Generating 3D Anisotropic Centroidal Voronoi Tessellations

Alexandre Marin, Alexandra Bac, Laurent Astart

World Academy of Science, Engineering and Technology
International Journal of Computer and Information Engineering
Vol:18, No:9, 2024

Generating 3D Anisotropic Centroidal Voronoi Tessellations

Alexandre Marin, Alexandra Bac, Laurent Astart

Abstract—New numerical methods for PDE resolution (such as Finite Volumes (FV) or Virtual Element Method (VEM)) open new needs in terms of meshing of domains of interest, and in particular polyhedral meshes have many advantages. One way to build such meshes consists in constructing Restricted Voronoi Diagrams (RVDs) whose boundaries respect the domain of interest. By minimizing a function defined for RVDs, the shapes of cells can be controlled, i.e. elongated according to user-defined directions or adjusted to comply with given aspect ratios (anisotropy) and density variations. In this paper, our contribution is threefold: first, we present a gradient formula for the Voronoi tessellation energy under a continuous anisotropy field. Second, we describe a meshing algorithm based on the optimisation of this function that we validate against state-of-the-art approaches. Finally, we propose a hierarchical approach to speed up our meshing algorithm.

Keywords—Anisotropic Voronoi Diagrams, Meshes for Numerical Simulations, Optimisation, Volumic Polyhedral Meshing.

I. Introduction

WITH the rapid expansion of computer graphics uses, new problems and needs are coming out. Among them, the generalization of numerical simulation entails new challenges.

Because of the drawbacks of tetrahedral and hex-dominant meshes and since finite element methods are limited to precise cell topologies (e.g. hexahedra and their degenerations), recent works on numerical simulation have developed schemes applicable to more general meshes. The Hybrid Mixed Mimetic (HMM) family or nonlinear finite volume schemes, reviewed in [1], are examples of such generalizations. Initially developed for standard geoscience meshes, those methods can actually be applied to a larger class of meshes, namely polyhedral meshes.

Polyhedral meshes are good candidates to simultaneously meet both of these criteria: reducing the number of cells and controlling their geometry to improve simulation stability and performances.

Generating a Voronoi diagram is the most famous way to tessellate a domain. In order to control the aspect ratio of cells, state-of-the-art works introduce an energy (called f_{CVT}), whose minimization provides control over the density and the anisotropy of the tessellation. However, such approaches remain computationally costly and complex (e.g. for technical reasons, this optimization requires a discretization of the volume of interest and of the anisotropy field).

Alexandre Marin is with IFPEN / Aix-Marseille University, France (corresponding author, e-mail: alexandre.marin1@laposte.net).
Alexandra Bac is with Aix-Marseille University, France (e-mail: alexandra.bac@univ-amu.fr).
Laurent Astart is with IFPEN, France (e-mail: laurent.astart@ifpen.fr).

Our contribution in this paper is threefold (as illustrated in Fig. 2). First we present an expression of ∇f_{CVT} for a continuous anisotropy field. Second, unlike state-of-the-art methods, we obtain a method (called "exact" in the sequel) which does not require any discretization of the anisotropy (in accordance with the volume of interest) and hence does not depend on the quality of the input mesh, see Fig. 1. Third, we present a heuristic hierarchical approach to accelerate and parallelize the computation of restricted anisotropic polyhedral tessellations. In order to validate this approach, we give qualitative and quantitative results and carry out a comparison with a pre-existent approach.

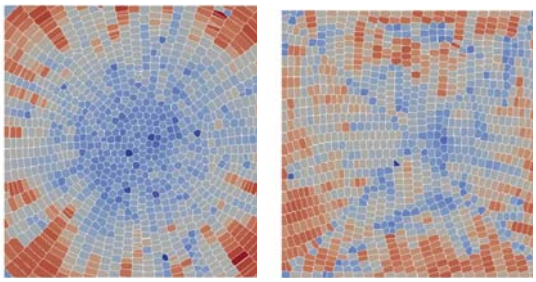

(a) (b)

Fig. 1 Results provided by our "exact" method (a) and by the method from [2] (b), for a Volume Of Interest (VOI) having 12 tetrahedra; the second result is faulty because of the approximation of the anisotropy over a coarse VOI

The paper is organised as follows: in Section II we present a brief review of the state of the art, then in Section III we introduce the whole theoretical context on which our work is based, in Section IV we introduce our main results and present our computational approach. In Section V we describe our validation tests, then in Section VI we present analyses and discuss our results. Our multi-level approach is introduced in Section IV-D. In Section VII we conclude and comment on possible perspectives.

II. State of the Art

A. Centroidal Voronoi Tessellations

Amongst all Voronoi diagrams that can be generated over a given domain and for a given number of seeds, some contain cells with more interesting shapes and orientations: they are called Centroidal Voronoi Tessellations (CVTs).

Two versions of CVTs have been studied. First, [3] examined the simplest case of regular CVTs. The main limitation of the paper lies in the fact that axes of the chosen tensor, which define the quality of cells, do not depend on a variable in $\mathbb{R}^3$. Authors suggested that "boundary terms"

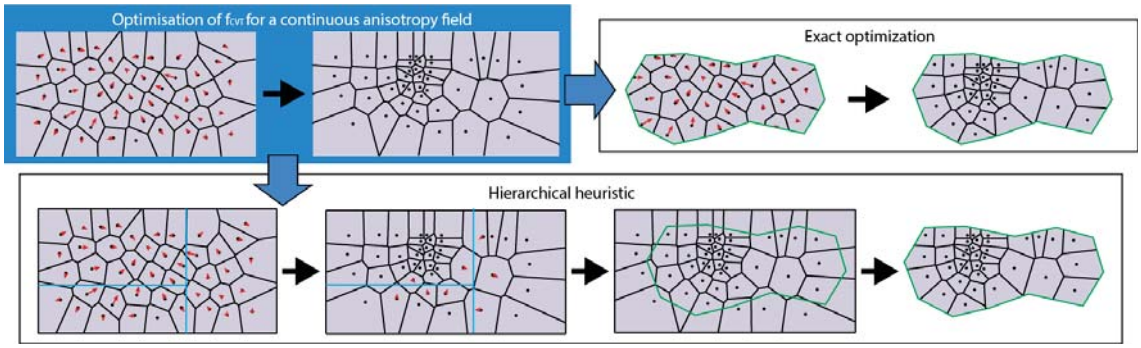

Fig. 2 Overview of our works (we represent the VOI in green): (a) our main tool an expression of ∇f_{CVT} under a continuous anisotropy field, (b) our "exact" optimisation method for restricted Voronoi diagrams (each step of the solver is applied to restricted Voronoi cells), (c) our four-step hierarchical heuristics

should appear in the derivation of an integral whose domain varies with the variables. However, no further proof was provided. Fortunately, in this symmetrical context, boundary terms vanish because they cancel each other out. Several other works considered CVTs with given densities. [4] provided and proved the gradient of such a f_{CVT} by invoking the REYNOLDS transport theorem. Nevertheless, it seems that this theorem is not completely adapted to f_{CVT} for Voronoi diagrams restricted to surfaces (RVDs): the surface transport theorem should have been used instead, but the results are correct as boundaries of Voronoi cells are planar. Authors of that paper also proved C^1-continuity of their generalized function f_{CVT} and showed how this result still holds when the integrand inside the expression of f_{CVT} is less regular. Another proof of the gradient was given in [5]. This gradient formula is quite similar to the one we introduce in the present paper, but we work in the more general setting of anisotropic 3D Voronoi tessellations.

Second, works introduced CVTs with controlled local anisotropy. Such tessellations have been largely studied (see [2], [6]–[10]) as they entail theoretical problems. State-of-the-art works provide various technical solutions to circumvent computational difficulties. [2] gave a formula to compute the gradient of CVTs, but only when the supplied anisotropy is piecewise constant on a tetrahedral mesh.

B. Other Approaches for Anisotropy

Let us point out that there are at least two other ways to get anisotropic "Voronoi" diagrams.

In [10], authors use an energy which depends on a convex function f to build Optimal Voronoi Tessellations (OVT). This energy is a discretization of the error between f and a continuous piecewise linear approximation over a Bregman diagram. However, setting user-defined anisotropies seems tricky because infering a smooth function f from a collection of Hessian matrices could be hard, and using such an approach would be delicate in practice. This last work is generalized by [11]: Bregman diagrams are replaced with Power diagrams and the input function f does not need to be convex, but optimal diagrams are no longer centroidal. Moreover, the minimized energy then depends on seed weights, and this double dependency makes the minimization tough.

Another way to build anisotropic Voronoi tessellations is the approach described in [9]: this work is inspired by the NASH theorem. Authors work in a space of dimension higher than 4: their method requires a discretization of a Riemannian metric and an optimisation step. However, due to the higher dimension of the working space, geometrical operations and predicates are more expensive.

C. Solvers

A few well-known solvers have been tested to get CVTs. Two of them proved to be efficient for CVT optimisation with density only. The Lloyd algorithm is a fixed-point algorithm which is regularly used in state-of-the-art works ([5], [10], [12]–[14]). Actually, in order to bring the anisotropic problem back to isotropic meshing (and hence to Lloyd relaxation), [15] embedded the problem into a space of higher dimension. As recalled by [14], the efficiency of the Lloyd method decreases with the size of the problem.

Quasi-Newton methods have also been used. Among them, L-BFGS (explained in [16]), which consumes little memory, seems to be very efficient for density variations as observed by [13].

III. THEORETICAL FRAME

Our work is based on Voronoi diagrams; in this section, we make some reminders and define Centroidal Voronoi Tessellations and their properties.

A. Reminders on Voronoi Diagrams

Let $\mathcal{V}$ be a non-empty finite set of points in the space $\mathbb{R}^3$ (called *seeds* in the sequel). Each point p is associated with the convex polyhedron V_p such that all points inside V_p are closer to p than to any other point of $\mathcal{V}$. Some of these regions are unbounded; moreover $\{V_p\,;\, p \in \mathcal{V}\}$ is a covering of $\mathbb{R}^3$. These polyhedra are called the *cells of the Voronoi diagram* $\mathscr{V}$ generated by seeds $\mathcal{V}$.

It is well known that the Voronoi diagram generated by the set of seeds $\mathcal{V}$ is the dual of the Delaunay triangulation of $\mathcal{V}$ (see [17] for an introduction). Actually, the most wide-spread way of computing the topology of a Voronoi diagram is first to build the Delaunay triangulation and then to calculate its dual.

However, most applications require bounded polyhedral meshes; the bounded domain studied Ω is called the *volume of interest* (VOI) with $\Omega \subset \mathbb{R}^3$ usually connected. The intersection of a Voronoi diagram $\mathscr{V}$ and Ω is called a *Restricted Voronoi Diagram* (abbreviated by RVD), and polyhedra $V_p \cap \Omega$ are called *Restricted Voronoi Cells* (RVCs). An efficient procedure to compute such restricted Voronoi diagrams can be found in [18].

B. Centroidal Voronoi Tessellations

Let $\mathscr{V}$ be the restricted Voronoi diagram generated by seeds $\mathcal{V} = \{x_i\,;\, i = 1 \dots N\}$ over a VOI Ω. The geometry of Voronoi cells only depends on the seeds. Controlling this geometry is clearly necessary in the context of polyhedral mesh generation. Various properties can be assessed on a given Voronoi tessellation, in particular: the shape ratios or orientation of cells and the distance between the centroid and the seed of cells. In order to quantify the quality of Voronoi diagrams, that is, to exhibit their ability to match user-defined features, "CVT functions" have been proposed:

$$f_{\text{CVT}} : \left(\mathbb{R}^3\right)^N \to \mathbb{R} :\ \mathbf{x} \mapsto \sum_{V_{x_i} \in \mathscr{V}} \mathcal{E}_i(\mathbf{x}) \tag{1}$$

where $\mathcal{E}_i$ is an "energy" function which associates the cell V_{x_i} with a positive real number. Such a function depends on the seeds $x_1,\ \dots,\ x_N$, i.e. on the variable

$$\mathbf{x} := (x_1,\ \dots,\ x_N) \in \mathbb{R}^{3N}. \tag{2}$$

Various properties of Voronoi diagrams can be assessed and controlled this way depending on the choice of the energy function:

- The energy:

$$\mathcal{E}_i^{\text{cent}} := \int_{V_{x_i}} \|y - x_i\|^2 \, \mathrm{d}y \tag{3}$$

  estimates the squared distance between the centroid and the seed of cells. Hence, minimizing the corresponding $f_{\text{CVT}}^{\text{cent}}$ function produces cells as compact as possible.
- Given $\rho : \Omega \to \mathbb{R}^+$ a real function encoding a density field over $\mathbb{R}^3$, let:

$$\mathcal{E}_i^{\text{dens}} := \int_{V_{x_i}} \rho(y) \cdot \|y - x_i\|^2 \, \mathrm{d}y. \tag{4}$$

  Minimizing the corresponding $f_{\text{CVT}}^{\text{dens}}$ produces compact cells whose density varies with ρ (i.e. the larger ρ, the denser the cells).

In both cases, a minimizer $\mathbf{x}$ of f_{CVT} produces Voronoi diagrams such that for all i, the seed x_i is the centroid of its cell V_{x_i}, hence the name *Centroidal Voronoi Tessellations* (or CVT for short). Moreover:

- Given $M : \Omega \to \mathcal{M}_3(\mathbb{R})$ a field of matrices encoding anisotropy, let:

$$\mathcal{E}_i^{\text{aniso}} := \int_{V_{x_i}} \|M[y - x_i]\|_p^p \, \mathrm{d}y \tag{5}$$

  where p is an even integer larger than 2. Minimizing $f_{\text{CVT}}^{\text{aniso}}$ tends to produce cells resembling ellipsoids (or parallelepipeds) and of shape ratio (i.e. ratio of extensions along axes) $\|M_j\|_2 / \|M_i\|_2$ along directions M_i and M_j, where M_k is the k-th row of M. This expression was introduced by LÉVY and MERLAND (e.g. [7] and [8]). Anisotropy is modeled through a Riemannian metric (encoded by M) over the domain. As a consequence: $M(x) = \Lambda(x) \cdot P(x)$, where $\Lambda(x)$ is a positive definite diagonal matrix and $P(x)$ is an orthogonal matrix.
  Note that the larger p, the closer from $\|\cdot\|_\infty$ the p-norm: hence orientations of facets of cells in a CVT appear more clearly as p increases. Last, with this expression, $|||M|||_p^p$ controls the density (in the sense of $f_{\text{CVT}}^{\text{dens}}$), where $|||M|||_p$ denotes the matrix norm subordinated to the p-norm.

C. Gradient and Regularity of f_{CVT}

Several state-of-the-art works provide estimates of the gradient of f_{CVT}. In [3], the authors compute the gradient of $f_{\text{CVT}}^{\text{cent}}$ by a direct derivation under the integral sign. However, authors point out that variations of cells along the minimization process must be considered and introduce a second integral term taking into account boundary variations. They also provide an estimate for the Hessian matrix of $f_{\text{CVT}}^{\text{cent}}$.

According to [13], $f_{\text{CVT}}^{\text{dens}}$ is C^2, provided that ρ is also C^2 and Ω is convex. If Ω is no longer convex, $f_{\text{CVT}}^{\text{dens}}$ remains C^2 almost everywhere. Moreover, the authors empirically checked that $f_{\text{CVT}}^{\text{dens}}$ is C^2 even if ρ is only continuous. [12], [13] both give the same expression of the gradient, which is proved by [19] for dimension 2. [5] relies on the REYNOLDS transport theorem to calculate the gradient of $f_{\text{CVT}}^{\text{dens}}$, not only for Voronoi diagrams but also for power diagrams.

Other works about Voronoi diagrams restricted to surfaces (see [4], [6], [20]) use analogous formulas.

Based on the first-order optimality criterion, one can easily check that in the absence of anisotropy, CVT functions reach their minimum when seeds are centroids of associated cells. Hence, the fixed-point algorithm, called the Lloyd algorithm, which consists in replacing each seed with the centroid of its associated cell, comes naturally: it is experimentally proved in [13] to converge.

To the best of our knowledge, there was no expression (nor rigorous proof) of the gradient of $f_{\text{CVT}}^{\text{aniso}}$ when the field M is continuous. In [4], a formula is given, but limited to Voronoi diagrams restricted to surfaces. Compared to our results, this is a narrower and therefore simpler context as it entails only surface integrals.

IV. OUR APPROACH

As illustrated in Fig. 2, based on our expression of $\nabla f_{\text{CVT}}^{\text{aniso}}$ for a continuous anisotropy field (represented in blue) we present two ways of computing a CVT. The first way (a) called "exact" approach in the sequel, consists in minimizing an energy function f_{CVT} by computing systematically all needed intersections between the VOI and the Voronoi diagram. The second way (b) called "hierarchical heuristic" in the sequel, is a multi-level approach which first computes the tessellation over a box containing the VOI, by a divide-and-conquer strategy, and then computes (and post-processes) the intersection with the VOI.

A. Gradient of f_{CVT} under a Continuous Anisotropy Field

In this section we introduce our main theoretical result: an expression of $\nabla f_{\text{CVT}}^{\text{aniso}}$ for M continuous.

In the sequel, Ω denotes the VOI, $\mathcal{V} = \{x_i\,;\, i = 1 \dots N\}$ the set of seeds, V_i the Voronoi cell associated with x_i, $\mathbb{V}(x_i)$ the set of neighbor seeds of x_i in the Voronoi diagram $\mathscr{V}$ and $\mathbf{x}$ the configuration $(x_1, \dots, x_N) \in \mathbb{R}^{3N}$.
Let us now state our main theoretical result:

Proposition 1. *Let M be a smooth field and let us assume: first that Ω is convex, second that the topology is locally invariant (i.e. given a configuration $\mathbf{x_0}$: for each $i \in \{1 \dots N\}$, there is an open neighborhood B^i of the i-th seed x_i of $\mathbf{x}_0$, such that, for all $\mathbf{x} \in \prod_i B^i$, the topology of $V_i(\mathbf{x})$ remains constant). Then, the gradient of f_{CVT} exists at $\mathbf{x}_0$ and is given by the following formula:*

$$\frac{\partial f_{\text{CVT}}}{\partial x_{ij}}(\mathbf{x}_0) =$$

$$p \int_{V_i(\mathbf{x}_0)} \left(e_j^3\right)^\top M^\top \begin{pmatrix} (M_1[x_i - y])^{p-1} \\ (M_2[x_i - y])^{p-1} \\ (M_3[x_i - y])^{p-1} \end{pmatrix} \mathrm{d}y \qquad (*)$$

$$+ \sum_{\substack{\Gamma_{ik} \\ x_k \in \mathbb{V}(x_i)}} \int_{\Gamma_{ik}} \left(\left\| M[x_i - z] \right\|_p^p - \left\| M[x_k - z] \right\|_p^p \right) \frac{(z_j - x_{ij})}{\|x_k - x_i\|} \, \mathrm{d}\sigma$$

where $i \in [1, N]$, $j \in [1, 3]$, z is a parametrization of the corresponding facet $\Gamma_{ik} := V_i(\mathbf{x}_0) \cap V_k(\mathbf{x}_0)$ in surface integrals, and e_j^3 is the j-th vector of the canonical basis of $\mathbb{R}^3$. Moreover, f_{CVT} has C^1-continuity at $\mathbf{x}_0$.

The proof is based on the generalized REYNOLDS theorem and is provided in the appendix, page , together with variants of Proposition 1.

B. Evaluation of f_{CVT} and ∇f_{CVT}

As illustrated in Fig. 2, minimization of f_{CVT} under a continuous anisotropy field is at the heart of our approach. In this section, we describe in more detail our related calculations and implementations.

Computing the value and the gradient of f_{CVT} entails calculations of integrals over the volume of interest. For this, at each step, we need to restrict the Voronoi diagram and hence to compute efficiently the intersections of Ω with the current Voronoi diagram. For efficiency reasons, we use the approach described by Yan and Lévy et al. in [18], [20]. It consists in two steps: first, computing the intersection of the current Voronoi diagram with the boundary of Ω, and second, computing the intersections between the tetrahedra of the VOI and Voronoi cells. These intersections are computed by propagation from 2D boundary intersections which have been determined during the first step. Note that internal Voronoi cells are not affected by these intersections and can be obtained directly from the Delaunay triangulation. This improvement is necessary to get a time-saving meshing algorithm. Since we use a continuous field M, it has no impact on the quality of our computations. But regarding discrete anisotropic fields which are piecewise constant on the VOI, finding all intersections of an internal Voronoi cell with simplices of the VOI is compulsory to compute precisely the energy.

Then, Voronoi cells, which have been decomposed into convex pieces by the intersection algorithm, are, in turn, tetrahedralized. Moreover, facets of RVCs which are not on the boundary of Ω are triangulated. Numerical integration is thus reduced to calculations on tetrahedra and triangles. These computations can be performed with high-order cubature rules on tetrahedra and triangles (which can be found in [21], [22]). Theses rules allow the exact calculation of integrals of polynomials on simplices.

C. Implementations of the "Exact" Method

This first approach produces a polyhedral mesh of an input tetrahedral mesh Ω, the VOI. It consists in optimizing seeds to obtain a Voronoi diagram which abides by a continuous anisotropic field M. An advantage of our contribution is that it does not depend on the quality of the input mesh, because no discretization of M is needed. Moreover, unlike pre-cited methods, the field M can be easily expressed and interpreted by users (no discretization), as described in subsection III-B.

Fig. 3 gives an overview of our "exact" approach, to compute an anisotropic CVT with an input power $p = p_0$ (the larger p, the more hexahedral the cells). As illustrated, p is gradually increased to improve the convergence (see subsection IV-C2 for details).

1) Solvers: The Lloyd algorithm converges when M is a multiple of identity, that is under density constraints only, with M of the form $M(y) = \alpha(y)\mathbb{I}_3$ with $\alpha : \mathbb{R}^3 \to \mathbb{R}$ a continuous function and $p = 2$. Actually, under anisotropy constraints, CVTs are no longer centroidal. So, the Lloyd algorithm should be theoretically ineffective for anisotropy (even for $p = 2$).

Then, to obtain the anisotropy, we use two gradient-based methods with preconditioning: L-BFGS and steepest descent.

We use an approximation of the Hessian matrix of f_{CVT}, obtained by deriving ∇f_{CVT} and by putting aside all surface integrals. The i-th diagonal block of this Hessian matrix $H(f_{\text{CVT}})(\mathbf{x}_0)$ is the following 3×3 matrix:

$$p(p-1) \int_{V_i} M^\top \operatorname*{diag}_{k=1,2,3} \left((M_k[x_i - y])^{p-2} \right) M \mathrm{d}\lambda \qquad (6)$$

which has the virtue of being symmetric, positive definite. Actually, this approximation is interesting as, should one want to be more precise, adding surface integrals like in [13] results in building a non-symmetric matrix (in the general case) which is almost singular in practice.

Another important issue in descent methods is the choice of the step. We use strong WOLFE conditions [16, p. 34], because if they are satisfied at each iteration of a descent method, convergence is ensured. All tested solvers (with preconditioning or not) managed to minimize f_{CVT} during our tests, although L-BFGS seems to be the most efficient method here.

Last, the stopping criterion of descents is based on the difference between two successive configurations. In order to estimate such displacement lengths, we have chosen the following norm on configurations: $\|\mathbf{x}\| := \max_i \|x_i\|_2$.

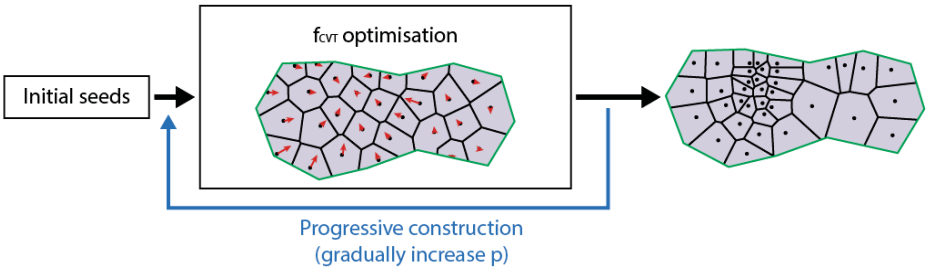

Fig. 3 Overview of our exact approach for restricted anisotropic Voronoi tessellation

2) Progressive Strategy: As previous sections show, f_{CVT} is C^1 but is not a priori C^2, and its evaluation, as well as the evaluation of its gradient, are computationally costly. As a consequence, a descent towards its minimum is delicate. In order to improve the convergence, as illustrated by Fig. 3, we have put in place a "progressive strategy":

1) Initial configuration $\mathbf{x}_0$: seeds are randomly placed in Ω with draws of probability density $(\int_\Omega \rho \, \mathrm{d}\lambda)^{-1}\rho$;
2) Lloyd iterations are performed, starting from $\mathbf{x}_0$, to abide by ρ: let us denote by $\mathbf{x}_{\text{Lloyd}}$ the minimizer obtained;
3) If $p > 2$, starting from $\mathbf{x}_{\text{Lloyd}}$, a preconditioned gradient-based method is used to minimize $f_{\text{CVT}}^{\text{aniso}}$ for $p = 2$. We thus obtain $\mathbf{x}_{\text{Opt2}}$;
4) Last, a solver is applied, starting from $\mathbf{x}_{\text{Opt2}}$ to minimize $f_{\text{CVT}}^{\text{aniso}}$ with norm powers gradually increasing from 2 to p.

D. Hierarchical Heuristic Method

Finding a local minimum of f_{CVT} is time-consuming hence, developing faster alternative methods to optimize f_{CVT} is necessary to produce large meshes. As a third contribution, we introduce a hierarchical approach, which capitalizes on a fast splitting and a fast merging procedure, together with a VOI conformity heuristic. Our hierarchical heuristic is summarized in Fig. 4.

First, let us point out that the most time-consuming step in previous "exact" method (as well as in state-of-the-art methods) is the computation, at each descent step, of the intersections between the Voronoi diagram and the VOI. Such computations can be largely accelerated when the VOI is a box. Therefore, we first compute the tessellation over Ω, a bounding box of the VOI, and introduce, as a last step, a conformity heuristic method (labelled "post-processing" in Fig. 4, and described in section IV-D3). Over the rectangular domain Ω, our approach introduces two complementary heuristics: an adaptive subdivision and a progressive merging.

Adaptive subdivision (labelled a) in Fig. 4) splits the original problem $\min_{x_i \in \Omega} f_{\text{CVT}}$ into subproblems $\min_{x_i \in \Omega_k} f_{\text{CVT}}$. The collection $\{\Omega_k; k = 1 \dots K\}$ is a covering of Ω (actually a kd-tree) such that all leaves have similar weights and are "as square as possible". Then optimisations of f_{CVT} over the leaves are carried out simultaneously and independently. We minimize f_{CVT} by optimizing seeds as described in previous sections. The resulting set of seeds produces a first approximate tessellation over Ω (labelled d) in Fig. 4).

Progressive merging (labelled e) in Fig. 4) is a bottom-up procedure which performs fast constrained re-optimization between kd-tree nodes.

Note that having several independent subproblems is profitable for two reasons: optimisation converges more quickly and parallelism can be used.

A last word about density and anisotropy fields: we assume these fields are defined over the bounding box of the volume of interest (which is often the case in practice when they are defined by continuous functions).

In the remainder of this section, we describe more precisely each step of this procedure.

1) Adaptive Subdivision: The number of seeds initially set in a domain is an important parameter which determines the final local size of the mesh (density and anisotropy actually control the ratio between cell size and shape). The goal of our adaptive subdivision is to produce subdomains $(\Omega_k)_k$ with similar weights (and secondarily, "as square as possible"). Each subdomain will eventually receive N initial seeds for further steps (N being constant). Hence, our algorithm yields many subproblems having similar complexities to make parallelism more efficient.

To this end, we build a kd-tree. The box Ω is recursively divided to produce subdomains $(\Omega_k)_k$ which satisfy both evoked constraints: similar weights and square-like shapes. At each step, we cut cells along the x, y or z-axis. The position of the cutting plane is computed so as to balance weights in both sub-cells. Let us assume that a given box Ω_k is being cut along the first dimension (i.e. the x-axis) during the recursive division algorithm. We can write this box as the Cartesian product of an interval $[a, b]$ by a rectangle R. Then the following function $\delta_1 : [a, b] \to \mathbb{R}^+$ computes for any $t \in [a, b]$ the (positive) difference between the weights in both sub-domains $[a, t] \times R$ and $[t, b] \times R$:

$$\delta_1(t) := \left| \int_{\Omega_k} \rho \, \mathrm{d}\lambda - 2 \int_{x \in [a,t]} \int_R \rho(x, y, z) \, \mathrm{d}y \, \mathrm{d}z \, \mathrm{d}x \right| . \quad (7)$$

We similarly define δ_2 and δ_3, for subdivisions along the second and third dimensions respectively. These functions are one-dimensional, continuous[1] and one easily shows that they are unimodal[2], hence derivative-free methods such as the golden-section method are enough to compute their minimum.

Let us point that given a parameter t, $\delta_i(t)$ can be evaluated quickly and efficiently. Indeed, the left integral (total weight

[1]provided that ρ is also continuous.
[2]provided that ρ is positive almost everywhere.

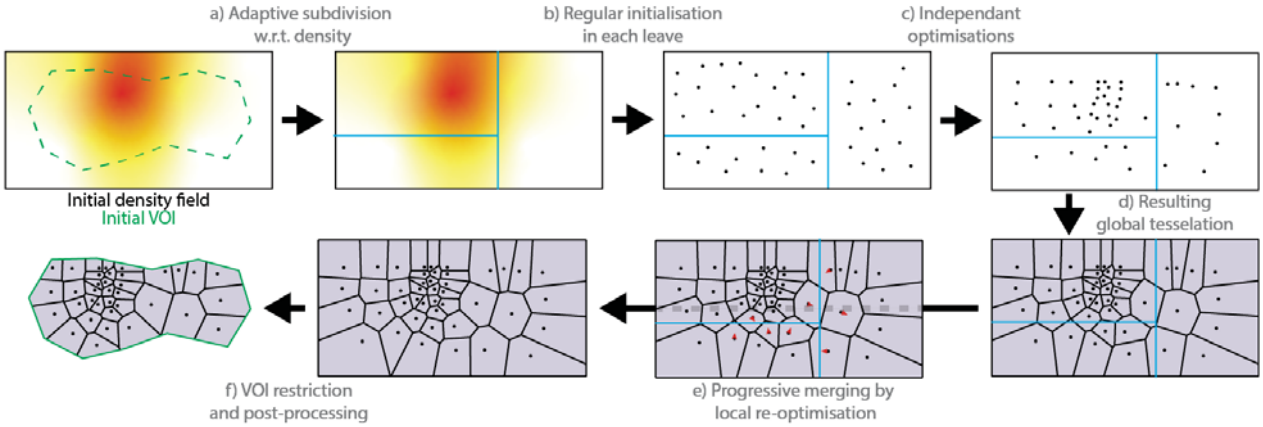

Fig. 4 Overview of our hierarchical heuristic approach

over Ω_k) is known from a previous subdivision, while the right integral can be efficiently evaluated using the VEGAS algorithm of LEPAGE dedicated to parallelepipedic domains.

Our hierarchical subdivision routine (labelled a) in Fig. 4) returns two bounding boxes Ω_1 and Ω_2 and two numbers of seeds N_1 and N_2 such that

$$\int_{\Omega_1} \rho \, \mathrm{d}\lambda \approx \int_{\Omega_2} \rho \, \mathrm{d}\lambda \quad \text{and} \quad N_1 \approx N_2 \ . \tag{8}$$

We actually perform three optimisations, one for each axis, i.e. we optimize δ_i with $i = 1 \dots 3$ and select the division which minimizes the worst shape ratio of both sub-boxes.

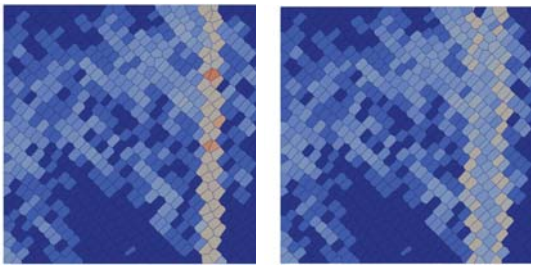

(a) (b)

Fig. 5 Tessellation obtained with the hierarchical approach for a single cutting plane, a constant anisotropy (with theoretical shape ratio 2) and Ω a "flat" parallelepiped: before local re-optimisation (a) vs. after (b)

2) Merges and Local Re-optimisation: Let us now give more details about merging and re-optimisation (labelled e) in Fig. 4. Given two brother nodes Ω_1 and Ω_2 of the binary tree and their respective point clouds x_1 and x_2, a new point cloud x is obtained by concatenation. It is actually close to a local minimizer of f_{CVT} over $\Omega := \Omega_1 \cup \Omega_2$. Indeed, away from the junction $\Omega_1 \cap \Omega_2$, as f_{CVT} minimization depends only on Voronoi neighborhoods, the cells are optimal. Therefore, only the cells close to the junction need to be adjusted; hence our local re-optimisation approach. It is a *constrained* local optimisation in order to preserve Voronoi the cells away from the junction.

Let us define it more precisely. Let $\bar{V} = \{\bar{x}_1, \dots, \bar{x}_m\}$ be a subset of seeds, considered for re-optimisation.

We respectively denote by $\mathcal{V}_1(\bar{V})$ and $\mathcal{V}_2(\bar{V})$ the 1-neighborhood and 2-neighborhood of seeds of $\bar{V}$ in the Voronoi tessellation generated by V. The goal of our local optimisation is to limit the complexity by moving only few seeds (actually seeds in $\bar{V}$). For that, only energies of the cells generated by seeds in $\bar{V} \cup \mathcal{V}_1(\bar{V})$ need to be computed. Last, we need seeds from $\mathcal{V}_2(\bar{V})$ to update Voronoi cells associated with $\mathcal{V}_1(\bar{V})$.

In order to guarantee the correctness of this approach, cells outside the 1-neighborhood of $\bar{V}$ must remain unchanged. In order to do so, we set the following constraint: for each seed $\bar{x}_k \in \bar{V}$, $\bar{x}_k$ must lie outside all circumspheres of the Delaunay triangulation of V which do not stem from a seed in $\bar{V}$ (we denote by $C_{\bar{V}}$ this set of circumspheres). In order to perform this constrained optimisation, we define a penalty function $p_i : \mathbb{R}^3 \to \mathbb{R}^+$ for each seed $\bar{x}_i$ in $\bar{V}$. When the seed $\bar{x}_i$ abides by the constraint, $p_i(\bar{x}_i)$ is 1, otherwise, its value is larger than 1. For seeds x_k which do not belong to $\bar{V}$, a constant penalty function $p_k \equiv 1$ is also defined. Thus, the local optimisation is carried out on a "restricted" version of f_{CVT}. Let $m = |\bar{V}|$, let $\mathbf{x} \in \mathbb{R}^{3N}$ and $\bar{\mathbf{x}} \in \mathbb{R}^{3m}$ be respectively coordinates of seeds associated with V and $\bar{V}$. The restricted version of f_{CVT} is:

$$f_{\mathrm{CVT}}^{\mathrm{restrict}}(\bar{\mathbf{x}}) = \sum_{x_i \in \bar{V} \cup \mathcal{V}_1(\bar{V})} p_i(x_i) \cdot \mathcal{E}_i^{\mathrm{aniso}}(\mathbf{x}) \ . \tag{9}$$

This local optimisation is computationally interesting: each evaluation of $f_{\mathrm{CVT}}^{\mathrm{restrict}}$ entails a "smaller" Voronoi diagram (with seeds in $\bar{V} \cup \mathcal{V}_1(\bar{V}) \cup \mathcal{V}_2(\bar{V})$) and only $| \bar{V} \cup \mathcal{V}_1(\bar{V}) |$ energies are evaluated.

The partial derivative of $f_{\mathrm{CVT}}^{\mathrm{restrict}}$ along the j-th coordinate of $\bar{x}_i \in \bar{V}$ at the configuration $\bar{\mathbf{x}}$ is given in Appendix VII-A.

Technically, p_i is computed as follows. We assign to each seed $\bar{x}_i \in \bar{V}$ of the initial configuration $\bar{\mathbf{x}}$, a constant ball B_i which is the larger ball centered on $\bar{x}_i$ which does not intersect a circumsphere of $C_{\bar{V}}$. Then p_i is simply defined as 1 inside B_i and quadratically tends to infinity outside.

Fig. 5 illustrates the result of this re-optimization along a single cutting plane (cells are colored according to the preconditioned gradient of f_{CVT}). For this example, Fig. 6 compares the energies of the "exact" minimization of

f_{CVT} (introduced in Section IV-C) with L-BFGS algorithm and of the hierarchical approach (before and after local re-optimisation). The blue line is the convergence history of the "exact" minimisation of f_{CVT} with L-BFGS algorithm (red bullet points indicate computations of the preconditioner and the orange line shows the minimum reached). Regarding the hierarchical method, the purple line and the green one are respectively the energies of the produced tessellation, before and after the local re-optimisation. We can notice that the energy f_{CVT} after the local re-optimisation is very close to the minimum reached by the "exact" (global) optimisation.

Tables I and II provide statistics on $\|\nabla f_{\mathrm{CVT}}\|$ (which somehow measure the non-minimality of a configuration) and shape ratios for all cells, both for the hierarchical results of Fig. 5 (before and after local optimisations) and for the "exact" optimisation. Both tables show that the local re-optimisation combined with previous steps of the hierarchical heuristic produce results similar to the "exact" method, both in terms of energy and shape ratios. Table II shows that for this constant anisotropy example, shape ratios converge to those obtained using the "exact" method. Moreover, worst cases are similar.

Actually, the gap with the expected result (the anisotropy field is constant with a shape ratio of 2) stems from the "flat" shape of the VOI where almost all cells are boundary cells.

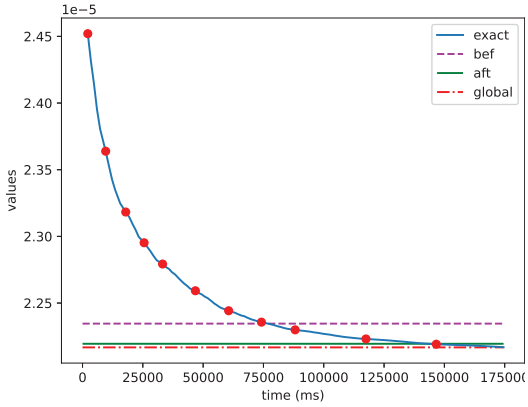

Fig. 6 For the example of Fig. 5, comparison of f_{CVT} energies

TABLE I
FOR THE EXAMPLE OF FIG. 5: STATISTICS ON $\|\nabla f_{\mathrm{CVT}}\|$ AT EACH SEED, AND ENERGIES OF THE TESSELLATIONS

-	Hierarchic (bef)	Hierarchic (aft)	Exact
Mean	$7.672 \cdot 10^{-6}$	$5.570 \cdot 10^{-6}$	$4.407 \cdot 10^{-6}$
Deviat.	$1.425 \cdot 10^{-5}$	$7.853 \cdot 10^{-6}$	$6.819 \cdot 10^{-6}$
Min	$8.625 \cdot 10^{-8}$	$8.625 \cdot 10^{-8}$	$6.889 \cdot 10^{-8}$
Max	$1.388 \cdot 10^{-4}$	$5.225 \cdot 10^{-5}$	$3.793 \cdot 10^{-5}$
Energy	$2.23457 \cdot 10^{-5}$	$2.21946 \cdot 10^{-5}$	$2.21675 \cdot 10^{-5}$

TABLE II
FOR THE EXAMPLE OF FIG. 5: STATISTICS ON SHAPE RATIOS

-	Hierarchic (bef)	Hierarchic (aft)	Exact
Mean	1.608	1.620	1.631
Deviation	0.184	0.173	0.170
Min	1.014	1.021	1.021
Max	2.433	2.433	2.358

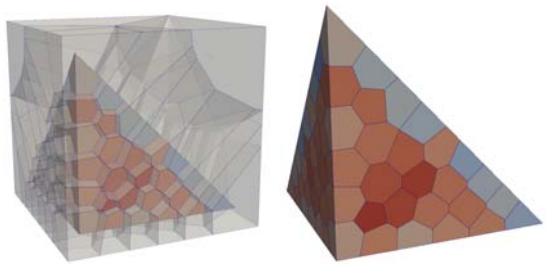

(a) (b)

Fig. 7 Result from the hierarchical method with a density modified according to the technique introduced in subsubsection IV-D3 (a), final result obtained after clipping (b)

3) Conformity Heuristic: In order to obtain conformity with respect to the VOI, state-of-the-art methods as well as our "exact" method intersect the VOI and the tessellation at each optimisation step, which is extremely costly. Our conformity heuristic relies on light pre-processing and post-processing procedures.

Our approach was designed in order to keep post-processing as simple as possible. It only consists in computing the intersection between the tessellation obtained through previous steps with the VOI. Then all seeds associated with clipped cells of low volume are discarded and we compute the resulting restricted tessellation (which naturally "fills the gaps" left by discarded seeds).

However, when the VOI is complex, the tessellation thus obtained may still contain cells of poor quality along the boundary. Pre-processing then consists in preparing data (anisotropy field or number of initial seeds) in order to make this simple post-processing efficient. Two heuristics can be used.

First, the underlying density ρ of the anisotropy field can be modified (progressively decreased outside the VOI): ρ is multiplied by a very low factor far from the VOI. Thus, the density forces seeds to stay in or close to the real VOI during optimisation. This approach is illustrated in Fig. 7 for an isotropic CVT over a tetrahedron.

Another technique consists in simply inserting more seeds to "fill space outside the VOI": if we want to produce roughly N cells in a VOI Ω included in a box B, the optimisation should be rather done with $\lfloor N \int_B \rho \, \mathrm{d}\lambda / \int_\Omega \rho \, \mathrm{d}\lambda \rfloor$ seeds.

V. VALIDATION

Our validation aims at comparing our approach with the most advanced state-of-the-art works, namely, the works by Lévy and Liu, in [2] (see Section II). In this paper, they provide a gradient ∇f_{CVT} in the restricted case of a field M which is piecewise constant on Ω. We compare the results obtained with our new general gradient formula against their approach and exhibit similar results in this restricted context. Besides this generalization of the gradient computation, we demonstrate that we are able to combine the three types of constraints (density, aspect ratio and orientation).

To carry out these comparisons, we build a continuous field M together with its discrete version on Ω (because Lévy and

Liu's approach necessitates a piecewise constant field). The software GMSH [23] has provided the fine discretized versions of these VOIs.

A. Methodology

Test parameters are: the density function ρ, the matrix-valued field of directions P, three "magnitude" functions μ_i, the power p and the number of seeds N. Then the field M is built as follows:

$$\forall y \in \Omega, M(y) = (\rho)^{1/p} \frac{1}{|||DP|||_p} DP \tag{10}$$

where $D(y) := \mathrm{diag}(\mu_1, \mu_2, \mu_3)$ and $P(y)$ is preferably orthogonal.

The discretized version of M is simply the field whose value in tetrahedron τ of Ω is the value of M at the centroid of τ. We follow the overall progressive approach described in Section IV-C2.

B. Test 1: Lévy's Test Case

We adapt to the 3D context the density function ρ from example 4 in [13] (only 2D in this paper). We add a "radial" anisotropy to induce facet orientations towards the center of the VOI, by setting $p = 6$ and $N = 3,000$.

C. Test 2: Density, Aspect Ratios and Orientations

We build a mesh of a sphere drilled with a cylinder C. Let d be the distance of y from the axis of C, then we choose $\mu_1 = 1$ while $\mu_2 = \mu_3$ are decreasing linear functions of d, whose values are in the interval $[\frac{1}{2}, 1]$. The density ρ is the inverse of a quadratic function of d. We take $p = 4$ to impose orientations of cells.

D. Test 3: Geological Case

This case shows possible applications of our meshing algorithm. Here we go from a simplified geological model, which was given by a tetrahedral mesh and a piecewise linear continuous stratigraphy. This stratigraphy functions defines a "watertight" model. To get a fine mesh close to its center, in order to simulate a well at this place, we choose a "quadratic" density function. We also define a radial anisotropy, turned towards the center. Each region is generated using our method with $N = 10,000$ seeds.

E. Test 4: Hierarchical Heuristics

We performed some tests to show results and performances of this approach. We first run three examples by setting $N = 100,000$:

- TR for a radial anisotropy in a torus,
- TC for a constant anisotropy in a torus,
- BD for the density from Test 1 in a box.

The density in BD has been rotated to challenge our hierarchical method and avoid interactions between the symmetries of the model and subdivision planes.

Tessellations are shown in Fig. 11 (for the hierarchical method, the depth of trees is limited to 8).

F. Test 5: Heuristic Method in a Refined Box

We take anisotropies from Test 4 again (i.e. radial, constant and periodic anisotropies) but this time, on a parallelepipedic domain. To carry out comparisons, we build a tetrahedral mesh of this box by constructing a homogeneous Cartesian mesh with $20 \times 20 \times 20$ hexahedra which have been tetrahedralized.

We compare running times for Lévy's method (T_{Lev}) and our hierarchical approach. More precisely, we run our method with depths limited to 1 and 4 respectively (in order to assess the impact of the subdivision approach). $T_{\mathrm{H}\,1}$ and $T_{\mathrm{H}\,4}$ stand for the respective running times.

Note that for this test, the heuristic method with depth 1 does not completely equal the "exact" method, since our hierarchical method performs the intersection between the VOI and the Voronoi diagram only once (during the post-processing step).

VI. Results and Analysis

A. Qualitative Results

With regard to test 1, CVTs on Fig. 8 seem to comply with the density ρ since we obtain the same patterns as in [13]. If the reader looks carefully, some cells are oriented towards the center of the mesh. Moreover, a lot of facets look like quadrangles.

Regarding test 2, CVTs on Fig. 9 make radial anisotropy, shape ratios and effects of the density clearly appear. Cross sections show that cells are smaller near the axis of the hole, while longitudinal sections show that the farther from the axis cells are, the more lengthened they are.

Regarding test 3, although all shape ratios have been set to 1, we can also observe on Fig. 10 that cells are aligned according to concentric circles, and volumes of cells strongly decrease near the well.

Regarding test 4, results produced by hierarchical heuristics are very similar to ones obtained thanks to the "exact" method. There are differences but they remain invisible to the naked eye. Meshes have the same pattern in Fig. 11 in three cases. Moreover, there is no sign of cutting planes anywhere, even for the periodic anisotropy example (case BD).

B. Analysis

a) Exact approach: According to Table III, our "exact" method and Lévy and Liu's method (LL for short) seem to be complementary. For simple anisotropy fields (e.g. piecewise constant fields), LL outperforms the "exact" method, while the converse holds for more complex fields.

Actually, regarding LL, for each descent step, energy computations are negligible against computations of the restricted Voronoi diagrams. However, to abide by density ρ, the mesh of the VOI must be refined to catch density variations (and hence strongly refined for complex anisotropy fields). The complexity resulting from intersections computation is thus higher. This could explain the ratio 0.4 observed in table III.

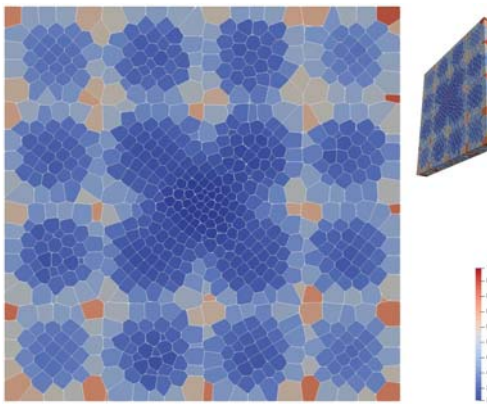

(a)

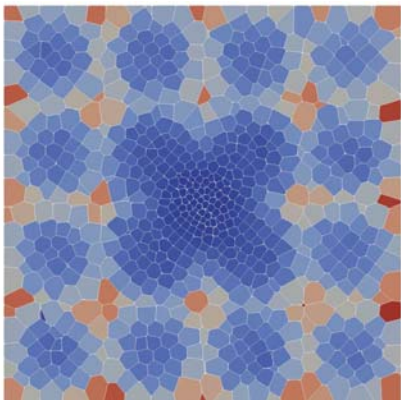

(b)

Fig. 8 Test 1: CVTs given by the preconditioned L-BFGS method, with our gradient (a) and with LÉVY's method (b), $N = 3,000$; cells are colored according to volumes

b) Hierarchical approach: As far as using the hierarchical approach is concerned (Tables III and IV provide data on running times for tests 4 and 5), it is always more efficient than the "exact" method and almost always faster than LÉVY and LIU's approach. Actually, the TR example (torus with radial anisotropy) exhibits the limits of our hierarchical approach. When the volume of the VOI is low against that of the bounding box, the extra cost resulting from the tessellation of the full bounding box may penalize performances (or not, see for instance, the TC case). In both the TR and TC examples, optimisation is carried out with the second pre-processing method and uses $326,000$ seeds.

On the contrary, when the VOI is (almost) parallelepipedic (such as in test 5), the hierarchical approach outperforms all other methods. Experimentally, the most efficient subdivision builds a kd-tree with $1,000$ seeds by leaf (actually in our

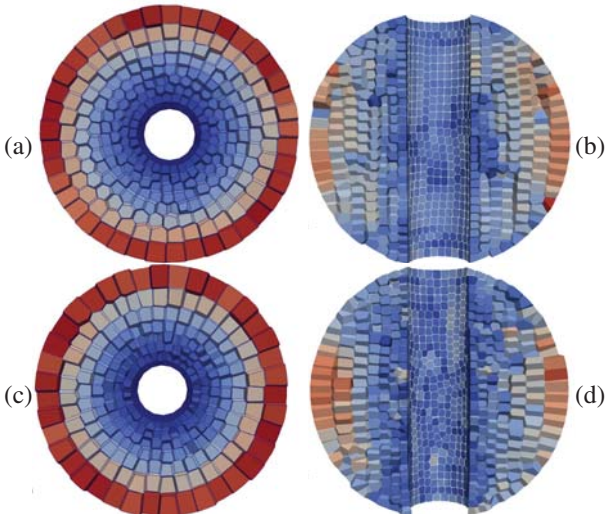

Fig. 9 Test 2: Cross section and longitudinal section of the CVT obtained by preconditioned L-BFGS with our method ((a), (b)) and LÉVY's gradient ((c), (d)), $N = 5,000$; cells are colored according to volumes ((a), (c)) and to shape ratios ((b), (d))

tests, we voluntarily bounded the depth in order to be able to evaluate the impact of depth on running times).

Let us also point out that tests have been executed on standard laptops with 6 cores; the ratio should be even more in favor of the hierarchical approach on computers with more cores.

c) Conformity along the VOI and hierarchical approach: In spite of our conformity heuristic (with both pre-processing and post-processing procedures), boundary cells in produced tessellations do not fully comply with anisotropy constraints. However, this is also the case in "exact" approaches (both LÉVY and LIU's and ours). Actually, wether anisotropy constraints are satisfiable or not along the boundary depends on the coherence between the anisotropy field and the boundary geometry (and not only on the optimisation method).

Experimentally, tessellations produced by our hierarchical fairly fulfill anisotropy along the VOI. The first pre-processing approach, which decreases the density away from the VOI gives better results (as illustrated in Fig. 7).

However, in order to get CVTs with better-shaped boundary cells, conformity constraints should be considered along the whole optimisation process, like in [24].

TABLE III
TEST 4: RATIOS BETWEEN RUNNING TIMES ON TR, TC AND BD EXAMPLES WITH $N = 100,000$ SEEDS FOR LÉVY'S METHOD (T_{Lev}), OUR "EXACT" METHOD (T_{Exact}) AND OUR HIERARCHICAL METHOD WITH DEPTH 8 ($T_{\text{H}\,8}$)

Example	$T_{\text{Exact}}/T_{\text{Lev}}$	$T_{\text{H}\,8}/T_{\text{Exact}}$	$T_{\text{H}\,8}/T_{\text{Lev}}$
TR	1.84	0.75	1.38
TC	2.03	0.26	0.52
BD	0.4	0.26	0.1

VII. CONCLUSION

This paper presents an expression of ∇f_{CVT} for a continuous anisotropy field, which we carefully prove. Starting from this

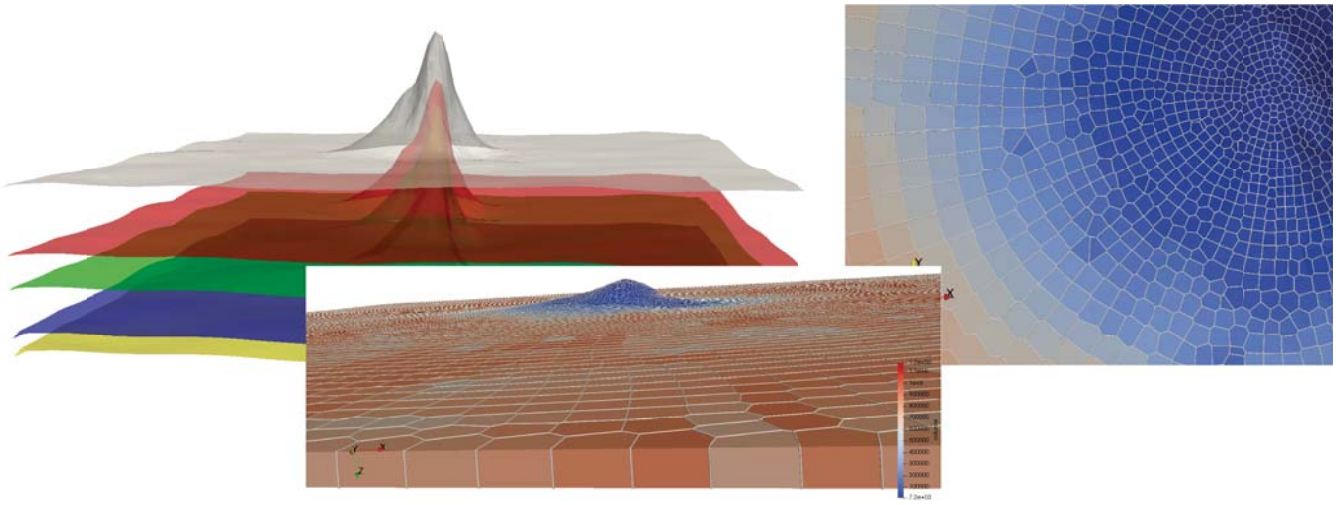

(a) (b)

Fig. 10 Test 3: Model horizons (a), zoom of top of the resulting mesh (b) and its three-quarter view (c)

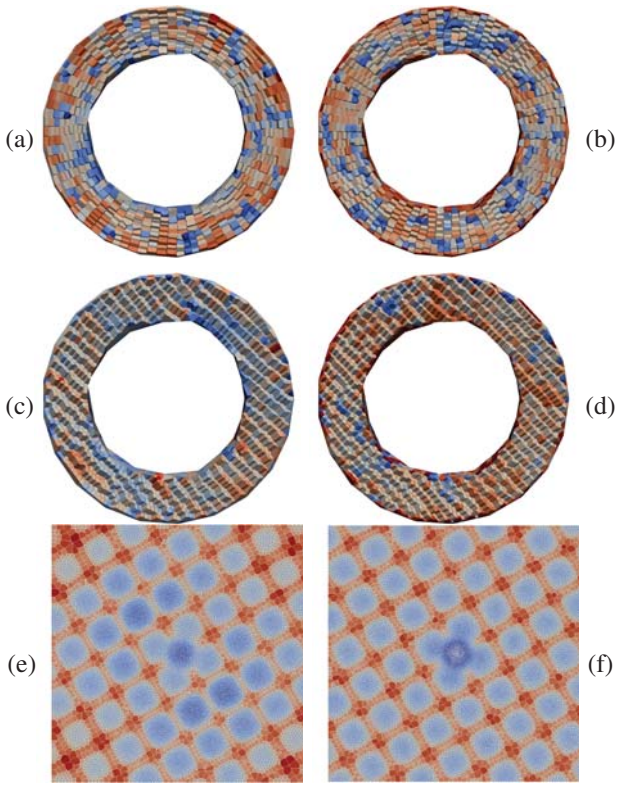

Fig. 11 Test 4: Exact method ((a), (c), (e)) vs. hierarchical heuristics ((b), (d), (f)) for TR test (radial anisotropy in a torus - (a), (b)), TC (constant anisotropy in a torus - (c), (d)), BD (periodic density and anisotropy from Test 1 - (e), (f))

formula, we propose two approaches to generate anisotropic Voronoi tessellations of volumes of interest. Our first approach, called "exact", uses this gradient to compute a tessellation minimizing f_{CVT} for such a continuous anisotropy. Then, in order to speed up and to allow the parallelization of optimisation, we introduce a hierarchical divide-and-conquer approach.

TABLE IV
TEST 5: RATIOS BETWEEN RUNNING TIMES ON RADIAL, CONSTANT AND PERIODIC ANISOTROPIES WITH $N = 10,000$ SEEDS AND PARALLELEPIPEDIC VOI FOR LÉVY'S METHOD (T_{Lev}) AND OUR HIERARCHICAL METHOD WITH DEPTHS 1 ($T_{\mathrm{H}\,1}$) AND 4 ($T_{\mathrm{H}\,4}$) RESPECTIVELY

Example	$T_{\mathrm{H}\,1}/T_{\mathrm{Lev}}$	$T_{\mathrm{H}\,4}/T_{\mathrm{H}\,1}$	$T_{\mathrm{H}\,4}/T_{\mathrm{Lev}}$
Radial	0.39	0.2	0.08
Constant	0.22	0.43	0.1
Periodic	0.014	0.15	$2.2 \cdot 10^{-3}$

Then, we have shown that our methods and LÉVY and LIU's method lead to equivalent results, that is, to CVTs where volumes of cells, shape ratios of cells and orientations of facets abide by an anisotropic matrix-valued field M. The "exact" method proves faster for complex anisotropy fields. The hierarchical method generally outperforms all other methods at the cost of a poorer conformity to the anisotropy field near the boundary. Furthermore, unlike state-of-the-art methods, the quality of our results does not depend on the discretization of the VOI with regard to the considered anisotropy field, as illustrated on Fig. 1.

In the general case, the efficient construction of a continuous anisotropy field M from user data, which can integrate density, shape ratios and orientation constraints, is not a trivial issue. Neither is the extension of such a field to the bounding box of the VOI. These questions will be addressed in a future work.

Moreover, we intend to improve our hierarchical approach by integrating anisotropy constraints along the boundary to enhance conformity.

APPENDIX

Detailed Proof of ∇f_{CVT} and the Regularity of f_{CVT}

Here are the hypotheses required for the proofs:

1) VOI definition: Ω is a conformal simplicial complex.
2) VOI convexity: Ω is convex.
3) Uniqueness of the Delaunay triangulation in an open neighborhood of a configuration $\mathbf{x}_0$ (which is ensured

when seeds are in *general position*, i.e. no five seeds lie on a common empty sphere). Moreover, seeds from $\mathbf{x}_0$ are pairwise different.

4) Non-degenerate configurations: let us consider a Voronoi diagram generated by $\mathbf{x}_0$. For each *original* Voronoi cell V_i, if a facet, an edge or a vertex e of V_i intersects a boundary d-simplex s of Ω with $d < 3$, the affine hull of element e must not contain and must not lie in the affine hull of s.
5) Non-degenerate configurations (2): let us consider a Voronoi diagram generated by $\mathbf{x}_0$. For each *original* Voronoi cell V_i, if a facet, an edge or a vertex e of V_i intersects a d-simplex s of Ω with $d < 3$, the affine hull of element e must not contain and must not lie in the affine hull of s.
6) Local invariability of the topology of the RVD given by a configuration $\mathbf{x_0}$: for each $i \in \{1 \dots N\}$, there is an open neighborhood B^i of the i-th seed x_i of $\mathbf{x}_0$, such that, for all $\mathbf{x} \in \prod_i B^i$, the topology of $V_i(\mathbf{x})$ remains constant.

Assumptions 1 and 3 are recurrent in our results. Assumptions 4 and 5 are quite similar but the later is more restrictive and will be required when M is piecewise constant on Ω. Last, assumption 6 is necessary to prove the expression of ∇f_{CVT}. Note that conjunction of assumptions 3 and 4 is roughly equivalent to assumption 6.

In order to prove our results, we will need integration and derivation formulas which are well-known for simple cases, e.g. for smooth domains that do not depend on variables. According to [25], many notions coming from differential geometry can be generalized to *manifolds with corners*: such domains are similar to manifolds with boundary, but they can also have "corners" which are topologically equivalent to tetrahedron corners. Paper [26] proves that the *differentiation lemma* still holds for manifolds with corners. In particular, tetrahedra are manifolds with corners and hence, the differentiation lemma indirectly holds for polyhedra. The REYNOLDS transport theorem is a particular case of this differentiation lemma, as the STOKES theorem also holds for manifolds with corners (see [25]). Hence, the following generalized theorem holds:

Theorem (Generalized REYNOLDS theorem)**.** *Let I be an open interval which contains 0. Let $(V(t))_{t\in I}$ be a family of domains indexed on I such that:*

- *For any $t \in I$, $V(t)$ is a manifold with corners of dimension 3 of $\mathbb{R}^3$.*
- *For any t, there exists a smooth homeomorphism φ_t between $V_0 := V(0)$ and $V(t)$.*

Let us denote the boundary of $V(t)$ by $S(t)$ and let $\vec{n}$ be the outward normal vector field of $S(t)$.

Let $f : \mathbb{R}^3 \times I \to \mathbb{R} : (x,t) \mapsto f(x,t)$ be a function such that:

- *For any $t \in I$, $f(\cdot,t)$ is smooth and integrable on $V(t)$.*
- *For any $x \in V_0$, $t \mapsto f(\varphi_t(x),t)\,|\det J\varphi_t|$ admits a derivative on I.*
- *There exists a smooth integrable function $g : V_0 \to \mathbb{R}$ such that, for any $x \in V_0$ and $t \in I$*

$$\left|\frac{\partial}{\partial t}\left(f(\varphi_t(x),t)|\det J\varphi_t|\right)\right| \leqslant g(x)\ . \qquad (**)$$

Then,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\int_{V(t)} f(x,t)\ \mathrm{d}x\right)\bigg|_{t=0} = \qquad (11)$$

$$\int_{V_0} \frac{\partial}{\partial t}\left(f(x,t)\right)\bigg|_{t=0} \mathrm{d}x + \int_{S(0)} f(z,0)v\cdot\vec{n}\ \mathrm{d}\sigma \qquad (12)$$

where v is the derivative of a parametrization of $S(t)$ with respect to t at the point 0.

The latter theorem enables to prove Proposition 1:

Proof:

Let $\mathbf{x}_0$ be a configuration satisfying assumption 6. Let $\{B^i;\ i = 1 \dots N\}$ be neighborhoods as defined in assumption 6, and let $B := \prod_i B^i$.

As a consequence of assumption 6, RVCs associated with configurations in B are either empty or polyhedra of dimension 3 (that is, different from polygons, edges and vertices).

Let x_i be a seed of the configuration $\mathbf{x}_0$ with $i \in \{1 \dots N\}$, and let $V_i(\mathbf{x_0})$ be its associated cell.

By assumption 6, if $V_i(\mathbf{x}_0)$ is empty, then for all $\mathbf{x} \in B$, $V_i(\mathbf{x})$ is empty. So, the energy of $V_i(\mathbf{x})$ equals zero and is a C^1-function of seeds.

If $V_i(\mathbf{x}_0)$ is not empty, its $\mathcal{H}$-representation (as an intersection of half-spaces defined by the pair of an outward normal and a point) is defined by two kinds of half-spaces:

- half-spaces $\left((x_k - x_i),\ \dfrac{x_i + x_k}{2}\right)$ for all neighbors x_k of x_i in the Voronoi neighborhood of x_i;
- half-spaces of the form (N, b), for each boundary triangle t of Ω which is incident to V_i, where N is the outward unit normal of t and b is a point of t.

These half-spaces are continuous functions of seeds.

Let o be a point in $\mathring{V}_i(\mathbf{x}_0)$. By continuity of half-spaces, there is a neighborhood B' of $\mathbf{x}_0$ such that $B' \subset B$ and for all $\mathbf{x} \in B'$, $o \in \mathring{V}_i(\mathbf{x})$. In the sequel we assume B is actually such a B'.

Because V_i is convex, one can build S_i, the simplicial decomposition of V_i made of tetrahedra connecting o to the facets of V_i. By assumption 6, as the topology of cells is locally invariant, S_i is also a topological simplicial decomposition of $V_i(\mathbf{x})$ for any $\mathbf{x} \in B$. In the sequel, given a topological tetrahedron $T \in S_i$, we will denote by $T(\mathbf{x})$ its geometrical realization in $V(\mathbf{x})$.

It can be shown that vertices of $V_i(\mathbf{x})$ are C^1-functions of $\mathbf{x}$ (see [2], note that assumptions 3 and 4 are necessary). Let T be a tetrahedron of S_i. For all $\mathbf{x} \in B$, as $T(\mathbf{x}_0)$ and $T(\mathbf{x})$ are never degenerate, there is an invertible affine map $\psi_{\mathbf{x}}$ between $T(\mathbf{x}_0)$ and $T(\mathbf{x})$. Then, for any smooth function f on $T(\mathbf{x})$:

$$\int_{T(\mathbf{x})} f\ \mathrm{d}\lambda = \int_{T(\mathbf{x}_0)} (f\circ\psi_{\mathbf{x}})\cdot|J\psi_{\mathbf{x}}|\,\mathrm{d}\lambda. \qquad (13)$$

As Voronoi vertices are C^1-functions of seeds, coefficients of $\psi_{\mathbf{x}}$ depend continuously on $\mathbf{x}$ and integrand $(\mathbf{x}, y) \mapsto (f \circ \psi_{\mathbf{x}})(y) \cdot |J\psi_{\mathbf{x}}|$ is continuous. Since $T(\mathbf{x}_0)$ has finite measure and since $B \times T(\mathbf{x}_0)$ is a compact set, $\mathbf{x} \mapsto \int_{T(\mathbf{x})} f \, \mathrm{d}\lambda$ is continuous at $\mathbf{x}_0$. That last argument can be applied to each tetrahedron $T \in S_i$ and for each seed x_i, and this is why f_{CVT} is continuous at $\mathbf{x}_0$.

Let us now prove the gradient formula. As previously, i denotes a seed, T a tetrahedron in S_i, a fixed topological simplicial decomposition of V_i and $T(\mathbf{x})$ its geometrical realization in $V(\mathbf{x})$.

Given $k \in \{1, \ldots, N\}$ and $j \in \{1, 2, 3\}$, let us consider the (k, j)-th variable of configurations (in particular, let us set $t_0 = (\mathbf{x}_0)_{(k,j)}$). We now consider "freezing" all the variables of $\mathbf{x}_0$ but the (k, j)-th one. Let $\hat{\mathbf{x}}_t^{kj}$ be the configuration $\mathbf{x}_0$ where the variable (k, j) is replaced with $t \in \mathbb{R}$. We thus get an invertible affine map $\varphi_t = \psi_{\hat{\mathbf{x}}_t^{kj}}$ between $T(\mathbf{x}_0)$ and $T(\hat{\mathbf{x}}_t^{kj})$, whose coefficients depend continuously on t. As $(T(\hat{\mathbf{x}}_t^{kj}))_{t \in I}$ (for some interval I) is a collection of (non-degenerate) tetrahedra which are manifolds with corners of dimension 3, the REYNOLDS theorem can be applied, provided that inequality $(**)$ holds.

Let us consider the i-th seed $\hat{x}_i(t)$ of the configuration $\hat{\mathbf{x}}_t^{kj}$. Likewise, let x_i be the i-th seed of configuration $\mathbf{x}_0$. Let I be a close bounded interval which contains t_0, such that topologies of RVCs generated by $\hat{\mathbf{x}}_t^{kj}$ remain constant. Let f be the function

$$f : T(\mathbf{x}_0) \times I \to \mathbb{R} : f(y, t) = \|M[\hat{x}_i(t) - y]\|_p^p \, . \quad (14)$$

As stated before, for all $y \in T(\mathbf{x}_0)$, $t \mapsto \varphi_t(y)$ is a C^1-function, so the same applies to its Jacobian matrix with respect to y and to $|\det J\varphi_t|$. Since f is also a C^1-function, for all $y \in T(\mathbf{x}_0)$, the following function

$$\tilde{f} : (y, t) \mapsto f(\varphi_t(y), t) |\det J\varphi_t| \quad (15)$$

is a C^1-function of t.

For all $y \in T(\mathbf{x}_0)$ and for any θ in I, the derivative of the latter function with respect to t at the point (y, θ) is:

$$\frac{\partial}{\partial t}(|\det J\varphi_t|)(\theta) \cdot f(\varphi_\theta(y), \theta) \quad (16)$$

$$+ \; |\det J\varphi_\theta| \cdot \frac{\partial}{\partial t}(f(\varphi_t(y), t))(\theta) \quad (17)$$

which is a continuous function of (y, t): indeed, since φ_t is affine, $|\det J\varphi_\theta|$ and $\frac{\partial}{\partial t}(|\det J\varphi_t|)(\theta)$ do not depend on y and other factors are continuous with respect to (y, t).

As $T(\mathbf{x}_0)$ has finite measure, since $T(\mathbf{x}_0) \times I$ is a compact set and $(y, t) \mapsto \partial_t \tilde{f}(y, t)$ is continuous, the existence of a function g which satisfies $(**)$ is ensured (in the REYNOLDS theorem, replace I with any open interval I' such that $t_0 \in I' \subset I$).

By summing all formulas which come from the REYNOLDS theorem for all tetrahedra of S_i, the partial derivative of $\mathcal{E}_i$ with respect to the (k, j) variable at t_0 becomes:

$$\int_{V_i(\mathbf{x}_0)} \frac{\partial f}{\partial t}(y, t_0) \, \mathrm{d}y + \int_{\partial V_i(\mathbf{x}_0)} f(y, t_0) v \cdot n \, \mathrm{d}\sigma \quad (18)$$

where v is the derivative of a parametrization of facets of V_i with respect to t at the point t_0, and where n is the outward normal unit vector.

Let us distinguish three cases:

First case: $i = k$

The chain rule gives

$$\frac{\partial f}{\partial t}(y, t_0) = p \left(e_j^3\right)^\top M^\top \begin{pmatrix} (M_1[x_i - y])^{p-1} \\ (M_2[x_i - y])^{p-1} \\ (M_3[x_i - y])^{p-1} \end{pmatrix}. \quad (19)$$

Second case: $x_k \in \mathbb{V}(x_i)$

The derivative $\partial f / \partial t$ is null. Let Γ be a facet of V_i which is different from the facet $V_i \cap V_k$. Suppose that this facet Γ is shared by V_i and a RVC V_m for some $m \in [1, N] \backslash \{k\}$. If no Voronoi vertex of V_m depends on $\hat{x}_k$, then $v = 0$. If some Voronoi vertices of V_m do depend on $\hat{x}_k$, then $v \cdot n = 0$ because the affine hull of Γ only depends on $\hat{x}_i$ and $\hat{x}_m$.

Otherwise: The derivative of the energy of V_i is zero because $\partial f / \partial t$ and v are null.

Note that, in any case, if Γ is a boundary facet included in $\partial\Omega$, then $v = 0$. That is why surface integrals are expressed in terms of integrals on facets shared by x_i and its neighbors in the Delaunay triangulation.

By summing derivatives of all energies (with a special focus on orientations of surface integrals), we obtain:

$$\frac{\partial f_{\mathrm{CVT}}}{\partial x_{ij}}(\mathbf{x}_0) =$$

$$p \int_{V_i(\mathbf{x}_0)} \left(e_j^3\right)^\top M^\top \begin{pmatrix} (M_1[x_i - y])^{p-1} \\ (M_2[x_i - y])^{p-1} \\ (M_3[x_i - y])^{p-1} \end{pmatrix} \mathrm{d}y + \quad (***)$$

$$\sum_{\substack{\Gamma_{ik} \\ x_k \in \mathbb{V}(x_i)}} \int_{\Gamma_{ik}} \left(\left\| M[x_i - z] \right\|_p^p - \left\| M[x_k - z] \right\|_p^p \right) \left(\frac{\partial z}{\partial x_{ij}} \cdot n \right) \mathrm{d}\sigma$$

where z is a parametrization of $\Gamma_{ik} := V_i(\mathbf{x}_0) \cap V_k(\mathbf{x}_0)$ and n is the outward normal unit vector of Γ, that is, $n = (x_k - x_i)/\|x_k - x_i\|$ by virtue of the PErpendicular BIsector property of Voronoi diagrams.

Surface integrals can be expressed more explicitly. Let Γ be the facet shared by two RVCs $V_i(\mathbf{x}_0)$ and $V_k(\mathbf{x}_0)$. Let us define $c_{ik} = (x_i + x_k)/2$ (which belongs to the plane defined by Γ but not necessarily to Γ). Let us denote by S_Γ the collection of triangles defined by edges of Γ and c_{ik}. When $c_{ik} \notin \Gamma$, it is not a manifold triangulation of Γ, but for any function f_Γ defined on Γ (even if it means extending f_Γ by the value 0 where it is not defined):

$$\int_\Gamma f_\Gamma \, \mathrm{d}\sigma = \sum_{\tau \in S_\Gamma} \int_\tau f_\Gamma \, \mathrm{d}\sigma. \quad (20)$$

As a consequence, the second integral of equation $(***)$ can be decomposed as a sum of integrals over the triangles $\tau \in S_\Gamma$.

Let us now consider such a triangle $\tau := (v_1, v_2, c_{ik})$. If Δ denotes the standard 2-simplex, the map $z : \Delta \to \mathbb{R}^3$ defined by $z : (s, t) \mapsto c_{ik} + su_1 + tu_2$ with $u_a := v_a - c_{ik}, a \in \{1, 2\}$, is a parametrization of τ (if the orientation is not correct,

inverting u_1 and u_2 is enough). Normal n is a differentiable function of seeds and

$$\frac{\partial n}{\partial x_{ij}}(x_i, x_k) = \frac{-1}{\|x_k - x_i\|} e_j^3 + \frac{x_{kj} - x_{ij}}{\|x_k - x_i\|^2} n. \tag{21}$$

Hence, as n is the unit normal vector of Γ:

$$\frac{\partial z}{\partial x_{ij}} \cdot n = \partial_{x_{ij}}(z \cdot n) - z \cdot \partial_{x_{ij}} n = \frac{z_j - x_{ij}}{\|x_k - x_i\|}. \tag{22}$$

Using this equality in $(***)$, we get the final expression of ∇f_{CVT}.

Last, let us prove the continuity of the gradient at $\mathbf{x}_0$. We use the terminology of [6] again: "inner terms" of ∇f_{CVT} are integrals on Voronoi cells and "boundary terms" refer to surface integrals. Inner terms are still continuous functions of seeds (by arguments similar to the proof of the continuity of f_{CVT}).

For boundary terms, let us consider a facet Γ. Let $\mathcal{T}_\Gamma$ be a triangulation of Γ whose triangles only have Voronoi vertices as corners: again such triangles are functions of seeds. Let h be the function

$$\mathbb{R}^3 \times \mathbb{R}^{3N} : (y, z_1, \dots, z_N) \mapsto \left(\|M[z_i - y]\|_p^p - \|M[z_k - y]\|_p^p\right) \frac{y_j - z_{ij}}{\|z_k - z_i\|}. \tag{23}$$

Let us take $t(\mathbf{x}) = (v_1, v_2, v_3) \in \mathcal{T}_\Gamma$. Then the function H

$$H : \mathbf{x} \mapsto \int_{t(\mathbf{x})} h(\cdot, \mathbf{x}) \, \mathrm{d}\sigma \tag{24}$$

where $\mathbf{x}$ denotes the configuration $(z_1, \dots, z_N)$, can be also written

$$\mathbf{x} \mapsto 2 \int_\Delta (h \circ y) \, |t(\mathbf{x})| \, \mathrm{d}\lambda \tag{25}$$

where y is a parametrization of $t(\mathbf{x})$ defined on Δ. Since Δ has finite measure, since $\Delta \times B$ is compact and the map

$$\Delta \times B : ((s,t), \mathbf{x}) \mapsto 2|t(\mathbf{x})| \; h(y(s,t), z_1, \dots, z_N) \tag{26}$$

is continuous, H is continuous at $\mathbf{x}_0$.

So, the boundary terms of the gradient of f_{CVT} are continuous at $\mathbf{x}_0$, and so is ∇f_{CVT}. ∎

Proposition 2. *Proposition 1 still holds when assumptions 2 and 4 are replaced with assumption 5.*

Proof: Similar arguments can be used, but instead of considering cells V_i, convex decompositions of cells $V_i \cap \tau$ with $\tau \in \Omega$ must be used in the proof. The gradient expression $(*)$ is obtained by summing the gradients of all energies over all these pieces of RVCs. ∎

Now, we no longer suppose assumptions 2 and 4, which are replaced with assumption 5.

Proposition 3. *Proposition 2 still holds when M is piecewise constant or piecewise smooth on Ω. By "piecewise" constant or smooth, we respectively mean that, for each tetrahedron τ of Ω, M is supposed to be constant on τ or the restriction $M_{|\tau}$ can be extended into a smooth function defined over the closed tetrahedron $\overline{\tau}$.*

Proof: If M is piecewise constant on Ω, then for any tetrahedron τ of Ω, the restriction of M to a piece of an RVC $V_i \cap \tau$ is constant, and is *a fortiori* smooth. Since Proposition 1 is proved by considering each cell independently of the others, $(*)$ still holds. ∎

Last, we have the following conjecture:

Conjecture. *If M is only continuous, Proposition 1 still holds.*

Proof: Indeed, if M is continuous, it is a uniform limit of a sequence of smooth fields $(M_n)_n$. Let us denote by f_{CVT_n} the function f_{CVT} where M is replaced with M_n. We conjecture that $((f_{\text{CVT}_n})(\mathbf{x}_0))_n$ converges towards $(f_{\text{CVT}})(\mathbf{x}_0)$, and that $(t \mapsto (\partial x_{ij})(f_{\text{CVT}_n})(\hat{\mathbf{x}}_{ij}^t))_n$ uniformly converges towards the formula of Proposition 1 (on some open interval which contains x_{ij}, after having replaced some variables with other ones). From that, we can deduce that f_{CVT} has a derivative with respect to x_{ij} which is of the form expected. ∎

A. Expression of $\nabla f_{\text{CVT}}^{restrict}$

Let V be a set of seeds and let $\bar{V}$ be a subset of V. Let $\mathbf{x}$ and $\bar{\mathbf{x}}$ be respectively configurations associated with V and $\bar{V}$.

The partial derivative of $f_{\text{CVT}}^{\text{restrict}}$ along with the j-th coordinate of $\bar{x}_i \in \bar{V}$ at the configuration $\bar{\mathbf{x}}$ is then given by:

$$p_i(\bar{x}_i) \cdot p \int_{V_i(\mathbf{x})} \left(e_j^3\right)^\top M^\top \begin{pmatrix} (M_1[\bar{x}_i - y])^{p-1} \\ (M_2[\bar{x}_i - y])^{p-1} \\ (M_3[\bar{x}_i - y])^{p-1} \end{pmatrix} \mathrm{d}y \tag{27}$$

$$+ \left(\partial_j p_i\right)(\bar{x}_i) \cdot \mathcal{E}_i(\mathbf{x}) \tag{28}$$

$$+ \sum_{\substack{\Gamma_{ik} \\ x_k \in \mathbb{V}(\bar{x}_i)}} \int_{\Gamma_{ik}} \Big(p_i(\bar{x}_i) \big\|M[\bar{x}_i - z]\big\|_p^p - p_k(x_k) \big\|M[x_k - z]\big\|_p^p \Big) \frac{(z_j - \bar{x}_{ij})}{\|x_k - \bar{x}_i\|} \, \mathrm{d}\sigma \tag{29}$$

where z is a parametrization of the corresponding facet $\Gamma_{ik} := V_i(\mathbf{x}) \cap V_k(\mathbf{x})$ in surface integrals, and e_j^3 is the j-th vector of the canonical basis of $\mathbb{R}^3$.

References

[1] J. Droniou, "Finite Volume Schemes For Diffusion Equations: Introduction to and Review of Modern Methods," *Mathematical Models and Methods in Applied Sciences*, vol. 24, no. 8, pp. 1575–1619, 2014. [Online]. Available: https://hal.archives-ouvertes.fr/hal-00813613

[2] Levy, Bruno and Liu, Yang, "L_p Centroidal Voronoi Tessellation and its Applications," *ACM TRANSACTIONS ON GRAPHICS*, vol. 29, no. 4, JUL 2010.

[3] M. Iri, K. Murota, and T. Ohya, *A Fast Voronoi-Diagram Algorithm with Applications to Geographic Optimization*. , 01 2006, ch. , pp. 273–288.

[4] V. Nivoliers and B. Levy, "Approximating Functions on a Mesh with Restricted Voronoi Diagrams," *COMPUTER GRAPHICS FORUM*, vol. 32, no. 5, pp. 83–92, AUG 2013.

[5] F. de Goes, K. Breeden, V. Ostromoukhov, and M. Desbrun, "Blue Noise through Optimal Transport," *ACM TRANSACTIONS ON GRAPHICS*, vol. 31, no. 6, NOV 2012.

[6] V. Nivoliers, B. Levy, and C. Geuzaine, "Anisotropic and Feature Sensitive Triangular Remeshing using Normal Lifting," *JOURNAL OF COMPUTATIONAL AND APPLIED MATHEMATICS*, vol. 289, pp. 225–240, DEC 1 2015.

[7] R. Merland, "Génération de grilles de type volumes finis : adaptation à un modèle structural, pétrophysique et dynamique," Ph.D. dissertation, Université de Lorraine, Apr. 2013.

[8] R. Merland, B. Lévy, G. Caumon, and P. Collon-Drouaillet, "Building Centroidal Voronoi Tessellations for Flow Simulation in Reservoirs using Flow Information," in , vol. 1, 2011, p. 153 – 163.

[9] Z. Zhong, W. Wang, B. Lévy, J. Hua, and X. Guo, "Computing a High-Dimensional Euclidean Embedding from an Arbitrary Smooth Riemannian Metric," *ACM Trans. Graph.*, vol. 37, no. 4, jul 2018.

[10] M. Budninskiy, B. Liu, F. de Goes, Y. Tong, P. Alliez, and M. Desbrun, "Optimal Voronoi Tessellations with Hessian-Based Anisotropy," *ACM Trans. Graph.*, vol. 35, no. 6, dec 2016.

[11] Y. Xiao, Z. Chen, J. Cao, Y. J. Zhang, and C. Wang, "Optimal power diagrams via function approximation," *Computer-Aided Design*, vol. 102, pp. 52–60, 2018, proceeding of SPM 2018 Symposium.

[12] Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi Tessellations: Applications and Algorithms," *SIAM Review*, vol. 41, no. 4, p. 637 – 676, 1999.

[13] Y. Liu, W. Wang, B. Lévy, F. Sun, D.-M. Yan, L. Lu, and C. Yang, "On Centroidal Voronoi Tessellation – Energy Smoothness and Fast Computation," *ACM Transactions on Graphics*, vol. 28, no. 4, p. 1 – 17, 2009.

[14] J. Hateley, H. Wei, and L. Chen, "Fast Methods for Computing Centroidal Voronoi Tessellations," *Journal of Scientific Computing*, vol. 63, 04 2014.

[15] B. Lévy and N. Bonneel, "Variational Anisotropic Surface Meshing with Voronoi Parallel Linear Enumeration," in *Proceedings of the 21st International Meshing Roundtable*. Springer Berlin Heidelberg, 2012, pp. 349–366. [Online]. Available: https://hal.inria.fr/hal-00804558

[16] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer New York, NY, 2006.

[17] H. Edelsbrunner, *Geometry and Topology for Mesh Generation*, ser. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2001.

[18] D.-M. Yan, W. Wang, B. Levy, and Y. Liu, "Efficient Computation of 3D Clipped Voronoi Diagram," in *ADVANCES IN GEOMETRIC MODELING AND PROCESSING, PROCEEDINGS*, ser. Lecture Notes in Computer Science, B. Mourrain, S. Schaefer, and G. Xu, Eds., vol. 6130, 2010, pp. 269–282, 6th International Conference on Geometric Modeling and Processing (GMP 2010), Castro Urdiales, SPAIN, JUN 16-18, 2010.

[19] J. Cortes, S. Martinez, and F. Bullo, "Spatially-distributed Coverage Optimization and Control with Limited-range Interactions," *ESAIM-CONTROL OPTIMISATION AND CALCULUS OF VARIATIONS*, vol. 11, no. 4, pp. 691–719, 2005.

[20] D.-M. Yan, B. Levy, Y. Liu, F. Sun, and W. Wang, "Isotropic Remeshing with Fast and Exact Computation of Restricted Voronoi Diagram," *COMPUTER GRAPHICS FORUM*, vol. 28, no. 5, SI, pp. 1445–1454, JUL 2009, 7th Eurographics Symposium on Geometry Processing (SGP), Berlin, GERMANY, JUL 15-17, 2009.

[21] J. Jaśkowiec and N. Sukumar, "High-Order Cubature Rules for Tetrahedra," *International Journal for Numerical Methods in Engineering*, vol. 121, no. 11, pp. 2418–2436, 2020.

[22] M. A. Taylor, B. A. Wingate, and L. P. Bos, "Several New Quadrature Formulas for Polynomial Integration in the Triangle," 2005. [Online]. Available: https://arxiv.org/abs/math/0501496

[23] C. Geuzaine and J.-F. Remacle, "Gmsh: a Three-Dimensional Finite Element Mesh Generator with Built-In Pre- and Post-Processing Facilities," *International Journal for Numerical Methods in Engineering*, pp. 1309–1331, 2009.

[24] R. Merland, G. Caumon, B. Levy, and P. Collon, "Voronoi Grids Conformal to 3D Structural Features," *Computational Geosciences*, vol. 18, pp. 1–11, 08 2014.

[25] J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed., ser. Graduate Texts in Mathematics. Springer New York, NY, 2012.

[26] M. Reddiger and B. Poirier, "The Differentiation Lemma and the Reynolds Transport Theorem for Submanifolds with Corners," 2019. [Online]. Available: https://arxiv.org/abs/1906.03330