Tetrahedral mesh improvement using moving mesh smoothing, lazy searching flips, and RBF surface reconstruction

Franco Dassi[1], Lennard Kamenski[2], Patricio Farrell[2], Hang Si[2]

submitted: February 21, 2017

[1] Dipartimento di Matematica e Applicazioni
     Università degli Studi di Milano-Bicocca
     Via Cozzi 53
     20125 Milano, Italy
     E-Mail: franco.dassi@unimib.it

[2] Weierstrass Institute
     Mohrenstr. 39
     10117 Berlin
     Germany
     E-Mail: lennard.kamenski@wias-berlin.de
             patricio.farrell@wias-berlin.de
             hang.si@wias-berlin.de

Abstract

Given a tetrahedral mesh and objective functionals measuring the mesh quality which take into account the shape, size, and orientation of the mesh elements, our aim is to improve the mesh quality as much as possible. In this paper, we combine the *moving mesh smoothing*, based on the integration of an ordinary differential equation coming from a given functional, with the *lazy flip* technique, a reversible edge removal algorithm to modify the mesh connectivity. Moreover, we utilize *radial basis function* (RBF) surface reconstruction to improve tetrahedral meshes with curved boundary surfaces. Numerical tests show that the combination of these techniques into a mesh improvement framework achieves results which are comparable and even better than the previously reported ones.

1   Introduction

The key mesh improvement operations considered in this work are *smoothing*, which moves the mesh vertices, *flipping*, which changes the mesh topology without moving the mesh vertices, and a *smooth boundary reconstruction*. Previous work shows that the combination of smoothing and flipping achieves better results than if applied individually [16, 31]. In this paper, we combine the recently developed flipping and smoothing methods into one mesh improvement scheme and apply them in combination with a smooth boundary reconstruction via radial basis functions.

Mesh smoothing improves the mesh quality by improving vertex locations, typically through Laplacian smoothing or some optimization-based algorithm. Most commonly used mesh smoothing methods are Laplacian smoothing and its variants [11, 34], where a vertex is moved to the geometric center of its neighboring vertices. While economic, easy to implement, and often effective, Laplacian smoothing guarantees neither a mesh quality improvement nor mesh validity. Alternatives are optimization-based methods that are effective with respect to certain mesh quality measures such as the ratio of the area to the sum of the squared edge lengths [1], the ratio of the volume to a power of the sum of the squared face areas [36], the condition number of the Jacobian matrix of the affine mapping between the reference element and physical elements [15], or various other measures [16, 33, 32, 2]. Most of the optimization-based methods are local and sequential, combining Gauss-Seidel-type iterations with location optimization problems over each patch. There is also a parallel algorithm that solves a sequence of independent subproblems [14].

In our scheme, we employ the moving mesh PDE (MMPDE) method, defined as the gradient flow equation of a meshing functional (an objective functional in the context of optimization) to move the mesh continuously in time. Such a functional is typically based on error estimation or physical and geometric considerations. Here, we consider a functional based on the equidistribution and alignment conditions [21] and employ the recently developed direct geometric discretization [23] of the underlying meshing functional on simplicial meshes. Compared to the aforementioned mesh smoothing methods, the considered method has several advantages: it can be easily parallelized, it is based on a continuous functional for which the existence of minimizers is known, the functional controlling the mesh shape

and size has a clear geometric meaning, and the nodal mesh velocities are given by a simple analytical matrix form. Moreover, the smoothed mesh will stay valid if it was valid initially [24].

Flipping is the most efficient way to locally improve the mesh quality and it has been extensively addressed in the literature [30, 16, 17, 31]. In the simplest case, the basic flip operations, such as 2-to-3, 3-to-2, and 4-to-4 flips, are applied as long as the mesh quality can be improved. The more effective way is to combine several basic flip operations into one edge removal operation, which extends the 3-to-2 and 4-to-4 flips. This operation removes the common edge of $n \geq 3$ adjacent tetrahedra by replacing them with $m = 2n - 4$ new tetrahedra (the so-called n-to-m flip). There are at most C_{n-2} possible variants to remove an edge by a n-to-m flip, where $C_n = \frac{(2n)!}{(n+1)!\,n!}$ is the Catalan number. If n is small (e.g., $n < 7$), one can enumerate all possible cases, compute the mesh quality for each case, and then pick the optimal one. Another way is to use dynamic programming to find the optimal configuration. However, the number of cases increases exponentially and finding the optimal solution with brute force is very time-consuming.

In this paper, we propose the so-called *lazy searching flips*. The key idea is to automatically explore sequences of flips to remove a given edge in the mesh. If a flip sequence leads to a configuration which does not improve the mesh quality, the algorithm reverses this sequence and explores another one (see Section 3 and Figs 2a to 2c). Once an improvement is found, the algorithms stops the search and returns without exploring the remaining possibilities.

When considering more arbitrary meshes (which may not be piecewise planar), we need to make sure that new nodes are added in a consistent way. To achieve this we use *RBF surface reconstruction* as introduced in [3]. Radial basis functions are a very useful tool in the context of higher-dimensional interpolation as they dispense with the expensive generation of a mesh [13, 29, 40]. Here, we will employ them to approximate the underlying continuous surface so that we can project nodes onto it as proposed in [6, 5]. This problem turns out to be very challenging for meshes with arbitrary boundary. Hence, we begin with a relatively simple mesh. For more complicated examples we first refine the boundary by using the RBF reconstruction and projection method and then keep the boundary nodes fixed while interior nodes may move.

In this paper, we provide a detailed numerical study of a combination of the MMPDE smoothing with the lazy searching flips and RBF surface reconstruction. More specifically, we compare the results of the whole algorithm with `Stellar` [31], `CGAL` [39] and `mmg3d` [8]. We also compare the lazy searching flips and the MMPDE smoothing with the flipping and smoothing procedures provided by `Stellar`.

2   The moving mesh PDE smoothing scheme

The key idea of this smoothing scheme is to move the mesh vertices via a moving mesh equation, which is formulated as the gradient system of an energy functional (the MMPDE approach). Originally, the method was developed in the continuous setting [27, 28]. In this paper, we use its discrete form [23, 24, 26], for which the mesh vertex velocities are expressed in a simple, analytical matrix form, which makes the implementation more straightforward to parallelize.

2.1   Moving mesh smoothing

Consider a polygonal (polyhedral) domain $\Omega \subset \mathbb{R}^d$ with $d \geq 1$. Let $\mathcal{T}_h$ denote the simplicial mesh as well as $\#\mathcal{N}_h$ and $\#\mathcal{T}_h$ the numbers of its vertices and elements, respectively. Let K be a generic

mesh element and $\hat{K}$ the reference element taken as a regular simplex with volume $|\hat{K}| = 1/\#\mathcal{T}_h$. Further, let F'_K be the Jacobian matrix of the affine mapping $F_K \colon \hat{K} \to K$ from the reference element $\hat{K}$ to a mesh element K. For notational simplicity, we denote the inverse of the Jacobian by $\mathbb{J}_K$, i.e., $\mathbb{J}_K := (F'_K)^{-1}$ (see Fig. 1).
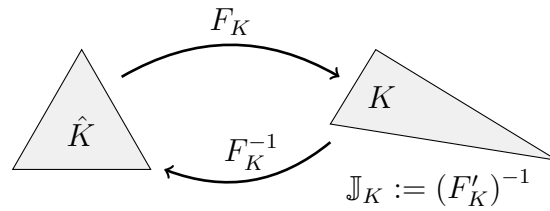

Figure 1: Reference element $\hat{K}$, mesh element K, and the corresponding mappings F_K and F_K^{-1}.

Then, the mesh $\mathcal{T}_h$ is uniform if and only if

$$|K| = \frac{|\Omega|}{\#\mathcal{T}_h} \quad \text{and} \quad \frac{1}{d} \operatorname{tr}\left(\mathbb{J}_K^T \mathbb{J}_K\right) = \det\left(\mathbb{J}_K^T \mathbb{J}_K\right)^{\frac{1}{d}} \quad \forall K \in \mathcal{T}_h. \tag{1}$$

The first condition requires all elements to have the same size and the second requires all elements to be shaped similarly to $\hat{K}$ (these conditions are the simplified versions of the equidistribution and alignment conditions [22, 28]).

The corresponding energy functional for which the minimization will result in a mesh satisfying Eq. (1) as closely as possible is

$$I_h = \sum_K |K| \, G\left(\mathbb{J}_K, \det \mathbb{J}_K\right) \tag{2}$$

with

$$G(\mathbb{J}, \det \mathbb{J}) = \theta\left(\operatorname{tr}\left(\mathbb{J}\mathbb{J}^T\right)\right)^{\frac{dp}{2}} + (1 - 2\theta)\, d^{\frac{dp}{2}} (\det \mathbb{J})^p, \tag{3}$$

where $\theta \in (0, 0.5]$ and $p > 1$ are dimensionless parameters (in Section 6, we use $\theta = 1/3$ and $p = 3/2$). This is a specific choice and other meshing functionals are possible. The interested reader is referred to [25] for a numerical comparison of meshing functionals for variational mesh adaptation.

In Eq. (2), I_h is a Riemann sum of a continuous functional for variational mesh adaptation based on equidistribution and alignment [21] and depends on the vertex coordinates $\boldsymbol{x}_i$, $i = 1, \ldots, \#\mathcal{N}_h$. The corresponding vertex velocities $\boldsymbol{v}_i$ for the mesh movement are defined as

$$\boldsymbol{v}_i := \frac{d\boldsymbol{x}_i}{dt} = -\left(\frac{\partial I_h}{\partial \boldsymbol{x}_i}\right)^T, \quad i = 1, \ldots, \#\mathcal{N}_h, \tag{4}$$

where the derivatives $\frac{d\boldsymbol{x}_i}{dt}$ are considered to be row vectors.

2.2 Vertex velocities and the mesh movement

The vertex velocities $\boldsymbol{v}_i$ can be computed analytically [23, Eqs (39) to (41)] using scalar-by-matrix differentiation [23, Sect. 3.2]. Denote the vertices of K and $\hat{K}$ by $\boldsymbol{x}_j^K$ and $\hat{\boldsymbol{x}}_j$, $j = 0, \ldots, d$, and define the element edge matrices as

$$E_K = [\boldsymbol{x}_1^K - \boldsymbol{x}_0^K, \ldots, \boldsymbol{x}_d^K - \boldsymbol{x}_0^K],$$
$$\hat{E} = [\hat{\boldsymbol{x}}_1 - \hat{\boldsymbol{x}}_0, \ldots, \hat{\boldsymbol{x}}_d - \hat{\boldsymbol{x}}_0].$$

Note, that $\hat{E}E_K^{-1} = \mathbb{J}_K$. Then, the local mesh velocities are given element-wise [23, Eqs (39) and (41)] by

$$\begin{bmatrix} (\boldsymbol{v}_1^K)^T \\ \vdots \\ (\boldsymbol{v}_d^K)^T \end{bmatrix} = -G_K E_K^{-1} + E_K^{-1}\frac{\partial G_K}{\partial \mathbb{J}}\hat{E}E_K^{-1} + \frac{\partial G_K}{\partial \det \mathbb{J}}\frac{\det(\hat{E})}{\det(E_K)}E_K^{-1}, \tag{5}$$

$$(\boldsymbol{v}_0^K)^T = -\sum_{j=1}^{d}(\boldsymbol{v}_j^K)^T,$$

where

$$G_K = G(\mathbb{J}_K, \det\mathbb{J}_K),$$

$$\frac{\partial G_K}{\partial \mathbb{J}} = \frac{\partial G}{\partial \mathbb{J}}(\mathbb{J}_K) = dp\theta\big(\mathrm{tr}(\mathbb{J}_K\mathbb{J}_K^T)\big)^{\frac{dp}{2}-1}\mathbb{J}_K^T,$$

$$\frac{\partial G_K}{\partial \det\mathbb{J}} = \frac{\partial G}{\partial \det\mathbb{J}}(\det\mathbb{J}_K) = p(1-2\theta)d^{\frac{dp}{2}}(\det\mathbb{J}_K)^{p-1}$$

are the derivatives of G with respect to its first and second argument [23, Example 3.2] evaluated at $\mathbb{J} = \mathbb{J}_K$ and $\det(\mathbb{J}) = \det\mathbb{J}_K$.

The moving mesh equation (4) becomes

$$\frac{d\boldsymbol{x}_i}{dt} = \sum_{K\in\omega_i}|K|\,\boldsymbol{v}_{i_K}^K, \quad i=1,\ldots,\#\mathcal{N}_h, \tag{6}$$

where ω_i is the patch of the vertex $\boldsymbol{x}_i$ and i_K is the local index of $\boldsymbol{x}_i$ on K.

During smoothing, we use the current vertex locations as the initial position and integrate Eq. (6) for a time period (with the proper modification for the boundary vertices, see Section 2.3). The connectivity is kept fixed during the smoothing step. In our examples in Section 6, we use the explicit Runge-Kutta Dormand-Prince ODE solver [9].

The moving mesh governed by Eq. (6) will stay nonsingular if it is nonsingular initially: the minimum height and the minimum volume of the mesh elements will stay bounded from below by a positive number depending only on the initial mesh and the number of the elements [24].

2.3   Velocity adjustment for the boundary vertices

The velocities of the boundary vertices need to be modified. If $\boldsymbol{x}_i$ is a fixed boundary vertex, then its velocity is set to zero Otherwise, $\boldsymbol{x}_i$ is allowed to move along a boundary curve or a surface represented by the zero level set of a function ϕ and its velocity is modified so that its normal component along the curve (surface) is zero:

$$\nabla\phi(\boldsymbol{x}_i)\cdot\frac{\partial\boldsymbol{x}_i}{\partial t} = 0.$$

For the special case of a piecewise linear complex (PLC) [35] the velocity adjustment is straightforward:

$$\begin{aligned} \text{facet vertices:}\quad &\text{project the velocity onto the facet plane,} \\ \text{segment vertices:}\quad &\text{project the velocity onto the segment line,} \\ \text{corner vertices:}\quad &\text{set the velocity to zero.} \end{aligned}$$

For a general non-polygonal or non-polyhedral domain, a simple way to adjust the boundary vertices is to move the vertex and then project it onto the boundary to which it belongs, which proved to work well for simple surface geometries (see Section 6.2). However, for complicated geometries, this simple projection can fail and a more reliable approach is needed.

3 Lazy searching flips

In this section, we explain how to remove an edge and how to reverse the removal using flips. In addition, we present the lazy searching algorithm which can be used to improve the quality of a mesh.

3.1 Edge removal and its inverse

A basic edge removal algorithm performs a sequence of elementary 2-to-3 and 3-to-2 flips [38]. We extend this algorithm by allowing the flip sequence to be reversed. Our algorithm saves the flips online and it uses no additional memory.

Let $[a, b] \in \mathcal{T}_h$ be an edge with endpoints a and b and $A[0, \ldots, n-1]$ be the array of $n \geq 3$ tetrahedra in $\mathcal{T}_h$ sharing $[a, b]$. For simplicity, we assume that $[a, b]$ is an interior edge of $\mathcal{T}_h$, so that all tetrahedra in A can be ordered cyclically such that the two tetrahedra $A[i]$ and $A[(i + 1) \mod n]$ share a common face. The index i takes values in $\{0, 1, \ldots, n-1\}$. Throughout this section, additions involving indices will be modulo n.

Given such an array A of n tetrahedra, we want to find a sequence of flips that will remove the edge $[a, b]$. Moreover, we also want to be able to reverse this sequence in order to return to the original state.

Our edge removal algorithm includes two subroutines

$$[\mathrm{done}, m] := \mathtt{flipnm}(A[0, \ldots, n-1], \mathtt{level}),$$
$$\mathtt{flipnm\_post}(A[0, \ldots, n-1], m)$$

with an array A (of length n) of tetrahedra and an integer $\mathtt{level}$ defining the maximum recursive level as input.

flipnm executes "forward" flips to remove the edge $[a, b]$. It returns a Boolean value indicating whether the edge is removed or not and an integer m ($3 \leq m \leq n$). If the edge is not removed ($\mathrm{done} = \mathrm{FALSE}$), m indicates the current size of A (initially, $m := n$).

flipnm_post must be called immediately after $\mathtt{flipnm}$. It releases the memory allocated in $\mathtt{flipnm}$ and can perform "backward" flips to undo the flip sequence performed by $\mathtt{flipnm}$.

The basic subroutine $\mathtt{flipnm}(A[0, \ldots, n-1], \mathtt{level})$ consists of the following three steps:

Step 1. Return $\mathrm{done} = \mathrm{TRUE}$ if $n = 3$ and $\mathtt{flip32}$ is possible for $[a, b]$ and $\mathrm{done} = \mathrm{FALSE}$ otherwise.

Step 2. For each $i \in \{0, \ldots, n-1\}$ try to remove the face $[a, b, p_i]$ by $\mathtt{flip23}$. If it is successfully flipped, reduce the size of A by 1. Update $A[0, \ldots, n-2]$ so that it contains the current set of tetrahedra sharing the edge $[a, b]$. Reuse the last entry, $A[n-1]$, to store the information of this $\mathtt{flip23}$ (see Figure 2c). It then (recursively) calls $\mathtt{flipnm}(A[0, \ldots, n-2], \mathtt{level})$. When no face can be removed, go to Step 3.

Step 3. If `level > 0`, try to remove an edge adjacent to $[a, b]$ using `flipnm`. For each $i \in \{0, \ldots, n-1\}$, let $[x, y]$ be given by either edge $[a, p_i]$ or edge $[b, p_i]$. Initialize an array $B[0, \ldots, n_1 - 1]$ of $n_1 \geq 3$ tetrahedra sharing $[x, y]$ and call `flipnm`$(B[0, \ldots, n_1 - 1], \text{level} - 1)$. If $[x, y]$ is successfully removed, reduce $|A|$ by 1. Update $A[0, \ldots, n-2]$ to contain the current set of tetrahedra sharing the edge $[a, b]$. Reuse the last entry, $A[n-1]$, to store the information of this `flipnm` and the address of the array B (to be able to release the occupied memory later). Then (recursively) call `flipnm`$(A[0, \ldots, n-2])$. Otherwise, if $[x, y]$ is not removed, call `flipnm_post`$(B[0, \ldots, n_1 - 1], m_1)$ to free the memory. Return `done = FALSE` if no edge can be removed.

Since `flipnm` is called recursively, not every face and edge should be flipped in Steps 2 and 3. In particular, if B is allocated, i.e., `flipnm` is called recursively, we skip flipping faces and edges belonging to the tetrahedra in $A \cap B$.

In the simplest case, that is, ignoring the option to reverse the flips, `flipnm_post`$(A[0, \ldots, n-1], m)$ simply walks through the array A from $A[m]$ to $A[n-1]$ and checks if a `flipnm` flip has been saved. If so, the saved array address B is extracted and its memory is released.

In Step 2 there are at most $\binom{n}{n-3}/(n-3)!$ different flip sequences, depending on the specific choice of faces in A. Each individual flip sequence is equivalent to a sequence of the n vertices (apexes) in the link of the edge $[a, b]$. We reuse the entries of A to store each flip sequence. After a 2-to-3 flip, the number of the tetrahedra in array A is reduced by one (two tetrahedra out, one terahedron in), since only one of the three new tetrahedra contains the edge $[a, b]$. The remaining tetrahedra are shifted by one in the list, so that the last entry, $A[n-1]$, can be used to store this flip (cf. Fig. 2c). In particular, the following information is saved:

- a flag indicating a 2-to-3 flip;

- the original position i, meaning that the face $[a, b, p_i]$ is flipped.

Both is compressed and stored in the entry $A[n-1]$ (note that a flag requires just a few bits of space). This stored data allows us to perform the reversal of a 2-to-3 flip as follows:

- use A and the position i with

$$A[i-1] = [a, b, p_{i-1}, p_{i+1}]$$

  to locate the three tetrahedra sharing the edge $[p_{i-1}, p_{i+1}]$: $[p_{i-1}, p_{i+1}, a, b]$, $[p_{i-1}, p_{i+1}, b, p_i]$, and $[p_{i-1}, p_{i+1}, p_i, a]$;

- perform a 3-to-2 flip on these three tetrahedra;

- insert two new tetrahedra into the array A:

$$A[i-1] = [a, b, p_i, p_{i-1}],$$
$$A[i] = [a, b, p_i, p_{i+1}].$$

In Step 3, if the selected edge $[a, p_i]$ is removed, the sequence of flips to remove $[a, p_i]$ is stored in B. We then use the last entry $A[n-1]$ to store this sequence of flips. In particular, the following information is saved:
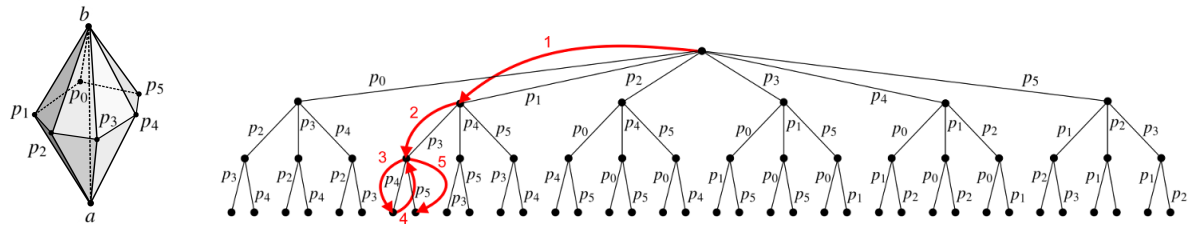
(a) The initial state.

(b) The lazy flip search tree for removing the edge $[a, b]$. p_i identifies the face $[a, b, p_i]$ which is flipped via a 2-to-3 flip. The search path is highlighted with arrows.
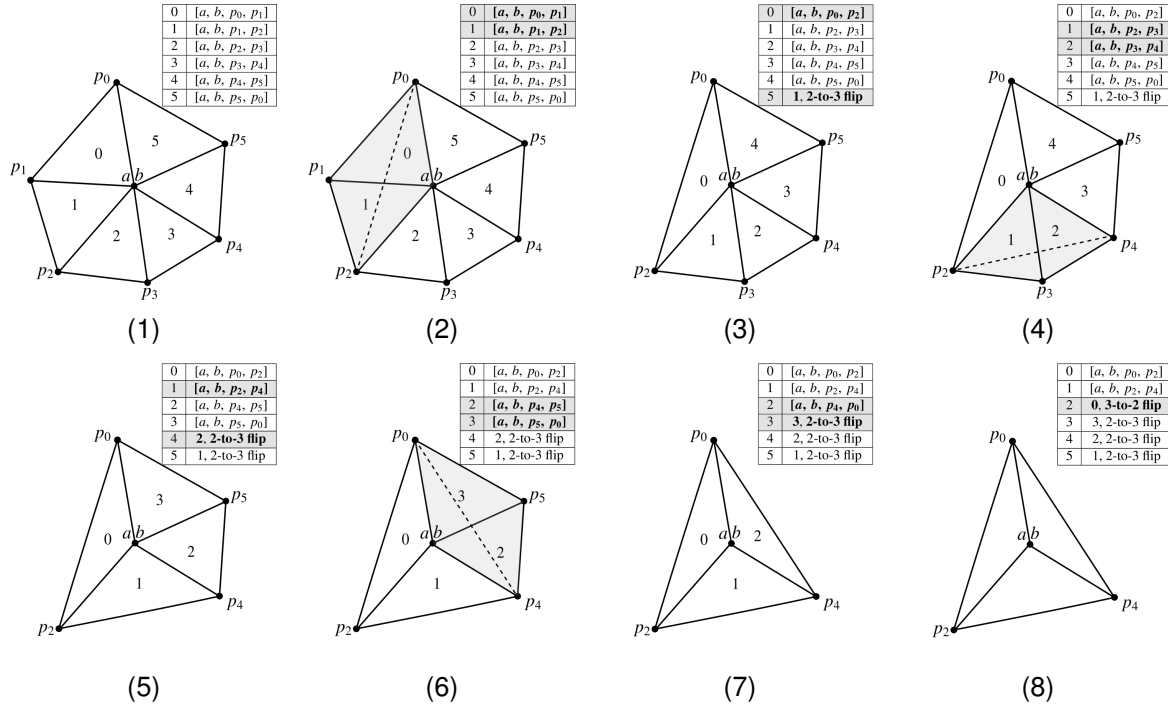
(c) The sequence of flips. The edge $[a, b]$ is represented by one vertex in the center (except (8)). A face $[a, b, p_i]$ is represented by an edge. Arrays attached to each figure show the current content of A. (1) $n = 5$ tetrahedra share the edge $[a, b]$. In (2) and (3), $[a, b, p_1]$ is removed by a 2-to-3 flip. In (4) and (5), a 2-to-3 flip is performed on $[a, b, p_3]$. In (6) and (7), $[a, b, p_5]$ is removed by a 2-to-3 flip. In (8), the edge $[a, b]$ is removed by a 3-to-2 flip.

Figure 2: An example of an edge removal by a sequence of flips.

- ■ a flag indicating that this entry stores the flip sequence to remove the edge $[a, p_i]$;

- ■ the original position i, i.e., the edge $[a, p_i]$ is flipped;

- ■ the address of the array B which stores the flip sequence.

This information allows us to reverse this sequence of flips exactly.

3.2 Lazy searching flips

During the mesh improvement process we perform flips to improve the mesh quality. Let us consider the case when it becomes necessary to remove an edge. The maximum possible number of flips for an edge removal is the Catalan number C_{n-2} (n is the size of A). Hence, the direct search for the optimal solution is only meaningful if n is very small. In most situations, an edge may not be flipped

if we restrict ourselves to adjacent faces of the edge. Our strategy is to search and perform the flips as long as they improve the current mesh quality. Our lazy searching scheme is not restricted by the number n and can be extended to adjacent edges.

The lazy searching flip scheme is like a walk in a k-ary search tree (a rooted tree with at most k children at each node, see Fig. 2b). The root represents the edge $[a, b]$ to be flipped and each of the tree nodes represents either an adjacent face $[a, b, p_i]$ or an adjacent edge $[a, p_i]$ or $[b, p_i]$ of $[a, b]$. The edges of the tree represent our search paths. In particular, the directed edge from `level` l to $l + 1$ represents either a `flip23` or a `flipnm`, and the reversed edge represents the inverse operation. The tree depth is given by the parameter `level`.

At `level` > 0, in order to to decide if an adjacent face $[a, b, p_i]$ should be flipped, we check if $[a, b, p_i]$ is flippable and make sure that this flip improves the local mesh quality. Note that we need to check only two of the three new tetrahedra: $[a, p_{i-1}, p_i, p_{i+1}]$ and $[b, p_{i-1}, p_i, p_{i+1}]$. The tetrahedron $[a, b, p_{i-1}, p_{i+1}]$ will be involved in the later flips, and will be flipped if the edge $[a, b]$ is flipped.

Once an improvement is found, the algorithm moves on to the next edge without exploring other possibilities.

4 Radial basis functions to handle curved boundaries

We describe in this section how to project the mesh on a smooth surface in order to deal with curved boundaries. We achieve this with the help of radial basis functions (RBFs), see [13, 29, 40].

4.1 Basic concepts and examples

Let $\mathcal{P}_m(\mathbb{R}^d)$ denote the space of d variate polynomials with absolute degree at most m and dimension $q := \dim \mathcal{P}_m(\mathbb{R}^d) = \binom{m-1+d}{d}$. For a basis $p_1, \ldots, p_q$ of this space, define the $N \times q$ polynomial matrix P_X through its ij^{th} entry,

$$p_{ij} = p_i(\mathbf{x}_j),$$

where $\mathbf{x}_j \in X$ and $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\} \subseteq \mathbb{R}^d$ denotes a data set. The function Φ is called *conditionally positive definite of order m* if the quadratic form

$$\mathbf{c}^T A_{\Phi,X} \mathbf{c}$$

for the distance matrix $A_{\Phi,X}$ with its ij^{th} entry defined by

$$(A_{\Phi,X})_{ij} = \Phi(\mathbf{x}_i - \mathbf{x}_j),$$

is positive for all data sets X and for all $\mathbf{c} \in \mathbb{R}^N \setminus \{\mathbf{0}\}$ which additionally satisfy the constraint $P_X^T \mathbf{c} = \mathbf{0}$.

Conditionally positive functions of order m are also conditionally positive definite for any order higher than m. Hence, the order shall denote the smallest positive integer m. A conditionally positive definite function of order $m = 0$ is called positive definite.

One speaks of radial basis functions if one additionally assumes that Φ is a radial function, i.e., there

exists a function $\phi\colon \mathbb{R}_{\geq 0} \to \mathbb{R}$ such that $\Phi(\mathbf{x}) = \phi(\|\mathbf{x}\|)$. Common examples of RBFs include:

$$\begin{aligned}
\text{Gaussian:} \quad & e^{-\|\mathbf{x}\|^2}, \\
\text{Multiquadric:} \quad & \sqrt{1 + \|\mathbf{x}\|^2}, \\
\text{Inverse Multiquadric:} \quad & 1/\sqrt{1 + \|\mathbf{x}\|^2}, \\
\text{Polyharmarmonic Spline:} \quad & \|\mathbf{x}\|^3.
\end{aligned}$$

For the numerical examples in this paper, we exclusively use the polyharmonic spline $\|\mathbf{x}\|^3$ (Fig. 3) which is conditionally positive of order 2.
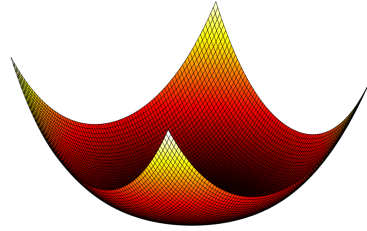

Figure 3: The polyharmonic spline $\|\mathbf{x}\|^3$.

We assume now that the interpolant $s\colon \mathbb{R}^d \to \mathbb{R}$ is given by a linear combination of translated radial basis function, augmented by a polynomial part, i.e.

$$s(\mathbf{x}) = \sum_{j=1}^{N} \alpha_j \Phi(\mathbf{x} - \mathbf{x}_j) + \sum_{k=1}^{q} \beta_k p_k(\mathbf{x}). \tag{7}$$

The interpolation conditions will then determine the linear system one needs to solve. For positive definite functions, the linear system is positive definite by construction. Hence the coefficients can be determined uniquely. It is also not difficult to verify that the interpolation conditions for conditionally positive definite functions lead to a uniquely solvable system, see [40, Theorem 8.21] for details. In the case of conditionally positive definite functions, it is known that at least $N - q$ eigenvalues of the matrix $A_{\Phi,X}$ are positive [40, Section 8.1].

4.2   Surface reconstruction with RBFs

We will assume that the surface Γ is given implicitly by the zero level set of some function $F\colon \Omega \subseteq \mathbb{R}^3 \to \mathbb{R}$, i.e.

$$\Gamma = \left\{ (x, y, z)^T \in \Omega \mid F(x, y, z) = 0 \right\}, \tag{8}$$

for some bounded domain Ω.

We cannot simply assume that the target function (which we wish to interpolate) is the zero level set of the function F since the right-hand side of the linear system one needs to solve would vanish which in turn implies that the coefficients vanish as well. Carr et al. [3] therefore made the additional assumption that the normal vectors are known. Then one can also prescribe on-surface and off-surface points. Assume that the points on the surface are denoted with $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ and the corresponding normal vectors with $M = \{\mathbf{n}_1, \ldots, \mathbf{n}_N\}$. We define the surface interpolation problem

$$\begin{aligned}
s(\mathbf{x}_i) = F(\mathbf{x}_i) = 0, \qquad & 1 \leq i \leq N \\
s(\mathbf{x}_i + \varepsilon\mathbf{n}_i) = F(\mathbf{x}_i + \varepsilon\mathbf{n}_i) = \varepsilon, \quad & N+1 \leq i \leq 2N
\end{aligned} \tag{9}$$

for some parameter $\varepsilon > 0$. Since the right-hand side of the linear system does not vanish anymore, we find a nontrivial solution. Recently, this surface interpolation technique was combined with the higher dimensional embedding technique [6, 5] to construct curvature-aligned anisotropic surface meshes. In this context the data set X corresponds to the vertices of the mesh.

4.3 Projection onto the reconstructed surface

In order to project a point onto the smooth surface reconstruction. We follow an idea by Hartmann [19] who provided a robust algorithm to project a point onto a surface defined via the zero level set of a function. This procedure is a combination of orthogonal projections on tangent planes as well as tangent parabolas. It requires only first order derivatives and uses a steepest descent method. The combination of this projection method with RBF surface reconstruction (Fig. 4) has also been discussed in [6, 5].
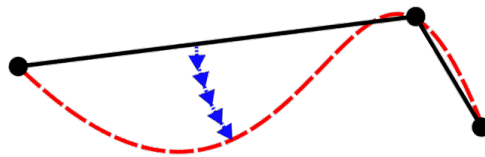

Figure 4: Projection onto a smooth surface for the example of edge splitting: the introduced edge midpoint is projected onto the RBF surface reconstruction of the mesh via a steepest descent method.

5 Mesh improvement strategy

The goal of the proposed algorithm is to obtain a new isotropic mesh whose elements come "as close as possible" to the equilateral one. To achieve this goal, we combine the local and global mesh operations described in Sections 2 and 3.

5.1 Mesh quality

To say "as close as possible to an equilateral tetrahedron" is somewhat vague from a mathematical point of view. To have a more precise criterion, the majority of the mesh improvement algorithms define a computable quantity $q(K)$ which quantifies how far a tetrahedron K is from being equilateral [31, 8, 18, 20, 37, 12]. Here, we take into account the following two:

Aspect Ratio: This is one of the most classical ways to evaluate the quality of a tetrahedron. It is defined as

$$q_{ar}(K) := \sqrt{\frac{2}{3}} \frac{L}{h}, \tag{10}$$

where L is the longest edge and h is the shortest altitude of the tetrahedron K. By construction, $q_{ar}(K) \geq 1$ and an equilateral tetrahedron is characterized by $q_{ar}(K) = 1$.

Min-max Dihedral Angle: For each tetrahedron K we consider both the minimal and the maximal dihedral angles $\theta_{\min,K}$ and $\theta_{\max,K}$. An equilateral tetrahedron has $\theta_{\min,K} = \theta_{\max,K} = \arccos(1/3) \approx 70.56°$. Applying an operation that increases $\theta_{\min,K}$ or decreases $\theta_{\max,K}$ of a given tetrahedron K makes K "closer" to the equilateral shape. Note that this is not a classical

---

Algorithm 1 The proposed mesh improvement scheme.

IMPROVE$(\mathcal{T}_h^{ini}, \theta_{\lim})$

1: **repeat**
2:     **repeat**
3:         **repeat**
4:             **repeat**
5:                 **repeat**
6:                     *MMPDE-based smoothing*
7:                     *RBF surface reconstruction* $\rightarrow$ smooth and
8:                     *lazy flips*
9:                 **until** no point is moved or no flip is done or $Q(\mathcal{T}_h) \geq \theta_{\lim}$
                flip
10:                 remove the edges $l_{\mathbf{e}} < 0.5\, l_{\mathrm{ave}}$ $\rightarrow$ main loop
11:                 *lazy flips*
12:             **until** no edge is contracted or $Q(\mathcal{T}_h) \geq \theta_{\lim}$
13:             split the edges $l_{\mathbf{e}} > 1.5\, l_{\mathrm{ave}}$
14:             *RBF surface reconstruction*
15:             *lazy flips*
16:         **until** no edge is split or $Q(\mathcal{T}_h) > \theta_{\lim}$
17:         split the tetrahedra K such that $\theta_{\min,K} < \theta_{\lim}$
18:         *lazy flips*
19:     **until** no tetrahedron is removed or $Q(\mathcal{T}_h) > \theta_{\lim}$
20:     change the flip criterion for the *lazy flips*
21: **until** no operation is done in the main loop or $Q(\mathcal{T}_h) > \theta_{\lim}$

---

quality measure since we associate two quantities with each tetrahedron, which is one of the novel aspects of the proposed mesh improvement procedure.

These two quality measures refer to a single tetrahedron K of the mesh. However, the design of our mesh improvement scheme requires a quality measure for the *whole* mesh as a stopping criterion. To estimate the quality of the whole mesh, we define the global parameter

$$Q(\mathcal{T}_h) := \min_{K \in \mathcal{T}_h} \left(\theta_{\min,K} \right). \tag{11}$$

If we consider a target dihedral angle $\theta_{\lim}$ and obtain a mesh $\mathcal{T}_h$ with $Q(\mathcal{T}_h) > \theta_{\lim}$, then *all* dihedral angles are guaranteed to be greater than $\theta_{\lim}$.

5.2   The scheme

The inputs for the mesh improvement algorithm are a tetrahedral mesh $\mathcal{T}_h^{ini}$ of a PLC and a target minimum angle $\theta_{\lim}$. The output is a mesh $\mathcal{T}_h^{fin}$ where each element has a minimum dihedral angle greater than $\theta_{\lim}$.

The scheme is presented in Algorithm 1 and consists of five nested "**repeat ... until**" loops, whose stopping criterion depends on the operations done inside the loop and $Q(\mathcal{T}_h)$. We apply the MMPDE smoothing and the *lazy flip* in the most internal loop (lines 5 to 9). The *lazy flip* is also exploited in the

outer loops both on the whole mesh (lines 11, 15 and 18) and on the tetrahedra involved in the local operations (lines 10, 13 and 17).

It is possible to consider several flipping criteria for the *lazy flip*, which makes the design of the scheme flexible. We exploit this feature by using two objective functionals and changing the flipping criterion in each iteration of the outer loop (line 20) by

1   maximizing $\theta_{\min,K}$ and minimizing $\theta_{\max,K}$ (simultaneously),

2   minimizing the aspect ratio.

The stopping criterion is always based on the minimal dihedral angle, $Q(\mathcal{T}_h)$, and the number of operations done.

After a number of iterations both the flipping and the smoothing procedure can stagnate, i.e., the mesh $\mathcal{T}_h$ converges to a fixed configuration where neither flips nor smoothing can improve the quality of the mesh. Unfortunately, it is not a priori guaranteed that such a mesh satisfies the constraint on the target minimum dihedral angle $\theta_{\lim}$. To overcome this difficulty, we apply edge splitting, edge contraction, and point insertion when this stagnation occurs (lines 10, 13 and 17 in Algorithm 1).

For the edge contraction and splitting, we use the standard edge length criterion: we compute the average edge length l_{ave} of the actual mesh, contract the edges shorter than $0.5\,l_{\text{ave}}$ (line 10), and split (halve) the ones longer than $1.5\,l_{\text{ave}}$ (line 13). In line 17, we split a tetrahedron K with $\theta_{\min,K} < \theta_{\lim}$ via a standard 1-to-4 flip by placing the newly added point at the barycenter of K [10]. In this way, the algorithm constructs via flipping and smoothing a mesh satisfying $Q(\mathcal{T}_h) > \theta_{\lim}$. At the moment, we are not interested in optimizing these operations, we exploit them only to overcome the stagnation of the algorithm.

The MMPDE smoothing can be easily parallelized because the computation of nodal velocities (Eqs. (5) and (6)) requires local, element-wise computations which are independent from each other; we parallelize it with OpenMP [4] in order to accelerate the mesh improvement scheme. On the other hand, the *lazy flip* may propagate to neighbors and neighbors of neighbors, thus, it is complex and difficult to parallelize; in our tests we use a sequential implementation.

6   Numerical examples

We test the proposed mesh improvement algorithm and compare it with the mesh improvement algorithm of `Stellar` [31], the remeshing procedure of `CGAL` [39], and `mmg3d` [8]. We compare the histograms of the dihedral angles of final meshes, the minimal and the maximal dihedral angles $\theta_{\min,\mathcal{T}_h}$ and $\theta_{\max,\mathcal{T}_h}$, the mean dihedral angle $\mu_{\mathcal{T}_h}$, and its standard deviation $\sigma_{\mathcal{T}_h}$.

6.1   Piecewise linear complexes (PLCs)

To analyze the effectiveness of the proposed mesh improvement scheme in case of a piecewise linear complex domain, we consider the following three examples (for more PLC examples, see [7]):

■   RAND1 tetrahedral meshes of a cube generated by inserting randomly located vertices inside and on the boundary [31] (Fig. 6),
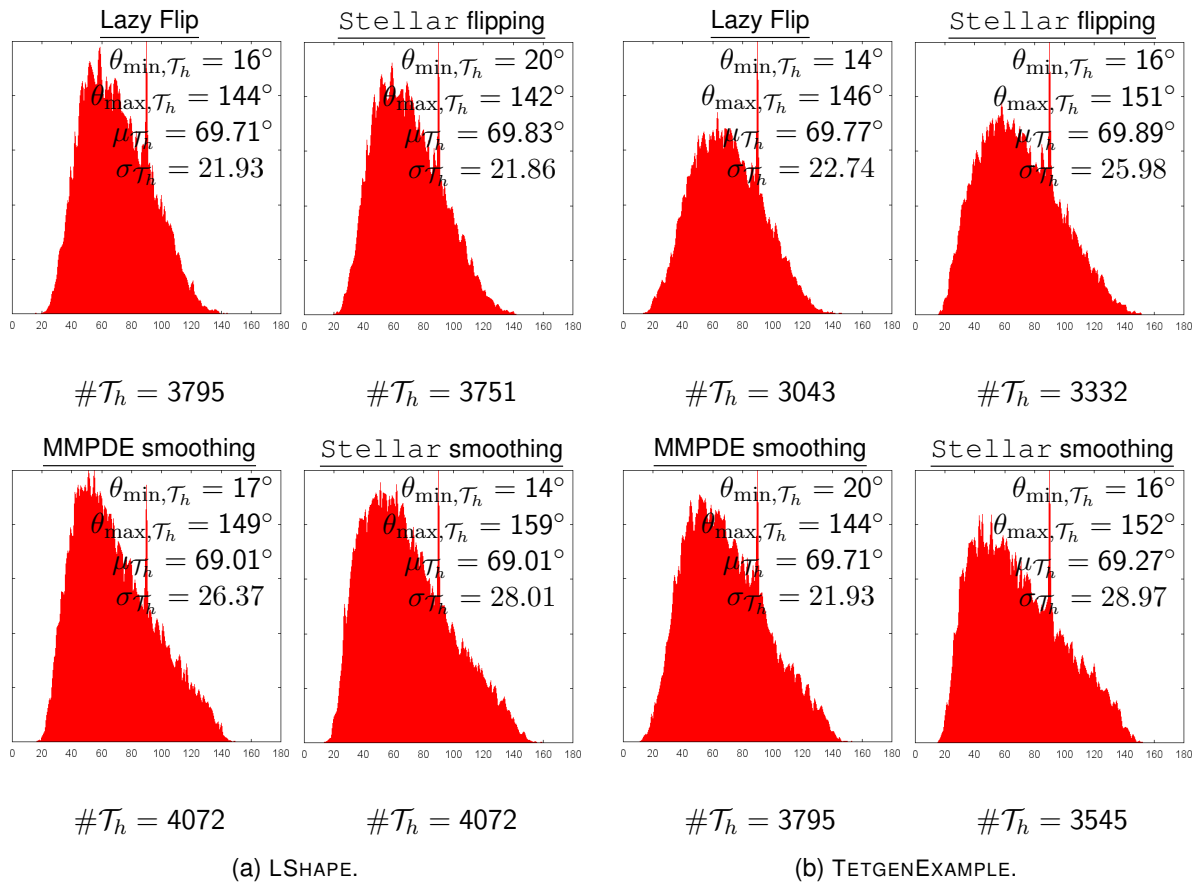
Figure 5: Comparison of flipping only (first row) and smoothing only (second row) for the initial meshes LSHAPE and TETGENEXAMPLE.

- LSHAPE is a tetrahedral mesh of an L-shaped PLC generated by `TetGen` [38] without optimizing the minimal dihedral angle (switches `-pa0.019`, Fig. 7),

- TETGENEXAMPLE is a tetrahedral example mesh of a non-convex PLC with a hole provided by `TetGen` (Fig. 8).

Smoothing and flipping by themselves. Before testing the full mesh improvement scheme, we test the effectiveness of the MMPDE smoothing and the *lazy flip* by themselves and employ smoothing and flipping separately, i.e., we improve a tetrahedral mesh exploiting *only* the flipping operation or the vertex smoothing. We compare our results with the ones provided by `Stellar` for the examples LSHAPE and TETGENEXAMPLE.

The results of the *lazy flip* are comparable to the `Stellar` flips (Fig. 5, first row). However, the MMPDE smoothing is better than its counterpart in `Stellar` (Fig. 5, second row): in both examples it achieves larger $\theta_{\min,\mathcal{T}_h}$, noticeably smaller $\theta_{\max,\mathcal{T}_h}$, and a smaller standard deviation of the mean dihedral angle.

Full scheme. We compare the whole scheme with the mesh improvement algorithm of `Stellar` [31], the remeshing procedure of `CGAL` [39], and `mmg3d` [8] (Figs 6 to 8).

Although all methods provide good results, the new scheme is better: $\theta_{\min,\mathcal{T}_h}$ is larger than the value obtained by `CGAL` or `mmg3d` and comparable to the value obtained by `Stellar`. Moreover, $\theta_{\max,\mathcal{T}_h}$ is smaller than the values obtained by `Stellar`, `CGAL`, or `mmg3d` in all examples but one in Fig. 8a.
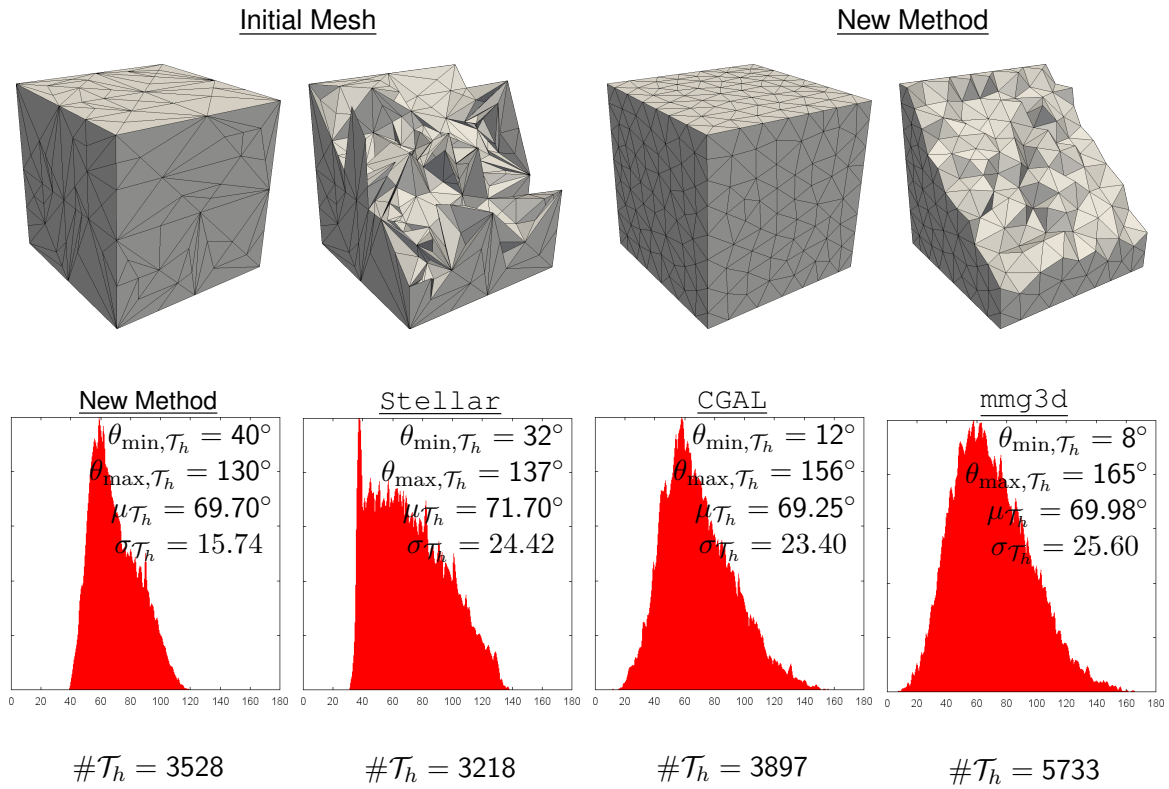
Figure 6: RAND1. The initial mesh with $\#\mathcal{T}_h = 5104$, the final (optimized) mesh, and statistics of dihedral angles for the final meshes.
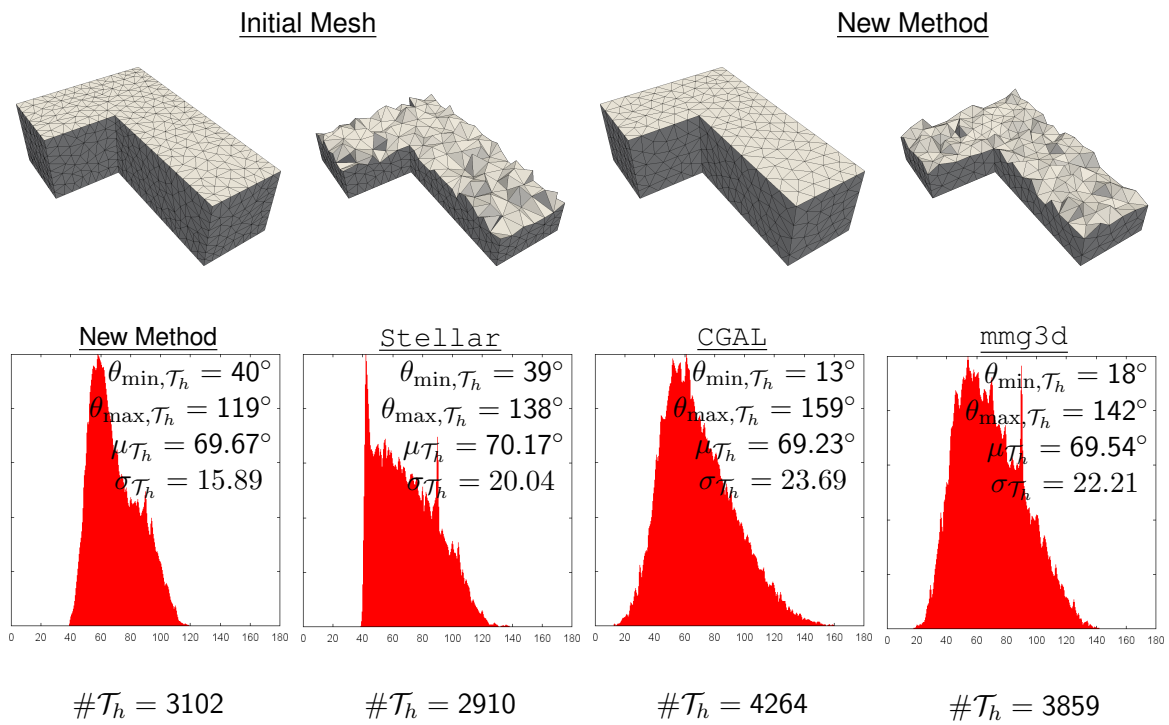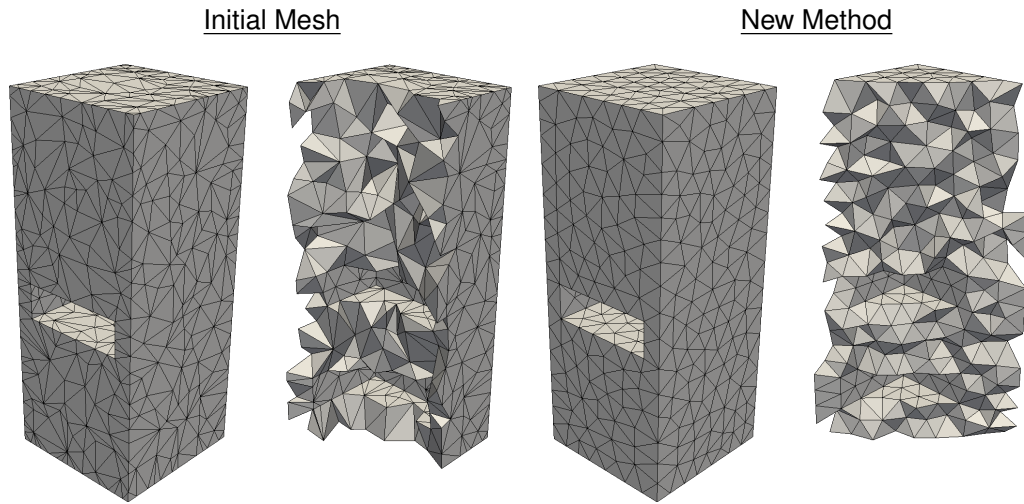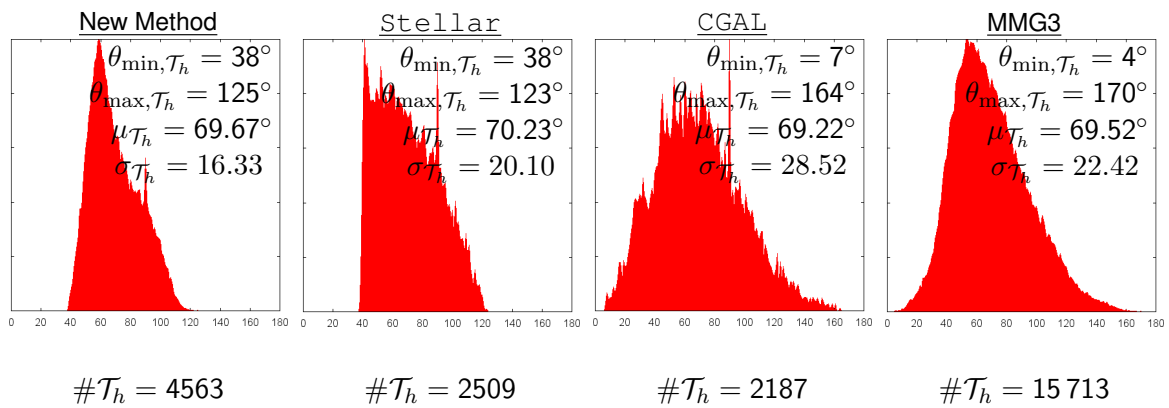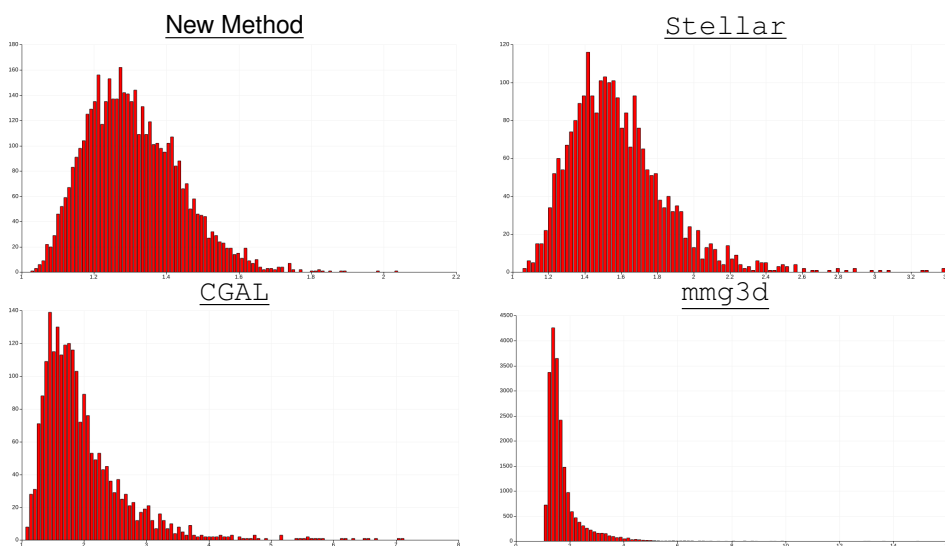

Figure 7: LSHAPE. The initial mesh with $\#\mathcal{T}_h = 4072$, the final (optimized) mesh, and statistics of dihedral angles for the final meshes.

Initial Mesh   New Method

(a) The initial mesh with $\#\mathcal{T}_h = 3545$ and the final optimized mesh.

New Method   Stellar   CGAL   MMG3

$\theta_{\min,\mathcal{T}_h} = 38°$ $\theta_{\min,\mathcal{T}_h} = 38°$ $\theta_{\min,\mathcal{T}_h} = 7°$ $\theta_{\min,\mathcal{T}_h} = 4°$

$\theta_{\max,\mathcal{T}_h} = 125°$ $\theta_{\max,\mathcal{T}_h} = 123°$ $\theta_{\max,\mathcal{T}_h} = 164°$ $\theta_{\max,\mathcal{T}_h} = 170°$

$\mu_{\mathcal{T}_h} = 69.67°$ $\mu_{\mathcal{T}_h} = 70.23°$ $\mu_{\mathcal{T}_h} = 69.22°$ $\mu_{\mathcal{T}_h} = 69.52°$

$\sigma_{\mathcal{T}_h} = 16.33$ $\sigma_{\mathcal{T}_h} = 20.10$ $\sigma_{\mathcal{T}_h} = 28.52$ $\sigma_{\mathcal{T}_h} = 22.42$

$\#\mathcal{T}_h = 4563$ $\#\mathcal{T}_h = 2509$ $\#\mathcal{T}_h = 2187$ $\#\mathcal{T}_h = 15\,713$

(b) Dihedral angle comparison for the final meshes.

New Method   Stellar

CGAL   mmg3d

(c) Aspect ratio comparison for the final meshes.

Figure 8: TETGENEXAMPLE. The initial mesh with $\#\mathcal{T}_h = 3545$, the final (optimized) mesh, and statistics of the dihedral angles and the aspect ratio.

Our method provides mean dihedral angles $\mu_{\mathcal{T}_h}$ around $69.6°$, which is close to the optimal value of $\arccos{(1/3)} \approx 70.56°$. Moreover, standard deviations $\sigma_{\mathcal{T}_h}$ are always smaller than the ones of other methods. Indeed, we get a distribution of dihedral angles close to the mean value. This quantitative consideration becomes clearer from the shape of the histograms in Figs 6 to 8.

For the TETGENEXAMPLE (Fig. 8) we also provide aspect ratio histograms (the results for the other examples are very similar and we omit them). The aspect ratio of an equilateral tetrahedron is equal to 1 and the more a tetrahedron is distorted and stretched the greater its aspect ratio becomes. Our method and `Stellar` clearly provide the best aspect ratio distribution. For our method, the vast majority of tetrahedra have an aspect ratio smaller than 1.8. The `Stellar` mesh is slightly worse with most of its tetrahedra having aspect ratios below 2.6.

6.2 Curved boundary domains

In the last part of this section, we experimentally demonstrate some examples with curved domains. We study two types of examples: one academic example for using the RBF surface reconstruction to project the boundary vertices on the smooth approximation of the discrete surface and two more complex examples with fixed boundary vertices.

First, we consider the discrete ellipsoid mesh (Fig. 9). Though it has a simple geometry, it requires some effort since the boundary is curved and no longer a PLC. The main challenge is to project the boundary vertices back onto the smooth surface if they leave it after a mesh improvement step. For this reason, we reconstruct the surface via RBFs (see Section 4.2) to assist the mesh optimization and project the moved (smoothed) boundary vertices to the reconstructed surface. Figure 9 shows that RBF reconstruction smoothes the initially rough surface approximation. The obtained tetrahedral mesh has high quality: the mean dihedral angle is close to the optimal value ($\approx 70.69°$) and the standard deviation of the dihedral angles is small ($\approx 18.16°$).

However, it has to be pointed out that complicated boundaries cannot be handled as easily as an ellipsoid and require more sophisticated methods.

In our next examples, we restrict ourselves to the case of fixed boundary vertices since `Stellar` does not handle curved surfaces described via an implicit function, start with a good isotropic triangular mesh as input, and keep the boundary vertices fixed for each of the algorithms.

Fixed boundary. The next two examples are meshes of a spinal bone and of an elephant (figs. 10 and 11). The initial surface meshes in both examples are constructed by means of the higher di-mensional embedding approach for surface mesh reconstruction [5] and their minimal face angles are approximately $33°$. The initial volume meshes are constructed by `TetGen` using the `-Y` flag to preserve the fixed boundary.

Figures 10 and 11 present the histograms of the dihedral angles of the resulting optimized meshes. In comparison to the PLC examples, where the geometry is simpler and the boundary vertices are allowed to move, the smallest dihedral angles for the spinal bone and the elephant examples are worse (smaller) than for the PLC examples. In comparison to `Stellar`, our algorithm achieves better values for $\theta_{\min,\mathcal{T}_h}$ and $\theta_{\max,\mathcal{T}_h}$, as well as a smaller mean deviation from the mean value.

These examples, too, show the "aggressive" nature of the `Stellar` mesh improvement algorithm, which aggressively removes vertices during the mesh improvement. In contrast, our mesh improvement scheme is able to produce a high-quality mesh while keeping the number of vertices close to the original input.
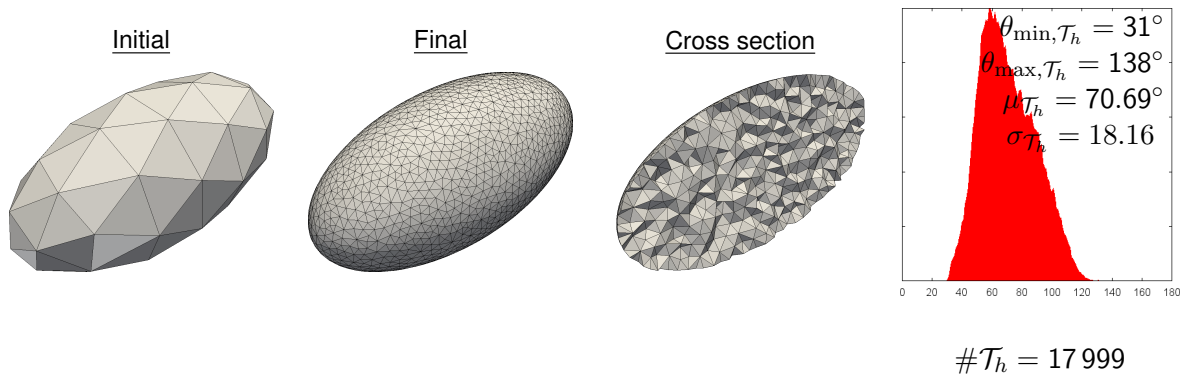
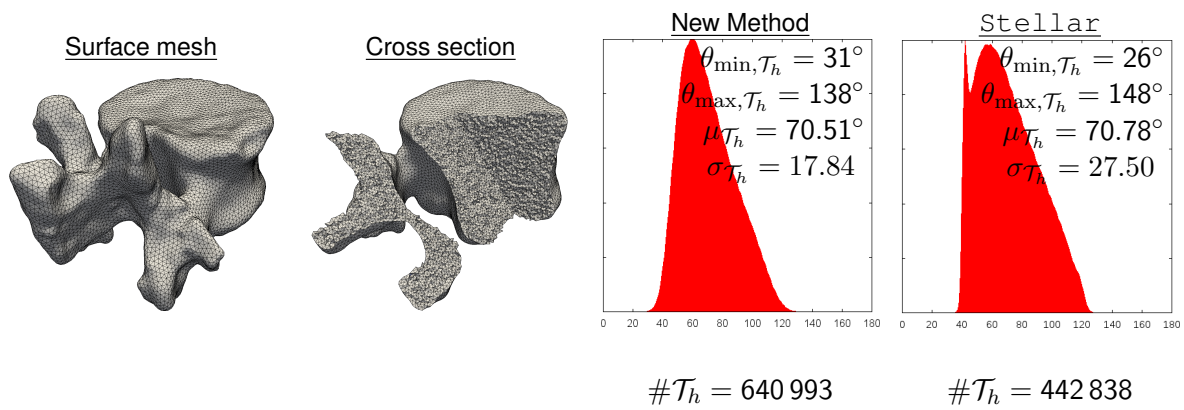Figure 9: Example of tetrahedral mesh improvement with a curved (reconstructed) surface.

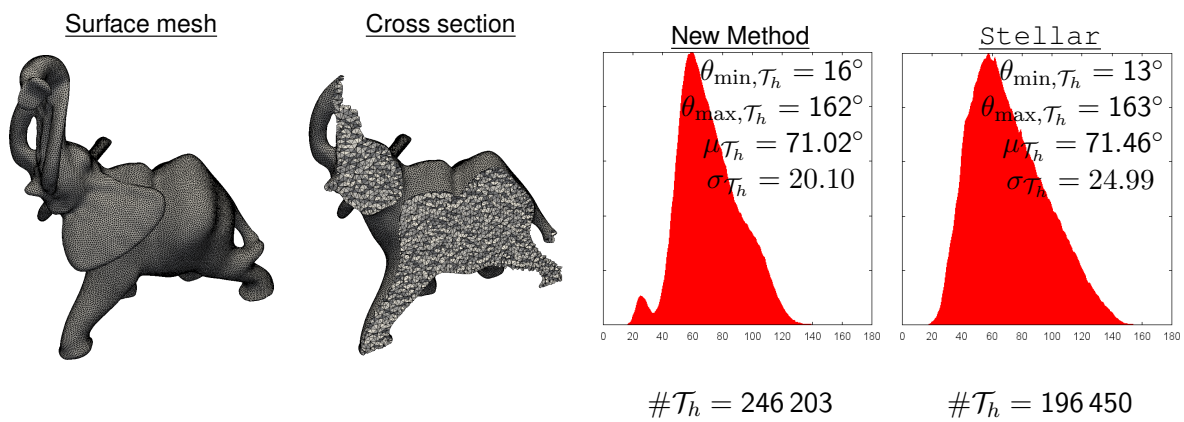Figure 10: Spine example: the initial mesh with $\#\mathcal{T}_h = 688\,420$ and the final optimized mesh.

Figure 11: Elephant example: the initial mesh with $\#\mathcal{T}_h = 260\,401$ and the final optimized mesh.

7 Conclusions

Mesh improvement is a challenging problem and we tackled it by combining several recently developed techniques, namely, moving mesh smoothing, lazy flipping, and RBF surface reconstruction. In comparison to the mesh improvement algorithm `Stellar` and the re-meshing procedures provided by `CGAL` and `mmg3d`, we obtain better results in terms of the distributions of dihedral angles for all considered examples. However, there are several directions in which this work could be extended.

First, for smooth and relatively simple boundaries, our approach works excellently but complicated curved boundaries pose a challenging problem. One possible solution could be the direct incorporation of the boundary description into the MMPDE smoothing scheme (parametrization) so that the boundary vertices will always stay on the surface. This will avoid the sometimes troublesome projection of vertices and velocities back onto the surface after a smoothing step.

Second, we need to find a more sophisticated method for edge contraction and splitting in order to improve the performance of both the MMPDE smoothing and the lazy flip.

Third, the MMPDE smoothing is based on the moving mesh method [23] which allows the definition of a metric field. Hence, the moving mesh smoothing can be extended to the adaptive and anisotropic setting.

References

[1] R. E. Bank. *PLTMG: a software package for solving elliptic partial differential equations*, volume 15 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994. Users' guide 7.0.

[2] S. A. Canann, J. R. Tristano, and M. L. Staten. An approach to combined Laplacian and optimization-based smoothing for triangular, quadrilateral, and quad-dominant meshes, in: Proceedings of the 7th International Meshing Rountable, 1998.

[3] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 67–76. ACM, 2001.

[4] L. Dagum and R. Menon. OpenMP: an industry standard api for shared-memory programming. *IEEE Comput. Sci. Eng.*, 5(1):46–55, 1998.

[5] F. Dassi, P. Farrell, and H. Si. An anisoptropic surface remeshing strategy combining higher dimensional embedding with radial basis functions. *Procedia Eng.*, 163:72–83, 2016. 25th International Meshing Roundtable.

[6] F. Dassi, P. Farrell, and H. Si. A novel remeshing scheme via radial basis functions and higher dimensional embedding. *WIAS Preprint 2265*, 2016.

[7] F. Dassi, L. Kamenski, and H. Si. Tetrahedral mesh improvement using moving mesh smoothing and lazy searching flips. *Procedia Eng.*, 163:302–314, 2016. 25th International Meshing Roundtable.

[8] C. Dobrzynski. MMG3D: User Guide. Technical Report RT-0422, INRIA, `https://hal.inria.fr/hal-00681813`, 2012.

[9] J. R. Dormand and P. J. Prince. A family of embedded Runge-Kutta formulae. *J. Comput. Appl. Math.*, 6(1):19–26, 1980.

[10] H. Edelsbrunner, M. J. Ablowitz, S. H. Davis, E. J. Hinch, A. Iserles, J. Ockendon, and P. J. Olver. *Geometry and Topology for Mesh Generation (Cambridge Monographs on Applied and Computational Mathematics)*. Cambridge University Press, New York, NY, USA, 2006.

[11] D. A. Field. Laplacian smoothing and Delaunay triangulations. *Comm. Appl. Num. Meth.*, 4:709–712, 1978.

[12] D. A. Field. Qualitative measures for initial meshes. *Internat. J. Numer. Methods Engrg.*, 47(4):887–906, 2000.

[13] B. Fornberg and N. Flyer. *A Primer on Radial Basis Functions with Applications to the Geosciences*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2015.

[14] L. Freitag, M. Jones, and P. Plassmann. A parallel algorithm for mesh smoothing. *SIAM J. Sci. Comput.*, 20(6):2023–2040, 1999.

[15] L. A. Freitag and P. M. Knupp. Tetrahedral mesh improvement via optimization of the element condition number. *Internat. J. Numer. Methods Engrg.*, 53(6):1377–1391, 2002.

[16] L. A. Freitag and C. Ollivier-Gooch. Tetrahedral mesh improvement using swapping and smoothing. *Internat. J. Numer. Methods Engrg.*, 40(21):3979–4002, 1997.

[17] P. George and H. Borouchaki. Back to edge flips in 3 dimensions, in: Proceedings of the 12th international meshing rountable, September 2003.

[18] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Internat. J. Numer. Methods Engrg.*, 79(11):1309–1331, 2009. `http://gmsh.info`.

[19] E. Hartmann. On the curvature of curves and surfaces defined by normal forms. *Comput. Aided Geom. Design*, 16(5):355 – 376, 1999.

[20] F. Hecht. New development in freefem++. *J. Numer. Math.*, 20(3-4):251–265, 2012.

[21] W. Huang. Variational mesh adaptation: isotropy and equidistribution. *J. Comput. Phys.*, 174(2):903–924, 2001.

[22] W. Huang. Mathematical principles of anisotropic mesh adaptation. *Commun. Comput. Phys.*, 1(2):276–310, 2006.

[23] W. Huang and L. Kamenski. A geometric discretization and a simple implementation for variational mesh generation and adaptation. *J. Comput. Phys.*, 301:322–337, 2015.

[24] W. Huang and L. Kamenski. On the mesh nonsingularity of the moving mesh PDE method, 2015. Submitted (in revision).

[25] W. Huang, L. Kamenski, and R. D. Russell. A comparative numerical study of meshing functionals for variational mesh adaptation. *J. Math. Study*, 48(2):168–186, 2015.

[26] W. Huang, L. Kamenski, and H. Si. Mesh smoothing: an MMPDE approach, 2015. Research Notes of the 24th International Meshing Roundtable.

[27] W. Huang, Y. Ren, and R. D. Russell. Moving mesh partial differential equations (MMPDES) based on the equidistribution principle. *SIAM J. Numer. Anal.*, 31(3):709–730, 1994.

[28] W. Huang and R. D. Russell. *Adaptive Moving Mesh Methods*, volume 174 of *Applied Mathematical Sciences*. Springer, New York, 2011.

[29] A. Iske. *Multiresolution Methods in Scattered Data Modelling*, volume 37 of *Lecture Notes in Computational Science and Engineering*. Springer, Berlin, 2004.

[30] B. Joe. Construction of three-dimensional improved-quality triangulations using local transformations. *SIAM J. Sci. Comput.*, 16(6):1292–1307, 1995.

[31] B. M. Klingner and J. R. Shewchuk. Aggressive tetrahedral mesh improvement. In *Proceedings of the 16th International Meshing Roundtable*, pages 3–23, Oct. 2007. `http://graphics.cs.berkeley.edu/papers/Klingner-ATM-2007-10`.

[32] P. M. Knupp. Applications of mesh smoothing: copy, morph, and sweep on unstructured quadrilateral meshes. *Internat. J. Numer. Methods Engrg.*, 45(1):37–45, 1999.

[33] P. M. Knupp. Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantities. Part I – A framework for surface mesh optimization. *Internat. J. Numer. Methods Engrg.*, 48:401–420, 2000.

[34] S. H. Lo. A new mesh generation scheme for arbitaar planar domains. *Int. J. Numer. Meth. Engng.*, 21:1403–1426, 1985.

[35] G. L. Miller. *Solving Irregularly Structured Problems in Parallel: 5th International Symposium, IRREGULAR'98 Berkeley, California, USA, August 9–11, 1998 Proceedings*, chapter Control volume meshes using sphere packing, pages 128–131. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998.

[36] M. Shephard and M. Georges. Automatic three-dimensional mesh generation by the finite octree technique. *Int. J. Numer. Meth. Engrg.*, 32:709–749, 1991.

[37] J. R. Shewchuk. What is a good linear finite element? Interpolation, conditioning, anisotropy, and quality measures, 2002.

[38] H. Si. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.*, 41(2):11:1–36, 2015. `http://tetgen.org`.

[39] The CGAL Project. *CGAL User and Reference Manual*, 4.8 edition, 2016. `https://doc.cgal.org/4.8`.

[40] H. Wendland. *Scattered Data Approximation*, volume 17 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2005.