Anisotropic Fast-Marching on Cartesian Grids Using Lattice Basis Reduction

Jean-Marie Mirebeau

Anisotropic Fast-Marching on cartesian grids using Lattice Basis Reduction[*]

Jean-Marie Mirebeau[†]

February 7, 2014

Abstract

We introduce a modification of the Fast Marching algorithm, which solves the anisotropic eikonal equation associated to an arbitrary continuous Riemannian metric $\mathcal{M}$, on a two or three dimensional domain. The algorithm has a complexity $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$, where N is the discrete domain cardinality. The logarithmic dependency in the maximum anisotropy ratio $\kappa(\mathcal{M})$ of the Riemannian metric allows to handle extreme anisotropies for a limited numerical cost. We prove the consistence of the algorithm, and illustrate its efficiency by numerical experiments. The algorithm relies on the computation at each grid point z of a special system of coordinates: a *reduced* basis of the lattice $\mathbb{Z}^m$, with respect to the symmetric positive definite matrix $\mathcal{M}(z)$ encoding the desired anisotropy at this point.

Introduction

The anisotropic Eikonal equation, or static Hamilton-Jacobi equation, is a Partial Differential Equation (PDE) which describes an elementary front propagation model: the speed of the front depends only on the front position and orientation. This PDE is encountered in numerous applications, such as motion planning control problems [23], modeling of bio-medical phenomena [21], and image analysis [19]. It was also recently used in the context of medical image analysis [4] for extracting vessels in two dimensional projections or three dimensional scans of the human body, and for performing virtual endoscopies. This application requires to solve a highly anisotropic generalized eikonal equation with a high resolution on a cartesian grid, at a computational cost compatible with user interaction. It is one of our key motivations.

This paper is devoted to the construction and the study of a new algorithm, Fast Marching using Lattice Basis Reduction (FM-LBR), designed to solve the anisotropic eikonal equation associated to a given Riemannian metric $\mathcal{M}$, and able to handle large or even extreme anisotropies. The domain must be of dimension two or three, and discretized on a cartesian grid. The FM-LBR, as its name indicates, is a variant of the classical Fast Marching algorithm [23, 26], an efficient method for solving the eikonal equation when the metric is isotropic (proportional at each point to the identity matrix). Lattice Basis Reduction [18] is a concept from discrete mathematics, used in the FM-LBR to produce sparse causal stencils for the discretization of the eikonal equation; it allows to benefit in an optimal way of the interplay between the Riemannian geometric structure of the PDE, and the arithmetic structure of the discretization grid. A similar technique is used in [7] to construct sparse non-negative stencils for anisotropic diffusion.

[†]CNRS, University Paris Dauphine, UMR 7534, Laboratory CEREMADE, Paris, France.

In order to illustrate the specificity of our approach, we need to introduce some notation. Denote by S_m^+ the collection of $m \times m$ symmetric positive definite matrices, and associate to each $M \in S_m^+$ the norm $\|u\|_M := \sqrt{\langle u, Mu \rangle}$ on $\mathbb{R}^m$. Consider a bounded open domain $\Omega \subset \mathbb{R}^m$, equipped with a Riemannian metric $\mathcal{M} \in C^0(\overline{\Omega}, S_m^+)$. We address the anisotropic eikonal equation: find the unique viscosity [14] solution $D : \overline{\Omega} \to \mathbb{R}$ of

$$\begin{cases} \|\nabla D(z)\|_{\mathcal{M}(z)^{-1}} = 1 & \text{for almost every } z \in \Omega, \\ D = 0 & \text{on } \partial\Omega. \end{cases} \tag{1}$$

See the end of §1 for different boundary conditions, and algorithmic restrictions on the dimension m. Introduce the Riemannian length of a Lipschitz path $\gamma : [0, 1] \to \overline{\Omega}$:

$$\text{length}(\gamma) := \int_0^1 \|\gamma'(t)\|_{\mathcal{M}(\gamma(t))} dt, \tag{2}$$

and denote by $D(x, y)$ the length of the shorted path joining $x, y \in \overline{\Omega}$, also referred to as the Riemannian distance between these points. The PDE (1) admits an optimal control interpretation: $D(x) = \min\{D(x, y); y \in \partial\Omega\}$ is the minimal distance from $x \in \Omega$ to the boundary.

Consider a discrete set $Z \subset \mathbb{R}^m$, which in our case will be a cartesian grid. Discretizations of (1) take the form of a fixed point problem: find $d : Z \to \mathbb{R}$ such that

$$\begin{cases} d(z) = \Lambda(d, z) & \text{for all } z \in Z \cap \Omega, \\ d(z) = 0 & \text{for all } z \in Z \setminus \Omega. \end{cases} \tag{3}$$

This formulation involves the Hopf-Lax update operator $\Lambda(d, z)$ [23, 5, 11, 8], which mimics at the discrete level Belmann's optimality principle associated to the optimal control interpretation of (1). See Appendix B for details on this principle, the approximations underlying its discretization, and their accuracy. The definition of $\Lambda(d, x)$ involves a mesh (or *stencil*) $V(x)$ of a small neighborhood of $x \in Z \cap \Omega$, with vertices on Z, and reads:

$$\Lambda(d, x) := \min_{y \in \partial V(x)} \|x - y\|_{\mathcal{M}(x)} + I_{V(x)} d(y), \tag{4}$$

where I_V denotes piecewise linear interpolation on a mesh V. (In this paper, a mesh in $\mathbb{R}^m$ is a finite collection $\mathcal{T}$ of m-dimensional non-flat simplices, which is conforming in the sense that the intersection $T \cap T'$ of any $T, T' \in \mathcal{T}$ is the convex hull of their common vertices.)

Numerical solvers of the eikonal equation differ by (i) the construction of the stencils $V(z)$, $z \in \Omega$, and (ii) the approach used to solve the system (3), which is inspired by the algorithms of Bellmann-Ford or of Dijkstra used for computing distances in graphs, instead of continuous domains. The algorithm presented in this paper, FM-LBR, belongs to the category of Dijkstra inspired algorithms with static stencils, and among these is the first one to guarantee a uniform upper bound on the stencil cardinality, independently of the Riemannian metric $\mathcal{M}$. The anisotropy ratio $\kappa(M)$ of a matrix $M \in S_m^+$, and the maximum anisotropy $\kappa(\mathcal{M})$ of the Riemannian metric $\mathcal{M}$, are defined as follows:

$$\kappa(M) := \max_{\|u\| = \|v\| = 1} \frac{\|u\|_M}{\|v\|_M} = \sqrt{\|M\| \|M^{-1}\|}, \qquad \kappa(\mathcal{M}) := \max_{z \in \overline{\Omega}} \kappa(\mathcal{M}(z)). \tag{5}$$

We denote by $N := \#(Z \cap \Omega)$ the cardinality of the discrete domain.

- **Bellman-Ford inspired algorithms.** The discrete fixed point problem (3) is solved via Gauss-Seidel iteration: the replacement rule $d(z_k) \leftarrow \Lambda(d, z_k)$ is applied for $k =$

$0, 1, 2, \dots$ to a mutable map $\delta : Z \cap \Omega \to \mathbb{R}$, until a convergence criterion is met. In the fast sweeping methods, see [25] and references therein, the sequence of points $(z_k)_{k \geq 0}$ enumerates repeatedly the lines and the columns of $Z \cap \Omega$. Alternatively this sequence is obtained via a priority queue in the Adaptive Gauss-Seidel Iteration (AGSI) of Bornemann and Rasch [5]. The stencil $V(z)$ of a point $z \in Z \cap \Omega$ is usually the offset by z of a fixed stencil V given at the origin, such as those illustrated on Figure 1.

Fast sweeping methods have $\mathcal{O}(\lambda(\mathcal{M})N)$ complexity when the metric $\mathcal{M}$ is *isotropic* (proportional to the identity at each point), but this result does not extend to anisotropic Riemannian metrics, see [30] for the proof and the expression of $\lambda(\mathcal{M})$. The AGSI has complexity $\mathcal{O}(\mu(\mathcal{M})N^{1+\frac{1}{m}})$, for arbitrary anisotropic Riemannian metrics, where $\mu(\mathcal{M})$ is a non explicit constant which depends on global geometrical features of the metric [5]. The AGSI is a popular, simple, and quite efficient method, which is included for comparison in our numerical tests.

- **Dijkstra inspired algorithms.** The system (3) is solved in a single pass, non-iteratively, using an ordering of $Z \cap \Omega$ determined at run-time. This is possible provided the Hopf-Lax update operator satisfies the so-called "causality property", see Proposition 1.1, which can be ensured if the stencil $V(z)$ of each $z \in Z \cap \Omega$ satisfies some geometrical properties depending on $\mathcal{M}(z)$, see Definition 1.2. The different Dijkstra inspired methods are characterized by the construction of the stencils $V(z)$, in contrast with Bellman-Ford inspired methods which are characterized by the choice of the sequence $(z_k)_{k \geq 0}$. Solving the system (3) with a Dijkstra inspired algorithm has complexity $\mathcal{O}(\mu(\mathcal{M})N \ln N)$, where $\mu(\mathcal{M})$ is an upper bound on the cardinality of the stencils (the number of simplices they are built of).

  In the Ordered Upwind Method (OUM) of Sethian and Vladimirsky [23, 29], the stencils are constructed at run-time; their cardinality is bounded by $\mathcal{O}(\kappa(\mathcal{M})^m)$ and drops to $\mathcal{O}(\kappa(\mathcal{M})^{m-1})$ as $N \to \infty$. In contrast, the stencils are constructed during a preprocessing step and then static in the Monotone Acceptance Ordered Upwind Method (MAOUM) of Alton and Mitchell [2]; their cardinality is bounded by $\mathcal{O}(\kappa(\mathcal{M})^m)$. The FM-LBR introduced in the present work uses a similar approach, except that the stencils cardinality is $\mathcal{O}(1)$, fully independent of the Riemannian metric $\mathcal{M}$. The complexity estimates are thus $\mathcal{O}(\kappa(\mathcal{M})^m N \ln N)$ for the OUM and the MAOUM (asymptotically $\mathcal{O}(\kappa(\mathcal{M})^{m-1}N \ln N)$ for the OUM), and $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$ for our approach the FM-LBR, where the second term in the complexity accounts for the stencil construction.

The above mentioned algorithms are consistent for the anisotropic eikonal equation associated to an arbitrary continuous Riemannian metric $\mathcal{M} \in C^0(\overline{\Omega}, S_m^+)$, in the sense that the discrete output d_h of the algorithm executed on the grid $Z_h := h\mathbb{Z}^2$, of scale $h > 0$, converges to the viscosity solution D of the continuous problem (1) as $h \to 0$. Some more specialized variants of the fast marching algorithm are only consistent for a restricted set of metrics, but can be executed nonetheless with an arbitrary anisotropic metric $\mathcal{M}$; in that case the discrete system (3) may not be solved, and the numerical results are variable, see §3. For instance the original fast marching algorithm [26] is consistent if $\mathcal{M}(z)$ is proportional to the identity matrix for each $z \in \Omega$, and more generally if $\mathcal{M}(z)$ is a diagonal matrix. In addition to these cases of isotropy and axis-aligned anisotropy, some variants are also consistent if the metric anisotropy $\kappa(\mathcal{M})$ is smaller than a given bound κ_0, see [9] and Figure 1. Our numerical experiments §3 include for comparison one of these methods: Fast Marching using the 8 point stencil (FM-8, center left stencil on Figure 1), which is popular in applications [4] thanks to its short computation times and despite the lack of convergence guarantee for arbitrary metrics. Depending on the
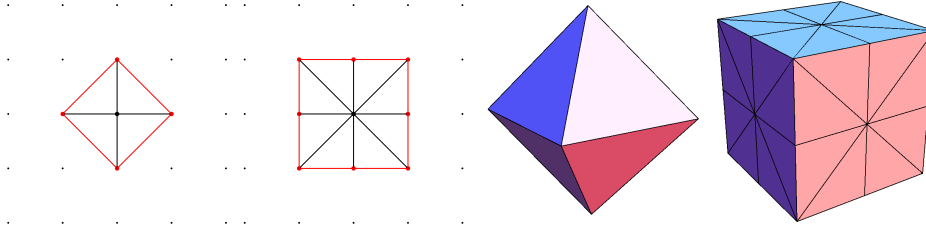
3

Figure 1: Some classical stencils used in the discretization of two dimensional (left) or three dimensional (right) eikonal equations. The meshes are M-acute, a property implying discrete causality, see Definition 1.2, for matrices M which are diagonal or of anisotropy ratio $\kappa(M)$ bounded by respectively 1, $1+\sqrt{2}$, 1, $(\sqrt{3}+1)/2$ (from left to right).

implementation [20, 23, 26], involving either a sorted list or a bucket sort, these methods have complexity $\mathcal{O}(N \ln N)$ or $\mathcal{O}(\Upsilon(\mathcal{M})N)$, where

$$\Upsilon(\mathcal{M}) := \sqrt{\max_{z \in \Omega} \|\mathcal{M}(z)\| \max_{z' \in \Omega} \|\mathcal{M}(z')^{-1}\|}.$$

In the applications for which our method is intended, one typically has $\ln(N) \lesssim \kappa(\mathcal{M}) \leq \Upsilon(\mathcal{M}) \ll N$, in such way that the complexity $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$ of the proposed method is comparable to $\mathcal{O}(N \ln N)$ and smaller than $\mathcal{O}(\Upsilon(\mathcal{M})N)$. In summary, the FM-LBR combines the universal consistency (i.e. for any Riemannian metric) of the AGSI, OUM and MAOUM, with a quasi-linear complexity just as the original Fast Marching algorithm.

Remark. *Each solver of the eikonal equation comes with a specific construction of the stencils $V(z)$, $z \in \Omega$. The highly efficient, but specialized, approach used in the FM-LBR limits its potential for generalization, see the end of §1. Static adaptive stencils also have a memory impact, see Remark 1.10. Since the Hopf-Lax update operator (4) depends on the stencils, the discrete solution d of (3) is scheme dependent, and so is its accuracy. See Appendix B for a heuristic accuracy analysis, and [15] for the case of constant metrics. Numerical experiments §3, on application inspired test cases, show that the FM-LBR accuracy is competitive.*

1 Sparse causal stencils for the anisotropic eikonal equation

Our main contribution is the construction of discretization stencils for anisotropic eikonal equations, which have a uniformly bounded cardinality and preserve a structural property of the PDE: causality, inherited from its interpretation as a deterministic control problem. Discrete causality, allowing to solve the fixed point system (3) in a single pass, is the following property.

Proposition 1.1 (Causality property, J.A. Sethian, A. Vladimirsky, Appendix of [23]). *Let $x \in \mathbb{R}^m$, and let V be a finite mesh of a neighborhood of x. Let $M \in S_m^+$ and let us assume that $\langle y - x, M(z - x) \rangle \geq 0$ for any vertices y, z of a common face of ∂V (acuteness condition). Consider a discrete map d defined on the vertices of V, and the optimization problem*

$$\Lambda := \min_{y \in \partial V} \|x - y\|_M + \mathrm{I}_V \, \mathrm{d}(y),$$

where I denotes piecewise linear interpolation. Then the minimum defining Λ is attained on a k-face $[y_1, \cdots, y_k]$ of ∂V, with $k \leq m$, such that $\Lambda > \mathrm{d}(y_i)$ for all $1 \leq i \leq k$.

A mesh V satisfying the geometric acuteness condition of Proposition 1.1, is called a causal stencil, at the point x, and with respect to M.

Consider the directed graph G, associated to the fixed point problem (3), where for all $x, y \in Z$ we place an arrow $x \to y$ iff the minimum defining $\mathrm{d}(x) = \Lambda(\mathrm{d}, x)$ is attained on a face of the stencil $\partial V(x)$ containing y. The causality property, states that the presence of an arrow $x \to y$ implies $\mathrm{d}(x) > \mathrm{d}(y)$; in particular the graph G has no cycles, hence the system of equations (3) does not feature any dependency loop. The Fast-Marching algorithm [26] traces back theses dependencies, determining at run-time an ordering $(x_i)_{i=1}^N$, of the discrete domain $Z \cap \Omega$, such that the distances $(\mathrm{d}(x_i))_{i=1}^N$ are increasing. We reproduce this method for completeness, see Algorithm 1, but refer to its original introduction [26], or its use with alternative static adaptive stencils in the MAOUM [1], for the proof that it solves the discrete system (3).

Algorithm 1 The Fast Marching algorithm, with static stencils adapted to a metric.

Input: The values $\mathcal{M}(z)$ of a Riemannian metric, for all $z \in Z \cap \Omega$.
Construct causal stencils $V(z)$, with respect to $\mathcal{M}(z)$, at all points $z \in Z \cap \Omega$.
Construct the reversed stencils, defined by $V[y] := \{z \in Z \cap \Omega;\ y$ is a vertex of $V(z)\}$.
Initialize a (mutable) table $\mathrm{d}: Z \to \mathbb{R}$, to $+\infty$ on $Z \cap \Omega$, and to 0 elsewhere.
Initialize a (mutable) boolean table $b: Z \to \{trial, accepted\}$ with $b(y) \leftarrow trial$ iff $V[y] \neq \emptyset$.
While there remains a *trial* point (i.e. $y \in Z$ such that $b(y) = trial$) **do**
    Denote by y a *trial* point which minimizes d, and set $b(y) \leftarrow accepted$.
    For all $x \in V[y]$, set $\mathrm{d}(x) \leftarrow \min\{\mathrm{d}(x), \Lambda(\mathrm{d}, x; b, y)\}$.
Output: The map $\mathrm{d}: Z \to \mathbb{R}$.

We denoted by $\Lambda(\mathrm{d}, x; b, y)$ the modification of the Hopf-Lax update operator (4) in which the minimum is only taken over faces (of any dimension) of $\partial V(x)$ which vertices (i) contain y, and (ii) are all *accepted*. Regarding the FM-LBR complexity $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$, we refer for details to the classical analysis in [26, 23, 1] and simply point out that (i) each FM-LBR causal stencil costs $\mathcal{O}(\ln \kappa(\mathcal{M}))$ to construct, (ii) maintaining a list of $\Omega \cap Z$, sorted by increasing values of the mutable map d, costs $\mathcal{O}(\ln N)$ for each modification of a single value of d, with a proper heap sort implementation, and (iii) the optimization problem defining the Hopf-Lax update (4), or its variant $\Lambda(\mathrm{d}, x; b, y)$, has an explicit solution: the minimum associated to each face of $\partial V(x)$ is the root of a simple univariate quadratic polynomial, see Appendix of [23]. Memory usage is discussed in detail in Remark 1.10.

As announced, we limit our attention to PDE discretizations on cartesian grids, of the form

$$Z = hR(\xi + \mathbb{Z}^m), \tag{6}$$

where $h > 0$ is a scaling parameter, R a rotation, and $\xi \in \mathbb{R}^m$ an offset (in (6) and (7) we abusively apply geometric transformations not only to points, but also to sets of points and meshes). The use of an unbounded grid (6) is a mathematical artifact aimed to simplify the exposition; only points of Z close to Ω (precisely: such that $V[y] \neq \emptyset$) play an active role in Algorithm 1. We may construct a causal stencil at $x \in Z$ with respect to $M \in S_m^+$, by suitably scaling, translating and rotating an $R^{\mathrm{T}} M R$-acute mesh $\mathcal{T}$, defined below:

$$V = x + h\,R\,\mathcal{T}. \tag{7}$$

Definition 1.2. *An M-acute mesh, where $M \in S_m^+$, is an m-dimensional mesh $\mathcal{T}$ covering a neighborhood of the origin, and such that the vertices $v_0, \cdots, v_m$ of any simplex $T \in \mathcal{T}$ satisfy (i) $v_0 = 0$, (ii) $(v_1, \cdots, v_m)$ form a basis of $\mathbb{Z}^m$, and (iii) $\langle v_i, Mv_j \rangle \geq 0$ for all $1 \leq i < j \leq m$.*

5

The condition $|\det(v_1, \cdots, v_m)| = 1$ ensures that 0 (resp. x) is the only grid point in the interior of the domain covered by $\mathcal{T}$ (resp. V), hence that information does not "fly over" some grid points when solving (3). Applying the next proposition to the classical stencils of Figure 1, we obtain that they are causal for Riemannian metrics of limited anisotropy.

Proposition 1.3. *Let $\mathcal{T}$ be an m-dimensional mesh satisfying the requirements of Definition 1.2, except (iii) (which does not make sense without a given matrix $M \in S_m^+$). Let*

$$\kappa(\mathcal{T}) := \sqrt{\frac{1 + \gamma(\mathcal{T})}{1 - \gamma(\mathcal{T})}}, \quad where \quad \gamma(\mathcal{T}) := \min_{T, (u,v)} \frac{\langle u, v \rangle}{\|u\| \|v\|}, \tag{8}$$

and where the minimum in $\gamma(\mathcal{T})$ is taken among all non-zero vertices u, v of a common simplex $T \in \mathcal{T}$. The mesh $\mathcal{T}$ is M-acute for any $M \in S_m^+$ such that $\kappa(M) \leq \kappa(\mathcal{T})$.

Proof. Let u, v be two non-zero vertices of a common simplex $T \in \mathcal{T}$, and let $M \in S_m^+$. Let $u' := u/\|u\|$ and let $v' := v/\|v\|$. By construction one has

$$\|u' + v'\|^2 = 2(1 + \langle u', v' \rangle) \geq 2(1 + \gamma(\mathcal{T})), \qquad \|u' - v'\|^2 = 2(1 - \langle u', v' \rangle) \leq 2(1 - \gamma(\mathcal{T})).$$

Let us assume for contradiction that $\langle u, Mv \rangle < 0$, which implies that $\|u' + v'\|_M < \|u' - v'\|_M$. Observing that

$$\kappa(M)^2 = \|M\| \|M^{-1}\| \geq \frac{\|u' - v'\|_M^2}{\|u' - v'\|^2} \frac{\|u' + v'\|^2}{\|u' + v'\|_M^2} > \frac{1 + \gamma(\mathcal{T})}{1 - \gamma(\mathcal{T})},$$

we obtain that $\kappa(M) > \kappa(\mathcal{T})$, which concludes the proof of this proposition. $\square$

Our construction of M-acute meshes relies on special coordinate systems in the grid $\mathbb{Z}^m$, adapted to the anisotropic geometry encoded by $M \in S_m^+$.

Definition 1.4 (Bases and superbases). *A basis of $\mathbb{Z}^m$ is an m-plet $(e_1, \cdots, e_m) \in (\mathbb{Z}^m)^m$ such that $|\det(e_1, \cdots, e_m)| = 1$. A superbase of $\mathbb{Z}^m$ is an $m + 1$-plet $(e_0, \cdots, e_m)$ such that $e_0 + \cdots + e_m = 0$, and $(e_1, \cdots, e_m)$ is a basis of $\mathbb{Z}^m$.*

Definition 1.5. *A superbase $(e_0, \cdots, e_m)$ of $\mathbb{Z}^m$ is said to be M-obtuse iff $\langle e_i, M e_j \rangle \leq 0$ for all $0 \leq i < j \leq m$.*

There exists for each $M \in S_m^+$, $m \in \{2, 3\}$ at least one M-obtuse superbase [6]. The construction of 2D and 3D M-obtuse superbases relies on lattice basis reduction algorithms [12, 18] (hence the name of our numerical scheme) and has cost $\mathcal{O}(\ln \kappa(M))$, see §2. We arrive at the main contribution of this paper: the FM-LBR stencils, which are causal and of bounded cardinality.

Proposition 1.6 (The FM-LBR acute meshes). *Let $M \in S_m^+$, and let $(e_0, \cdots, e_m)$ be an M-obtuse superbase of $\mathbb{Z}^m$ (if one exists). An M-acute mesh $\mathcal{T}(M)$ is obtained by collecting the simplices of vertices $(\sum_{i=0}^{k-1} b_i)_{k=0}^m$ associated to all $(m+1)!$ permutations $(b_i)_{i=0}^m$ of $(e_i)_{i=0}^m$. It has $2^{m+1} - 1$ vertices.*

Proof. Proof of the properties of the simplices. Let $(b_i)_{i=0}^m$ be a permutation of $(e_i)_{i=0}^m$, and let $v_k := \sum_{i=0}^{k-1} b_i$, for all $0 \leq k \leq m$. Clearly (i) $v_0 = 0$, and (ii) $|\det(v_1, \cdots, v_m)| =$
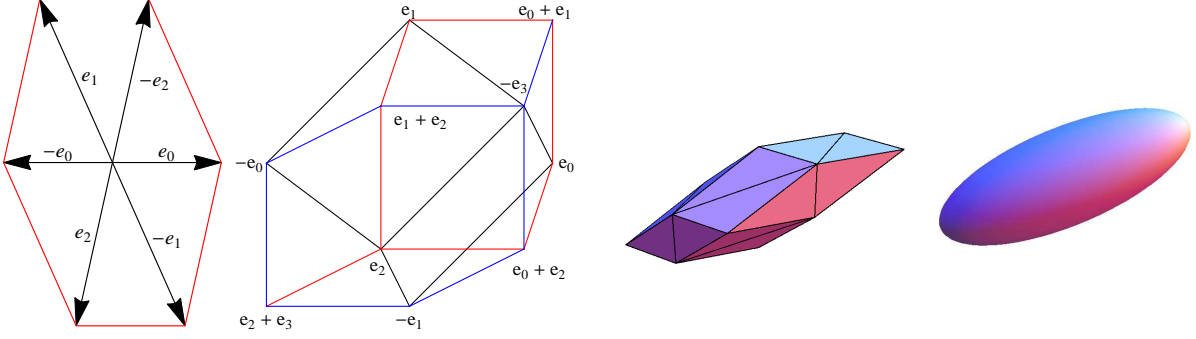
Figure 2: Connectivity of the 2D (left) and 3D (center left, interior edges omitted, missing boundary edges are symmetric w.r.t. the origin) FM-LBR meshes $\mathcal{T}(M)$, see Proposition 1.6 (note that $-e_i = \sum_{j \neq i} e_j$). Mesh $\mathcal{T}(M)$ (center right) associated to $M \in S_3^+$ of eigenvalues $3^2, 1, 1$, and eigenvector $(3, 1, 2)$ for the first eigenvalue. Right: unit ball for the norm $\|\cdot\|_M$.

$|\det(b_1, \cdots, b_m)| = |\det(e_1, \cdots, e_m)| = 1$. Acuteness condition (iii): for any $1 \leq k < l \leq m$ one has $v_l = -\sum_{j=l}^m b_j$ by Definition 1.4, hence by Definition 1.5

$$\langle v_k, M v_l \rangle = -\sum_{0 \leq i < k} \sum_{l \leq j \leq m} \langle e_i, M e_j \rangle \geq 0.$$

Proof that $\mathcal{T}(M)$ is a conforming mesh covering a neighborhood of the origin. Consider the $m+1$-dimensional Kuhn simplices $T_\sigma := \{(\lambda_0, \cdots, \lambda_m); 0 \leq \lambda_{\sigma(0)} \leq \lambda_{\sigma(1)} \leq \cdots \leq \lambda_{\sigma(m)} \leq 1\}$, associated to all permutations σ of $\{0, \cdots, m\}$, which form a conforming mesh $\mathcal{T}_0$ of $[0, 1]^{m+1}$. The linear map $A : \mathbb{R}^{m+1} \to \mathbb{R}^m$, defined by $A(\lambda_0, \cdots, \lambda_m) := -\sum_{i=0}^m \lambda_i e_i$, has a kernel generated by $v_1 := (1, \cdots, 1)$. It sends $[0, 1]^{m+1}$ onto a neighborhood of $0 \in \mathbb{R}^m$, and transforms $\mathcal{T}_0$ into an m-dimensional conforming mesh $\mathcal{T}_M$ of this neighborhood, collapsing onto 0 the common edge $[0, v_1]$ of the simplices T_σ. Other vertices of $[0, 1]^{m+1}$ have pairwise distinct images by A, since their difference is not proportional to v_1; hence $\mathcal{T}_M$ has $2^{m+1} - 1$ vertices (one less than $\mathcal{T}_0$). Noticing the identity $-\sum_{i=0}^m \lambda_i e_i = \sum_{k=1}^m (\lambda_{\sigma(k)} - \lambda_{\sigma(k-1)}) \sum_{i=0}^{k-1} e_{\sigma(k)}$, we find that the image of T_σ by A is the simplex of $\mathcal{T}(M)$ associated to the permuted superbase $(e_{\sigma(i)})_{i=0}^m$, hence $\mathcal{T}_M = \mathcal{T}(M)$, which concludes the proof. $\square$

Obtuse superbases are similarly used in [7] to produce sparse non-negative stencils for Anisotropic Diffusion (AD-LBR scheme), with different results: 3D stencils have 12 non-zero vertices in [7], and 14 here. This illustrates the versatility of this concept, which can be used to design stencils satisfying various geometric properties: an acuteness condition for the FM-LBR (implying the scheme causality), and the non-negative decomposition of a tensor for the AD-LBR (guaranteeing the scheme monotonicity).

Definition 1.7. *A family of meshes $(\mathcal{T}(x))_{x \in \Omega}$, defined for all points of an open domain $\Omega \subset \mathbb{R}^m$ is admissible iff there exists two constants $0 < r$ and $R < \infty$ such that: for each $x \in \Omega$*

- *The mesh $\mathcal{T}(x)$ covers neighborhood of 0, and its vertices belong to $\mathbb{Z}^m$.*

- *(Boundedness) The vertices e of $\mathcal{T}(x)$ satisfy $\|e\| \leq R$.*

- *(Stability) There exists a basis $\mathcal{B}(x)$ of $\mathbb{Z}^m$ which elements, and their opposites, are vertices of $\mathcal{T}(y)$ for all $y \in \Omega$ such that $\|x - y\| < r$.*
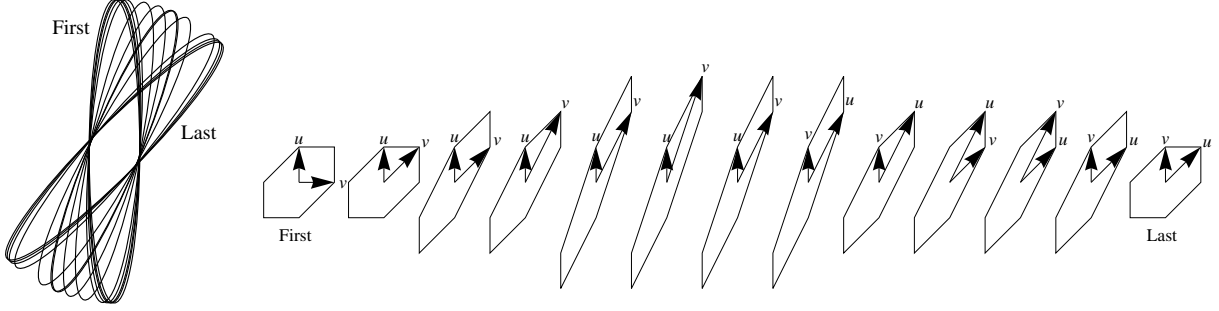
Figure 3: The unit sphere $\{u \in \mathbb{R}^2; \|u\|_M = 1\}$, an M-reduced basis (u, v), and the boundary of the FM-LBR mesh $\mathcal{T}(M)$, for some $M \in S_2^+$ of anisotropy ratio $\kappa(M) = 6$, and eigenvector $(\cos\theta, \sin\theta)$, $\theta \in [\pi/4, \pi/2]$, associated to the small eigenvalue.

A fortunate but non-trivial fact is that the FM-LBR family of meshes $(\mathcal{T}(\mathcal{M}(x)))_{x \in \Omega}$ is admissible, see Proposition 1.8 and Figures 3 and 4. This property implies the FM-LBR consistence, using Proposition 1.9 which is a minor extension of the convergence results in [26, 5].

Proposition 1.8. *Let $\Omega \subset \mathbb{R}^m$ be open and bounded, and let $\mathcal{M} \in C^0(\overline{\Omega}, S_m^+)$. If $m \leq 3$ then the FM-LBR family of meshes $(\mathcal{T}(\mathcal{M}(x)))_{x \in \Omega}$ is admissible (with Boundedness constant $R = C_m \kappa(\mathcal{M})$, $C_2 = 2$, $C_3 = 4$). More generally, if $m \leq 4$ then any family of meshes $(\mathcal{T}(x))_{x \in \Omega}$ such that $\mathcal{T}(x)$ is $\mathcal{M}(x)$-acute for all $x \in \Omega$, satisfies the (Stability) property.*

Proposition 1.9. *Let $\Omega \subset \mathbb{R}^m$ be an open bounded set, equipped with a Riemannian metric $\mathcal{M} \in C^0(\overline{\Omega}, S_m^+)$, and an admissible family $(\mathcal{T}(z))_{z \in \Omega}$ of meshes. For all $h > 0$ let $Z_h := h\mathbb{Z}^2$, and for all $z \in Z_h \cap \Omega$ consider the stencil $V_h(z) := z + h\mathcal{T}(z)$. Then the solutions $d_h : Z_h \to \mathbb{R} \cup \{+\infty\}$ of the discrete system (3) converge uniformly as $h \to 0$ to the viscosity solution D of the eikonal equation (1):*

$$\lim_{h \to 0} \max_{z \in Z_h \cap \Omega} |d_h(z) - D(z)| = 0. \tag{9}$$

To sum up the FM-LBR strengths, this algorithm is universally consistent, has a competitive accuracy w.r.t. alternative methods, and its computational cost is almost unaffected by the Riemannian metric anisotropy. For fairness we discuss below the potential downsides of our original stencil construction, which limits the potential for generalization (efficiency is as often at the cost of specialization), and impacts memory usage.

- (Finsler metrics) The FM-LBR only applies to the anisotropic eikonal equation associated to a Riemannian metric, while other methods such as the AGSI, OUM, MAOUM [5, 23, 2] can handle more general Finsler metrics. Indeed the structure $\| \cdot \|_{\mathcal{M}(z)}$ of the local norm at z associated to a Riemannian metric is required in the FM-LBR stencil construction Proposition 1.6, which involves an $\mathcal{M}(z)$-obtuse superbase. Finsler metrics are in contrast defined by arbitrary asymmetric norms $| \cdot |_z$, depending continuously on z.

  Constructing causal static stencils is an active subject of research in the case of Finsler metrics. A characterization obtained in [27], was used in [16] to develop the FM-ASR (Fast Marching using Anisotropic Stencil Refinement), which is close in spirit and in efficiency to the FM-LBR, but has a different application range, since it handles Finsler metrics on two dimensional domains. It was also observed in [22, 23, 1] that the canonical stencil

8

(Figure 1, left) is causal for Finsler metrics featuring only axis-aligned anisotropy, in such way that the original fast marching algorithm can be applied.

- (Domain discretization) The FM-LBR requires a cartesian grid discretization. In contrast an important research effort [3, 5, 10, 13, 23] has been devoted to the more difficult setting of meshed domains, with an unstructured set of vertices. These methods are natural candidates for the computation of distances on a manifold $\mathcal{S} \subset \mathbb{R}^m$, while the FM-LBR would require a (collection of) local chart(s) $\varphi : \Omega \to \mathcal{S}$ equipped with the metric $\mathcal{M}(z) := \nabla\varphi(z)^{\mathrm{T}}\nabla\varphi(z)$. The FM-LBR heavily relies on the cartesian grid arithmetic structure, through the concept of obtuse superbases.

- (Boundary conditions) The null boundary conditions chosen in the eikonal equation (1) can be replaced with Dirichlet data $D_0 : \partial\Omega \to \mathbb{R}$, of 1-Lipschitz regularity [14] with respect to the Riemannian distance $D(\cdot, \cdot)$, see (2). In that event, Algorithm 1 requires to extend the boundary data D_0 to a ghost layer, containing all grid points $y \in Z \setminus \Omega$ such that the reverse stencil $V[y]$ is non-empty. The FM-LBR uses large stencils, of euclidean radius $\mathcal{O}(h\kappa(\mathcal{M}))$ on a grid (6) of scale h, which complicates this extension in contrast with e.g. the AGSI [5] of stencil radius $\mathcal{O}(h)$.

  Outflow boundary conditions, on a portion $\Gamma \subsetneq \partial\Omega$ of the domain's boundary, are natural in applications, see §3. They are implemented by excluding in the definition (4) of $\Lambda(\mathrm{d}, x)$, faces of $\partial V(x)$ containing an exterior vertex $y \in Z \setminus \Omega$ close to Γ. If x lies in a corner of Ω, and if the stencil $V(x)$ is strongly anisotropic, then it may happen that all the vertices of $\partial V(x)$ lie *outside* Ω, so that the solution of (3) satisfies $\mathrm{d}(x) = +\infty$. (In our experiments §3, this happened in the square domain's corners when test case 1 was rotated by an angle $\theta \in [0.56, 0.61]$ radians. These four infinite values were rejected when estimating numerical errors.) Despite these minor inconveniences, the FM-LBR behaves remarkably well in our numerical experiments, §3 and Appendix B, with these more general boundary conditions.

- (Dimension) The FM-LBR causal stencils construction, see Proposition 1.6, is limited to domains of dimension 2 and 3, because some matrices $M \in S_4^+$ do not admit any M-obtuse superbase [18]. An alternative construction of M-acute meshes of uniformly bounded cardinality is proposed in [15], for all $M \in S_4^+$; this cardinality is not small unfortunately, with 768 simplices.

  The FM-LBR however extends in a straightforward manner to Riemannian metrics $\mathcal{M}$ having a block diagonal structure, with blocks of size 1, 2 or 3, using the following construction. For $i \in \{1, 2\}$ let m_i be a positive integer, let $M_i \in S_{m_i}^+$, and let $\mathcal{T}_i$ be an M_i-acute mesh. Let $m := m_1 + m_2$ and let $M \in S_m^+$ be the matrix of diagonal blocks M_1, M_2. An M-acute mesh $\mathcal{T}$ is obtained by collecting the m-dimensional simplices of vertices $(0, 0)$, $(u_1, 0), \cdots, (u_{m_1}, 0)$, $(0, v_1), \cdots, (0, v_{m_2})$, where the simplices of vertices $0, u_1, \cdots, u_{m_1}$ and $0, v_1, \cdots, v_{m_2}$ belong to $\mathcal{T}_1$ and $\mathcal{T}_2$ respectively.

  Block diagonal metrics are not uncommon in the context of medical imaging [4], as they inherit the cartesian product structure of the fast marching domain: $\Omega = \Omega_0 \times \Omega_1$, where Ω_0 is a physical domain of dimension ≤ 3, and Ω_1 is an abstract parameter domain of dimension ≤ 2.

Outline. Further insight on the FM-LBR is given in §2, including the proof of Propositions 1.8 and 1.9. Numerical experiments are presented in §3. In addition, Appendix A describes
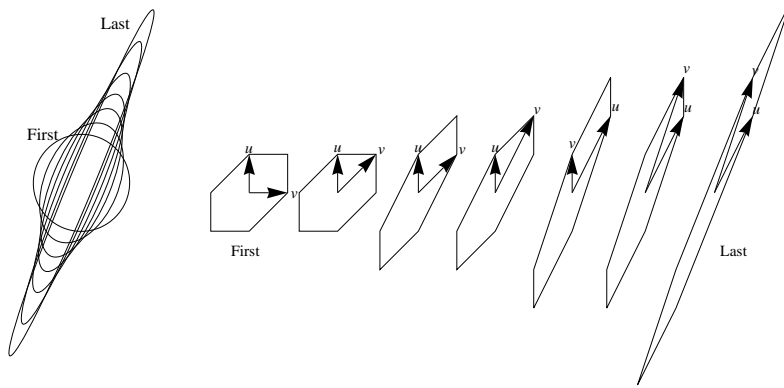
Figure 4: The unit sphere $\{u \in \mathbb{R}^2; \|u\|_M = 1\}$, an M-reduced basis (u, v), and the boundary of the FM-LBR mesh $\mathcal{T}(M)$, for some $M \in S_2^+$ of anisotropy ratio $\kappa(M)$ ranging from 1 to 15, and eigenvector $(\cos(3\pi/8), \sin(3\pi/8))$ associated to the small eigenvalue.

a robust minimal path extraction method for the FM-LBR and other Dijkstra inspired solvers of the eikonal equation. A heuristic analysis of the FM-LBR accuracy, and a last numerical experiment, appear in Appendix B.

Remark 1.10 (Memory requirements). *The memory requirements of numerical methods for the eikonal equation, such as the AGSI, the OUM and the FM-LBR, are dominated by (I) storing the discrete solution* d *and the Riemannian metric* $\mathcal{M}$*, sampled on the discrete domain* $\Omega \cap Z$*, and (II) storing the graph structure underlying the numerical scheme. Point (I) requires two tables of* N *and* $Nm(m+1)/2$ *reals, which may be represented in 32 bit (single precision) or 64 bit (double precision) format. The storage cost for the metric can be avoided if it has an analytical expression.*

Point (II) can be avoided for the AGSI and the OUM when these methods are executed on a mesh with a trivial periodic structure, which is the case in our experiments. For the FM-LBR, Point (II) amounts to storing the non-empty reverse stencils $V[y]$*, at all points of* $Y := \{y \in Z; V[y] \neq \emptyset\}$*, since the direct ones can be recomputed individually on demand for a minor cost. The set* Y *the union of* $\Omega \cap Z$ *and of a thin boundary layer (in our experiments,* $Y = \Omega \cap Z$ *due to the use of outflow boundary conditions). The chosen data structure uses two tables: one of vectors (the differences* $x - y$*, for* $x \in V[y]$*,* $y \in Y$*, enumerated consecutively), and one of* $\#(Y) \approx N$ *integers (the start, for each* $y \in Y$*, of the description of* $V[y]$ *in the previous table). We represent integers in 32bit format, and vector components in 8bit format, since these are small integers by construction.*

Summing up, we find that the memory requirements of the FM-LBR are larger than those of the AGSI or the OUM (on a grid), by a factor ranging from 2 *(metric and solution stored in double precision), to* 9 *(analytical metric, solution stored in single precision), through* 3 *(metric and solution stored in single precision), in two dimensions. Respectively, in three dimensions, from* 2.6 *to* 24*, through* 4.3*.*

2   Analysis of the FM-LBR

We introduce in §2.1 the concepts of Lattice Basis Reduction. They are used in §2.2 to estimate the construction cost of the FM-LBR meshes $\mathcal{T}(M)$, and to prove their admissibility in the sense of Definition 1.7, as announced in Proposition 1.8. We finally prove in §2.3 the announced

convergence result Proposition 1.9.

2.1 Introduction to Lattice Basis Reduction

We briefly introduce the framework of (low dimensional) Lattice Basis Reduction, used in the next subsection to construct and demonstrate the properties of M-acute meshes. See [18] and references therein for more details on this rich theory, from which we use only one result: Theorem 2.2 stated below. We denote by $b_1\mathbb{Z} + \cdots + b_k\mathbb{Z}$ the sub-lattice of $\mathbb{Z}^m$ generated by $b_1, \cdots, b_k \in \mathbb{Z}^m$:

$$b_1\mathbb{Z} + \cdots + b_k\mathbb{Z} := \{b_1 z_1 + \cdots + b_k z_k; \ z_1, \cdots, z_k \in \mathbb{Z}\}.$$

If $k = 0$ then the above sum equals $\{0\}$ by convention.

Definition 2.1. *Let $1 \leq m \leq 4$ and let $M \in S_m^+$. A basis $(b_1, \cdots, b_m)$ of $\mathbb{Z}^m$ is said M-reduced iff for all $1 \leq k \leq m$:*

$$b_k \in \operatorname{argmin}\{\|z\|_M; \ z \in \mathbb{Z}^m \setminus (b_1\mathbb{Z} + \cdots + b_{k-1}\mathbb{Z})\}. \tag{10}$$

If $M \in S_m^+$ is a diagonal matrix of coefficients $(\lambda_1, \cdots, \lambda_m)$, and if $0 < \lambda_{\sigma(1)} \leq \cdots \leq \lambda_{\sigma(m)}$ for some permutation σ, then the permutation $(e_{\sigma(i)})_{i=1}^m$ of the canonical basis of $\mathbb{Z}^m$ is an M-reduced basis of $\mathbb{Z}^m$. See Figures 2, 3, 4, for some examples of M-reduced bases associated to non-diagonal matrices M. In dimension $m \geq 5$ there exists matrices $M \in S_m^+$ such that no basis of $\mathbb{Z}^m$ satisfies the relations (10), see [18] (these relations state that $\|b_i\|_M$ equals the i-th Minkowski's minimum $\lambda_i(M)$). Minkowski's reduction [18] is the natural generalization of Definition 2.1 in dimension $m \geq 5$.

Theorem 2.2 (Nguyen, Stelhé, 2009)**.** *There exists an algorithm which, given a matrix $M \in S_m^+$ as input, $1 \leq m \leq 4$, produces an M-reduced basis of $\mathbb{Z}^m$ and has the numerical cost $\mathcal{O}(1 + \ln \kappa(M))$.*

Proof. The proof is contained in [18], and we only point out here the precise reference within the paper and the slight differences in notations. The algorithm described in [18] takes as input a basis $(b_1, \cdots, b_m)$ (here: the canonical basis of $\mathbb{R}^m$) of a lattice L (here: $\mathbb{Z}^m$), and its Gram matrix with respect to some scalar product (here: the Gram matrix is M). The algorithm outputs a greedy reduced basis of the lattice L, a notion which coincides with Minkowski's reduction if $m \leq 4$ (Lemma 4.3.2 in [18]), which itself coincides with Definition 2.1 if $m \leq 4$.

The main loop of the iterative algorithm is executed at most the following number of times (Theorem 6.0.5 in [18]):

$$\mathcal{O}\left(1 + \ln \max_{1 \leq i \leq m} \|b_i\|_M - \ln \min_{u \in L} \|u\|_M\right),$$

hence $\mathcal{O}(1 + \ln \|M\|^{\frac{1}{2}} - \ln \|M^{-1}\|^{-\frac{1}{2}}) = \mathcal{O}(1 + \ln \kappa(M))$ times in our setting. The complexity of each of these iterations is dominated by a closest vector search, described in Theorem 5.0.4 in [18], which consists of the inversion of a $k \times k$ Gram matrix, where $1 \leq k \leq m - 1$, and an exhaustive search among $\mathcal{O}(1)$ candidate vectors. In terms of elementary operations $(+, -, \times, /)$ among reals, each iteration of this algorithm thus has cost $\mathcal{O}(1)$, and the overall cost is the number of iterations $\mathcal{O}(1 + \ln \kappa(M))$.

Note that an important part of the discussion in [18] is devoted to the special case where the vectors $(b_1, \cdots, b_m)$ have *large integer* coefficients, the Gram matrix is computed with respect to the standard euclidean scalar product, and the complexity of an elementary operation

11

$(+, -, \times, /)$ among integers is not $\mathcal{O}(1)$ but depends on the size of these integers. This more subtle notion of complexity, named bit complexity, is not relevant in our setting. $\square$

The two dimensional version of the algorithms mentioned in Theorem 2.2 dates back to Lagrange [12], and mimicks the search for the greatest common divisor of two integers. This algorithm uses only a pair (u, v) of (mutable) variables in $\mathbb{Z}^2$, initialized as the canonical basis of $\mathbb{R}^2$. The pair (u, v) becomes an M-reduced basis at the end of the following loop, which takes at most $\mathcal{O}(\ln \kappa(M))$ iterations. Round denotes rounding to a closest integer.

$$\textbf{Do } (u, v) \leftarrow (v, \ u - \text{Round}(\langle u, Mv \rangle / \|v\|_M^2)\, v),$$
$$\textbf{While } \|u\|_M > \|v\|_M.$$

We end this subsection with a basic estimate of the norms and scalar products of the elements of M-reduced bases.

Proposition 2.3. *Let* $1 \leq m \leq 4$, *let* $M \in S_m^+$ *and let* $(b_1, \cdots, b_m)$ *be an M-reduced basis of* $\mathbb{Z}^m$. *Then for any* $1 \leq i \leq m$

$$\|b_i\| \leq \kappa(M), \quad and \quad \|b_i\|_M \leq \kappa(M)\|b_1\|_M. \tag{11}$$

For any integer combination $z \in b_1 \mathbb{Z} + \cdots + b_{i-1} \mathbb{Z} + b_{i+1} \mathbb{Z} + \cdots + b_m \mathbb{Z}$, *of the basis elements distinct from* b_i, *one has*

$$2|\langle b_i, Mz \rangle| \leq \|z\|_M^2. \tag{12}$$

Proof. Proof of 11. Let $(e_j)_{j=1}^m$ denote the canonical basis of $\mathbb{R}^m$. By a dimensionality argument there exists $1 \leq j \leq m$ such that $e_j \notin b_1 \mathbb{Z} + \cdots + \cdots b_{i-1} \mathbb{Z}$. By construction (10) we obtain $\|b_i\|_M \leq \|e_j\|_M \leq \|M\|^{\frac{1}{2}}$ since $\|e_j\| = 1$, hence $\|b_i\| \leq \kappa(M)$ as announced in (11). Observing that $\|b_1\|_M \geq \|M^{-1}\|^{-\frac{1}{2}}$ since $\|b_1\| \geq 1$, and recalling that $\|b_i\|_M \leq \|M\|^{\frac{1}{2}}$, we obtain the second announced estimate.

Proof of (12). Remark that $b_i + z \notin b_1 \mathbb{Z} + \cdots + b_{i-1} \mathbb{Z}$, since otherwise the basis element b_i would be a linear combination of $b_1, \cdots, b_{i-1}, b_{i+1}, \cdots, b_m$. Definition 2.1 thus implies:

$$\|b_i\|_M^2 \leq \|b_i + z\|_M^2 = \|b_i\|_M^2 + 2\langle b_i, Mz \rangle + \|z\|_M^2,$$

hence $-2\langle b_i, Mz \rangle \leq \|z\|_M^2$. Likewise $2\langle b_i, Mz \rangle \leq \|z\|_M^2$, which concludes the proof. $\square$

2.2 Properties of M-acute meshes

The FM-LBR stencil construction has numerical cost $\mathcal{O}(1 + \ln \kappa(M))$.

The construction of an FM-LBR mesh $\mathcal{T}(M)$, $M \in S_m^+$, has unit cost (for fixed m) given an M-obtuse superbase, see Proposition 1.6. We give below a unit cost construction of an M-obtuse superbase given an M-reduced basis, in dimension $m \in \{2, 3\}$, which by Theorem 2.2 can itself be obtained at cost $\mathcal{O}(1 + \ln \kappa(M))$.

Dimension 2. Let (b_1, b_2) be an M-reduced basis. Up to replacing b_2 with its opposite, we may assume that $\langle b_1, Mb_2 \rangle \leq 0$. Using (12) we obtain $2|\langle b_1, Mb_2 \rangle| \leq \|b_1\|_M^2$, hence $\langle Mb_1, -b_1 - b_2 \rangle \leq 0$ and likewise $\langle Mb_2, -b_1 - b_2 \rangle \leq 0$, so that $(b_1, b_2, -b_1 - b_2)$ is an M-obtuse superbase.

Dimension 3: we reproduce without proof the construction of [7]. Let (b_1, b_2, b_3) be the elements of an M-reduced basis, permuted and signed so that $|\langle b_1, Mb_2 \rangle| \leq \min\{-\langle b_1, Mb_3 \rangle, -\langle b_2, Mb_3 \rangle\}$. An M-obtuse basis is given by $(b_1, b_2, b_3, -b_1 - b_2 - b_3)$ if $\langle b_1, Mb_2 \rangle \leq 0$, and $(-b_1, b_2, b_1 + b_3, -b_2 - b_3)$ otherwise.

Radius of the FM-LBR stencils.

Assume that an FM-LBR acute mesh $\mathcal{T}(M)$ is built from an M-obtuse superbase obtained as in the previous paragraph[1] from an M-reduced basis $(b_i)_{i=1}^m$. Then one easily checks that its vertices have the form $e = \sum_{i=1}^m \varepsilon_i b_i$, where $\varepsilon_i \in \{-1, 0, 1\}$ and $(b_i)_{i=1}^m$ is the used M-reduced basis. Hence the vertices norm $\|e\| \leq \sum_{i=1}^m \|b_i\| \leq m\kappa(M)$ obeys the announced bound.

An M-reduced basis of $\mathbb{Z}^m$ contains by definition *small* vectors with respect to the norm $\|\cdot\|_M$: the smallest linearly independent ones with integer coordinates. As a result, the FM-LBR meshes $\mathcal{T}(M)$ have a small radius with respect to the norm $\|\cdot\|_M$. In constrast, these meshes can be large from the euclidean perspective. This has consequences on the FM-LBR accuracy, see Appendix B.

Stability of M-acute meshes.

Consider the following distance on the set S_m^+ of symmetric positive definite matrices: for all $M, N \in S_m^+$

$$\text{dist}(M, N) := \sup_{u \neq 0} |\ln \|u\|_M - \ln \|u\|_N| . \tag{13}$$

This distance allows to compare the norms of vectors multiplicatively, in contrast with the classical operator norm which is tailored for additive comparisons. Indeed denoting $\alpha := \text{dist}(M, N)$ and $\beta := \|M^{\frac{1}{2}} - N^{\frac{1}{2}}\|$, one has for all $u \in \mathbb{R}^2$ such that $\|u\| = 1$

$$e^{-\alpha} \leq \|u\|_M / \|u\|_N \leq e^\alpha, \quad \text{and} \quad -\beta \leq \|u\|_M - \|u\|_N \leq \beta.$$

The next lemma establishes a lower bound on the $\|\cdot\|_M$ norm of points with integer coordinates outside of an N-reduced mesh, when the matrices $M, N \in S_m^+$ are close enough.

Lemma 2.4. *Let $M, N \in S_m^+$, with $m \leq 4$. Let $u_1, \cdots, u_m$ be an arbitrary M-reduced basis of $\mathbb{Z}^m$, and let $\mathcal{T}$ be an N-reduced mesh. Consider a point $z \in \mathbb{Z}^m$ which is not a vertex of $\mathcal{T}$. Then there exists $1 \leq l \leq m$ such that*

$$z \in u_1\mathbb{Z} + \cdots + u_l\mathbb{Z} \quad and \quad \|z\|_M^2 \, e^{4\,\text{dist}(M,N)} \geq \|u_l\|_M^2 + \|u_1\|_M^2.$$

Proof. Since $\mathcal{T}$ covers a neighborhood of the origin, there exists a simplex $T \in \mathcal{T}$ and a real $\lambda > 0$ such that $\lambda z \in T$. Denoting by $v_1, \cdots, v_m$ the non-zero vertices of T, there exists therefore non-negative reals $\alpha_1, \cdots, \alpha_m \in \mathbb{R}_+$ such that $z = \alpha_1 v_1 + \cdots + \alpha_m v_m$. Since $(v_1, \cdots, v_m)$ form a basis of $\mathbb{Z}^m$, these coefficients are integers.

Up to reordering the vertices $v_1, \cdots, v_m$, we may assume that $\alpha_1, \cdots, \alpha_k$ are positive, and that $\alpha_{k+1}, \cdots, \alpha_m$ are zero, for some $1 \leq k \leq m$. Let l be the smallest integer such that $z \in u_1\mathbb{Z} + \cdots + u_l\mathbb{Z}$. There exists $1 \leq i \leq k$ such that $v_i \notin u_1\mathbb{Z} + \cdots + u_{l-1}\mathbb{Z}$, hence $\|v_i\|_M \geq \|u_l\|_M$; other vertices satisfy $\|v_j\|_M \geq \|u_1\|_M$, $1 \leq j \leq k$, since they are non-zero and have integer coordinates. Therefore, denoting $\delta := \text{dist}(M, N)$

$$e^{4\delta}\|z\|_M^2 \geq e^{2\delta}\|z\|_N^2 = e^{2\delta}\left(\sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_N^2 + 2 \sum_{1 \leq i < j \leq k} \alpha_i \alpha_j \langle v_i, N v_j \rangle \right)$$

$$\geq e^{2\delta} \sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_N^2 \geq \sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_M^2 \geq \|u_l\|_M^2 + \left(\sum_{1 \leq i \leq k} \alpha_i^2 - 1 \right) \|u_1\|_M^2.$$

Observing that $\alpha_1^2 + \cdots + \alpha_k^2 \geq 2$, since z is not a vertex of T, we conclude the proof. $\square$

---

[1] This implementation is used in our numerical experiments. See Corollary 2.6 for a proof with no assumption.

We prove in the next corollary that the vertices of an M-acute mesh contain the elements of an N-reduced basis, and their opposites, when the matrices M, N are sufficiently close. This result, and the compactness of $\overline{\Omega}$, immediately implies the second point of Proposition 1.8.

Corollary 2.5. *Let $M, N \in S_m^+$, with $m \leq 4$, be such that*

$$\text{dist}(M, N) < \ln(1 + \kappa(M)^{-2})/4. \tag{14}$$

Let $(b_1, \cdots, b_m)$ be an M-reduced basis of $\mathbb{Z}^m$, and let $\mathcal{T}$ be an N-reduced mesh. Then $b_1, \cdots, b_m$ and $-b_1, \cdots, -b_m$ are vertices of $\mathcal{T}$.

Proof. Let $1 \leq l \leq m$. We have $b_l \notin b_1 \mathbb{Z} + \cdots + b_{l-1}\mathbb{Z}$, since $(b_i)_{i=1}^m$ is a basis, and Proposition 2.3 implies

$$\|b_l\|_M^2 e^{4\,\text{dist}(M,N)} < \|b_l\|_M^2 + \kappa(M)^{-2}\|b_l\|_M^2 \leq \|b_l\|_M^2 + \|b_1\|_M^2.$$

Hence b_l is a vertex of $\mathcal{T}$, and likewise $-b_l$, by Lemma 2.4. $\qquad\square$

We finally estimate the radius of the FM-LBR meshes $\mathcal{T}(M)$, see Proposition 1.6, in terms of the condition number of the matrix $M \in S_m^+$. This concludes the proof of Proposition 1.8.

Corollary 2.6. *Let $M \in S_m^+$, with $m \in \{2, 3\}$. Then any vertex e of $\mathcal{T}(M)$ satisfies $\|e\| \leq C_m \kappa(M)$, with $C_2 := 2$ and $C_3 := 4$.*

Proof. Given m linearly independent vertices $(b_i)_{i=1}^m$ of $\mathcal{T}(M)$, one can express any other vertex under the form $e = \sum_{i=1}^m \alpha_i b_i$. A simple check by exhaustive enumeration shows that $\sum_{i=1}^m |\alpha_i| \leq C_m$ (these coefficients are independent of M), hence $\|e\| \leq C_m \max\{\|b_i\|; 1 \leq i \leq m\}$. Applying Corollary 2.5 with $M = N$, we find that the vertices of $\mathcal{T}(M)$ contain an M-reduced basis $(b_i)_{i=1}^m$, which by Proposition 2.3 satisfies $\|b_i\| \leq \kappa(M)$ for all $1 \leq i \leq m$. This concludes the proof. $\qquad\square$

2.3 Convergence of the FM-LBR

We prove in this section the uniform convergence of the discrete system (3) solutions towards the anisotropic eikonal PDE (1) solution, under the assumptions of Proposition 1.9 and with its notations. Following the steps of [5], we begin with a discrete Lipschitz regularity estimate for the maps $d_h : Z_h \to \mathbb{R}$.

Lemma 2.7. *There exists constants $h_0 > 0$ and $C_0 < \infty$ such that for all $0 < h \leq h_0$ and all $x, y \in Z_h$ one has*

$$|d_h(x) - d_h(y)| \leq C_0 \|x - y\|. \tag{15}$$

Proof. We prove below that $|d_h(x) - d_h(y)| \leq C_1 h$ when $\|x - y\| = h$, in other words when x and y are neighbors on the grid $Z_h := h\mathbb{Z}^m$. This immediately implies $|d_h(x) - d_h(y)| \leq C_1 \|x - y\|_1$, where $\|(\lambda_1, \cdots, \lambda_m)\|_1 := \sum_{i=1}^m |\lambda_i|$, hence also (15) with $C_0 := C_1 \sqrt{m}$.

If $x \notin \Omega$ and $y \notin \Omega$, then $d_h(x) = d_h(y) = 0$, and the result is proved. Up to exchanging x and y, we may therefore assume that $x \in \Omega$. Let $B := \mathcal{B}(x)$ be the basis corresponding to the property (Stability) of the admissible family of meshes $(\mathcal{T}(x))_{x \in \Omega}$, see Definition 1.7. Each element e of B is a vertex of $\mathcal{T}(x)$, hence satisfies $\|e\| \leq R$, see Definition 1.7 (Boundedness). We abusively regard B as an $m \times m$ matrix which columns are the basis elements, and observe that $\|B\| \leq R\sqrt{m}$ and $|\det(B)| = 1$. Therefore $\|B^{-1}\| \leq \|B\|^{m-1}/|\det B| \leq C_2 := (R\sqrt{m})^{m-1}$.

Since we assumed $\|x - y\| = h$, we have $y = x + \varepsilon h e_j$, for some $\varepsilon \in \{-1, 1\}$, $1 \leq j \leq m$, and where $(e_j)_{j=1}^m$ denotes the canonical basis of $\mathbb{R}^m$. Hence denoting by $(b_i)_{i=1}^m$ the elements

of the basis B, and by $(\alpha_{ij})_{i,j=1}^m$ the coefficients of B^{-1}, we obtain $y = x + \varepsilon h \sum_{i=1}^m \alpha_{ij} b_i$. Let $s := \sum_{i=1}^m |\alpha_{ij}|$, and let $(x_k)_{k=0}^s$ be a finite sequence of points of Z_h such that $x_0 := x$, $x_s := y$, and $x_{k+1} - x_k \in \{\pm h b_i; 1 \le i \le m\}$ for all $0 \le k < s$.

By Cauchy-Schwartz's inequality one has $s/\sqrt{m} \le (\sum_{i=1}^m \alpha_{ij}^2)^{\frac{1}{2}} \le \|B^{-1}\| \le C_2$, hence $\|x_k - x\| \le hRC_2\sqrt{m}$ for any $0 \le k \le s$. We choose the upper grid scale bound h_0 so that this constant smaller than the radius r involved in property (Stability) of Definition 1.7. If $x_k \in \Omega$, then by property (Stability) a grid point x_l, with $|k - l| \le 1$ and $0 \le l \le s$, is a vertex of the stencil $V_h(x)$. Hence inserting x_l in (4) we obtain $\mathrm{d}_h(x_k) = \Lambda_h(\mathrm{d}_h, x_k) \le \|x_k - x_l\|_{\mathcal{M}(x)} + \mathrm{d}_h(x_l) \le hRC_3 + \mathrm{d}_h(x_l)$, with $C_3 := \max\{\|\mathcal{M}(z)\|^{\frac{1}{2}}; z \in \overline{\Omega}\}$. If $x_k \notin \Omega$ then $\mathrm{d}_h(x_k) = 0$, hence obviously $\mathrm{d}_h(x_k) \le \mathrm{d}_h(x_l)$. Exchanging the roles of of k and l we obtain $|\mathrm{d}_h(x_k) - \mathrm{d}_h(x_l)| \le hRC_3$, hence by the triangular inequality $|\mathrm{d}_h(x) - \mathrm{d}_h(y)| \le hRC_3 s \le hRC_3 C_2 \sqrt{m}$, which concludes the proof. $\qquad \square$

The rest of the proof is only sketched, since it amounts to a minor adaptation of Theorem 11 in [5]. The only difference lies in the presence, in [5], of a canonical interpolation operator on Ω (piecewise linear interpolation on a prescribed mesh). Denote by D_h the bilinear interpolant[2] of d_h on the grid Z_h, and observe that $|\mathrm{D}_h(x) - \mathrm{D}_h(y)| \le KC_0\|x - y\|$, for all $x, y \in \mathbb{R}^2$ and all $0 < h \le h_0$, where K is an absolute constant depending only on the interpolation scheme, and C_0, h_0, are the constants of Lemma 2.7. Note also that $\mathrm{supp}(\mathrm{D}_h) \subset \{z + e; z \in \overline{\Omega}, \|e\| \le h\sqrt{m}\}$, hence by Lipschitz regularity D_h is bounded uniformly independently of h, and therefore by Arzelà-Ascoli's theorem the family $(\mathrm{D}_h)_{0 < h \le h_0}$ is pre-compact. Considering an arbitrary converging sub-sequence $\mathrm{D}_{h(n)}$, with $h(n) \to 0$ as $n \to \infty$, one observes that the limit is supported on $\overline{\Omega}$, and applying the arguments of Theorem 11 in [5] that it is a viscosity solution of the eikonal PDE (1). Uniqueness of such a solution [14] implies the pointwise convergence $\mathrm{D}_h(x) \to \mathrm{D}(x)$, as $h \to 0$, for all $x \in \overline{\Omega}$. Finally the announced uniform convergence (9) follows from the uniform KC_0-Lipschitz regularity of D_h, for all $0 < h \le h_0$.

3   Numerical experiments

We compare numerically the FM-LBR with two popular solvers (AGSI, FM-8) of the eikonal equation, which enjoy a reputation of simplicity and efficiency in applications, and with the recent and closely related MAOUM. The Adaptive Gauss Seidel Iteration[3] (AGSI) [5] produces numerical approximations which are guaranteed to converge towards the solution of the continuous anisotropic eikonal equation as one refines the computation grid[4], for an arbitrary continuous Riemannian metric $\mathcal{M}$. Fast Marching using the 8 point stencil (FM-8, stencil illustrated on Figure 1, center left) does not offer this convergence guarantee, but has a quasi-linear complexity $\mathcal{O}(N \ln N)$, in contrast with the super-linear complexity $\mathcal{O}(\mu(\mathcal{M})N^{1+\frac{1}{m}})$ of the AGSI. Fast Marching using Lattice Basis Reduction[5] (FM-LBR) aims to offer the best of both worlds: a convergence guarantee, and fast computation times[6].

---

[2] Other interpolation schemes could be used, such as piecewise linear interpolation on a trivial periodic mesh, provided one can control the Lipschitz regularity constant and the support of the interpolated function.

[3] As suggested in [5], the stopping criterion tolerance for the AGSI iterations is set to 10^{-8}.

[4] The grid is triangulated with a trivial periodic mesh, for the AGSI and the MAOUM. As a result the AGSI uses a 6 point stencil.

[5] The FM-LBR stencil $V(z)$, at each grid point $z \in \Omega \cap Z$, is built (7) from the $\mathcal{M}(z)$-reduced mesh $\mathcal{T}(\mathcal{M}(z))$ of Proposition 1.6, except if the matrix $\mathcal{M}(z)$ is detected to be exactly diagonal. In that case we use the standard 4 vertices neighborhood in 2D (resp. 6 vertices in 3D), which is an $\mathcal{M}(z)$-reduced mesh, see Figure 1 (left and center right). This modification has little impact on accuracy or CPU time, but avoids to pointlessly break symmetry.

[6] Note that the FM-LBR memory requirement is higher than that of the AGSI and FM-8, see Remark 1.10.
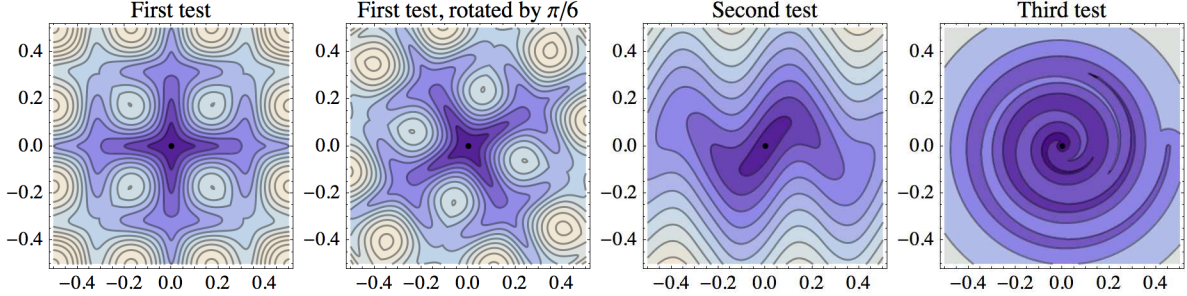
Figure 5: Level lines of the solutions of the two dimensional test cases.

We also implemented the Monotone Acceptance Ordered Upwind Method (MAOUM) [2], a Dijkstra inspired method using static stencils, like the FM-LBR. The difference between these two methods is that the MAOUM stencil[7] $V(z)$ at a grid point $z \in \Omega \cap Z$, is isotropic, only depends on the anisotropy ratio $\kappa(\mathcal{M}(z))$, and its boundary has cardinality $\mathcal{O}(\kappa(\mathcal{M}(z))^{m-1})$; in contrast the FM-LBR stencil is anisotropic, aligned with the ellipse defined by $\mathcal{M}(z)$, and of cardinality $\mathcal{O}(1)$. The MAOUM stencils were precomputed and stored in a look-up table, resulting in a complexity $\mathcal{O}(\kappa(\mathcal{M})N \ln N)$ for this algorithm in 2D.

We consider application inspired test cases, which violate some of the simplifying assumptions used in our convergence analysis Proposition 1.9: they involve outflow boundary conditions, non-trivial Dirichlet boundary conditions in Appendix B, and discontinuous Riemannian metrics in cases 3 and 4. Their exact technical description is given in Remark 3.1.

The first test is a distance computation on a parametrized surface, considered in [27]. As shown on Figure 6, the FM-LBR is the fastest in terms of CPU time[8], but is less accurate than the AGSI or the FM-8. Rotating this test case by the angle $\theta = \pi/6$, and conducting the same experiment, shows a different story: the numerical errors of the AGSI and the FM-8 increase by a factor larger than 5, while the FM-LBR, unaffected, is now the most accurate method, see Figure 6. The FM-LBR cuts L^∞ and L^1 numerical errors by 40% in comparison with the AGSI and the MAOUM, and CPU time by 85%, while the FM-8 produces even larger errors.

Figure 6 shows that the FM-LBR offers the best accuracy for more grid orientations θ than its alternatives. The maximal error and averaged error with respect to θ are also in favor of the FM-LBR. The strong dependence of the AGSI and the FM-8 accuracy on the test case orientation is puzzling, and contrasts with the more consistent behavior of the MAOUM and the FM-LBR. The author's heuristic and personal interpretation of this phenomenon, which is open to debate and put into question by a reviewer, is that this test case is for $\theta = 0$ dominated by (close to) axis-aligned anisotropy. The AGSI and the FM-8, which are based on small and fixed stencils, benefit from this special configuration; the FM-8 also works well for $\theta = \pi/4$, because its stencil includes the four diagonals.

The second benchmark, discussed in [28, 23], is inspired by seismic imaging. There is no bias here towards axis-aligned anisotropy. As shown on the table Figure 6, the FM-LBR takes

---

[7]The stencils for the MAOUM are here built using the ComputeUpdateSet stencil construction routine described and used in all the numerical experiments of [2]. The paper [2] also outlines sufficient conditions (called δ-NGA or DRB) for anisotropic stencils to be causal, but no explicit anisotropic stencil construction.

[8]All timings obtained on a 2.4Ghz Core 2 Duo, using a single core. Timings of the FM-LBR include the stencil construction, which typically accounts for 25%.
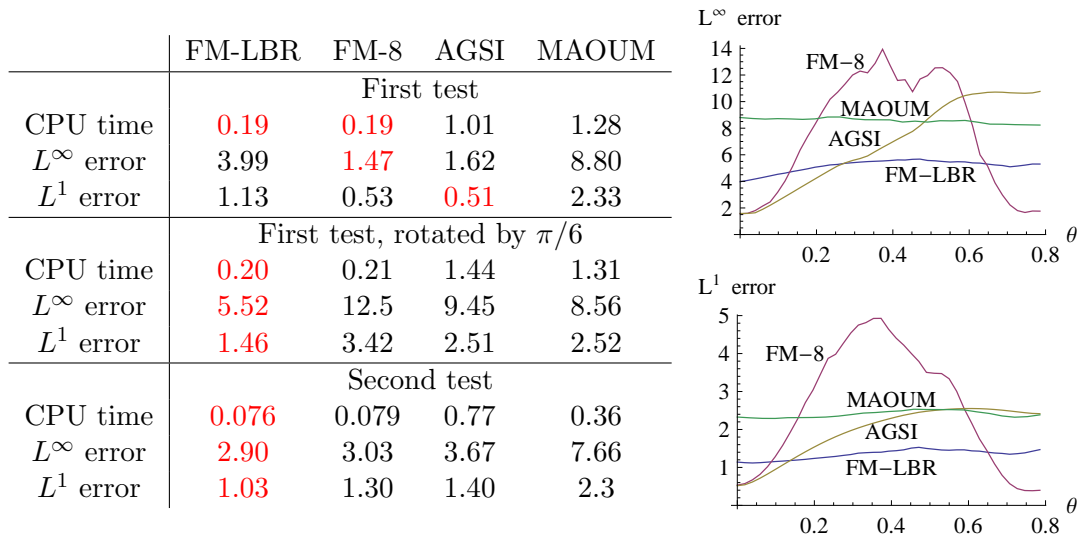
	FM-LBR	FM-8	AGSI	MAOUM
	First test			
CPU time	0.19	0.19	1.01	1.28
L^∞ error	3.99	1.47	1.62	8.80
L^1 error	1.13	0.53	0.51	2.33
	First test, rotated by $\pi/6$			
CPU time	0.20	0.21	1.44	1.31
L^∞ error	5.52	12.5	9.45	8.56
L^1 error	1.46	3.42	2.51	2.52
	Second test			
CPU time	0.076	0.079	0.77	0.36
L^∞ error	2.90	3.03	3.67	7.66
L^1 error	1.03	1.30	1.40	2.3

Figure 6: Tables of CPU time in seconds, L^∞ error an averaged L^1 error (left). Accuracy, in the first test rotated by an angle $\theta \in [0, \pi/4]$ (this interval is enough, since the dependence in θ is $\pi/2$-periodic and even). In average over theta, CPU times are $0.21s$, $0.20s$, $1.37s$, $1.31s$, L^∞ errors 5.16, 7.64, 6.86, 8.57 and averaged L^1 errors 1.34, 2.58, 1.95, 2.40 for the FM-LBR, FM-8, AGSI and MAOUM respectively. *All errors are multiplied by* 100, *for better readability.*
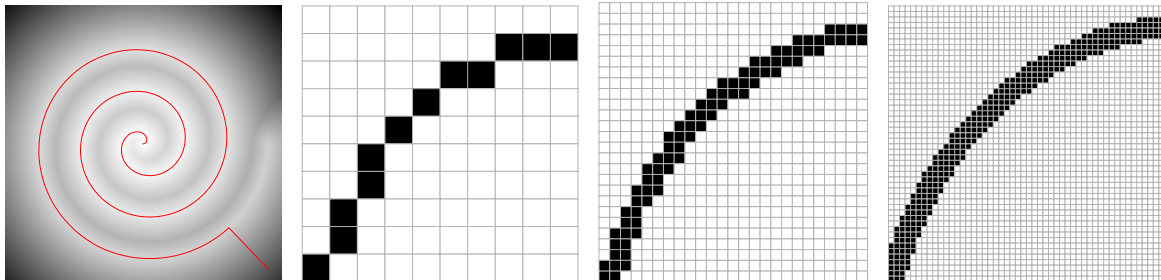

Figure 7: Reference solution for the third test case (left). The Riemannian metric $\mathcal{M}$ is anisotropic only on a thin band along a spiraling curve, wide of a few grid points. Detail at resolutions $n \times n$, where n equals 200 (center left), 500 (center right) and 1000 (right).
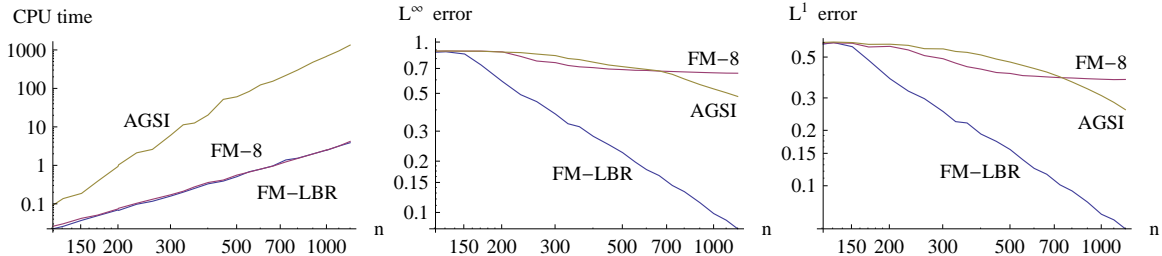
17

Figure 8: CPU Time (left, in seconds), L^∞ error (center), and averaged L^1 error (right) of the FM-LBR, FM-8 and AGSI, at several resolutions ranging from 120 to 1200 (log-log scale).

a smaller CPU time and offers a better accuracy than its alternatives. Note that one can also construct configurations in which anisotropy is not axis-aligned *and* the AGSI is more accurate than the FM-LBR. In some cases, the accuracy advantage of the AGSI even grows unboundedly as the anisotropy ratio of the metric tends to infinity. A heuristic analysis and prediction of this phenomenon is presented in Appendix B, where it is illustrated with a fifth test case.

The third [4] and fourth test cases are relevant benchmarks if one's objective is to use fast marching methods for the segmentation of tubular structures, in medical images or volume data respectively. The FM-LBR reveals its full potential, and stands out as the only practical option, in this more difficult setting which involves a discontinuous and highly anisotropic metric. The Riemannian metric tensor $\mathcal{M}(z)$ is the identity matrix, except at points z on the neighborhood of a curve Γ, where $\mathcal{M}(z)$ has a small eigenvalue δ_0^2 associated to an eigenvector tangent to Γ, and the other eigenvalue is 1 on the orthogonal space. The shortest path joining the point $(1,-1)$ (resp. $(0,0,3)$) to the origin is extracted via "gradient descent on the Riemannian manifold $(\Omega, \mathcal{M})$", see Appendix A for details:

$$\gamma'(t) = -\mathcal{M}(\gamma(t))^{-1}\nabla \mathrm{D}(\gamma(t)). \tag{16}$$

By construction of the Riemannian metric $\mathcal{M}$, traveling close and tangentially to the curve Γ is cheap. This is reflected by the level lines of D, and by the allure of the minimal path, see Figures 5, 7 and 10. Heuristically, this path joins the neighborhood of the curve Γ in straight line, almost orthogonally, and then follows it. The alignment of the minimal path with the direction of anisotropy, observed in this test case, is not an uncommon phenomenon. The FM-LBR presumably benefits a lot from this behavior in terms of accuracy, since its stencils typically provide a good "angular resolution" in the direction of anisotropy, see Figures 3, 4, 7. Since in addition the stencil radii remain rather small for most anisotropy orientations, see [15] for details, usually most updates for points in the "fast band" come from the fast band when the FM-LBR is run on these examples. When the fast band is missed however, accuracy degrades and zig-zag artifacts appear in the extracted path, see Figure 11.

Third test case [4], in 2D. The different method's performance is illustrated on Figure 8, except for the MAOUM which showed a poor accuracy, presumably due to the huge stencils it generated. The CPU time/resolution curve of the AGSI shows a stronger slope than for the one pass solvers, presumably reflecting its super-linear complexity $\mathcal{O}(\mu(\mathcal{M})N^{\frac{3}{2}})$. The L^∞ and L^1 error curves suggest that the FM-8 is not consistent in this test case, contrary to the FM-LBR and the AGSI. At the resolution 1000×1000, typical in image analysis, the FM-LBR cuts the L^∞ error by 80% and the L^1 error by 75% with respect to the AGSI, while reducing CPU time from 11

18

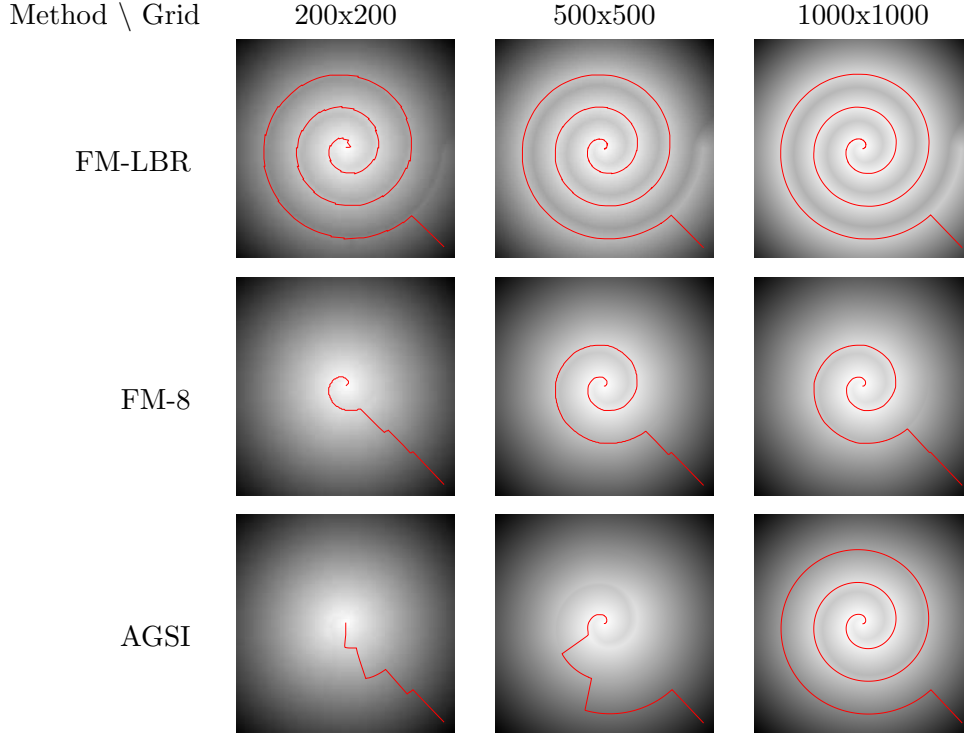| Method \ Grid	200x200	500x500	1000x1000
FM-LBR			
FM-8			
AGSI			

Figure 9: Visual comparison of the accuracy of three algorithms, at three resolutions, in the 2D test case. Qualitatively, the approximate geodesic has the right behavior for a resolution as low as 170×170 with the FM-LBR, and 1000×1000 with the AGSI. This is presumably never the case for the FM-8, which is not consistent here.

minutes to 2.5 seconds (!). As illustrated on Figure 9, the better accuracy of the FM-LBR in this test case effectively translates into a better extraction of minimal paths. The reason for the, un-rivaled, performance of the FM-LBR in this specific test case is partly elucidated in Appendix B.

Fourth test case, in 3D. CPU time was 105s for the FM-LBR, while the AGSI took 480s and failed to recover the minimal path presented on Figure 10 (center) (a straight line joining the two endpoints was obtained instead). The FM-LBR is capable of addressing a large scale (more than 10 millions grid points), strongly anisotropic ($\kappa(\mathcal{M}) = 50$) three dimensional shortest path problem, with a good accuracy and within reasonable CPU time on standard laptop computer.

Remark 3.1. *The four application inspired test cases. Outflow boundary conditions, except for* $D(0) = 0$*; hence the solution is the Riemannian distance to the origin:* $D(z) = D(z, 0)$*. Numerical errors with respect to a* 4000×4000 *reference solution, bi-linearly interpolated, obtained in cases 1,2, with the AGSI, and in case 3 with the FM-ASR [16].*

1. *(Geometry processing [28],* $\kappa(\mathcal{M}) \simeq 5.1$*) Compute the Riemannian distance from the origin* $(0, 0, 0)$ *on parametric surface of height map* $z(x, y) := (3/4) \sin(3\pi x) \sin(3\pi y)$*. Riemannian metric:* $\mathcal{M}(x, y) = \mathrm{Id} + \nabla z(x, y) \nabla z(x, y)^{\mathrm{T}}$*. Coordinates* (x, y) *restricted to the square* $[-0.5, 0.5]^2$*, discretized on a* 292×292 *grid, or in a second step to this square rotated by the indicated angle* θ*.*

2. *(Seismic imaging [28, 23],* $\kappa(\mathcal{M}) = 4$*) The metric* $\mathcal{M}(x, y)$ *has the two eigenvalues* 0.8^{-2}*,*
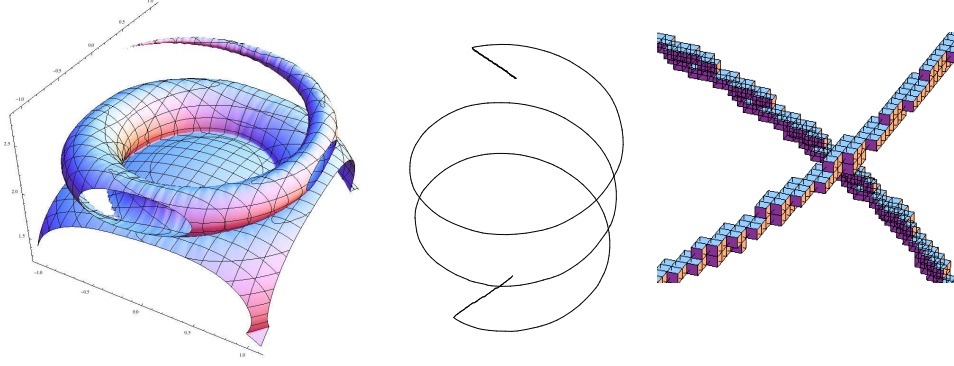
Figure 10: Results of the FM-LBR in the fourth, 3D, test case. Iso-surface $\{\mathrm{d}(z) = 2\}$ (left), and shortest path joining the points $(0, 0, 0)$ and $(3, 0, 0)$ (center). Detail of the discrete points (represented by small cubes), in the neighborhood of the curve $\Gamma(t) = (\cos \omega_0 t, \sin \omega_0 t, t)$, for which the Riemannian metric is not euclidean (right).

> 0.2^{-2}, *the former associated to the eigenvector* $(1, (\pi/2) \cos(4\pi x))$. *Domain* $[0.5, 0.5]^2$, *discretized on a* 193×193 *grid.*

> 3. *(Tubular segmentation [4], $\kappa(\mathcal{M}) = 100$) Define the curve* $\Gamma(t) = t(\cos \omega_0 t, \sin \omega_0 t)$, $t \in [0, 1]$. *Set* $\mathcal{M}(z) = \mathrm{Id}$, *except if there exists* $0 \le t \le 1$ *and* $0 \le r \le r_0$ *such that* $z = \Gamma(t) + r(\cos \omega_0 t, \sin \omega_0 t)$. *In that case* $\mathcal{M}(z)$ *has the eigenvalues* δ_0^2 *and* 1, *the former with eigenvector* $\Gamma'(t)$. *Parameters:* $\omega_0 := 6\pi$, $r_0 := \delta_0 := 0.01$. *Domain:* $[-1, 1]^2$, *grid sizes* $n \times n$ *with* $120 \le n \le 1200$.

> 4. *(Tubular segmentation, $\kappa(\mathcal{M}) = 50$, 3D) Define the curve* $\Gamma(t) = (\cos \omega_0 t, \sin \omega_0 t, t)$, *with* $\omega_0 := (5/2)\pi$. *Set* $\mathcal{M}(z) = \mathrm{Id}$, *except if there exists* $t, \lambda, \mu \in \mathbb{R}$ *such that* $z = \Gamma(t) + (\lambda \cos \omega_0 t, \lambda \sin \omega_0 t, \mu)$ *and* $\lambda^2 + \mu^2 \le (r_0/2)^2$. *In that case* $\mathcal{M}(z)$ *has the eigenvalues* δ_0^2 *and* 1, *the former with eigenvector* $\Gamma'(t)$ *and the latter with multiplicity* 2. *Parameters:* $\omega_0 := (5/2)\pi$, $\delta_0 = r_0 = 0.02$. *Domain:* $[-1.1, 1.1]^2 \times [0, 3]$, *grid size* $200 \times 200 \times 272$.

Conclusion

The FM-LBR, introduced in this paper, combines the Fast Marching algorithm with a concept from discrete geometry named Lattice Basis Reduction. It has the following strongpoints. (I, Convergence) The FM-LBR is consistent for the anisotropic eikonal equation associated to any continuous Riemannian metric, of arbitrary anisotropy. (II, Complexity) It has a numerical cost comparable to classical isotropic Fast Marching, independently of the problem anisotropy. (III, Accuracy) The accuracy of the FM-LBR is competitive in general, and striking in test cases, related to tubular segmentation in medical images, where the Riemannian metric has a pronounced anisotropy close to and tangentially to a curve.

These qualities come at the price of the specialization of the FM-LBR: (i) the Riemannian metric may not be replaced with a more general Finsler metric, see [16] for an adaptation to this setting in 2D, (ii) the domain needs to be discretized on a cartesian grid, and (iii) of dimension 2 or 3. Hopefully these requirements are met in many applications, and future work will be devoted to the application of the proposed algorithm in the context of medical image processing.
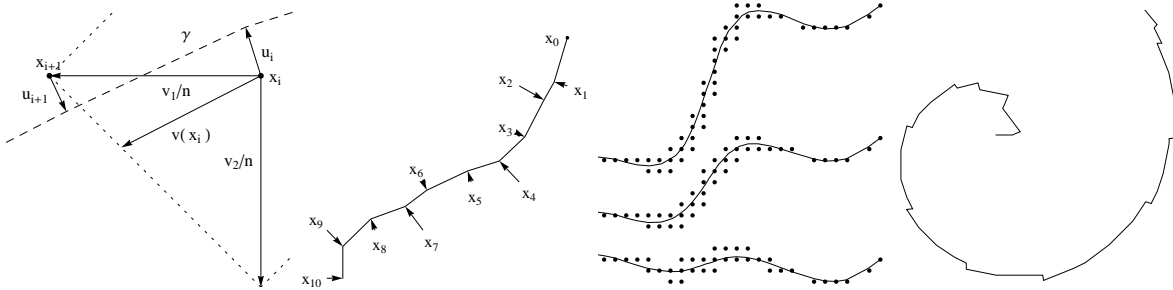
Figure 11: Left: Notations for the minimal path computation; the contour of the stencil $V(x_i)$ is shown dotted. Center left: Grid points $x_0, \cdots, x_r \in \Omega \cap Z$, corrections $u_0, \cdots, u_r \in \mathbb{R}^m$ shown as arrows, and piecewise linear path γ. Center right: Grid Points $(x_i)_{i=1}^r$ and piecewise linear path γ in the second test case at resolution 100×100, using the FM-LBR. Right: In hard test cases, combining strong anisotropy, metric discontinuity, and low grid resolution, the extracted path exhibits zig-zag artifacts (detail of the third test case at resolution $n = 200$).

A Robust extraction of minimal paths

Obtaining the shortest path joining two given points is essential in motion planning control problems [2], as well as in the envisioned application to tubular structure centerline extraction [4]. This involves solving the Ordinary Differential Equation (ODE) (16), a task less trivial than it seems. The author is conscious that a comparison of minimal paths, as on Figure 9, reflects the properties of the ODE solver (and the time spent adjusting its sometimes numerous parameters), as much as those of the eikonal solver. This is done nevertheless due to the importance of minimal paths in applications. Eikonal solvers based on the discrete fixed point problem (3), such as the FM-LBR, FM-8 and AGSI, provide at each grid point $x \in \Omega \cap Z$ an estimate $d(x)$ of the distance $D(x)$, and in addition an estimate $v(x)$ of the direction and orientation of the distorted negative gradient $-\mathcal{M}(x)^{-1} \nabla D(x)$, of the form:

$$v(x) := \sum_{1 \leq j \leq k} \alpha_j (z_j - x), \tag{17}$$

where the integer $1 \leq k \leq m$, the *positive* coefficients $(\alpha_j)_{j=1}^k$ and the vertices $(z_j)_{j=1}^k$ of $\partial V(z)$ are the barycentric coordinates of the point $y \in \partial V(x)$ achieving the minimum in the Hopf-Lax update operator (4), in the face of $\partial V(x)$ containing y and of minimal dimension.

From this point, a typical approach to solve (16) is to extend the values of d or v to the continuous domain Ω via an interpolation procedure, and then to use a black box ODE solver or a Runge Kutta method. Note that the accuracy usually expected from these high order methods is mostly doomed, since the discretization (3) of the eikonal equation is only first order, and since the vector field $\mathcal{M}^{-1} \nabla D$ is discontinuous both at "caustics" and discontinuities of $\mathcal{M}$. A more significant issue is that computations frequently get stuck, despite the use of state of the art and/or commercial interpolation methods and ODE solvers, see e.g. [2] Figure 5.10.

We propose a method for the computation of minimal paths, which trades high order accuracy for robustness, and never gets stuck if the eikonal solver is Dijkstra inspired. It takes advantage of the specific form (3) of the discretization of the eikonal equation, and does not rely on black box routines. It is also parameter free: there is not interpolation order or gradient step to adjust.

Algorithm 2 Minimal path computation, starting from a given grid point $x_0 \in \Omega \cap Z$.

Initialisation: $x \leftarrow x_0$, $u \leftarrow 0$ (mutable variables).

While x is not an initial source point, **do**

    Denote by $z_1, \cdots, z_k$ the grid points appearing in the expression (17) of $v(x)$.

    Find $\lambda \in \mathbb{R}_+$, $1 \leq j \leq k$, which minimize $\|x + u + \lambda v(x) - z_j\|$.

    Perform updates $u \leftarrow x + u + \lambda v(x) - z_j$ and $x \leftarrow z_j$.

The successive iterations of Algorithm 2 generate grid points $x_0, \cdots, x_r \in Z$, and small correcting offsets $u_0, \cdots, u_r \in \mathbb{R}^m$. If the causality property holds, see Proposition 1.1, then the values $(\mathrm{d}(x_i))_{i=0}^r$ are strictly decreasing, so that the algorithm is guaranteed to terminate. The piecewise linear path $\gamma : [0, r] \to \Omega$, parametrized so that $\gamma(i) := x_i + u_i$, satisfies the differential equation

$$\gamma'(t) = \lambda_{\lfloor t \rfloor} v(x_{\lfloor t \rfloor}), \tag{18}$$

for non-integer $t \in [0, r]$, where the constants $(\lambda_i)_{i=0}^{r-1}$ are the minimizers in the second step of the while loop, and the vector $v(z)$, $z \in \Omega \cap Z$, is defined in (17). The particularity of our path extraction method is that the direction field v is not evaluated on the curve γ, but at the points $(x_i)_{i=0}^r$ which remain close as shown by the next proposition (the involved exponent 3 seems to be either over-estimated or a rare worst case scenario, in view of the experiments Figure 11).

Proposition A.1. *Let $\Omega \subset \mathbb{R}^2$ be an open bounded set, equipped with a Riemannian metric $\mathcal{M} \in C^0(\overline{\Omega}, S_2^+)$, and an admissible family $(\mathcal{T}(z))_{z \in \Omega}$ of meshes with "Boundedness" constant R. (See Definition 1.7, $R \lesssim \kappa(\mathcal{M})$ for the FM-LBR). Consider a cartesian grid Z of scale h, equipped with the corresponding stencils, see (6) and (7), and solve the discrete system (3). Let $\gamma \in C^0([0, r], \Omega)$ be a path extracted with Algorithm 2 and parametrized as in (18). Then for all $t \in [0, r]$*

$$\|\gamma(t) - x_{\lfloor t \rfloor}\| \leq ChR^3, \tag{19}$$

where C is an absolute constant (i.e. independent of Ω, $(\mathcal{T}(z))_{z \in \Omega}$, $\mathcal{M}$, or the path origin).

The rest of this appendix is devoted to the proof. Consider a fixed $0 \leq i < r$, and observe that for all $t \in [i, i+1[$ one has, since γ is linear on this interval and since $\gamma(i) = x_i + u_i$:

$$\|\gamma(t) - x_{\lfloor t \rfloor}\| \leq \max\{\|\gamma(i+1) - x_i\|, \ \|\gamma(i) - x_i\|\} \leq \max\{\|u_{i+1}\| + \|x_{i+1} - x_i\|, \ \|u_i\|\}. \tag{20}$$

In order to avoid notational clutter, we denote $x := x_i$, $x' := x_{i+1}$, $u := u_i$, $u' := u_{i+1}$ and $v := v(x_i)$. Let $k \in \{1, 2\}$ and let z_j, $1 \leq j \leq k$, be the vertices of $\partial V(x)$ appearing in the expression (17) of the discrete negative gradient v.

We assume without loss of generality that the grid is $Z := \mathbb{Z}^2$, so that $v_j := z_j - x$ is a vertex of $\partial \mathcal{T}(x)$, for all $1 \leq j \leq k$. Hence $\|v_j\| \leq R$ by (Boundedness) and if $k = 2$ then $\det(v_1, v_2)$ is a non-zero integer. In particular $\|x' - x\| = \|v_j\| \leq R$, for some $1 \leq j \leq k$. By construction (second step of the while loop in the path computation), we have for any $\lambda \in \mathbb{R}_+$ and any $1 \leq j \leq k$

$$\|u'\| \leq \|u + \lambda v - v_j\|. \tag{21}$$

We prove in the following an upper bound of the form $\|u'\| \leq \max\{\|u\|, CR^3\}$, which by an immediate induction argument implies $\|u_i\| \leq CR^3$ for any $1 \leq i \leq r$. Using (19), our previous

estimate on $\|x_{i+1} - x_i\|$, and rescaling by a factor h, we obtain as announced (19). Case $k = 1$: choosing $\lambda = 1$, $j = 1$, and observing that $v = v_1$, we obtain $\|u'\| \leq \|u + 1 \times v - v_1\| = \|u\|$. The second case begins with a lemma.

Lemma A.2. *Let $v_1, v_2 \in I\!\!R^2$, and let $v := \alpha_1 v_1 + \alpha_2 v_2$, with $\alpha_1, \alpha_2 > 0$. Let $\mu := \sqrt{2}$, and let $w_1 := v_1 \mu - v_2/\mu$, $w_2 := v_2 \mu - v_1/\mu$. Then one can choose $\lambda \in I\!\!R_+$ and $1 \leq j \leq 2$, such that $v_j - \lambda v$ is positively proportional to any of the following vectors: w_1, w_2, $-w_1$ or $-w_2$, with a proportionality constant $0 < \nu \leq \mu$.*

Proof. In the case of w_1, choose $j := 1$, $\lambda := 1/(\alpha_1 + \mu^2 \alpha_2)$, so that $\nu = 1/(\mu + \mu^{-1} \alpha_1/\alpha_2) \leq 1/\mu$. In the case of $-w_1$, choose $j := 2$, $\lambda := 1/(\alpha_2 + \mu^{-2} \alpha_1)$, so that $\nu = 1/(\mu^{-1} + \mu \alpha_2/\alpha_1) \leq \mu$. The cases of w_2 and $-w_2$ are similar. $\square$

Case $k = 2$. We use the notations of the lemma. Denoting by A the matrix of lines w_1, w_2, there exists $1 \leq k \leq 2$ and $\varepsilon \in \{-1, 1\}$ such that

$$\sqrt{2}\langle u, \varepsilon w_k \rangle \geq \sqrt{\langle u, w_1 \rangle^2 + \langle u, w_2 \rangle^2} = \|Au\| \geq \|A^{-1}\|^{-1}\|u\|. \tag{22}$$

The norm $\|A^{-1}\|^{-1}$ is estimated as follows. $|\det A| = (\mu^2 - \mu^{-2})|\det(v_1, v_2)| \geq \mu^2 - \mu^{-2}$, since $\det(v_1, v_2)$ is a non-zero integer. On the other hand $\|w_j\| \leq (\mu + \mu^{-1}) \max\{\|v_1\|, \|v_2\|\} \leq C_0 R$ with $C_0 = 2(\mu + \mu^{-1})$, hence $\|A\| \leq \sqrt{\|w_1\|^2 + \|w_2\|^2} \leq \sqrt{2} C_0 R$. Finally $\|A^{-1}\|^{-1} = |\det A|/\|A\| \geq C_1/R$ where $C_1 := (\mu - \mu^{-1})/(\sqrt{2} C_0)$.

We next choose λ and j, using the previous lemma, so that $v_j - \lambda v = \nu \varepsilon w_k$, with ε, k as in (22), and where $0 < \nu \leq \mu := \sqrt{2}$. Hence, using (21):

$$\begin{aligned}
\|u'\|^2 &\leq \|u - (v_j - \lambda v)\| = \|u - \nu \varepsilon w_k\|^2 = \|u\|^2 - 2\nu \langle u, \varepsilon w_k \rangle + \nu^2 \|w_k\|^2 \\
&\leq \|u\|^2 - 2\nu \|A^{-1}\|^{-1}\|u\| + \nu^2 \|w_k\|^2 \leq \|u\|^2 - 2C_1\|u\|/R + C_0^2 R^2.
\end{aligned}$$

If $\|u\| \geq C_2 R^3$, with $C_2 := C_0^2/(2C_1)$, then $\|u'\| \leq \|u\|$. If $\|u\|$ is below this bound, then choosing $\lambda = 0$ in (21) yields $\|u'\| \leq \|u\| + \|v_1\| \leq C_2 R^3 + 2R$. Thus $\|u'\| \leq \max\{\|u\|, CR^3\}$ as announced, which concludes the proof.

B Stencil size, discretization errors and metric regularity

Experience in the discretization of Partial Differential Equations (PDEs) tells that robust and accurate numerical schemes are usually based on small, localized and isotropic stencils. The FM-LBR, which achieves causality in the eikonal equation by the use of long range, sparse and highly anisotropic stencils, seems to violate this principle. We give in this section a *heuristic* analysis of its accuracy, which explains its excellent performance in the third and fourth tests (once a source of puzzlement for the author and the reviewers), but also exposes some weaknesses.

Let us emphasize that the computational domain Ω is equipped with two geometries.

- The *extrinsic* Euclidean geometry, inherited from the embedding $\Omega \subset I\!\!R^m$.

- The *intrinsic* Riemannian geometry, given by the Riemannian metric $\mathcal{M}$ on Ω, which is part of eikonal PDE structure.

The AGSI stencils are small, localized, with respect to the extrinsic Euclidean distance. The FM-LBR stencil at a point $z \in \Omega$ is built from an $\mathcal{M}(z)$-reduced basis, see §2.2, which consists of the smallest linearly independent vectors of $\mathbb{Z}^m$ in the local norm $\|\cdot\|_{\mathcal{M}(z)}$. Hence this stencil is by construction small and localized *in the twisted perspective* of the intrinsic Riemannian

distance. We refer to [15] for a quantitative estimate of the size of the FM-LBR stencil, from this perspective, in average over all orientations of the discretization grid. To better reflect the shapes of these stencils, which (as much as possible) adapt their orientation and aspect ratio to the metric $\mathcal{M}$, but keep a constant volume (they are built of $(m+1)!$ simplices of volume $h^m/m!$), we introduce a normalized metric $\widehat{\mathcal{M}}$: for all $z \in \Omega$

$$\widehat{\mathcal{M}}(z) := \det(\mathcal{M}(z))^{-\frac{1}{m}} \mathcal{M}(z), \tag{23}$$

so that $\det(\widehat{\mathcal{M}}(z)) = 1$ identically. For all $x, y \in \overline{\Omega}$ we denote by $\mathrm{D}(x, y)$ (resp. $\widehat{\mathrm{D}}(x, y)$) the Riemannian distance (2) on Ω associated to $\mathcal{M}$ (resp. $\widehat{\mathcal{M}}$).

Bellman's optimality principle applied to the solution D of the eikonal equation (1), reads: for all $x \in V \subset \Omega$

$$\mathrm{D}(x) = \min_{y \in \partial V} \mathrm{D}(x, y) + \mathrm{D}(y).$$

The Hopf-Lax update operator (4) reflects this identity at discrete level, up to two approximations:

$$\mathrm{D}(x, y) \approx \|x - y\|_{\mathcal{M}(x)}, \tag{24}$$

$$\mathrm{D}(y) \approx \mathrm{I}_{V(x)}\,\mathrm{d}(y), \tag{25}$$

where $V(x)$ denotes the stencil at x (which is given under the form of a triangulation of a neighborhood of x), I_V the linear interpolation operator on V, and $y \in \partial V$.

If one ignores the issue of the scheme causality, then the choice of stencil should be dictated by the local regularity properties of the quantities that are approximated on it. Discretization (24) amounts to approximating the metric $\mathcal{M}$ with the constant $\mathcal{M}(x)$ on the stencil $V(x)$. Such piecewise constant approximation errors are controlled by Lipschitz regularity constants. The natural distance on S_m^+ is defined in (13), but the distance on Ω should be chosen appropriately so as to reflect the geometry of the computational stencils. Indeed, the AGSI (resp. FM-LBR) stencil at a point $x \in \Omega$ is heuristically not much different from ball centered at x and of radius the grid scale h, defined with respect to the euclidian distance (resp. Riemannian distance $\widehat{D}$). The AGSI stencil should therefore be preferred if the metric $\mathcal{M}$ has a small Lipschitz regularity constant K_0 with respect to the extrinsic Euclidean distance:

$$\mathrm{dist}(\mathcal{M}(x), \mathcal{M}(y)) \leq K_0 \|x - y\|. \tag{26}$$

On the other hand, the FM-LBR stencil is more suitable for metrics which have a small Lipschitz regularity constant K_1 with respect to the intrinsic (up to the normalization (23)) distance $\widehat{D}$:

$$\mathrm{dist}(\mathcal{M}(x), \mathcal{M}(y)) \leq K_1 \widehat{D}(x, y). \tag{27}$$

In the fifth numerical example below, the most accurate of these two methods can indeed be guessed from the ratio K_0/K_1. Regularity conditions of the form (27) arise naturally in the study of anisotropic mesh generation, see Part III of [17]. The Riemannian metric involved in the third and fourth numerical tests of this section, inspired by applications to tubular structure segmentation [4], varies slowly in the direction of the eigenvector associated to the small eigenvalue of $\mathcal{M}(z)$ (the direction of the tube), but quickly (in fact discontinuously) in the orthogonal direction. Thus, although discontinuous, it is heuristically not far from satisfying (27), which explains the exceptional performance of the FM-LBR on these specific examples.

The interpolation error (25) is harder to estimate, yet in favor of the FM-LBR stencil let us mention the intrinsic 1-Lipschitz regularity of the solution: for all $x, y \in \Omega$

$$|D(x) - D(y)| \leq D(x, y).$$

In the special case of a constant metric, where (24) is exact and all error comes from (25), the FM-LBR stencil offers the best accuracy, see [15].

The following illustrative example was proposed by A. Vladimirsky: let $M \in S_m^+$, let $u \in \mathbb{R}^m$, and let $\mathcal{M} : \mathbb{R}^m \to S_m^+$ be the Riemannian metric defined by

$$\mathcal{M}(z) := M \exp(\langle u, z \rangle). \tag{28}$$

Assume for normalization that $\det(M) = 1$, so that $\widehat{\mathcal{M}} = M$ identically, and $\widehat{D}(x, y) = \|x - y\|_M$ for all $x, y \in \mathbb{R}^m$. Then, with $D := M^{-1}$:

$$\text{dist}(\mathcal{M}(x), \mathcal{M}(y)) = |\langle u, x - y \rangle| \leq \|u\| \|x - y\|,$$
$$\text{dist}(\mathcal{M}(x), \mathcal{M}(y)) = |\langle u, x - y \rangle| \leq \|u\|_D \|x - y\|_M = \|u\|_D \widehat{D}(x, y).$$

The Lipschitz regularity constants are therefore $K_0 = \|u\|$, and $K_1 = \|u\|_D$. The discretization (24) is hence likely more accuracte with the AGSI stencil if $\|u\| \leq \|u\|_D$, and with the FM-LBR stencil otherwise. Defining for all $z \in \mathbb{R}^m$

$$\text{d}(z) := \exp(\langle u, z \rangle)/\|u\|_D, \tag{29}$$

we observe that this map is unimodular: $\|\nabla \text{d}(z)\|_{\mathcal{M}(z)^{-1}} = 1$. The value $\text{d}(z)$ can also be regarded as the geodesic distance from z to a point at infinity in the direction of $-Du$. The characteristic curves of the solution are parallel straight lines, of direction Du. Although the present discussion is about accuracy rather than CPU time, let us mention that thanks to this special property, and as pointed out by A. Vladimirsky, the AGSI converges in a single pass in the numerical tests below (provided its priority queue is suitably initialized: the upwind boundary points z must be sorted by increasing values of their scalar product $\langle z, Du \rangle$ with the characteristics direction Du). As a result, and in contrast with §3, the AGSI CPU time is here only *half* of the FM-LBR CPU time.

Our fifth and last numerical test, involves a metric depending on three parameters $\kappa \geq 1$, $\theta, \varphi \in \mathbb{R}$: denoting $e_\theta := (\cos\theta, \sin\theta)$, and $(x, y)^\perp := (y, -x)$,

$$\mathcal{M}(z; \kappa, \theta, \varphi) := M(\kappa, \theta) \exp(\langle z, e_\varphi \rangle), \quad \text{with } M(\kappa, \theta) := \kappa e_\theta e_\theta^\mathrm{T} + \kappa^{-1} e_\theta^\perp (e_\theta^\perp)^\mathrm{T}. \tag{30}$$

For each matrix $M = M(\theta, \varphi)$, given by its condition number κ, and anisotropy direction θ, we consider different unit vectors $u = e_\varphi$, given by their angle φ with respect to the x-axis. For $\varphi = \theta$ the FM-LBR is favored, since K_0/K_1 takes its maximal value $\sqrt{\kappa}$. For $\varphi = -\pi/4$, the AGSI is favored, since K_0/K_1 is close to its minimal value $1/\sqrt{\kappa}$, and the interpolation error (25) vanishes. The domain is the $[-2, 2]^2$ square discretized on a 100×100 grid. The FM-LBR stencil is included in a square of $(2w + 1) \times (2w + 1)$ pixels, where w respectively equals 2, 3, and 5 for the three different pairs (κ, θ) in Table 1. The boundary condition (29) is applied on a the upwind part of the square boundary, on a layer of width w. The numerical tests presented in Table 1 are typical of the authors experience, and tend to agree with the above heuristical error analysis. Note however that the FM-LBR performance is unexpectedly bad[9] in case *, and unexpectdly good in cases †.

---

[9]The FM-LBR in that case produces large errors close to the *downwind* boundary, presumably due to its large stencils. Removing a 5 pixel band on the boundary yields in this case the L^∞ errors: 32.3 (FM-LBR) and 45.8 (AGSI), in favor of the FM-LBR as predicted from the Lipschitz constants ratio.

κ	θ	φ	K_0/K_1 $= \|u\|/\|u\|_D$	L^1 error FM-LBR	AGSI	L^∞ error FM-LBR	AGSI
3	$\pi/3$	$\pi/3$	1.73	2.78	6.96	40.2	60.5
		$\pi/6$	1	2.80	3.07	17.1	18.5
		$-\pi/4$	0.59	3.95	1.45	37.4	13.7
10	$3\pi/8$	$3\pi/8$	3.16	3.74	9.45	38.4	79.5
		1.48	1	1.34	2.03	7.44	11.4
		$-\pi/4$	0.34	2.92	0.83	27.3	7.82
30	$\pi/8$	$\pi/8$	5.48	3.89	6.62	94.3*	61.0
		0.57	1	$0.91^\dagger$	2.55	$4.62^\dagger$	12.9
		$-\pi/4$	0.20	3.24	0.49	29.4	4.51

Table 1: Metric $z \mapsto \mathcal{M}(z; \kappa, \theta, \varphi)$, see (30), on the $[-2,2]^2$ square discretized on a 100×100 grid. The most accurate algorithm, among the AGSI and the FM-LBR, can in most cases be predicted by the Lipschitz constants ratio K_0/K_1, at least when it is far from 1. The star points out an exception. CPU time approximatively 10 ms for the AGSI, and 20 ms for the FM-LBR. Numerical errors multiplied by 100 for better readability.

The accuracy advantage of the AGSI is larger than a factor 6 for the last set of parameters in Table 1, and it may of course grow unboundedly as the anisotropy ratio κ tends to infinity, for suitable angles θ, φ. Anisotropy does therefore, sometimes, play against the FM-LBR accuracy.

References

[1] K. Alton, I. M. Mitchell, *Fast Marching Methods for Stationary Hamilton-Jacobi Equations with Axis-Aligned Anisotropy*, SIAM Journal of Numerical Analysis, 47:1, pp. 363–385, 2008.

[2] K. Alton, I. M. Mitchell, *An Ordered Upwind Method with Precomputed Stencil and Monotone Node Acceptance for Solving Static Hamilton-Jacobi Equations*, Journal of Scientific Computing, 51:2, pp. 313–348, 2012.

[3] T. J. Barth, J. A. Sethian, *Numerical schemes for the Hamilton-Jacobi and level set equations on triangulated domains*, Journal of Computational Physics, 145(1), 1-40, 1998.

[4] F. Benmansour, L. D. Cohen, *Tubular Structure Segmentation Based on Minimal Path Method and Anisotropic Enhancement*, International Journal of Computer Vision, 92(2), 192-210, 2010.

[5] F. Bornemann, C. Rasch, *Finite-element Discretization of Static Hamilton-Jacobi Equations based on a Local Variational Principle*, Computing and Visualization in Science, 9(2), 57-69, 2006.

[6] J. H. Conway, N. J. A. Sloane, *Low-dimensional lattices. VI. Voronoi reduction of three-dimensional lattices.*, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences 436.1896, 55-68, 1992.

[7] J. Fehrenbach, and J.-M. Mirebeau, *Sparse Non-Negative Stencils for Anisotropic Diffusion*, Journal of Mathematical Imaging and Vision, available online, to appear in print, 2013.

[8] R. Gonzales, E. Rofman,*On Deterministic Control Problems: an Approximate Procedure for the Optimal Cost, I, the Stationary Problem*, SIAM Journal on Control and Optimization, 23, 2, pp. 242-266, 1985.

[9] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pélégrini-Issac, H. Benali, *Accurate Anisotropic Fast Marching for Diffusion-Based Geodesic Tractography*, International Journal of Biomedical Imaging, 2008.

[10] R. Kimmel, J. A. Sethian, *Computing geodesic paths on manifolds*, Proceedings of the National Academy of Sciences USA, 95(15), 8431-8435, 1998.

[11] H.J. Kushner, P.G. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, Academic Press, New York, 1992.

[12] J. L. Lagrange, *Recherches d'arithmétique*, Nouveaux Mémoires de lAcadémie de Berlin, 1773.

[13] X.-G. Li, W. Yan, C. K. Chan, *Numerical schemes for Hamilton-Jacobi equations on unstructured meshes*, Numerische Mathematik, 94(2), 315-331, 2003.

[14] P.L. Lions, *Generalized solutions of Hamilton-Jacobi equations*, Pitman, Boston, 1982.

[15] J.-M. Mirebeau, *On the Accuracy of Anisotropic Fast Marching*, preprint available on Arxiv, 2012.

[16] J.-M. Mirebeau, *Efficient Fast Marching with Finsler Metrics*, Numerische Mathematic, available online, to appear in print, 2013.

[17] J.-M. Mirebeau, *Adaptive and anisotropic finite element approximation : Theory and algorithms*, PhD thesis.

[18] P. Q. Nguyen, and D. Stehlé, *Low-dimensional lattice basis reduction revisited*, ACM Transactions on Algorithms, Article 46, 2009.

[19] G. Peyré, M. Péchaud, R. Keriven, L. D. Cohen, *Geodesic Methods in Computer Vision and Graphics*, Foundations and Trends in Computer Graphics and Vision, 5(3-4), 197-397, 2010.

[20] C. Rasch, T. Satzger, *Remarks on the $\mathcal{O}(N)$ Implementation of the Fast Marching Method*, IMA Journal of Numerical Analysis, 29, 806-813, 2009.

[21] M. Sermesant, E. Konukoglu, H. Delingette, *An anisotropic multi-front fast marching method for real-time simulation of cardiac electrophysiology*, Proc of Functional Imaging and Modeling of the Heart, 2007.

[22] J.A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision and Materials Sciences*, Cambridge University Press, 1996.

[23] J. A. Sethian, A. Vladimirsky, *Ordered Upwind Methods for Static Hamilton-Jacobi Equations: Theory and Algorithms*, SIAM Journal of Numerical Analysis, 41(1), 325-363, 2003.

[24] J. A. Sethian, A. Vladimirsky, *Fast methods for the Eikonal and related Hamilton-Jacobi equations on unstructured meshes*, Proceedings of the National Academy of Sciences, 97(11), 5699-5703, 2000.

[25] Y.-H.R. Tsai, L.-T. Cheng, S. Osher, and H.-K. Zhao, *Fast sweeping algorithms for a class of Hamilton-Jacobi equations*, SIAM Journal on Numerical Analysis, 41:2, pp.659-672, 2003.

[26] J. N. Tsitsiklis, *Efficient algorithms for globally optimal trajectories*, IEEE Transactions on Automatic Control, 40(9), 1528-1538, 1995.

[27] A. Vladimirsky, *Label-setting methods for Multimode Stochastic Shortest Path problems on graphs*, Mathematics of Operations Research 33(4), pp. 821-838, 2008.

[28] A. Vladimirsky, *Fast methods for static Hamilton-Jacobi Partial Differential Equations*, PhD Thesis, 2001.

[29] A. Vladimirsky, *Static PDEs for Time-Dependent Control Problems*, Interfaces and Free Boundaries, 8(3), pp. 281-300, 2006.

[30] H. Zhao, *A Fast Sweeping Method for Eikonal Equations*, Mathematics of Computation, 74(250), 603-627, 2005.