}$$

Proof Let N and N' be two metrics that witness the (γ_0, M) -cospherical configuration (τ, p) , such that point p belongs to the interior of the N -circumball $B_N(\tau)$ while p does not belong to the interior of the N' -circumball $B_{N'}(\tau)$. Let $c_N, c_{N'}$ denote respectively the N and N' -circumcenters of τ . Then, using Lemma 6.1,

$$\begin{aligned}
 d_M(p, c) &\leq d_M(p, c_N) + d_M(c_N, c) \\
 &\leq \gamma_0 d_N(p, c_N) + f_d(\rho_0, \sigma_0, \gamma_0)r \\
 &\leq \gamma_0 h_d^+(\rho_0, \sigma_0, \gamma_0)r + f_d(\rho_0, \sigma_0, \gamma_0)r \\
 &\leq [\gamma_0^2 + (1 + \gamma_0^2) f_d(\rho_0, \sigma_0, \gamma_0)] r \\
 &= g_d^+(\rho_0, \sigma_0, \gamma_0)
 \end{aligned} \tag{24}$$

and

$$\begin{aligned}
 d_M(p, c) &\geq d_M(p, c_N) - d_M(c_N, c) \\
 &\geq \frac{1}{\gamma_0} d_N(p, c_N) - f_d(\rho_0, \sigma_0, \gamma_0)r \\
 &\geq \frac{1}{\gamma_0} h_d^-(\rho_0, \sigma_0, \gamma_0)r - f_d(\rho_0, \sigma_0, \gamma_0)r \\
 &\geq \left[\frac{1}{\gamma_0^2} - \left(1 + \frac{1}{\gamma_0^2}\right) f_d(\rho_0, \sigma_0, \gamma_0) \right] r \\
 &= g_d^-(\rho_0, \sigma_0, \gamma_0)
 \end{aligned} \tag{25}$$

Inequalities (24) and (25) are just another way to state Lemma 6.3. $\square$

Let τ be a k -simplex of a star $S(v)$, write r_v for its M_v -circumradius, and consider a refinement point p to be taken from the M_v -picking region $P_v(\tau)$ of τ . Point p is required to lie outside any forbidden region. After considering the case of a forbidden region associated to a sliver in the previous section, we consider now the case of a forbidden region $W_v(\tau')$ associated to a hitting set τ' of $d + 1$ sites that form with p a small M -cospherical configuration for a metric M close to M_p . Again (see Definition 4.2), by a metric M close to M_p , we mean such that $\gamma(M, M_p) \leq \gamma_0$, and by small configuration, we mean a configuration whose M -circumradius is smaller than βr_v . For convenience, τ' will denote both the subset of sites and the simplex formed by its convex hull. Note that τ' is not required to be a simplex appearing in some current star.

Lemma 6.4 (Forbidden regions due to cospherical configurations) *The M_v -volume of the forbidden region $W_v(\tau')$ is bounded from above as follows*

$$\text{Vol}_v(W_v(\tau')) \leq \omega(\rho_0, \sigma_0, \gamma_0) \beta^d r_v^d,$$

where $\omega(\rho_0, \sigma_0, \gamma_0)$ is a function that tends to 0 when σ_0 tends to 0 and γ_0 tends to 1 in such a way that $\frac{\gamma_0 - 1}{\sigma_0^d}$ tends to 0.

Proof As for the proof of Lemma 6.2, we prove that the forbidden region $W_v(\tau')$ is included in a M_v -spherical shell $\mathcal{S}_v(c'_v, r_v^+, r_v^-)$ enclosed between two M_v -spheres centered at c'_v , the M_v -circumcenter of τ' . For the same reason as in the proof of Lemma 6.2, there exists a metric M satisfying $\gamma(M, M_v) \leq \gamma_0^2$ such that τ' forms with p a (γ_0, M) -cospherical configuration. Let c', r'

be respectively the M -circumcenter and the M -circumradius of τ' . Applying Lemmas 6.1 and 6.3 to τ' , we get:

$$\begin{aligned} d_v(p, c'_v) &\leq d_v(p, c') + d_v(c', c'_v) \\ &\leq \gamma_0^2 d_M(p, c') + \gamma_0^2 d_M(c', c'_v) \\ &\leq \gamma_0^2 g_d^+(\rho_0, \sigma_0, \gamma_0) r' + \gamma_0^2 f_d(\rho_0, \sigma_0, \gamma_0^2) r' \stackrel{\text{def}}{=} \lambda^+ r', \end{aligned} \quad (26)$$

where $\lambda^+ = \gamma_0^2 g_d^+(\rho_0, \sigma_0, \gamma_0) + \gamma_0^2 f_d(\rho_0, \sigma_0, \gamma_0^2)$. Similarly,

$$\begin{aligned} d_v(p, c'_v) &\geq d_v(p, c') - d_v(c', c'_v) \\ &\geq \frac{1}{\gamma_0^2} d_M(p, c') - \gamma_0^2 d_M(c', c'_v) \\ &\geq \frac{1}{\gamma_0^2} g_d^-(\rho_0, \sigma_0, \gamma_0) r' - \gamma_0^2 f_d(\rho_0, \sigma_0, \gamma_0^2) r' \stackrel{\text{def}}{=} \lambda^- r', \end{aligned} \quad (27)$$

where $\lambda^- = \frac{1}{\gamma_0^2} g_d^-(\rho_0, \sigma_0, \gamma_0) - \gamma_0^2 f_d(\rho_0, \sigma_0, \gamma_0^2)$.

By definition of a hitting set, the quasi-spherical configuration (τ', p) is required to be small. Specifically, r' has to be at most βr_v . It follows that the forbidden region $W_v(\tau')$ associated to τ' is included in the M_v -spherical shell $\mathcal{S}_v(c'_v, r_v^+, r_v^-)$ enclosed by the two M_v -spheres centered at c'_v of radii $r_v^+ = \lambda^+ \beta r_v$ and $r_v^- = \lambda^- \beta r_v$.

The M_v -volume of the spherical shell $\mathcal{S}_v(c'_v, r_v^+, r_v^-)$ and thus the M_v -volume of $W_v(\tau')$ is at most

$$\phi_d \left(\frac{r_v^+ + r_v^-}{2} \right)^{d-1} (r_v^+ - r_v^-)$$

where ϕ_d is a constant that depends only on the dimension d , and

$$\begin{aligned} r_v^+ + r_v^- &= \left[\left(\gamma_0^4 + \frac{1}{\gamma_0^4} \right) + \left(\gamma_0^2 (1 + \gamma_0^2) - \frac{1}{\gamma_0^2} \left(1 + \frac{1}{\gamma_0^2} \right) \right) f_d(\rho_0, \sigma_0, \gamma_0) \right] \beta r_v \\ r_v^+ - r_v^- &= \left[\left(\gamma_0^4 - \frac{1}{\gamma_0^4} \right) + \left(\gamma_0^2 (1 + \gamma_0^2) + \frac{1}{\gamma_0^2} \left(1 + \frac{1}{\gamma_0^2} \right) \right) f_d(\rho_0, \sigma_0, \gamma_0) + 2\gamma_0^2 f_d(\rho_0, \sigma_0, \gamma_0^2) \right] \beta r_v \end{aligned} \quad (28)$$

Hence, we can write $\text{Vol}_v(W_v(\tau')) \leq \omega(\rho_0, \sigma_0, \gamma_0) \beta^d r_v^d$. Moreover, it follows from the remark after the circumsphere Lemma 6.1 that $\omega(\rho_0, \sigma_0, \gamma_0)$ tends to 0 when σ_0 tends to 0 and γ_0 tends to 1 in such a way that $\frac{\gamma_0 - 1}{\sigma_0^{k'-1}}$ tends to 0. $\square$

Bounding the number of forbidden regions

Lemma 6.5 Assume that a refinement point is searched in the M_v -picking region $P_v(\tau)$ of a d -simplex in the star S_v , and write $K_v(\tau)$ for the set of hitting subsets of $P_v(\tau)$. If the algorithm parameters β, δ , the size bound α_0 and the distorsion bound γ_0 satisfy the relation

$$\alpha_0 \gamma_0 (\delta + 2\beta \gamma_0^2) < 1 \quad (29)$$

then the cardinality of $K_v(\tau)$ is bounded by a constant K that depends on α_0, β, δ and γ_0 and remains bounded when γ_0 tends to 1.

Proof First observe that the cardinality of each hitting subset τ' in $K_v(\tau)$ is at most $d + 1$. To bound the cardinality of $K_v(\tau)$, we bound the cardinality of the set $Q_v(\tau)$ of vertices that may be part of a hitting set τ' .

Let q be a vertex of $Q_v(\tau)$. The slivers or quasi-cospherical configurations to avoid are required to be small, i.e. to have a M -circumradius smaller than βr_v , for a metric M such that $\gamma(M, M_v) \leq \gamma_0^2$. Therefore, the M -distance from q to p is at most $2\beta r_v$. Moreover, if c_v denotes as usual the M_v -circumcenter of the simplex τ to be refined,

$$\begin{aligned} d_v(c_v, q) &\leq d_v(c_v, p) + d_v(p, q) \\ &\leq \delta r_v + \gamma_0^2 d_M(p, q) \quad (\text{using Hypothesis 2.2}) \\ &\leq (\delta + 2\gamma_0^2 \beta) r_v \quad (\text{since } v \text{ and } c_v \text{ both belong to } B_v(\tau)) \end{aligned}$$

We have $r_v \leq \alpha_0 \text{sf}(c_v)$ since, when a point is searched in the picking region of a simplex, rule (1) does not apply anymore. Hence, the inequality above becomes

$$d_v(c_v, q) \leq l_1 \text{sf}(c_v), \quad \text{where } l_1 = \alpha_0 (\delta + 2\gamma_0^2 \beta). \quad (30)$$

We now use inequality (20) that bounds from below the distance from a site to the other sites of V . We have for any vertex $q' \in Q_v(\tau)$

$$\begin{aligned} s(q, q') &\geq \frac{C}{(1+C)\Gamma} \text{sf}(q), \\ &\geq \frac{C}{(1+C)\Gamma\gamma_0} [\text{sf}(c_v) - d_{c_v}(c_v, q)] \\ &\geq \frac{C}{(1+C)\Gamma\gamma_0} [\text{sf}(c_v) - \gamma_0 d_v(c_v, q)] \end{aligned}$$

The same bound plainly holds for $d_v(q, q')$. Now, using inequality (30), we get

$$d_v(q, q') \geq \frac{C}{1+C} \frac{1 - \gamma_0 l_1}{\Gamma\gamma_0} \text{sf}(c_v)$$

Let us write

$$l_2 = \frac{C}{1+C} \frac{1 - \gamma_0 l_1}{\Gamma\gamma_0} = \frac{C}{1+C} \frac{1 - \alpha_0 \gamma_0 (\delta + 2\gamma_0^2 \beta)}{\Gamma\gamma_0}$$

which is positive when condition (29) is satisfied. We then have $d_v(q, q') \geq l_2 \text{sf}(c_v)$. This last inequality and inequality (30) show that the M_v -balls centered at the vertices of $Q_v(\tau)$ and with radii $l_2 \text{sf}(c_v)/2$ are disjoint and all contained in the M_v -ball $B(c_v, (l_1 + l_2/2) \text{sf}(c_v))$.

A volume argument then shows that the cardinality of $Q_v(\tau)$ is bounded by $(1 + 2l_1/l_2)^d$. By considering all possible simplices with vertices in $Q_v(\tau)$, we get a bound on the number $|K_v(\tau)|$ of forbidden regions we need to avoid when picking a refinement point in $P_v(\tau)$

$$|K_v(\tau)| \leq |Q_v(\tau)|^{d+1} \leq 2^{d+1} (1 + 2l_1/l_2)^{d(d+1)}$$

The lemma is proved by taking

$$\begin{aligned} K &= 2^{d+1} (1 + 2l_1/l_2)^{d(d+1)} \\ &= 2^{d+1} \left[1 + 2 \left(\frac{1+C}{C} \right) \Gamma_{\gamma_0} \left(\frac{\alpha_0(\delta + 2\gamma_0^2\beta)}{1 - \alpha_0\gamma_0(\delta + 2\gamma_0^2\beta)} \right) \right]^{d(d+1)} \end{aligned}$$

□

Proof of the picking lemma

When a refinement point p has to be picked in the picking region $P_v(\tau)$ of some d -simplex τ in star S_v , the M_v -volume of the picking region $P_v(\tau)$ is $\delta^d r_v^d(\tau) u_d$ where u_d is the volume of the unit Euclidean ball of dimension d .

To be valid, the refinement point has to lie outside the forbidden regions. In the previous lemmas, we have bounded the M_v -volume of the forbidden regions. More precisely, in Lemma 6.2, we gave a bound on the volume of the forbidden region associated to small k -sliver and, in Lemma 6.4, we gave a bound on the volume of the forbidden region associated to a small quasi-cospherical configuration. The total number of possible forbidden regions has then be bounded in Lemma 6.5 by a constant K .

A valid refinement point exists in $P_v(\tau)$ if the volume of the picking region exceeds the total volume of the forbidden regions which is guaranteed if both following conditions holds:

$$K \mu_{k'}(\rho_0, \sigma_0, \gamma_0) \beta^d \leq \delta^d u_d, \quad k' = 1, \dots, d \quad (31)$$

$$K \omega(\rho_0, \sigma_0, \gamma_0) \beta^d \leq \delta^d u_d \quad (32)$$

Remark. If α_0, β, δ , have been fixed, and if γ_0 remains bounded from above, K remains also bounded. Moreover, since $\mu_{k'}$ and ω can be made arbitrarily small when σ_0 tends to 0 and γ_0 tends to 1 in such a way that $(\gamma_0 - 1)/\sigma_0^d$ tends to 0, inequalities (31) and (32) are satisfied if one choose σ_0 small enough and γ_0 close enough to 1.

To conclude, we discuss how to choose the parameters of the algorithm. We first choose $\delta < 1$. Then ρ_0 and β are chosen such that conditions (18) and (19) of Theorem 5.6 proving that the algorithm terminates, are satisfied for some reasonable value of γ_0 , say $\gamma_0 = 1.5$. These equations will remain satisfied for any smaller value of γ_0 . Then, we choose parameter α_0 small enough so as to satisfy condition (29) of Lemma 6.5 proving that the number of forbidden regions to take into account at each call of the `Pick_valid()` procedure, is bounded for the value $\gamma_0 = 1.5$. This equation will also be satisfied for any smaller value of γ_0 . At last, we choose σ_0 and γ_0 to satisfy equations (31) and (32) ensuring the picking lemma.

7 Conclusion

We have proposed a new class of anisotropic meshes, the so-called anisotropic Delaunay meshes. These meshes conform to a given metric field, can be defined in any dimension, and keep locally the nice properties of Delaunay meshes. We also described an algorithm to generate such meshes in any dimension d . Differently from other methods that have been proposed in dimensions higher than 2, our algorithm produces meshes with a precise characterization and theoretical guarantees.

The algorithm is simple and has been implemented for $d = 2$ and 3 using the CGAL library [1]. Figure 7 shows the output of the algorithm on a 3-dimensional ball where the metric is stretched horizontally in the left part and vertically in the right part. The metric field varies slowly on the figure on the left and rapidly on the figure on the right. In this example, we did not enforce any size bound, so that the refinement is only governed by the need to remove inconsistencies. As expected, the mesh density depends on the distortion of the metric. The line where the eigenvectors exchange their eigenvalues is clearly visible on the figure on the right. Further experimental results will be reported elsewhere.

By placing anisotropic meshes in the realm of Delaunay meshes, our framework allows to benefit from recent advances in isotropic mesh generation. In particular, our work can be extended in the following directions.

- *Domains with complex boundaries.* Our analysis extends to the case of polyhedra of $\mathbb{R}^3$ with no sharp edges as shown in [8]. A simple variant of our algorithm can be used to generate anisotropic Delaunay meshes of surfaces [6]. Combining the present work and [6], it is possible to mesh 3-dimensional domains bounded by complex smooth boundaries in a way similar to what has been done for isotropic meshes [29]. The case of piecewise smooth boundaries is another extension that can be done easily by adapting the work of Cheng, Dey and Levine [14, 13, 16]. The error analysis of Mirebeau for the approximation of cartoon functions is also relevant in this context [28].
- *Mesh optimization.* As demonstrated in the isotropic case, local optimization can greatly improve the quality of meshes [34]. In particular, ODT (optimal Delaunay triangulations) [11, 4] nicely improve Delaunay meshes generated by refinement. Since our mesh is locally Delaunay, similar techniques can be applied in our anisotropic framework.
- *Parallelization.* Our algorithm computes the stars independently and then look for inconsistencies among neighboring stars. It can therefore be parallelized rather naturally.

A more difficult extension that we let for future work would be to allow our approach to adapt to time varying metric fields. This is important in the context of numerical PDEs where the mesh is not fixed in advance but is dynamically updated based on the knowledge on the exact solution one gains during the solution process.

Acknowledgments

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

References

- [1] CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
- [2] F. Alauzet. Size gradation control of anisotropic meshes. *Finite Elements in Analysis and Design*, 46(1-2):181–202, 2010.
- [3] F. Alauzet, A. Loseille, A. Dervieux, and P. Frey. Multi-dimensional continuous metric for mesh adaptation. In *Proc. 15th Int. Meshing Roundtable*, pages 191–214, 2006.

Figure 5: Two examples of anisotropic meshes produced by our algorithm.

- [4] P. Alliez, D. Cohen-Steiner, M. Desbrun, P. Schröder, and M. Yvinec. Variational tetrahedral meshing. In *Proceedings SIGGRAPH*, 2005. To appear.
- [5] I. Babuska and W.C. Rheinboldt. Error estimates for adaptive finite element computations. *SIAM Journal on Numerical Analysis*, pages 736–754, 1978.
- [6] J.D. Boissonnat, K.L. Shi, and M. Yvinec. Anisotropic Delaunay meshes of surfaces. *submitted*, 2011.
- [7] J.D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic diagrams: Labelle and Shewchuk approach revisited. *Theoretical Computer Science*, 408(2-3):163–173, 2008.
- [8] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Locally uniform anisotropic meshing. In *Proceedings of the 24th Annual Symposium on Computational Geometry*, pages 270–277, 2008.
- [9] Houman Borouchaki, Paul Louis George, Frédéric Hecht, Patrick Laug, and Eric Saltel. Delaunay mesh generation governed by metric specifications. part I algorithms. *Finite Elem. Anal. Des.*, 25(1-2):61–83, 1997.
- [10] Frank Bossen and Paul Heckbert. A pliant method for anisotropic mesh generation. In *5th International Meshing Roundtable*, October 1996.
- [11] L. Chen, P. Sun, and J. Xu. Optimal anisotropic meshes for minimizing interpolation errors in l^p -norm. *Mathematics of Computation*, 76(257):179, 2007.
- [12] Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Rephael Wenger. Anisotropic surface meshing. In *In SODA'06 : Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 202–211, New York, NY, USA, 2006. ACM.

- [13] S.W. Cheng, T.K. Dey, and J. Levine. Theory of a practical Delaunay meshing algorithm for a large class of domains. *Algorithms, Architecture and Information Systems Security*, B. Bhattacharya, S. Sur-Kolay, S. Nandy, and A. Bagchi, Eds, 3:17–41, 2008.
- [14] S.W. Cheng, T.K. Dey, and J.A. Levine. A practical Delaunay meshing algorithm for a large class of domains*. In *Proceedings of the 16th International Meshing Roundtable*, pages 477–494. Springer, 2008.
- [15] E.F. D’azevedo. Optimal triangular mesh generation by coordinate transformation. *SIAM Journal on Scientific and Statistical Computing*, 12:755, 1991.
- [16] T.K. Dey and J.A. Levine. Delaunay meshing of piecewise smooth complexes without expensive predicates. *Algorithms*, 2(4):1327–1349, 2009.
- [17] C. Dobrzynski and P. Frey. Anisotropic delaunay mesh adaptation for unsteady simulations. *Proceedings of the 17th International Meshing Roundtable*, pages 177–194, 2008.
- [18] Q. Du and D. Wang. Anisotropic centroidal voronoi tessellations and their applications. *SIAM Journal on Scientific Computing*, 26(3):737–761, 2005.
- [19] Herbert Edelsbrunner, Xiang-Yang Li, Gary Miller, Andreas Stathopoulos, Dafna Talmor, Shang-Hua Teng, Alper Üngör, and Noel Walkington. Smoothing and cleaning up slivers. In *STOC ’00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 273–277, New York, NY, USA, 2000. ACM.
- [20] P.J. Frey and F. Alauzet. Anisotropic mesh adaptation for cfd computations. *Computer methods in applied mechanics and engineering*, 194(48-49):5068–5082, 2005.
- [21] Francois Labelle and Jonathan Richard Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SCG’ 03 : Proceedings of the nineteenth annual symposium on Computational geometry*, pages 191–200, New York, NY, USA, 2003. ACM Press.
- [22] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Symposium on Computational Geometry*, pages 341–349, 2000.
- [23] Xiang-Yang Li. Generating well-shaped d-dimensional Delaunay meshes. *Theor. Comput. Sci.*, 296(1):145–165, 2003.
- [24] Xiang-Yang Li and Shang-Hua 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. Society for Industrial and Applied Mathematics, 2001.
- [25] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. Biting ellipses to generate anisotropic mesh. In *8th International Meshing Roundtable*, October 1999.
- [26] A. Loseille and F. Alauzet. Optimal 3d highly anisotropic mesh adaptation based on the continuous mesh framework. *Proceedings of the 18th International Meshing Roundtable*, pages 575–594, 2009.
- [27] A. Loseille, A. Dervieux, and F. Alauzet. Fully anisotropic goal-oriented mesh adaptation for 3d steady Euler equations. *Journal of Computational Physics*, 229(8):2866–2897, 2010.

-
- [28] J-M. Mirebeau. *Anisotropic finite element approximation. Theory and algorithms*. PhD thesis, Université Pierre et Marie Curie, Paris 6, 2010.
 - [29] S. Oudot, L. Rineau, and M. Yvinec. Meshing volumes bounded by smooth surfaces. In *Proceedings of the 14th International Meshing Roundtable*, pages 203–219. Springer, 2005.
 - [30] Jessica Schoen. Robust, guaranteed-quality anisotropic mesh generation. M.S. thesis, University of California at Berkeley, 2008.
 - [31] Jonathan Richard 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.
 - [32] Richard 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.
 - [33] 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.
 - [34] J. Tournois, C. Wormser, P. Alliez, and M. Desbrun. Interleaving Delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Transactions on Graphics (TOG)*, 28(3):1–9, 2009.

Appendix

7.1 Proof of the Sliver Lemma (3.2)

Proof (Sliver lemma) In this proof, all lengths, volumes and angles are measured with respect to the metric M . We denote by r and $r(v)$ the circumradii of τ and $\tau(v)$ respectively, by V and $V(v)$ their respective volumes and by e and $e(v)$ the lengths of their respective shortest edge. Let d be the distance from v to the affine hull $\text{aff}(\tau(v))$ of $\tau(v)$ and let d' be the distance from v to the sphere $\text{aff}(\tau(v)) \cap \mathcal{C}(v)$.

Using the fact that τ is a sliver, we have

$$V = \frac{1}{k} d V(v) \leq \sigma_0^k e^k,$$

which yields

$$d < \frac{k\sigma_0^k e^k}{V(v)}.$$

As $\tau(v)$ is a face of τ , we have $e \leq e(v)$, and, since $\tau(v)$ is not a sliver, $V(v) \geq \sigma_0^{k-1} e(v)^{k-1}$. Then,

$$\begin{aligned} d &< \frac{k\sigma_0^k e^k}{\sigma_0^{k-1} e(v)^{k-1}}, \\ &\leq k\sigma_0 e(v) \\ &\leq 2k\sigma_0 r(v), \end{aligned}$$

which proves the first part of the lemma.

To bound the distance d' , we consider the 2-plane through v and the centers c and c' of the circumspheres $\mathcal{C}$ and $\mathcal{C}(v)$ of τ and $\tau(v)$ respectively. See Figure 6. Let p be the projection of v on the affine envelope $\text{aff}(\tau(v))$ and let p' be the projection of v on the sphere $\text{aff}(\tau(v)) \cap \mathcal{C}(v)$. Thus $d = \|vp\|$ and $d' = \|vp'\|$. Let q be the point where the ray issued from c that passes through c' intersects $\mathcal{C}$. Let $\varphi = \widehat{pp'v}$ and $\theta = \widehat{qcp'}$. Observe that $\frac{d}{d'} = \sin \varphi$ and $\sin \theta = \frac{r(v)}{r} \geq \frac{r(v)/e(v)}{r/e} \geq \frac{1}{2\rho_0}$, because $r(v) \geq e(v)/2$ and the radius-edge ratio $\frac{r}{e}$ of τ is smaller than ρ_0 .

We distinguish two cases according to the position of c and v with respect to the affine hull $\text{aff}(\tau(v))$ of $\tau(v)$.

In the first case (Figure 6, left part), c and v are on different sides of $\text{aff}(\tau(v))$. We have $\varphi \geq \frac{\theta}{2}$ and therefore

$$d' = \frac{d}{\sin \varphi} \leq \frac{d}{\sin(\frac{\theta}{2})} \leq \frac{2k\sigma_0 r(v)}{\sin(\frac{1}{2} \arcsin \frac{1}{2\rho_0})} \leq \frac{\pi k\sigma_0 r(v)}{\frac{1}{2} \arcsin \frac{1}{2\rho_0}} \leq 4\pi k\rho_0\sigma_0 r(v)$$

where we have made use of the first part of the lemma and of the fact that $\frac{2}{\pi}u \leq \sin u \leq u$ for any $u \in [0, \frac{\pi}{2}]$ and $u \leq \arcsin u$ for $u \in [0, 1]$

In the second case (Figure 6, right part), c and v are on the same side of $\text{aff}(\tau(v))$. Then, $\varphi \geq \theta$ and

$$d' = \frac{d}{\sin \varphi} \leq \frac{d}{\sin \theta} \leq 4k\rho_0\sigma_0 r(v),$$

which ends the proof. $\square$

Figure 6: For the proof of the sliver lemma.

7.2 Proof of the Circumsphere Lemma (6.1)

Proof We first prove the circumsphere lemma when τ is a d -simplex. The case of a k -simplex will easily follow by applying the same argument in the affine hull of τ .

7.2.1 Computing the circumcenters

Let $\tau = (p_0, \dots, p_d)$ be a d -simplex. Since the M_v -circumcenter, c_v , of τ is at equal M_v -distance from all the vertices of τ , we have $d_v^2(c_v, p_i) = r_v^2$ for $i = 0, \dots, d$. Therefore,

$$(p_i - c_v)^t F_v^t F_v (p_i - c_v) = (p_0 - c_v)^t F_v^t F_v (p_0 - c_v) \quad i = 1, \dots, d.$$

Equivalently, we have for $i = 0, \dots, d$

$$\begin{aligned} ((p_i - p_0) + (p_0 - c_v))^t F_v^t F_v ((p_i - p_0) + (p_0 - c_v)) &= (p_0 - c_v)^t F_v^t F_v (p_0 - c_v) \\ \Leftrightarrow (p_i - p_0)^t F_v^t F_v (p_i - p_0) &= 2(p_i - p_0)^t F_v^t F_v (c_v - p_0) \end{aligned}$$

Writing $P = (p_1 - p_0, \dots, p_d - p_0)$ for the square matrix whose columns are the vectors $p_i - p_0$, $i = 1, \dots, d$, and $\text{Diag}(A)$ for the column matrix whose elements are the elements of the main diagonal of a square matrix A , the last equation becomes

$$\text{Diag}(P^t F_v^t F_v P) = 2 P^t F_v^t F_v (c_v - p_0),$$

from which we get the position of c_v with respect to the position of the vertices of τ

$$c_v - p_0 = \frac{1}{2} (F_v^t F_v)^{-1} P^{-t} \text{Diag}(P^t F_v^t F_v P). \quad (33)$$

An equivalent formula gives the M_w -circumcenter c_w of τ .

7.2.2 Bounding the distance between c_v and c_w

In the following, we choose a coordinate system in which $M_v = F_v^t F_v$ and F_v are identity matrices. (Equivalently, we could assume without loss of generality that M_v is the Euclidean metric since the distance $d_v(c_v, c_w)$ is only related to the relative distortion between M_v and M_w .) Then, we deduce from (33) :

$$c_v - p_0 = \frac{1}{2} P^{-t} q \quad (34)$$

$$c_v - c_w = \frac{1}{2} [P^{-t} q - (F_w^t F_w)^{-1} P^{-t} q'] \quad (35)$$

where

$$\begin{aligned} q &= \text{Diag}(P^t P), \\ q' &= \text{Diag}(P^t F_w^t F_w P). \end{aligned}$$

We further write

$$c_v - c_w = \frac{1}{2} [I - (F_w^t F_w)^{-1}] P^{-t} q + \frac{1}{2} (F_w^t F_w)^{-1} P^{-t} (q - q'),$$

where I is the identity matrix. By our choice of the coordinate system, the M_v -norm of a vector x is just the Euclidean norm $\|x\|$ of its coordinates in this reference system. Therefore,

$$d_v(c_v, c_w) = \|c_v - c_w\| \leq \frac{1}{2} \|I - (F_w^t F_w)^{-1}\| P^{-t} q + \frac{1}{2} \|(F_w^t F_w)^{-1}\| P^{-t} (q - q'). \quad (36)$$

The following claim provides bounds for the two terms on the right hand side of (36).

Claim 7.1

$$\|(I - (F_w^t F_w)^{-1}) P^{-t} q\| \leq 2(\gamma_0^2 - 1) r_v. \quad (37)$$

$$\|(F_w^t F_w)^{-1} P^{-t} (q - q')\| \leq \gamma_0^2 (\gamma_0^2 - 1) \frac{2^{d+1}}{d} \frac{\rho_0^d}{\sigma_0^d} r_v. \quad (38)$$

Proof Writing $\|A\|$ for the Euclidean norm of a matrix A , ($\|A\| = \sup_{\|x\|=1} \|Ax\|$), we have,

$$\|(I - (F_w^t F_w)^{-1}) P^{-t} q\| \leq \|(I - (F_w^t F_w)^{-1})\| \|P^{-t} q\|.$$

$F_w^t F_w$ is a symmetric square matrix with eigenvalues in the interval $[\frac{1}{\gamma_0^2}, \gamma_0^2]$. The absolute values of the eigenvalues of matrix $I - (F_w^t F_w)^{-1}$ are thus at most $\gamma_0^2 - 1$. Moreover, from (34), $\|P^{-t} q\| = 2d_v(c, p_0)$ is just twice the M_v -circumradius of τ , which proves inequality (37).

To prove (38), we write

$$\|(F_w^t F_w)^{-1} P^{-t} (q - q')\| \leq \|(F_w^t F_w)^{-1}\| \|P^{-t}\| \|q - q'\|. \quad (39)$$

We will bound the three terms on the right hand side of (39). We first note that

$$\|(F_w^t F_w)^{-1}\| \leq \gamma_0^2. \quad (40)$$

Then, for $\|P^{-t}\|$, we use the fact that $\|P^{-t}\| \leq \|P^{-t}\|_\infty$ where $\|P^{-t}\|_\infty$ is the maximum absolute value of any entry in P^{-t} . Each entry in P^{-t} is a cofactor of matrix P^t divided by the determinant of P^t . The determinant of P^t is $d!$ times the M_v -volume of τ . Each entry in P^t is a coordinate of some $p_i - p_0$ and therefore less than $\|p_i - p_0\| \leq 2r_v$, which implies that each cofactor of P^t is at most $(d-1)!(2r_v)^{d-1}$. Therefore,

$$\begin{aligned} \|P^{-t}\| &\leq \|P^{-t}\|_\infty \\ &\leq \frac{(d-1)!(2r_v)^{d-1}}{d! \text{Vol}_v(\tau)} \\ &\leq \frac{2^{d-1}}{d} \frac{\rho_0^{d-1}}{\sigma_0^d e_v}, \end{aligned} \quad (41)$$

where e_v is the M_v -length of the shortest (for M_v) edge of τ . We now bound $\|q - q'\|$:

$$\begin{aligned} \|q - q'\| &= \|\text{Diag}(P^t P) - \text{Diag}((P^t F_w^t F_w P))\| \\ &\leq \|\text{Diag}(P^t P) - \text{Diag}((P^t F_w^t F_w P))\|_\infty \\ &\leq \max_i \left| d_v(p_i, p_0)^2 - d_w(p_i, p_0)^2 \right| \\ &\leq 4(\gamma_0^2 - 1) r_v^2 \\ &\leq 4(\gamma_0^2 - 1) \rho_0 e_v r_v \end{aligned} \quad (42)$$

Inequalities (39), (40), (41) and (42) yield (38) which achieves to prove claim 36. $\square$

We finally get from (36), (37) and (38) :

$$\begin{aligned} d_v(c_v, c_w) &\leq (\gamma_0^2 - 1)r_v + \frac{1}{2}\gamma_0^2(\gamma_0^2 - 1) \frac{2^{d+1}}{d} \frac{\rho_0^d}{\sigma_0^d} r_v \\ &\leq \left[1 + \frac{2^d}{d} \frac{\gamma_0^2 \rho_0^d}{\sigma_0^d} \right] (\gamma_0^2 - 1) r_v \end{aligned}$$

This ends the proof of the first part of Lemma 6.1 in the case of a d -simplex.

7.2.3 Bounding the circumradius r_w

Let p be a vertex of τ . We have $r_v = d_v(c_v, p)$ and $r_w = d_w(c_w, p)$. Since metric M_v satisfies the triangular inequality,

$$d_v(c_w, p) - d_v(c_v, c_w) \leq d_v(c_v, p) \leq d_v(c_w, p) + d_v(c_v, c_w).$$

Then, using the fact $\gamma(M_v, M_w) \leq \gamma_0$ and the first part of Lemma 6.1,

$$\begin{aligned} \frac{d_w(c_w, p)}{\gamma_0} - f_k(\gamma_0)r_v &\leq r_v \leq \gamma_0 d_w(c_w, p) + f_k(\gamma_0)r_v \\ \frac{r_w}{\gamma_0} - f_k(\gamma_0)r_v &\leq r_v \leq \gamma_0 r_w + f_k(\gamma_0)r_v. \end{aligned}$$

Therefore,

$$\frac{r_v}{\gamma_0} (1 - f_k(\gamma_0)) \leq r_w \leq r_v \gamma_0 (1 + f_k(\gamma_0)), \quad (43)$$

which proves the second part of Lemma 6.1 in the case of a d -simplex.

7.2.4 The case of a k -simplex

In the case of a k -simplex τ , the circumcenters c_v and c_w belong to the k -dimensional subspace that is the affine hull, $\text{aff}(\tau)$, of τ . If $\mathcal{C}(v)$ and $\mathcal{C}(w)$ are respectively the M_v and M_w circumspheres of τ , the above proof applies verbatim to the spheres $\text{aff}(\tau) \cap \mathcal{C}(v)$ and $\text{aff}(\tau) \cap \mathcal{C}(w)$ that are the circumspheres of τ in the subspace $\text{aff}(\tau)$. This yields the proof of Lemma 6.1 in the case of a k -simplex. $\square$

Contents

1	Introduction	3
2	Preliminaries	5
2.1	Anisotropic Metric	5
2.2	Metric Field and Distortion	6
2.3	Sizing field	6
3	Stars and Refinement	7
3.1	Stars and Inconsistencies	8
3.2	Slivers	8
3.3	Quasi-Cosphericity	9
4	Meshing Algorithm	10
4.1	Algorithm Outline	10
4.2	Picking Region and Hitting Sets	11
4.3	Quality of the final mesh	14
5	Termination of the Algorithm	14
6	Proof of the picking lemma 4.3	19
7	Conclusion	25
7.1	Proof of the Sliver Lemma (3.2)	30
7.2	Proof of the Circumsphere Lemma (6.1)	31
7.2.1	Computing the circumcenters	31
7.2.2	Bounding the distance between c_v and c_w	32
7.2.3	Bounding the circumradius r_w	33
7.2.4	The case of a k -simplex	34

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

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399